

[2016-New Ensure Pass 70-464 Exam By Training GreatExam New VCE And PDF Dumps (21-40)

We are all well aware that a major problem in the IT industry is that there is a lack of quality study materials. Our exam preparation material provides you everything you will need to take a certification examination. Our Microsoft 70-464 Exam will provide you with exam questions with verified answers that reflect the actual exam. These questions and answers provide you with the experience of taking the actual test. High quality and value for the 70-464 Exam. 100% guarantee to pass your Microsoft 70-464 exam and get your Microsoft certification. QUESTION 21 You have a database hosted on SQL Azure. You are developing a script to create a view that will be used to update the data in a table. The following is the relevant portion of the script. (Line numbers are included for reference only.) You need to ensure that the view can update the data in the table, except for the data in Column1. Which code segment should you add at line 06?

```
01 CREATE VIEW View1
02 AS
03 SELECT
04 ...
05 WHERE Column1 = 'City1'
06 www.greatexam.com
```

A. WITH ENCRYPTION B. WITH VIEW_METADATA C. WITH CHECK OPTION D. WITH SCHEMABINDING

Answer: C Explanation: The question concerning the view that has a clause "WHERE Column1 = 'City1'" is wrong. That's not what the CHECK option is made for. Actually you will be able to update ONLY the rows satisfied by that WHERE clause, that is, only the rows with the Column1 being 'City1'. None of the answers are valid from that question. You need a trigger to achieve that.

<http://msdn.microsoft.com/en-us/library/ms187956.aspx> QUESTION 22 You have an application that uses a view to access data from multiple tables. You need to ensure that you can insert rows into the underlying tables by using the view. What should you do?

A. Define the view by using the SCHEMABINDING option. B. Define the view by using the CHECK option. C. Create an INSTEAD OF trigger on the view. D. Materialize the view. Answer: C Explanation:

<http://msdn.microsoft.com/en-us/library/ms180800.aspx> <http://msdn.microsoft.com/en-us/library/ms187956.aspx> QUESTION 23 You have a SQL Azure database. You execute the following code: The Sales.Customers table will contain 100,000 rows. You expect the FaxNumber column to contain a null value for 70 percent of the rows. You need to create an index to support Sales.CustomersByFaxNumber. The solution must minimize the disk storage requirements. Which code segment should you execute?

```
CREATE SCHEMA Sales;
GO
```

```
CREATE TABLE Sales.Customers
(
    CustomerID int IDENTITY(1,1) PRIMARY KEY,
    FaxNumber char(10) SPARSE NULL,
    CustomerName varchar(100) NOT NULL,
    EmailAddress varchar(100) NOT NULL
);
GO
```

```
CREATE PROCEDURE Sales.CustomersByFaxNumber
    @FaxNumber char(10)
AS
SELECT CustomerID,
    CustomerName
FROM Sales.Customers
WHERE FaxNumber = @FaxNumber www.greatexam.com
```

A. CREATE INDEX IX_Customers ON Customers (FaxNumber) WHERE FaxNumber IS NOT NULL B. CREATE INDEX IX_Customers ON Customers (FaxNumber) WITH FILLFACTOR=0 C. CREATE INDEX IX_Customers ON Customers (CustomerName) INCLUDE (FaxNumber) D. CREATE INDEX IX_Customers ON Customers (FaxNumber) E. CREATE INDEX IX_Customers ON Customers (FaxNumber) WHERE FaxNumber IS NULL Answer: A QUESTION 24 You run the following code: You need to ensure that the root node of the XML data stored in the Details column is <Order_Details>. What should you implement? More than one answer choice may achieve the goal. Select the BEST answer.

```
CREATE TABLE dbo.Orders
(
    Id int CONSTRAINT PK_Order_Id PRIMARY KEY,
    Amount decimal,
    Details xml
);
www.greatexam.com
```

A. an XML index
 B. an XML schema collection
 C. a user-defined data type
 D. a Data Definition Language (DDL) trigger
 E. a data manipulation language (DML) trigger
 Answer: B
 Explanation: <http://msdn.microsoft.com/en-us/library/ms187856.aspx>
QUESTION 25
 Drag and Drop Question
 You have a table named Customers that has a clustered index defined on the ID column. You write a script to create a stored procedure. You need to complete the script for the stored procedure. The solution must minimize the number of locks and deadlocks. What should you do? To answer, drag the appropriate option to the correct location in the answer area. (Answer choices may be used once, more than once, or not at all.)

READ COMMITTED	CREATE PROCEDURE Proc1 (@ParamID int)
SERIALIZABLE	AS
WITH(UPDLOCK)	SET TRANSACTION ISOLATION LEVEL
WITH(XLOCK)	BEGIN TRANSACTION
	DECLARE @var as NCHAR(10)
	Select @var = Name
	FROM dbo.Customers
	WHERE ID = @ParamID
	...
	UPDATE dbo.Customers
	SET Name = @var
	WHERE ID = @ParamID
	COMMIT TRANSACTION;
	GO

www.greatexam.com

Answer:

SERIALIZABLE	CREATE PROCEDURE Proc1 (@ParamID int)
	AS
WITH(XLOCK)	SET TRANSACTION ISOLATION LEVEL
	READ COMMITTED
	BEGIN TRANSACTION
	DECLARE @var as NCHAR(10)
	Select @var = Name
	FROM dbo.Customers
	WITH(UPDLOCK)
	WHERE ID = @ParamID
	...
	UPDATE dbo.Customers
	SET Name = @var
	WHERE ID = @ParamID
	COMMIT TRANSACTION;
	GO

www.greatexam.com

QUESTION 26
 You are creating a table named Orders. You need to ensure that every time a new row is added to the Orders table, a user-defined function is called to validate the row before the row is added to the table. What should you use? More than one answer choice may achieve the goal. Select the BEST answer.
 A. a FOREIGN KEY constraint
 B. a data manipulation language (DML) trigger
 C. a DEFAULT constraint
 D. a CHECK constraint
 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>
QUESTION 27
 You execute the following code

```
CREATE TABLE UserInfo
(
    ID int NOT NULL IDENTITY (1, 1)
    CONSTRAINT PK_UserInfo PRIMARY KEY
    UserName varchar(100) NOT NULL,
    Manager varchar(100) NULL,
    HireDate date NOT NULL,
    PerformanceReviewScore int NULL
);

You have a stored procedure that includes the following code:

SELECT UserName, PerformanceReviewScore
FROM UserInfo
WHERE Manager = 'Ben Smith';
```

You need to create a covering index on UserInfo. Which code segment should you execute?

- A.

```
CREATE NONCLUSTERED INDEX [IX_Covering_Inf]
(
  [UserName] ASC,
  [PerformanceReviewScore] ASC,
);
```
- B.

```
CREATE NONCLUSTERED INDEX [IX_Covering_Inf]
(
  [Manager] ASC
);
```
- C.

```
CREATE NONCLUSTERED INDEX [IX_Covering_Inf]
(
  [UserName] ASC,
  [Manager] ASC
);
```
- D.

```
CREATE NONCLUSTERED INDEX [IX_Covering_Inf]
(
  [Manager] ASC,
  [PerformanceReviewScore] ASC,
  [UserName] ASC
);
```

www.gre

A. Option AB. Option BC. Option CD. Option D Answer: D QUESTION 28 You plan to execute the following code: You need to identify how many rows will be in dbo.Table1 after you execute the code. How many rows should you identify?

```
CREATE TABLE dbo.Table1
(
  datavalue varchar(100)
);
GO

BEGIN TRANSACTION;

INSERT INTO Table1
(
  datavalue
)
VALUES
(
  'A'
);

BEGIN TRANSACTION;
INSERT INTO Table1
(
  datavalue
)
VALUES
(
  'B'
);
COMMIT TRANSACTION;

INSERT INTO Table1
(
  datavalue
)
VALUES
(
  'C'
);

ROLLBACK TRANSACTION;
Go
```

A. 3B. 2C. 1D. 0 Answer: D QUESTION 29 You have an index for a table in a SQL Azure database. The database is used for Online Transaction Processing (OLTP). You discover that the index consumes more physical disk space than necessary. You need to minimize the amount of disk space that the index consumes. What should you set from the index options? A. STATISTICS_NORECOMPUTE = ONB. STATISTICS_NORECOMPUTE = OFFC. FILLFACTOR = 0D. FILLFACTOR = 80 Answer: C Explanation: <http://msdn.microsoft.com/en-us/library/ms177459.aspx>

<http://msdn.microsoft.com/en-us/library/ms188783.aspx> QUESTION 30 Drag and Drop Question You execute the following code

```
CREATE TABLE Customers
(
  id int
  name nvarchar(50)
);
```

You discover that the Customers table was created in the dbo schema. You need to create a code segment to move the table to another schema named Schema2. What should you create? To answer, drag the appropriate code segments to the correct location in the answer area. (Answer choices may be used once, more than once, or not at all.)

Code Segments	Answer Area			
ALTER SCHEMA	Code	Code	Code	Code
ALTER TABLE				
dbo				
dbo.Customers				
EXEC sp_rename				
TRANSFER				
Schema2				

www.greate

Answer:

Code Segments	Answer Area
<input type="text"/>	ALTER SCHEMA Schema2 TRANSFER dbo.Customers
ALTER TABLE	
dbo	
EXEC sp_rename	
<input type="text"/>	
<input type="text"/>	

www.greatexam.com

Explanation: Note: ALTER SCHEMA newschema TRANSFER oldschema.TABLE This will transfer the table defined under "oldschema" and transfer it to "newschema". QUESTION 31 You plan to design an application that temporarily stores data in a SQL Azure database. You need to identify which types of database objects can be used to store data for the application. The solution must ensure that the application can make changes to the schema of a temporary object during a session. Which type of objects should you identify? A. common table expressions (CTEs) B. table variables C. temporary tables D. temporary stored procedures

Answer: C Explanation: <http://msdn.microsoft.com/en-us/library/ms175972.aspx>

<http://msdn.microsoft.com/en-us/library/ms189084.aspx> <http://msdn.microsoft.com/en-us/library/ms175010.aspx>

<http://msdn.microsoft.com/en-us/library/bb510489.aspx> <http://msdn.microsoft.com/en-us/library/ms187926.aspx>

<http://zacksfiasco.com/post/2010/01/21/SQL-Server-Temporary-Stored-Procedures.aspx> 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
<input type="checkbox"/> read committed	SP1 <input type="checkbox"/> read uncommitted
<input type="checkbox"/> repeatable read	SP2 <input type="checkbox"/> serializable
<input type="checkbox"/> snapshot	

www.greatexam.com

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);
GO

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
```

www.greatexam.com

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.Table1
(
  Column1 int PRIMARY KEY,
  Column2 varchar(50) SPARSE
);
```

www.greatexam.com

A. CREATE INDEX IX_Table1 ON Table1 (Column2) WITH FILLFACTOR=0 B. CREATE INDEX IX_Table1 ON Table1 (Column1) INCLUDE (Column2) C. CREATE INDEX IX_Table1 ON Table1 (Column2) WHERE Column2 IS NULL D. 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> QUESTION 38 You have a SQL Azure database. You need to identify which keyword must be used to create a view that will be indexed. Which keyword should you identify? A. VIEW_METADATA B. SCHEMABINDING C. DEFAULT D. DISTINCT Answer: B
Explanation:

<http://msdn.microsoft.com/en-us/library/ms187956.aspx> <http://msdn.microsoft.com/en-us/library/ms191432.aspx> QUESTION 39

Drag and Drop Question You are designing two stored procedures named Procedure1 and Procedure2. You identify the following requirements: - Procedure1 must take a parameter that ensures that multiple rows of data can pass into the stored procedure. - Procedure2 must use business logic that resides in a Microsoft .NET Framework assembly. You need to identify the appropriate technology for each stored procedure. Which technologies should you identify? To answer, drag the appropriate technology to the correct stored procedure in the answer area. (Answer choices may be used once, more than once, or not at all.)

Technologies

Common language runtime (CLR)
Extensible Markup Language (XML)
a table-valued parameter (TVP)

Answer:

Technologies	Answer Area
	Procedure1 a table-valued parameter (TVP)
Extensible Markup Language (XML)	Procedure2 Common language runtime (CLR)

www.greatexam.com

QUESTION 40 You create a view by using the following code www.greatexam.com

```
CREATE VIEW dbo.View1
WITH VIEW_METADATA
AS
SELECT t1.col1, t1.col2, t2.*
FROM dbo.Table1 AS t1 JOIN dbo.Table2 AS t2 ON t1.col1=t2.col2;
```

Several months after you create the view, users report that the view has started to return unexpected results. You discover that the design of Table2 was modified since you created the view. You need to ensure that the view returns the correct results. Which code segment should you run

- A. EXEC sp_refreshview @viewname = 'dbo.View1';
- B. ALTER dbo.View1 WITH SCHEMABINDING, VIEW_METADATA AS SELECT t1.col1, t1.col2, t2.* FROM dbo.Table1 AS t1 JOIN dbo.Table2 AS t2 ON t1.col1=t2.col2;
- C. EXEC sp_refreshsqlmodule @name = 'dbo.Table2';
- D. DROP dbo.View1; GO CREATE dbo.View1 WITH SCHEMABINDING, VIEW_METADATA AS SELECT t1.col1, t1.col2, t2.* FROM dbo.Table1 AS t1 JOIN dbo.Table2 AS t2 ON t1.col1=t2.col2; www.greatexam.com

A. Option AB. Option BC. Option CD. Option D Answer: A The Microsoft 70-464 questions and answers in PDF on GreatExam are the most reliable study guide for 70-464 exam. Comparing with others', our 70-464 practice test is more authoritative and complete. We provide the latest full version of 70-464 PDF and VCE dumps with new real questions and answers to ensure your 70-464 exam 100% pass. <http://www.greatexam.com/70-464-exam-questions.html>