



70-433^{Q&As}

TS: Microsoft SQL Server 2008, Database Development

Pass Microsoft 70-433 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass4itsure.com/70-433.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





QUESTION 1

You have created an assembly that utilizes unmanaged code to access external resources. You need to deploy the assembly with the appropriate permissions.

Which permission set should you use?

- A. SAFE
- B. UNSAFE
- C. EXTERNAL_ACCESS
- D. Default permission set

Correct Answer: B

QUESTION 2

You administer a Microsoft SQL Server 2008 database named AdventureWorks that contains a table named Production.Product. The table contains a primary key named PK_Product_ProductID and a non-clustered index named AK_Product_ProductNumber. Both indexes have been created on a single primary partition.

The table has the following definition:

```
CREATE TABLE [Production].[Product] (  
    [ProductID] [int] IDENTITY(1,1) NOT NULL,  
    [Name] [nvarchar](50) NOT NULL,  
    [ProductNumber] [nvarchar](25) NOT NULL,  
    [Color] [nvarchar](15) NULL,  
    [Class] [nchar](2) NULL,  
    [Style] [nchar](2) NULL,  
    [ModifiedDate] [datetime] NOT NULL,  
    CONSTRAINT [PK_Product_ProductID] PRIMARY KEY CLUSTERED  
    (  
        [ProductID] ASC  
    ) ON [PRIMARY]  
) ON [PRIMARY]  
GO
```



The index has the following definition: A. Option A

```
CREATE UNIQUE NONCLUSTERED INDEX [AK_Product_ProductNumber] ON [Production].[Product]  
 (  
    [ProductNumber] ASC  
 ) ON [PRIMARY]  
GO
```

The Production.Product table contains 1 million rows.





- A. ALTER DATABASE [AdventureWorks] SET AUTO_UPDATE_STATISTICS ON
- B. UPDATE STATISTICS Production.Product
- C. UPDATE INDEX AK_Product_ProductNumber ON Production.Product SET (STATISTICS_NORECOMPUTE = ON)
- D. ALTER STATISTICS Production.Product
- E. ALTER INDEX AK_Product_ProductNumber ON Production.Product REBUILD
- F. ALTER INDEX AK_Product_ProductNumber ON Production.Product REBUILD Partition = 1
- G. ALTER INDEX AK_Product_ProductNumber ON Production.Product REORGANIZE
- H. EXEC sys.sp_configure 'index create memory', 1
- I. SELECT * FROM STATS WHERE name='AK_Product_ProductNumber'
- J. SELECT * FROM sys.indexes where name=N'Production.Product'
- K. SELECT * FROM SYS.STATS WHERE name='AK_Product_ProductNumber'
- L. SELECT * FROM sys.dm_db_index_physical_stats (DB_ID(), OBJECT_ID(N'Production.Product'),NULL, NULL, NULL)
- M. SELECT * FROM sys.dm_db_index_operational_stats (DB_ID(), OBJECT_ID(N'Production.Product'),NULL, NULL)
- N. CREATE STATISTICS ProductClass_Stats
ON Production.Product (Name, ProductNumber, Class)
WHERE Class <> null
- O. CREATE STATISTICS ProductClass_Stats
ON Production.Product (Name, ProductNumber, Class)
WHERE Class is not null
- P. CREATE STAT ProductClass_Stats
ON Production.Product (Name, ProductNumber, Class)
WHERE Class is not null
WITH SAMPLE 100 PERCENT
- Q. DBCC SHOW_STATISTICS ('Production.Product', AK_Product_ProductNumber)



B. Option B



C. Option C

D. Option D

E. Option E

F. Option F

G. Option G

H. Option H

I. Option I

J. Option J

K. Option K

L. Option L

M. Option M

N. Option N

O. Option O

P. Option P

Q. Option Q

Correct Answer: Q

QUESTION 3

You have tables named Products and OrderDetails. The Products table has a foreign key relationship with the OrderDetails table on the ProductID column. You have the following Transact-SQL batch:

```
BEGIN TRY BEGIN TRANSACTION DELETE FROM Products WHERE ProductID = 5; BEGIN TRANSACTION  
INSERT INTO OrderDetails ( OrderID, ProductID, Quantity ) VALUES ( 1234, 5, 12 ); COMMIT TRANSACTION  
COMMIT TRANSACTION END TRY BEGIN CATCH ROLLBACK TRANSACTION PRINT ERROR_MESSAGE(); END  
CATCH
```

You need to analyze the result of executing this batch.

What should be the expected outcome?

- A. 1. The product will be deleted from the Products table.
- 2. The order details will be inserted into the OrderDetails table.
- B. 1. The product will be deleted from the Products table.
- 2. The order details will not be inserted into the OrderDetails table.
- C. 1. The product will not be deleted from the Products table.



2. The order details will be inserted into the OrderDetails table.

D. 1. The product will not be deleted from the Products table.

2. The order details will not be inserted into the OrderDetails table.

Correct Answer: D

```
ROLLBACK { TRAN | TRANSACTION }
```

```
[ transaction_name | @tran_name_variable
```

```
| savepoint_name | @savepoint_variable ]
```

```
[ ; ]
```

transaction_name

Is the name assigned to the transaction on BEGIN TRANSACTION. When nesting transactions, transaction_name must be the name from the outermost BEGIN TRANSACTION statement.

savepoint_name

Is savepoint_name from a SAVE TRANSACTION statement. Use savepoint_name when a conditional rollback should affect only part of the transaction. ROLLBACK TRANSACTION without a savepoint_name or transaction_name rolls back

to the beginning of the transaction. When nesting transactions, this same statement rolls back all inner transactions to the outermost BEGIN TRANSACTION statement. In both cases, ROLLBACK TRANSACTION decrements the

@@TRANCOUNT system function to 0. ROLLBACK TRANSACTION savepoint_name does not decrement @@TRANCOUNT.

A transaction cannot be rolled back after a COMMIT TRANSACTION statement is executed, except when the COMMIT TRANSACTION is associated with a nested transaction that is contained within the transaction being rolled back. In this

instance, the nested transaction will also be rolled back, even if you have issued a COMMIT TRANSACTION for it.

SQL Server 2008 error handling best practice

```
CREATE PROCEDURE SaveTranExample
```

```
@InputCandidateID INT
```

```
AS
```

```
-- Detect whether the procedure was called from an active transaction and save that for later use. -- In the procedure, @hasOuterTransaction = 0 means there was no active transaction -- and the procedure started one.
```

```
-- @hasOuterTransaction > 0 means an active transaction was started before the -- procedure was called.
```

```
DECLARE @hasOuterTransaction BIT = CASE WHEN @@TRANCOUNT > 0 THEN 1 ELSE END;
```

```
-- Save points need unique names if modules can nest otherwise you can rollback -- to the wrong save point. The solution is to use a GUID to name the save points. DECLARE @rollbackPoint nchar(32) = REPLACE(CONVERT(NCHAR(36),
```



```
NEWID()), N\\'-'', N\\'\\'); IF @hasOuterTransaction > 0
```

```
BEGIN
```

```
-- Procedure called when there is an active transaction. -- Create a savepoint to be able to roll back only the work done  
in the procedure if there is an error.
```

```
SAVE TRANSACTION @rollbackPoint;
```

```
END
```

```
ELSE
```

```
-- Procedure must start its own transaction.
```

```
BEGIN TRANSACTION @rollbackPoint;
```

```
-- Modify database.
```

```
BEGIN TRY
```

```
-- Do work;
```

```
DELETE HumanResources.JobCandidate
```

```
WHERE JobCandidateID = @InputCandidateID;
```

```
-- Get here if no errors; must commit
```

```
-- any transaction started in the
```

```
-- procedure, but not commit a transaction
```

```
-- started before the transaction was called.
```

```
IF @hasOuterTransaction = 0
```

```
BEGIN
```

```
-- @hasOuterTransaction = 0 means no transaction was started before the procedure was called. -- The procedure  
must commit the transaction it started.
```

```
COMMIT TRANSACTION;
```

```
END
```

```
END TRY
```

```
BEGIN CATCH
```

```
-- An error occurred;
```

```
-- If the transaction is still valid
```

```
IF XACT_STATE() = 1
```

```
-- The XACT_STATE function can return the following values:
```



- 1 An open transaction exists that can be either committed or rolled back.
- 0 There is no open transaction.
- 1 An open transaction exists, but it is in a doomed state. Due to the type of error that was raised, the transaction can only be rolled back.

BEGIN

-- Because the syntax for ROLLBACK TRANSACTION is the same for the transaction and for a savepoint --
(ROLLBACK TRANSACTION [transaction_name | @tran_name_variable | savepoint_name | @savepoint_variable]) --
we can write

the following:

```
ROLLBACK TRANSACTION @rollbackPoint;
```

-- In case @rollbackPoint has the name of a transaction, roll back to the beginning of the transaction.

-- In case @rollbackPoint has the name of a savepoint, roll back to the savepoint.

END;

```
ELSE IF XACT_STATE() = -1
```

```
BEGIN
```

```
IF @hasOuterTransaction = 0
```

```
BEGIN
```

```
-- Transaction started in procedure.
```

```
-- Roll back complete transaction.
```

```
ROLLBACK TRANSACTION;
```

```
END
```

```
-- If the transaction is uncommittable, a rollback to the savepoint is not allowed -- because the savepoint rollback writes  
to the log. Just return to the caller, which -- should roll back the outer transaction.
```

```
END
```

```
-- Execute Standard module error handler;
```

```
-- After the appropriate rollback, echo error information to the caller.
```

```
DECLARE @ErrorMessage NVARCHAR(4000);
```

```
DECLARE @ErrorSeverity INT;
```

```
DECLARE @ErrorState INT;
```

```
SELECT @ErrorMessage = ERROR_MESSAGE();
```

```
SELECT @ErrorSeverity = ERROR_SEVERITY();
```



```
SELECT @ErrorState = ERROR_STATE();  
  
RAISERROR (@ErrorMessage, -- Message text.  
  
@ErrorSeverity, -- Severity.  
  
@ErrorState -- State.  
  
);  
  
END CATCH  
  
GO
```

QUESTION 4

You are writing a query that returns a list of products that have grossed more than \$10,000.00 during the year 2007. You need to insert the following filter expression into the query. $SUM([Order Details].UnitPrice * [Order Details].Quantity) > 10000$ Into which clause should you insert this expression?

- A. ON
- B. WHERE
- C. HAVING
- D. GROUP BY

Correct Answer: C

QUESTION 5

You administer a Microsoft SQL Server 2008 database for an inventory management system. The application contains a product table that has the following definition:



```
CREATE TABLE [Production].[Product] (  
    [ProductID] [int] IDENTITY(1,1) NOT NULL,  
    [Name] [nvarchar] (50) NOT NULL,  
    [ProductNumber] [nvarchar] (25) NOT NULL,  
    [Color] [nvarchar] (15) NULL,  
    [Class] [nchar] (2) NULL,  
    [Style] [nchar] (2) NULL,  
    [Active] [bit] NOT NULL,  
    [ModifiedDate] [datetime] NOT NULL,  
    CONSTRAINT [PK_Product_ProductID] PRIMARY KEY CLUSTERED  
    (  
        [ProductID] ASC  
    ) ON [PRIMARY]  
) ON [PRIMARY]  
GO
```



You want to add a new field to the Product table to meet the following requirements:

Allows user-specified information that will be added to records in the Product table. Supports the largest storage size needed for the field. Uses the smallest data type necessary to support the domain of values that will be entered by users.

You need to add a field named User_Data_1 to support integer values ranging from -10 through 10.

Which SQL statement should you use?

- A. ALTER TABLE [Production].[Product] ADD [User_Data_1] TINYINT
- B. ALTER TABLE [Production].[Product] ADD [User_Data_1] SMALLINT
- C. ALTER TABLE [Production].[Product] ADD [User_Data_1] INT
- D. ALTER TABLE [Production].[Product] ADD [User_Data_1] BIGINT
- E. ALTER TABLE [Production].[Product] ADD [User_Data_1] BIT
- F. ALTER TABLE [Production].[Product] ADD [User_Data_1] NUMERIC(11,6)
- G. ALTER TABLE [Production].[Product] ADD [User_Data_1] NUMERIC(6,11)
- H. ALTER TABLE [Production].[Product] ADD [User_Data_1] NUMERIC(5,6)
- I. ALTER TABLE [Production].[Product] ADD [User_Data_1] SMALLMONEY
- J. ALTER TABLE [Production].[Product] ADD [User_Data_1] MONEY
- K. ALTER TABLE [Production].[Product] ADD [User_Data_1] CHAR(100)
- L. ALTER TABLE [Production].[Product] ADD [User_Data_1] VARCHAR(100)
- M. ALTER TABLE [Production].[Product] ADD [User_Data_1] NCHAR(100)
- N. ALTER TABLE [Production].[Product] ADD [User_Data_1] NVARCHAR(100)
- O. ALTER TABLE [Production].[Product] ADD [User_Data_1] SMALLDATETIME



P. ALTER TABLE [Production].[Product] ADD [User_Data_1] DATETIME

Q. ALTER TABLE [Production].[Product] ADD [User_Data_1] DATETIME2

R. ALTER TABLE [Production].[Product] ADD [User_Data_1] DATE

Correct Answer: B

QUESTION 6

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.

Correct Answer: CD

As a database is changed by such actions as adding indexes or changing data in indexed columns, the original query plans used to access its tables should be optimized again by recompiling them. This optimization happens automatically the first time a stored procedure is run after Microsoft SQL Server is restarted. It also occurs if an underlying table used by the stored procedure changes. But if a new index is added from which the stored procedure might benefit, optimization does not happen until the next time the stored procedure is run after SQL Server is restarted. In this situation, it can be useful to force the stored procedure to recompile the next time it executes. SQL Server provides three ways to force a stored procedure to recompile: The sp_recompile system stored procedure forces a recompile of a stored procedure the next time that it is run. It does this by deleting the existing plan from the procedure cache forcing a new plan to be created the next time that the procedure is run. Creating a stored procedure that specifies the WITH RECOMPILE option in its definition indicates that SQL Server does not cache a plan for this stored procedure; the stored procedure is recompiled every time that it is executed. Use of this option is uncommon and causes the stored procedure to execute more slowly, because the stored procedure must be recompiled every time that it is executed. You can force the stored procedure to be recompiled by specifying the WITH RECOMPILE option when you execute the stored procedure. A. CREATE and ALTER PROCEDURE syntax for CLR Stored Procedure does not have RECOMPILE option. B. The RECOMPILE option is ignored for procedures created with FOR REPLICATION. C. ENCRYPTION option and RECOMPILE option can go together.

QUESTION 7

You are creating a function that references a table. You need to prevent the table from being dropped.

Which option should you use when you create the function?

- A. WITH ENCRYPTION
- B. WITH EXECUTE AS
- C. WITH SCHEMABINDING



D. WITH RETURNS NULL ON NULL INPUT

Correct Answer: C

QUESTION 8

You have a table named Stores that has an XML column named OpenHours. This column contains the opening and closing times.

...

You need to write a query that returns a list of stores and their opening time for a specified day.

Which code segment should you use?

- A. DECLARE @Day VARCHAR(10) = 'Tuesday' SELECT StoreName, OpenHours.value('\/hours[1]\/@open', 'time') FROM Stores WHERE OpenHours.value('\/hours[1]\/@dayofweek', 'varchar(20)') = @Day
- B. DECLARE @Day VARCHAR(10) = 'Tuesday' SELECT StoreName, OpenHours.value('\/hours[1]\/@open', 'time') FROM Stores WHERE OpenHours.exist('\/hours[@dayofweek=sql:variable("@Day")]') = 1
- C. DECLARE @Day VARCHAR(10) = 'Tuesday' SELECT Storename, OpenHours.query('\/data\/hours[@dayofweek=sql:variable("@Day")]\/@open') FROM Stores
- D. DECLARE @Day VARCHAR(10) = 'Tuesday' SELECT StoreName, OpenHours.value('\/hours[1]\/@dayofweek=sql:variable("@Day")]\/@open', 'time') FROM Stores

Correct Answer: C

```
CREATE TABLE Stores(
StoreName VARCHAR(10)NOT NULL,
OpenHours [xml] NULL,
CONSTRAINT [PK_Stores] PRIMARY KEY CLUSTERED (StoreName)) GO
INSERT INTO Stores (StoreName, OpenHours)
VALUES
(
'Store1',
' '), ('Store2',
' ')
DECLARE @Day VARCHAR
(10) = 'Tuesday' SELECT Storename, OpenHours.query('\/data\/hours[@dayofweek=sql:variable("@Day")]\/@open') FROM Stores GO
```

**QUESTION 9**

You are using Microsoft SQL Server 2008 Enterprise Edition. You need to maintain a history of all data modifications made to a table, including the type of modification and the values modified.

Which tracking method should you use?

- A. Database Audit
- B. Change Tracking
- C. C2 Audit Tracing
- D. Change Data Capture

Correct Answer: D

Database Audit Auditing an instance of the SQL Server Database Engine or an individual database involves tracking and logging events that occur on the Database Engine. SQL Server audit lets you create server audits, which can contain server audit specifications for server level events, and database audit specifications for database level events. Audited events can be written to the event logs or to audit files.

Change Data Capture Change data capture provides historical change information for a user table by capturing both the fact that DML changes were made and the actual data that was changed. Changes are captured by using an asynchronous process that reads the transaction log and has a low impact on the system.

Change Tracking Change tracking captures the fact that rows in a table were changed, but does not capture the data that was changed. This enables applications to determine the rows that have changed with the latest row data being obtained directly from the user tables. Therefore, change tracking is more limited in the historical questions it can answer compared to change data capture. However, for those applications that do not require the historical information, there is far less storage overhead because of the changed data not being captured.

Topic 3, Volume C

QUESTION 10

You are using TRY...CATCH error handling.

You need to raise an error that will pass control to the CATCH block.

Which severity level should you use?

- A. 0
- B. 9
- C. 10
- D. 16

Correct Answer: D

A TRY...CATCH construct catches all execution errors that have a severity higher than 10 that do not close the



database connection. Severity levels from 20 through 25 are considered fatal. If a fatal severity level is encountered, the client

connection is terminated after receiving the message, and the error is logged in the error and application logs.

Severity levels less than 0 are interpreted as 0. Severity levels greater than 25 are interpreted as 25.

QUESTION 11

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

Correct Answer: D

SQL Server Collation Name consists of several parts, one of them is responsible for CaseSensitivity CI specifies case-insensitive, CS specifies case-sensitive.

BIN specifies the binary sort order to be used.

So, because we want case-sensitive location, B and C are not suitable. Latin1_General_BIN use binary sort order, but we want linguistical sort order (according to rules of language), not based on the code point values of characters.

QUESTION 12

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;

Correct Answer: B



Fill Factor

The fill-factor option is provided for fine-tuning index data storage and performance. When an index is created or rebuilt, the fill-factor value determines the percentage of space on each leaf-level page to be filled with data, reserving the

remainder on each page as free space for future growth. For example, specifying a fill-factor value of 60 means that 40 percent of each leaf-level page will be left empty, providing space for index expansion as data is added to the underlying

table.

Page Splits

A correctly chosen fill-factor value can reduce potential page splits by providing enough space for index expansion as data is added to the underlying table. When a new row is added to a full index page, the Database Engine moves

approximately half the rows to a new page to make room for the new row. This reorganization is known as a page split. A page split makes room for new records, but can take time to perform and is a resource intensive operation. Also, it can

cause fragmentation that causes increased I/O operations.

Fragmentation

Fragmentation breaks down to physical and logical fragmentation (or, internal and external fragmentation).

Physical/Internal fragmentation is caused when there is wasted space in the index caused by page splits, deletes, FILLFACTOR and

PAD_INDEX. Logical/ External fragmentation is when the pages that make up the leaf levels of the index are not in good order. This is usually caused by page splits. Both cause performance problems. Internal fragmentation causes problems

because the indexes get bigger and bigger requiring more and more pages to store the data, which means that reads from the index slow things down. External fragmentation causes problems because when the data needs to be read from

the index it has to skip all over the place following the links between pages, again, slowing things down.

The SQL Server Database Engine automatically maintains indexes whenever insert, update, or delete operations are made to the underlying data. Over time these modifications can cause the information in the index to become scattered in

the database (fragmented). Fragmentation exists when indexes have pages in which the logical ordering, based on the key value, does not match the physical ordering inside the data file. You can remedy index fragmentation by either

reorganizing an index or by rebuilding an index. To reorganize one or more indexes, use the ALTER INDEX statement with the REORGANIZE clause.

Reorganizing an index defragments the leaf level of clustered and nonclustered indexes on tables and views by physically reordering the leaf-level pages to match the logical order (left to right) of the leaf nodes. Having the pages in order

improves index-scanning performance. The index is reorganized within the existing pages allocated to it; no new pages are allocated.

Rebuilding an index drops the index and creates a new one. In doing this, fragmentation is removed, disk space is reclaimed by compacting the pages using the specified or existing fill factor setting, and the index rows are reordered in



contiguous pages (allocating new pages as needed). This can improve disk performance by reducing the number of page reads required to obtain the requested data.

You can use the CREATE INDEX with the DROP_EXISTING clause or ALTER INDEX with the REBUILD clause statements to set the fill-factor value for individual indexes. EXEC sys.sp_configure 'fill factor (%)', '60'; will modify the server

default fill factor value.

[70-433 VCE Dumps](#)

[70-433 Exam Questions](#)

[70-433 Braindumps](#)



To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

Try our product !

100% Guaranteed Success

100% Money Back Guarantee

365 Days Free Update

Instant Download After Purchase

24x7 Customer Support

Average 99.9% Success Rate

More than 800,000 Satisfied Customers Worldwide

Multi-Platform capabilities - [Windows](#), [Mac](#), [Android](#), [iPhone](#), [iPod](#), [iPad](#), [Kindle](#)

We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications. You can view Vendor list of All Certification Exams offered:

<https://www.pass4itsure.com/allproducts>

Need Help

Please provide as much detail as possible so we can best assist you.

To update a previously submitted ticket:



 <p>One Year Free Update Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p>	 <p>Money Back Guarantee To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>	 <p>Security & Privacy We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.</p>
---	---	--

Any charges made through this site will appear as Global Simulators Limited.

All trademarks are the property of their respective owners.

Copyright © pass4itsure, All Rights Reserved.