

[2016-New Free Share GreatExam Microsoft 70-465 VCE Dumps With New Update Exam Questions (1-20)

How to 100% pass 70-465 exam? GreatExam provides the guaranteed 70-465 exam preparation material to boost up your confidence in 70-465 exam. Successful candidates have provided their reviews about our 70-465 dumps. Now GreatExam supplying the new version of 70-465 VCE and PDF dumps. We ensure our 70-465 exam questions are the most complete and authoritative compared with others', which will ensure your 70-465 exam pass. QUESTION 1 Drag and Drop Question You need to recommend a backup process for an Online Transaction Processing (OLTP) database. The process must meet the following requirements:- Ensure that if a hardware failure occurs, you can bring the database online with a minimum amount of data loss.- Minimize the amount of administrative effort required to restore any lost data. What should you include in the recommendation? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

- Create a database snapshot.
- Perform a differential backup of the database every night.
- Ship the logs to a secondary server.
- Perform a backup of the transaction log every hour.
- Set the database to the simple recovery model.
- Set the database to the full recovery model.
- Perform a weekly full backup of the database.

Answer

Create a database snapshot.	Set the database to the full recovery model.
Perform a weekly full backup of the database.	Perform a differential backup of the database every night.
Ship the logs to a secondary server.	Perform a backup of the transaction log every hour.
Set the database to the simple recovery model.	

www.greatexam.com

QUESTION 2 You use SQL Server 2014 Enterprise Edition. Your database contains a partitioned table named AuditData. AuditData is partitioned by year. Partition 1 contains data from the year 2010 and prior. Management has decided to archive all AUDITDATA records from 2010 and prior. Management wants the records to be removed from the database entirely and provided to the backup team as a zipped text file. The data must no longer reside in the database. There is very little tolerance for performance degradation in your environment. You need to remove all 2010 and prior data from the AuditData table by using the least amount of system resources possible. Develop the solution by selecting and arranging the required SQL actions in the correct order. You may not need all of the actions

SQL Actions	Answer Area
Drop Table	
Select Into	
Switch Partition	
Move Partition	
Merge Range	
BCP	
Split Range	
Create Table	
Delete Partition	
Drop Partition	

www.greatexam.com

Answer:

SQL Actions	Answer Area
Drop Table	Create Table
Select Into	Split Range
Switch Partition	Select Into
Move Partition	BCP
Merge Range	Drop Table
BCP	Delete Partition
Split Range	
Create Table	
Delete Partition	
Drop Partition	

www.greatexam.com

Explanation:Note:- Create a new partitioned table with the partition function you want, and then insert the data from the old table into the new table by using an INSERT INTO...SELECT FROM statement.- SPLIT RANGE (boundary_value)Adds one partition to the partition function. boundary_value determines the range of the new partition, and must differ from the existing boundary ranges of the partition function. Based on boundary_value, the Database Engine splits one of the existing ranges into two.Of these two, the one where the new boundary_value resides is considered the new partition.- BCP can be used top produce the zipped text file.- Example:plitting a partition of a partitioned table or index into two partitionsThe following example creates a partition function to partition a table or index into four partitions. ALTER PARTITION FUNCTION splits one of the partitions into two to create a total of five partitions.CREATE PARTITION FUNCTION myRangePF1 (int)AS RANGE LEFT FOR VALUES (1, 100, 1000); GO--Split the partition between boundary_values 100 and 1000--to create two partitions between boundary_values 100 and 500 --and between boundary_values 500 and 1000.ALTER PARTITION FUNCTION myRangePF1 ()SPLIT RANGE (500);

QUESTION 3Drag and Drop QuestionYou are planning to deploy a database to Windows Azure SQL Database.You need to design a stored procedure to update rows.The stored procedure must meet the following requirements:- If more than one row is updated, an error must be raised to the application and the update must be discarded.- The stored procedure must be designed to maximize concurrency.What should you include in the design? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order

- Raise an error and roll back the transaction if the row count is greater than 1.
- Issue a SELECT statement to count the number of rows.
- Perform the update.
- Set the isolation level to serializable.
- Read the @@ROWCOUNT system variable.
- Begin an explicit transaction.
- Commit the transaction.

www.greatexam.com

Answer

Actions	Work Area
Raise an error in a catch block.	Begin an explicit transaction.
Commit the transaction in a finally block.	Perform the update in a try block.
Read the @@ROWCOUNT system variable.	Read the @@ROWCOUNT system variable.
Perform the update in a try block.	Raise an error and roll back the transaction if the row count is less than 1.
Raise an error and roll back the transaction if the row count is less than 1.	Commit the transaction in a finally block.
Issue a SELECT statement to count the number of rows.	
Set the isolation level to serializable.	
Begin an explicit transaction.	

www.greatexam.com

Explanation:Box 1: Begin an explicit transaction.Box 2: Perform an update in a try block.Box 3: Read the @@ROWCOUNT system

variable.Box 4: Raise an error and roll back the transaction if the row count is less than 1.Box 5: Commit the transaction in a finally block.Note:- Read Committed is SQL Server's default isolation level.- @@ROWCOUNT returns the number of rows affected by the last statement.- Using TRY...CATCH in a transactionThe following example shows how a TRY...CATCH block works inside a transaction.The statement inside the TRY block generates a constraint violation error.- BEGIN TRANSACTION;- BEGIN TRY-- Generate a constraint violation error.DELETE FROM Production.ProductWHERE ProductID = 980;END TRYBEGIN CATCH SELECTERROR_NUMBER() AS ErrorNumber,ERROR_SEVERITY() AS ErrorSeverity,ERROR_STATE() AS ErrorState ,ERROR_PROCEDURE() AS ErrorProcedure,ERROR_LINE() AS ErrorLine,ERROR_MESSAGE() AS ErrorMessage;IF @@TRANCOUNT > 0ROLLBACK TRANSACTION;END CATCH;IF @@TRANCOUNT > 0COMMIT TRANSACTION;GO

QUESTION 4You are designing a monitoring application for a new SQL Server 2014 instance.You need to recommend a solution to generate a report that displays the 10 most frequent wait types that occur for the instance.What should you include in the recommendation? More than one answer choice may achieve the goal. Select the BEST answer. A. The SQL Server error logB. The sys.dm_os_wait_stats dynamic management viewC. The DBCC SQLPERF(WAITSTATS) commandD. SQL Server Profiler

Answer: BExplanation:sys.dm_os_wait_statsReturns information about all the waits encountered by threads that executed.You can use this aggregated view to diagnose performance issues with SQL Server and also with specific queries and batches.Columns include:waiting_tasks_countNumber of waits on this wait type.This counter is incremented at the start of each wait.

QUESTION 5 Drag and Drop QuestionYou have two servers named SQL1 and SQL2 that have SQL Server 2012 installed.SQL1 contains a database that is mirrored asynchronously to SQL2.The database contents are updated once a month.You need to upgrade the database to SQL Server 2014.The solution must minimize downtime.Which upgrade steps should you recommend? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order

- Fail over.
- Fail back.
- Upgrade SQL1.
- Upgrade SQL2.
- Establish a mirror.
- Break the mirror.

Answer

Fail over.	Break the mirror.
Fail back.	Upgrade SQL1.
	Upgrade SQL2.
	Establish a mirror.

www.greatexam.com

Explanation:Box 1: Fail over.Box 2: Upgrade SQL1.Box 3: Upgrade SQL2.Box 4: Establish a mirror.Note:- To perform the rolling upgradeStep 1: For each mirroring session whose mirror server instance has just been upgraded, wait for the session to synchronize. Then, connect to the principal server instance, and manually fail over the session.Step 2: Upgrade each server instance that is now the mirror server in all mirroring sessions in which it is a partner.Step 3: Resume the mirroring session.- When upgrading server instances to SQL Server 2014, you can reduce downtime for each mirrored database to only a single manual failover by performing a sequential upgrade, known as a rolling upgrade. A rolling upgrade is a multi-stage process that in its simplest form involves upgrading the server instance that is currently acting as the mirror server in a mirroring session, then manually failing over the mirrored database, upgrading the former principal server, and resuming mirroringReference: Minimize Downtime for Mirrored Databases When Upgrading Server Instances QUESTION 6You deploy a database by using SQL Server 2012.The database contains a table named Table1.You need to recommend a solution that meets the following requirements:- Stores the most recent data from Table1 by using the fastest storage solution possible.- Stores the historical data from Table1 by using a slower storage

solution. What should you recommend? A. partitioned views B. a database snapshot C. change data capture D. table partitioning
 Answer: D
QUESTION 7 Drag and Drop Question You plan to deploy SQL Server 2014. Your company identifies the following monitoring requirements for the database:- An e-mail message must be sent when a user logs in.- An e-mail message must be sent if CPU utilization exceeds 90 percent. You need to identify which feature meets each monitoring requirement. Which features should you identify? To answer, drag the appropriate feature to the correct monitoring requirement in the answer area

- Policy-Based Management
- a SQL Server Agent alert
- SQL Server Integration Services (SSIS)
- trace flags
- triggers

Answer

Policy-Based Management	An e-mail message must be sent when a user logs in.	Policy-Based Management
a SQL Server Agent alert	An e-mail message must be sent if CPU utilization exceeds 90 percent.	a SQL Server Agent alert
SQL Server Integration Services (SSIS)		
trace flags		
triggers		

www.greatexam.com

QUESTION 8 You are creating a database that will store usernames and passwords for an application. You need to recommend a solution to store the passwords in the database. What should you recommend? More than one answer choice may achieve the goal. Select the BEST answer. A. One-way encryption B. Transparent Data Encryption (TDE) C. Encrypting File System (EFS) D. Reversible encryption
 Answer: C
 Explanation:- Transparent Data Encryption (TDE) is a special case of encryption using a symmetric key. TDE encrypts an entire database using that symmetric key called the database encryption key. The database encryption key is protected by other keys or certificates which are protected either by the database master key or by an asymmetric key stored in an EKM module.- SQL Server provides the following mechanisms for encryption:- Transact-SQL functions- Asymmetric keys- Symmetric keys- Certificates- Transparent Data Encryption
QUESTION 9 Drag and Drop Question You have a SQL Azure database named Database1. You need to design the schema for a table named table1. Table1 will have less than one million rows. Table1 will contain the following information for each row

Column	Description
ID	An incremental numeric value used to identify the row
Name	A string in English
Code	An alphanumeric code that has five characters
ModifiedDate	The date of the last modification

www.greatexam.com

The solution must minimize the amount of space used to store each row. Which data types should you recommend for each column?
 To answer, drag the appropriate data type to the correct column in the answer area

Data Types		Answer Area
int	ID	Data type
bigint	Name	Data type
varchar	Code	Data type
nvarchar	ModifiedDate	Data type
char		
smalldatetime		
date		

www.greatexam.com

Answer

Data Types	Answer Area
<input type="text"/>	ID <input type="text" value="int"/>
<input type="text" value="bigint"/>	Name <input type="text" value="varchar"/>
<input type="text" value="nvarchar"/>	Code <input type="text" value="char"/>
<input type="text" value="smalldatetime"/>	ModifiedDate <input type="text" value="date"/>

www.greatexam.com

Explanation: <http://msdn.microsoft.com/en-US/library/ms187752.aspx> QUESTION 10 Drag and Drop Question You need to recommend the actions that are required to partition a table. In which order should the four actions be performed? To answer, move the actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Answer Area
<input type="text" value="Create a partition scheme."/>	
<input type="text" value="Create a partition function."/>	
<input type="text" value="Create filegroups."/>	
<input type="text" value="Create the table."/>	

www.gre

Answer

Actions	Answer Area
<input type="text" value="Create a partition scheme."/>	<input type="text" value="Create filegroups."/>
<input type="text" value="Create a partition function."/>	<input type="text" value="Create a partition function."/>
<input type="text" value="Create filegroups."/>	<input type="text" value="Create a partition scheme."/>
<input type="text" value="Create the table."/>	<input type="text" value="Create the table."/>

www.greatexam.com

Explanation: <http://technet.microsoft.com/en-us/library/ms188730.aspx> QUESTION 11 You have a SQL Server 2014 environment that contains 20 servers. The corporate security policy states that all SQL Server 2014 instances must meet specific security standards. You need to recommend a management strategy for the SQL Server 2014 servers. What should you include in the recommendation? More than one answer choice may achieve the goal. Select the BEST answer. A. Multi server jobs B. Policy-Based Management C. Common criteria compliance D. Maintenance plans Answer: B Explanation: Policy-Based Management is a system for managing one or more instances of SQL Server. When SQL Server policy administrators use Policy-Based Management, they use SQL Server Management Studio to create policies to manage entities on the server, such as the instance of SQL Server, databases, or other SQL Server objects. QUESTION 12 Drag and Drop Question You are designing a database for a university. The database will contain two tables named Classes and Enrollment that have the following specifications:- Classes will store brochures in the XPS format. The brochures must be structured in folders and must be accessible by using UNC paths.- Enrollment will store information about students and their classes. Performance must be enhanced for queries of the current enrollments. You need to identify which SQL Server technology meets the specifications of each table. Which technologies should you identify? To answer, drag the appropriate technology to the correct table in the answer area

Technologies	Answer Area
<input type="text" value="FileStream"/>	Classes <input type="text"/>
<input type="text" value="FileTable"/>	Enrollment <input type="text"/>
<input type="text" value="Partitioned tables"/>	
<input type="text" value="Partitioned views"/>	

www.greatex

Answer

Explanation:- The FileTable feature brings support for the Windows file namespace and compatibility with Windows applications to the file data stored in SQL Server. FileTable lets an application integrate its storage and data management components, and provides integrated SQL Server services-including full-text search and semantic search-over unstructured data and metadata. In other words, you can store files and documents in special tables in SQL Server called FileTables, but access them from Windows applications as if they were stored in the file system, without making any changes to your client applications. QUESTION 13 You deploy a database by using SQL Server 2014. The database contains a table named Table1. You need to recommend a solution to track all of the deletions executed on Table1. The solution must minimize the amount of custom code required. What should you recommend?

A. Change data capture B. Statistics C. A trigger D. Master Data Services Answer: A Explanation: Change data capture is designed to capture insert, update, and delete activity applied to SQL Server tables, and to make the details of the changes available in an easily consumed relational format. The change tables used by change data capture contain columns that mirror the column structure of a tracked source table, along with the metadata needed to understand the changes that have occurred.

QUESTION 14 You have four databases that are accessed by using an Online Transaction Processing (OLTP) application. The databases are stored on a server named SQL1 that has SQL Server 2014 installed. You plan to deploy an additional server that has SQL Server 2014 installed. You need to design a high-availability solution for the databases that meets the following requirements:- If SQL1 fails, the databases must be available.- Users must be able to run reports against a secondary copy of the databases. What should you include in the design? More than one answer choice may achieve the goal. Select the BEST answer. A. AlwaysOn availability groups B. Database mirroring C. Log shipping D. Failover Clustering Answer: A Explanation: The AlwaysOn Availability Groups feature is a high-availability and disaster-recovery solution that provides an enterprise-level alternative to database mirroring. Introduced in SQL Server 2012, AlwaysOn Availability Groups maximizes the availability of a set of user databases for an enterprise. An availability group supports a failover environment for a discrete set of user databases, known as availability databases, that fail over together.

QUESTION 15 You are troubleshooting an application that runs a query. The application frequently causes deadlocks. You need to identify which transaction causes the deadlock. What should you do? More than one answer choice may achieve the goal. Select the BEST answer. A. Query the sys.dm_exec_requests dynamic management view. B. Create a trace in SQL Server Profiler that contains the Deadlock graph event. C. Query the sys.dm_exec_sessions dynamic management view. D. Create an extended events session to capture deadlock information. Answer: D Explanation: Troubleshooting deadlocks You have been receiving reports from users indicating that certain applications are returning deadlock errors. To maximize the effectiveness of troubleshooting these problems, you decide to focus on the deadlocks that are hit most frequently. You create an Extended Events session that:- Configures deadlock event tracking for the session.- Specifies a target that aggregates based on an identifier for the deadlock. You run the Extended Events session, and after additional deadlocks are reported you are able to obtain aggregated deadlock information, along with the complete XML deadlock graph for each source. Using this information you are able to pin point the most common deadlocks and start working on a solution.

QUESTION 16 Drag and Drop Question You have a server named SQL1 that has SQL Server 2012 installed. SQL1 hosts a database named Database1. Database1 contains a table named Table1. Table1 is partitioned across five filegroups based on the Date field. The schema of Table1 is configured as shown in the following table

Column	Data type
ID	Bigint
Account	Bigint
Amount	Decimal
TransactionType	Int
TransactionDate	Date

Table1 contains the indexes shown in the following table

Index	Type	Columns
PK_Table1	Clustered, primary key	ID, TransactionDate
IX_Account	Nonclustered	Account
IX_Type	Nonclustered	TransactionType
IX_Date	Nonclustered	TransactionDate
IX_Amount	Nonclustered	Amount

You need to recommend an index strategy to maximize performance for the queries that consume the indexes available to Table1. Which type of index storage should you recommend? To answer, drag the appropriate index storage type to the correct index in the answer area

Index Storage Types

Aligned

Nonaligned

Answer Area

IX_Type

Index Storage Type

IX_Account

Index Storage Type

IX_DATE

Index Storage Type

IX_Amount

Index Storage Type

Answer

Index Storage Types

Aligned

Nonaligned

Answer Area

IX_Type

Aligned

IX_Account

Nonaligned

IX_DATE

Nonaligned

IX_Amount

Nonaligned

Explanation: Index Storage Type- Designing a partitioned index independently (unaligned) of the base table can be useful in the following cases:- The base table has not been partitioned.- The index key is unique and it does not contain the partitioning column of the table.- You want the base table to participate in collocated joins with more tables using different join columns. QUESTION 17 Drag and Drop Question You plan to deploy SQL Server 2014. Your company identifies the following monitoring requirements:- Tempdb must be monitored for insufficient free space.- Deadlocks must be analyzed by using Deadlock graphs. You need to identify which feature meets each monitoring requirement. Which features should you identify? To answer, drag the appropriate feature to the correct monitoring requirement in the answer area

Features

dynamic management view

Activity Monitor

Resource Governor

SQL Trace

Answer Area

Tempdb must be monitored for insufficient free space.

Feature

Deadlocks must be analyzed by using Deadlock graphs.

Feature

Answer:

Output as PDF file has been powered by [[Universal Post Manager](#)] plugin from [www.ProfProjects.com](#)

| Page 7/9 |

Features	Answer Area
dynamic management view	Tempdb must be monitored for insufficient free space. Deadlocks must be analyzed by using Deadlock graphs.
Activity Monitor	
Resource Governor	
SQL Trace	

dynamic management view
 SQL Trace

www.greatexam.com

Explanation:- You can use the sys.dm_db_file_space_usage dynamic management view to monitor the disk space used by the user objects, internal objects, and version stores in the tempdb files. Additionally, to monitor the page allocation or deallocation activity in tempdb at the session or task level, you can use the sys.dm_db_session_space_usage and sys.dm_db_task_space_usage dynamic management views. These views can be used to identify large queries, temporary tables, or table variables that are using a large amount of tempdb disk space. - Use SQL Server Profiler to identify the cause of a deadlock. A deadlock occurs when there is a cyclic dependency between two or more threads, or processes, for some set of resources within SQL Server. Using SQL Server Profiler, you can create a trace that records, replays, and displays deadlock events for analysis. - SQL Server Profiler and SQL Server Management Studio use a deadlock wait-for graph to describe a deadlock. The deadlock wait-for graph contains process nodes, resource nodes, and edges representing the relationships between the processes and the resources. Reference: Troubleshooting Insufficient Disk Space in tempdb Reference: Analyze Deadlocks with SQL Server Profiler

QUESTION 18 You plan to create a database. The database will be used by a Microsoft .NET application for a special event that will last for two days. During the event, data must be highly available. After the event, the database will be deleted. You need to recommend a solution to implement the database while minimizing costs. The solution must not affect any existing applications. What should you recommend? More than one answer choice may achieve the goal. Select the BEST answer.

A. SQL Server 2014 Enterprise
 B. SQL Server 2014 Standard
 C. SQL Azure
 D. SQL Server 2014 Express with Advanced Services

Answer: B

Explanation: Programmability (AMO, ADOMD.Net, OLEDB, XML/A, ASSL) supported by Standard and Enterprise editions only. Reference: Features Supported by the Editions of SQL Server 2014

QUESTION 19 You are designing a database named DB1. Changes will be deployed to DB1 every Wednesday night. You need to recommend a strategy to deploy the changes to DB1. The strategy must meet the following requirements: - The strategy must not disrupt backup operations. - DB1 must be unavailable to users while the changes are deployed. - You must be able to undo quickly the entire operation. What should you recommend? More than one answer choice may achieve the goal. Select the BEST answer.

A. Perform a copy-only database backup before the changes are deployed. If the deployment fails, restore the database to another server and recover the original objects from the restored database.
 B. Create a database snapshot. If the deployment fails, recover the objects from the database snapshot.
 C. Create a database snapshot. If the deployment fails, revert the database to the database snapshot.
 D. Perform a full database backup before the changes are deployed. If the deployment fails, restore the database to another server and recover the original objects from the restored database.

Answer: C

QUESTION 20 Drag and Drop Question You plan to deploy SQL Server 2014. You are designing two stored procedures named SP1 and SP2 that have the following requirements: - Prevent data read by SP1 from being modified by other active processes. - Prevent SP2 from performing dirty reads. You need to recommend the isolation level for each stored procedure. The solution must maximize concurrency. Which isolation levels should you recommend? To answer, drag the appropriate isolation level to the correct stored procedure in the answer area

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

www.greatexam.com

Answer	Isolation Levels	Answer Area
	<input type="checkbox"/> read uncommitted	SP1 <input type="checkbox"/> repeatable read
	<input type="checkbox"/> serializable	SP2 <input type="checkbox"/> read committed

www.greatexam.com

Explanation:- REPEATABLE READ This isolation level includes the guarantees given by SNAPSHOT isolation level. In addition, REPEATABLE READ guarantees that for any row that is read by the transaction, at the time the transaction commits the row has not been changed by any other transaction. Every read operation in the transaction is repeatable up to the end of the transaction. - Committed Read is SQL Server's default isolation level. It ensures that an operation will never read data another application has changed but not yet committed. The Microsoft 70-465 questions and answers from GreatExam are the most reliable guide for Microsoft exam. We offer the latest 70-465 PDF and VCE dumps with new version VCE player for free download, and the newest 70-465 practice test ensures your exam 100% pass. A large number of successful candidates have shown a lot of faith in our 70-465 exam dumps. If you want pass the Microsoft 70-465 exam, please choose GreatExam.

<http://www.greatexam.com/70-465-exam-questions.html>