

Microsoft Official Exam Center New Changed 70-469 Questions In PDF & VCE Format Offered By Braindump2go For Free Download (81-90)

We never believe in second chances and Braindump2go brings you the best 70-469 Exam Preparation Materials which will make you pass in the first attempt. We guarantee all questions and answers in our 70-469 Dumps are the latest released, we check all exam dumps questions from time to time according to Microsoft Official Center, in order to guarantee you can read the latest questions! Vendor: Microsoft Exam Code: 70-469 Exam Name: Recertification for MCSE: Data Platform Keywords: 70-469 Exam Dumps, 70-469 Practice Tests, 70-469 Practice Exams, 70-469 Exam Questions, 70-469 PDF, 70-469 VCE Free, 70-469 Book, 70-469 E-Book, 70-469 Study Guide, 70-469 Braindump, 70-469 Prep Guide

Compared Before Buying Microsoft 70-469 PDF & VCE!		
Pass4sure	Braindump2go	Test King
	100% Pass OR Money Back	
281 Q&As – Practice	292 Q&As – Real Questions	281 Q&As – Practice
\$124.99	\$99.99	\$171.32
No Discount	Coupon Code: BDNT2014	No Discount

QUESTION 81 You need to recommend a disk monitoring solution that meets the business requirements. What should you include in the recommendation? A. An audit B. A dynamic management view C. A maintenance plan D. A SQL Server Agent alert
 Answer: B Explanation: Dynamic Management Views and Functions (Transact-SQL)

QUESTION 82 You need to recommend a solution to improve the performance of usp_UpdateInventory. The solution must minimize the amount of development effort. What should you include in the recommendation? A. A table variable B. A common table expression C. A subquery D. A cursor
 Answer: A Explanation: * Scenario: Database2 will contain a stored procedure named usp_UpdateInventory. Usp_UpdateInventory will manipulate a table that contains a self-join that has an unlimited number of hierarchies. * A table variable can be very useful to store temporary data and return the data in the table format. table * Example: The following example uses a self-join to find the products that are supplied by more than one vendor. Because this query involves a join of the ProductVendor table with itself, the ProductVendor table appears in two roles. To distinguish these roles, you must give the ProductVendor table two different aliases (pv1 and pv2) in the FROM clause. These aliases are used to qualify the column names in the rest of the query. This is an example of the self-join Transact-SQL statement: USE AdventureWorks2008R2; GO SELECT DISTINCT pv1.ProductID, pv1.VendorID FROM Purchasing.ProductVendor pv1 INNER JOIN Purchasing.ProductVendor pv2 ON pv1.ProductID = pv2.ProductID AND pv1.VendorID <> pv2.VendorID ORDER BY pv1.ProductID
 Incorrect: Not B: Using a CTE offers the advantages of improved readability and ease in maintenance of complex queries. The query can be divided into separate, simple, logical building blocks. These simple blocks can then be used to build more complex, interim CTEs until the final result set is generated.

QUESTION 83 You need to recommend a solution for the deployment of SQL Server 2014. The solution must meet the business requirements. What should you include in the recommendation? A. Create a new instance of SQL Server 2014 on the server that hosts the SQL Server 2008 instance. B. Upgrade the existing SQL Server 2008 instance to SQL Server 2014. C. Deploy two servers that have SQL Server 2014 installed and implement Failover Clustering. D. Deploy two servers that have SQL Server 2014 installed and implement database mirroring.
 Answer: C Explanation: Scenario: The databases must be available if the SQL Server service fails. Failover Clustering Overview Windows Server Failover Clustering (WSFC) with SQL Server

QUESTION 84 You need to recommend a solution to synchronize Database2 to App1_Db1. What should you recommend? A. Change data capture B. Snapshot replication C. Master Data Services D. Transactional replication
 Answer: D Explanation: Scenario: * Data from Database2 will be accessed periodically by an external application named Application1. The data from Database2 will be sent to a database named App1_Db1 as soon as changes occur to the data in Database2. * All data in Database2 is recreated each day and does not change until the next data creation process.

QUESTION 85 You need to recommend a database reporting solution that meets the business requirements. What should you include in the recommendation? A. Data collection B. Performance Monitor C. A maintenance plan D. A dynamic management view
 Answer: A Explanation: * Scenario: System administrators must be able to run real-time reports on disk usage. * The data collector provides an historical report for each of the System Data collection sets. Each of the following reports use data that is stored in the management data warehouse: Disk Usage Summary Query Statistics History Server Activity History You can use these reports to obtain information for monitoring system capacity and troubleshooting system performance. System Data Collection Set Reports

QUESTION 86 You need to recommend an isolation level for usp_UpdateOrderDetails. Which isolation level should recommend? A. Read committed B. Repeatable read C. Read uncommitted D. Serializable
 Answer: B Explanation: * Scenario: Database1 will also contain a stored procedure named

usp_UpdateOrderDetails. The stored procedure is used to update order information. The stored procedure queries the Orders table twice each time the procedure executes. The rows returned from the first query must be returned on the second query unchanged along with any rows added to the table between the two read operations.* REPEATABLE READS specifies that statements cannot read data that has been modified but not yet committed by other transactions and that no other transactions can modify data that has been read by the current transaction until the current transaction completes. Transaction Isolation Levels Case Study 8 - Contoso Ltd (Question 87 - Question 96) Overview Application Overview Contoso, Ltd., is the developer of an enterprise resource planning (ERP) application. Contoso is designing a new version of the ERP application. The previous version of the ERP application used SQL Server 2008 R2. The new version will use SQL Server 2014. The ERP application relies on an import process to load supplier data. The import process updates thousands of rows simultaneously, requires exclusive access to the database, and runs daily. You receive several support calls reporting unexpected behavior in the ERP application. After analyzing the calls, you conclude that users made changes directly to the tables in the database. Tables The current database schema contains a table named OrderDetails. The OrderDetails table contains information about the items sold for each purchase order. OrderDetails stores the product ID, quantities, and discounts applied to each product in a purchase order. The product price is stored in a table named Products. The Products table was defined by using the SQL_Latin1_General_CP1_CI_AS collation. A column named ProductName was created by using the varchar data type. The database contains a table named Orders. Orders contains all of the purchase orders from the last 12 months. Purchase orders that are older than 12 months are stored in a table named OrdersOld. Stored Procedures The current version of the database contains stored procedures that change two tables. The following shows the relevant portions of the two stored procedures:

```
CREATE PROC Sales.Proc1
AS
BEGIN TRAN
UPDATE Sales.Table1 ...
UPDATE Sales.Table2 ...
COMMIT TRAN
GO
CREATE PROC Sales.Proc2
AS
BEGIN TRAN
UPDATE Sales.Table2 ...
UPDATE Sales.Table1 ...
COMMIT TRAN
GO
```

Customer Problems Installation Issues The current version of the ERP application requires that several SQL Server logins be set up to function correctly. Most customers set up the ERP application in multiple locations and must create logins multiple times. Index Fragmentation Issues Customers discover that clustered indexes often are fragmented. To resolve this issue, the customers defragment the indexes more frequently. All of the tables affected by fragmentation have the following columns that are used as the clustered index key:

Column	Data type
id	uniqueidentifier
lastModified	datetime
modifiedBy	varchar(200)

Backup Issues Customers who have large amounts of historical purchase order data report that backup time is unacceptable. Search Issues Users report that when they search product names, the search results exclude product names that contain accents, unless the search string includes the accent. Missing Data Issues Customers report that when they make a price change in the Products table, they cannot retrieve the price that the item was sold for in previous orders. Query Performance Issues Customers report that query performance degrades very quickly. Additionally, the customers report that users cannot run queries when SQL Server runs maintenance tasks. Import Issues During the monthly import process, database administrators receive many supports call from users who report that they cannot access the supplier data. The database administrators want to reduce the amount of time required to import the data. Design Requirements File Storage Requirements The ERP database stores scanned documents that are larger than 2 MB. These files must only be accessed through the ERP application. File access must have the best possible read and write performance. Data Recovery Requirements If the import process fails, the database must be returned to its prior state immediately.

Security Requirements You must provide users with the ability to execute functions within the ERP application, without having direct access to the underlying tables. Concurrency Requirements You must reduce the likelihood of deadlocks occurring when Sales.Proc1 and Sales.Proc2 execute. QUESTION 87 You need to recommend a solution that resolves the missing data issue. The solution must minimize the amount of development effort. What should you recommend? A. Denormalize the Products table. B. Denormalize the OrderDetails table. C. Normalize the OrderDetails table. D. Normalize the Products table. Answer: D
 Explanation: * Scenario: / Missing Data Issues Customers report that when they make a price change in the Products table, they cannot retrieve the price that the item was sold for in previous orders. / The current database schema contains a table named OrderDetails. The OrderDetails table contains information about the items sold for each purchase order. OrderDetails stores the product ID, quantities, and discounts applied to each product in a purchase order. The product price is stored in a table named Products. [http://technet.microsoft.com/en-us/library/aa933055\(v=sql.80\).aspx](http://technet.microsoft.com/en-us/library/aa933055(v=sql.80).aspx) QUESTION 88 You need to recommend a solution that addresses the index fragmentation and index width issue. What should you include in the recommendation? (Each correct answer presents part of the solution. Choose all that apply.) A. Change the data type of the lastModified column to smalldatetime. B. Remove the lastModified column from the clustered index. C. Change the data type of the modifiedBy column to tinyint. D. Change the data type of the id column to bigint. E. Remove the modifiedBy column from the clustered index. F. Remove the id column from the clustered index. Answer: BE Explanation: Scenario: Index Fragmentation Issues Customers discover that clustered indexes often are fragmented. To resolve this issue, the customers defragment the indexes more frequently. All of the tables affected by fragmentation have the following columns that are used as the clustered index key:

Column	Data type
id	uniqueidentifier
lastModified	datetime
modifiedBy	varchar(200)

QUESTION 89 You need to recommend a solution that reduces the time it takes to import the supplier data. What should you include in the recommendation? A. Enable instant file initialization. B. Reorganize the indexes. C. Disable Resource Governor. D. Enable Auto Update Statistics. Answer: CE Explanation: * The ERP application relies on an import process to load supplier data. The import process updates thousands of rows simultaneously, requires exclusive access to the database, and runs daily. QUESTION 90 You need to recommend a solution that meets the data recovery requirement. What should you include in the recommendation? A. A differential backup. B. A transaction log backup. C. Snapshot isolation. D. A database snapshot. Answer: D
 Braindump2go is one of the Leading 70-469 Exam Preparation Material Providers Around the World! We Offer 100% Money Back Guarantee on All Products! Feel Free In Downloading Our New Released 70-469 Real Exam Questions!

Compared Before Buying Microsoft 70-469 PDF &		
Pass4sure	Braindump2go	T
281 Q&As - Practice	292 Q&As - Real Questions 100% Pass OR Money Back	281 Q&As
\$124.99	\$99.99	\$171.32
No Discount	Coupon Code: BDNT2014	No Discou

<http://www.braindump2go.com/70-469.html>