

Free Braindump2go Microsoft 70-450 VCE Exam Questions and Answers Download (11-20)

Braindump2go New Published Microsoft 70-450 Dumps PDF Contains the latest questions from Microsoft Exam Center! 100% Certification got guaranteed! Vendor: Microsoft Exam Code: 70-450 Exam Name: PRO: Designing, Optimizing and Maintaining a Database Administrative Solution Using Microsoft SQL Server 2008 Keywords: [70-450 Exam Dumps](#), [70-450 Practice Tests](#), [70-450 Practice Exams](#), [70-450 Exam Questions](#), [70-450 PDF](#), [70-450 VCE Free](#), [70-450 Book](#), [70-450 E-Book](#), [70-450 Study Guide](#), [70-450 Braindump](#), [70-450 Prep Guide](#)

Compared Before Buying Microsoft 70-450 PDF & VCE!		
Pass4sure	Braindump2go	Test King
	100% Pass OR Money Back	
Not In Stock	132 Q&As – Real Questions	Not In Stock
/	\$99.99	/
/	Coupon Code: BDNT2014	/

Microsoft 70-450 Dumps VCE Download: <http://www.braindump2go.com/70-450.html> QUESTION 11 You administer a SQL Server 2008 infrastructure. You plan to design a maintenance strategy for a mission-critical database that includes a large table named Orders. The design plan includes index maintenance operations. You must design the strategy after considering the following facts:- The Orders table in the database is constantly accessed.- New rows are frequently added to the Orders table.- The average fragmentation for the clustered index of the Orders table is less than 2 percent.- The Orders table includes a column of the xml data type. You need to implement the strategy so that the performance of the queries on the table is optimized. What should you do? A. Drop the clustered index of the Orders table. B. Rebuild the clustered index of the Orders table offline once a month. C. Reorganize the clustered index of the Orders table by decreasing the fill factor. D. Exclude the clustered index of the Orders table from scheduled reorganizing or rebuilding operations. Answer: D Explanation: Since the clustered index never has any significant fragmentation, there's no reason to rebuild or reorganize it. QUESTION 12 You administer a SQL Server 2008 infrastructure. An instance contains a database that includes a large table named OrderDetails. The application queries only execute DML statements on the last three months data. Administrative audits are conducted monthly on data older than three months. You discover the following performance problems in the database. The performance of the application queries against the OrderDetail table is poor. The maintenance tasks against the database, including index defragmentation, take a long time. You need to resolve the performance problems without affecting the server performance. What should you do? A. Create a database snapshot for the OrderDetails table every three months. Modify the queries to use the current snapshot. B. Create an additional table named OrderDetailsHistory for data older than three months. Partition the OrderDetails and OrderDetailsHistory tables in two parts by using the OrderDate column. Create a SQL Server Agent job that runs every month and uses the ALTER TABLE...SWITCH Transact-SQL statement to move data that is older than three months to the OrderDetailsHistory table. C. Create an additional table named OrderDetailsHistory for data older than three months. Create a SQL Server Agent job that runs the following Transact-SQL statement every month. INSERT INTO OrderDetailsHistory SELECT * FROM OrderDetails WHERE DATEDIFF(m, OrderDate, GETDATE ()) > 3 D. Create an additional table named OrderDetailsHistory for data older than three months. Use the following Transact-SQL statement. CREATE TRIGGER trgMoveData ON OrderDetails AFTER INSERT AS INSERT INTO OrderDetailsHistory SELECT * FROM OrderDetails WHERE DATEDIFF(m, OrderDate, GETDATE()) > 3 Answer: B QUESTION 13 You administer a SQL Server 2008 infrastructure. Humongous Insurance has 20 branch offices that store customer data in SQL Server 2008 databases. Customer data that is stored across multiple database instances has to be security compliant. You plan to design a strategy for custom policies by using the Policy-Based Management feature. Custom policies are in XML format. The strategy must meet the following requirements:- Custom policies are distributed to all instances.- The policies are enforced on all instances. You need to implement the strategy by using the least amount of administrative effort. What should you do? A. Use a configuration server. B. Use the Distributed File System Replication service. C. Distribute the policies by using Group Policy Objects. D. Distribute the policies by using the Active Directory directory service Answer: A Explanation: Configuration servers are the original name for central management servers, which allow the administration and enforcement of SQL Server 2008 policies for multiple servers to be centralized. QUESTION 14 You administer a SQL Server 2008 infrastructure. Developers in your company have rights to author policies. A test server is used to develop and test the policies. The Policy-Based Management feature generates SQL Server Agent alerts when a policy is violated. The developers are able to create and modify policies, but are unable to test policy violation alerts.

You need to grant the necessary permission to the developers to test the policies. You also need to comply with the least privilege principle when you grant the permission. What should you do? A. Add the developers to the sysadmin server role. B. Grant the ALTER TRACE permission to the developers. C. Add the developers to the PolicyAdministratorRole role in the MSDB database. D. Grant the EXECUTE permission on the sys.sp_syspolicy_execute_policy stored procedure to the developers. Answer: B

Explanation: <http://technet.microsoft.com/en-us/library/bb510667.aspx> Additional Considerations About Alerts Be aware of the following additional considerations about alerts:- Alerts are raised only for policies that are enabled. Because On demand policies cannot be enabled, alerts are not raised for policies that are executed on demand.- If the action you want to take includes sending an e-mail message, you must configure a mail account. We recommend that you use Database Mail. For more information about how to set up Database Mail, see How to: Create Database Mail Accounts (Transact-SQL).- Alert security: When policies are evaluated on demand, they execute in the security context of the user. To write to the error log, the user must have ALTER TRACE permissions or be a member of the sysadmin fixed server role. Policies that are evaluated by a user that has less privileges will not write to the event log, and will not fire an alert. The automated execution modes execute as a member of the sysadmin role. This allows the policy to write to the error log and raise an alert.

QUESTION 15 You administer a SQL Server 2008 infrastructure. You plan to design a solution to obtain hardware configurations, such as the number of processors on a computer and the processor type of all SQL Server 2008 computers. The solution must meet the following requirements:- It is hosted on the central computer.- It can verify hardware configurations for multiple servers. You need to select a technology that meets the requirements by using the minimum amount of development effort. What should you do? A. Use the Invoke-Sqlcmd cmdlet in SQL Server PowerShell cmdlet. B. Define policies based on conditions by using the ExecuteSql function. C. Define policies based on conditions by using the ExecuteWQL function. D. Use the Windows Management Instrumentation (WMI) provider for the server events. Answer: C

Explanation: ExecuteWQL is a relatively straightforward way to query operating system data from SQL server. It can then be stored in a database for analysis. The Invoke-Sqlcmd cmdlet is a powershell cmdlet for executing sql commands. It doesn't apply well to the question. WMI is a driver extension with scripting language and could theoretically be used to accomplish the goal, but with a much more complex development process. EXECUTESQL is a SQL command for running a pre-built SQL statement

QUESTION 16 You administer a SQL Server 2008 instance for your company. Your company has a team of database administrators. A team of application developers create SQL Server 2008 Integration Services (SSIS) packages on the test server in a shared project. One of the packages requires a fixed cache file. On completion of development, the packages will be deployed to the production server. Only the database administrators can access the production server. You need to ensure that the application developers can deploy the project successfully to the production server. What should you do? A. Use the Import and Export Wizard to save packages. B. Create a deployment utility for the SSIS project. C. Create a direct package configuration for each package. D. Create an indirect package configuration for all packages. Answer: B

Explanation: This is a strange question. The underlying lesson is that deployment utilities make SSIS package deployment easier, especially in situations where limited access may be available. Direct and indirect package configurations are explained on MSDN. Official source: <http://msdn.microsoft.com/en-us/library/ms141682.aspx>

QUESTION 17 You administer a SQL Server 2008 infrastructure. You plan to design an infrastructure for a new application. The application has the following requirements:- Users can connect to an instance named SQLSERVER1.- SQLSERVER1 is linked to a server named SQLSERVER2.- SQLSERVER1 and SQLSERVER2 run on different computers.- The SQL Server instances use only Windows authentication. You need to configure the infrastructure to ensure that the distributed queries are executed in the Windows security context of the login. Which two actions should you perform? (Each correct answer presents part of the solution. Choose two.) A. Configure all servers to use the Shared Memory network protocol. B. Register a server principal name (SPN) for SQLSERVER1 and SQLSERVER2. C. Use the local computer account as a service account for SQLSERVER1 and SQLSERVER2. D. Create a map for each SQL login from SQLSERVER1 to SQLSERVER2 and use the impersonate option. E. Ensure that the two instances use the same Windows account for the Microsoft SQL Service. Create the link so that each account uses the current security context. Answer: B

Explanation: <http://msdn.microsoft.com/en-us/library/ms189580%28v=SQL.100%29.aspx>

QUESTION 18 You administer two SQL Server 2008 instances named Instance1 and Instance2. Instance1 contains the Sales database, and Instance2 contains the Accounts database. A procedure in the Sales database starts a transaction. The procedure then updates the Sales.dbo. Order table and the Accounts.dbo. OrderHistory table through a linked server. You need to ensure that the transaction uses a two-phase commit. What should you do? A. Configure the linked server to use distributed transactions. B. Configure a Service Broker to enable the appropriate transaction control. C. Ensure that the linked server is appropriately configured for delegation. D. Ensure that the linked server is appropriately configured for impersonation. Answer: A

Explanation: A distributed transaction can be executed via a linked server using BEGIN DISTRIBUTED TRANSACTION. This will guarantee that the entire batch will complete or fail.

Two-phase commit is a transaction protocol employed by the MSDTC to ensure this is possible. The service broker is for distributing heavy workloads across multiple database instances. Delegation and impersonation are for passing credentials from one server/instance to another. Delegation passes windows credentials; Impersonation passes a SQL Server login.

QUESTION 19 You administer a SQL Server 2008 instance named Instance1 at the New York central site. Your company has a sales team to fulfill purchase orders for customer requests. The sales team uses portable computers to update data frequently in a local database. When the portable computers connect to the central site, the local database must be synchronized with a database named Sales. You plan to create a replication model to replicate the local database to the Sales database. The replication model must meet the following requirements:

- Data conflicts are handled when multiple users update the same data independently.
- The sales team cannot update sensitive data such as product price.
- The sales team can synchronize data at scheduled times and on demand also.

You need to identify the best model to replicate data by using minimum development efforts. What should you do?

A. Use merge replication along with each portable computer that is set up as a subscriber.
B. Use snapshot replication along with each portable computer that is set up as a subscriber.
C. Use transactional replication along with each portable computer that is set up as a publisher.
D. Use SQL Server Integration Services (SSIS) to push data changes and pull updates to the Sales database along with the SSIS packages, on demand.

Answer: A
Explanation: <http://www.codeproject.com/KB/database/sql2005-replication.aspx>

QUESTION 20 You are planning to upgrade a database application that uses merge replication. The table currently has a column type of UNIQUEIDENTIFIER and has a DEFAULT constraint that uses the NEWID() function. A new version of the application requires that the FILESTREAM datatype be added to a table in the database. The data type will be used to store binary files. Some of the files will be larger than 2 GB in size. While testing the upgrade, you discover that replication fails on the articles that contain the FILESTREAM data. You find out that the failure occurs when a file object is larger than 2 GB. You need to ensure that merge replication will continue to function after the upgrade. You also need to ensure that replication occurs without errors and has the best performance. What should you do? (More than one answer choice may achieve the goal. Select the BEST answer.)

A. Drop and recreate the table that will use the FILESTREAM data type.
B. Change the DEFAULT constraint to use the NEWSEQUENTIALID() function.
C. Place the table that will contain the FILESTREAM data type on a separate filegroup.
D. Use the sp_changemergearticle stored procedure and set the @stream_blob_columns option to true for the table that will use the FILESTREAM data type.

Answer: D
Explanation: <http://msdn.microsoft.com/en-us/library/bb895334.aspx>

Considerations for Merge Replication

If you use FILESTREAM columns in tables that are published for merge replication, note the following considerations:

- Both merge replication and FILESTREAM require a column of data type uniqueidentifier to identify each row in a table. Merge replication automatically adds a column if the table does not have one. Merge replication requires that the column have the ROWGUIDCOL property set and a default of NEWID() or NEWSEQUENTIALID(). In addition to these requirements, FILESTREAM requires that a UNIQUE constraint be defined for the column. These requirements have the following consequences:
- If you add a FILESTREAM column to a table that is already published for merge replication, make sure that the uniqueidentifier column has a UNIQUE constraint. If it does not have a UNIQUE constraint, add a named constraint to the table in the publication database. By default, merge replication will publish this schema change, and it will be applied to each subscription database. For more information about schema changes, see Making Schema Changes on Publication Databases.
- If you add a UNIQUE constraint manually as described and you want to remove merge replication, you must first remove the UNIQUE constraint; otherwise, replication removal will fail.
- By default, merge replication uses NEWSEQUENTIALID() because it can provide better performance than NEWID(). If you add a uniqueidentifier column to a table that will be published for merge replication, specify NEWSEQUENTIALID() as the default.

Merge replication includes an optimization for replicating large object types. This optimization is controlled by the @stream_blob_columns parameter of sp_addmergearticle. If you set the schema option to replicate the FILESTREAM attribute, the @stream_blob_columns parameter value is set to true. This optimization can be overridden by using sp_changemergearticle. This stored procedure enables you to set @stream_blob_columns to false. If you add a FILESTREAM column to a table that is already published for merge replication, we recommend that you set the option to true by using sp_changemergearticle.

Enabling the schema option for FILESTREAM after an article is created can cause replication to fail if the data in a FILESTREAM column exceeds 2 GB and there is a conflict during replication. If you expect this situation to arise, it is recommended that you drop and re-create the table article with the appropriate FILESTREAM schema option enabled at creation time.

Merge replication can synchronize FILESTREAM data over an HTTPS connection by using Web Synchronization. This data cannot exceed the 50 MB limit for Web Synchronization; otherwise, a run-time error is generated.

2015 Latest Released Microsoft 70-450 Exam Dumps Free Download From Braindump2go Now! All Questions and Answers are checked again by Braindump2go Experts Team, 100% Real Questions and Correct Answers Guaranteed! Full Money Back Guarantee Show our Confidence in helping you have a 100% Success of Exam 70-450! Just have a try!

Compared Before Buying Microsoft 70-450 PDF & VCE!		
Pass4sure	Braindump2go	Test King
	100% Pass OR Money Back	
Not In Stock	132 Q&As – Real Questions	Not In Stock
/	\$99.99	/
/	Coupon Code: BDN2014	/

<http://www.braindump2go.com/70-450.html>