

Download Microsoft 70-450 Dumps Latest Version Right Now From Braindump2go (31-40)

Quick and Effective Microsoft 70-450 Exam Preparation Options - Braindump2go new released 70-450 Exam Dumps Questions! Microsoft Official 70-450 relevant practice tests are available for Instant downloading at Braindump2go! PDF and VCE Formates, easy to use and install! 100% Success Achievement 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 & | | |
|--|--------------------------------------|-----------------|
| Pass4sure | Braindump2go | |
| | 100% Pass OR Money Back | |
| Not In Stock | 132 Q&As – Real Questions | Not In S |
| / | \$99.99 | / |
| / | Coupon Code: BDN2014 | / |

Microsoft 70-450 Dumps VCE Download: <http://www.braindump2go.com/70-450.html> QUESTION 31 You are a professional level SQL Sever 2008 Database Administrator. A solution will be implemented to have hardware configurations obtained. The hardware configuration obtained should include the number of processors on a computer and the processor type of all SQL Server 2008 computers. The requirements listed below should be satisfied by the solution. The solution should be hosted on the central computer. And the hardware configurations for multiple servers should be verified by the solution. A technology should be selected to satisfy the requirements, and the least amount of development effort should be utilized. Which action should you perform to finish the task? A. To finish the task, you should utilize the Invoke-Sqlcmd cmdlet in SQL Server PowerShell cmdlet. B. To finish the task, you should utilize the ExecuteWQL function to define policies based on conditions. C. To finish the task, you should utilize the Windows Management Instrumentation (WMI) provider for the server events. D. To finish the task, you should utilize the ExecuteSql function to define policies based on conditions. Answer: B Explanation: Policy-Based Management is a system for managing one or more instances of SOL Server 2008. When SOL Server policy administrators use Policy-Based Management, they use SOL Server Management Studio to create policies to manage entities on the server, such as the instance of SOL Server, databases, or other SOL Server objects. Policy-Based Management has three components: - Policy management- Policy administrators create policies, - Explicit administration Administrators select one or more managed targets and explicitly check that the targets comply with a specific policy, or explicitly make the targets comply with a policy. Evaluation modes There are four evaluation modes, three of which can be automated: On demand. This mode evaluates the policy when directly specified by the user. On change: prevent. This automated mode uses DDL triggers to prevent policy violations. Important: If the nested triggers server configuration option is disabled, OI! change: pI'BV6T II will not work correctly. Policy-Based Management relies on DDL triggers to detect and roll back DDL operations that do not comply with policies that use this evaluation mode. Removing the Policy-Based Management DDL triggers or disabling nest triggers, will cause this evaluation mode to fail or perform unexpectedly. On change: log only. This automated mode uses event notification to evaluate a policy when a relevant change is made. On schedule. This automated mode uses a SOL Server Agent job to periodically evaluate a policy. When automated policies are not enabled, Policy-Based Management will not altect system performance. 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 32 You are a professional level SQL Sever 2008 Database Administrator. A database is hosted by the instance. You can see the backup strategy utilized by the database from table listed below: A batch process is run on Saturday, and the process lasts from 21:00 hr to 23:00 hr; you find that the batch process takes no effect because certain data is modified some data by a user at 21:05 hr on the same day. The database should be recovered to its state at the start of the batch process in the shortest time. Which action will you perform?

| Backup Type | Full Database | Differential | Trans |
|---------------------------|-----------------------------|---------------------------|---------|
| Frequency | Tuesday, Thursday, Saturday | Monday, Wednesday, Friday | Hourly |
| Time of the backup | 21:00 | 1:00 | Hourly |
| Time taken for the backup | 60-90 minutes | 30-45 minutes | Less th |

A. The full database backup that was performed on Saturday should be recovered. B. The full database backup that was performed on Thursday should be recovered. And then, all transaction logs should be recovered from the time of the full backup on Thursday and should be stopped at 21:00 hr on Saturday. C. The full database backup that was performed on Thursday should be recovered. And then the differential backup that was performed on Friday should be recovered. At last, all transaction logs should be recovered from the time of the differential backup on Friday and should be stopped at 21:00 hr on Saturday. D. The full database backup that was performed on Saturday should be recovered. And then all transaction logs should be recovered from the time of the full backup on Saturday and should be stopped at 21:00 hr on Saturday. Answer: C Explanation: Because we want the database restored to Saturday 21:00, because the job start exactly at the same time than the full backup, we can't use it. So, we have to restore the last full backup then the differential backup and finally the transactions logs. QUESTION 33 You are a professional level SQL Sever 2008 Database Administrator. The largest table in a database is partitioned into four equal-sized parts. The first two partitions are seldom queried, while the third partition is regularly queried by several reports. Meanwhile, the fourth partition is heavily queried and regularly modified. The disk space of the table should be cut to the least and the query performance should be with the least impact. What will you do? (Choose more than one.)

| Backup Type | Full Database | Differential | Transaction Log |
|---------------------------|-----------------------------|---------------------------|---------------------|
| Frequency | Tuesday, Thursday, Saturday | Monday, Wednesday, Friday | Hourly |
| Time of the Backup | 21:00 | 21:00 | Hourly |
| Time taken for the backup | 60-90 minutes | 30-45 minutes | Less than 5 minutes |

A. PAGE compression should be enabled for the first two partitions. B. ROW compression should be enabled for the last partition. C. ROW compression should be enabled for the third partition. D. ROW compression should be enabled for the first two partitions. E. PAGE compression should be enabled for the last partition. F. PAGE compression should be enabled for the third partition. Answer: A Explanation: The two first partition are seldom queried, this means that we can compress to the maximum the data in the partition. Page compression will be use as it removes duplicate pattern and compress to the max the data. The Third partition is regularly queried, this means that the partition should be @ht. Thus, the ROW compression is more adequate for this partition. QUESTION 34 You are a professional level SQL Sever 2008 Database Administrator. The capacity planning information is required by your company. A long-term infrastructure monitoring strategy should be implemented. Which two actions should be contained in the strategy? (Choose more than one) A. System monitor counters should be reviewed regularly. B. After deployment of the production databases, the system should be baselined. C. All databases should be backed up every day. D. The system log and the application log should be cleared every hour. E. A maintenance plan that rebuilds indexes every week should be created. Answer: A B Explanation: After the production deployment, we will set a baseline of system counter in order to have metrics about the health of the server. This is useful in order to query on a regular basis the status of those counter. When big deviation from the counters are seen, further investigation could be done to solve potential problem. QUESTION 35 You are a professional level SQL Sever 2008 Database Administrator. It is on a four-processor, quad-core server. CPU pressure is regularly experienced by the server. An extremely large mission-critical database is consisted by the instance, and it is utilized constantly. The online index rebuilds should not occupy all available CPU cycles. From the following four configurations, which one should you utilize? A. You should utilize affinity mask. B. You should utilize max degree of parallelism. C. You should utilize affinity I/O mask. D. You should utilize optimize for ad hoc workloads. Answer: B Explanation: When SQL Server runs on a computer with more than one microprocessor or CPU, it detects the best degree of parallelism, that is, the number of processors employed to run a single statement, for each parallel plan execution. You can use the max degree of parallelism option to limit the number of processors to use in parallel plan execution. The default value of 0 uses all available processors. Set max degree of parallelism to 1 to suppress parallel plan generation. Set the value to a number greater than 1 (up to a maximum of 64) to restrict the maximum number of processors used by a single query execution. If a value greater than the number of available processors is specified, the actual number of available processors is used. If the computer has only one processor, the max degree of parallelism value is ignored, Overriding max degree of parallelism SQL Server will consider parallel execution plans for queries, index data definition language (DDL) operations, and static and keyset-driven cursor population. For more information about parallel execution, see Degree of Parallelism. You can override the max degree of parallelism value in queries by specifying the MAXDOP query hint in the query statement. For more information, see Query Hints (Transact-SQL). Index operations that create or rebuild an index, or that drop a clustered index, can be resource intensive. You can override the max degree of parallelism value for index operations by specifying the MAXDOP index option in the index statement. The MAXDOP value is applied to the statement at execution time and is not stored in the index metadata. For more information, see Configuring Parallel Index Operations. In addition to queries and index operations, this option also controls the parallelism of DBCC CHECKTABLE, DBCC CHECKDB, and DBCC CHECKFILEG

ROU P. You can disable parallel execution plans for these statements by using trace flag 2528. For more information, see Trace Flags (Transact-SQL). Note: If the affinity mask option is not set to the default, it may restrict the number of processors available to SQL Server on symmetric multiprocessing (SMP) system.

QUESTION 36 You are a professional level SQL Server 2008 Database Administrator. A large database is hosted by the instance. The company utilizes the following backup strategy for the database: - First, the performance of full database backups is implemented weekly. - Secondly, the performance of differential backups is implemented daily. - Thirdly, the performance of transaction log backups is implemented hourly. According to the restoring plan, an unscheduled full backup should be performed. A full database backup should be performed, and the scheduled backup strategy should not be interrupted. From the following four options, which one should be utilized to finish the task? A. To finish the task, SKIP should be utilized. B. To finish the task, NOINIT should be utilized. C. To finish the task, MIRROR TO should be utilized. D. To finish the task, COPY_ONLY should be utilized. Answer: D Explanation: COPY_ONLY is a backup option which will not affect the normal backup sequence. NOINIT specifies that the backup should be appended to an existing set. SKIP works in conjunction with INIT/NOINIT and skips the check for media name and expiration. MIRROR TO will make a secondary copy of the backup. ----- A copy-only backup is a SQL Server backup that is independent of the sequence of conventional SQL Server backups. Usually, taking a backup changes the database and affects how later backups are restored, However, occasionally, it is useful to take a backup for a special purpose without affecting the overall backup and restore procedures for the database. For this purpose, copy-only backups were introduced SQL Server 2005. The types of copy-only backups are as follows: Copy-only full backups {all recovery models} A copy-only full backup cannot serve as a differential base or differential backup and does not affect the differential base. Copy-only log backups {full recovery model and bulk-logged recovery model only} A copy-only log backup preserves the existing log archive point and, therefore, does not affect the sequencing of regular log backups. Copy-only log backups are typically unnecessary. Instead, you can create another routine, current log backup (using WITH NORECOVERY), and then use that backup together with all other previous log backups that are required for the restore sequence. However, a copy-only log backup can be created for performing an online restore. The transaction log is never truncated after a copy-only backup. Copy-only backups are recorded in the is_copy_only column of the backupset table.

QUESTION 37 You are a professional level SQL Server 2008 Database Administrator. A business-critical database is hosted by the instance, and the database must be constantly available to the users with no data loss. Filestream data is contained by the database. A high-availability solution should be designed for the site. Which solution should be utilized? A. Database snapshot should be utilized. B. Failover clustering should be utilized. C. Asynchronous database mirroring should be utilized. D. Synchronous database mirroring with a witness server should be utilized. Answer: B Explanation: Failover clustering in SQL Server provides high-availability support for an entire SQL Server instance. For example, you can configure a SQL Server instance on one node of a failover cluster to fail over to any other node in the cluster during a hardware failure, operating system failure, or a planned upgrade. A failover cluster is a combination of one or more nodes (servers) with two or more shared disks, known as a resource group. The combination of a resource group, along with its network name, and an internet protocol (IP) address that makes up the clustered application or server, is referred to as a failover cluster or a failover cluster instance. A SQL Server failover cluster appears on the network as if it were a single computer, but has functionality that provides failover from one node to another if the current node becomes unavailable. A failover cluster appears on the network as a normal application or single computer, but it has additional functionality that increases its availability. Failover clustering has a new architecture and new workflow for all Setup scenarios in SQL Server 2008. The two options for installation are Integrated installation and Advanced Enterprise installation. Integrated installation creates and configures a single-node SQL Server failover cluster instance. Additional nodes are added using add node functionality in Setup. For example, for Integrated installation, you run Setup to create a single-node failover cluster. Then, you run Setup again for each node you want to add to the cluster. Advanced Enterprise installation consists of two steps. The Prepare step prepares all nodes of the failover cluster to be operational. Nodes are defined and prepared during this initial step. After you prepare the nodes, the Complete step is run on the active node--the node that owns the shared disk--to complete the failover cluster instance and make it operational.

When to Use Failover Clustering Use failover clustering to: - Administer a failover cluster from any node in the clustered SQL Server configuration. For more information, see [Installing a SQL Server 2008 Failover Cluster](#). - Allow one failover cluster node to fail over to any other node in the failover cluster configuration. For more information, see [Installing a SQL Server 2008 Failover Cluster](#). - Configure Analysis Services for failover clustering. For more information, see [How to: install Analysis Services on a failover cluster](#). - Execute full-text queries by using the Microsoft Search service with failover clustering. For more information, see [Using SQL Server Tools with Failover Clustering](#).

QUESTION 38 You are a professional level SQL Server 2008 Database Administrator. New SQL Server 2008 Integration Services (SSIS) packages will be deployed, and the packages are created by the SQL Server Import and Export Wizard. The packages will utilize stored user names and passwords to

connect to external data sources. Since you are the technical support of your company, you are required to confirm that unauthorized users can not get access to the user names and passwords. Which action will you perform? A. No other data should be saved together with the packages, and then files should be saved to the most secure directory. B. No sensitive data should be saved together with the packages, and then the packages should be saved in the SQL Server 2008 server. C. The SQL Server 2008 server should be utilized to store the packages, and administrate the security in the SQL Server 2008 server. D. A common password should be set to make a highest level of security for all information in the package when each package is saved. Answer: C

Explanation: <http://msdn.microsoft.com/en-us/library/ms141747.aspx?ppud=4> There are several options for package protection. C) is the equivalent of "Rely on server storage", which allows SQL server to handle the security. A) and B) will cause the package to fail, and D) is just a terrible idea.

QUESTION 39 You are a professional level SQL Sever 2008 Database Administrator. A new database named Engineering is deployed. Large documents are managed by the Engineering database, and the documents will be revised occasionally. A table structure should be designed to permit fast read access. The storage space requirements should be cut to the least. What will you do to achieve the goal? (Choose more than one.) A. To achieve the goal, row-level compression should be utilized on the document table. B. To achieve the goal, NTFS file system compression should be utilized on the volume. C. To achieve the goal, varbinary(MAX) data type should be utilized with FILESTREAM storage. D. To achieve the goal, row-level compression should be enabled on all columns that the vardecimal() data type is utilized. Answer: BC

Explanation: 1. FILESTREAM Much of the data that is created every day is unstructured data, such as text documents, images, and videos, This unstructured data is often stored outside the database, separate from its structured data. This separation can cause data management complexities. Or, if the data is associated with structured storage, the file streaming capabilities and performance can be limited. FILESTREAM

integrates the SOL Server Database Engine with an NTFS file system by storing varbinary(max) binary large object (BLOB) data as files on the file system. Transact-SOL statements can insert, update, query, search, and back up FILESTREAIVI data. Win32 tile system interfaces provide streaming access to the data. FILESTFiEAlv1 uses the NT system cache for caching file data. This helps reduce any effect that FILESTREAM data might have on Database Engine performance. The SQL Server butter pool is not used; therefore, this memory is available for query processing. Note: FILESTREAM data is not encrypted even when transparent data encryption is enabled.

2. NTFS compression NTFS compression should be used to compress the files on the NTFS file system

QUESTION 40 You are a professional level SQL Sever 2008 Database Administrator. A monitoring solution should be designed for the instance to have object utilization statistics monitored. A list of first 10 objects should be identified by the solution for each of the following two components:- Most frequently executed stored procedures and functions- Long running Transact-SQL statements

The monitoring solution should be implemented with minimum performance effect, and the least amount of administrative effort should be utilized. Which action will you perform to accomplish the task? A. To accomplish the task, a System Monitor counter log should be utilized. B. To accomplish the task, a client-side SQL Server Profiler trace should be utilized. C. To accomplish the task, a server-side SQL Server Profiler trace should be utilized. D. To accomplish the task, dynamic management views should be utilized. Answer: D

Explanation: Dynamic management views and functions return server state information that can be used to monitor the health of a server instance, diagnose problems, and tune performance. Important: Dynamic management views and functions return internal, implementation-specific state data. Their schemas and the data they return may change in future releases of SOL Server.

Therefore, dynamic management views and functions in future releases may not be compatible with the dynamic management views and functions in SQL Server 2008. There are two types of dynamic management views and functions:-

Server-scoped dynamic management views and functions. These require VIEW SERVER STATE permission on the server.-

Database-scoped dynamic management views and functions. These require VIEW DATABASE STATE permission on the database.

Braindump2go New Released Premium 70-450 Exam Dumps Guarantee You a 100% Exam Success Or We Promise Full Money Back! Download Microsoft 70-450 Exam Dumps Full Version From Braindump2go Instantly!

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

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