

## [25-June-2018 Update Exam 70-764 VCE Dumps and 70-764 PDF Dumps from PassLeader

Valid 70-764 Dumps shared by PassLeader for Helping Passing 70-764 Exam! PassLeader now offer the newest 70-764 VCE dumps and 70-764 PDF dumps, the PassLeader 70-764 exam questions have been updated and ANSWERS have been corrected, get the newest PassLeader 70-764 dumps with VCE and PDF here: <https://www.passleader.com/70-764.html> (365 Q&As Dumps --> 388 Q&As Dumps --> 412 Q&As Dumps)

BTW, DOWNLOAD part of PassLeader 70-764 dumps from Cloud Storage:

[https://drive.google.com/open?id=0B-ob6L\\_QjGLpN3N6eHJ6Z2EzZWc](https://drive.google.com/open?id=0B-ob6L_QjGLpN3N6eHJ6Z2EzZWc)

NEW QUESTION 301

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have a data warehouse that stores sales data. One fact table has 100 million rows. You must reduce storage needs for the data warehouse. You need to implement a solution that uses column-based storage and provides real-time analytics for the operational workload.

Solution: You remove any clustered indexes and load the table for processing.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Clustered columnstore tables offer both the highest level of data compression as well as the best overall query performance. Clustered columnstore tables will generally outperform clustered index or heap tables and are usually the best choice for large tables. For these reasons, clustered columnstore is the best place to start when you are unsure of how to index your table.

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-tables-overview>

NEW QUESTION 302

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

Your company is developing a new business intelligence application that will access data in a Microsoft Azure SQL Database instance. All objects in the instance have the same owner. A new security principal named BI\_User requires permission to run stored procedures in the database. The stored procedures read from and write to tables in the database. None of the stored procedures perform IDENTIFY\_INSERT operations or dynamic SQL commands. The scope of permissions and authentication of BI\_User should be limited to the database. When granting permissions, you should use the principle of least privilege. You need to create the required security principals and grant the appropriate permissions.

Solution: You run the following Transact-SQL statement in the database:

```
CREATE USER BI_User WITH PASSWORD = 'Pa$$w0rd'  
GRANT EXECUTE TO BI_User  
EXEC sp_addrolemember 'db_datawriter', 'BI_user'
```

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:



more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

You are deploying a Microsoft SQL Server architecture to support a new mission-critical application. The application includes a dedicated reporting component that performs read-only operations against the application's databases. You need to implement a solution that meets the following requirements:

- Provide maximum uptime for the databases.
- Include automatic failover in the event of a hardware problem on the primary server.
- Separate the reporting workload from the read/write transactional processing workload and contain real-time data.

Modifications to the application to support the new architecture are not permitted. What should you implement?

- A. a Microsoft Azure Stretch Database
- B. log shipping
- C. an Always On Availability Group with all replicas in synchronous-commit mode
- D. a file share witness
- E. a SQL Server failover cluster instance (FCI)
- F. a Windows Cluster with a shared-nothing architecture
- G. an Always On Availability group with secondary replicas in asynchronous-commit mode

Answer: C

Explanation:

Synchronous-commit mode emphasizes high availability over performance, at the cost of increased transaction latency.

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/availability-modes-always-on-availability-groups?view=sql-server-2017>

NEW QUESTION 306

Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.

Start of repeated scenario.

You are a database administrator for a company that has on-premises Microsoft SQL Server environment. There are two domains in separate forests. There are no trust relationships between the domains. The environment hosts several customer databases, and each customer uses a dedicated instance running SQL Server 2016 Standard edition. The customer environments are shown in the following table:

Customer	Domain	Description
AdventureWorks Cycles	DomainB	The environment includes a database named <b>Adventureworks</b> that contains a single schema named <b>ADVSchema</b> . You must implement auditing for all objects in the <b>ADVSchema</b> schema. You must also implement auditing to record access to data that is considered sensitive by the company.
Tailspin Toys	DomainA	Tailspin Toys has a database named <b>TSpinDB</b> . Tailspin Toys requires a custom application that monitors <b>TSpinDB</b> and captures information over time about which database objects are accessed and how frequently they are accessed.
Contoso, Ltd.	DomainB	The environment has a database named <b>ConDB</b> and is also running SQL Server Reporting Services (SSRS).
Wingtip Toys	DomainA	Wingtip Toys has a database named <b>WingDB</b> . All tables in the database have indexes. Users report system response time is slow during peak activity periods. You observe that the performance issues are related to locking.  Wingtip Toys receives data updates from suppliers each week. You must implement a process for importing the data into <b>WingDB</b> . You must use minimal logging and minimize data loss during the import process.
Wide World Importers	DomainB	The environment includes a database named <b>WDWDB</b> . Neither auditing nor statistics are configured for <b>WDWDB</b> . You must log any deletion of views and all database record update operations.

End of repeated scenario.



#### NEW QUESTION 308

You manage a Microsoft SQL Server environment. You plan to encrypt data when you create backups. You need to configure the encryption options for backups. What should you configure?

- A. a certificate
- B. an MD5 hash
- C. an SHA-256 hash
- D. an AES 256-bit key

Answer: D

Explanation:

To encrypt a backup we need to configure an encryption algorithm (supported encryption algorithms are: AES 128, AES 192, AES 256, and Triple DES) and an encryptor (a certificate or asymmetric key).

<https://www.mssqltips.com/sqlservertip/3145/sql-server-2014-backup-encryption/>

#### NEW QUESTION 309

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

A company has an on-premises Microsoft SQL Server environment. SQL Server backups should be stored as Microsoft Azure blob pages. The connection process from the SQL Server instances to Azure should be encrypted. You need to store backups as Azure blob pages. Which option should you use?

- A. backup compression
- B. backup encryption
- C. file snapshot backup
- D. mirrored backup media sets
- E. SQL Server backup to URL
- F. SQL Server Managed Backup to Azure
- G. tail-log backup
- H. back up and truncate the transaction log

Answer: F

Explanation:

SQL Server Managed Backup to Microsoft Azure manages and automates SQL Server backups to Microsoft Azure Blob storage. You can choose to allow SQL Server to determine the backup schedule based on the transaction workload of your database. Or you can use advanced options to define a schedule. The retention settings determine how long the backups are stored in Azure Blob storage.

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/sql-server-managed-backup-to-microsoft-azure?view=sql-server-2017>

#### NEW QUESTION 310

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

A company has several Microsoft SQL Server database in Microsoft Azure. One database experiences a storage failure, and pages that store critical database metadata are corrupted. You need to perform an offline restore of the database's pages. Which option should you use first?

- A. backup compression
- B. backup encryption
- C. file snapshot backup
- D. mirrored backup media sets
- E. SQL Server backup to URL
- F. SQL Server Managed Backup to Azure
- G. tail-log backup
- H. back up and truncate the transaction log



SQL Server Managed Backup to Microsoft Azure supports point in time restore for the retention time period specified.

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/sql-server-managed-backup-to-microsoft-azure?view=sql-server-2017>

#### NEW QUESTION 313

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

A company has a Microsoft SQL Server environment in Microsoft Azure. The databases are stored directly in Azure blob storage. The company uses a complex backup process. You need to simplify the backup process. Future restores should not require differential or multiple incremental logs to perform a restore. You need to design a backup solution for the SQL Server instances. Which option should you use?

- A. backup compression
- B. backup encryption
- C. file snapshot backup
- D. mirrored backup media sets
- E. SQL Server backup to URL
- F. SQL Server Managed Backup to Azure
- G. tail-log backup
- H. back up and truncate the transaction log

Answer: C

Explanation:

SQL Server File-snapshot backup uses Azure snapshots to provide nearly instantaneous backups and quicker restores for database files stored using the Azure Blob storage service. This capability enables you to simplify your backup and restore policies.

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/file-snapshot-backups-for-database-files-in-azure?view=sql-server-2017>

#### NEW QUESTION 314

You have an application that queries a database. Users report that the application is slower than expected. You discover that several server process identifiers (SPIDs) have PAGELATCH\_UP and PAGELATCH\_EX waits. The resource descriptions of the SPIDs contains 2:1:1. You need to resolve the issue. What should you do?

- A. Allocate additional processor cores to the server.
- B. Add files to the file group of the application database.
- C. Reduce the fill factor of all clustered indexes.
- D. Add data files to tempdb.

Answer: D

Explanation:

PAGELATCH contention in tempdb is typically on allocation bitmaps and occurs with workloads with many concurrent connections creating and dropping small temporary tables (which are stored in tempdb). Assuming that the temporary tables are needed for performance, the trick is to have multiple data files for tempdb so that the allocations are done round-robin among the files, the contention is split over multiple PFS pages, and so the overall contention goes down.

<https://sqlperformance.com/2015/10/sql-performance/knee-jerk-wait-statistics-pagelatch>

#### NEW QUESTION 315

A company has an on-premises Microsoft SQL Server environment and Microsoft Azure SQL Database instances. The environment hosts several customer databases. A customer that uses an on-premises instance reports that queries take a long time to complete. You need to reconfigure table statistics so that the query optimizer can use the optimal query execution plans available. Which Transact-SQL segment should you use?

- A. sys.index\_columns
- B. UPDATE STATISTICS
- C. CREATE STATISTICS
- D. SET AUTO\_CREATE\_STATISTICS ON

Answer: D



