

[25/June/2018 Updated Latest PassLeader 365q 70-764 Exam Dumps Collection

New Updated 70-764 Exam Questions from PassLeader 70-764 PDF dumps! Welcome to download the newest PassLeader 70-764 VCE dumps: <http://www.passleader.com/70-764.html> (365 Q&As) Keywords: 70-764 exam dumps, 70-764 exam questions, 70-764 VCE dumps, 70-764 PDF dumps, 70-764 practice tests, 70-764 study guide, 70-764 braindumps, Administering a SQL Database Infrastructure Exam P.S. New 70-764 dumps PDF: https://drive.google.com/open?id=0B-ob6L_QjGLpN3N6eHJ6Z2EzZWc >> New 70-761 dumps PDF: https://drive.google.com/open?id=0B-ob6L_QjGLpaEZzRVFnOE9OenM >> New 70-762 dumps PDF: https://drive.google.com/open?id=0B-ob6L_QjGLpN3RVQ25sVUM5dkU >> New 70-765 dumps PDF: https://drive.google.com/open?id=0B-ob6L_QjGLpZHIHSG5KM09xUms >> New 70-767 dumps PDF: https://drive.google.com/open?id=0B-ob6L_QjGLpcXXXWU14dHhIUVk >> New 70-768 dumps PDF: https://drive.google.com/open?id=0B-ob6L_QjGLpeXAxaUJkWEZnVIU 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: One method of creating multiple lines of defense around your database is to implement all data access using stored procedures or user-defined functions. You revoke or deny all permissions to underlying objects, such as tables, and grant EXECUTE permissions on stored procedures. This effectively creates a security perimeter around your data and database objects.

<https://docs.microsoft.com/en-us/dotnet/framework/data/adonet/sql/managing-permissions-with-stored-procedures-in-sql-server> NEW QUESTION 303 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: `CREATE USER BI_User WITH PASSWORD = 'Pa$$w0rd'
GRANT EXECUTE TO BI_User` Does the

solution meet the goal?
A. Yes
B. No Answer:
A Explanation:
One method of creating multiple lines of defense around your database is to implement all data access using stored procedures or user-defined functions. You revoke or deny all permissions to underlying objects, such as tables, and grant EXECUTE permissions on stored procedures. This effectively creates a security perimeter around your data and database objects.

<https://docs.microsoft.com/en-us/dotnet/framework/data/adonet/sql/managing-permissions-with-stored-procedures-in-sql-server>

NEW QUESTION 304 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. You are migrating a set of databases from an existing Microsoft SQL Server instance to a new instance. You need to complete the migration while minimizing administrative effort and downtime. Which should you implement? A. log shipping B. an Always On Availability Group with all replicas in synchronous-commit mode C. a file share witness D. a SQL Server failover cluster instance (FCI) E. a Windows Cluster with a shared-nothing architecture F. an Always On Availability Group with secondary replicas in asynchronous-commit mode Answer: A Explanation: SQL Server Log shipping allows you to automatically send transaction log backups from a primary database on a primary server instance to one or more secondary databases on separate secondary server instances. The transaction log backups are applied to each of the secondary databases individually.

<https://docs.microsoft.com/en-us/sql/database-engine/log-shipping/about-log-shipping-sql-server?view=sql-server-2017> NEW

QUESTION 305 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. 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 company that imports bicycles. You can see it is connected to the Internet.
Tailspin Toys	DomainA	Tailspin Toys requires capturing images that are used in their products.
Contoso, Ltd.	DomainB	The company runs a large office.
Wingtip Toys	DomainA	Wingtip Toys data retrieval is slow and unreliable. Wingtip Toys is a toy company.
Wide World Importers	DomainB	The company neither imports nor exports. You can see it is updated frequently.

End of repeated scenario. You need to monitor WingDB and gather information for troubleshooting issues. Which two tools should you use? (Each correct answer presents a complete solution. Choose two.) A. sys.dm_tran_locks
B. sp_lock C. sys.dm_tran_active_snapshot_database_transactions
D. Activity Monitor E. sp_monitor Answer: BD Explanation: The performance issues is related to locking. B: sp_lock reports snapshot information about locks, including the object ID, index ID, type of lock, and type or resource to which the lock applies. D: The Activity Monitor in SQL Server Management Studio is useful for ad hoc views of current activity and graphically displays information about: ? Processes running on an instance of SQL Server. ? Blocked processes. ? Locks. ? User activity. Incorrect: Not E: System Monitor primarily tracks resource usage, such as the number of buffer manager page requests in use, enabling you to monitor server performance and activity using predefined objects and counters or user-defined counters to monitor events. System Monitor (Performance Monitor in Microsoft Windows NT 4.0) collects counts and rates rather than data about the events (for example, memory usage, number of active transactions, number of blocked locks, or CPU activity). You can set thresholds on specific counters to generate alerts that notify operator.

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/performance-monitoring-and-tuning-tools?view=sql-server-2017>

NEW QUESTION 307 You have a database named DB1 that contains two tables. You need to encrypt one column in each

table by using the Always Encrypted feature. The solution must support groupings on encrypted columns. Which two actions should you perform? (Each correct answer presents part of the solution. Choose two.) A. Encrypt both columns by using deterministic encryption. B. Provision a symmetric key by using Transact-SQL.

C. Encrypt both columns by using randomized encryption. D. Provision column master keys and column encryption keys by using Microsoft SQL Server Management Studio (SSMS). Answer: AD Explanation:

A: Use deterministic encryption for columns that will be used as search or grouping parameters, for example a government ID number. Deterministic encryption always generates the same encrypted value for any given plain text value. Using deterministic encryption allows point lookups, equality joins, grouping and indexing on encrypted columns. D: Always Encrypted uses two types of keys: column encryption keys and column master keys. A column encryption key is used to encrypt data in an encrypted column. A column master key is a key-protecting key that encrypts one or more column encryption keys. Incorrect: Not B: A column encryption key (CEK), is a content encryption key (i.e. a key used to protect data) that is protected by a CMK. All Microsoft CMK store providers encrypt CEKs by using RSA with Optimal Asymmetric Encryption Padding (RSA-OAEP) with the default parameters specified by RFC 8017 in Section A.2.1. Not C: Randomized encryption uses a method that encrypts data in a less predictable manner. Randomized encryption is more secure, but prevents searching, grouping, indexing, and joining on encrypted columns.

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine?view=sql-server-2017>

17 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 Answer: G Explanation: An unbroken chain of log backups must be available, up to the current log file, and they must all be applied to bring the page up to date with the current log file. A tail-log backup captures any log records that have not yet been backed up (the tail of the log) to prevent work loss and to keep the log chain intact. Before you can recover a SQL Server database to its latest point in time, you must back up the tail of its transaction log. The tail-log backup will be the last backup of interest in the recovery plan for the database.

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/tail-log-backups-sql-server?view=sql-server-2017#TailLog Scenarios>

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/restore-pages-sql-server?view=sql-server-2017#Restrictions>

NEW QUESTION 311 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 that has multiple databases. A database named DB1 has multiple online file groups. It is configured to use the full recovery model. A full backup is performed nightly and transaction logs are performed on the hour. A large number of records are accidentally deleted at 17:20. You need to perform a point-in-time recovery. 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 Answer: G Explanation: To back up the tail of the log (that is, the active log), check Back up the tail of the log, and leave database in the restoring state. A tail-log backup is taken after a failure to back up the tail of the log in order to prevent work loss. Back up the active log (a tail-log backup) both after a failure, before beginning to restore the database, or when failing over to a secondary database. Selecting this option is equivalent to specifying the NORECOVERY option in the BACKUP LOG statement of Transact-SQL.

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/back-up-a-transaction-log-sql-server?view=sql-server-2017>

NEW QUESTION 312 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. You need to ensure that you can restore a database to a specific point in time between backups while minimizing the number of Azure storage containers required. 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 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 contain 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 Explanation:

https://docs.microsoft.com/en-us/sql/t-sql/statements/alter-database-transact-sql-set-options?view=sql-server-2017#auto_update_statistics NEW QUESTION 316 Hotspot A company has an on-premises Microsoft SQL Server environment and Microsoft Azure SQL Database instances. The environments host several customer databases. You configure an Always On Availability Group for a customer. You must create log reports for the customer that detail when the log is flushed to disk on the primary and secondary replica. You need to develop a report containing the requested information. In the table below, identify the log type that you should use for each replica. (NOTE: Make only one selection in each column. Each correct selection is worth one point.)

Log type
Log capture
Log hardened
Log receive/Log cache
Log flush

Answer:

Log type	Flush on primary	Flush on secondary
Log capture	<input type="radio"/>	<input type="radio"/>
Log hardened	<input type="radio"/>	<input checked="" type="radio"/>
Log receive/Log cache	<input type="radio"/>	<input type="radio"/>
Log flush	<input checked="" type="radio"/>	<input type="radio"/>

Explanation: ? Flush on primary: Log flush Log flush. Log data is generated and flushed to disk on the primary replica in preparation for replication to the secondary replica. It then enters the send queue. ? Flush on secondary: Log hardened The log is flushed on the secondary replica, and then a notification is sent to the primary replica to acknowledge completion of the transaction.

Incorrect: ? Not Log capture Log capture. Logs for each database are captured on the primary replica, compressed, and sent to the corresponding queue on the secondary replica. This process runs continuously as long as database replicas are connecting. If this process is not able to scan and enqueue the messages quickly enough, the log send queue continues to grow. ? Not Log receive/Log cache Log receive/Log cache. Each secondary replica gets messages from the primary replica and then caches the messages. http://www.futas.net/ora/doc/SQL_Server_2016_Higher_availability_eBook_EN_US.pdf NEW QUESTION 317

Drag and Drop You are designing a high availability (HA) environment for a company that has three office locations. Details of the services deployed at each office are shown in the table below:

Location	SQL Server edition	Number of databases	Number of nodes
Seattle	Enterprise	200	4
New York	Standard	50	2
Atlanta	Enterprise	120	4

You need to maximize availability, minimize data loss, and minimize downtime in the event of a failure. Which solution should you implement for each location? (To answer, drag the appropriate solutions to the correct locations. Each solution may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.)

Solutions

Always On Availability Groups SQL server
SQL Server failover cluster instance
Database mirroring
SQL Server replication

Answer Area	
Location	Solution
Seattle	Solution
New York	Solution
Atlanta	Solution

www.passleader.com

Answer:

Solutions

Always On Availability Groups SQL server
SQL Server failover cluster instance
Database mirroring
SQL Server replication

Answer Area	
Location	Solution
Seattle	Always On Availability Groups SQL server
New York	Database mirroring
Atlanta	Always On Availability Groups SQL server

www.passleader.com

Explanation: The Always On availability groups feature is a high-availability and disaster-recovery solution that provides an enterprise-level alternative to database mirroring.

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/always-on-availability-groups-sql-server?view=sql-server-2017> NEW QUESTION 318 ?? Download the newest PassLeader 70-764 dumps from passleader.com now! 100% Pass Guarantee! 70-764 PDF dumps & 70-764 VCE dumps: <http://www.passleader.com/70-764.html> (365 Q&As) (New Questions Are 100% Available and Wrong Answers Have Been Corrected! Free VCE simulator!) P.S. New 70-764 dumps PDF:

https://drive.google.com/open?id=0B-ob6L_QjGLpN3N6eHJ6Z2EzZWc >> New 70-761 dumps PDF:

https://drive.google.com/open?id=0B-ob6L_QjGLpaEZzRVFnOE9OenM >> New 70-762 dumps PDF:

https://drive.google.com/open?id=0B-ob6L_QjGLpN3RVQ25sVUM5dkU >> New 70-765 dumps PDF:

https://drive.google.com/open?id=0B-ob6L_QjGLpZHIHSG5KM09xUms >> New 70-767 dumps PDF:

https://drive.google.com/open?id=0B-ob6L_QjGLpcZXZWU14dHhIUVk >> New 70-768 dumps PDF:

https://drive.google.com/open?id=0B-ob6L_QjGLpeXAxaUJkWEZnVIU