



70-764^{Q&As}

Administering a SQL Database Infrastructure

Pass Microsoft 70-764 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.passapply.com/70-764.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





QUESTION 1

You work as a Database Administrator (DBA) for a company named ABC.com.

The company uses a Microsoft SQL Server 2012 infrastructure.

You are configuring a highly-available database solution using an AlwaysOn availability group on two servers running SQL Server 2012. The two servers are in separate datacenters.

The two datacenters are connected by a WAN link with a network latency of more than 200ms.

Which of the following failover types should you configure for the availability group?

- A. You should configure the asynchronous manual failover failover type.
- B. You should configure the synchronous manual failover failover type.
- C. You should configure the synchronous automatic failover failover type.
- D. You should configure the Asynchronous automatic failover failover type.

Correct Answer: A

QUESTION 2

You need to validate rows before they are added to a table every time a row is added using a user-defined function. What should you use? More than one answer may achieve the goal. Select the BEST answer.

- A. DML Trigger
- B. Default constraint
- C. Foreign key
- D. CHECK constraint

Correct Answer: D

QUESTION 3

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 Adventureworks that contains a single schema named ADVSchema . You must implement auditing for all objects in the ADVSchema 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 TSpinDB . Tailspin Toys requires a custom application that monitors TSpinDB 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 ConDB and is also running SQL Server Reporting Services (SSRS).
Wingtip Toys	DomainA	Wingtip Toys has a database named WingDB . 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 WingDB . You must use minimal logging and minimize data loss during the import process.
Wide World Importers	DomainB	The environment includes a database named WDWDB . Neither auditing nor statistics are configured for WDWDB . You must log any deletion of views and all database record update operations.

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.

- A. sys.dm_tran_locks
- B. sp_lock
- C. sys.dm_tran_active_snapshot_database_transactions
- D. Activity Monitor
- E. sp_monitor

Correct Answer: BD

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 Answers:

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

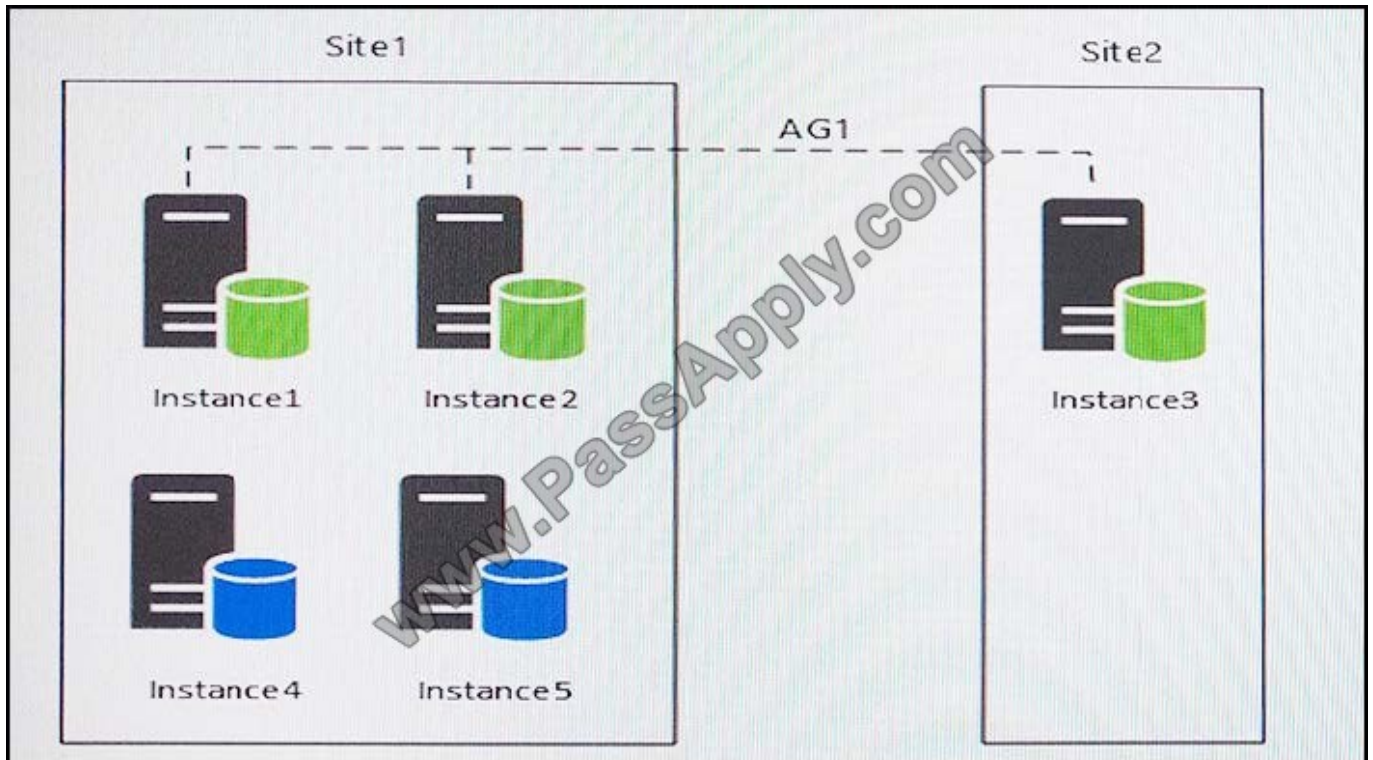
activity). You can set thresholds on specific counters to generate alerts that notify operator

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

QUESTION 4

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.

You have five servers that run Microsoft Windows 2012 R2. Each server hosts a Microsoft SQL Server instance. The topology for the environment is shown in the following diagram.



You have an Always On Availability group named AG1. The details for AG1 are shown in the following table.

Instance	Node type
Instance1	Primary
Instance2	Synchronous readable secondary
Instance3	Asynchronous readable secondary

Instance1 experiences heavy read-write traffic. The instance hosts a database named OperationsMain that is four terabytes (TB) in size. The database has multiple data files and filegroups. One of the filegroups is read_only and is half of the total database size.

Instance4 and Instance5 are not part of AG1. Instance4 is engaged in heavy read-write I/O.

Instance5 hosts a database named StagedExternal. A nightly BULK INSERT process loads data into an empty table that has a rowstore clustered index and two nonclustered rowstore indexes.

You must minimize the growth of the StagedExternal database log file during the BULK INSERT operations and perform point-in-time recovery after the BULK INSERT transaction.

Changes made must not interrupt the log backup chain.

You plan to add a new instance named Instance6 to a datacenter that is geographically distant from Site1 and Site2. You must minimize latency between the nodes in AG1.

All databases use the full recovery model. All backups are written to the network location \\SQLBackup\. A separate process copies backups to an offsite location. You should minimize both the time required to restore the databases and the space required to store backups. The recovery point objective (RPO) for each instance is shown in the following table.



Instance	Recovery point objective
Instance 1	5 minutes
Instance 2	5 minutes
Instance 3	5 minutes
Instance 4	60 minutes
Instance 5	24 hours

Full backups of OperationsMain take longer than six hours to complete. All SQL Server backups use the keyword COMPRESSION.

You plan to deploy the following solutions to the environment. The solutions will access a database named DB1 that is part of AG1.

Reporting system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader role. The user has EXECUTE permissions on the database. Queries make no changes to the data.

The queries must be load balanced over variable read-only replicas.

Operations system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader and db_datawriter roles.

The user has EXECUTE permissions on the database. Queries from the operations system will perform both DDL and DML operations.

The wait statistics monitoring requirements for the instances are described in the following table.

Instance	Description
Instance1	Aggregate wait statistics since the last server restart.
Instance4	Identify the most prominent wait types for all the commands originating from a session, between session connections, or between application pool resets.
Instance5	Identify all the wait types for queries currently running on the server.

You need to propose a new process for the StagedExternal database.

Which five actions should you recommended be performed in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:



Actions

- Drop all nonclustered indexes on the target table.
- Create a transaction log backup. Change the recovery model of **StagedExternal** to **SIMPLE**.
- Run the nightly import process.
- Change the recovery model of **StagedExternal** to **SIMPLE**.
- Change the recovery model of **StagedExternal** to **FULL**. Create a transaction log backup.
- Drop all clustered and nonclustered indexes on the target table.
- Recreate any dropped indexes on the target table.
- Create a transaction log backup. Change the recovery model of **StagedExternal** to **BULK_LOGGED**.

Answer Area



Correct Answer:



Actions

- Drop all nonclustered indexes on the target table.
- Create a transaction log backup. Change the recovery model of **StagedExternal** to **SIMPLE**.
- Change the recovery model of **StagedExternal** to **SIMPLE**.

Answer Area

- Create a transaction log backup. Change the recovery model of **StagedExternal** to **BULK_LOGGED**.
- Drop all clustered and nonclustered indexes on the target table.
- Run the nightly import process.
- Change the recovery model of **StagedExternal** to **FULL**. Create a transaction log backup.
- Recreate any dropped indexes on the target table



From scenario: Instance5 hosts a database named StagedExternal. A nightly BULK INSERT process loads data into an empty table that has a rowstore clustered index and two nonclustered rowstore indexes.

You must minimize the growth of the StagedExternal database log file during the BULK INSERT operations and perform point-in-time recovery after the BULK INSERT transaction.

Changes made must not interrupt the log backup chain.

All databases use the full recovery model.

References: [https://technet.microsoft.com/en-us/library/ms190421\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms190421(v=sql.105).aspx)

QUESTION 5

Background

Corporate Information



Fabrikam, Inc. is a retailer that sells electronics products on the Internet. The company has a headquarters site and one satellite sales office. You have been hired as the database administrator, and the company wants you to change the architecture of the Fabrikam ecommerce site to optimize performance and reduce downtime while keeping capital expenditures to a minimum. To help with the solution, Fabrikam has decided to use cloud resources as well as on-premise servers.

Physical Locations

All of the corporate executives, product managers, and support staff are stationed at the headquarters office. Half of the sales force works at this location. There is also a satellite sales office. The other half of the sales force works at the satellite office in order to have sales people closer to clients in that area. Only sales people work at the satellite location.

Problem Statement

To be successful, Fabrikam needs a website that is fast and has a high degree of system uptime. The current system operates on a single server and the company is not happy with the single point of failure this presents. The current nightly

backups have been failing due to insufficient space on the available drives and manual drive cleanup often needing to happen to get past the errors. Additional space will not be made available for backups on the HQ or satellite servers.

During your investigation, you discover that the sales force reports are causing significant contention.

Configuration

Windows Logins

The network administrators have set up Windows groups to make it easier to manage security. Users may belong to more than one group depending on their role. The groups have been set up as shown in the following table:

Group	Members
OurDomain\Management	All corporate executives
OurDomain\SalesStaff	All sales people
OurDomain\ProductionStaff	All product managers and support staff
OurDomain\AllUsers	Everyone
OurDomain\CustomerSupport	Customer support representatives

Server Configuration The IT department has configured two physical servers with Microsoft Windows Server 2012 R2 and SQL Server 2014 Enterprise Edition and one Windows Azure Server. There are two tiers of storage available for use by database files only a fast tier and a slower tier. Currently the data and log files are stored on the fast tier of storage only. If a possible use case exists, management would like to utilize the slower tier storage for data files. The servers are configured as shown in the following table:



Location	Server
Company headquarters	HQ_Server
Satellite sales office	Satellite_Server
Microsoft Windows Azure (cloud)	Cloud_File Server

Database

Currently all information is stored in a single database called ProdDB, created with the following script:

```
CREATE DATABASE ProdDB  
GO  
ALTER DATABASE ProdDB SET RECOVERY SIMPLE  
GO
```

The Product table is in the Production schema owned by the ProductionStaff Windows group. It is the main table in the system so access to information in the Product table should be as fast as possible. The columns in the Product table are defined as shown in the following table:

Column	Data type
ProductID	INT
ProductName	VARCHAR(100)
ProductDescription	VARCHAR(MAX)
ProductPrice	SMALLMONEY
QuantityOnHand	INT
ProductCost	SMALLMONEY
ProductSupplierID	INT

The SalesOrderDetail table holds the details about each sale. It is in the Sales schema owned by the SalesStaff Windows group. This table is constantly being updated, inserted into, and read. The columns in the SalesOrderDetail table are defined as shown in the following table:

Column	Data type
SalesOrderDetailID	INT
ProductID	INT
SalePrice	SMALLMONEY
SaleQuantity	INT

Database Issues

The current database does not perform well. Additionally, a recent disk problem caused the system to go down, resulting in lost sales revenue. In reviewing the current system, you found that there are no automated maintenance procedures.

The database is severely fragmented, and everyone has read and write access.



Requirements

Database

The database should be configured to maximize uptime and to ensure that very little data is lost in the event of a server failure. To help with performance, the database needs to be modified so that it can support in-memory data, specifically

for the Product table, which the CIO has indicated should be a memory optimized table. The auto-update statistics option is set off on this database. Only product managers are allowed to add products or to make changes to the name,

description, price, cost, and supplier. The changes are made in an internal database and pushed to the Product table in ProdDB during system maintenance time. Product managers and others working at the headquarters location also should

be able to generate reports that include supplier and cost information.

Customer data access

Customers access the company's website to order products, so they must be able to read product information such as name, description, and price from the Product table. When customers place orders, stored procedures called by the website

update product quantity on-hand values. This means the product table is constantly updated at random times.

Customer support data access

Customer support representatives need to be able to view and not update or change product information. Management does not want the customer support representatives to be able to see the product cost or any supplier information.

Sales force data access

Sales people at both the headquarters office and the satellite office must generate reports that read from the Product and SalesOrderDetail tables. No updates or inserts are ever made by sales people. These reports are run at random times

and there can be no reporting downtime to refresh the data set except during the monthly maintenance window.

The reports that run from the satellite office are process intensive queries with large data sets. Regardless of which office runs a sales force report, the SalesOrderDetail table should only return valid, committed order data; any orders not yet

committed should be ignored.

Historical Data

The system should keep historical information about customers who access the site so that sales people can see how frequently customers log in and how long they stay on the site.

The information should be stored in a table called Customer Access. Supporting this requirement should have minimal impact on production website performance.

Backups

The recovery strategy for Fabrikam needs to include the ability to do point in time restores and minimize the risk of data loss by performing transaction log backups every 15 minutes.

Database Maintenance



The company has defined a maintenance window every month when the server can be unavailable. Any maintenance functions that require exclusive access should be accomplished during that window.

Project milestones completed

Revoked all existing read and write access to the database, leaving the schema ownership in place.

Configured an Azure storage container secured with the storage account name MyStorageAccount with the primary access key StorageAccountKey on the cloud file server.

SQL Server 2014 has been configured on the satellite server and is ready for use.

On each database server, the fast storage has been assigned to drive letter F:,

and the slow storage has been assigned to drive letter D:.

You need to create the CustomerAccess table to support the reporting and performance requirements.

Develop the solution by selecting and arranging the required code blocks in the correct order. You may not need all of the code blocks.

Select and Place:



Code Blocks

```
CREATE TABLE [CustomerAccess]  
ON [CustomerAccessFG] (  

```

```
CustomerAccessID INT IDENTITY(1,1)  
NOT NULL PRIMARY KEY  
, CustomerID INT NOT NULL  
, LoginDate DATETIME NOT NULL  
, LogoffDate DATETIME NULL  

```

```
CREATE TABLE [CustomerAccess] (  

```

```
ALTER DATABASE [ProdDB] ADD FILE  
(NAME = N'ProdDB_CustomerAccess',  
FILENAME = N'F:\Data  
\ProdDB_CustomerAccess.ndf') TO  
FILEGROUP [CustomerAccessFG]
```

```
ALTER DATABASE [ProdDB] ADD FILE  
(NAME = N'ProdDB_CustomerAccess',  
FILENAME = N'D:\Data  
\ProdDB_CustomerAccess.ndf') TO  
FILEGROUP [CustomerAccessFG]
```

```
) ON [ProdDB_CustomerAccess]
```

```
) ON [CustomerAccessFG]
```

```
ALTER DATABASE [ProdDB] ADD FILE  
(NAME = N'ProdDB_CustomerAccess',  
FILENAME = N'D:\Data  
\ProdDB_CustomerAccess.ndf')
```

```
)
```

```
ALTER DATABASE [ProdDB] ADD  
FILEGROUP [CustomerAccessFG]
```

Answer Area

Correct Answer:



Code Blocks

```
CREATE TABLE [CustomerAccess] (
```

```
ALTER DATABASE [ProdDB] ADD FILE  
(NAME = N'ProdDB_CustomerAccess',  
FILENAME = N'D:\Data  
\ProdDB_CustomerAccess.ndf') TO  
FILEGROUP [CustomerAccessFG]
```

```
) ON [ProdDB_CustomerAccess]
```

```
ALTER DATABASE [ProdDB] ADD FILE  
(NAME = N'ProdDB_CustomerAccess',  
FILENAME = N'D:\Data  
\ProdDB_CustomerAccess.ndf')
```

```
)
```

Answer Area

```
ALTER DATABASE [ProdDB] ADD  
FILEGROUP [CustomerAccessFG]
```

```
ALTER DATABASE [ProdDB] ADD FILE  
(NAME = N'ProdDB_CustomerAccess',  
FILENAME = N'F:\Data  
\ProdDB_CustomerAccess.ndf') TO  
FILEGROUP [CustomerAccessFG]
```

```
CREATE TABLE [CustomerAccess]  
ON [CustomerAccessFG] (
```

```
CustomerAccessID INT IDENTITY(1,1)  
NOT NULL PRIMARY KEY  
, CustomerID INT NOT NULL  
, LoginDate DATETIME NOT NULL  
, LogoutDate DATETIME NULL
```

```
) ON [CustomerAccessFG]
```

[Latest 70-764 Dumps](#)

[70-764 PDF Dumps](#)

[70-764 Study Guide](#)



To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

Try our product !

100% Guaranteed Success
100% Money Back Guarantee
365 Days Free Update
Instant Download After Purchase
24x7 Customer Support
Average 99.9% Success Rate
More than 800,000 Satisfied Customers Worldwide
Multi-Platform capabilities - [Windows](#), [Mac](#), [Android](#), [iPhone](#), [iPod](#), [iPad](#), [Kindle](#)

We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications. You can view Vendor list of All Certification Exams offered:

<https://www.passapply.com/allproducts>

Need Help

Please provide as much detail as possible so we can best assist you.
To update a previously submitted ticket:



 <p>One Year Free Update Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p>	 <p>Money Back Guarantee To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>	 <p>Security & Privacy We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.</p>
---	---	--

Any charges made through this site will appear as Global Simulators Limited.
All trademarks are the property of their respective owners.
Copyright © passapply, All Rights Reserved.