

[New 70-762 Dumps Exam Pass 100% ! Braindump2go 70-762 PDF Dumps 120Q Instant Download] 82-92

2018/August Braindump2go Microsoft 70-762 Exam Dumps with PDF and VCE New Updated! Following are some new 70-762 Real Exam Questions: 1. | 2018 Latest 70-762 Exam Dumps (PDF & VCE) 120Q&As

Download: <https://www.braindump2go.com/70-762.html> | 2018 Latest 70-762 Exam Questions & Answers

Download: <https://drive.google.com/drive/folders/0B75b5xYLjSSNajNKbVh2RV9IZIU?usp=sharing> QUESTION 82 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 have a database named DB1. There is no memory-optimized filegroup in the database. You have a table and a stored procedure that were created by running the following Transact-SQL statements:

```
CREATE TABLE Employee
(
    EmployeeId int NOT NULL PRIMARY KEY,
    FirstName varchar(20),
    LastName varchar(20),
    Status char(1),
    Address varchar(100),
    Department int NOT NULL
);

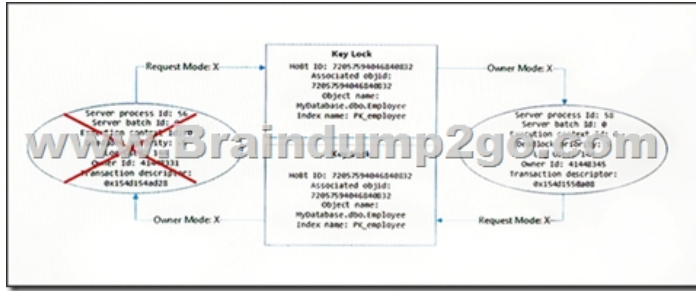
CREATE PROCEDURE usp_GetEmployeesByDepartment
(
    @LastName varchar(20)
)
AS
BEGIN
    SELECT e.FirstName, e.LastName,
    FROM Employee e
    JOIN Department d on e.DepartmentId = d.DepartmentId
    WHERE e.Status = 'T' AND e.LastName = @LastName
END;
```

The Employee table is persisted on disk. You add 2,000 records to the Employee table. You need to create an index that meets the following requirements: Optimizes the performance of the stored procedure. Covers all the columns required from the Employee table. Uses FirstName and LastName as included columns. Minimizes index storage size and index key size. What should you do? A. Create a clustered index on the table. B. Create a nonclustered index on the table. C. Create a nonclustered filtered index on the table. D. Create a clustered columnstore index on the table. E. Create a nonclustered columnstore index on the table. F. Create a hash index on the table. Answer: B Explanation: [https://technet.microsoft.com/en-us/library/jj835095\(v=sql.110\).aspx](https://technet.microsoft.com/en-us/library/jj835095(v=sql.110).aspx) QUESTION 83 Background You have a database named HR1 that includes a table named Employee. You have several read-only, historical reports that contain regularly changing totals. The reports use multiple queries to estimate payroll expenses. The queries run concurrently. Users report that the payroll estimate reports do not always run. You must monitor the database to identify issues that prevent the reports from running. You plan to deploy the application to a database server that supports other applications. You must minimize the amount of storage that the database requires. Employee Table You use the following Transact-SQL statements to create, configure, and populate the Employee table:

```
CREATE TABLE dbo.Employee
(
    EmployeeID INT PRIMARY KEY,
    LastName varchar(50),
    FirstName varchar(50),
    DepartmentID int,
    HireDate datetime,
    TerminationDate datetime,
    SupervisorID int,
    CostCenterNumber int,
    EmployeeStatus int,
    DepartmentID int
) GO
CREATE INDEX IX_1 on dbo.Employee (LastName, FirstName, DepartmentID) INCLUDE (HireDate)
CREATE INDEX IX_2 on dbo.Employee (LastName) INCLUDE (EmployeeID, FirstName, DepartmentID)
CREATE INDEX IX_3 on dbo.Employee (LastName, FirstName) INCLUDE (DepartmentID)
CREATE INDEX IX_4 on dbo.Employee (LastName, FirstName) INCLUDE (HireDate, DepartmentID)
GO
INSERT INTO Employee (EmployeeID, LastName, CostCenterNumber) VALUES(1001, 'Employee A', 3001001)
INSERT INTO Employee (EmployeeID, LastName, CostCenterNumber) VALUES(1002, 'Employee B', 3001001)
GO
```

Application You have an application that updates the Employees table. The application calls the following stored procedures simultaneously and asynchronously: UspA: This stored procedure updates only the EmployeeStatus column. UspB: This stored procedure updates only the EmployeePayRate column. The application uses views to control access to data. Views must meet the following requirements: Allow user access to all columns in the tables that the view accesses. Restrict updates to only the rows that

the view returns.Exhibit



You are analyzing the performance of the database environment. You discover that locks that are held for a long period of time as the reports are generated. You need to generate the reports more quickly. The database must not use additional resources. What should you do?
 A. Update the transaction level of the report query session to READPAST.
 B. Modify the report queries to use the UNION statement to combine the results of two or more queries.
 C. Set the READ_COMMITTED_SNAPSHOT database option to ON.
 D. Update the transaction level of the report query session to READ UNCOMMITTED.
 Answer: D
 Explanation: Transactions running at the READ UNCOMMITTED level do not issue shared locks to prevent other transactions from modifying data read by the current transaction. This is the least restrictive of the isolation levels.

[https://technet.microsoft.com/en-us/library/ms173763\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms173763(v=sql.105).aspx)
 QUESTION 84 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 questions 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 database that is 130 GB and contains 500 million rows of data. Granular transactions and mass batch data imports change the database frequently throughout the day. Microsoft SQL Server Reporting Services (SSRS) uses the database to generate various reports by using several filters. You discover that some reports time out before they complete. You need to reduce the likelihood that the reports will time out.
 Solution: You partition the largest tables.
 Does this meet the goal?
 A. Yes
 B. No
 Answer: B

QUESTION 85 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 questions 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 database that is 130 GB and contains 500 million rows of data. Granular transactions and mass batch data imports change the database frequently throughout the day. Microsoft SQL Server Reporting Services (SSRS) uses the database to generate various reports by using several filters. You discover that some reports time out before they complete. You need to reduce the likelihood that the reports will time out.
 Solution: You change the transaction log file size to expand dynamically in increments of 200 MB.
 Does this meet the goal?
 A. Yes
 B. No
 Answer: B

QUESTION 86 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 questions 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 database that is 130 GB and contains 500 million rows of data. Granular transactions and mass batch data imports change the database frequently throughout the day. Microsoft SQL Server Reporting Services (SSRS) uses the database to generate various reports by using several filters. You discover that some reports time out before they complete. You need to reduce the likelihood that the reports will time out.
 Solution: You create a file group for the indexes and a file group for the data files. You store the files for each file group on separate disks.
 Does this meet the goal?
 A. Yes
 B. No
 Answer: A

Explanation: Consider creating two additional File Groups: Tables and Indexes. It is best not to put your stuff in PRIMARY as that is where SQL SERVER stores all of its data and meta-data about your objects. You create your Table and Clustered Index (as that is the data for the table) on [Tables] and all Non-Clustered indexes on [Indexes].
 QUESTION 87 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 questions 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 3-TB database. The database server has 64 CPU cores. You plan to migrate the database to Microsoft Azure SQL Database. You need to select the service tier for the Azure SQL database. The solution must meet or exceed the current processing capacity.
 Solution: You select the Premium service tier.
 Does this meet the goal?
 A. Yes
 B. No
 Answer: A

Explanation: Premium service is required for 3 TB of storage. Single database DTU and storage limits

	Basic	Standard	Premium
Maximum DTUs	5	3000	4000

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-service-tiers-dtu> QUESTION 88 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 questions 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 3-TB database. The database server has 64 CPU cores. You plan to migrate the database to Microsoft Azure SQL Database. You need to select the service tier for the Azure SQL database. The solution must meet or exceed the current processing capacity. Solution: You select the Standard service tier. Does this meet the goal? A. Yes B. No Answer: B Explanation: Premium service is required for 3 TB of storage. Single database DTU and storage limits

	Basic	Standard	Premium
Maximum DTUs	5	3000	4000

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-service-tiers-dtu> QUESTION 89 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 questions 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 3-TB database. The database server has 64 CPU cores. You plan to migrate the database to Microsoft Azure SQL Database. You need to select the service tier for the Azure SQL database. The solution must meet or exceed the current processing capacity. Solution: You select the Basic service tier. Does this meet the goal? A. Yes B. No Answer: B Explanation: Premium service is required for 3 TB of storage. Single database DTU and storage limits

	Basic	Standard	Premium
Maximum DTUs	5	3000	4000

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-service-tiers-dtu> QUESTION 90 Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals. You have a table that has a clustered index and a nonclustered index. The indexes use different columns from the table. You have a query named Query1 that uses the nonclustered index. Users report that Query1 takes a long time to report results. You run Query1 and review the following statistics for an index seek operation:

Index Seek (NonClustered)
Scan a particular range of rows from a nonclustered index
Physical Operation
Logical Operation
Actual Execution Mode
Actual Number of Rows
Actual Number of Batches
Estimated I/O Cost
Estimated Operator Cost
Estimated CPU Cost
Estimated Subtree Cost
Estimated Number of Executions
Number of Executions
Estimated Number of Rows
Estimated Row Size
Actual Rebinds
Actual Rewinds
Ordered
Node ID

You need to resolve the performance issue. Solution: You drop the nonclustered index. Does the solution meet the goal? A. Yes B. No Answer: B Explanation: We see Actual Number of Row is 3571454, while Estimated Number of Rows is 0. This indicates that the statistics are old, and need to be updated. QUESTION 91 Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals. You have a table that has a clustered index and a nonclustered index. The indexes use different columns from the table. You have a query named Query1 that uses the nonclustered index. Users report that Query1 takes a long time to report results. You run Query1 and review the following statistics for an index seek operation:

Index Seek (NonClustered)	
Scan a particular range of rows from a nonclustered index.	
Physical Operation	Index Seek
Logical Operation	Index Seek
Actual Execution Mode	Row
Actual Number of Rows	3571454
Actual Number of Batches	0
Estimated I/O Cost	0.0093577
Estimated Operator Cost	0.0107304 (0%)
Estimated CPU Cost	0.0013727
Estimated Subtree Cost	0.0107304
Estimated Number of Executions	1
Number of Executions	8
Estimated Number of Rows	0
Estimated Row Size	19 B
Actual Rebinds	0
Actual Rewinds	0
Ordered	True
Node ID	100

You need to resolve the performance issue. Solution: You defragment both indexes. Does the solution meet the goal? A. Yes B. No
 Answer: B Explanation: We see Actual Number of Row is 3571454, while Estimated Number of Rows is 0. This indicates that the statistics are old, and need to be updated.
 QUESTION 92 Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals. You have a table that has a clustered index and a nonclustered index. The indexes use different columns from the table. You have a query named Query1 that uses the nonclustered index. Users report that Query1 takes a long time to report results. You run Query1 and review the following statistics for an index seek operation:

Index Seek (NonClustered)	
Scan a particular range of rows from a nonclustered index.	
Physical Operation	Index Seek
Logical Operation	Index Seek
Actual Execution Mode	Row
Actual Number of Rows	3571454
Actual Number of Batches	0
Estimated I/O Cost	0.0093577
Estimated Operator Cost	0.0107304 (0%)
Estimated CPU Cost	0.0013727
Estimated Subtree Cost	0.0107304
Estimated Number of Executions	1
Number of Executions	8
Estimated Number of Rows	0
Estimated Row Size	19 B
Actual Rebinds	0
Actual Rewinds	0
Ordered	True
Node ID	100

You need to resolve the performance issue. Solution: You update statistics for the nonclustered index. Does the solution meet the goal? A. Yes B. No
 Answer: A Explanation: We see Actual Number of Row is 3571454, while Estimated Number of Rows is 0. This indicates that the statistics are old, and need to be updated. !!!RECOMMEND!!!
 2018 Latest 70-762 Exam Dumps (PDF & VCE) 120Q&As Download: <https://www.braindump2go.com/70-762.html>
 2018 Latest 70-762 Study Guide Video: YouTube Video: [YouTube.com/watch?v=pCSFKe6Is58](https://www.youtube.com/watch?v=pCSFKe6Is58)