

## [New 70-767 Dumps 100% Valid 70-767 VCE and PDF Dumps 318Q Provided by Braindump2go[288-304]

2018/October Braindump2go 70-767 Exam Dumps with PDF and VCE New Updated Today! Following are some new 70-767 Real Exam Questions: 1. | 2018 Latest 70-767 Exam Dumps (VCE & PDF) 308Q&As

Download: <https://www.braindump2go.com/70-767.html> | 2. | 2018 Latest 70-767 Exam Questions & Answers Download:

<https://drive.google.com/drive/folders/0B75b5xYLjSSNN1RSdIN6Z0VwRjg?usp=sharing> QUESTION 288 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 information about products, sales, and orders for a manufacturing company. The instance contains a database that has two tables named SalesOrderHeader and SalesOrderDetail. SalesOrderHeader has 500,000 rows and SalesOrderDetail has 3,000,000 rows. Users report performance degradation when they run the following stored procedure:

```
CREATE PROCEDURE Sales.GetRecentSales (@date datetime)
AS BEGIN
    IF @date is NULL
        SET @date = DATEADD(MONTH, -1, (SELECT MAX(ORDERDATE) FROM Sales.SalesOrderHeader))
    SELECT * FROM Sales.SalesOrderDetail
    WHERE h.SalesOrderID = d.SalesOrderID
    AND h.OrderDate > @date
END
```

You need to optimize performance. Solution: You run the following Transact-SQL statement:

```
CREATE STATISTICS Stat1
WITH FULLSCAN
```

Does the solution meet the goal? A. Yes B. No Answer: A Explanation: UPDATE STATISTICS updates query optimization statistics on a table or indexed view. FULLSCAN computes statistics by scanning all rows in the table or indexed view. FULLSCAN and SAMPLE 100 PERCENT have the same results.

<https://docs.microsoft.com/en-us/sql/t-sql/statements/update-statistics-transact-sql?view=sql-server-2017> QUESTION 289

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 plan to deploy a Microsoft SQL server that will host a data warehouse named DB1. The server will contain four SATA drives configured as a RAID 10 array. You need to minimize write contention on the transaction log when data is being loaded to the database. Solution: You replace the SATA disks with SSD disks. Does this meet the goal? A. Yes B. No Answer: B Explanation: A data warehouse is too big to store on an SSD. Instead you should place the log file on a separate drive.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/policy-based-management/place-data-and-log-files-on-separate-drives?view=sql-server-2017> QUESTION 290 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 plan to deploy a Microsoft SQL server that will host a data warehouse named DB1. The server will contain four SATA drives configured as a RAID 10 array. You need to minimize write contention on the transaction log when data is being loaded to the database. Solution: You add more data files to DB1. Does this meet the goal? A. Yes B. No Answer: B Explanation: There is no performance gain, in terms of log throughput, from multiple log files. SQL Server does not write log records in parallel to multiple log files. Instead you should place the log file on a separate drive.

References: <https://www.red-gate.com/simple-talk/sql/database-administration/optimizing-transaction-log-throughput/>

<https://docs.microsoft.com/en-us/sql/relational-databases/policy-based-management/place-data-and-log-files-on-separate-drives?view=sql-server-2017> QUESTION 291 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 plan to deploy a Microsoft SQL server that will host a data warehouse named DB1. The server will contain four SATA drives configured as a RAID 10 array. You need to minimize write contention on the transaction log when data is being loaded to the database. Solution: You configure the server to automatically delete the transaction logs nightly. Does this meet the goal? A. Yes B. No Answer: B Explanation: You should place the log file on a separate drive. References:

<https://www.red-gate.com/simple-talk/sql/database-administration/optimizing-transaction-log-throughput/>

<https://docs.microsoft.com/en-us/sql/relational-databases/policy-based-management/place-data-and-log-files-on-separate-drives?view=sql-server-2017> QUESTION 292 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 Microsoft Azure SQL Data Warehouse instance. You run the following Transact-SQL statement:

```
SELECT CustomerKey, SUM(SalesAmt) TotalSales
FROM sales.FactOrders
GROUP BY CustomerKey
```

The query fails to return results. You need to determine why the query fails. Solution: You run the following Transact-SQL statements:

```
SELECT CustomerKey, SUM(SalesAmt) TotalSales
FROM sales.FactOrders
GROUP BY CustomerKey
OPTION (LABEL = 'TotalSales')
SELECT TOP 1 status, total_elapsed_time, submit_time
FROM sys.dm_pdw_exec_requests
WHERE [label] = 'TotalSales'
ORDER BY submit_time
```

Does the solution meet the goal? A. Yes B. No Answer: A Explanation: To use submit\_time we must use sys.dm\_pdw\_exec\_requests table, which holds information about all requests currently or recently active in SQL Data Warehouse. It lists one row per request/query. Label is an Optional label string associated with some SELECT query statements.

<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-pdw-exec-requests-transact-sql?view=aps-pdw-2016-au7> QUESTION 293 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 Microsoft Azure SQL Data Warehouse instance. You run the following Transact-SQL statement:

```
SELECT CustomerKey, SUM(SalesAmt) TotalSales
FROM sales.FactOrders
GROUP BY CustomerKey
```

The query fails to return results. You need to determine why the query fails. Solution: You run the following Transact-SQL statement:

```
SELECT TOP 1 status, total_elapsed_time, submit_time
FROM sales.FactOrders
WHERE [label] = 'TotalSales'
ORDER BY submit_time
```

Does the solution meet the goal? A. Yes B. No Answer: B Explanation: To use submit\_time we must use sys.dm\_pdw\_exec\_requests table.

<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-pdw-exec-requests-transact-sql?view=aps-pdw-2016-au7>QUESTION 294Note: 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 Microsoft Azure SQL Data Warehouse instance. You run the following Transact-SQL statement:


```
SELECT CustomerKey, SUM(SalesAmt) TotalSales
FROM sales.FactOrders
GROUP BY CustomerKey
```

The query fails to return results.You need to determine why the query fails.Solution: You run the following Transact-SQL statements:

```
SELECT CustomerKey, SUM(SalesAmt) TotalSales
FROM sales.FactOrders
GROUP BY CustomerKey
OPTION (QueryId = 'Totalsales')
SELECT TOP 1 status, total_elapsed_time, submit_time
FROM sys.dm_pdw_exec_requests
WHERE [QueryId] = 'Totalsales'
ORDER BY submit_time
```

Does the solution meet the goal?A. YesB. NoAnswer: BExplanation:We must use Label, not QueryID in the WHERE clause.

<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-pdw-exec-requests-transact-sql?view=aps-pdw-2016-au7>QUESTION 295Note: 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 uses Microsoft SQL Server to deploy a data warehouse to an environment that has a SQL Server Analysis Services (SSAS) instance. The data warehouse includes the Fact.Order table as shown in the following table definition. The table has no indexes.



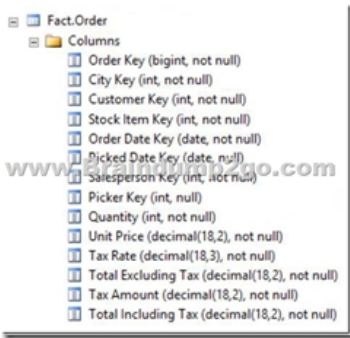
Column Name	Data Type	Nullability
Order Key	bigint	not null
City Key	int	not null
Customer Key	int	not null
Stock Item Key	int	not null
Order Date Key	date	not null
Picked Date Key	date	null
Salesperson Key	int	not null
Picker Key	int	null
Quantity	int	not null
Unit Price	decimal(18,2)	not null
Tax Rate	decimal(18,3)	not null
Total Excluding Tax	decimal(18,2)	not null
Tax Amount	decimal(18,2)	not null
Total Including Tax	decimal(18,2)	not null

You must minimize the amount of space that indexes for the Fact.Order table consume. You run the following queries frequently. Both queries must be able to use a columnstore index:

```
SELECT AVG([Tax Amount]) AS [Average Tax Amount]
FROM Fact.Order
WHERE [Order Date Key] BETWEEN '20150701' AND '20151231'
SELECT SUM([Total Excluding Tax]) AS [Total Revenue]
FROM Fact.Order
WHERE [Order Date Key] BETWEEN '20150701' AND '20151231'
```

You need to ensure that the queries complete as quickly as possible.Solution: You create a measure for the Fact.Order table.Does the solution meet the goal?A. YesB. NoAnswer: BExplanation:You should use a columnstore index.Columnstore indexes are the standard for storing and querying large data warehousing fact tables. This index uses column-based data storage and query processing to achieve gains up to 10 times the query performance in your data warehouse over traditional row-oriented storage.

<https://docs.microsoft.com/en-us/sql/relational-databases/indexes/columnstore-indexes-overview?view=sql-server-2017>QUESTION 296Note: 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 uses Microsoft SQL Server to deploy a data warehouse to an environment that has a SQL Server Analysis Services (SSAS) instance. The data warehouse includes the Fact.Order table as shown in the following table definition. The table has no indexes.



You must minimize the amount of space that indexes for the Fact.Order table consume. You run the following queries frequently. Both queries must be able to use a columnstore index: You need to ensure

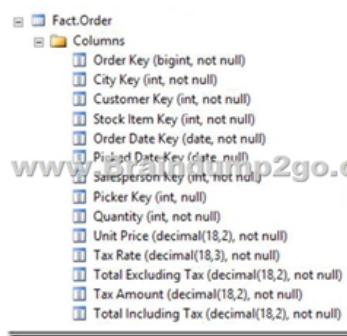
```
SELECT AVG([Tax Amount]) AS [Average Tax Amount]
FROM Fact.Order
WHERE [Order Date Key] BETWEEN '20150701' AND '20151231'

SELECT SUM([Total Excluding Tax]) AS [Total Revenue]
FROM Fact.Order
WHERE [Order Date Key] BETWEEN '20150701' AND '20151231'
```

that the queries complete as quickly as possible. Solution: You create two nonclustered indexes. The first includes the [Order Date Key] and [Tax Amount] columns. The second will include the [Order Date Key] and [Total Excluding Tax] columns. Does the solution meet the goal? A. Yes B. No Answer: B Explanation: You should use a columnstore index. Columnstore indexes are the standard for storing and querying large data warehousing fact tables. This index uses column-based data storage and query processing to achieve gains up to 10 times the query performance in your data warehouse over traditional row-oriented storage.

<https://docs.microsoft.com/en-us/sql/relational-databases/indexes/columnstore-indexes-overview?view=sql-server-2017>

QUESTION 297 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 uses Microsoft SQL Server to deploy a data warehouse to an environment that has a SQL Server Analysis Services (SSAS) instance. The data warehouse includes the Fact.Order table as shown in the following table definition. The table has no indexes.



You must minimize the amount of space that indexes for the Fact.Order table consume. You run the following queries frequently. Both queries must be able to use a columnstore index:

```
SELECT AVG([Tax Amount]) AS [Average Tax Amount]
FROM Fact.Order
WHERE [Order Date Key] BETWEEN '20150701' AND '20151231'

SELECT SUM([Total Excluding Tax]) AS [Total Revenue]
FROM Fact.Order
WHERE [Order Date Key] BETWEEN '20150701' AND '20151231'
```

You need to ensure that the queries complete as quickly as possible. Solution: You create one columnstore index that includes the [Order Date Key], [Tax Amount], and [Total Excluding Tax] columns. Does the solution meet the goal? A. Yes B. No Answer: A Explanation: You should use a columnstore index. Columnstore indexes are the standard for storing and querying large data warehousing fact tables. This index uses column-based data storage and query processing to achieve gains up to 10 times the query performance in your data warehouse over traditional row-oriented storage.

<https://docs.microsoft.com/en-us/sql/relational-databases/indexes/columnstore-indexes-overview?view=sql-server-2017>

QUESTION 298 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 Microsoft SQL server that has Data Quality Services (DQS) installed. You need to review the completeness and the uniqueness of the data stored in the matching policy. Solution: You profile the data. Does this meet the goal? A. Yes B. No Answer: B Explanation: Use a matching rule.

<https://docs.microsoft.com/en-us/sql/data-quality-services/create-a-matching-policy?view=sql-server-2017> QUESTION 299

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 Microsoft SQL server that has Data Quality Services (DQS) installed. You need to review the completeness and the uniqueness of the data stored in the matching policy. Solution: You create a matching rule. Does this meet the goal? A. Yes B. No Answer: B Explanation: Use a matching rule, and use completeness and uniqueness data to determine what weight to give a field in the matching process. If there is a high level of uniqueness in a field, using the field in a matching policy can decrease the matching results, so you may want to set the weight for that field to a relatively small value. If you have a low level of uniqueness for a column, but low completeness, you may not want to include a domain for that column.

<https://docs.microsoft.com/en-us/sql/data-quality-services/create-a-matching-policy?view=sql-server-2017> QUESTION 300

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 Microsoft SQL server that has Data Quality Services (DQS) installed. You need to review the completeness and the uniqueness of the data stored in the matching policy. Solution: You modify the weight of the domain in the matching rule. Does this meet the goal? A. Yes B. No Answer: A Explanation: Use a matching rule, and use completeness and uniqueness data to determine what weight to give a field in the matching process. If there is a high level of uniqueness in a field, using the field in a matching policy can decrease the matching results, so you may want to set the weight for that field to a relatively small value. If you have a low level of uniqueness for a column, but low completeness, you may not want to include a domain for that column.

<https://docs.microsoft.com/en-us/sql/data-quality-services/create-a-matching-policy?view=sql-server-2017> QUESTION 301

You have a server that has Data Quality Services (DQS) installed. You create a matching policy that contains one matching rule. You need to configure the Similarity of Similar percentage that defines a match. Which similarity percentage will always generate a similarity score of 0? A. 55 B. 80 C. 70 D. 75 Answer: A Explanation: The minimum similarity between the values of a field is 60%. If the calculated matching score for a field of two records is less than 60, the similarity score is automatically set to 0.

<https://docs.microsoft.com/en-us/sql/data-quality-services/create-a-matching-policy?view=sql-server-2017> QUESTION 302

You need to ensure that a downstream system can consume data in a Master Data Services (MDS) system. What should you configure? A. a Data Collector B. a knowledge base C. a matching policy D. a subscription view Answer: D Explanation: Subscription views to consume your master data.

<https://docs.microsoft.com/en-us/sql/master-data-services/master-data-services-overview-mds?view=sql-server-2017>

QUESTION 303 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 that has change data capture enabled. A Microsoft SQL Server Integration Services (SSIS) job runs once weekly. The job loads changes from DB1 to a data warehouse by querying the change data capture tables. You remove the Integration Services job. You need to stop tracking changes to the database. The solution must remove all the change data capture configurations from DB1. Which stored procedure should you execute? A. catalog.deploy\_project B. catalog.restore\_project C. catalog.stop\_operation D. sys.sp\_cdc\_add\_job E. sys.sp\_cdc\_change\_job F. sys.sp\_cdc\_disable\_db G. sys.sp\_cdc\_enable\_db H. sys.sp\_cdc\_stop\_job Answer: F Explanation: sys.sp\_cdc\_disable\_db disables change data capture for all tables in the database currently enabled. All system objects related to change data capture, such as change tables, jobs, stored procedures and functions, are dropped.

<https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sys-sp-cdc-disable-db-transact-sql?view=sql-server-2017> 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 have a database named DB1 that has change

data capture enabled. A Microsoft SQL Server Integration Services (SSIS) job runs once weekly. The job loads changes from DB1 to a data warehouse by querying the change data capture tables. Users report that an application that uses DB1 is suddenly unresponsive. You discover that the Integration Services job causes severe blocking issues in the application. You need to ensure that the users can run the application as quickly as possible. Your SQL Server login is a member of only the `ssis_admin` database role. Which stored procedure should you execute? A. `catalog.deploy_project` B. `catalog.restore_project` C. `catalog.stop_operation` D. `sys.sp_cdc_add_job` E. `sys.sp_cdc_change_job` F. `sys.sp_cdc_disable_db` G. `sys.sp_cdc_enable_db` H. `sys.sp_cdc_stop_job` Answer: E Explanation: `sys.sp_cdc_change_job` modifies the configuration of a change data capture cleanup or capture job in the current database.

<https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sys-sp-cdc-change-job-transact-sql?view=sql-server-2017>!!!RECOMMEND!!!1.[2018 Latest 70-767 Exam Dumps (VCE & PDF) 308Q&As

Download:<https://www.braindump2go.com/70-767.html>2.[2018 Latest 70-767 Study Guide Video: YouTube Video:

[YouTube.com/watch?v=WKRvjLr9hww](https://www.youtube.com/watch?v=WKRvjLr9hww)