

## [March-2019-NewBraindump2go 70-765 Exam Dumps and 70-765 VCE 192Q Free Offered

**2019/March Braindump2go 70-765 Exam Dumps with PDF and VCE New Updated Today! Following are some new 70-765 Real Exam Questions:]1.[2019 Latest 70-765 Exam Dumps (PDF & VCE) Instant Download:**

**<https://www.braindump2go.com/70-765.html>2.[2019 Latest 70-765 Exam Questions & Answers Instant Download:**

**<https://drive.google.com/drive/folders/0B75b5xYLjSSNTnR6dFR2U3A5cFk?usp=sharing>**  
New QuestionNote: 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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.You are tuning the performance of a virtual machines that hosts a Microsoft SQL Server instance.The virtual machine originally had four CPU cores and now has 32 CPU cores.The SQL Server instance uses the default settings and has an OLTP database named db1. The largest table in db1 is a key value store table named table1.Several reports use the PIVOT statement and access more than 100 million rows in table1.You discover that when the reports run, there are PAGELATCH\_IO waits on PFS pages 2:1:1, 2:2:1, 2:3:1, and 2:4:1 within the tempdb database.You need to prevent the PAGELATCH\_IO waits from occurring.Solution: You add more tempdb databases.Does this meet the goal?A. YesB. NoAnswer: BExplanation:From SQL Server's perspective, you can measure the I/O latency from sys.dm\_os\_wait\_stats. If you consistently see high waiting for PAGELATCH\_IO, you can benefit from a faster I/O subsystem for SQL Server.A cause can be poor design of your database - you may wish to split out data located on 'hot pages', which are accessed frequently and which you might identify as the causes of your latch contention. For example, if you have a currency table with a data page containing 100 rows, of which 1 is updated per transaction and you have a transaction rate of 200/sec, you could see page latch queues of 100 or more. If each page latch wait costs just 5ms before clearing, this represents a full half-second delay for each update. In this case, splitting out the currency rows into different tables might prove more performant (if less normalized and logically structured).References:

**<https://www.mssqltips.com/sqlservertip/3088/explanation-of-sql-server-io-and-latches/>**  
New QuestionNote: 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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.You are tuning the performance of a virtual machines that hosts a Microsoft SQL Server instance.The virtual machine originally had four CPU cores and now has 32 CPU cores.The SQL Server instance uses the default settings and has an OLTP database named db1. The largest table in db1 is a key value store table named table1.Several reports use the PIVOT statement and access more than 100 million rows in table1.You discover that when the reports run, there are PAGELATCH\_IO waits on PFS pages 2:1:1, 2:2:1, 2:3:1, and 2:4:1 within the tempdb database.You need to prevent the PAGELATCH\_IO waits from occurring.Solution: You add more files to db1.Does this meet the goal?A. YesB. NoAnswer: AExplanation:From SQL Server's perspective, you can measure the I/O latency from sys.dm\_os\_wait\_stats. If you consistently see high waiting for PAGELATCH\_IO, you can benefit from a faster I/O subsystem for SQL Server.A cause can be poor design of your database - you may wish to split out data located on 'hot pages', which are accessed frequently and which you might identify as the causes of your latch contention. For example, if you have a currency table with a data page containing 100 rows, of which 1 is updated per transaction and you have a transaction rate of 200/sec, you could see page latch queues of 100 or more. If each page latch wait costs just 5ms before clearing, this represents a full half-second delay for each update. In this case, splitting out the currency rows into different tables might prove more performant (if less normalized and logically structured).References:

**<https://www.mssqltips.com/sqlservertip/3088/explanation-of-sql-server-io-and-latches/>**  
New QuestionNote: 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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.You are migrating an on-premises Microsoft SQL Server instance to SQL Server on a Microsoft Azure virtual machine. The instance has 30 databased that consume a total of 2 TB of disk space. The instance sustains more than 30,000 transactions per second.You need to provision storage for the virtual machine. The storage must be able to support the same load as the on-premises deployment.Solution: You create one storage account that has 30 containers. You create a VHD in each container.Does this meet the goal?A. YesB. NoAnswer: BExplanation:Each Storage Account handles up to 20.000 IOPS, and 500TB of data.References:  
**<https://www.tech-coffee.net/understand-microsoft-azure-storage-for-virtual-machines/>**  
New QuestionYou have a SQL server

2016 server that has several databases. You are designing a performance monitoring strategy for the server. You need to recommend a solution to log and track the performance of store procedures. The solution must minimize the resource usage on the server. What should you include in the recommendation? A. SQL server extended events B. Activity monitor C. SQL server profiler D. SQL Trace

**Answer: D** New Question You need to recommend a solution to minimize the amount of disk resources used by Query1.

What should you recommend? A. Rebuild the index by using FILLFACTOR=100 B. Add quantity as an included column to IX\_TransHistory\_ProductID C. Refactor the query into a partitioned view D. Add quantity as a filtered index

**Answer: D** New Question You have a database named DB1 that has the following:- a Table named Customers that contains a list of customers - a table named Invoice that contains invoice information. There is foreign key relationship between Invoice and Customer. You need to create a stored procedure that will retrieve a list of all the invoice for a set of customers. The stored procedure must allow the caller to specify an unlimited list of customers. What should you use to create the stored procedure? A. An optional parameter B. A table-valued parameter C. A nonclustered index D. A Multiple active -result set(MARS)

**Answer: B** New Question You have a database that contains a table named Customers. You need to retrieve a set of customers by using multiple nested SELECT statements. The solution must meet the following requirements:- The result can be used in a FROM statement- The result must be reusable. What should you use to retrieve the set of customers? A. A CLR stored procedure B. A view C. An unsigned store procedure D. A scalar function

**Answer: B** New Question You have a SQL Server 2016 database that contains a large table named table1. Table1 is split into three filegroups. Each filegroups is on a separate server. You need to write a query against Table1 that includes data from all three filegroups. What should you use? A. A filtered index B. A columnstore index C. A distributed reply controller D. A distributed view

**Answer: D** New Question You create two databases named DB1 and DB2 in DB1, you create a table named table1. In DB2 you create table2. In DB1 you create a view named View1. View1 reads data from table1 and table2.

You grant a group named Group1 the SELECT permission to View1. You need to recommend a solution to ensure that the members of Group1 can execute View1 Successfully. What should you include in recommendation? A. Signed store procedure B. Mixed mode authentication C. Cress-database ownership chaining D. Windows authentication

**Answer: C** Explanation:

<https://www.mssqltips.com/sqlservertip/1782/understanding-cross-database-ownership-chaining-in-sql-server/> New Question

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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are migrating an on-premises Microsoft SQL Server instance to SQL Server on a Microsoft Azure virtual machine. The instance has 30 databased that consume a total of 2 TB of disk space. The instance sustains more than 30,000 transactions per second. You need to provision storage for the virtual machine. The storage must be able to support the same load as the on-premises deployment.

Solution: You create 30 storage accounts that each has one container. You create a VHD in each container. Does this meet the goal? A. Yes B. No

**Answer: A** Explanation: Each Storage Account handles up to 20,000 IOPS, and 500TB of data. References: <https://www.tech-coffee.net/understand-microsoft-azure-storage-for-virtual-machines/>

New Question 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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are migrating an on-premises Microsoft SQL Server instance to SQL Server on a Microsoft Azure virtual machine. The instance has 30 databased that consume a total of 2 TB of disk space. The instance sustains more than 30,000 transactions per second. You need to provision storage for the virtual machine. The storage must be able to support the same load as the on-premises deployment.

Solution: You create one storage account that has one container. You create multiple VHDs in the container. Does this meet the goal? A. Yes B. No

**Answer: B** Explanation: Each Storage Account handles up to 20,000 IOPS, and 500TB of data. References: <https://www.tech-coffee.net/understand-microsoft-azure-storage-for-virtual-machines/>

!!!RECOMMEND!!!1. **[2019 Latest 70-765 Exam Dumps (PDF & VCE) Instant Download:**

<https://www.braindump2go.com/70-765.html>2. **[2019 Latest 70-765 Study Guide Video Instant Download:** YouTube Video:

[YouTube.com/watch?v=Qzw2xt-YEIO](https://www.youtube.com/watch?v=Qzw2xt-YEIO)