

December/2018 Braindump2go 70-762 Exam VCE and PDF 145Q Free Offer[Q121-Q132]

December/2018 Braindump2go 70-762 Exam Dumps with PDF and VCE New Updated Today! Following are some new 70-762 Real Exam Questions:]1.[2018 Latest 70-762 Exam Dumps (PDF & VCE) 145Q&As Download:

<https://www.braindump2go.com/70-762.html>2.[2018 Latest 70-762 Exam Questions & Answers Download:

<https://drive.google.com/drive/folders/0B75b5xYLjSSNajNKbVh2RV9IZIU?usp=sharing>
QUESTION 121Note: 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 increase the number of log files for the database. You store the log files across multiple disks.Does this meet the goal?A. YesB. NoAnswer: BExplanation:Instead, create a file group for the indexes and a file group for the data files.
QUESTION 122Note: 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 a database that contains the following tables: BlogCategory, BlogEntry, ProductReview, Product, and SalesPerson. The tables were created using the following Transact SQL statements:

```
CREATE TABLE BlogCategory
(
    CategoryID int PRIMARY KEY,
    CategoryName varchar(50)
);

CREATE TABLE BlogEntry
(
    EntryID int PRIMARY KEY,
    EntryTitle varchar(100),
    CategoryID int FOREIGN KEY REFERENCES BlogCategory(CategoryID)
);

CREATE TABLE Product
(
    ProductID int PRIMARY KEY,
    ProductName varchar(100),
    ReviewValue float
);

CREATE TABLE ProductReview
(
    ProductID int FOREIGN KEY REFERENCES Product(ProductID),
    ReviewerID int FOREIGN KEY REFERENCES SalesPerson(SalesID),
    ReviewText varchar(255)
);
```

You must modify the ProductReview Table to meet the following requirements:- The table must reference the ProductID column in the Product table- Existing records in the ProductReview table must not be validated with the Product table.- Deleting records in the Product table must not be allowed if records are referenced by the ProductReview table.- Changes to records in the Product table must propagate to the ProductReview table.You also have the following database tables: Order, ProductTypes, and SalesHistory. The transact-SQL statements for these tables are not available.You must modify the Orders table to meet the following requirements:- Create new rows in the table without granting INSERT permissions to the table.- Notify the sales person who places an order whether or not the order was completed.You must add the following constraints to the SalesHistory table:- a constraint on the SaleID column that allows the field to be used as a record identifier- a constant that uses the ProductID column to reference the Product column of the ProductTypes table- a constraint on the CategoryID column that allows one row with a null value in the column- a constraint that limits the SalePrice column to values greater than four- Finance department users must be able to retrieve data from the SalesHistory table for sales persons where the value of the SalesYTD column is above a certain threshold.You plan to create a memory-optimized table named SalesOrder. The table must meet the following requirements:- The table must hold 10 million unique sales orders.- The table must use checkpoints to minimize I/O operations and must not use transaction logging.- Data loss is acceptable.Performance for queries against the SalesOrder table that use WHERE clauses with exact equality operations must

be optimized. You need to modify the environment to meet the requirements for the Orders table. What should you create? A. an AFTER UPDATE trigger B. a user-defined function C. a stored procedure with output parameters D. an INSTEAD OF INSERT trigger
Answer: D
Explanation: From Question: You must modify the Orders table to meet the following requirements: Create new rows in the table without granting INSERT permissions to the table. Notify the sales person who places an order whether or not the order was completed.
References:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-trigger-transact-sql?view=sql-server-2017> QUESTION 123
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:

Scan a parti

Physical Operation
Logical Operation
Actual Execution
Actual Number of
Actual Number of
Estimated I/O Cost
Estimated Operation
Estimated CPU Cost
Estimated Subtree
Estimated Number
Number of Execution
Estimated Number
Estimated Row Size
Actual Rebinds
Actual Rewinds
Ordered
Node ID

You need to resolve the performance issue. Solution: You rebuild the clustered index. Does the solution meet the goal? A. Yes B. No
Answer: B
Explanation: The query uses the nonclustered index, so improving the clustered index will not help. We should update statistics for the nonclustered index. QUESTION 124
You need to create a view that can be indexed. You write the following statement.

```
01 create view myView2
02
03 as
04 select col1, col2
05 from dbo.test
06 go
```

What should you add at line 02? A. with check_option B. with recompile C. with view_metadata D. with schemabinding
Answer: D
Explanation: The following steps are required to create an indexed view and are critical to the successful implementation of the indexed view: 1. Verify the SET options are correct for all existing tables that will be referenced in the view. 2. Verify that the SET options for the session are set correctly before you create any tables and the view. 3. Verify that the view definition is deterministic. 4. Create the view by using the WITH SCHEMABINDING option. 5. Create the unique clustered index on the view.
References:

<https://docs.microsoft.com/en-us/sql/relational-databases/views/create-indexed-views?view=sql-server-2017> QUESTION 125
You have a nonpartitioned table that has a single dimension. The table is named dim.Products.Projections. The table is queried frequently by several line-of-business applications. The data is updated frequently throughout the day by two processes. Users report that when they query data from dim.Products.Projections, the responses are slower than expected. The issue occurs when a large number of rows are being updated. You need to prevent the updates from slowing down the queries. What should you do? A. Use the NOLOCK option. B. Execute the DBCC UPDATEUSAGE statement. C. Use the max worker threads option. D. Use a

table-valued parameter.E. Set SET ALLOW_SNAPSHOT_ISOLATION to ON.Answer: AExplanation:The NOLOCK hint allows SQL to read data from tables by ignoring any locks and therefore not being blocked by other processes.This can improve query performance, but also introduces the possibility of dirty reads.References:

<https://www.mssqltips.com/sqlservertip/2470/understanding-the-sql-server-nolock-hint/>QUESTION 126Your company runs end-of-the-month accounting reports. While the reports run, other financial records are updated in the database.Users report that the reports take longer than expected to run.You need to reduce the amount of time it takes for the reports to run. The reports must show committed data only.What should you do?A. Use the NOLOCK option.B. Execute the DBCC UPDATEUSAGE statement.C. Use the max worker threads option.D. Use a table-valued parameter.E. Set SET ALLOW_SNAPSHOT_ISOLATION to ON.F. Set SET XACT_ABORT to ON.G. Execute the ALTER TABLE T1 SET (LOCK_ESCALATION = AUTO); statement.H. Use the OUTPUT parameters.Answer: EExplanation:Snapshot isolation enhances concurrency for OLTP applications.Once snapshot isolation is enabled, updated row versions for each transaction are maintained in tempdb. A unique transaction sequence number identifies each transaction, and these unique numbers are recorded for each row version. The transaction works with the most recent row versions having a sequence number before the sequence number of the transaction. Newer row versions created after the transaction has begun are ignored by the transaction.References:

<https://docs.microsoft.com/en-us/dotnet/framework/data/adonet/sql/snapshot-isolation-in-sql-server>QUESTION 127You have several real-time applications that constantly update data in a database. The applications run more than 400 transactions per second that insert and update new metrics from sensors.A new web dashboard is released to present the data from the sensors. Engineers report that the applications take longer than expected to commit updates.You need to change the dashboard queries to improve concurrency and to support reading uncommitted data.What should you do?A. Use the NOLOCK option.B. Execute the DBCC UPDATEUSAGE statement.C. Use the max worker threads option.D. Use a table-valued parameter.E. Set SET ALLOW_SNAPSHOT_ISOLATION to ON.F. Set SET XACT_ABORT to ON.G. Execute the ALTER TABLE T1 SET (LOCK_ESCALATION = AUTO); statement.H. Use the OUTPUT parameters.Answer: AExplanation:The NOLOCK hint allows SQL to read data from tables by ignoring any locks and therefore not being blocked by other processes.This can improve query performance, but also introduces the possibility of dirty reads.Incorrect Answers:F: When SET XACT_ABORT is ON, if a Transact-SQL statement raises a run-time error, the entire transaction is terminated and rolled back.G: DISABLE, not AUTO, would be better.There are two more lock escalation modes: AUTO and DISABLE.The AUTO mode enables lock escalation for partitioned tables only for the locked partition. For non-partitioned tables it works like TABLE.The DISABLE mode removes the lock escalation capability for the table and that is important when concurrency issues are more important than memory needs for specific tables.Note: SQL Server's locking mechanism uses memory resources to maintain locks. In situations where the number of row or page locks increases to a level that decreases the server's memory resources to a minimal level, SQL Server's locking strategy converts these locks to entire table locks, thus freeing memory held by the many single row or page locks to one table lock. This process is called lock escalation, which frees memory, but reduces table concurrency.References:

<https://www.mssqltips.com/sqlservertip/2470/understanding-the-sql-server-nolock-hint/>QUESTION 128You have multiple stored procedures inside a transaction.You need to ensure that all the data modified by the transaction is rolled back if a stored procedure causes a deadlock or times out.What should you do?A. Use the NOLOCK option.B. Execute the DBCC UPDATEUSAGE statement.C. Use the max worker threads option.D. Use a table-valued parameter.E. Set SET ALLOW_SNAPSHOT_ISOLATION to ON.F. Set SET XACT_ABORT to ON.G. Execute the ALTER TABLE T1 SET (LOCK_ESCALATION = AUTO); statement.H. Use the OUTPUT parameters.Answer: FExplanation:When SET XACT_ABORT is ON, if a Transact-SQL statement raises a run-time error, the entire transaction is terminated and rolled back.References:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/set-xact-abort-transact-sql?view=sql-server-2017>QUESTION 129You run the following Transact-SQL statements: Records m

```
CREATE TABLE Customers (
    CustomerID INT NOT NULL IDENTITY PRIMARY KEY CLUSTERED,
    CustomerName NVARCHAR (100) UNIQUE NOT NULL
)

CREATE TABLE Orders (
    OrderID INT NOT NULL IDENTITY PRIMARY KEY CLUSTERED,
    CustomerID INT NOT NULL FOREIGN KEY REFERENCES Customers (CustomerID)
)

CREATE VIEW v_CustomerOrder
AS SELECT
    b.CustomerName, a.OrderID, a.OrderDate,
    (SELECT COUNT(*) FROM Orders c WHERE c.CustomerID = a.CustomerID) AS CustomerOrderCount
FROM Orders a
INNER JOIN Customers b ON a.CustomerID = b.CustomerID
```

ust only be added to the Orders table by using the view. If a customer name does not exist, then a new customer name must be

created. You need to ensure that you can insert rows into the Orders table by using the view. A. Add the CustomerID column from the Orders table and the WITH CHECK OPTION statement to the view. B. Create an INSTEAD OF trigger on the view. C. Add the WITH SCHEMABINDING statement to the view statement and create a clustered index on the view. D. Remove the subquery from the view, add the WITH SCHEMABINDING statement, and add a trigger to the Orders table to perform the required logic. Answer: A Explanation: The WITH CHECK OPTION clause forces all data-modification statements executed against the view to adhere to the criteria set within the WHERE clause of the SELECT statement defining the view. Rows cannot be modified in a way that causes them to vanish from the view. References: <http://www.informit.com/articles/article.aspx?p=130855&seqNum=4> QUESTION 130 You run the following Transact-SQL statement:

```
CREATE TABLE OrderLines(  
    OrderLineID INT NOT NULL IDENTITY PRIMARY KEY CLUSTERED,  
    OrderID INT NOT NULL,  
    StockItemID INT NOT NULL,  
    Description NVARCHAR(100) NOT NULL,  
    Quantity INT NOT NULL,  
    UnitPrice DECIMAL(18, 2) NULL  
)
```

There are multiple unique OrderID values. Most of the UnitPrice values for the same OrderID are different. You need to create a single index seek query that does not use the following operators: Nested loop Sort Key lookup Which Transact-SQL statement should you run? A. CREATE INDEX IX_OrderLines_1 ON OrderLines (OrderID, UnitPrice) INCLUDE (Description, Quantity) B. CREATE INDEX IX_OrderLines_1 ON OrderLines (OrderID, UnitPrice) INCLUDE (Quantity) C. CREATE INDEX IX_OrderLines_1 ON OrderLines (OrderID, UnitPrice, Quantity) D. CREATE INDEX IX_OrderLines_1 ON OrderLines (UnitPrice, OrderID) INCLUDE (Description, Quantity) Answer: A Explanation: An index with nonkey columns can significantly improve query performance when all columns in the query are included in the index either as key or nonkey columns. Performance gains are achieved because the query optimizer can locate all the column values within the index; table or clustered index data is not accessed resulting in fewer disk I/O operations. Note: All data types except text, ntext, and image can be used as nonkey columns. Incorrect Answers: C: Redesign nonclustered indexes with a large index key size so that only columns used for searching and lookups are key columns. D: The most unique column should be the first in the index. References:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-index-transact-sql?view=sql-server-2017> QUESTION 131 You are developing an ETL process to cleanse and consolidate incoming data. The ETL process will use a reference table to identify which data must be cleansed in the target table. The server that hosts the table restarts daily. You need to minimize the amount of time it takes to execute the query and the amount of time it takes to populate the reference table. What should you do? A. Convert the target table to a memory-optimized table. Create a natively compiled stored procedure to cleanse and consolidate the data. B. Convert the reference table to a memory-optimized table. Set the DURABILITY option to SCHEMA_AND_DATA. C. Create a native compiled stored procedure to implement the ETL process for both tables. D. Convert the reference table to a memory-optimized table. Set the DURABILITY option to SCHEMA_ONLY. Answer: D Explanation: If you use temporary tables, table variables, or table-valued parameters, consider conversions of them to leverage memory-optimized tables and table variables to improve performance. In-Memory OLTP provides the following objects that can be used for memory-optimizing temp tables and table variables: Memory-optimized tables Durability = SCHEMA_ONLY Memory-optimized table variables Must be declared in two steps (rather than inline): CREATE TYPE my_type AS TABLE ...; , then DECLARE @mytablevariable my_type;. References:

<https://docs.microsoft.com/en-us/sql/relational-databases/in-memory-oltp/faster-temp-table-and-table-variable-by-using-memory-optimization?view=sql-server-2017> QUESTION 132 You are designing a stored procedure for a database named DB1. The following requirements must be met during the entire execution of the stored procedure: - The stored procedure must only read changes that are persisted to the database. - SELECT statements within the stored procedure should only show changes to the data that are made by the stored procedure. You need to configure the transaction isolation level for the stored procedure. Which Transact-SQL statement or statements should you run? A. SET TRANSACTION ISOLATION LEVEL READ UNCOMMITTED ALTER DATABASE DB1 SET READ_COMMITTED_SNAPSHOT ON B. SET TRANSACTION ISOLATION LEVEL READ COMMITTED ALTER DATABASE DB1 SET READ_COMMITTED_SNAPSHOT OFF C. SET TRANSACTION ISOLATION LEVEL SERIALIZABLE SET TRANSACTION ISOLATION LEVEL READ UNCOMMITTED ALTER DATABASE SET READ_COMMITTED_SNAPSHOT OFF Answer: B Explanation: READ COMMITTED specifies that statements cannot read data that has been modified but not committed by other transactions. This prevents dirty reads. Data can be changed by other transactions between individual statements within the current transaction, resulting in nonrepeatable reads or phantom data. This option is the

SQL Server default. Incorrect Answers: A, D: READ UNCOMMITTED specifies that statements can read rows that have been modified by other transactions but not yet committed. References:

<https://docs.microsoft.com/en-us/dotnet/framework/data/adonet/sql/snapshot-isolation-in-sql-server>!!!RECOMMEND!!!

1. | 2018 Latest 70-762 Exam Dumps (PDF & VCE) 145Q&As Download: <https://www.braindump2go.com/70-762.html> 2. | 2018

Latest 70-762 Study Guide Video: YouTube Video: [YouTube.com/watch?v=60BAyOf9gTo](https://www.youtube.com/watch?v=60BAyOf9gTo)