


[2017-New-Exams70-762 PDF and VCE(Full Version)70q Download in Braindump2go[Q31-Q40

2017 Apr. New Microsoft 70-762 Exam Dumps with PDF and VCE Updated! Free Instant Download 70-762 PDF and VCE Dumps from www.Braindump2go.com Today! 100% Real Exam Questions! 100% Exam Pass Guaranteed! 1. |2017 New Version 70-762 Exam Dumps (PDF & VCE) 70Q&As Download: <http://www.braindump2go.com/70-762.html> 2. |2017 New Version 70-762 Exam Questions & Answers Download: <https://1drv.ms/f/s!AvI7wzKf6QBjghMIDqu1BwVHuSOI> QUESTION 31 Drag and Drop Question You have a database named Sales that contains the following database tables. Customer, Order, and Products. The Products table and the order table shown in the following diagram.



The Customer table includes a column that stores the date for the last order that the customer placed. You plan to create a table named Leads. The Leads table is expected to contain approximately 20,000 records. Storage requirements for the Leads table must be minimized. You need to begin to modify the table design to adhere to third normal form. Which column should you remove for each table? To answer? drag the appropriate column names to the correct locations. Each column name may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content. Answer: QUESTION 32 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 NOT NULL PRIMARY KEY,
    CategoryName nvarchar (20)
);

CREATE TABLE BlogEntry
(
    Entry int NOT PRIMARY KEY,
    Entrytitle nvarchar (50),
    Category int NOT NULL FOREIGN KEY REFERENCES BlogCategory
(CategoryID)
);

CREATE TABLE dbo.ProductReview
(
    ProductID int NOT NULL PRIMARY KEY,
    Product int NOT NULL,
    Review varchar (1000) NOT NULL
);

CREATE TABLE dbo.Product
(
    ProductID int Identity(1,1) PRIMARY KEY,
    Name varchar(1000) NOT NULL
);

CREATE TABLE dbo.SalesPerson
(
    SalesPersonID int IDENTITY(1,1) PRIMARY KEY,
    Name varchar (1000) NOT NULL,
    SalesID Money
);
```

You must modify the ProductReview Table to meet the following requirements: 1. The table must reference the ProductID column in the Product table. 2. Existing records in the ProductReview table must not be validated with the Product table. 3. Deleting records in the Product table must not be allowed if records are referenced by the ProductReview table. 4. 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: 1. Create new rows in the table without granting INSERT permissions to the table. 2. 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 Sale Price 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 enable referential integrity for the ProductReview table. How should you complete the relevant Transact-SQL statement? To answer, select the appropriate Transact-SQL segments in the answer area.

Answer Area

Alter Table dbo.ProductReview

Braindump2go.com
REFERENCES Product
ON DELETE NO ACTION

Select two alternatives. A. For the first selection select: WITH CHECK. For the first selection select : WITH NOCHECK. For the second selection select: ON DELETE NO ACTION ON UPDATE CASCADE. For the second selection select: ON DELETE CASCADE ON UPDATE CASCADE. For the second selection select: ON DELETE NO ACTION ON UPDATE NO ACTION. For the second selection select: ON DELETE CASCADE ON UPDATE NO ACTION. Answer: BC QUESTION 33 Hotspot Question 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 NOT NULL PRIMARY KEY,
    CategoryName nvarchar (20)
);

CREATE TABLE BlogEntry
(
    Entry int NOT PRIMARY KEY,
    EntryTitle nvarchar (50),
    Category int NOT NULL FOREIGN KEY REFERENCES BlogCategory
(CategoryID)
);

CREATE TABLE dbo.ProductReview
(
    Product int NOT NULL,
    Review varchar (1000) NOT NULL
);

CREATE TABLE dbo.Product
(
    ProductID int Identity(1,1) PRIMARY KEY,
    Name varchar(1000) NOT NULL
);

CREATE TABLE dbo.SalesPerson
(
    SalesPersonID int IDENTITY(1,1) PRIMARY KEY,
    Name varchar (1000) NOT NULL,
    SalesID Money
)
```

You must modify the ProductReview Table to meet the following requirements: 1. The table must reference the ProductID column in the Product table. 2. Existing records in the ProductReview table must not be validated with the Product table. 3. Deleting records in the Product table must not be allowed if records are referenced by the ProductReview table. 4. 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: 1. Create new rows in the table without granting INSERT permissions to the table. 2. 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 Sale Price 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 create an object that allows finance users to be able to retrieve the required data. The object must not have a negative performance impact. How should you complete the Transact-SQL statements? To answer, select the appropriate Transact-SQL segments in the answer area. Answer: QUESTION 34 You have a reporting database that includes a non-partitioned fact table named Fact_Sales. The table is persisted on disk. Users report that their queries take a long time to complete. The system administrator reports that the table takes too much space in the database. You observe that there are no indexes defined on the table, and many columns have repeating values. You need to create the most efficient index on the table, minimize disk storage and improve reporting query performance. 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 column store index on the table. E. Create a nonclustered column store index on the table. F. Create a hash index on the table. Answer: D QUESTION 35 You have a database named DB1. The database does not use a memory-optimized filegroup. The database contains a table named Table1. The table must support the following workloads:

Workload	Type	Description
Reporting	Existing	The reporting workload must scan most of the records in the table to aggregate on a number of columns. A clustered columnstore index is not supported on this workload. Rows are identified by using two columns. The filter is variant on one of the two columns while constant on the other. Only a small number of records with a few columns are returned by the query.

You need to add the most efficient index to support the new OLTP workload, while not deteriorating the existing Reporting query performance. 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 column store index on the table. E. Create a nonclustered column store index on the table. F. Create a hash index on the table. Answer: C

QUESTION 36 Drag and Drop Question You are evaluating the performance of a database environment. You must avoid unnecessary locks and ensure that lost updates do not occur. You need to choose the transaction isolation level for each data scenario. Which isolation level should you use for each scenario? To answer, drag the appropriate isolation levels to the correct scenarios. Each isolation may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content. Answer: QUESTION 37 Drag and Drop Question You have two database tables. Table1 is a partitioned table and Table 2 is a nonpartitioned table. Users report that queries take a long time to complete. You monitor queries by using Microsoft SQL Server Profiler. You observe lock escalation for Table1 and Table 2. You need to allow escalation of Table1 locks to the partition level and prevent all lock escalation for Table2. Which Transact-SQL statement should you run for each table? To answer, drag the appropriate Transact-SQL statements to the correct tables. Each command may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content. Answer: QUESTION 38 Drag and Drop Question You have a database that contains three encrypted stored procedures named dbo.Proc1, dbo.Proc2 and dbo.Proc3. The stored procedures include INSERT, UPDATE, DELETE and BACKUP DATABASE statements. You have the following requirements:- You must run all the stored procedures within the same transaction.- You must automatically start a transaction when stored procedures include DML statements.- You must not automatically start a transaction when stored procedures include DDL statements. You need to run all three stored procedures. Which four Transact-SQL segments should you use to develop the solution? To answer, move the appropriate Transact-SQL segments to the answer area and arrange them in the correct order. Answer: QUESTION 39 Hotspot Question You are profiling a frequently used database table named UserEvents. The READ_COMMITTED_SNAPSHOT database option is set to OFF. In the trace results, you observe that lock escalation occurred for one stored procedure even though the number of locks in the database did not exceed memory or configuration thresholds. Events details are provided in the following table:

Attribute	Value
EventClass	LockEscalation
Mode	5 - X (X represents exclusive)
ObjectID	274800017 (objectID of table UserEvents)
Type	5 - OBJECT
TextData	SET LOCKTIME ON; SET TRANSACTION ISOLATION LEVEL READ COMMITTED; LABEL_DELETE: DELETE FROM UserEvents WHERE (ROWLOCK) WHERE EventType = @EventType; END;

You need to modify the uspDeleteEvents stored procedure to avoid lock escalation. How should you modify the stored procedure? To answer, select the appropriate Transact-SQL segments in the answer area. Answer: QUESTION 40 You have a database that is experiencing deadlock issues when users run queries. You need to ensure that all deadlocks are recorded in XML format. What should you do? A. Create a Microsoft SQL Server Integration Services package that uses sys.dm_tran_locks. B. Enable trace flag 1224 by using the Database Consistency Checker (BDCC). C. Enable trace flag 1222 in the startup options for Microsoft SQL Server. D. Use the Microsoft SQL Server Profiler Lock: Deadlock event class. Answer: C !!!RECOMMEND!!! 1. | 2017 New Version 70-762 Exam Dumps (PDF & VCE) 70Q&As Download: <http://www.braindump2go.com/70-762.html> 2. | 2017 New Version 70-762 Study Guide Video: YouTube Video: [YouTube.com/watch?v=nYdYpxuZ0DU](https://www.youtube.com/watch?v=nYdYpxuZ0DU)