

[New 70-762 Dumps 100% Success-Braindump2go 70-762 Exam Questions PDF 120Q Instant Download[93-103

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>2. |2018 Latest 70-762 Exam Questions & Answers

Download:<https://drive.google.com/drive/folders/0B75b5xYLjSSNajNKbVh2RV9IZIU?usp=sharing> QUESTION 93 You have a reporting application that uses a table named Table1. You deploy a new batch update process to perform updates to Table1. The environment is configured with the following properties:- The database is configured with the default isolation setting.- The application and process use the default transaction handling. You observe the application cannot access any rows that are in use by the process. You have the following requirements:- Ensure the application is not blocked by the process.- Ensure the application has a consistent view of the data- Ensure the application does not read dirty data. You need to resolve the issue and meet the requirements with the least amount of administrative effort. What should you do? A. Enable the database for the ALLOW_SNAPSHOT_ISOLATION isolation level. Modify the application for the SERIALIZABLE isolation level. B. Enable the database for the READ_COMMITTED_SNAPSHOT isolation level. C. Enable the application for the WITH (NOLOCK) hint. D. Enable the database for the ALLOW_SNAPSHOT_ISOLATION isolation level. Modify the application and the update process for the SNAPSHOT isolation level. Answer: B Explanation: Snapshot isolation must be enabled by setting the ALLOW_SNAPSHOT_ISOLATION ON database option before it is used in transactions. This activates the mechanism for storing row versions in the temporary database (tempdb). READ COMMITTED is the default isolation level for SQL Server. It prevents dirty reads by specifying that statements cannot read data values that have been modified but not yet committed by other transactions. Other transactions can still modify, insert, or delete data between executions of individual statements within the current transaction, resulting in non-repeatable reads, or "phantom" data. Incorrect Answers: A: SERIALIZABLE is the most restrictive isolation level, because it locks entire ranges of keys and holds the locks until the transaction is complete. It encompasses REPEATABLE READ and adds the restriction that other transactions cannot insert new rows into ranges that have been read by the transaction until the transaction is complete.

<https://docs.microsoft.com/en-us/dotnet/framework/data/adonet/sql/snapshot-isolation-in-sql-server> QUESTION 94 You run the following Transact-SQL following statement:

```
CREATE TABLE Customer (  
    CustomerId INT IDENTITY (1, 1) PRIMARY KEY,  
    Code CHAR(5) NOT NULL,  
    FirstName VARCHAR (50) NOT NULL,  
    LastName VARCHAR (50) NOT NULL  
)
```

Customer records may be inserted individually or in bulk from an application. You observe that the application attempts to insert duplicate records. You must ensure that duplicate records are not inserted and bulk insert operations continue without notifications. Which Transact-SQL statement should you run? A. CREATE UNIQUE NONCLUSTERED INDEX IX_Customer_Code ON Customer (Code) WITH (ONLINE = OFF) B. CREATE UNIQUE INDEX IX_CUSTOMER_Code O Customer (Code) WITH (IGNORE_DUP_KEY = ON) C. CREATE UNIQUE INDEX IX Customer Code ON Customer (Code) WITH (IGNORE DUP KEY =OFF) D. CREATE UNIQUE NONCLUSTERED INDEX IX_Customer_Code ON Customer (Code) E. CREATE UNIQUE NONCLUSTERED INDEX IX_Customer_Code ON Customer (Code) WITH (ONLINE = ON) Answer: B Explanation: IGNORE_DUP_KEY = { ON | OFF } specifies the error response when an insert operation attempts to insert duplicate key values into a unique index. The IGNORE_DUP_KEY option applies only to insert operations after the index is created or rebuilt. The option has no effect when executing CREATE INDEX, ALTER INDEX, or UPDATE. The default is OFF. Incorrect Answers: ONLINE = { ON | OFF } specifies whether underlying tables and associated indexes are available for queries and data modification during the index operation. The default is OFF.

<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-index-transact-sql?view=sql-server-2017> QUESTION 95 You suspect deadlocks on a database. Which two trace flags in the Microsoft SQL Server error log should you locate? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point. A. 1204 B. 1211 C. 1222 D. 2528 E. 3205 Answer: AC Explanation: Trace flag 1204 returns the resources and types of locks participating in a deadlock and also the current command affected. Trace flag 1222 returns the resources and types of locks that are participating in a deadlock and also the current command affected, in an XML format that does not comply with any XSD schema.

<https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-traceon-trace-flags-transact-sql?view=sql-server>

-2017QUESTION 96 You have the following stored procedure that is called by other stored procedures and applications:

```
CREATE PROCEDURE U
AS
BEGIN
    EXEC Proc=FA
    DECLARE @Statu
    SELECT @Statu
    @CustomerId
END
```

You need to modify the stored procedure to meet the following requirements: Always return a value to the caller. Return 0 if @Status is NULL. Callers must be able to use @Status as a variable. Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Replace NULL values with 0. Add a PRINT statement to return @Status.

B. Add a RETURN statement.

C. Replace NULL values with 0. Add an output parameter to return @Status.

D. Replace NULL values with 0. Add a SELECT statement to return @Status.

E. Add a PRINT statement.

F. Add a SELECT statement to return @Status.

G. Add an output parameter to return @Status.

Answer: BC

Explanation: There are three ways of returning data from a procedure to a calling program: result sets, output parameters, and return codes.

<https://docs.microsoft.com/en-us/sql/relational-databases/stored-procedures/return-data-from-a-stored-procedure?view=sql-server-2017>

QUESTION 97 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 the series. Information and details provided in a question apply only to that question. You are developing an application to track customer sales. You need to create an object that meets to following requirements:- Run managed code packaged in an assembly that was created in the Microsoft.NET Framework and uploaded in Microsoft SQL Server.- Run written a transaction and roll back if a failure occurs.- Run when a table is created or modified. What should you create?

A. extended procedure

B. CLR procedure

C. user-defined procedure

D. DML trigger

E. DDL trigger

F. scalar-valued function

G. table-valued function

Answer: B

Explanation: The common language runtime (CLR) is the heart of the Microsoft .NET Framework and provides the execution environment for all .NET Framework code. Code that runs within the CLR is referred to as managed code. With the CLR hosted in Microsoft SQL Server (called CLR integration), you can author stored procedures, triggers, user-defined functions, user-defined types, and user-defined aggregates in managed code.

<https://docs.microsoft.com/en-us/dotnet/framework/data/adonet/sql/introduction-to-sql-server-clr-integration>

QUESTION 98 You are creating the following two stored procedures:- A natively-compiled stored procedure- An interpreted stored procedure that accesses both disk-based and memory-optimized tables Both stored procedures run within transactions. You need to ensure that cross-container transactions are possible. Which setting or option should you use?

A. the SET TRANSACTION_READ_COMMITTED isolation level for the connection

B. the SERIALIZABLE table hint on disk-based tables

C. the SET MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT=ON option for the database

D. the SET MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT=OFF option for the database

Answer: C

Explanation: Provide a supported isolation level for the memory-optimized table using a table hint, such as WITH (SNAPSHOT). The need for the WITH (SNAPSHOT) hint can be avoided through the use of the database option MEMORY_OPTIMIZED_ELEVATE_TO_SNAPSHOT. When this option is set to ON, access to a memory-optimized table under a lower isolation level is automatically elevated to SNAPSHOT isolation. Incorrect Answers: B: Accessing memory optimized tables using the READ COMMITTED isolation level is supported only for autocommit transactions. It is not supported for explicit or implicit transactions.

<https://docs.microsoft.com/en-us/sql/relational-databases/in-memory-oltp/transactions-with-memory-optimized-tables?view=sql-server-2017>

QUESTION 99 You are developing a database reporting solution for a table that contains 900 million rows and is 103 GB. The table is updated thousands of times a day, but data is not deleted. The SELECT statements vary in the number of columns used and the amount of rows retrieved. You need to reduce the amount of time it takes to retrieve data from the table. The must prevent data duplication. Which indexing strategy should you use?

A. a nonclustered index for each column in the table

B. a clustered columnstore index for the table

C. a hash index for the table

D. a clustered index for the table and nonclustered indexes for nonkey columns

Answer: B

Explanation: Columnstore indexes are the standard for storing and querying large data warehousing fact tables. It uses column-based data storage and query processing to achieve up to 10x query performance gains in your data warehouse over traditional row-oriented storage. A clustered columnstore index is the physical storage for the entire table. Generally, you should define the clustered index key with as few columns as possible. A nonclustered index contains the index key values and row locators that point to the storage location of the table data. You can create multiple

nonclustered indexes on a table or indexed view. Generally, nonclustered indexes should be designed to improve the performance of frequently used queries that are not covered by the clustered index.

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

QUESTION 100Note: 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 the series.

Information and details provided in a question apply only to that question. You are developing an application to track customer sales. You need to create a database object that meets the following requirements:- Launch when table data is modified.- Evaluate the state a table before and after a data modification and take action based on the difference.- Prevent malicious or incorrect table data operations.- Prevent changes that violate referential integrity by cancelling the attempted data modification.- Run managed code packaged in an assembly that is created in the Microsoft.NET Framework and located into Microsoft SQL Server. What should you create? A. extended procedure B. CLR procedure C. user-defined procedure D. DDL trigger E. scalar-valued function F. table-valued function

Answer: B Explanation: You can create a database object inside SQL Server that is programmed in an assembly created in the Microsoft .NET Framework common language runtime (CLR). Database objects that can leverage the rich programming model provided by the CLR include DML triggers, DDL triggers, stored procedures, functions, aggregate functions, and types. Creating a CLR trigger (DML or DDL) in SQL Server involves the following steps: Define the trigger as a class in a .NET Framework-supported language. For more information about how to program triggers in the CLR, see CLR Triggers. Then, compile the class to build an assembly in the .NET Framework using the appropriate language compiler. Register the assembly in SQL Server using the CREATE ASSEMBLY statement. For more information about assemblies in SQL Server, see Assemblies (Database Engine). Create the trigger that references the registered assembly.

<https://msdn.microsoft.com/en-us/library/ms179562.aspx> **QUESTION 101** You use Microsoft SQL Server Profiler to evaluate a query named Query1. The Profiler report indicates the following issues:- At each level of the query plan, a low total number of rows are processed.- The query uses many operations. This results in a high overall cost for the query. You need to identify the information that will be useful for the optimizer. What should you do? A. Start a SQL Server Profiler trace for the event class Performance statistics in the Performance event category. B. Create one Extended Events session with the sqlserver.missing_column_statistics event added. C. Start a SQL Server Profiler trace for the event class Soft Warnings in the Errors and Warnings event category. D. Create one Extended Events session with the sqlserver.error_reported event added. Answer: A

Explanation: The Performance Statistics event class can be used to monitor the performance of queries, stored procedures, and triggers that are executing. Each of the six event subclasses indicates an event in the lifetime of queries, stored procedures, and triggers within the system. Using the combination of these event subclasses and the associated sys.dm_exec_query_stats, sys.dm_exec_procedure_stats and sys.dm_exec_trigger_stats dynamic management views, you can reconstitute the performance history of any given query, stored procedure, or trigger.


<https://docs.microsoft.com/en-us/sql/relational-databases/event-classes/performance-statistics-event-class?view=sql-server-2017>

QUESTION 102 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,
Rating int)
GO
CREATE INDEX IX_1 ON dbo.Employee (LastName)
CREATE INDEX IX_2 ON dbo.Employee (LastName)
CREATE INDEX IX_3 ON dbo.Employee (LastName)
CREATE INDEX IX_4 ON dbo.Employee (LastName)
GO
INSERT INTO Employee (EmployeeID, LastName, Co
INSERT INTO Employee (EmployeeID, LastName, Co
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 all FROM clauses of the DML statements to use the IGNORE_CONSTRAINTS table hint.
 B. Modify the report queries to use the UNION statement to combine the results of two or more queries.
 C. Apply a nonclustered index to all tables used in the report queries.
 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 103
 Drag and Drop Question
 You have a Microsoft Azure SQL Database named MyDb that uses server version V12. You plan to use Query Performance Insight to troubleshoot performance problems. The database query store is not enabled. You need to enable the database query store to meet the following requirements for the database:- Statistics must be aggregated every 15 minutes.- Query stores must use no more than 1,024 megabytes (MB) of storage.- Query information must be retained for at least 15 days.- Queries must be captured based on resource consumption. You connect to the database by using SQL Server Managements Studio. How should you complete the Transact-SQL statements? To answer, drag the appropriate Transact-SQL segments to the correct locations. Each Transact-SQL segment 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. NOTE: More than one combination of answer choices is correct. You will receive credit for any of the correct combinations you select. Each correct selection is worth one point.

Transact-SQL segments	Answer area
STALE_QUERY_THRESHOLD_DAYS	ALTER DATABASE MyDb QUERY_STORE = <input type="text" value="Transact-SQL segment"/>
OFF	ALTER DATABASE MyDb SET QUERY_STORE
AUTO	(
ALL	SIZE_Based_CLEANUP_MODE = <input type="text" value="Transact-SQL segment"/>
DATA_FLUSH_INTERVAL_SECONDS	CAPTURE_MODE = <input type="text" value="Transact-SQL segment"/>
INTERVAL_LENGTH_MINUTES	QUERY_CAPTURE_MODE = <input type="text" value="Transact-SQL segment"/>
ON	MAX_STORAGE_SIZE_MB = 1024
)

Answer:

Transact-SQL segments	Answer area
STALE_QUERY_THRESHOLD_DAYS	ALTER DATABASE MyDb QUERY_STORE = <input type="text" value="ON"/>
OFF	ALTER DATABASE MyDb SET QUERY_STORE
AUTO	(
ALL	SIZE_Based_CLEANUP_MODE = AUTO
DATA_FLUSH_INTERVAL_SECONDS	CAPTURE_MODE = <input type="text" value="ALL"/>
INTERVAL_LENGTH_MINUTES	QUERY_CAPTURE_MODE = <input type="text" value="ALL"/>
ON	MAX_STORAGE_SIZE_MB = 1024
)

Explanation: <https://msdn.microsoft.com/en-us/library/mt604821.aspx>!!!RECOMMEND!!!!1. | 2018 Latest 70-762 Exam Dumps (PDF & VCE) 120Q&As Download: <https://www.braindump2go.com/70-762.html> 2. | 2018 Latest 70-762 Study Guide Video: YouTube Video: [YouTube.com/watch?v=pCSFK6fIs58](https://www.youtube.com/watch?v=pCSFK6fIs58)