



Vendor: Microsoft

Exam Code: 70-762

Exam Name: Developing SQL Databases

Version: Demo

QUESTION 1

DRAG DROP

You are analyzing the performance of a database environment.

Applications that access the database are experiencing locks that are held for a large amount of time. You are experiencing isolation phenomena such as dirty, nonrepeatable and phantom reads.

You need to identify the impact of specific transaction isolation levels on the concurrency and consistency of data.

What are the consistency and concurrency implications of each transaction isolation level?

To answer, drag the appropriate isolation levels to the correct locations. Each isolation level 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.

Isolation levels

read committed

serializable

read uncommitted

repeatable read

<div style="font-size: 2em;">↑</div> <p style="writing-mode: vertical-rl; transform: rotate(180deg);">Concurrency</p>	Isolation Level	<div style="font-size: 2em;">↓</div> <p style="writing-mode: vertical-rl; transform: rotate(180deg);">Consistency</p>
	Isolation level	
	Isolation level	
	Isolation level	
	Isolation level	

Correct Answer:

Isolation levels

read committed

serializable

read uncommitted

repeatable read

<div style="font-size: 2em;">↑</div> <p style="writing-mode: vertical-rl; transform: rotate(180deg);">Concurrency</p>	Isolation Level	<div style="font-size: 2em;">↓</div> <p style="writing-mode: vertical-rl; transform: rotate(180deg);">Consistency</p>
	read uncommitted	
	read committed	
	repeatable read	
	serializable	

QUESTION 2

You are developing an application that connects to a database.

The application runs the following jobs:

Job	Transact-SQL statement	Description
JobA	Exec uspDeletePrevRecords	The stored procedure deletes all records from a table named tblBalanceTransactions that were created before the current month by using a single DELETE statement. Approximately 10 million records are deleted each time you run this stored procedure
JobB	Exec uspUpdateCurRecords	This stored procedure updates records in the tblBalanceTransaction table that were created in the current month. Only a few hundred records are updated each time you run this stored procedure.

The READ_COMMITTED_SNAPSHOT database option is set to OFF, and auto-content is set to ON. Within the stored procedures, no explicit transactions are defined.

If JobB starts before JobA, it can finish in seconds. If JobA starts first, JobB takes a long time to complete.

You need to use Microsoft SQL Server Profiler to determine whether the blocking that you observe in JobB is caused by locks acquired by JobA.

Which trace event class in the Locks event category should you use?

- A. LockAcquired
- B. LockCancel
- C. LockDeadlock
- D. LockEscalation

Correct Answer: A

Explanation:

The Lock: Acquiredevent class indicates that acquisition of a lock on a resource, such as a data page, has been achieved.

The Lock: Acquired and Lock:Released event classes can be used to monitor when objects are being locked, the type of locks taken, and for how long the locks were retained. Locks retained for long periods of time may cause contention issues and should be investigated.

QUESTION 3

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 independent of the other questions in this series. Information and details provided in a question apply only to that question.

You are a database developer for a company. The company has a server that has multiple physical disks. The disks are not part of a RAID array. The server hosts three Microsoft SQL Server instances. There are many SQL jobs that run during off-peak hours.

You observe that many deadlocks appear to be happening during specific times of the day.

You need to monitor the SQL environment and capture the information about the processes that are causing the deadlocks.

What should you do?

- A. Create a sys.dm_os_waiting_tasks query.
- B. Create a sys.dm_exec_sessions query.
- C. Create a PerformanceMonitor Data Collector Set.
- D. Create a sys.dm_os_memory_objects query.
- E. Create a sp_configure 'max server memory' query.
- F. Create a SQL Profiler trace.
- G. Create a sys.dm_os_wait_stats query.
- H. Create an Extended Event.

Correct Answer: F

Explanation:

To view deadlock information, the Database Engine provides monitoring tools in the form of two trace flags, and the deadlock graph event in SQL Server Profiler.

Trace Flag 1204 and Trace Flag 1222

When deadlocks occur, trace flag 1204 and trace flag 1222 return information that is captured in the SQL Server error log. Trace flag 1204 reports deadlock information formatted by each node involved in the deadlock. Trace flag 1222 formats deadlock information, first by processes and then by resources. It is possible to enable both trace flags to obtain two representations of the same deadlock event.

References: [https://technet.microsoft.com/en-us/library/ms178104\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms178104(v=sql.105).aspx)

QUESTION 4

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.

Correct Answer: C

Explanation:

When deadlocks occur, trace flag 1204 and trace flag 1222 return information that is captured in the SQL Server error log. Trace flag 1204 reports deadlock information formatted by each node involved in the deadlock. Trace flag 1222 formats deadlock information, first by processes and then by resources.

The output format for Trace Flag 1222 only returns information in an XML-like format.

References: [https://technet.microsoft.com/en-us/library/ms178104\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms178104(v=sql.105).aspx)

QUESTION 5

DRAG DROP

You have a trigger named CheckTriggerCreation that runs when a user attempts to create a trigger. The CheckTriggerCreation trigger was created with the ENCRYPTION option and additional proprietary business logic.

You need to prevent users from running the ALTER and DROP statements or the sp_tableoption stored procedure.

Which three Transact-SQL segments should you use to develop the solution?

To answer, move the appropriate Transact-SQL segments from the list of Transact-SQL segments to the answer area and arrange them in the correct order.

Commands

```
DISABLE TRIGGER CheckTriggerCreation ON
DATABASE;
```

```
ENABLE TRIGGER CheckTriggerCreation ON
DATABASE;
```

```
CREATE TRIGGER CheckTriggerCreation
ON DATABASE
FOR CREATE_TRIGGER
AS
    RAISERROR ('Error message', 10, 1)
    ROLLBACK
GO
```

```
CREATE TRIGGER CheckTableChanges
ON DATABASE
FOR ALTER_TABLE, DROP_TABLE
AS
    RAISERROR ('Error message', 10, 1)
    ROLLBACK
GO
```

```
DROP TRIGGER CheckTrigerCreation ON
DATABASE;
```

```
CREATE TRIGGER CheckTableChanges
ON DATABASE
FOR DDL_TABLE_EVENTS
AS
    RAISERROR ('Error message', 10, 1)
    ROLLBACK
GO
```

Answer Area



Correct Answer:

Commands

```
DISABLE TRIGGER CheckTriggerCreation ON
DATABASE;

ENABLE TRIGGER CheckTriggerCreation ON
DATABASE;

CREATE TRIGGER CheckTriggerCreation
ON DATABASE
FOR CREATE_TRIGGER
AS
    RAISERROR ('Error message', 10, 1)
    ROLLBACK
GO

CREATE TRIGGER CheckTableChanges
ON DATABASE
FOR ALTER_TABLE, DROP_TABLE
AS
    RAISERROR ('Error message', 10, 1)
    ROLLBACK
GO

DROP TRIGGER CheckTrigerCreation ON
DATABASE;

CREATE TRIGGER CheckTableChanges
ON DATABASE
FOR DDL_TABLE_EVENTS
AS
    RAISERROR ('Error message', 10, 1)
    ROLLBACK
GO
```

Answer Area

```
DISABLE TRIGGER CheckTriggerCreation ON
DATABASE;
CREATE TRIGGER CheckTableChanges
ON DATABASE
FOR DDL_TABLE_EVENTS
AS
    RAISERROR ('Error message', 10, 1)
    ROLLBACK
GO
ENABLE TRIGGER CheckTriggerCreation ON
DATABASE;
```

QUESTION 6

HOTSPOT

You are developing an app that allows users to query historical company financial data. You are reviewing email messages from the various stakeholders for a project.

The message from the security officer is shown in the Security Officer Email exhibit below.

TO: Database developer

From: Security Officer

Subject: SQL object requirements

We need to simplify the security settings for the SQL objects. Having a assign permissions at every object in SQL is tedious and leads to a problem. Documentation is also more difficult when we have to assign permissions at multiple levels. We need to assign the required permissions at one object, even though that object may be obtaining from other objects.

The message from the sales manager is shown in the Sales Manager Email exhibit below.

TO: Database developer

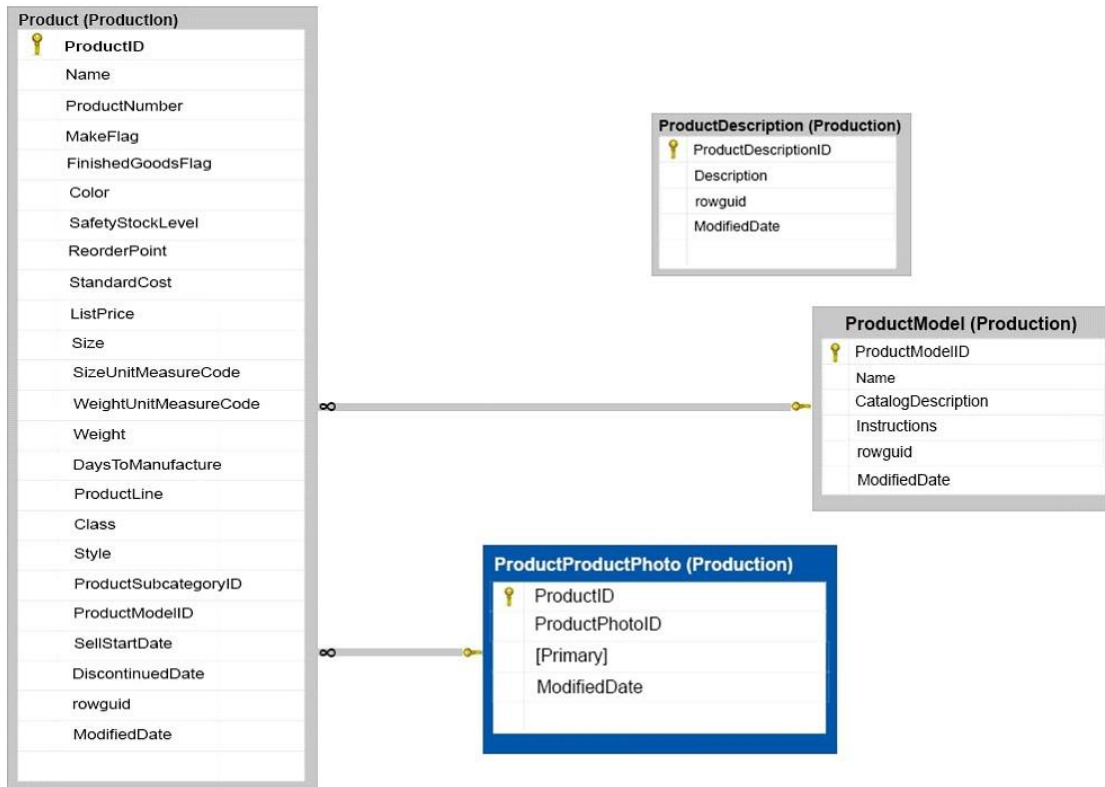
From: Sales Manager

Subject: Needed SQL objects

When creating objects for our use, they need to be flexible. We will be changing the base infrastructure frequently. We need components in SQL that will provide backward compatibility to

our front end applications as the environments change so that do not need to modify the front end applications. We need objects that can provide a filtered set of the data. The data may be coming from multiple tables and we need an object that can provide access to all of the data through a single object reference.

This is an example of the types of data we need to be able to have queries against without having to change the front end applications.



The message from the web developer is shown in the Web Developer Email exhibit below.

TO: Database developer

From: Web Developer

Subject: SQL Object component

Whatever you will be configuring to provide access to data in SQL, it needs to connect using the items referenced in this interface. We have been using this for a long time, and we cannot change this from end easily. Whatever objects are going to be used in SQL they must work using object types this interface references.

Database Name: 

Table Name: 

Column Name: 

You need to create one or more objects that meet the needs of the security officer, the sales manager and the web developer.

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

Answer Area

	Yes	No
You must create a stored procedure	<input type="radio"/>	<input type="radio"/>
You must create a trigger	<input type="radio"/>	<input type="radio"/>
You must create a view	<input type="radio"/>	<input type="radio"/>

Correct Answer:

Answer Area

	Yes	No
You must create a stored procedure	<input checked="" type="radio"/>	<input type="radio"/>
You must create a trigger	<input type="radio"/>	<input checked="" type="radio"/>
You must create a view	<input checked="" type="radio"/>	<input type="radio"/>

QUESTION 7

DRAG DROP

You are analyzing the performance of a database environment.

You suspect there are several missing indexes in the current database.

You need to return a prioritized list of the missing indexes on the current database.

How should you complete the Transact-SQL statement?

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.

Transact-SQL segments

-
-
-
-
-
-

Answer Area

```

SELECT so.name
      , (avg_total_user_cost * avg_user_impact) * (user_seeks + user_scans)
as Impact
      , mid.equality_columns
      , mid.inequality_columns
      , mid.included_columns
FROM [ ] Transact-SQL statement AS stats
INNER JOIN sys.dm_db_missing_index_groups AS mig ON stats.group_handle =
mig.index_group_handle
INNER JOIN sys.dm_db_missing_index_details AS mid ON mig.index_handle =
nid.index_handle
INNER JOIN sys.objects so WITH (nolock) ON mid.object_id = so.object_id
WHERE stats.group_handle IN (
SELECT TOP (5000) [ ] Transact-SQL statement
FROM [ ] Transact-SQL statement WITH (nolocks)
ORDER BY (avg_total_user_cost * avg_user_impact) * (user_seeks +
user_scans) DESC)
    
```

Correct Answer:

Transact-SQL segments

- sys.dm_db_missing_index_group_stats
- sys.dm_db_missing_index_details
- sys.dm_db_missing_index_stats
- sql_handle
- plan_handle
- group_handle

Answer Area

```

SELECT so.name
      , (avg_total_user_cost * avg_user_impact) * (user_seeks + user_scans)
as Impact
      , mid.equality_columns
      , mid.inequality_columns
      , mid.included_columns

FROM sys.dm_db_missing_index_group_stats AS stats
INNER JOIN sys.dm_db_missing_index_groups AS mig ON stats.group_handle =
mig.index_group_handle
INNER JOIN sys.dm_db_missing_index_details AS mid ON mig.index_handle =
mid.index_handle
INNER JOIN sys.objects so WITH (nolock) ON mid.object_id = so.object_id
WHERE stats.group_handle IN (

SELECT TOP (5) group_handle
FROM sys.dm_db_missing_index_group_stats AS H (nolocks)
ORDER BY (avg_total_user_cost * avg_user_impact) * (user_seeks +
user_scans) DESC)
    
```

QUESTION 8

DRAG DROP

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.

Isolation levels

- read committed
- serializable
- read uncommitted
- repeatable read

Answer area

Scenario	Isolation levels
Reading accurate data is top priority. Select statements will wait until any transaction that currently owns the data has been committed or rolled back before returning the value	Isolation level
Performance is top priority. The work and memory required by the Microsoft SQL Server lock manager is reduced	Isolation level
The same select statement is issued multiple times within a transaction and the same result are returned. New records are allowed to be inserted into the table referenced by the Select statement	Isolation level

Correct Answer:

Isolation levels

- read committed
- serializable
- read uncommitted
- repeatable read

Answer area

Scenario

Reading accurate data is top priority. Select statements will wait until any transaction that currently owns the data has been committed or rolled back before returning the value

Performance is top priority. The work and memory required by the Microsoft SQL Server lock manager is reduced

The same select statement is issued multiple times within a transaction and the same result are returned. New records are allowed to be inserted into the table referenced by the Select statement

Isolation levels

read committed

read uncommitted

serializable

QUESTION 9

DRAG DROP

You are monitoring a Microsoft Azure SQL Database.

The database is experiencing high CPU consumption.

You need to determine which query uses the most cumulative CPU.

How should you complete the Transact-SQL statement?

To answer, drag the appropriate Transact-SQL segments to the correct locations. Each Transact-SQL segment may be used once, more than one or not at all. You may need to drag the split bar between panes or scroll to view content.

Transact-SQL segments

- sys.dm_exec_query_stats o
- sys.dm_db_partition_stats o
- sys.dm_exec_sessions o
- sys.dm_tran_database_transactions o
- highest_cpu_queries.plan_handle DESC
- highest_cpu_queries.total_worker_time DESC
- q.objectid DESC
- q.number DESC

Answer Area

```

SELECT
    highest_cpu_queries.plan_handle,
    highest_cpu_queries.total_worker_time,
    q.dbid,
    q.objectid,
    q.number,
    q.encrypted,
    q.[text]
FROM
    (SELECT TOP 50
        o.plan_handle,
        o.total_worker_time
    FROM
        [Transact-SQL segment]

    ORDER BY o.total_worker_time desc) AS highest_cpu_queries
CROSS APPLY sys.dm_exec_sql_text(plan_handle) AS q

ORDER BY [Transact-SQL segment] ;
    
```

Correct Answer:

Transact-SQL segments

sys.dm_exec_query_stats o
 sys.dm_db_partition_stats o
 sys.dm_exec_sessions o
 sys.dm_tran_database_transactions o
 highest_cpu_queries.plan_handle DESC
 highest_cpu_queries.total_worker_time DESC
 q.objectid DESC
 q.number DESC

Answer Area

```
SELECT
    highest_cpu_queries.plan_handle,
    highest_cpu_queries.total_worker_time,
    q.dbid,
    q.objectid,
    q.number,
    q.encrypted,
    q.[text]
FROM
    (SELECT TOP 50
        o.plan_handle,
        o.total_worker_time
    FROM
        sys.dm_exec_query_stats o
    ORDER BY o.total_worker_time desc) AS highest_cpu_queries
CROSS APPLY sys.dm_exec_sql_text(plan_handle) AS q
ORDER BY highest_cpu_queries.total_worker_time DESC
```

QUESTION 10

Note: The 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 question in the series. Information and details provided in a question apply only to that question.

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 columnstore index on the table.
- E. Create a nonclustered columnstore index on the table.
- F. Create a hash index on the table.

Correct Answer: D

Explanation:

The columnstore index is 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, and up to 10x data compression over the uncompressed data size. A clustered columnstore index is the physical storage for the entire table.