

MS SQL Server Locks Interview Questions And Answers Guide.



Global Guideline.

<https://globalguideline.com/>



MS SQL Server Locks Job Interview Preparation Guide.

Question # 1

What is Optimistic concurrency?

Answer:-

Optimistic concurrency: It assumes that resource conflicts between multiple users are very unlikely to occur and thus allows transactions to execute without any locking mechanisms on the resources. It is only while changing the data that a check is made on resources if any conflicts have occurred. If there's a conflict then the application must read the data again and try to change it as well.

[Read More Answers.](#)

Question # 2

Do you know the isolation level that SQL Server support?

Answer:-

Isolation levels supported by SQL Server:

Read uncommitted: Lowest level of isolation

Read committed: Default

Repeatable read

Serializable: Highest level of isolation. All transactions are isolated from each other completely.

[Read More Answers.](#)

Question # 3

Do you know what are the differences between lost updates and uncommitted dependencies?

Answer:-

Lost updates

Last update overwrites other updates

Cannot report on data that does not exist.

Data is lost

Uncommitted dependencies

Access a row being updated by others.

May report on data that does not exist

Updates are lost

[Read More Answers.](#)

Question # 4

Explain what are the events recorded in a transaction log?

Answer:-

Events recorded in a transaction log:

Broker event category includes events produced by Service Broker.

Cursors event category includes cursor operations events.

CLR event category includes events fired by .Net CLR objects.

Database event category includes events of data.log files shrinking or growing on their own.

Errors and Warning event category includes SQL Server warnings and errors.

Full text event category include events occurred when text searches are started, interrupted, or stopped.

Locks event category includes events caused when a lock is acquired, released, or cancelled.

Object event category includes events of database objects being created, updated or deleted.

OleDb event category includes events caused by OleDb calls.

Performance event category includes events caused by DML operators.

Progress report event category includes Online index operation events.

Scans event category includes events notifying table/index scanning.

Security audit event category includes audit server activities.

Server event category includes server events.

Sessions event category includes connecting and disconnecting events of clients to SQL Server.

Stored procedures event category includes events of execution of Stored procedures.

Transactions event category includes events related to transactions.



MS SQL Server Locks Interview Questions And Answers

TSQL event category includes events generated while executing TSQL statements.
User configurable event category includes user defined events.

[Read More Answers.](#)

Question # 5

Explain optimistic and pessimistic concurrency?

Answer:-

Optimistic concurrency: - Assumes that a resource is likely to be available at all times. This means that resource locking is very unlikely. If a conflict occurs, the application must read the data and attempt the change again.

Pessimistic concurrency : - this locks the resources as and when required. A transaction can be assured to be completed unless a deadlock occurs

[Read More Answers.](#)

Question # 6

What is Pessimistic concurrency?

Answer:-

Pessimistic concurrency: Assumes that resource conflicts between multiple users are very likely to occur and hence locks resources as they are used by transactions for the duration of the transaction. A transaction is assured of successful completion unless a Deadlock occurs.

[Read More Answers.](#)

Question # 7

Explain isolation levels that SQL server supports?

Answer:-

SQL Server isolation levels:

READ COMMITTED: Shared locks are held while any data is being read.

READ UNCOMMITTED: Specifies isolation level 0 locking. There are thus no shared locks or exclusive locks. It is the least restrictive of all the isolation levels.

REPEATABLE READ: Locks are applied on all data being used by a query. However, new phantom rows can be inserted into the data set by another user and are included in later reads within the current transaction.

SERIALIZABLE: Issues a range lock on data set, preventing other users to update or insert data into dataset until the transaction is complete.

[Read More Answers.](#)

Question # 8

What are the different types of lock modes in SQL Server 2000?

Answer:-

Lock modes in SQL Server 2000:

Shared: Used for operations that read data. When shared lock is applied concurrent transactions can read data/resource but not modify it. This lock is released as soon as the transaction finishes.

Update: These locks are used when there is a need to update a row or page and later promote the update page lock to an exclusive lock before making the actual changes. Update locks are used to prevent deadlocks.

Exclusive: Used for data modification operations. No other transaction can read or modify data while Exclusive lock has been acquired.

Intent: These are used when SQL Server wants to acquire a shared or exclusive lock on some resource lower down in the hierarchy.

Schema: These are used when an operation dependent on the schema of a table is executing.

Bulk update: These are used during bulk copying of data into tables provided either TABLOCK hint is mentioned or table lock on bulk upload table option is set.

Key range: These are used by SQL Server to prevent phantom insertions/deletions into a set of records which are being accessed by a transaction.

[Read More Answers.](#)

Question # 9

Explain the various types of concurrency problem?

Answer:-

Concurrency problem:

Lost updates: this occurs when two or more transactions are trying to update same row without being aware of each other. The last update overwrites updates made by the other transactions, which results in lost data.

Uncommitted dependency (Dirty read): This occurs when a second transaction selects a row that is being updated by another transaction. This second transaction is reading data that may not have been committed.

Inconsistent Analysis (Nonrepeatable Read): This occurs when a second transaction accesses the same row several times and reads different data each time. It is similar to dirty read. However, here it reads committed data, but different data each time.

Phantom Reads: This occurs when a range of rows which is being read by a transaction is deleted or updated. The transaction's first read of the range of rows shows a row that no longer exists in the second or succeeding read, as a result of a deletion by a different transaction.

[Read More Answers.](#)

Question # 10

What is a dead lock?

Answer:-

When two computer programs sharing the same resource prevent each other from accessing the resource results in a deadlock. This deadlock results in finishing the program.

Example:

P1 requests R1 and receives it.

P2 requests R2 and receives it.

P1 requests resource R2 and is queued up, pending the release of R2.

P2 requests resource R1 and is queued up, pending the release of R1



Here, P and R is Program and Resource respectively.

[Read More Answers.](#)

Question # 11

Do you know what guidelines should be followed to help minimize deadlocks?

Answer:-

Guidelines to minimize deadlocks:-

Avoid user interaction in the transactions. The transaction must not rely on any inputs from the user.

The concurrent transactions must access data in the same order. There should be consistency in which the operations occur

Transactions must be short and simple to avoid deadlocks. Long transactions may block other necessary activities

A lower isolation level like read committed must be used. Using lower isolation levels reduces the possibility of holding shared locks for a longer time

Bound connections should be used. Here, two or more connections opened by the same application can assist each other.

[Read More Answers.](#)

Question # 12

Explain the various types of concurrency problem. I.e. Lost or buried updates, uncommitted dependency, inconsistent analysis, phantom read?

Answer:-

Types of concurrency problem:-

Lost or buried updates: - When the same row is selected for updates by two or more transactions and updates the row based on the value originally selected. Here, each transaction is unaware of the other transactions. The last update overwrites updates made by the other transactions, which results in lost data.

Uncommitted dependency: - here, a transaction reads data of another transaction which has not been committed yet. The value may be changed by the other transaction.

Inconsistent analysis: - here, the transaction reads the data inconsistently. This means that every time the data is read; different values are read. This is because another transaction is continuously updating the data.

Phantom Read: - Here, an insert or update is done on a row that belongs to some other transaction. Hence the transaction may read a row that may be deleted by some other transaction.

[Read More Answers.](#)

Question # 13

What is a live lock?

Answer:-

When a request for exclusive lock is denied again and again because a series of overlapping shared locks are interfering with each other and to adapt from each other they keep on changing the status, it is known as live lock.

[Read More Answers.](#)

Question # 14

Explain what is lock escalation and what is its purpose?

Answer:-

Lock escalation is when the system combines multiple locks into a higher level one. This is done to recover resources taken by the other finer granular locks. The system automatically does this. The threshold for this escalation is determined dynamically by the server.

Purpose:

To reduce system overhead by recovering locks

Maximize the efficiency of queries

Helps to minimize the required memory to keep track of locks.

[Read More Answers.](#)

Question # 15

Explain different types of lock modes in SQL Server 2000?

Answer:-

Different lock modes:

Shared (S): Mostly used for Read only operations like SELECT statements. It allows concurrent transactions to read data. No other transaction can modify the data until the lock is present. The lock is released as soon as the read is over.

Update locks (U): used to prevent dead locks. Used on resources that can be updated. Common forms of deadlocks occur when multiple sessions are reading, locking, and potentially updating resources later.

Exclusive (X): used for data modifications statements like INSERT, UPDATE or DELETE. This lock ensures multiple updates cant be made simultaneously.

Schema: Sch-M or Schema modification locks are used when an operation related to the table schema is being performed. Schema Stability or Sch-S locks are used when queries are being compiled.

Bulk update locks: Used when bulk copy is being performed. BU allows processes to bulk copy data concurrently into the same table while preventing other processes that are not bulk copying data from accessing the table.

[Read More Answers.](#)

Question # 16

What is SQL Server locking?

Answer:-

SQL Server has 3 main lock types:

Shared: Locks are compatible with other shared and update locks.

Update: Locks are compatible with shared locks.

Exclusive: Locks are not compatible with any other locks.

Apart from lock types, there are transaction isolation levels for managing security among transactions:



READ UNCOMMITTED
READ COMMITTED
REPEATABLE READ
SERIALIZABLE

SQL Server has some locking optimizer hints along with lock types:

NOLOCK
HOLDLOCK
UPDLOCK
TABLOCK
PAGLOCK
TABLOCKX
READCOMMITTED
READUNCOMMITTED
REPEATABLE READ
SERIALIZABLE
READPAST
ROWLOCK

[Read More Answers.](#)

Global Guideline - COM

MS SQL Server Most Popular Interview Topics.

- 1 : [Basic SQL Server Frequently Asked Interview Questions and Answers Guide.](#)
- 2 : [MS SQL Server Architecture Frequently Asked Interview Questions and Answers Guide.](#)
- 3 : [MS SQL Server Cursors Frequently Asked Interview Questions and Answers Guide.](#)
- 4 : [MS SQL Server 2008 Frequently Asked Interview Questions and Answers Guide.](#)
- 5 : [SQL Server Joins Frequently Asked Interview Questions and Answers Guide.](#)
- 6 : [MS SQL Server Functions Frequently Asked Interview Questions and Answers Guide.](#)
- 7 : [MS SQL Server Triggers Frequently Asked Interview Questions and Answers Guide.](#)
- 8 : [MS SQL Stored Procedures Frequently Asked Interview Questions and Answers Guide.](#)
- 9 : [SQL Server Indexes Frequently Asked Interview Questions and Answers Guide.](#)
- 10 : [SQL Server Data Mining Frequently Asked Interview Questions and Answers Guide.](#)

About Global Guideline.

Global Guideline is a platform to develop your own skills with thousands of job interview questions and web tutorials for fresher's and experienced candidates. These interview questions and web tutorials will help you strengthen your technical skills, prepare for the interviews and quickly revise the concepts. Global Guideline invite you to unlock your potentials with thousands of [Interview Questions with Answers](#) and much more. Learn the most common technologies at Global Guideline. We will help you to explore the resources of the World Wide Web and develop your own skills from the basics to the advanced. Here you will learn anything quite easily and you will really enjoy while learning. Global Guideline will help you to become a professional and Expert, well prepared for the future.

* This PDF was generated from <https://GlobalGuideline.com> at **November 29th, 2023**

* If any answer or question is incorrect or inappropriate or you have correct answer or you found any problem in this document then don't hesitate feel free and [e-mail us](#) we will fix it.

You can follow us on FaceBook for latest Jobs, Updates and other interviews material.
www.facebook.com/InterviewQuestionsAnswers

Follow us on Twitter for latest Jobs and interview preparation guides
<https://twitter.com/InterviewGuide>

Best Of Luck.

Global Guideline Team
<https://GlobalGuideline.com>
Info@globalguideline.com