

<div>Degrees of Isolation</div> <div>Every transaction has three characteristics: diagnostics_size, access_mode, and isolation_level.</div> <div>Diagnositics_size:</div> <div>The diagnostics_size determines the number of error condition that can be recorded for the transaction.</div> <div>Access_mode:</div> <div>There are two access_modes: READ ONLY and READ WRITE.</div> <div>If the access_mode is READ ONLY, the transaction is not allowed to modify the database. Thus INSERT, DELETE, UPDATE and CREATE statements cannot be executed. For transactions with READ ONLY access_mode, only shared locks need to be obtained, thereby increasing concurrency.</div> <div>If we have to execute one of commands INSERT, DELETE, UPDATE or CREATE, the access_mode should be set to READ WRITE.</div>	<div>Isolation_levels</div> <div>Most systems do not provide automatically serializability!!</div> <div><ul style="list-style-type: none">Implementors did not understand the issuesImplementors make a compromise between correctness and performance and provide options called <i>levels of isolation</i> (or degrees of isolation)</div> <div>The isolation_level controls the extent to which a given transaction is exposed to the actions of other transactions executing concurrently. By choosing one of four possible isolation_level settings, a user can obtain greater concurrency at the cost of increasing the transaction's exposure to other transaction's uncommitted changes.</div> <div>In SQL-92 the isolation levels are:</div> <div><ul style="list-style-type: none">READ UNCOMMITTEDREAD COMMITTEDREPEATABLE READSERIALIZABLE</div>	<div>Isolation_levels</div> <div><ul style="list-style-type: none">SERIALIZABLE – this isolation level ensures that T reads only the changes made by committed transactions, that no value read or written by T is changed by any other transaction until t is complete.</div> <div>In terms of lock-based implementation, a SERIALIZABLE means that the lock algorithm is two-phase and well formed.</div> <div><ul style="list-style-type: none">REPEATABLE READ - this isolation level ensures that T reads only the changes made by committed transactions, that no value read or written by T is changed by any other transaction until t is complete. However, T could experience the phantom phenomenon.</div> <div>In terms of lock-based implementation, a REPEATABLE READ means that the lock algorithm is two-phase and well formed. A REPEATABLE READ uses the same locking protocol as a SERIALIZABLE transaction except that it does not do index locking – it locks only individual objects - not sets of objects.</div>																				
<div>Isolation levels</div> <div><ul style="list-style-type: none">READ COMMITTED (<i>cursor stability</i>) - this isolation level ensures that T reads only the changes made by committed transactions, that no value written by T is changed by any other transaction until t is complete. However, a value read by T may well be modified by another transaction while T is in progress, and T is exposed to the phantom phenomenon.</div> <div>In terms of lock-based implementation, a READ COMMITTED means that the lock algorithm is two-phase with respect to write locks and well formed with respect to reads. In other words, all shared locks obtained by T are released immediately.</div> <div><ul style="list-style-type: none">READ UNCOMMITTED (browse)- T can read changes made to an object by an ongoing transaction. Moreover, the object can be changed further while T is in progress, and T is exposed to the phantom phenomenon.</div> <div>In terms of lock-based implementation, a READ UNCOMMITTED means a transaction T obtains write locks before writing data items, and holds these locks until the end, but does not obtain shared locks before reading data items.</div> <div>READ UNCOMMITTED is allowed only for read-only transactions – a transaction is required to have an access mode of READ ONLY.</div>	<div>Isolation levels</div> <table><tr><th>Isolation Level</th><th>Dirty Read</th><th>Unrepeatable Read</th><th>Phantom</th></tr><tr><td>READ UNCOMMITTED</td><td>maybe</td><td>maybe</td><td>maybe</td></tr><tr><td>READ COMMITTED</td><td>no</td><td>maybe</td><td>maybe</td></tr><tr><td>REPEATABLE READ</td><td>no</td><td>no</td><td>maybe</td></tr><tr><td>SERIALIZABLE</td><td>no</td><td>no</td><td>no</td></tr></table> <div>Why READ COMMITTED is called sometimes <i>Cursor Stability</i>?</div> <div>exec sql select balance into :balance from account where account_id=:id;</div> <div>balance=balance+10;</div> <div>exec sql update account set balance=:balance where account_id=:id;</div>	Isolation Level	Dirty Read	Unrepeatable Read	Phantom	READ UNCOMMITTED	maybe	maybe	maybe	READ COMMITTED	no	maybe	maybe	REPEATABLE READ	no	no	maybe	SERIALIZABLE	no	no	no	<div>exec sql declare cursor c for select balance from account where account_id=:id;</div> <div>exec sql open cursor c</div> <div>exec sql fetch c into :balance</div> <div>balance=balance+10;</div> <div>exec sql update account set balance=:balance where current of cusor c;</div> <div>exec sql close c;</div> <div><ul style="list-style-type: none">Most SQL-systems keep a shared lock on the record currently addressed by a cursor.</div> <div>The isolation and access-mode can be set using the SET TRANSACTION command. The following command declares the current to be SERIALIZABLE and READ ONLY:</div> <div>SET TRANSACTION ISOLATION LEVEL SERIALIZABLE READ ONLY</div>
Isolation Level	Dirty Read	Unrepeatable Read	Phantom																			
READ UNCOMMITTED	maybe	maybe	maybe																			
READ COMMITTED	no	maybe	maybe																			
REPEATABLE READ	no	no	maybe																			
SERIALIZABLE	no	no	no																			