### Isolation Levels From Zero to Hero!

It's the I in the sky

Friday, March 6, 2020

Pythian

toria adv. TATOR diction /'drk $\int(\partial n) n/n$ . ciation in speaking o dictio from dico dict-s dictionary /'dık [ənə book listing (usu. al explaining the words giving corresponding language. 2 reference terms of

#### isolation

The fact that something is separate and not connected to other things.

## Introduction

Two quick reviews of basic concepts

A quick review of items not related to this subject, but that must be present in any talk (1 of 2)



A quick review of items not related to this subject, but that must be present in any talk (2 of 2)



#### A quick review of items related to the subject of this talk

#### Database

- Computer software that allows users to interact with a set of related data and the way it is organized.
- Control access to data.
- Manages the definition of data organization.
- Update of data.
- Retrieval of data.
- Administration.

#### A quick review of items related to the subject of this talk

- Concurrency
- Consistency
- Transaction
- Isolation
- Durability

A quick review of items not related to this subject, but that must be present in any talk (2 of 2)



## ACID

Everybody needs a 303

#### ACID

- Atomicity (Transactions)
- Consistency (Follow rules and expected behaviors)
- Isolation (Concurrency and Consistency)
- Durability (Consistency)

#### Isolation

- Multiuser data manipulation.
- Rollback may lead to inconsistencies on other transactions.
- Data not committed can lead to inconsistencies on our transaction.

#### Phenomenon (Not a Travolta movie)

- Three types of inconsistencies (as defined by the standard):
  - P1 Dirty reads

I read data not committed, later being rolled back.

• P2 Non repeatable read.

I read data, later changed by another transaction. I read again.

• P3 Phantom reads

I read data, later more data added. I read again.

#### **Isolation Levels**

- Read uncommitted
- Read committed
- Repeatable read
- Serializable

Level	P1 – Dirty Read	P2 – Unrepeatable Read	P3 – Phantom Read
Read uncommitted	Possible	Possible	Possible
Read committed	Not possible	Possible	Possible
Repeatable read	Not possible	Not possible	Possible
Serializable	Not possible	Not possible	Not possible

#### What's wrong with Isolation Levels?

- Three types of inconsistencies (as defined by the standard):
  - P1 Dirty reads

I read data not committed, later being rolled back.

• P2 Non repeatable read.

I read data, later changed by another transaction. I read again.

• P3 Phantom

I read data, later more data added. I read again.

Is there anything missing here?

#### What's wrong with Isolation Levels?

- Where are the writes?
  - P1 Dirty **READS**

I read data not committed, later being rolled back.

• P2 Non repeatable **READ**.

I read data, later changed by another transaction. I read again.

• P3 Phantom READ

I read data, later more data added. I read again.

#### What's wrong with Isolation Levels?

- Transactions are not only reads.
- Already modified data can be overwritten.
- One direction isolation.
- No grants to control isolation levels.
- Missing Phenomena 0 and 4
  - P0 Dirty writes
  - P4 Lost update
- Missing isolation levels
  - Cursor stability
  - Snapshot isolation
- What about DDL?

## Concurrency control

Ordem e progresso

#### 2PL

- Oldest concurrency control mechanism.
- Two lock types:
  - Read lock
  - Write lock
- Two-phase locking.
  - Expanding phase: Lock acquisition.
  - Shrinking phase: Lock release.
- Pros: strict serializability.
- Cons: locks are expensive.

#### 2PL (MySQL8)

- Explicit locking using:
  - FOR UPDATE
  - FOR SHARE (LOCK IN SHARE MODE)
- Filter
  - OF <table\_name>
- No wait for lock
  - NOWAIT
  - SKIP LOCKED

#### **MVCC**

- Keep multiple versions of each modified row to satisfy isolation levels from current transactions.
- Where are "old" versions kept?
  - MySQL: undo log.
  - Postgres: same table.
- Pros
  - Less contention.
- Cons:
  - Fixes only read issues.
  - Background operations to clean old versions.

#### MVCC (MySQL8)

- Additional fields added (hidden)
  - DB\_TRX\_ID
  - DB\_ROLL\_PTR
- Undo segments
  - Store old data
  - Store inserted data
- Undo segments are required for:
  - Consistency
  - Atomicity
  - Isolation

## And Now For Something Completely Different

A Monty Python tribute

#### UNDO & REDO

- Undo logging
  - Transaction consistency
- Redo logging
  - Durability
  - Performance
- Double-write buffer
  - Durability
- 1 Update requires multiple "writes"
  - 1 w table (+1w redo +1w double-buffer)
  - 1 w undo (+1w redo +1w double-buffer)
- Optimized
  - Multiple writes per sync, smaller redo writes, sequential...

#### CURSOR: 2 "transactions" for the price of 1

- Usually SQL sentences are executed sequentially.
- Cursors are the exception:
  - Initial snapshot
  - Internal loop data modifications can be committed.
- Intra-transaction isolation

#### MySQL Repeatable Read (InnoDB)

- Transactions start at first data retrieval.
- Repeatable read actually implements snapshot isolation level
  - Consistent snapshot of all tables
  - Good for mysqldump.
  - Should be read-only to avoid completely lost updates
- Autocommit active by default.
- Better option
  - Read-committed
  - Explicit locking when required.



But I said, "No, no, no"

#### Isolation levels and concurrency

- The standard does not cover all the cases.
- MySQL Repeatable read goes beyond repeatable read.
- Beware of cursors (not found exception)
- Use explicit locking when needed.

#### One thing your DBA ignore

- It is not true that in MySQL it is not possible to lock reads.
- Serializable isolation level locks read!
  - T1 updates table A.
  - T2 updates table B.
  - T1 queries table B (locks)
  - T2 queries table A (deadlock detected)



But if you could just see the beauty. These things I could never describe. These pleasures a wayward distraction, this is my one lucky prize. Isolation, isolation, isolation. Isolation.

Isolation - Joy Division

## Thank you!

Questions?

#### Pep Pla

Born in Vinaròs, a small village near the Mediterranean and currently living in Barcelona.

Most of the time I'm busy with my three kids, my wife and our two cats.

And in my spare time I'm a DBC at Pythian, surrounded by some of the most brilliant DBAs in the world.



#### **ABOUT PYTHIAN**

Pythian's 400+ IT professionals help companies adopt and manage disruptive technologies to better compete





# Pythian

-56.45