

Large Language Models 2024 | April 11 2024 | Online

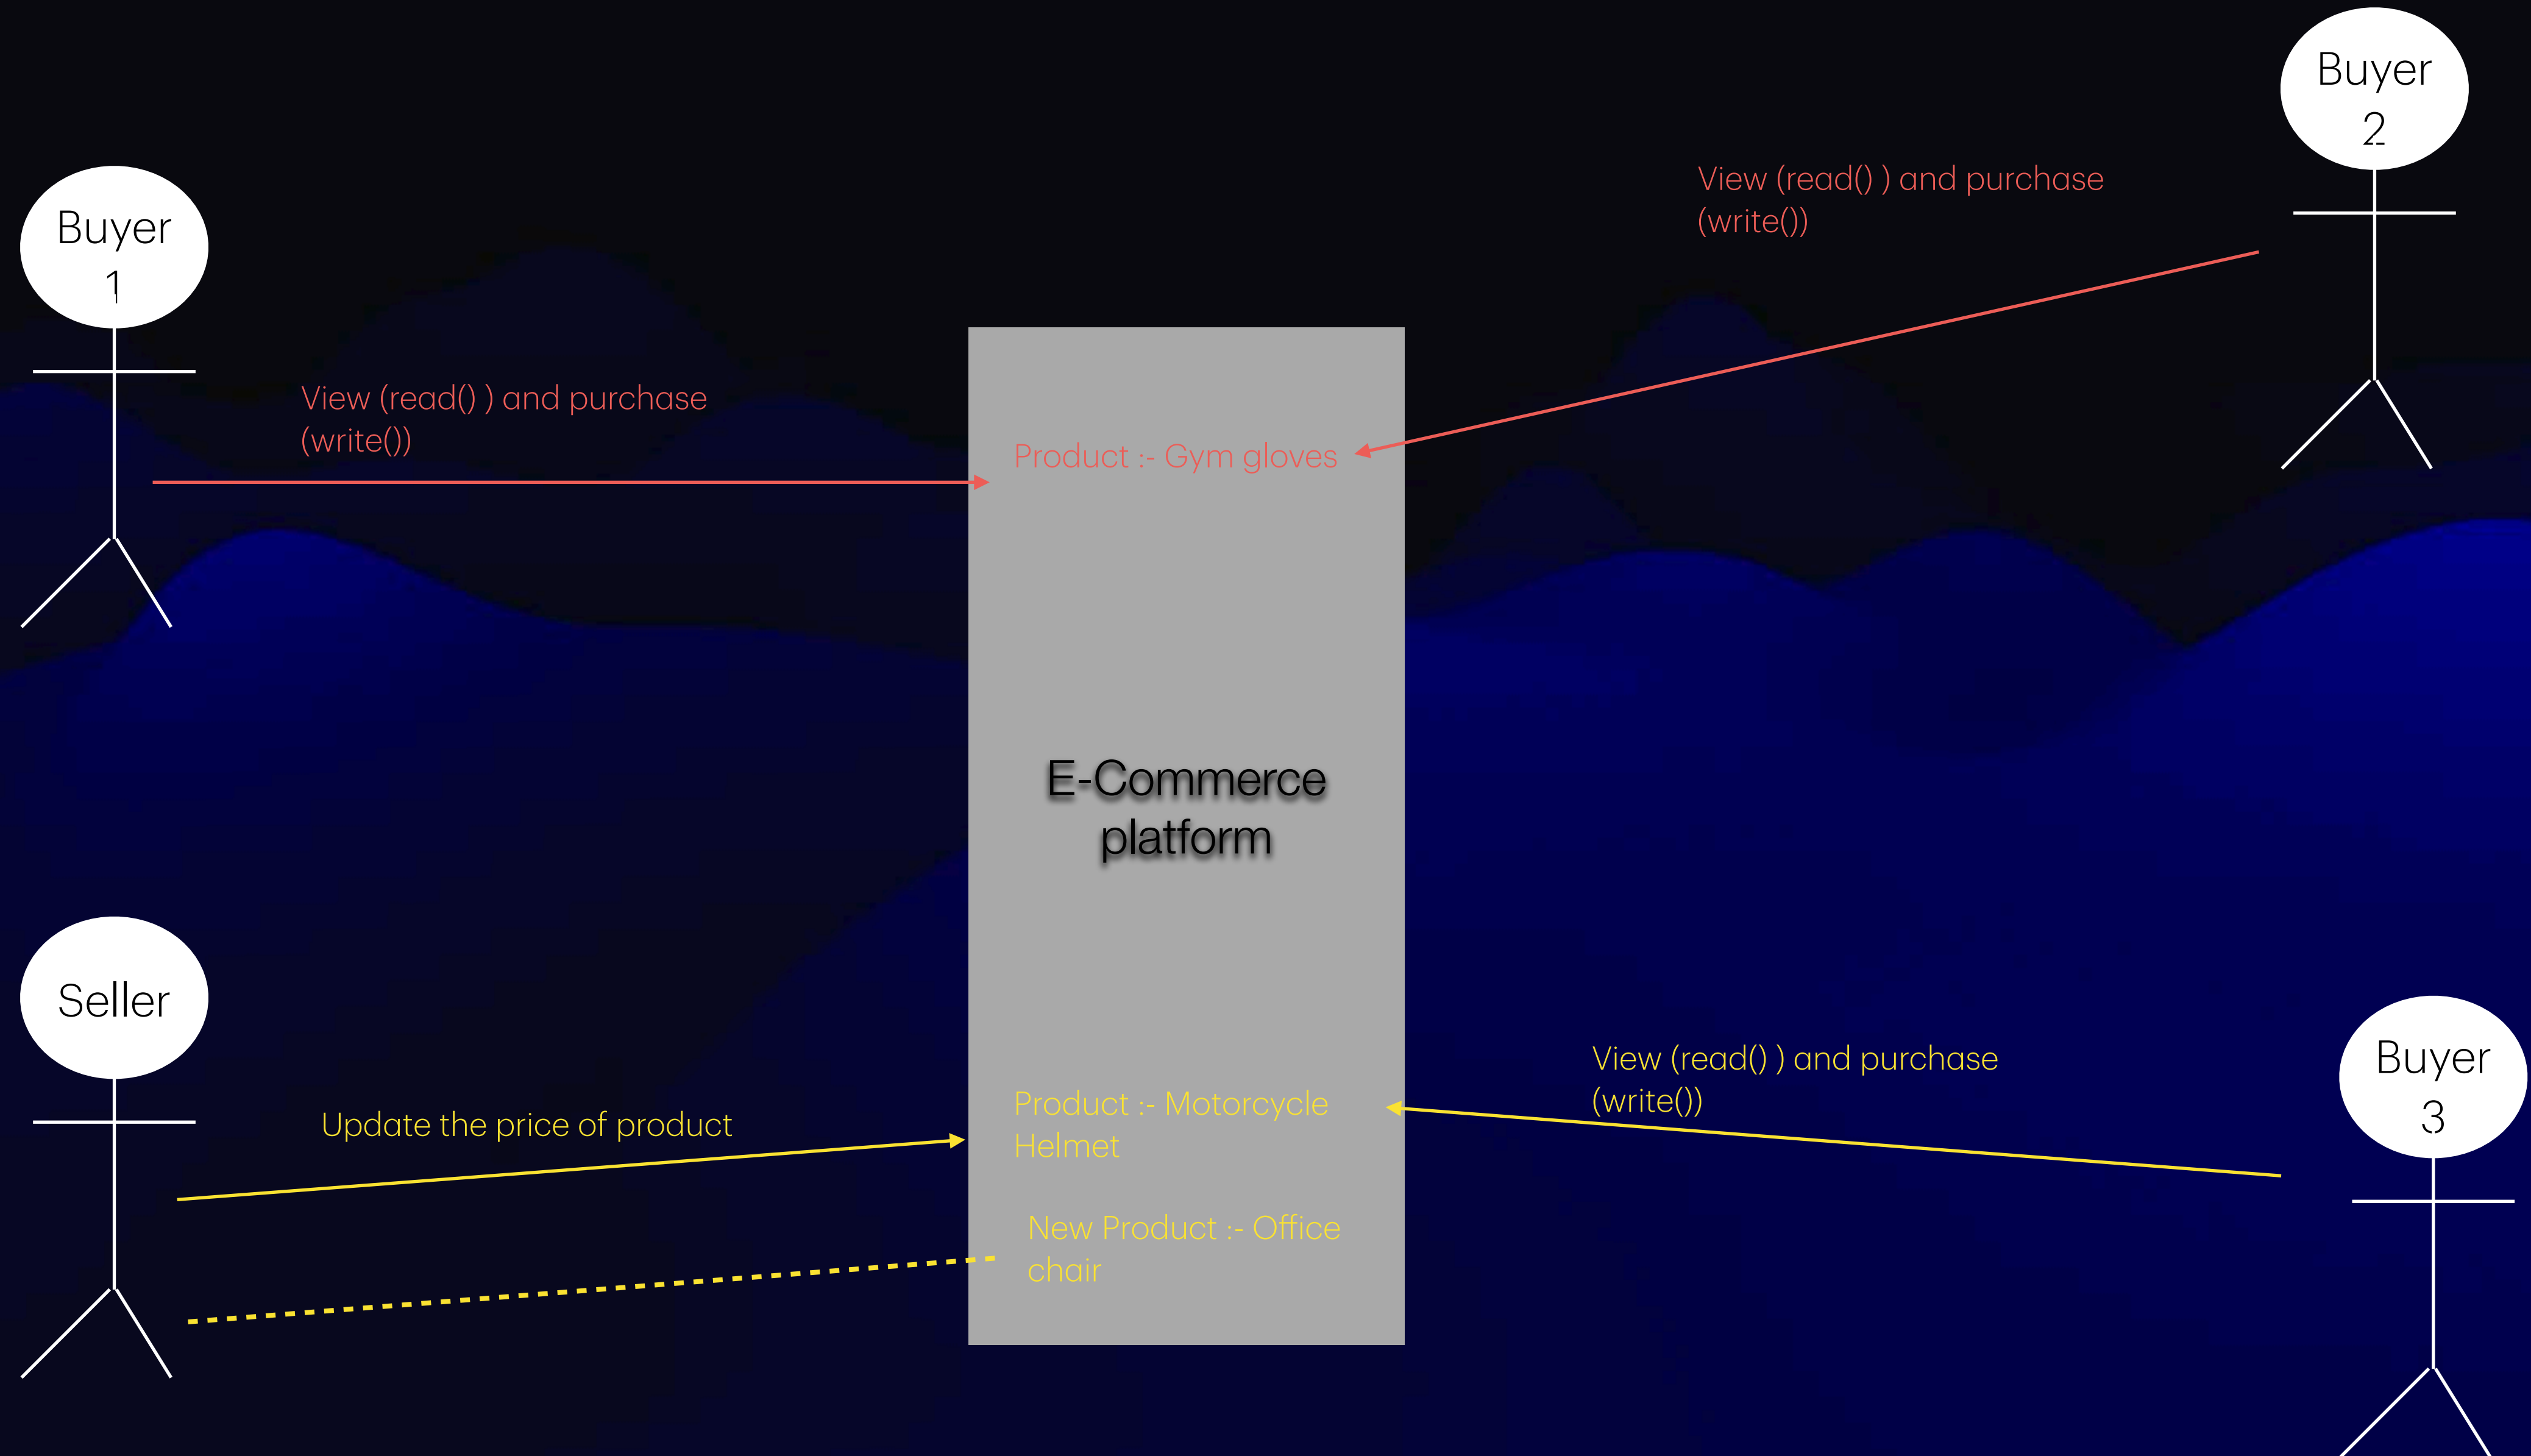
Isolation levels and partial failures in distributed systems with LLMs

Santosh Nikhil Kumar Adireddy
Software Engineering Lead, ByteDance

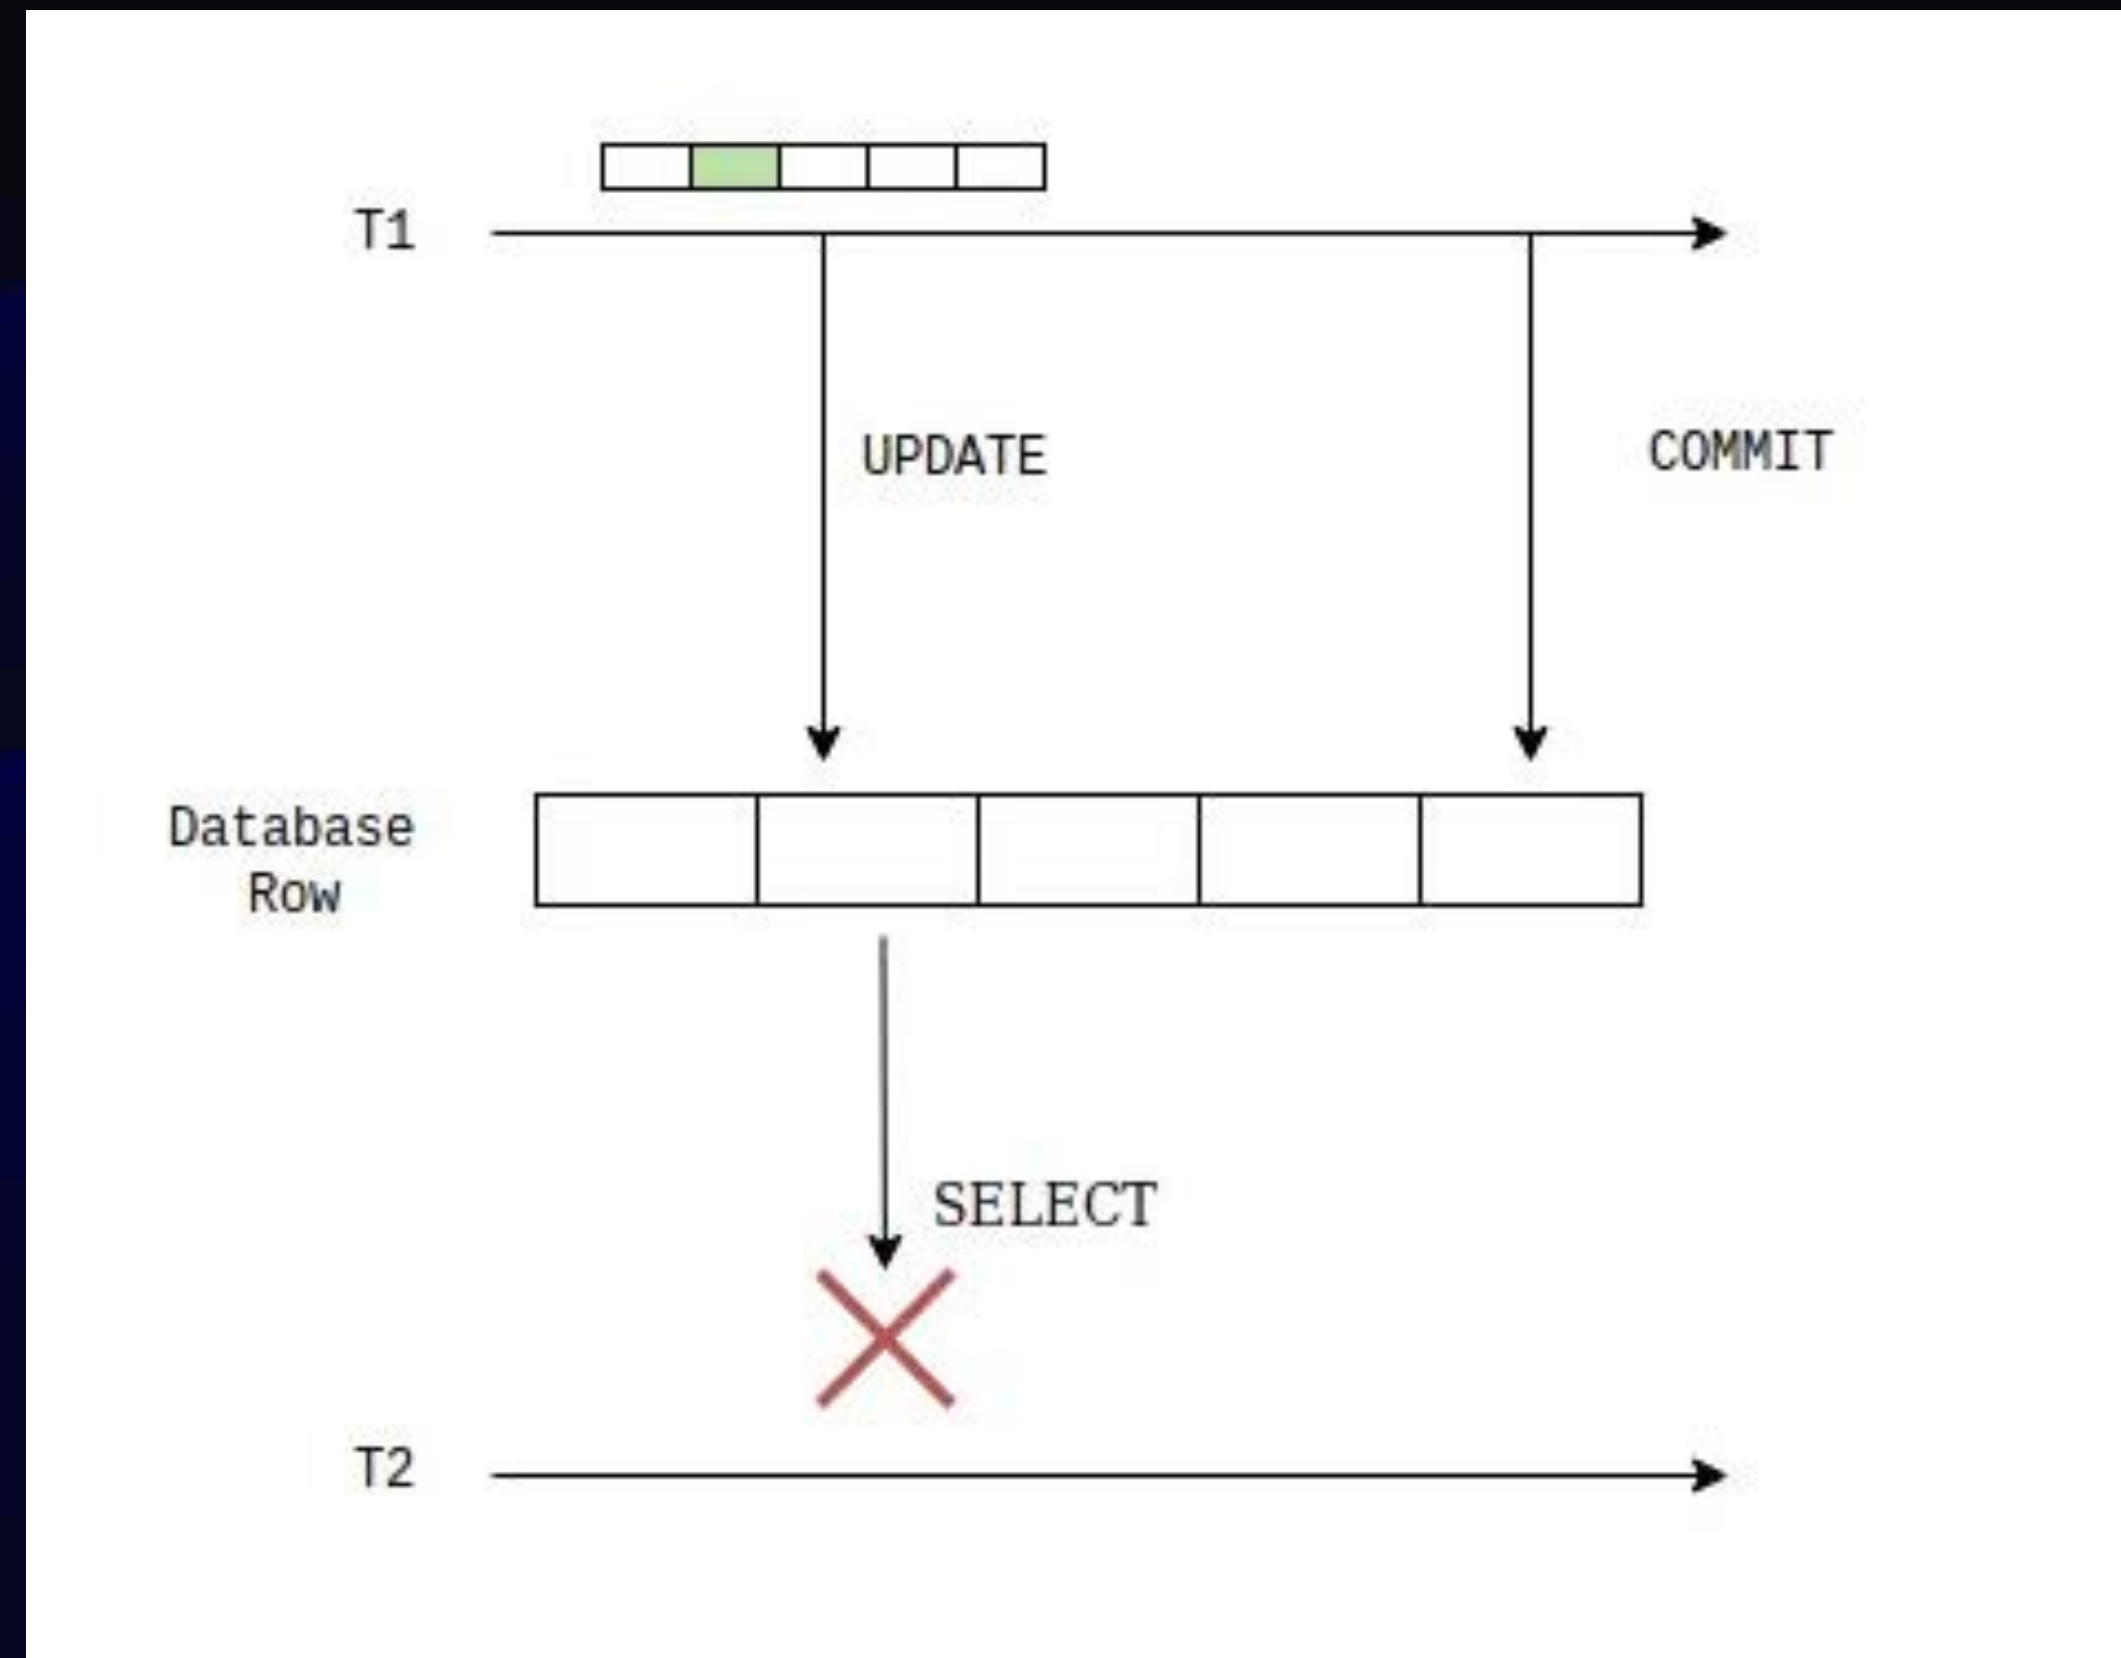
Isolation levels in Transactions

Tradeoff between consistency, concurrency and performance

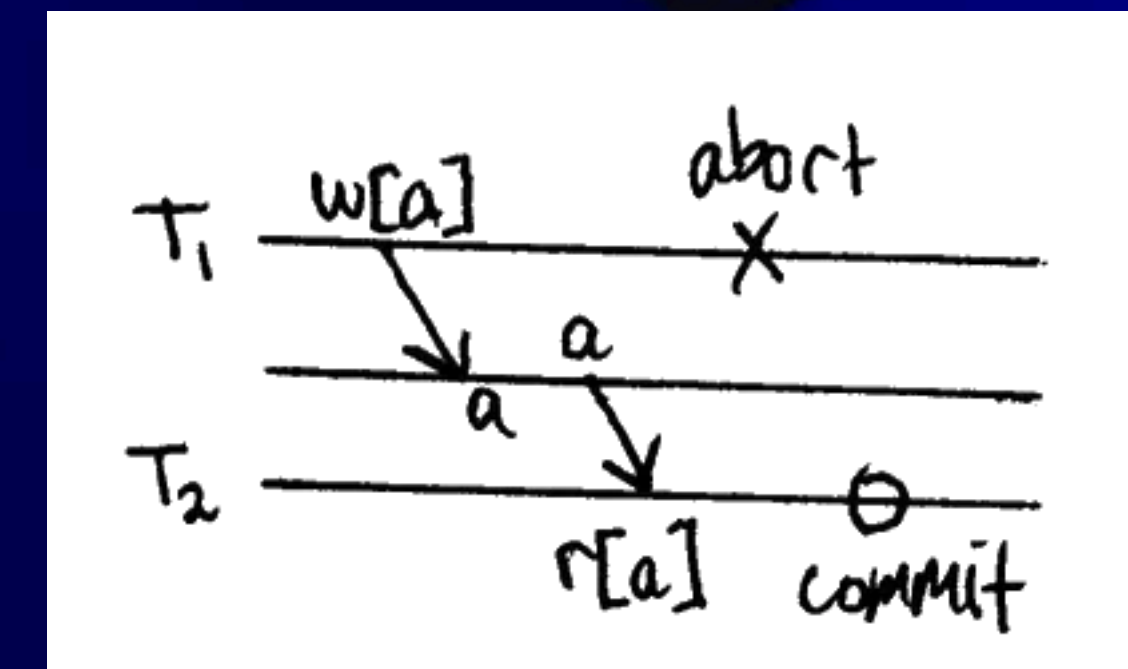
What is Isolation and Isolation level



Isolation level : Read Committed

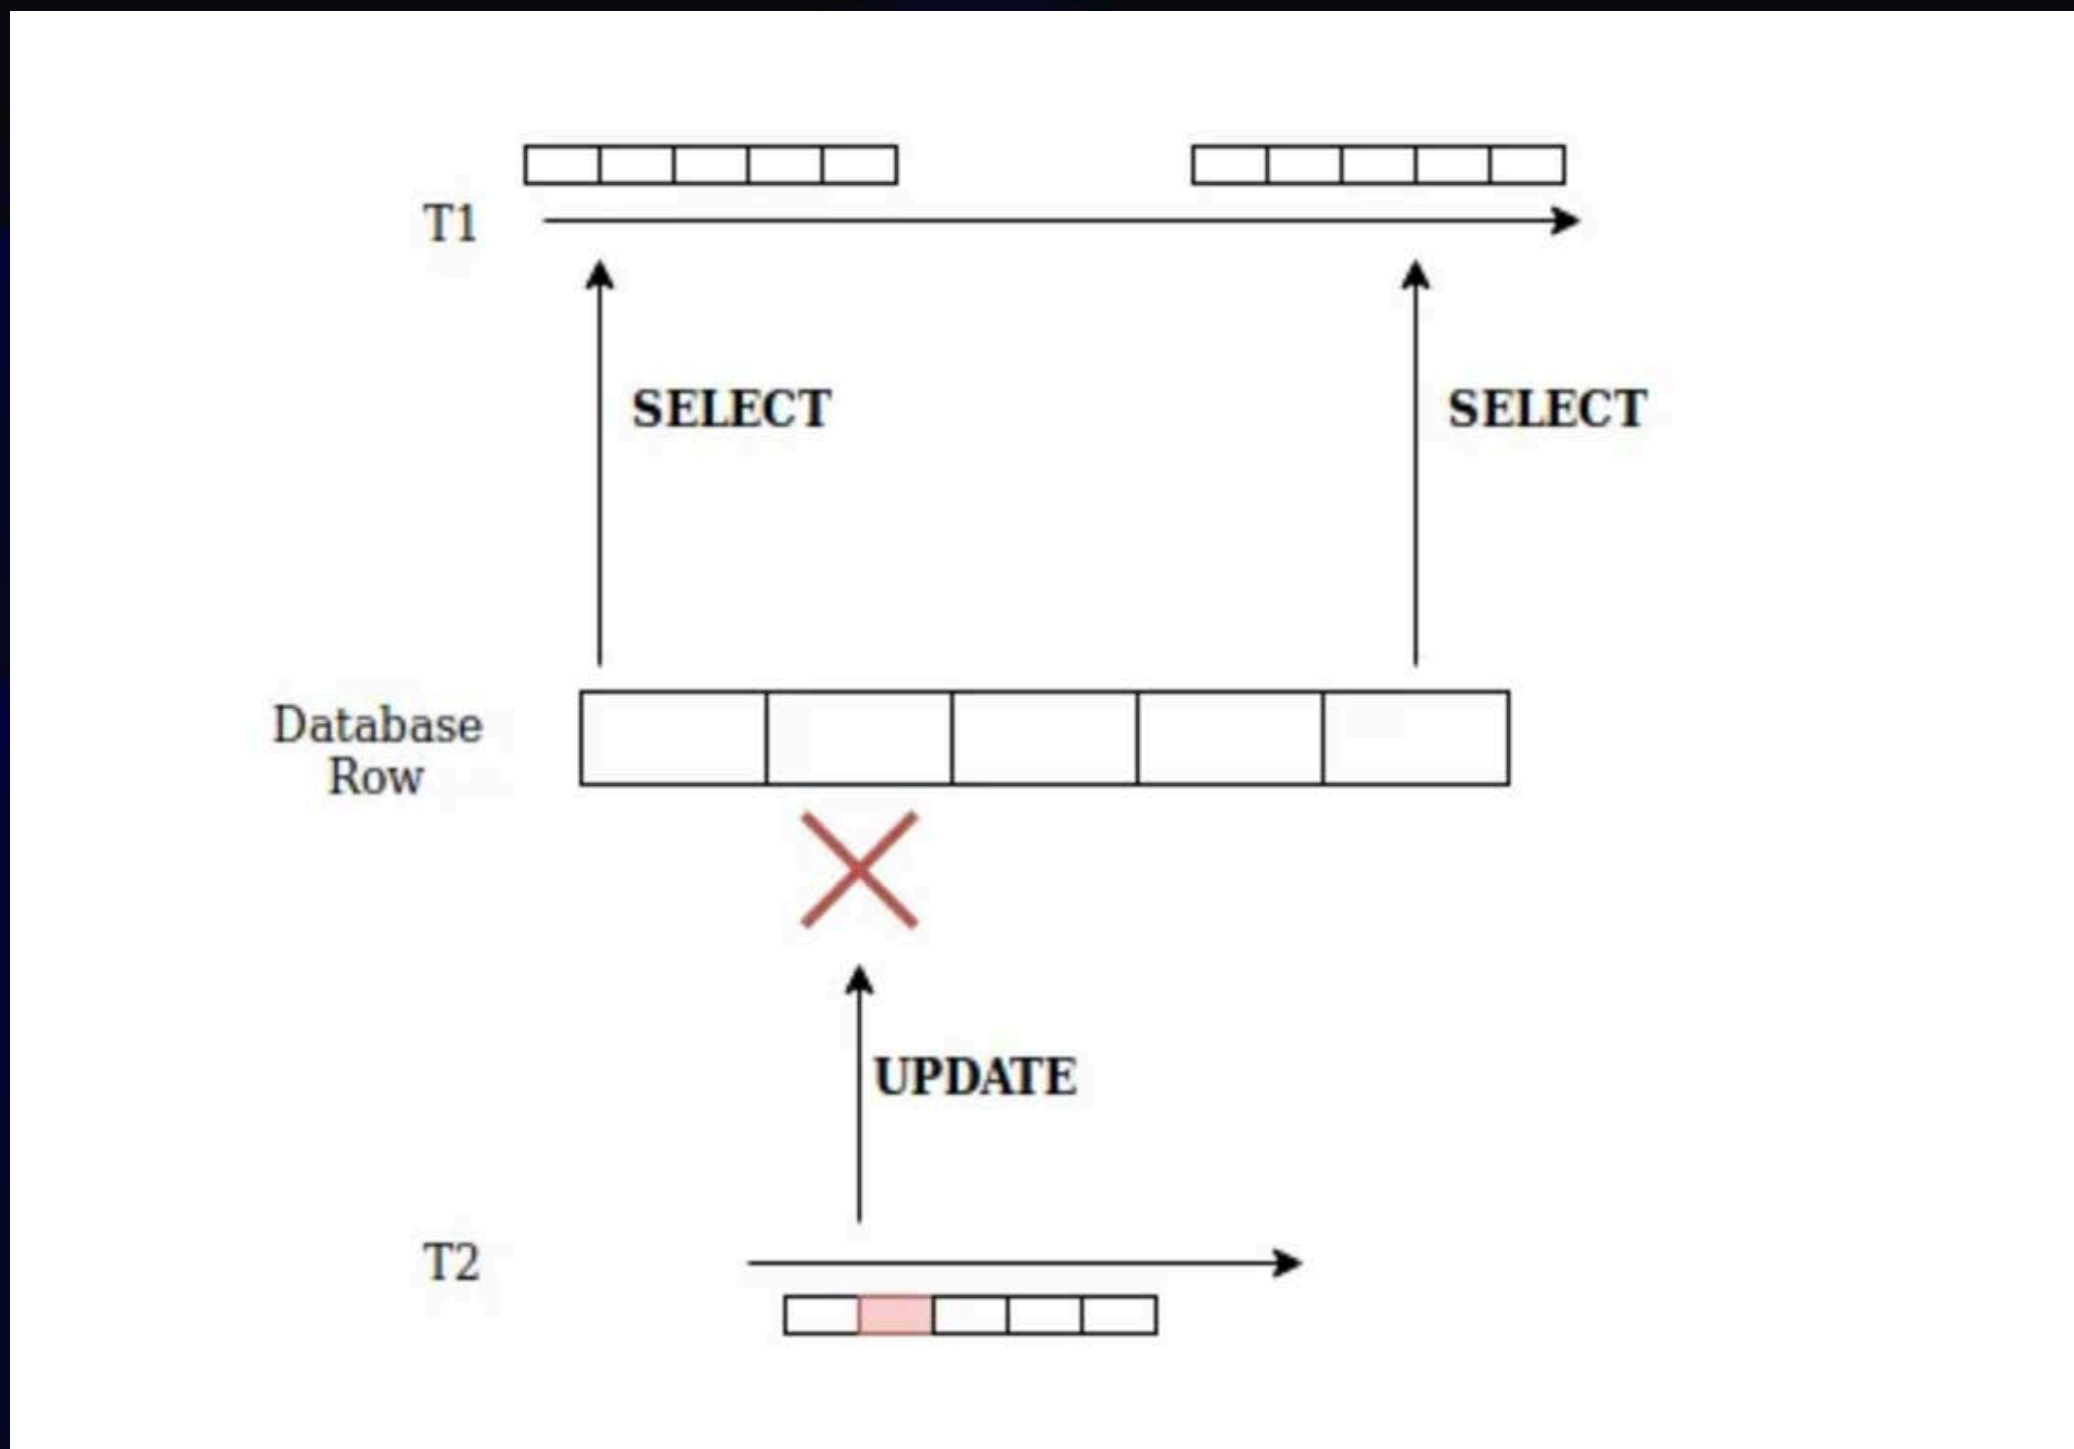


Dirty read

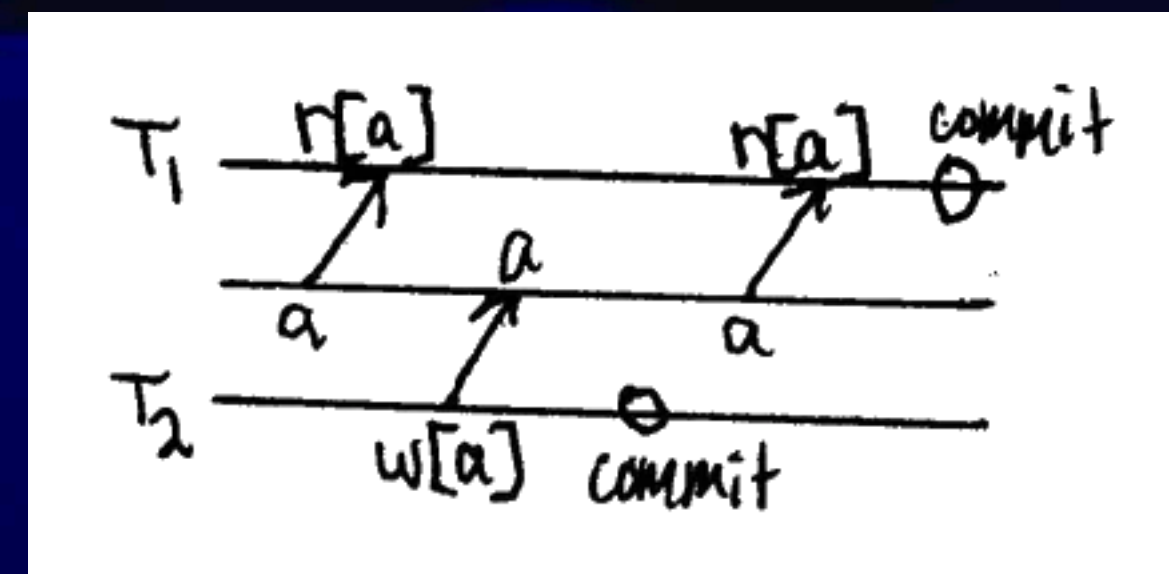


- Offers good concurrency with some trade-offs in consistency, good performance

Isolation level : Repeatable Read



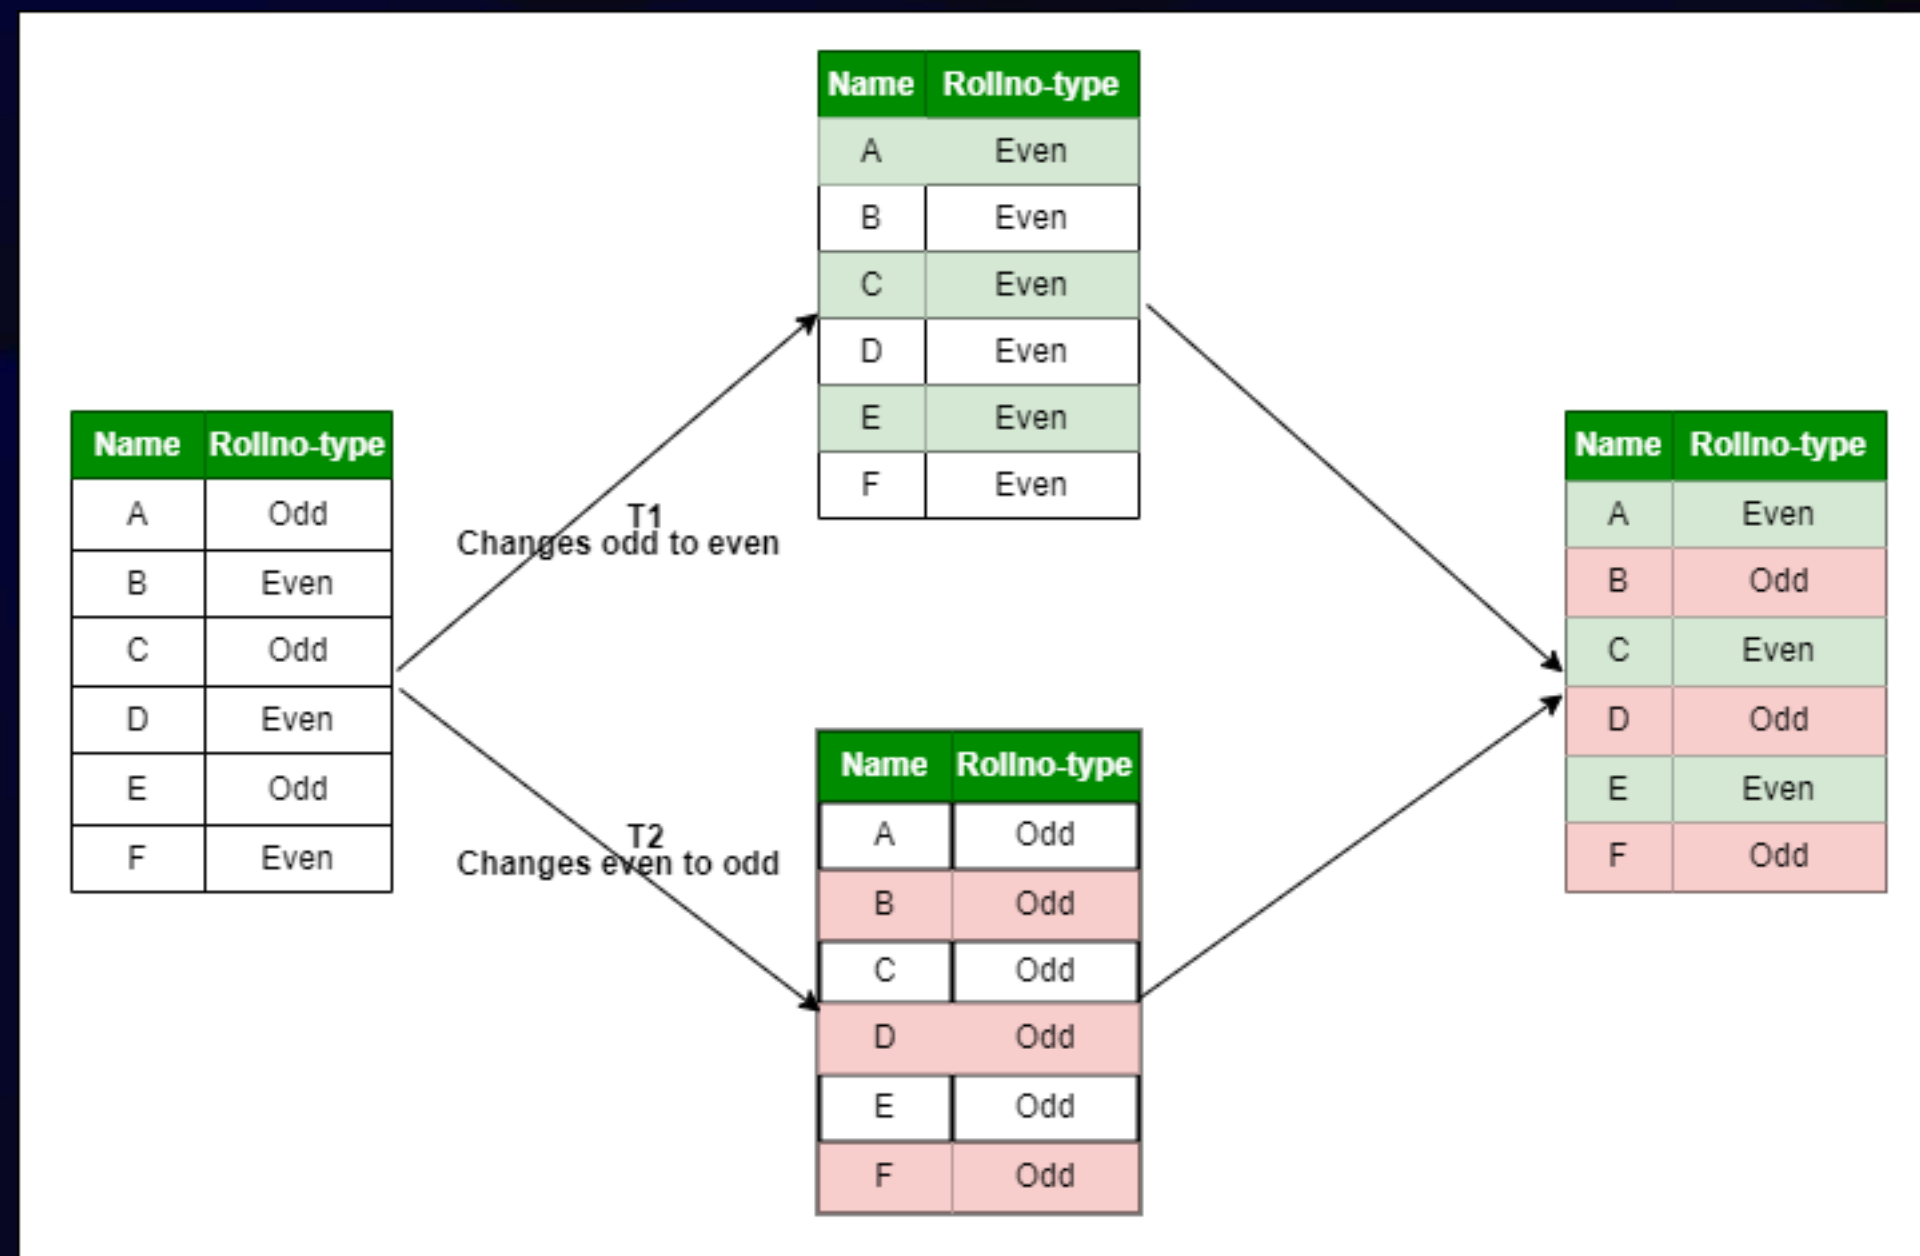
Non-repeatable read



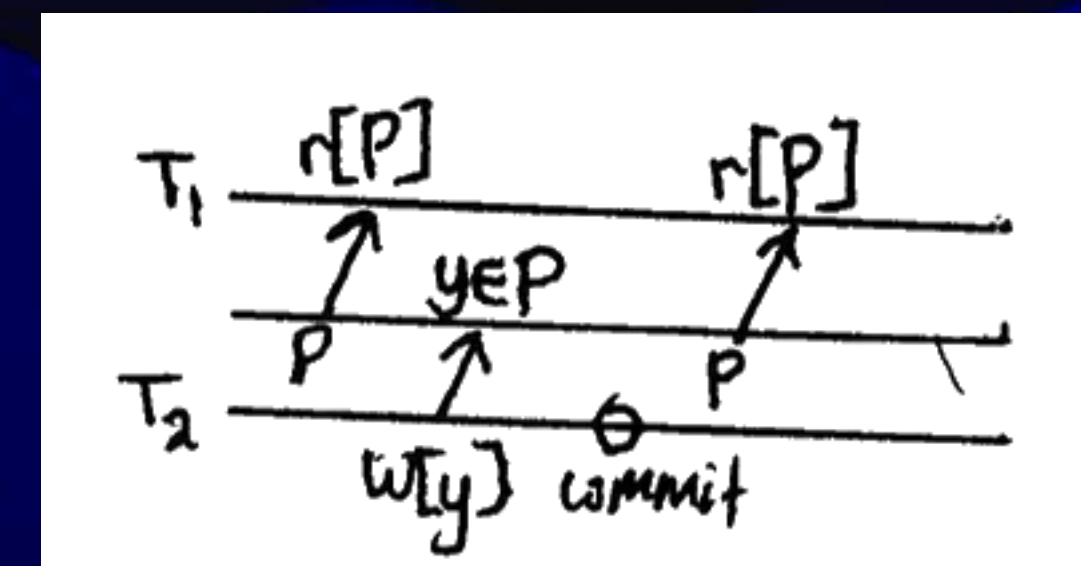
- Offers stronger consistency at the expense of reduced concurrency, bad performance

Isolation level : Snapshot Isolation

Readers never block writers and writers never block readers



Phantom read




- Offers strongest consistency with high concurrency, good performance in read-heavy apps

Partial Failures

Call for System Availability and Reliability

What are Partial Failures



1. Node failures:- Individual nodes may experience failures due to hardware or software issues

2. Network Partitions:- Communication between certain nodes may be disrupted, leading to network partitions or isolated segments.

Replication to handle partial failures

- Full Replication
- Partial Replication
- Sharding
- Replication Chains
- Primary-Backup Replication
- Quorum-based Replication