

Lock Modes

Mode	Used
Access Share Lock	SELECT
Row Share Lock	SELECT FOR UPDATE
Row Exclusive Lock	INSERT, UPDATE, DELETE
Share Lock	CREATE INDEX
Share Row Exclusive Lock	EXCLUSIVE MODE but allows ROW SHARE LOCK
Exclusive Lock	Blocks ROW SHARE LOCK and SELECT...FOR UPDA
Access Exclusive Lock	ALTER TABLE, DROP TABLE, VACUUM, and unquali

Transaction Status

pg_clog

XID Status flags

028	0	0	0	1	0	0	1	0
024	1	0	1	0	0	0	0	0
020	1	0	1	0	0	1	0	0
016	0	0	0	0	0	0	1	0
012	0	0	0	1	0	1	1	0
008	1	0	1	0	0	0	1	0
004	1	0	1	0	0	0	0	0
000	1	0	0	1	0	0	1	0

Transaction Id (XID)

00 In Progress
01 Aborted
10 Committed



xmin

xmax

Multi-Version Concurrency Control

- ▶ Each query sees only transactions completed before it started
- ▶ On query start, PostgreSQL records:
 - ▶ the transaction counter
 - ▶ all transaction id's that are in-process
- ▶ In a multi-statement transaction, a transaction's own previous queries are also visible
- ▶ The above assumes the default *read committed isolation level*

MVCC Tuple Requirements

- ▶ Visible tuples must have a creation transaction id that:
 - ▶ is a committed transaction
 - ▶ is less than the transaction counter stored at query start *and*
 - ▶ was not in-process at query start
- ▶ Visible tuples must *also* have an expire transaction id that:
 - ▶ is blank *or* aborted *or*
 - ▶ is greater than the transaction counter stored at query start *or*
 - ▶ was in-process at query start

MVCC Behavior

Cre	40
Exp	

INSERT

Cre	40
Exp	47

DELETE

Cre	64
Exp	78

old (delete)

UPDATE

Cre	78
Exp	

new (insert)

UPDATE is effectively a DELETE and an INSERT.

MVCC Examples

Create-Only

Cre 30
Exp
Cre 50
Exp
Cre 110
Exp

Visible

Sequential Scan

Invisible

Snapshot

Invisible

The highest-numbered committed transaction: 100

Open Transactions: 25, 50, 75

Create & Expire

Cre 30
Exp 80
Cre 30
Exp 75
Cre 30
Exp 110

Invisible

Visible

Visible

For simplicity, assume all other transactions are committed.

Internally, the creation xid is stored in the system column 'xmin', and expire in 'xmax'.

System Tables

