Replication Futures

Andres Freund
PostgreSQL Developer & Committer
Citus Data

citusdata

A few words

- UDR/BDR are mostly developed by 2ndQuadrant
- Most of my work on the topic was done while at 2ndQ
- Slides:

Demo Time

How do things fit together

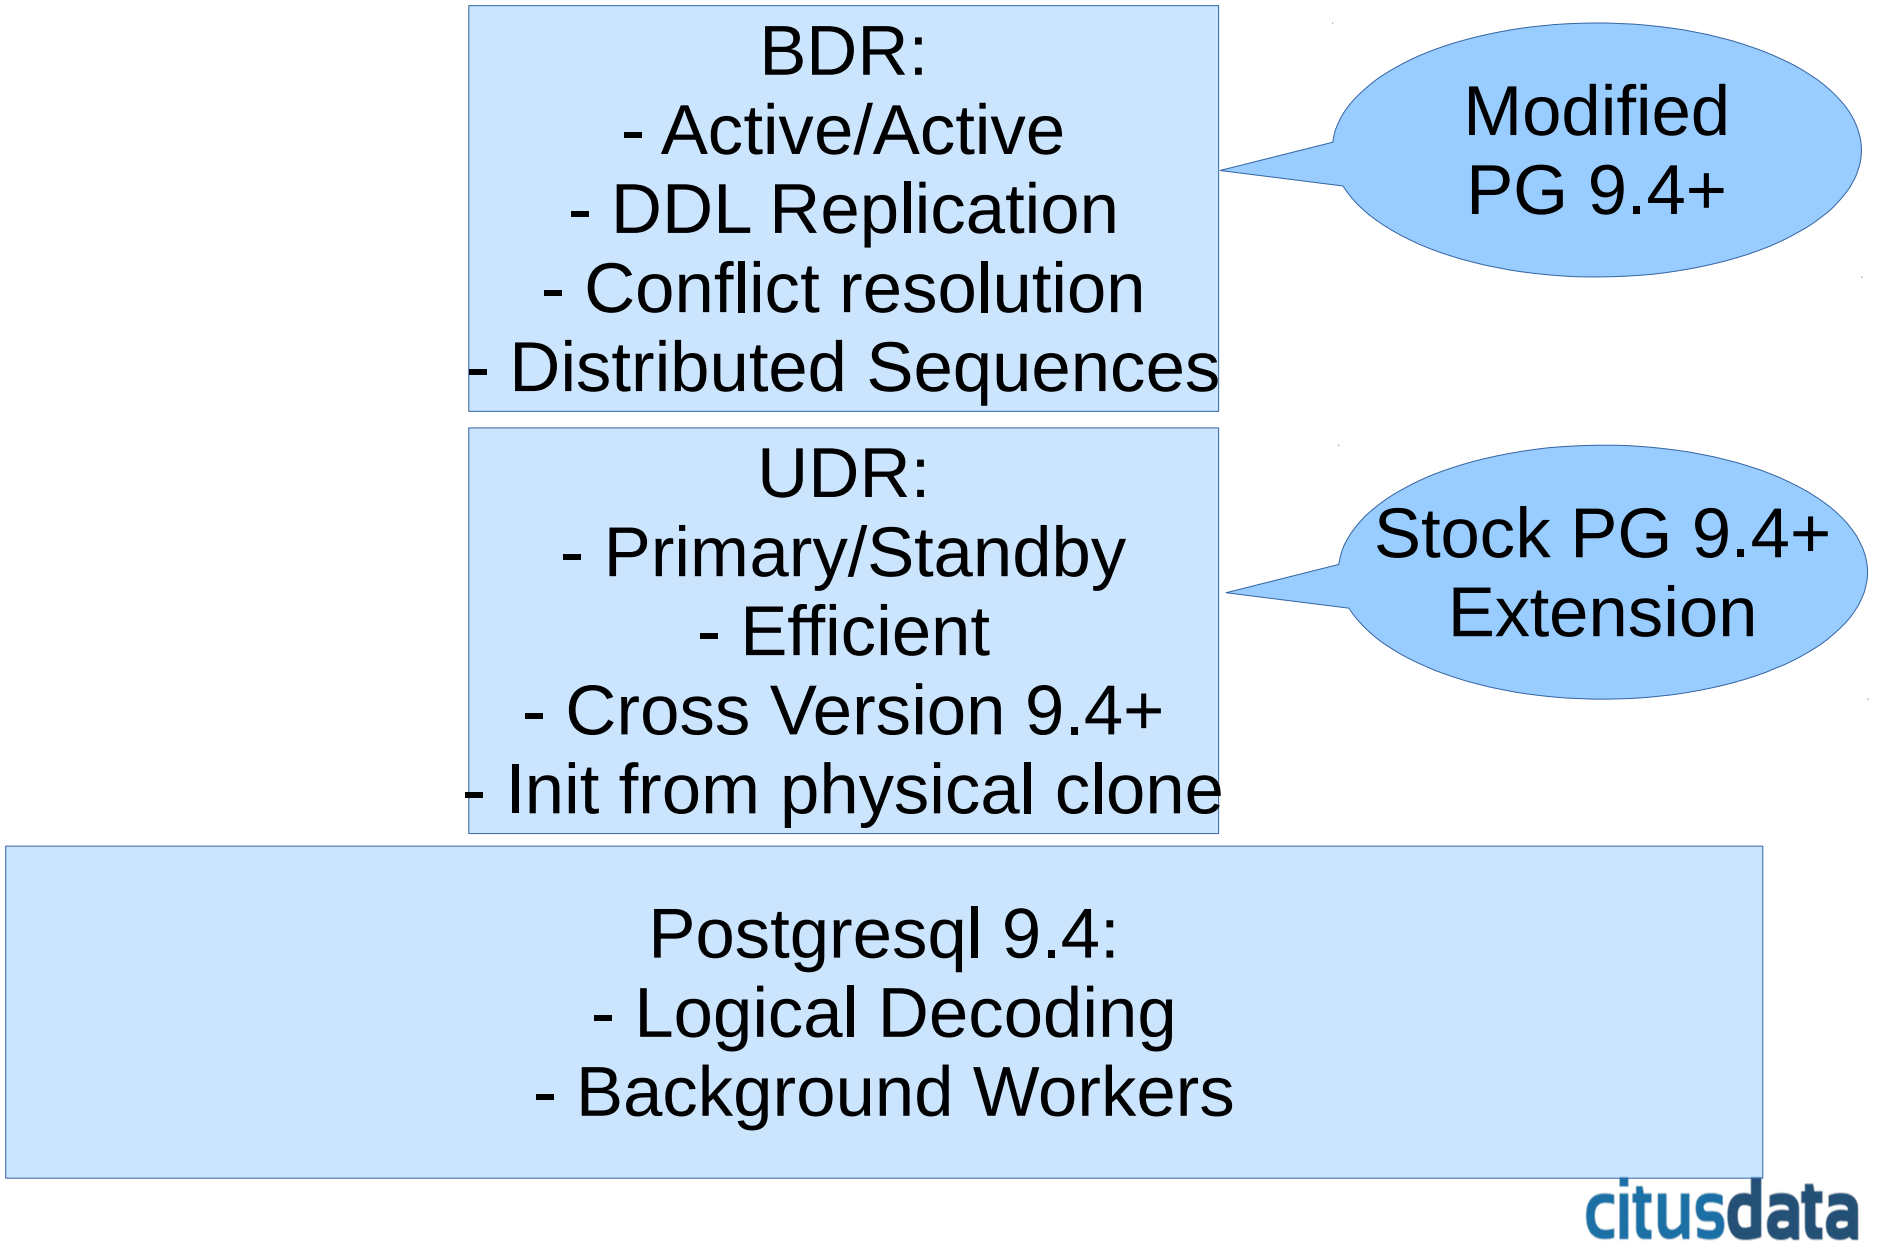

Recap: 9.4

- Background workers
- Logical decoding

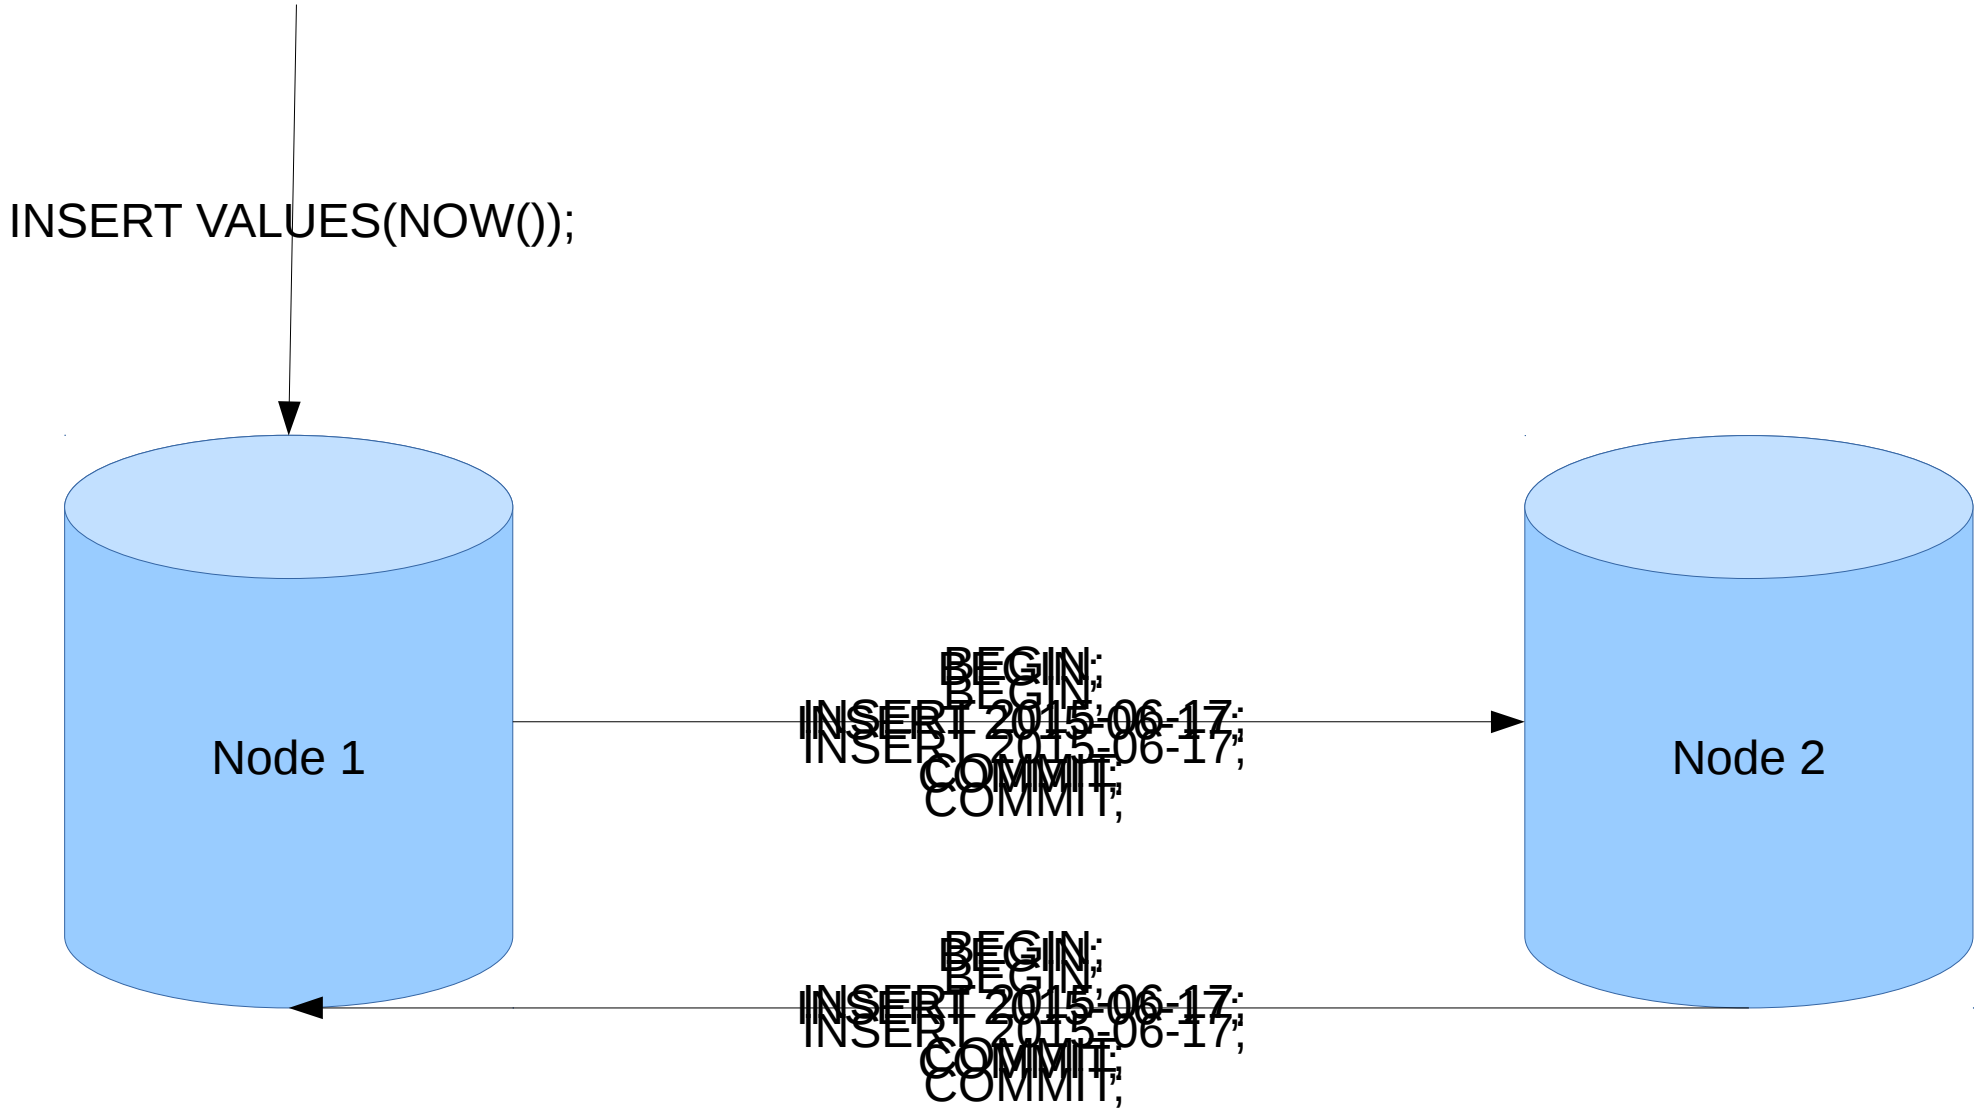
INSERT VALUES(NOW());
Node 1
BEGIN;
INSERT 2015-06-17;
COMMIT;
Node 2
BEGIN;
INSERT 2015-06-17;
COMMIT;

Replication Origin Tracking

- Identify nodes by name
- Associate session with node
- Output Plugin can decide based on origin

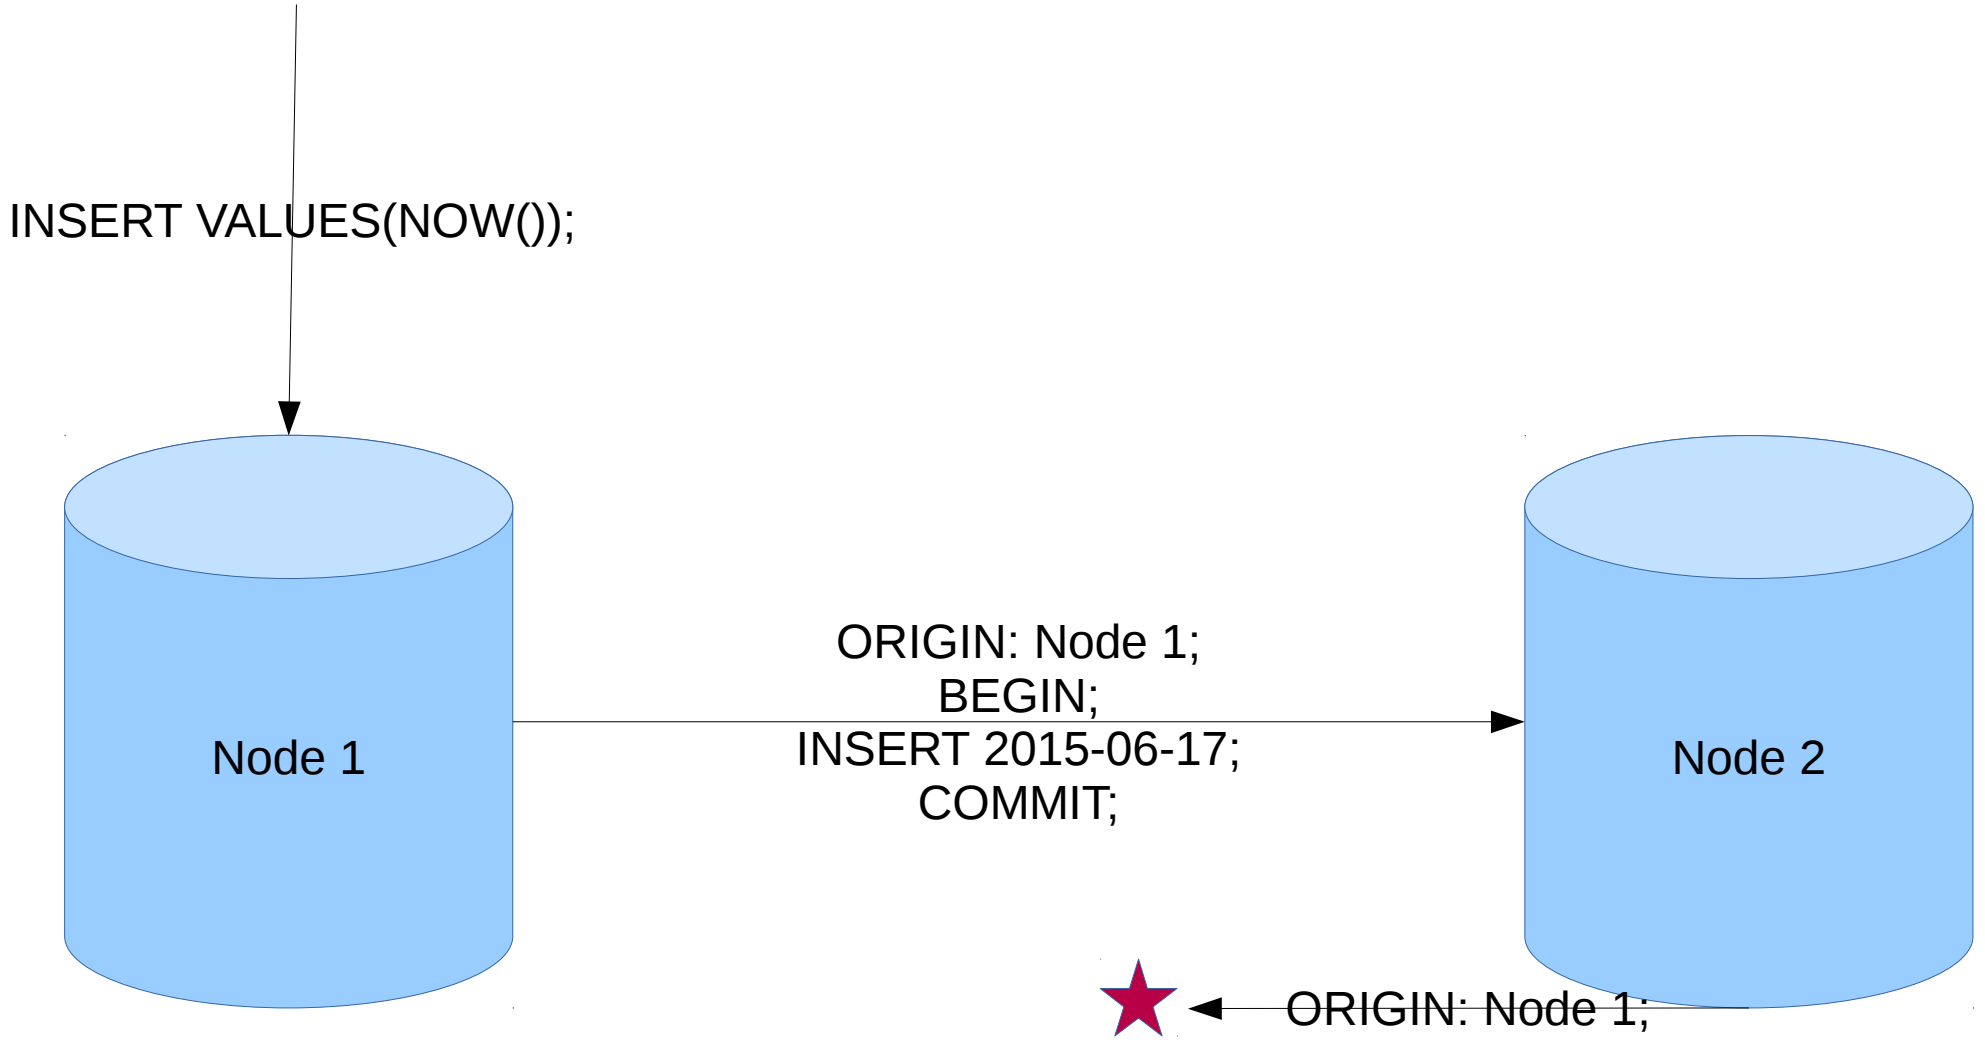
INSERT VALUES(NOW());
Node 1
ORIGIN: Node 1;
BEGIN;
INSERT 2015-06-17;
COMMIT;
Node 2
ORIGIN: Node 1;

Replication Progress Tracking

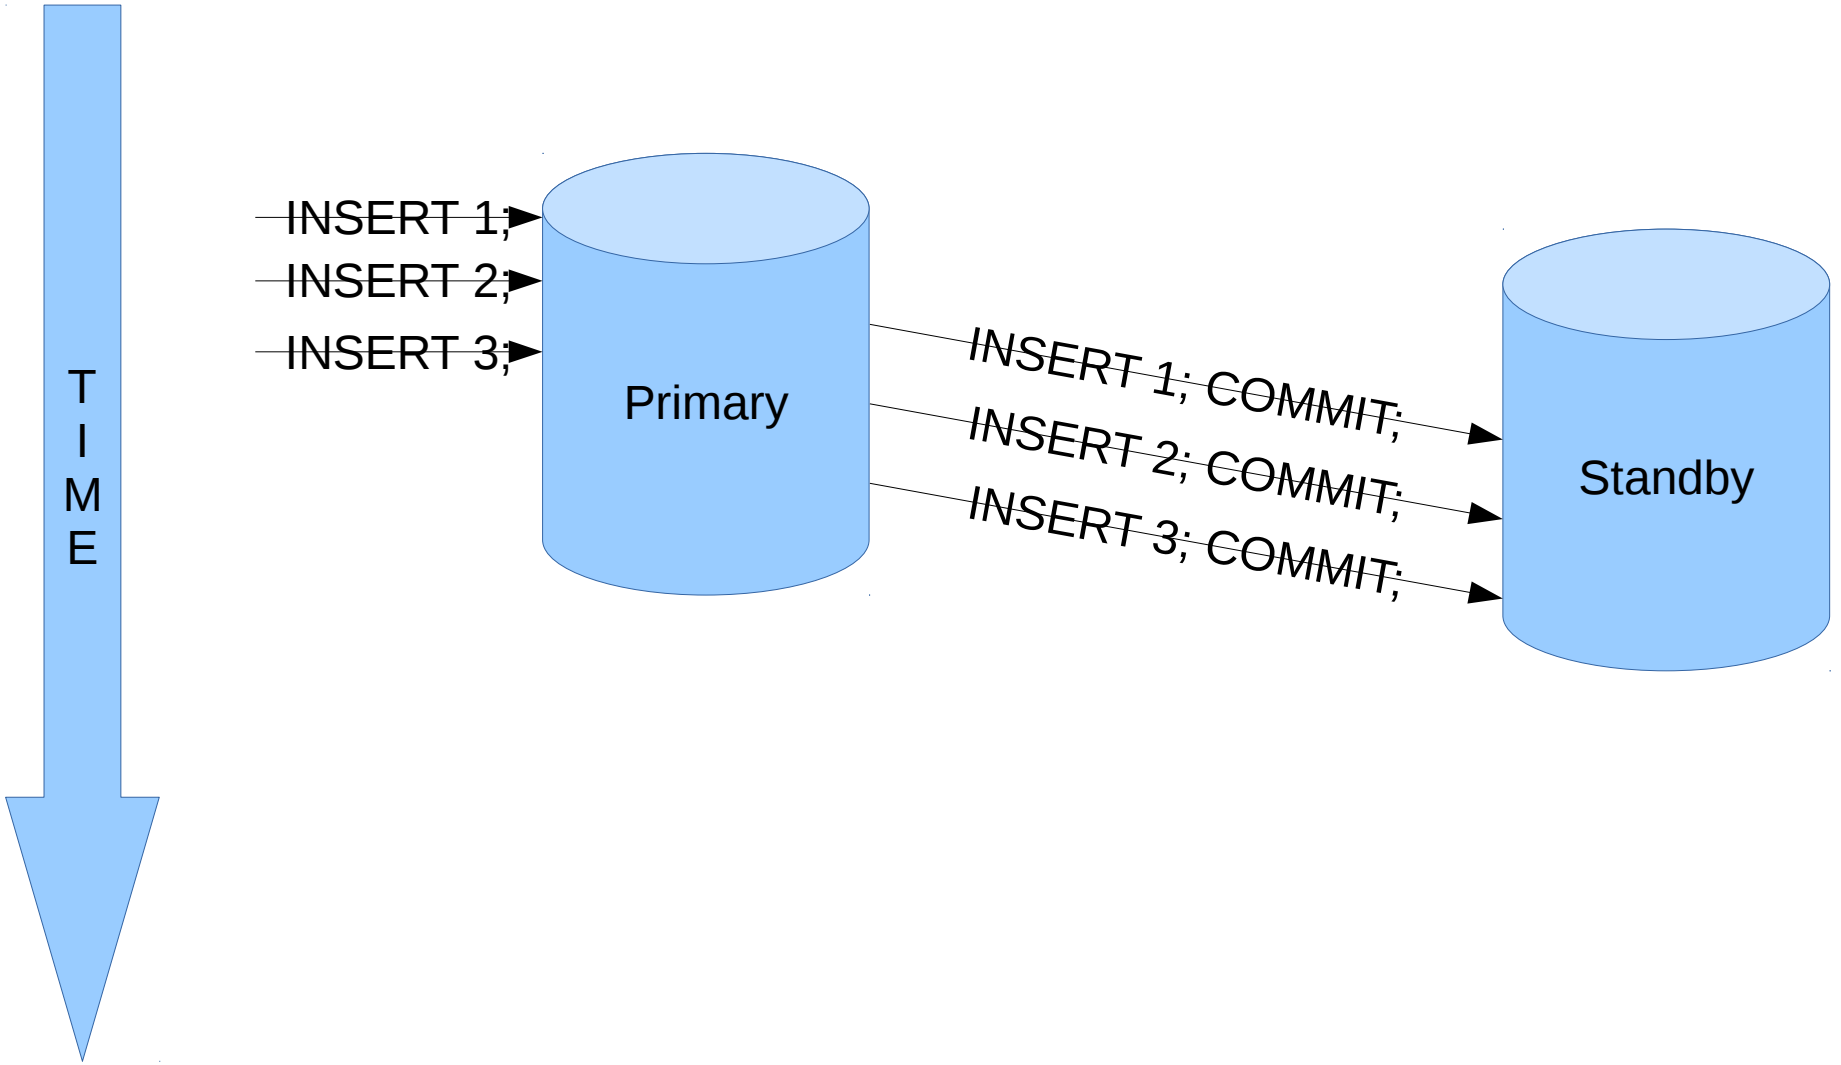

Replication Progress Tracking

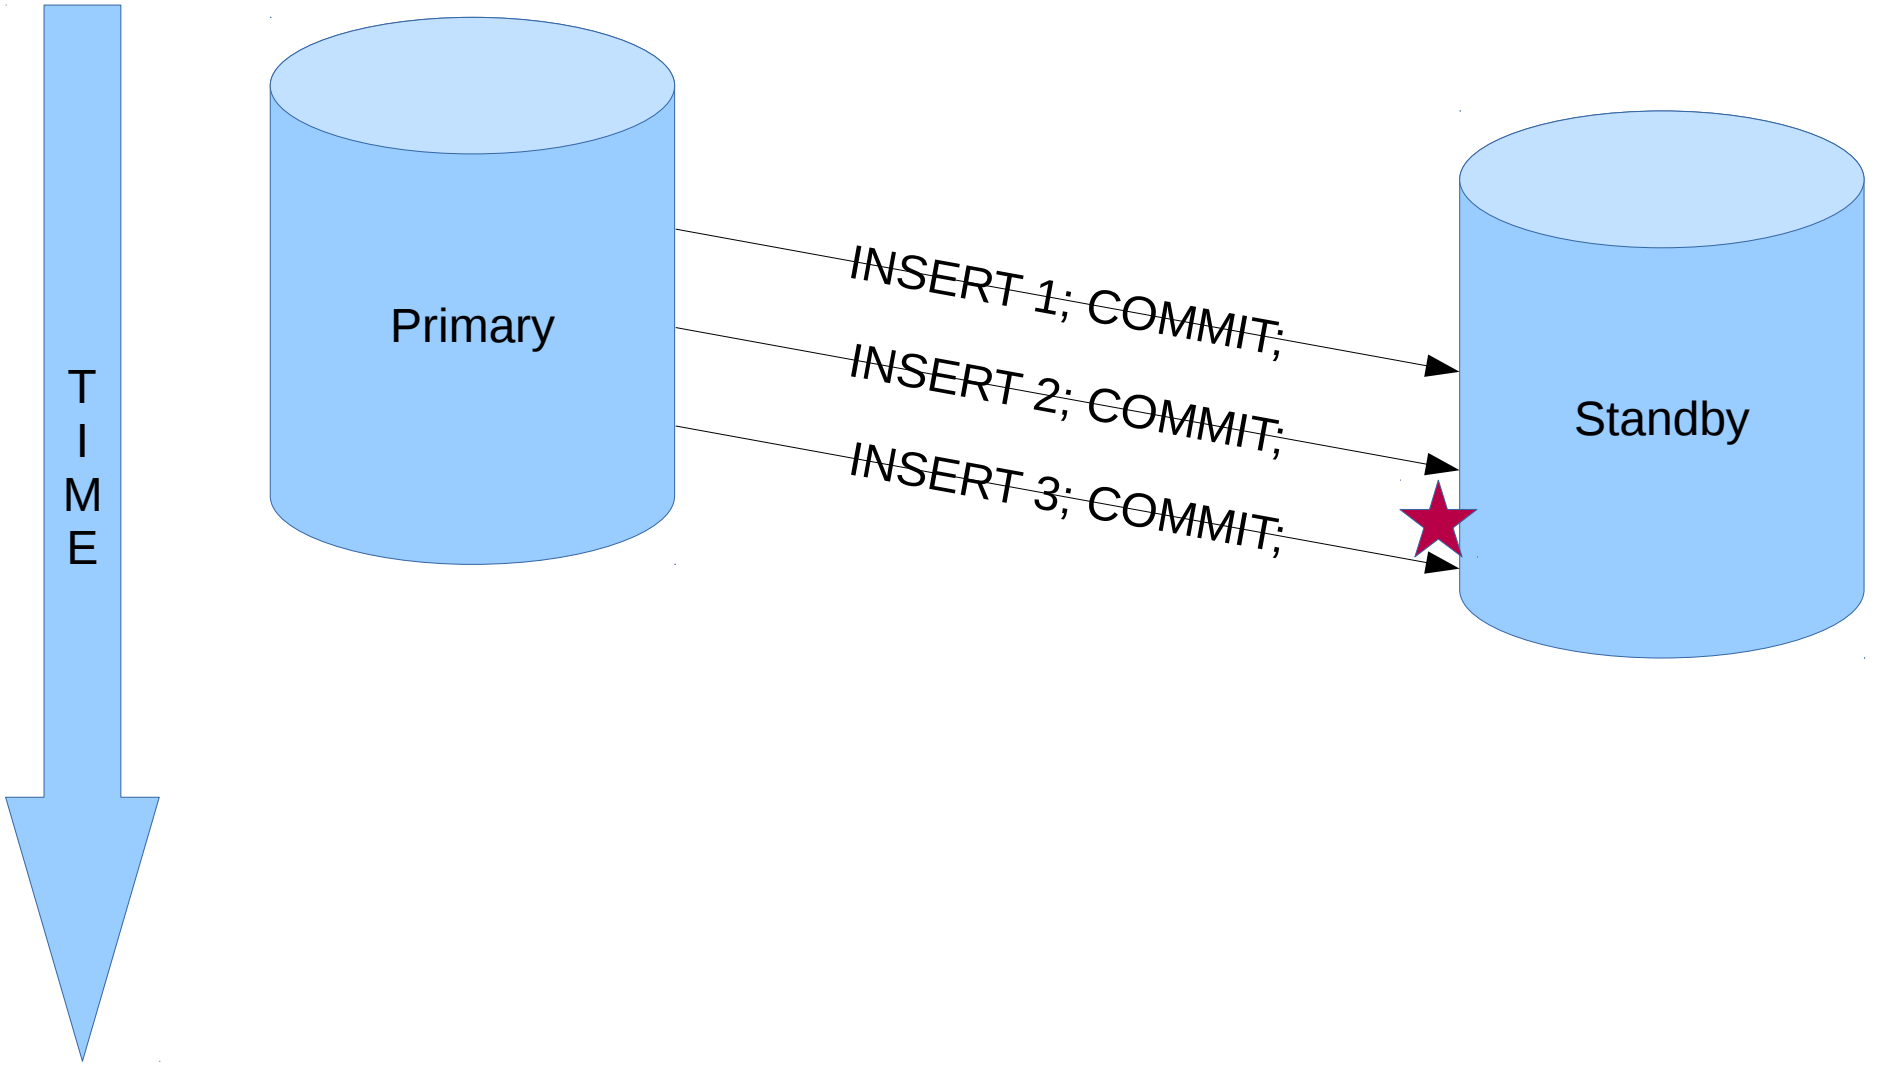

citusdata

Replication Progress Tracking

When Replaying:

```
REPLAYING_FROM: node-1
BEGIN;
REPLAYING XACT 3435/1734;
INSERT …;
COMMIT;
```

When restarting replay:

```
SELECT remote_lsn FROM pg_replication_origin_status
WHERE external_id = 'node-1';
```

DDL Replication

- DDL Replication needs normalization

  ALTER TABLE blarg
  ADD COLUMN foo sometype;

  =>

  ALTER COLUMN public.blarg

  ADD COLUMN foo schmema.sometype
- DDL Replication needs decomposition (e.g serial)

DDL Replication

- Basics integrated
- Actual normalization can be written in extension

Toolbox

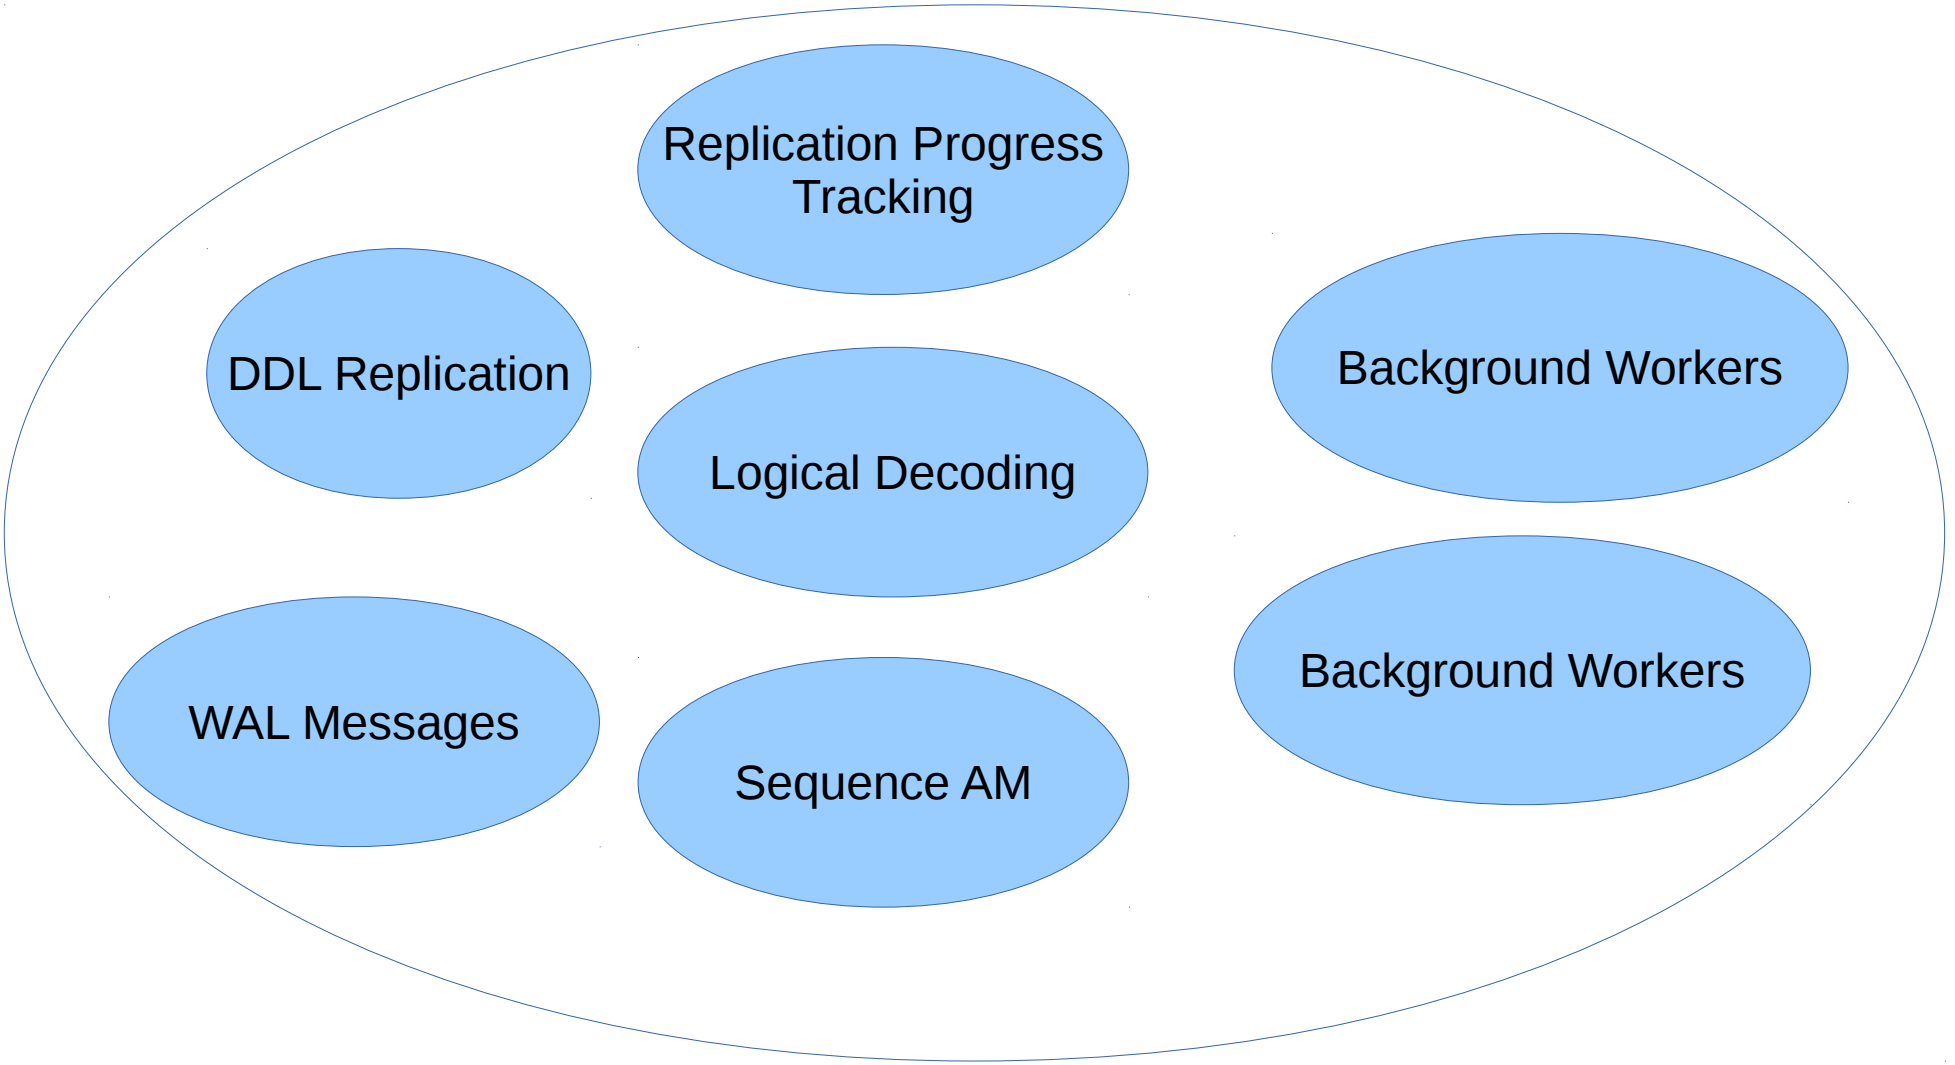
Replication Progress Tracking
DDL Replication
Logical Decoding
Background Workers
WAL Messages
Sequence AM
Background Workers

citusdata