



Choosing from the Many Flavors of PostgreSQL Replication

Richard Yen

PG Conf Silicon Valley 2019

Disclaimer

This is not an EDB talk!

Why Replication?

- Because one copy is not enough!
 - Fault-Tolerance
 - Scalability
 - Recoverability
- People want “real” data to work with

Which one is for me?

- Streaming replication
- Logical replication
- Slony
- Bucardo
- Londiste
- DRBD
- Corosync
- 2ndQ BDR
- pg_shard
- pglogical
- Daffodil
- Mammoth
- Pgpool
- EDB Replication Server
- Spectral Core Replicator Pro
- And More!

It all boils down to...

Trigger-based | WAL-based

Trigger-based Replication

- Relies on PostgreSQL's trigger mechanism to log all DML into event/shadow tables
- Data from shadow tables ingested by another program and replayed to other nodes
- Very reliable: replicating production environments since 2003

Trigger-based Replication

- Main players
 - Slony — written in C
 - Londiste — written in Python
 - Bucardo — written in Perl

Trigger-based Pros & Cons

- Select which tables to replicate
- Enables row-based/column-based replication
- Major-version upgrades
- Master-master replication (Bucardo only)

Trigger-based Pros & Cons

- Primary keys required
- Limited ability to replicate DDL
- Large object replication not available
- Trigger activity generates additional I/O
- Event tracking consumes network bandwidth
- Hard to scale up
- Infrastructure is brittle

WAL-based Replication

- WAL files are the mechanism to recover from crashes
- WAL files can be used to replay DML on remote servers
- Streaming WAL activity effectively creates an up-to-date clone

WAL-based Pros & Cons

- It's fast
- Easy to set up
- Requires no additional software, I/O, bandwidth
- Exact replica provides assurance for backups
- Exact replica means all data types are supported
- Synchronous replication is possible

WAL-based Pros & Cons

- All or nothing
- No major version upgrades
- Query cancellations
- Storage impact on primary
 - `wal_keep_segments`
 - replication slots

Middle Ground: Logical Replication

- Introduced in v. 10
- Based off Logical Decoding framework introduced in v. 9.4
- WAL files are reconstructed into SQL statements
- Paves the way to more granular replication filters
- Happy-medium

Best of Both Worlds?

- Like Trigger-based Replication:
 - Still need uniqueness/PK
 - DDL is not replicated
 - Sequences are not replicated
- Like WAL-based Replication
 - Still need to monitor storage impact
 - Can't create column-level filters (yet)

Additional Gotchas

- WAL events are unrolled into individual SQL queries
 - COPY statements are translated into INSERTs
 - UPDATE on 1000 rows are translated into 1000 individual UPDATEs
- TRUNCATEs are replicated, but not cascaded
- Partitioned tables cannot be easily replicated

Other options

- DRBD, Corosync, Windows Cluster
 - Basically, mirrors a disk or other block-level device
 - Doesn't support read-only standby
 - Use cases for these are somewhat rare
- Use at your own discretion

How do you choose?

It all depends on your requirements

Tolerances, privacy, etc.

Process of Elimination

Start with built-in Streaming Replication

- Super-simple to set up
- No additional software
- Best performance, least impact
- Good for most use-cases

Streaming **X** replication

Next, try Logical Replication

- Pick and choose which tables to replicate
- WAL-based replication performance
- Zero-downtime major-version upgrades
- Only available in v. 10 and later

Streaming Replication

Logical Replication

Next, try pglogical or EPRS

- Replication engines supporting Logical Decoding
 - pglogical
 - EnterpriseDB Postgres Replication Server 7 (EPRS7 — Limited Availability)
- Row-level and column-level filters
- Still WAL-based

Streaming Replication

Logical Replication

pglogical | EPRS7

Use a Trigger-based Solution

Use a Trigger-based Solution

- But please upgrade soon!
- Impacts performance
- Often difficult to administration
- Logical Replication is the future

RDS?

Try Bucardo

Multi-Master Replication?

Multi-Master Replication?

- Holy Grail of replication solutions?
- First, ask yourself if you REALLY need it
 - Consistency
 - Performance
 - Conflict Resolution
- Often, the other options are good enough



Multi-Master Replication

- Trigger-based
 - Bucardo
 - XDB by EnterpriseDB
- WAL-based
 - Bi-Directional Replication (BDR) by 2nd Quadrant
 - requires a custom compile/installation of PG
 - XDB/EPRS7 by EnterpriseDB

Q & A

Thank You!

Richard Yen

Principal Support Engineer, EnterpriseDB

support@enterprisedb.com

richyen.com