

# PostgreSQL Replication : (Almost) Everything You Want To Know

---

Devrim Gündüz

devrim@gunduz.org



DevrimGunduz



DevrimGunduzTR



# SELF INTRODUCTION

- Using Red Hat (and then Fedora) since 1996.
- Using PostgreSQL since 1998.
- Started building RPMs in 2002, took over the project in 2004.
- Planet PostgreSQL: 2004 → <https://planet.PostgreSQL.org>
- Responsible for PostgreSQL YUM and ZYPP repositories.
- Working at EnterpriseDB since 2011
- PostgreSQL Major Contributor
- Living in London, UK.
- **The Guy With The PostgreSQL Tattoo!**

# SOCIAL MEDIA

Please tweet:

#PostgreSQL

Please follow:

@PostgreSQL

# AS USUAL:

Please tweet:

#BlameMagnus

Please follow:

@BlameMagnus

# AS USUAL...



# UPCOMING EVENTS

- FOSDEM PGDay and FOSDEM 2020 – Brussels (Jan 31-Feb 2)
- Prague PostgreSQL Developer Day 2020 (5-6 Feb 2020)
- Nordic PGDay 2020 – Helsinki - 24 March 2020
- PGDay.Paris 2020 – Paris – 26 March 2020
- Swiss PGDay 2020 – 18/19 June 2020
- PGConf.DE Stuttgart TBA (AskAds©)
- **PostgreSQL Conference Europe 2020 BERLIN! October 20-23**
- More at: <https://www.postgresql.org/about/events/>

# History of replication in PostgreSQL

Time travel



# HISTORY OF REPLICATON IN POSTGRESQL

- Trigger based:
  - Slony
  - Bucardo
  - Londiste



# HISTORY OF REPLICATON IN POSTGRESQL

- 8.2: Warm standby
- 9.0: Initial release of in-core Streaming replication (Physical replication)
- 9.1: Synchronous replication, pg\_basebackup
- 9.2: Cascading replication
- 9.3: Follow timeline switch
- 9.4: Logical decoding, replication slots
- 9.6: Multiple sync replication, quorum, remote\_apply

# HISTORY OF REPLICATON IN POSTGRESQL

- 10: Initial release of in-core logical replication
  - pg\_basebackup: stream by default
  - Replication-ready by default
- 11: Logical replication supports TRUNCATE
- 12: Removal of recovery.conf!

# It all starts with WAL

## Some basics

# WHAT IS WAL?

- Write Ahead Log:
  - Logging of transactions
  - a.k.a. xlog in ancient times (transaction log),
  - 16 MB in most of the installations (can be configured, `--with-wal-segsize`)
  - v11+: `initdb` has a `--wal-segsize` parameter
  - `initdb --wal-segsize=64` ← in MB
  - 8 kB page size (can be configured, `--with-wal-blocksize` during `configure`)
  - `pg_resetwal --wal-segsize=64` ← in MB

# WHAT IS LSN?

- Log Sequence Number
  - Position of the record in WAL file.
  - Provides uniqueness for each WAL record.
- 64-bit integer (historically 2x32-bit) (We'll need this info soon)
- Per docs: "Pointer to a location in WAL file"
- LSN: Block ID + Segment ID (See next slides)
- During recovery, LSN on the page and LSN in the WAL file are compared.
- The larger one wins.

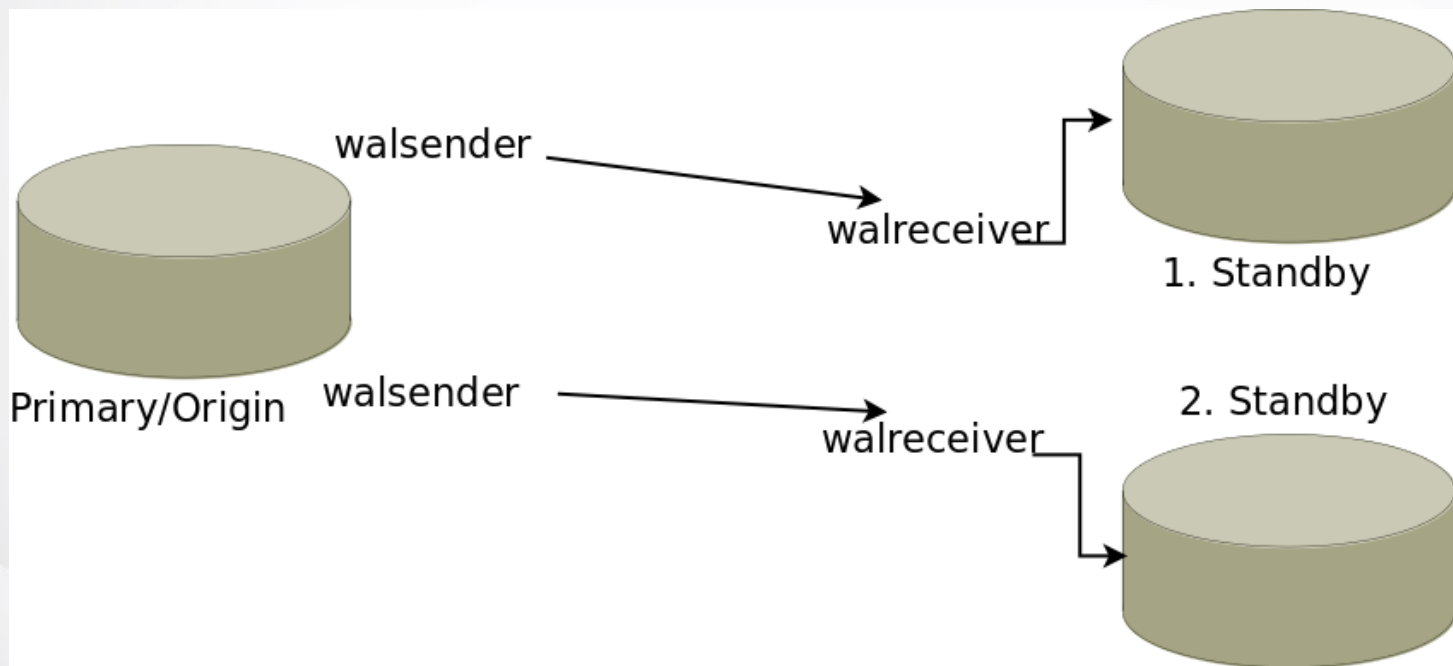
# WAL FILE NAMING

- 24 chars, hex.
  - 1st 8 chars: timelineID
    - 00000001 is the timelineID created by initdb
  - 2<sup>nd</sup> 8 chars: Block ID
  - 3<sup>rd</sup> 8 chars: Segment ID
- 000000010000000000000001 → 000000010000000000000002
- ... 0000000100000000000000FF → 000000010000000100000000
- ...and 0000000100000001000000FF → 000000010000000200000000

# Replication basics

## Terminology

# REPLICATION BASICS





# REPLICATION BASICS

- primary, origin
- standby, subscriber
- master, slave. Please.

# REPLICATION BASICS

- Base backup
- walsender
- Replication slot
- Logical decoding

# Replication parameters

## Primary server

# REPLICATION PARAMETERS: PRIMARY

- `synchronous_commit` (on, off, local, remote\_write, remote\_apply)
- `max_wal_senders`
- `wal_keep_segments`
- `synchronous_standby_names` (FIRST, ANY)

# Replication parameters

## Standby server

# REPLICATION PARAMETERS: STANDBY

- `primary_conninfo`
- `primary_slot_name`
- `promote_trigger_file`
- `hot_standby`
- `hot_standby_feedback`
- `recovery_min_apply_delay` (time delayed standby)

# Streaming replication

## General features

# STREAMING REPLICATION: GENERAL FEATURES

- Replication of whole cluster
- WAL-logged transactions are replicated
- Works on the PostgreSQL port
- No built-in auto failover/failback
  - Patroni, repmgr, PAF
  - Closed source solutions are also available



# STREAMING REPLICATION: REPLICATION USER

- Separate user for replication
  - `CREATE ROLE blamemagnus PASSWORD 'foobar' REPLICATION LOGIN;`
  - `pg_hba.conf`

# Streaming replication: Taking base backup

**pg\_basebackup**

# BASE BACKUP: BASICS

- **The** prerequisite for setting up streaming replication
- Should be on the same OS/patch level
- Physical backup of everything in the instance
- Taken from primary to standby(s)
- Also used for PITR/backup
- `pg_hba.conf` (on primary)

# BASE BACKUP: PG\_BASEBACKUP

- pg\_basebackup
  - -D
  - -Fp (default)
  - -R
  - -X stream (default)
  - -c fast/spread (default)
  - -C -S slot\_name
  - -P

# BASE BACKUP: PG\_BASEBACKUP

- `pg_basebackup`
  - `-p`
  - `-h`
  - `-U`

# BASE BACKUP: PG\_BASEBACKUP

- Example:

```
pg_basebackup -D /var/lib/pgsql/12/repdata -Fp -R -c  
fast -C -S blamemagnus -P -h 192.168.100.10 -p 5412 -  
U blamemagnus
```

# Replication configuration: Standby server

# REPLICATION CONFIGURATION: STANDBY

- recovery.conf < 12, postgresql.auto.conf and postgresql.conf >= 12
  - pg\_basebackup -R
  - application\_name in primary\_conninfo (useful, and also needed for sync replication)



# Logical replication

## General features

# LOGICAL REPLICATION: GENERAL FEATURES

- Not a replacement of streaming replication
- Different use cases
- `wal_level=logical`
- Different features
  - Replication between different major versions
  - Single table/database replication
  - Replication of set of tables
  - Writeable replica

# LOGICAL REPLICATION: RESTRICTIONS

- Schema/DDDL cannot be replicated
- Large objects are not replicated
- Sequences are not replicated
- Views, MV, foreign tables, partition root tables are not replicated

# LOGICAL REPLICATION: EXAMPLES

- `CREATE TABLE t1 (c1 int);`
- `CREATE PUBLICATION pgpub FOR TABLE t1;`
- Alternatives:
  - `CREATE PUBLICATION pgpub FOR TABLE t1,t2;`
  - `CREATE PUBLICATION pgpub FOR ALL TABLES;`
  - `CREATE PUBLICATION pgpub FOR TABLE t1`  
`WITH (publish = 'insert');`

# LOGICAL REPLICATION: EXAMPLES

- Create tables on standby first. `pg_dump --schema` will help.
- `CREATE SUBSCRIPTION pgsub CONNECTION 'dbname=postgres host=localhost user=blamemagnus port=5412' PUBLICATION pgpub;`
- Initial data copy is done.

# LOGICAL REPLICATION: EXAMPLES

- `CREATE TABLE t1 (c1 in);`
- `CREATE PUBLICATION pgpub FOR TABLE t1;`
- Alternatives:
  - `CREATE PUBLICATION pgpub FOR TABLE t1,t2;`
  - `CREATE PUBLICATION pgpub FOR ALL TABLES;`
  - `CREATE PUBLICATION pgpub FOR TABLE t1  
WITH (publish = 'insert');`

# Some important points

# TIPS

- Cascading replication?
- Issues on standby server
  - Replication delays
  - Network / hardware problems
- What happens when replica is dropped?



# Replication monitoring

# REPLICATION MONITORING: PRIMARY

- `pg_stat_replication`
- `pg_replication_slots`

# REPLICATION MONITORING: STANDBY

- `pg_stat_wal_receiver;`
- `pg_is_in_recovery()`

# PHOTO TIME

# @CheerPostgreSQL

# QUESTIONS & DISCUSSION

# THANK YOU

---

Devrim Gündüz  
devrim@gunduz.org  
Twitter: @DevrimGunduz



**EDB**<sup>™</sup>  
POSTGRES