Replication in Postgres

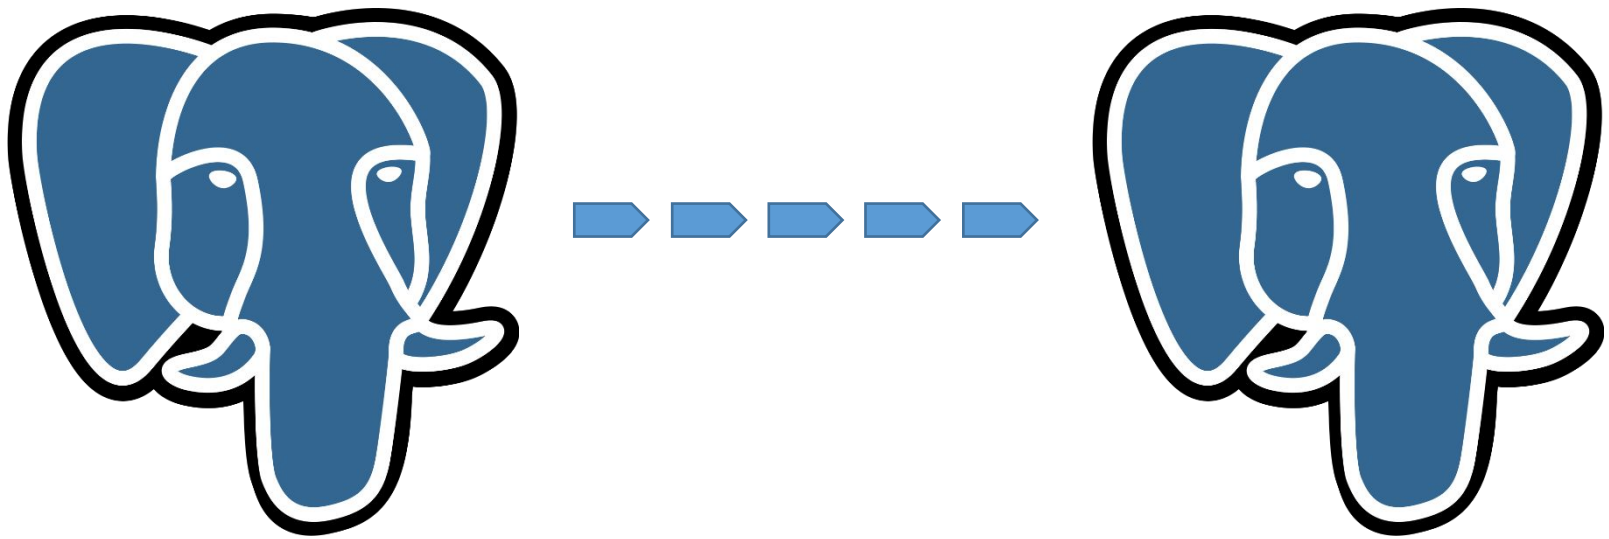

Agenda

- Replikasyon nedir? Neden ihtiyaç vardır?
- Log-Shipping nedir?
- High Availability'ye ve Load Balancing'e nasıl etkisi vardır?
- Failover anında bizi nasıl kurtarır?
- Core PostgreSQL Replikasyon nasıl yapılır ve tipleri nelerdir? Örnek topoloji.
- Streaming Replication ve avantajları nelerdir?
- Cascading Replication ve detayları nelerdir
- Kurulumdaki master ve standby'ın konfigurasyonu
- Replikasyon için ayarlanması gereken önemli parametreler hangileridir?
- Postgresql 10 ile gelen Logical Repikasyon ve Quorum Commit
- Kahoot Uygulaması ile yarışma→ https://kahoot.it/ veya Uygulamayı indir!!

What is Replication?

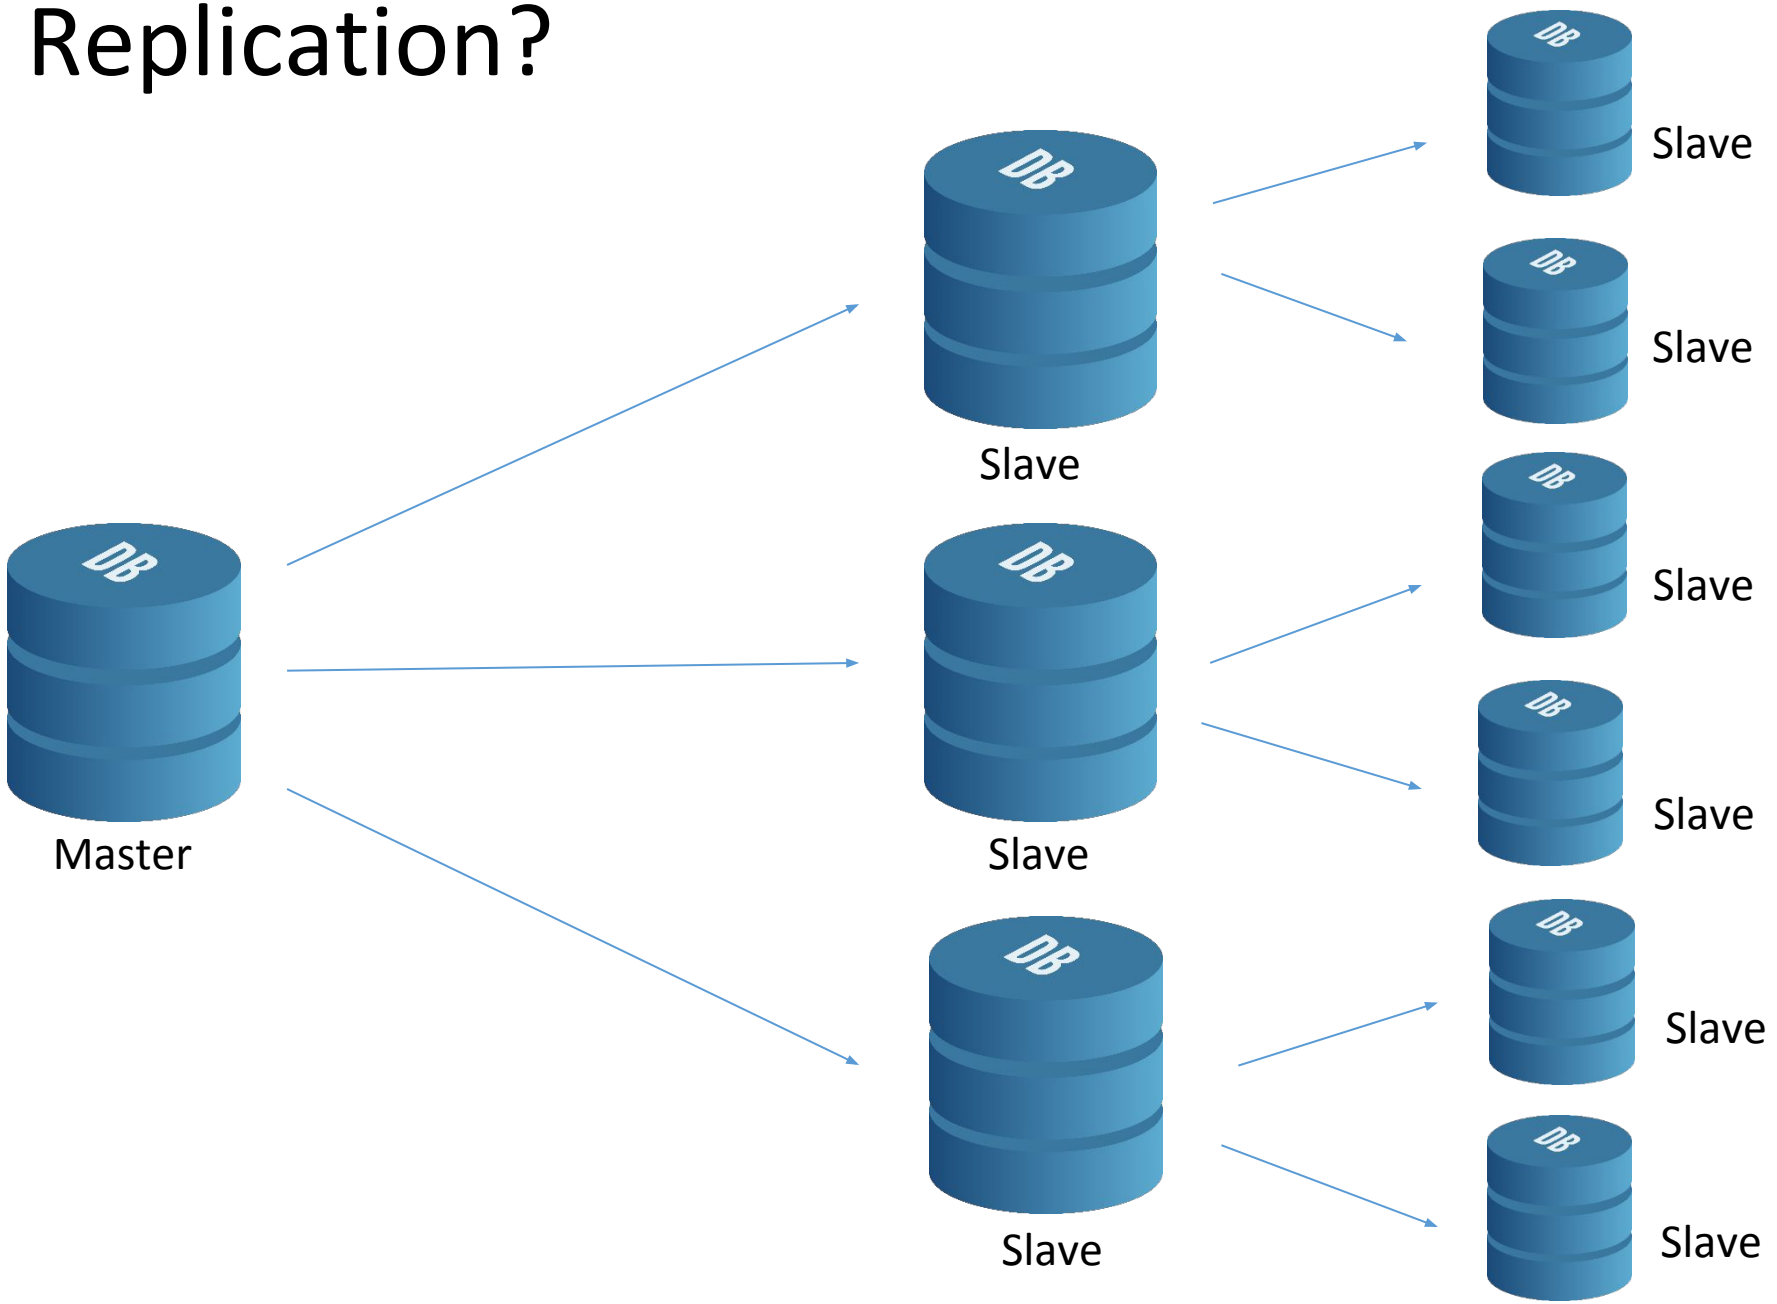

Replication Layers

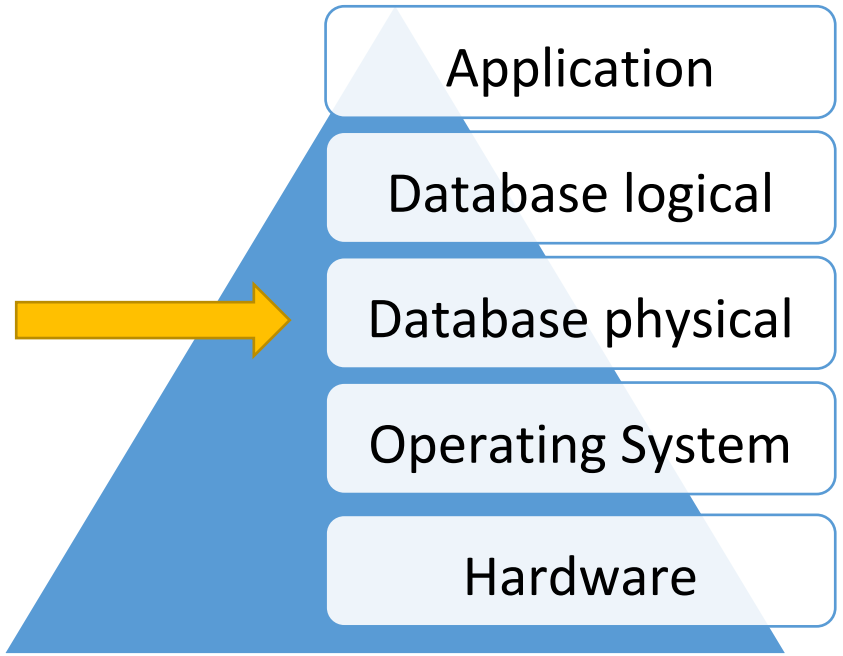

High Availability, Load Balancing, and Replication Feature Matrix

Table 25-1. High Availability, Load Balancing, and Replication Feature Matrix

Feature	Shared Disk Failover	File System Replication	Transaction Log Shipping	Trigger-Based Master-Standby Replication	Statement-Based Replication Middleware	Asynchronous Multimaster Replication	Synchronous Multimaster Replication
Most Common Implementation	NAS	DRBD	Streaming Repl.	Slony	pgpool-II	Bucardo	
Communication Method	shared disk	disk blocks	WAL	table rows	SQL	table rows	table rows and row locks
No special hardware required		•	•	•	•	•	•
Allows multiple master servers					•	•	•
No master server overhead	•		•		•		
No waiting for multiple servers	•		with sync off	•		•	
Master failure will never lose data	•	•	with sync on		•		•
Standby accept read-only queries			with hot	•	•	•	•
Per-table granularity				•		•	•
No conflict resolution necessary	•	•	•	•			•

https://www.postgresql.org/docs/9.5/static/different-replication-solutions.html

What is Replication?

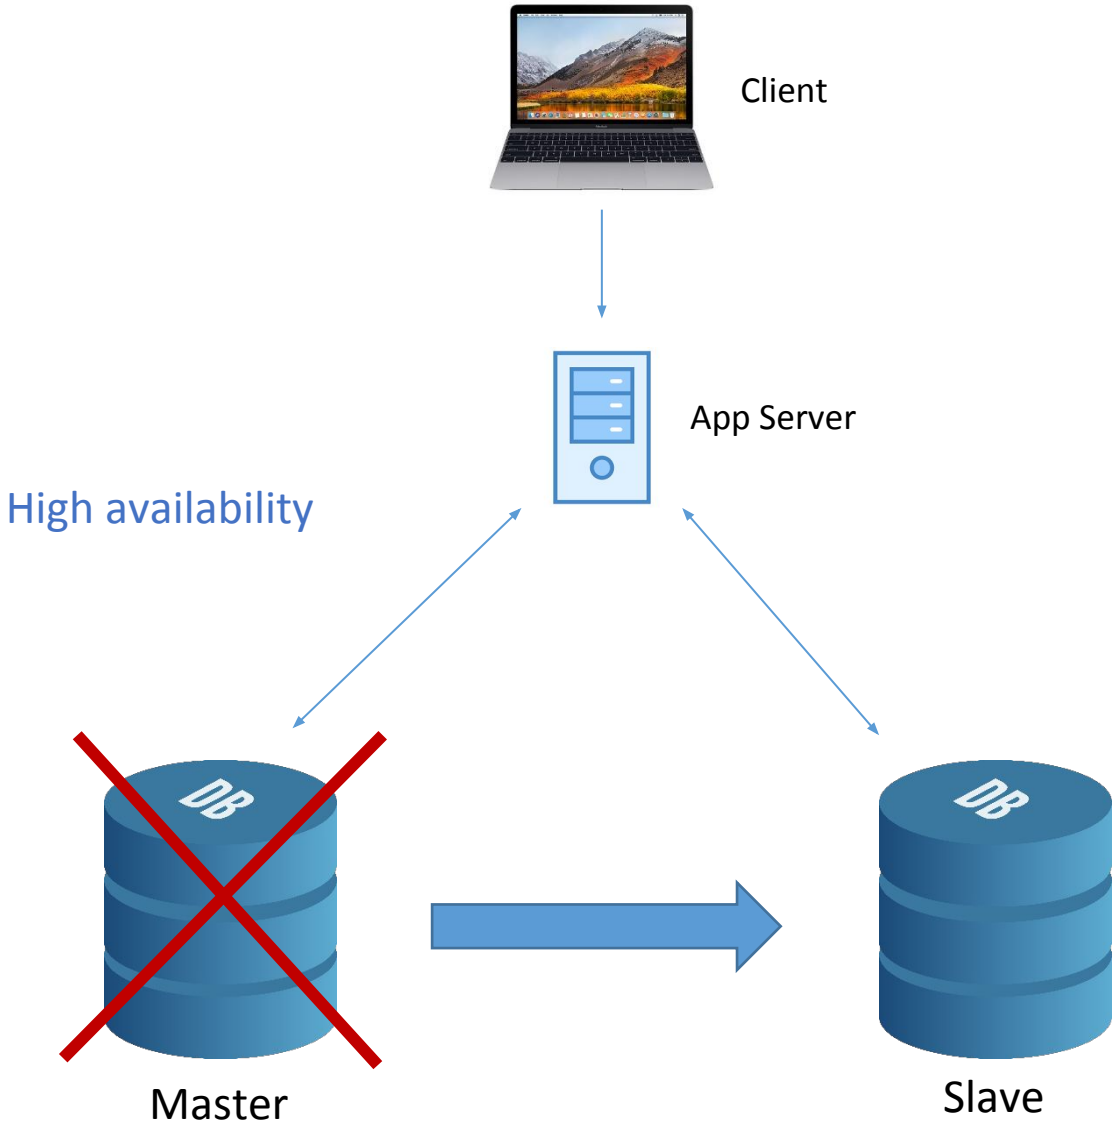

What is Replication?

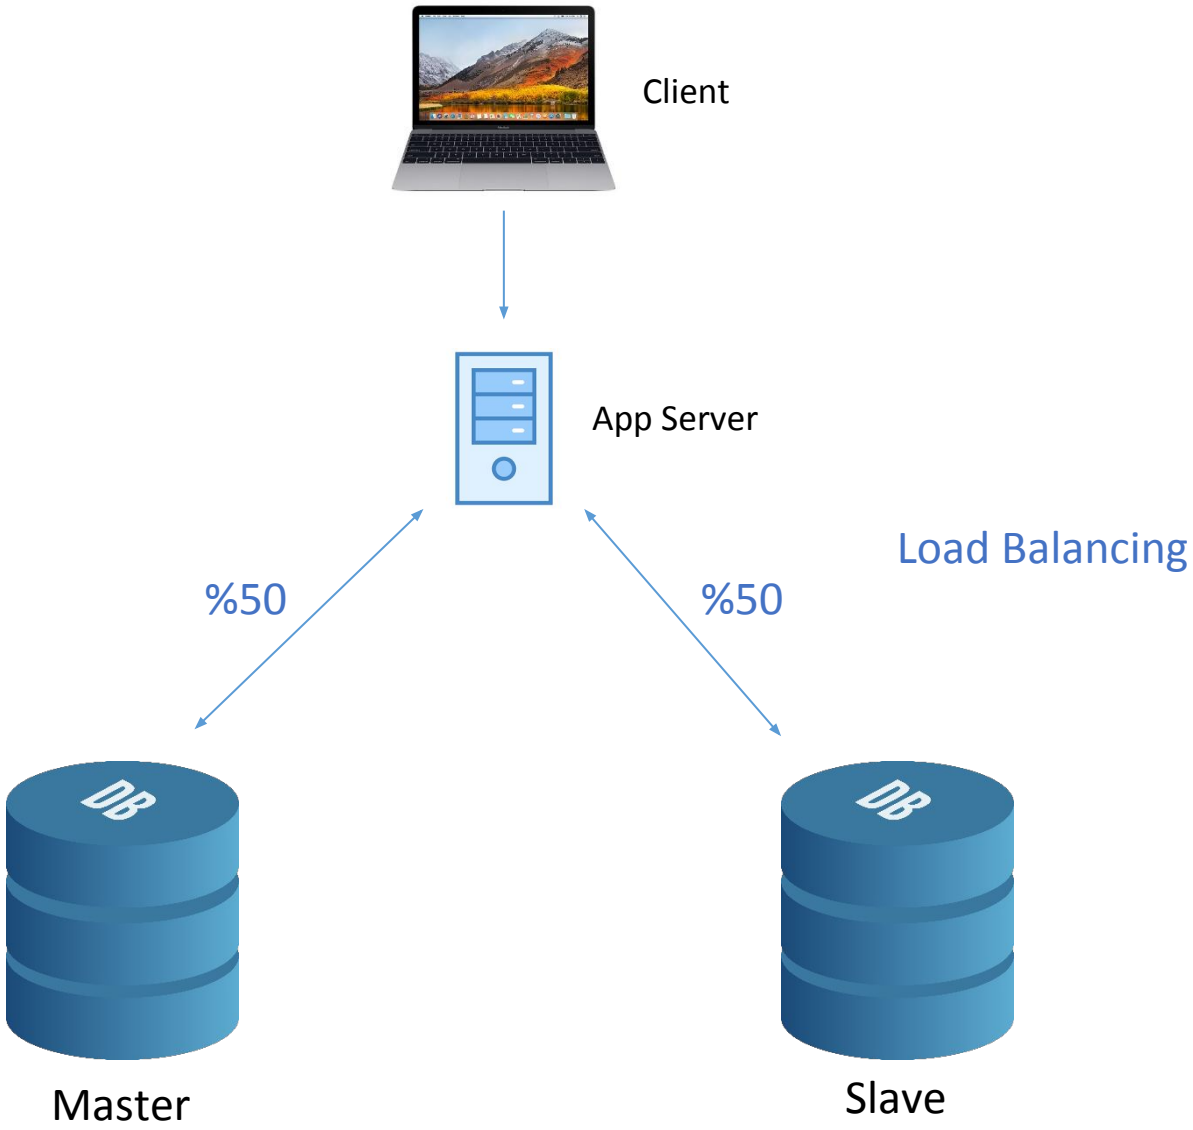

Database Physical

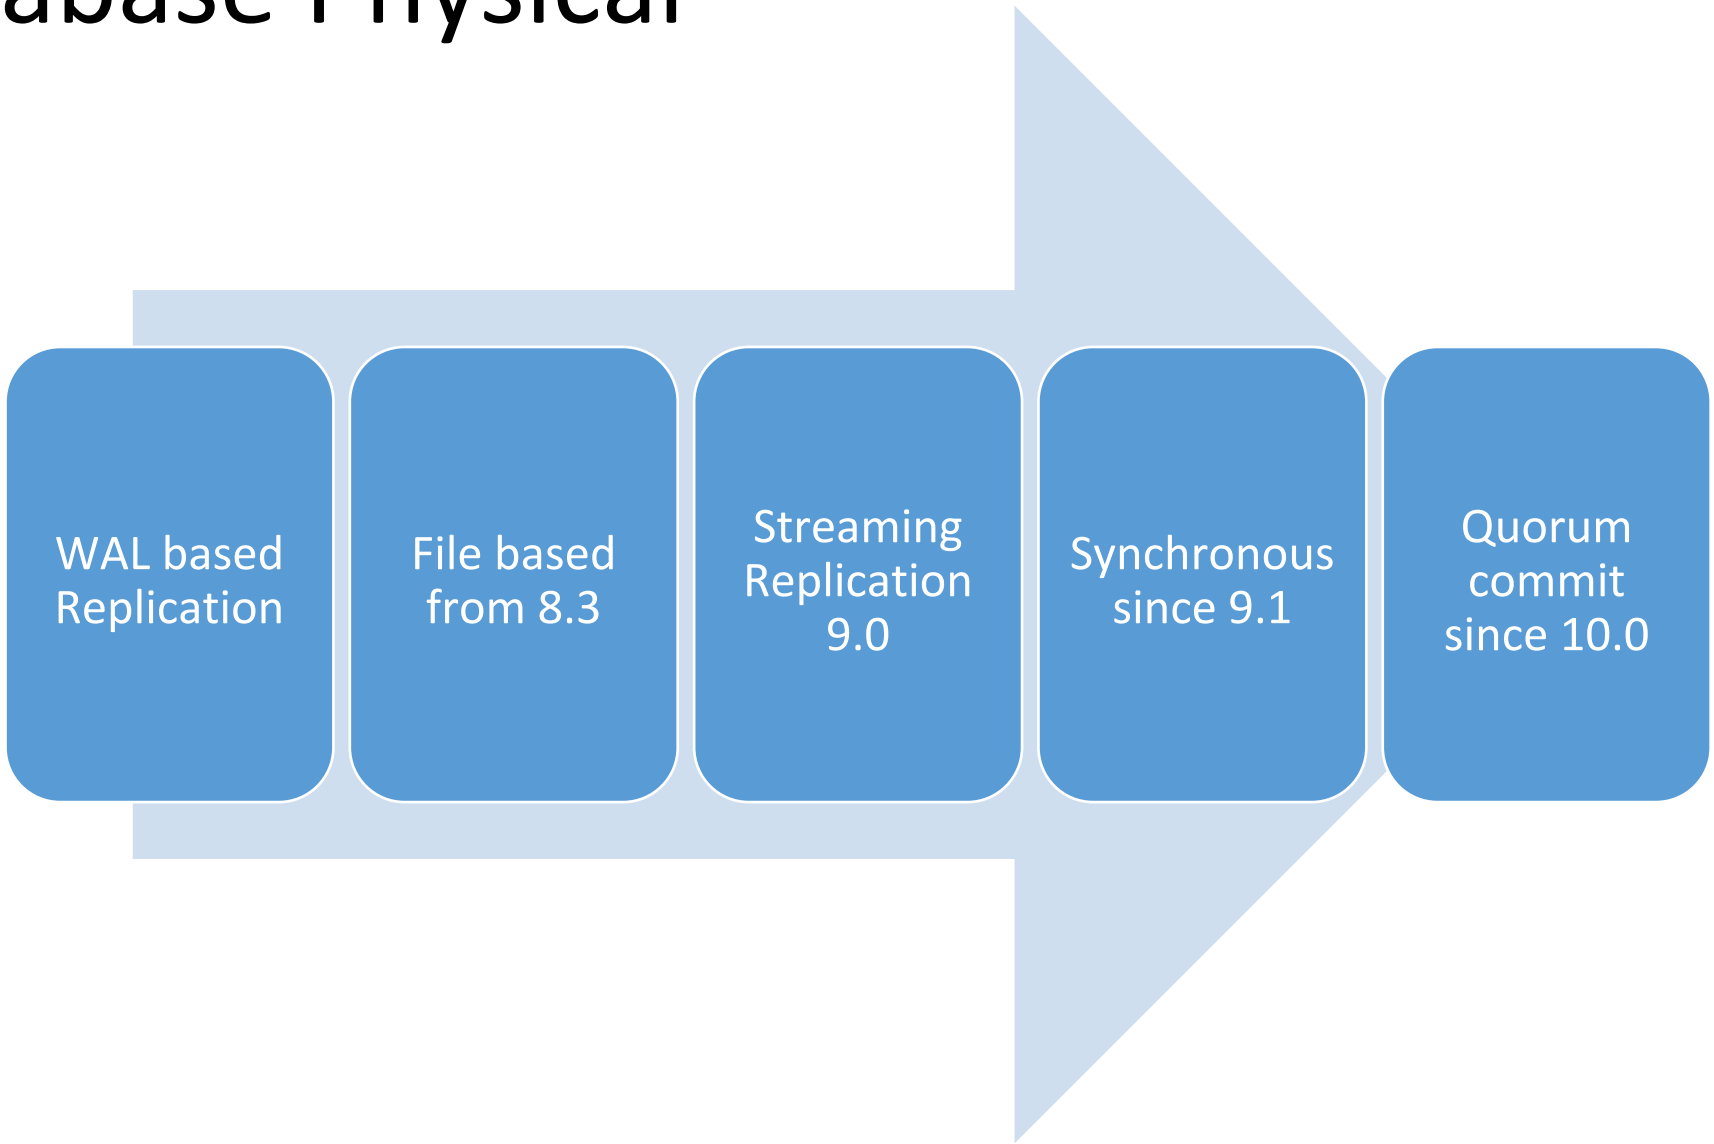

Streaming Replication

The primary and standby servers so that they are as similar as possible

1- Major PostgreSQL release levels is not possible

2- 32-bit to a 64-bit system will not work.

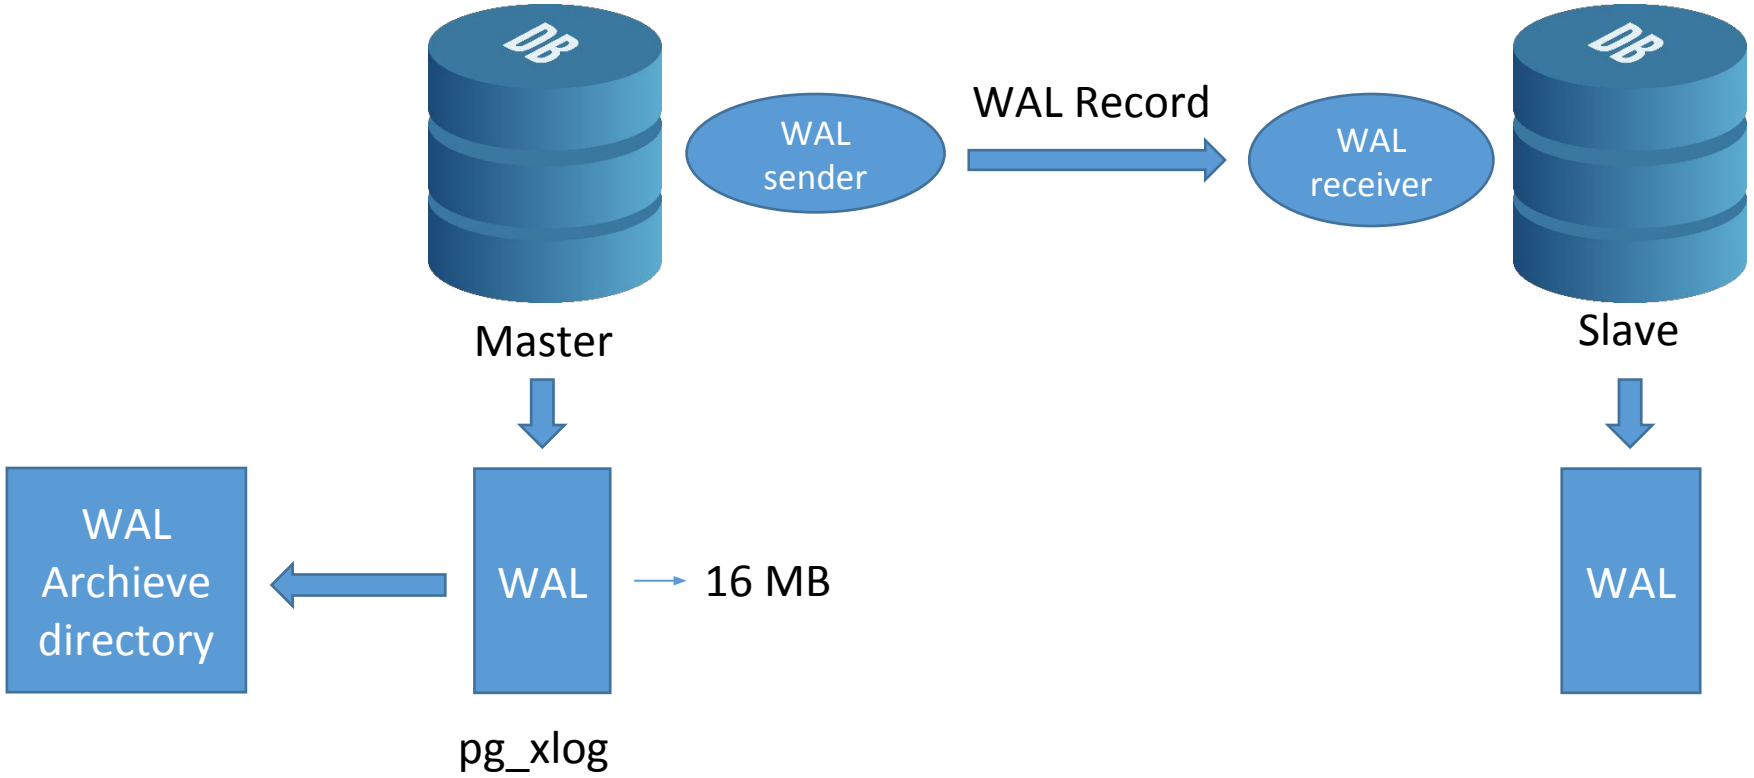

Streaming Replication

1- Async vs Sync

2- Hot Standby or not?

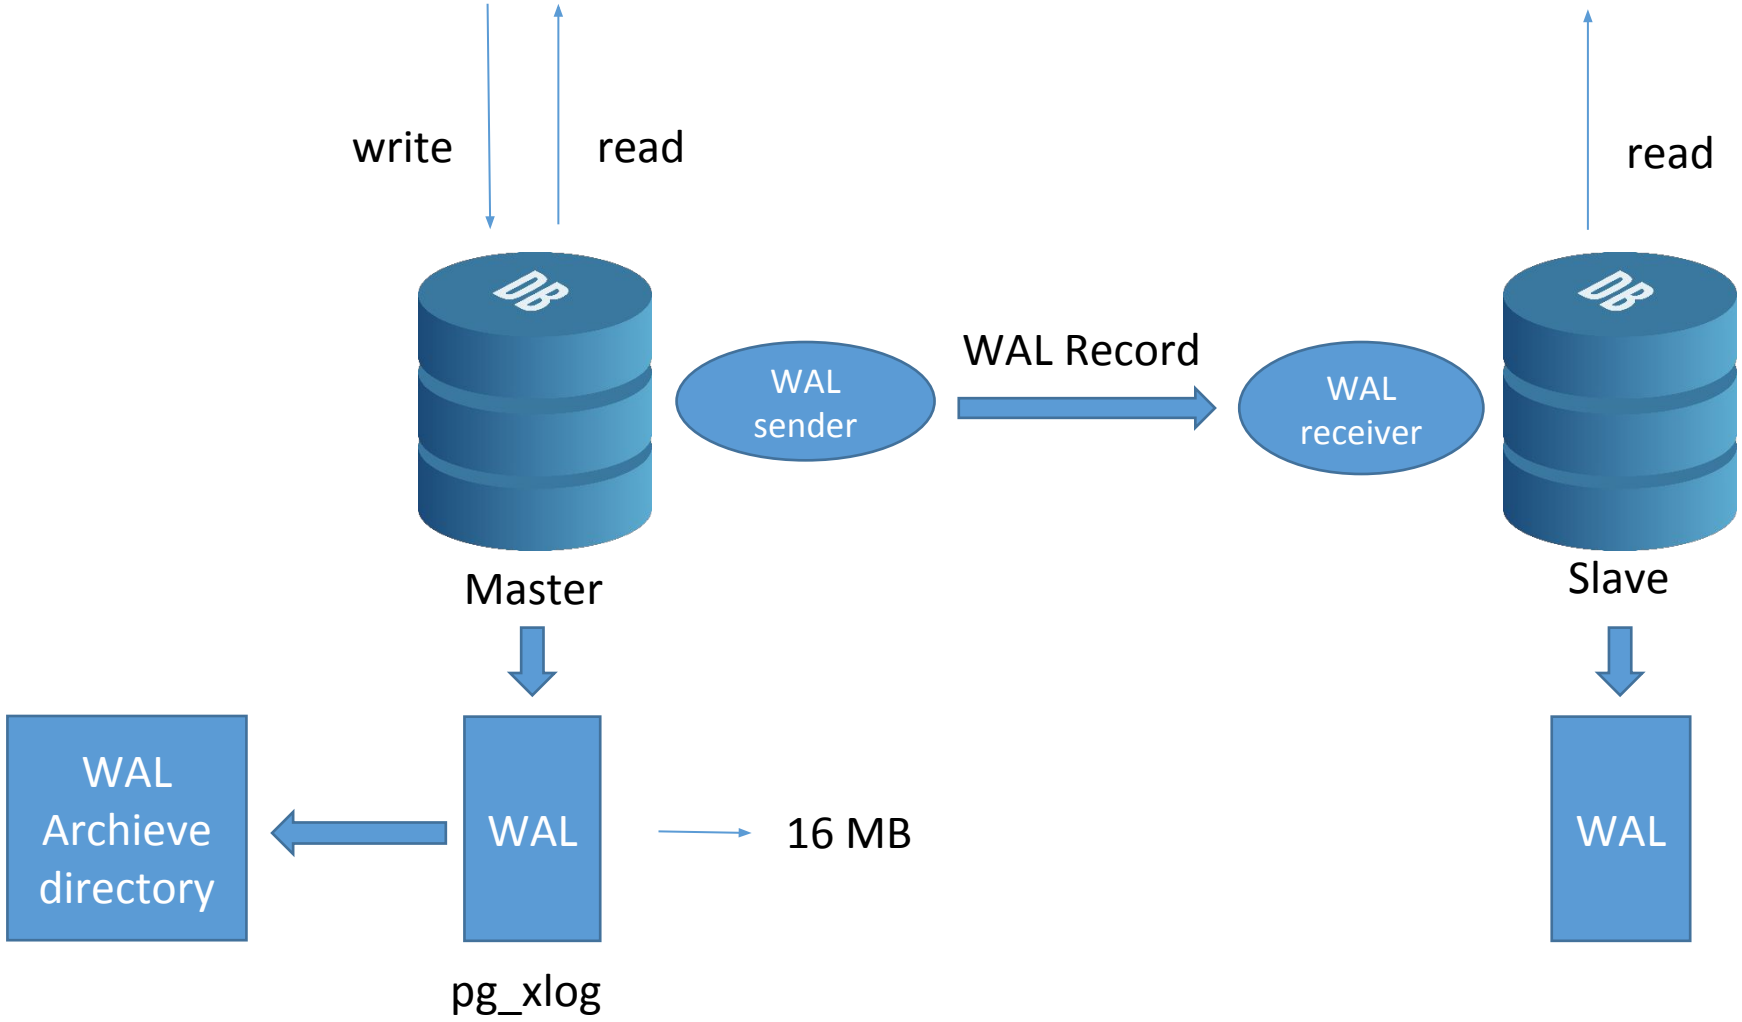

Hot Standy - postgresql.conf

- wal_level → determines how much information is written

wal_level=‘minimal’

wal_level=‘archive’

wal_level=‘hot_standby’

max_wal_senders

- max_wal_senders= specifies the maximum number of concurrent connections

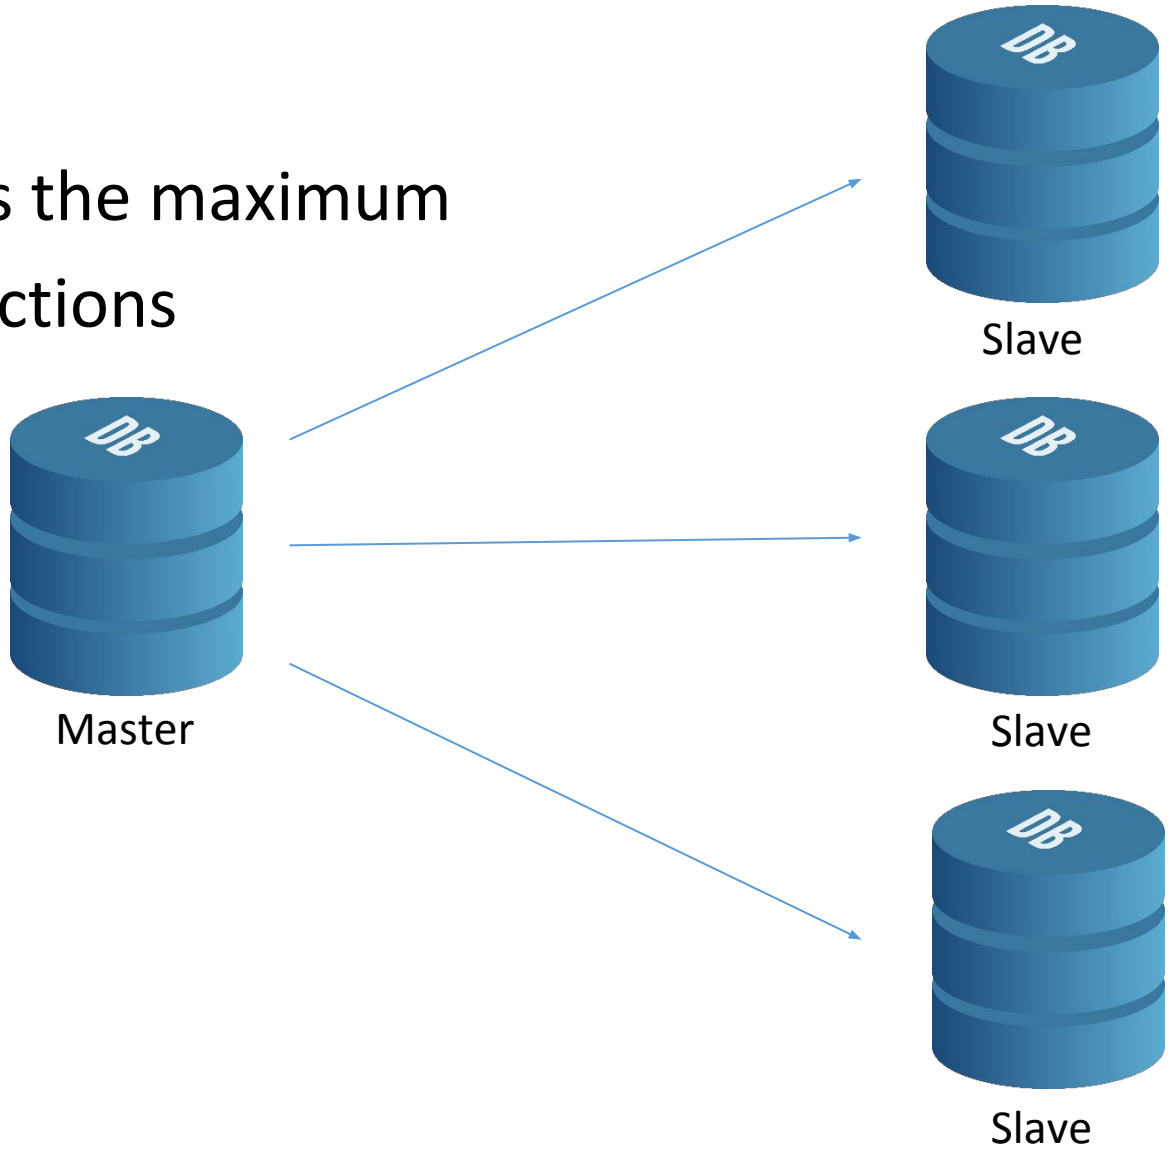

max_wal_segments

```
root@dfast-postgresql-prod:/var/lib/postgresql/9.5/main/pg_xlog# ls -ltr
total 327688
-rw------- 1 postgres postgres      311 Sep 16 02:54 000000010000001C5000000CB.00000028.backup
-rw------- 1 postgres postgres 16777216 Sep 18 14:33 000000010000001CC00000054
-rw------- 1 postgres postgres 16777216 Sep 18 14:39 000000010000001CC00000057
-rw------- 1 postgres postgres 16777216 Sep 18 14:41 000000010000001CC00000055
-rw------- 1 postgres postgres 16777216 Sep 18 14:43 000000010000001CC00000056
-rw------- 1 postgres postgres 16777216 Sep 18 14:45 000000010000001CC00000044
-rw------- 1 postgres postgres 16777216 Sep 18 14:46 000000010000001CC00000045
-rw------- 1 postgres postgres 16777216 Sep 18 14:49 000000010000001CC00000046
-rw------- 1 postgres postgres 16777216 Sep 18 14:52 000000010000001CC00000047
-rw------- 1 postgres postgres 16777216 Sep 18 14:53 000000010000001CC00000048
-rw------- 1 postgres postgres 16777216 Sep 18 14:55 000000010000001CC00000049
-rw------- 1 postgres postgres 16777216 Sep 18 14:57 000000010000001CC0000004A
-rw------- 1 postgres postgres 16777216 Sep 18 15:00 000000010000001CC0000004B
-rw------- 1 postgres postgres 16777216 Sep 18 15:01 000000010000001CC0000004C
-rw------- 1 postgres postgres 16777216 Sep 18 15:03 000000010000001CC0000004D
-rw------- 1 postgres postgres 16777216 Sep 18 15:04 000000010000001CC0000004E
-rw------- 1 postgres postgres 16777216 Sep 18 15:06 000000010000001CC0000004F
-rw------- 1 postgres postgres 16777216 Sep 18 15:08 000000010000001CC00000050
-rw------- 1 postgres postgres 16777216 Sep 18 15:10 000000010000001CC00000051
-rw------- 1 postgres postgres 16777216 Sep 18 15:12 000000010000001CC00000052
drwx------ 2 postgres postgres     4096 Sep 18 15:12 archive_status
-rw------- 1 postgres postgres 16777216 Sep 18 15:14 000000010000001CC00000053
root@dfast-postgresql-prod:/var/lib/postgresql/9.5/main/pg_xlog#
```

max_wal_segments

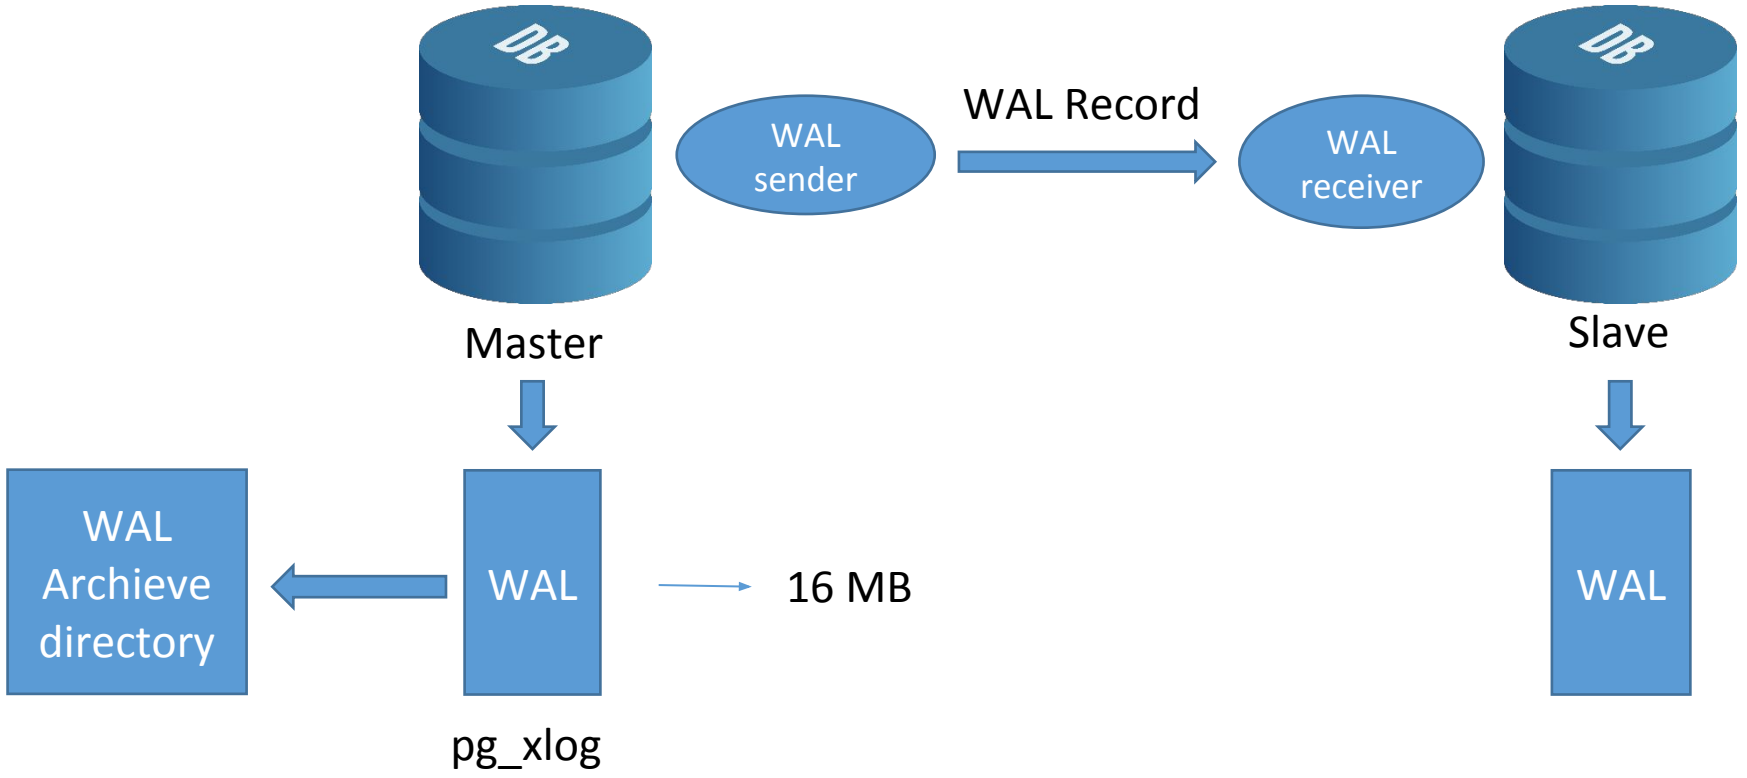

Replication User

- sudo -u postgres psql

Next, create a new user and role with the following command:

- postgres=#CREATE USER replica REPLICATION LOGIN ENCRYPTED PASSWORD ‘*******';

postgres=#\du

- You should see the following output:

```
                                   List of roles
 Role name |                         Attributes                         | Member of
-----------+------------------------------------------------------------+-----------
 postgres  | Superuser, Create role, Create DB, Replication, Bypass RLS | {}
 replica   | Replication                                                | {}
```

Hot Standby Configuration for Master

in postgresql.conf

- wal_level=hot_standby
- wal_keep_segment=20
- max_wal_sender=3
- archieve_mode=on
- archive_command = 'test ! -f /var/lib/postgresql/pg_log_archive/%f && cp %p /var/lib/postgresql/pg_log_archive/%f'

pg_hba.conf configuration for Master

For authentication:

host replication replica 10.70.82.60/32 md5

Hot standy configuration for slave

In Postgresql.conf

- hot standby=on

Below configuration in case of fail over

- archive_mode = on
- archive_command = 'test ! -f /var/lib/postgresql/pg_log_archive/%f && cp %p /var/lib/postgresql/pg_log_archive/%f'
- wal_keep_segment=20
- max_wal_sender=3

Syncronize Data from Master Server to Slave Server

On the slave server, stop the postgresql service:

- sudo systemctl stop postgresql and move existing data folder.

pg_basebackup -h 10.70.82.30 -U replica -D /var/lib/postgresql/9.5/main -P –xlog

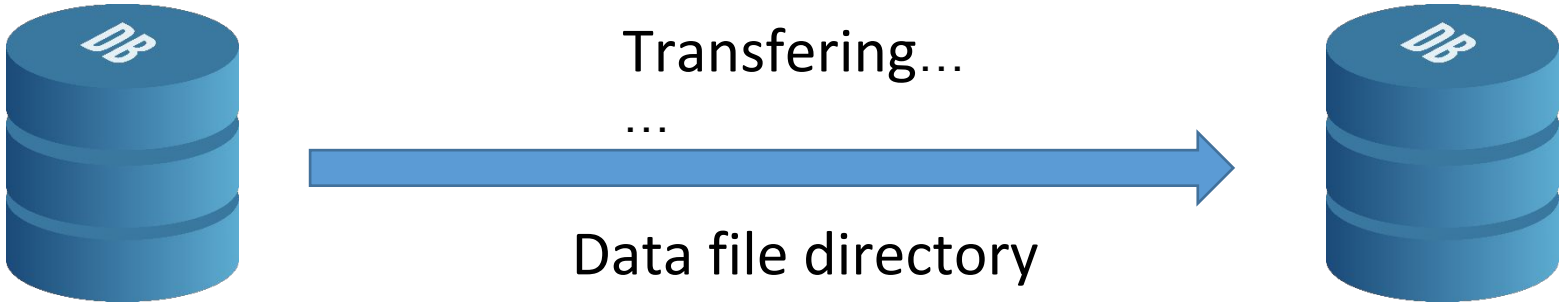

Recovery.conf file on standby

Datafile Directory→/var/lib/postgresql/9.5/main

```
standby_mode = 'on'
primary_conninfo = 'host=192.168.1.110 port=5432 user=replica
password=[redacted]'
restore_command = 'cp //var/lib/postgresql/9.4/main/archive/%f %p'
trigger_file = '/tmp/postgresql.trigger.5432'
```

Test Replication

Command→psql -x -c "select * from pg_stat_replication;"

```
postgres@dfast-postgresql:/var$ psql -x -c "select * from pg_stat_replication;"
-[ RECORD 1 ]----+--------------------------
pid              | 21104
usesysid         | 1426344
usename          | replica
application_name | walreceiver
client_addr      | 10.70.82.61
client_hostname  |
client_port      | 54546
backend_start    | 2018-08-17 11:32:08.9081+03
backend_xmin     |
state            | streaming
sent_location    | 1E8/D6DAFF48
write_location   | 1E8/D6DAFF48
flush_location   | 1E8/D6DAFF48
replay_location  | 1E8/D6DAFF48
sync_priority    | 0
sync_state       | async
```

Cascading Postgresql Replication

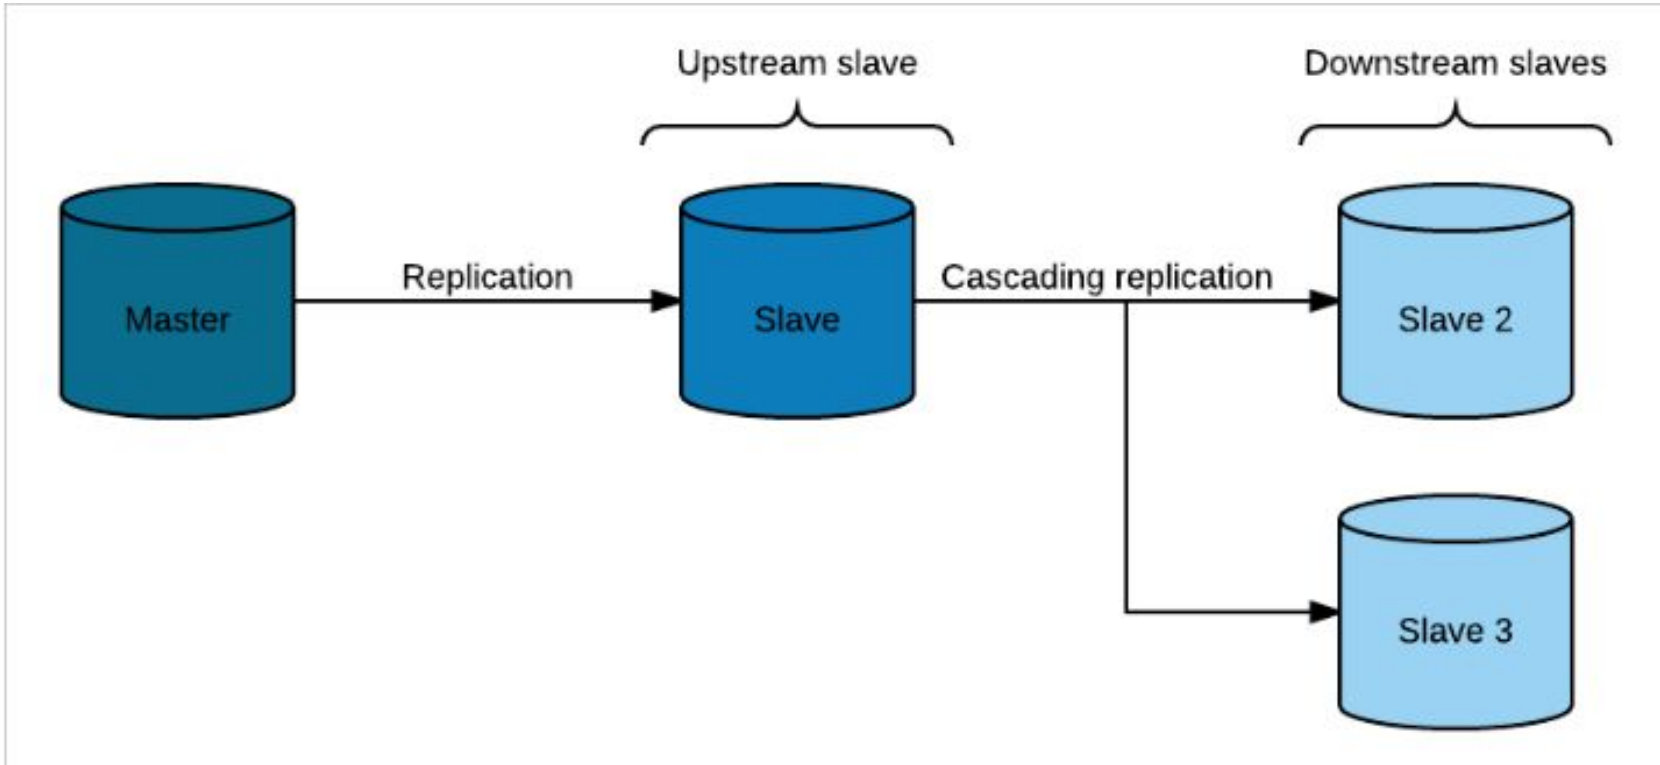

Topoloji

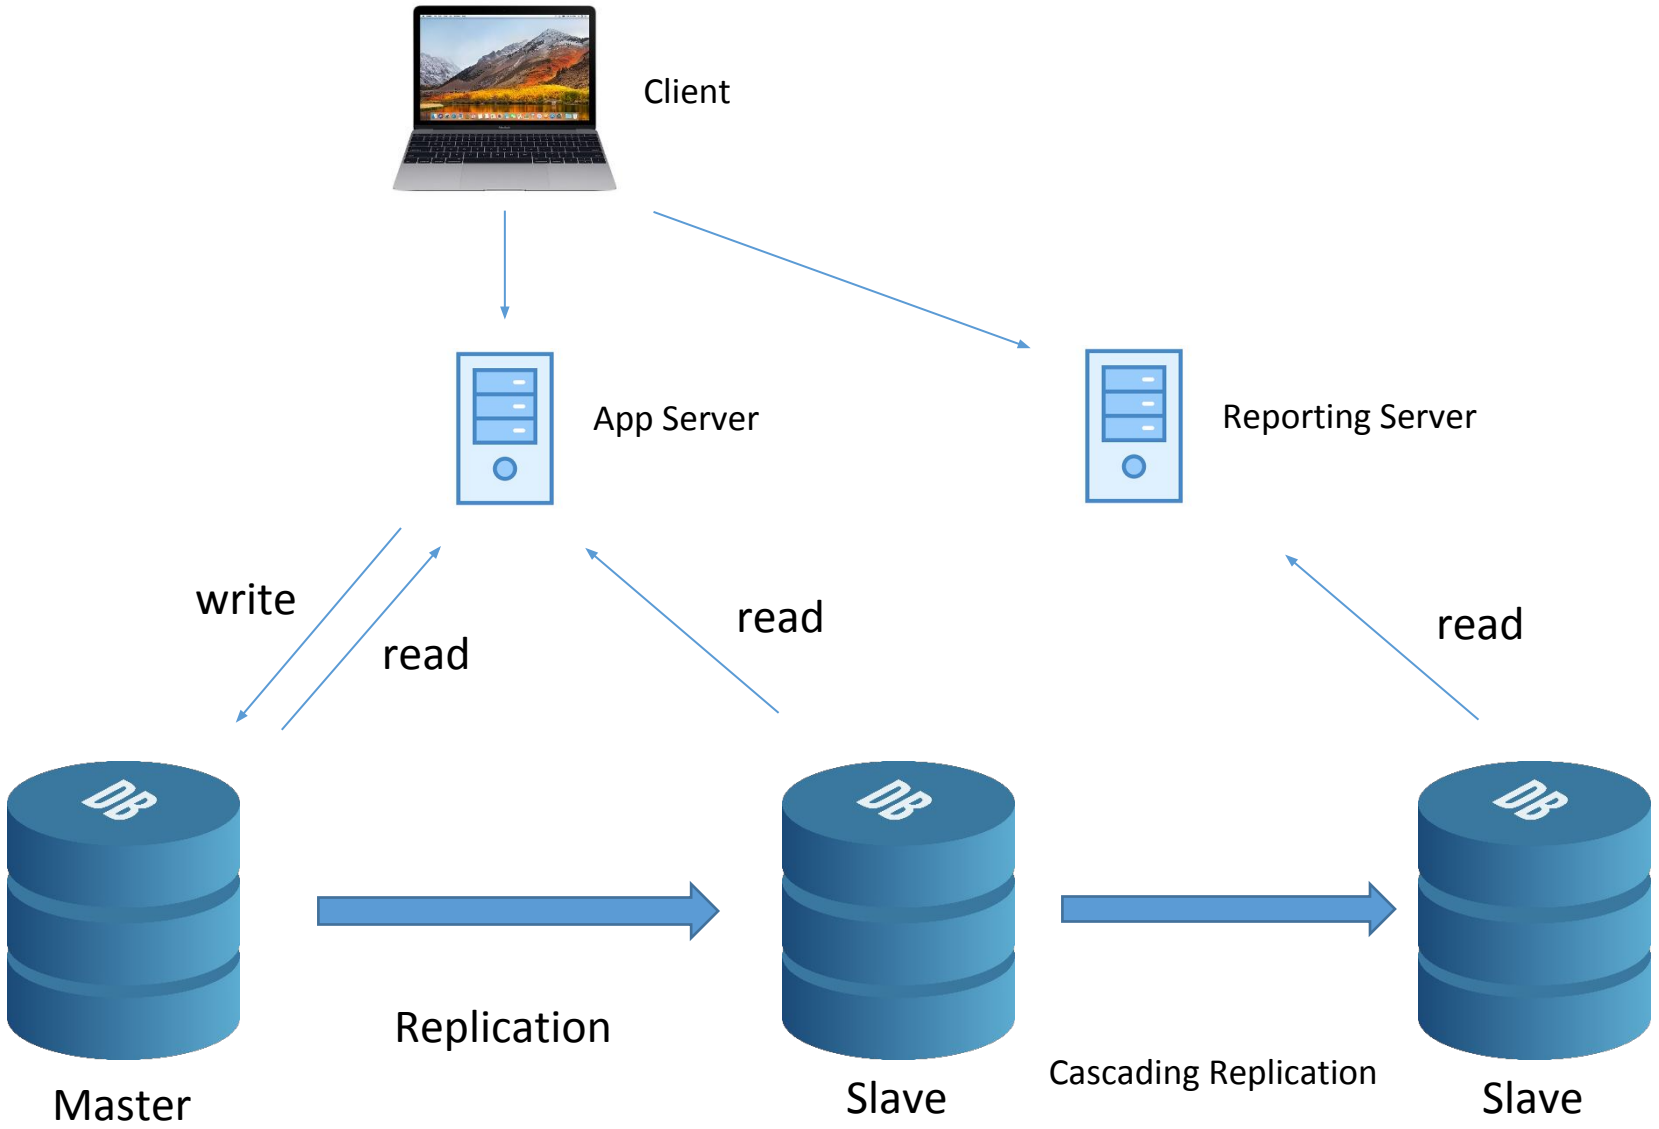
Client
App Server
Reporting Server
write
read
read
read
DB
DB
DB
Replication
Cascading Replication
Master
Slave
Slave

max_standy_archive_delay for standby

```
max_standby_archive_delay = 180s        # max delay before canceling queries
                                        # when reading WAL from archive;
                                        # -1 allows indefinite delay
max_standby_streaming_delay = 180s      # max delay before canceling queries
                                        # when reading streaming WAL;
                                        # -1 allows indefinite delay
```

Terminal

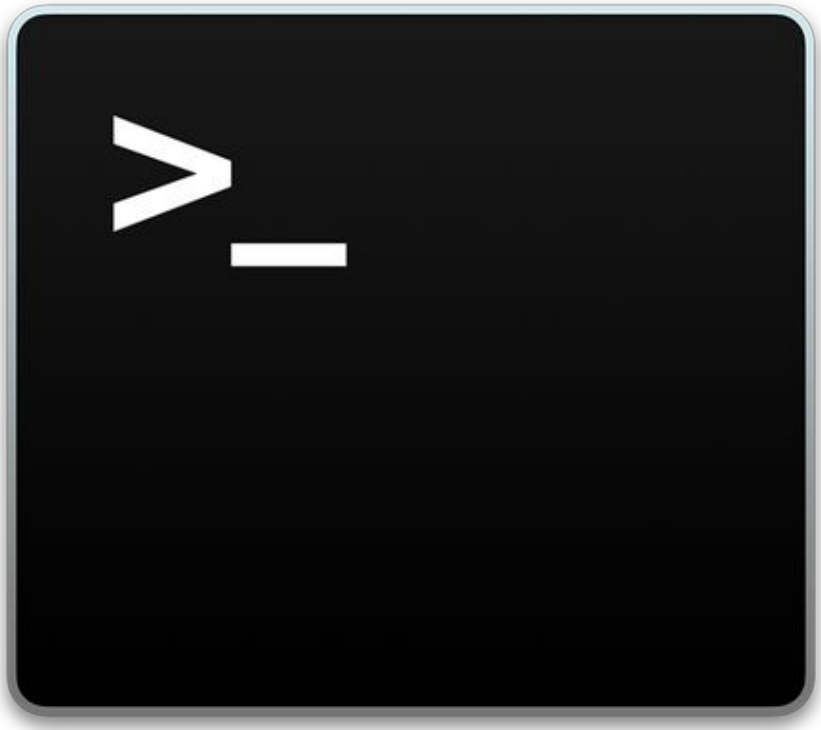

Logical replication with PostgreSQL 10

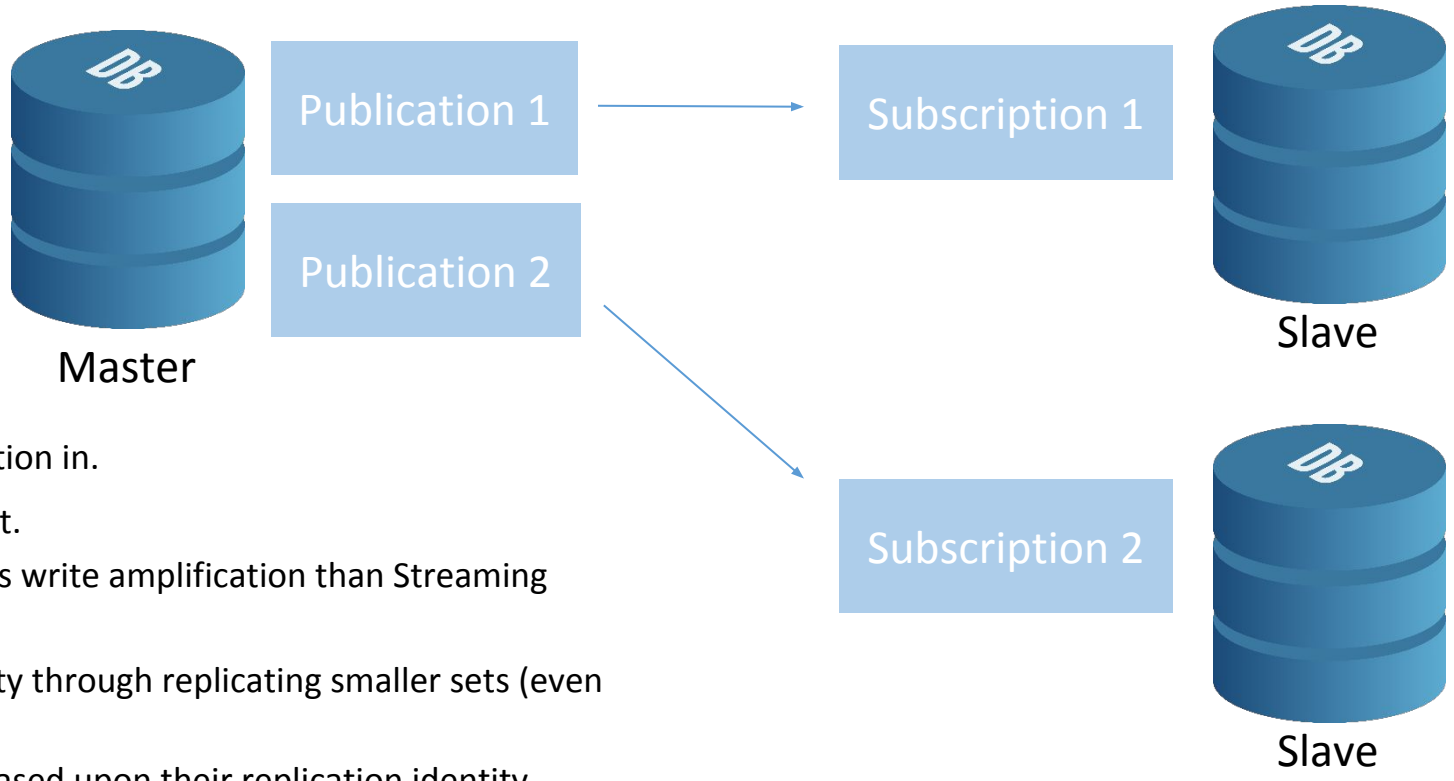

1- Block level replication out, row level replication in.

2- Logical Replication has cross-version support.

3- When compared, Logical Replication has less write amplification than Streaming Replication

4- Logical Replication provides storage flexibility through replicating smaller sets (even partitioned tables)

5- Logical Replication replicates data objects based upon their replication identity (generally a
primary key or unique index)

Limitations in 10.0

- does not replicate schema/DDL
- does not replicate sequences
- does not replicate TRUNCATE
- does not replicate Large Objects
- Tables must have primary key or unique key

Quorum Commit for Sync Replication

- 10.0: Quorum Commit

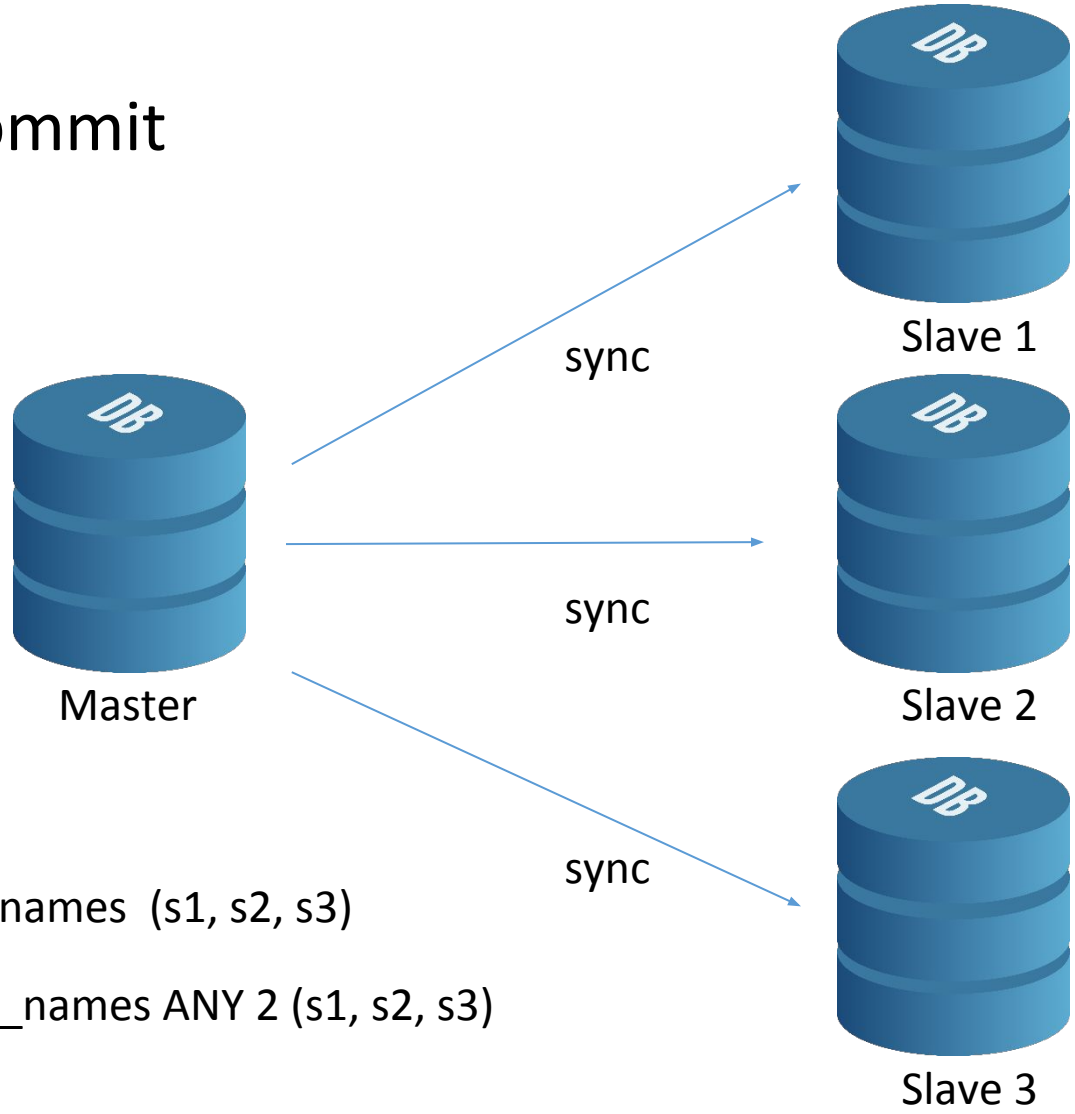

9.6→synchronous_standby_names (s1, s2, s3)

10.0→synchronous_standby_names ANY 2 (s1, s2, s3)


Kahoot!

Thank you

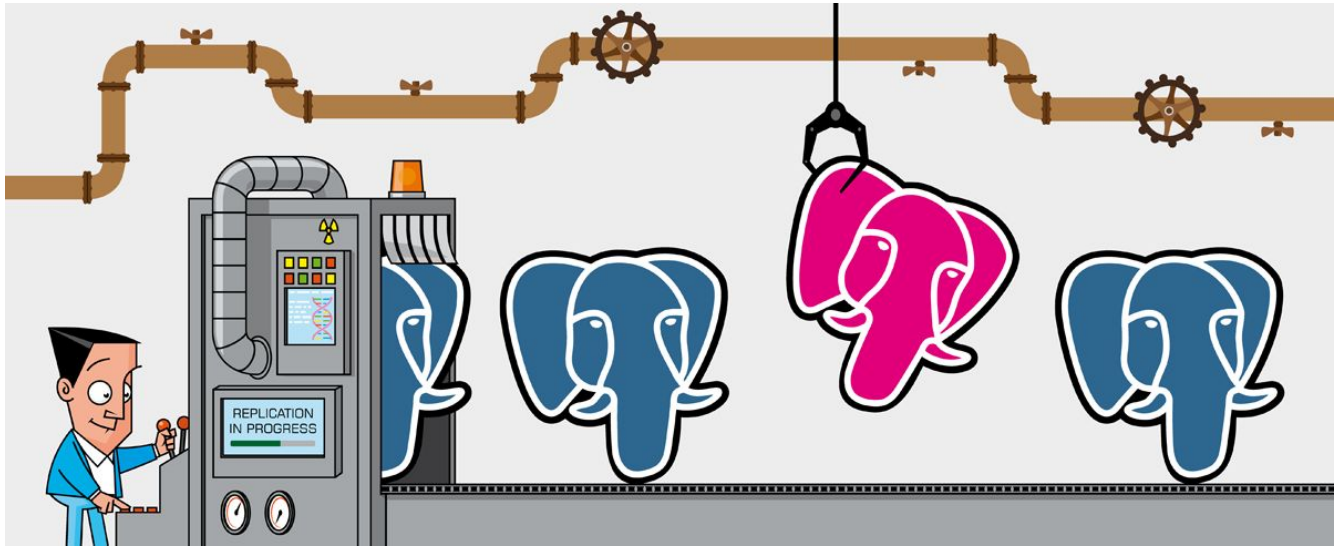

Fırat Güleç