

F O S D E M 2 0 2 1 , F E B 6 - 7 , O N L I N E

PostgreSQL Architectures in Production

Dimitri Fontaine
PostgreSQL Major Contributor

The Art of PostgreSQL

A decorative graphic consisting of several interlocking circles of varying sizes, arranged in a symmetrical, cross-like pattern. The circles are white and stand out against the dark purple background.

Turn Thousands of Lines
of Code into Simple Queries

A BOOK ABOUT POSTGRESQL BY DIMITRI FONTAINE

PostgreSQL for developers

Dimitri Fontaine
PostgreSQL Major Contributor

SECOND EDITION

The Art of PostgreSQL

Turn Thousands of Lines
of Code into Simple Queries

Special Discount

Use Code for 30%

“FOSDEM2021”

by Dimitri Fontaine


POSTGRES BUILD 2020, DEC 8-9, ONLINE

PostgreSQL Architectures in Production

Dimitri Fontaine
PostgreSQL Major Contributor



Postgres Architectures in Production



PostgreSQL

*The World's Most Advanced Open Source
Relational Database*




Architectures

Production Components and Dependencies



in Production
This is happening live!

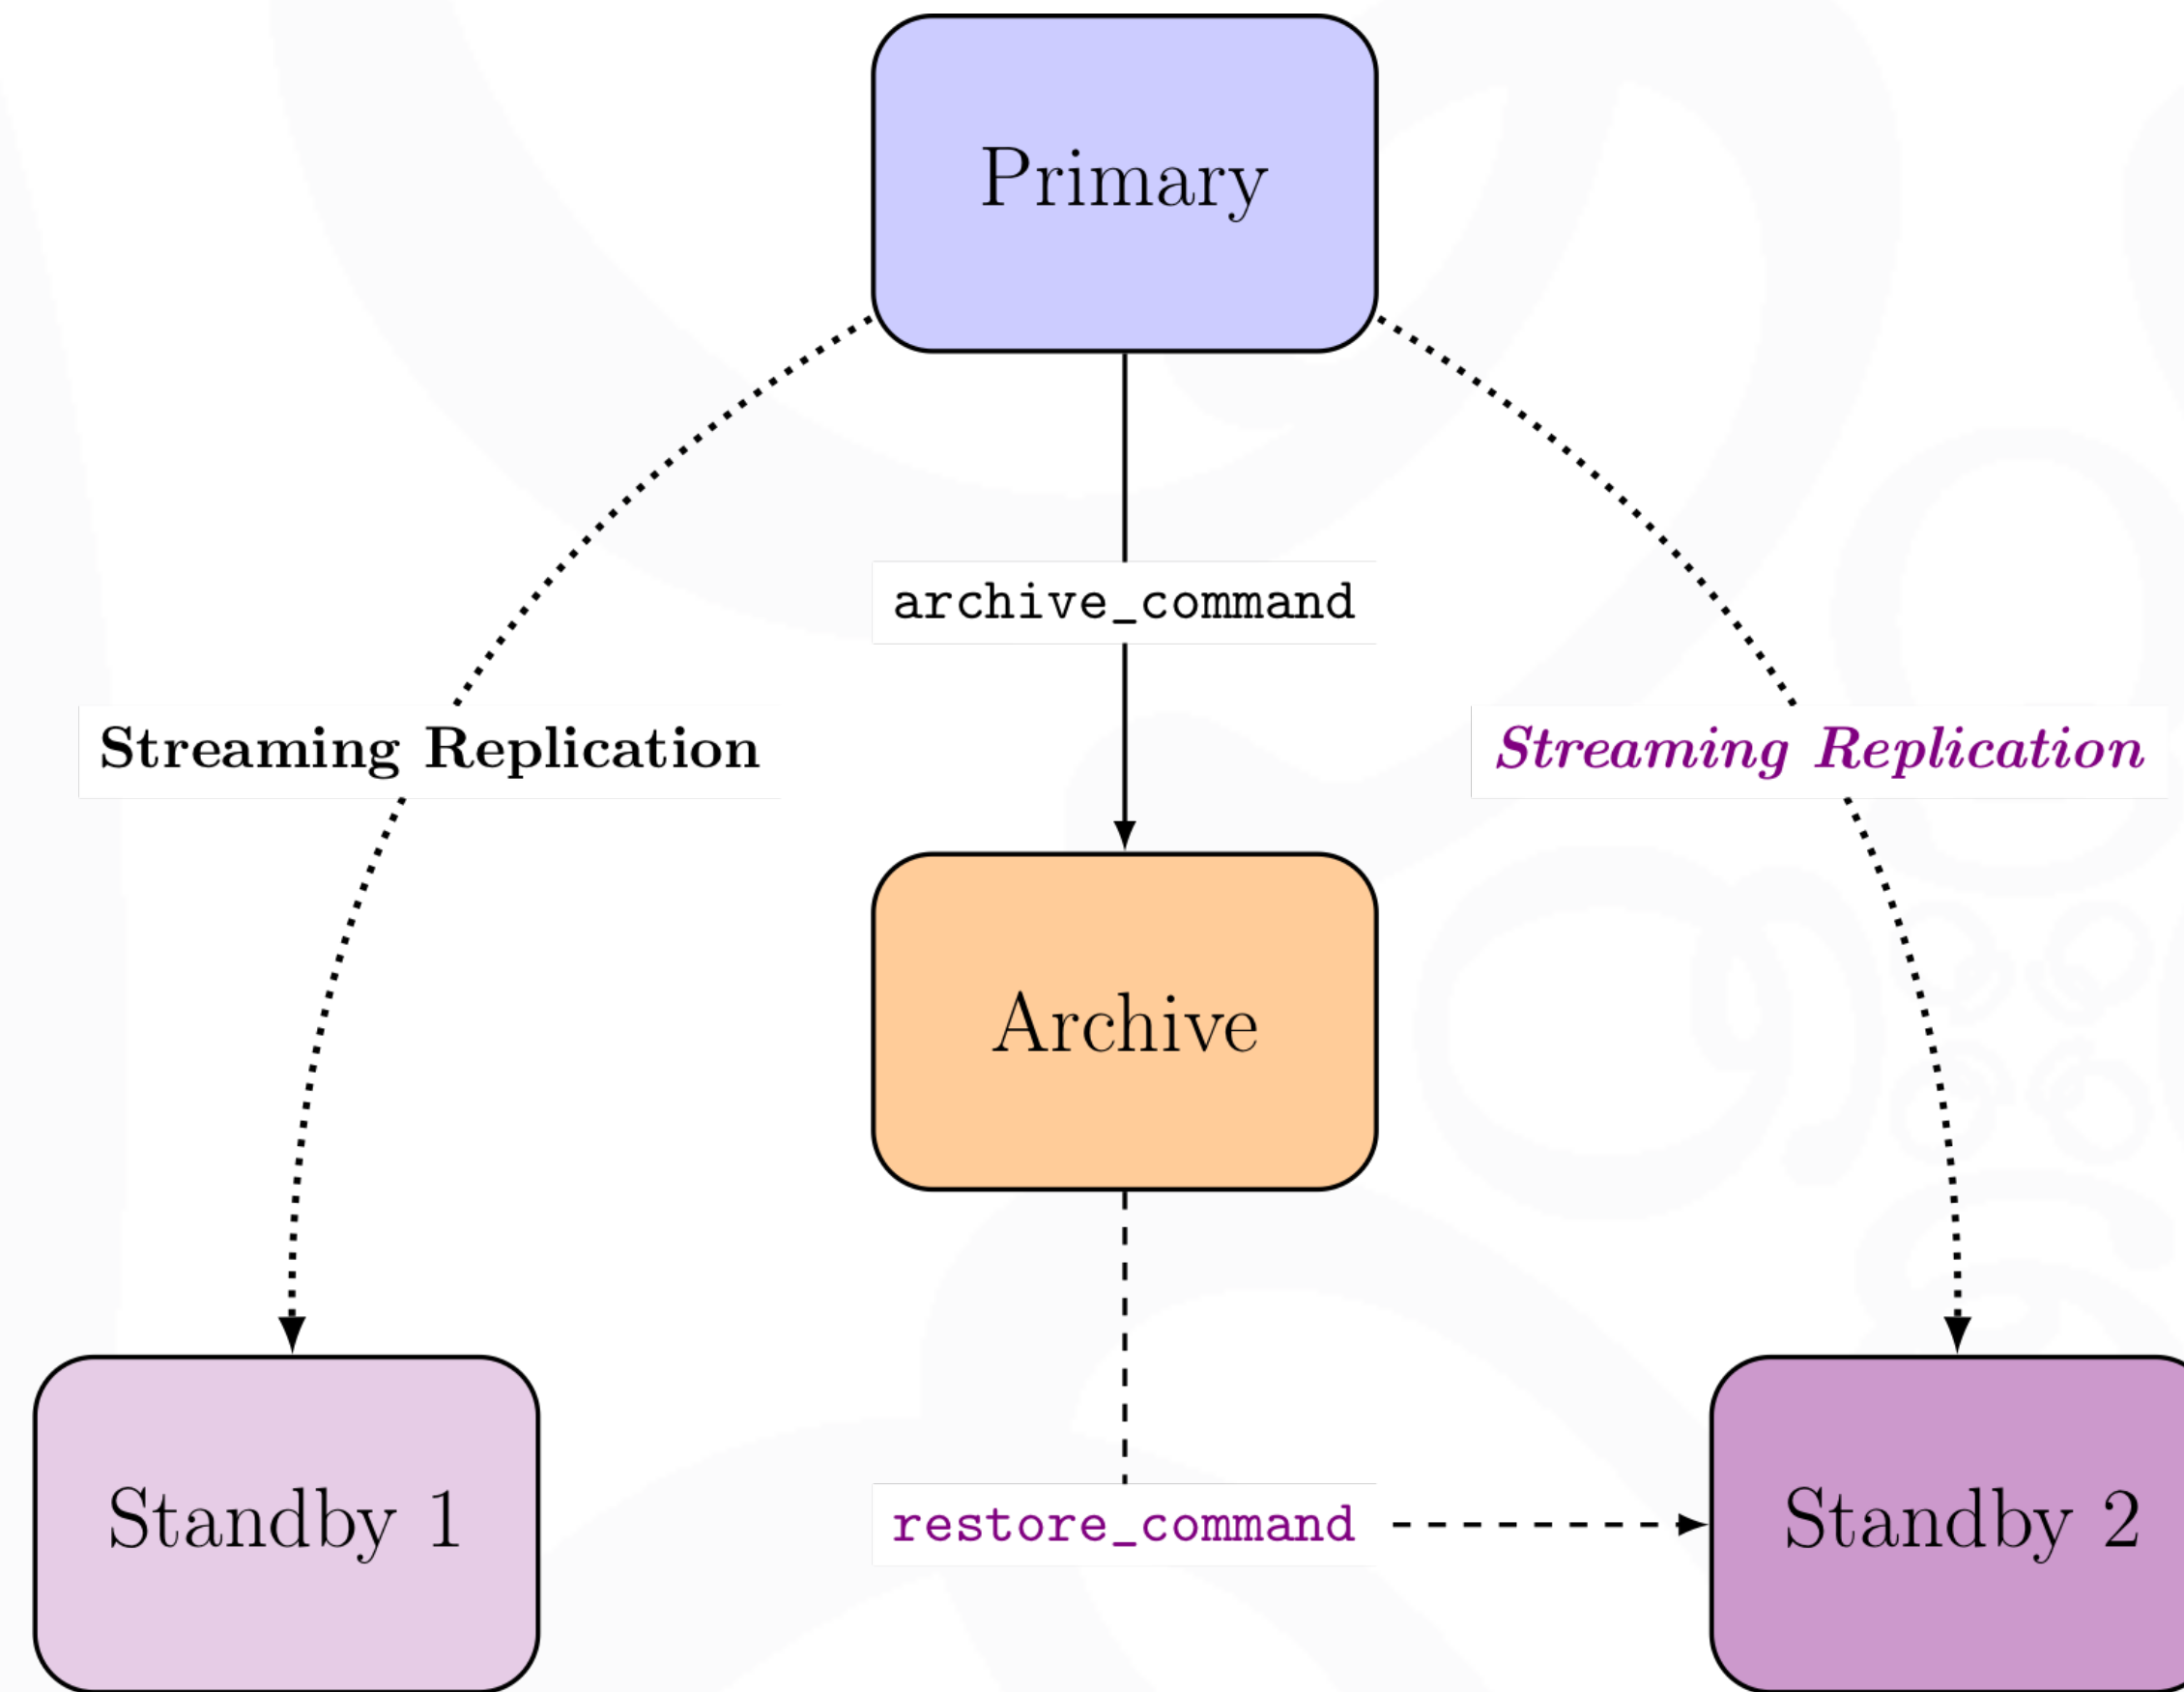
Availability: Service & Data



High Availability

- Maintain **service**
 - Failover
 - Automated?
- Maintain **data**
 - Disaster Recovery
 - Streaming Replication






Classic Postgres Architecture


Using Streaming Replication with Two Standby Nodes to enable Service and Data Availability

Postgres Failover




Included in Postgres

- *Streaming Replication*
 - Sync or Async
 - Per transaction setting
 - `ALTER ROLE vip`
`SET synchronous_commit`
`TO 'remote_apply';`
 - `ALTER ROLE ... SET ...`
 - `ALTER DATABASE ... SET ...`




Included in Postgres

- *Online Streaming changes*
 - *Replication Slots*
 - `pg_basebackup`
 - `pg_rewind`
 - Fast forward
 - Cascading Replication
 - Online Standby Promotion




Included in Postgres

- *Disaster Recovery basics: PITR*
- *Setup a “standby” with a target*
- *Archiving provided separately*



Included in Postgres


- *Client Side HA*
 - Multi-Host Connection String
 - Client still has to **reconnect** when connection is lost, of course



`postgres://hostA:5502,hostB:5503,hostC:5501/dbname?target_session_attrs=read-write&sslmode=require`


Not included in Postgres

- **Archiving**
 - Base Backups
 - WAL Archiving
 - Retention setup for Disaster Recovery




Not included in Postgres

- Online “**Membership**” changes
- Idea that **roles are DYNAMIC**
 - Ensure retention during failover
 - For WALs
 - And for base backups
 - Automated node role changes
- Restart to reconnect to new primary




Not included in Postgres

- *Configuration Management*
 - postgresql.conf
 - pg_hba.conf
 - *Some setting change require a restart*
- *Extension Management*
- *Postgres Upgrades*
 - Minor upgrade (scheduled restarts)
 - Major upgrades

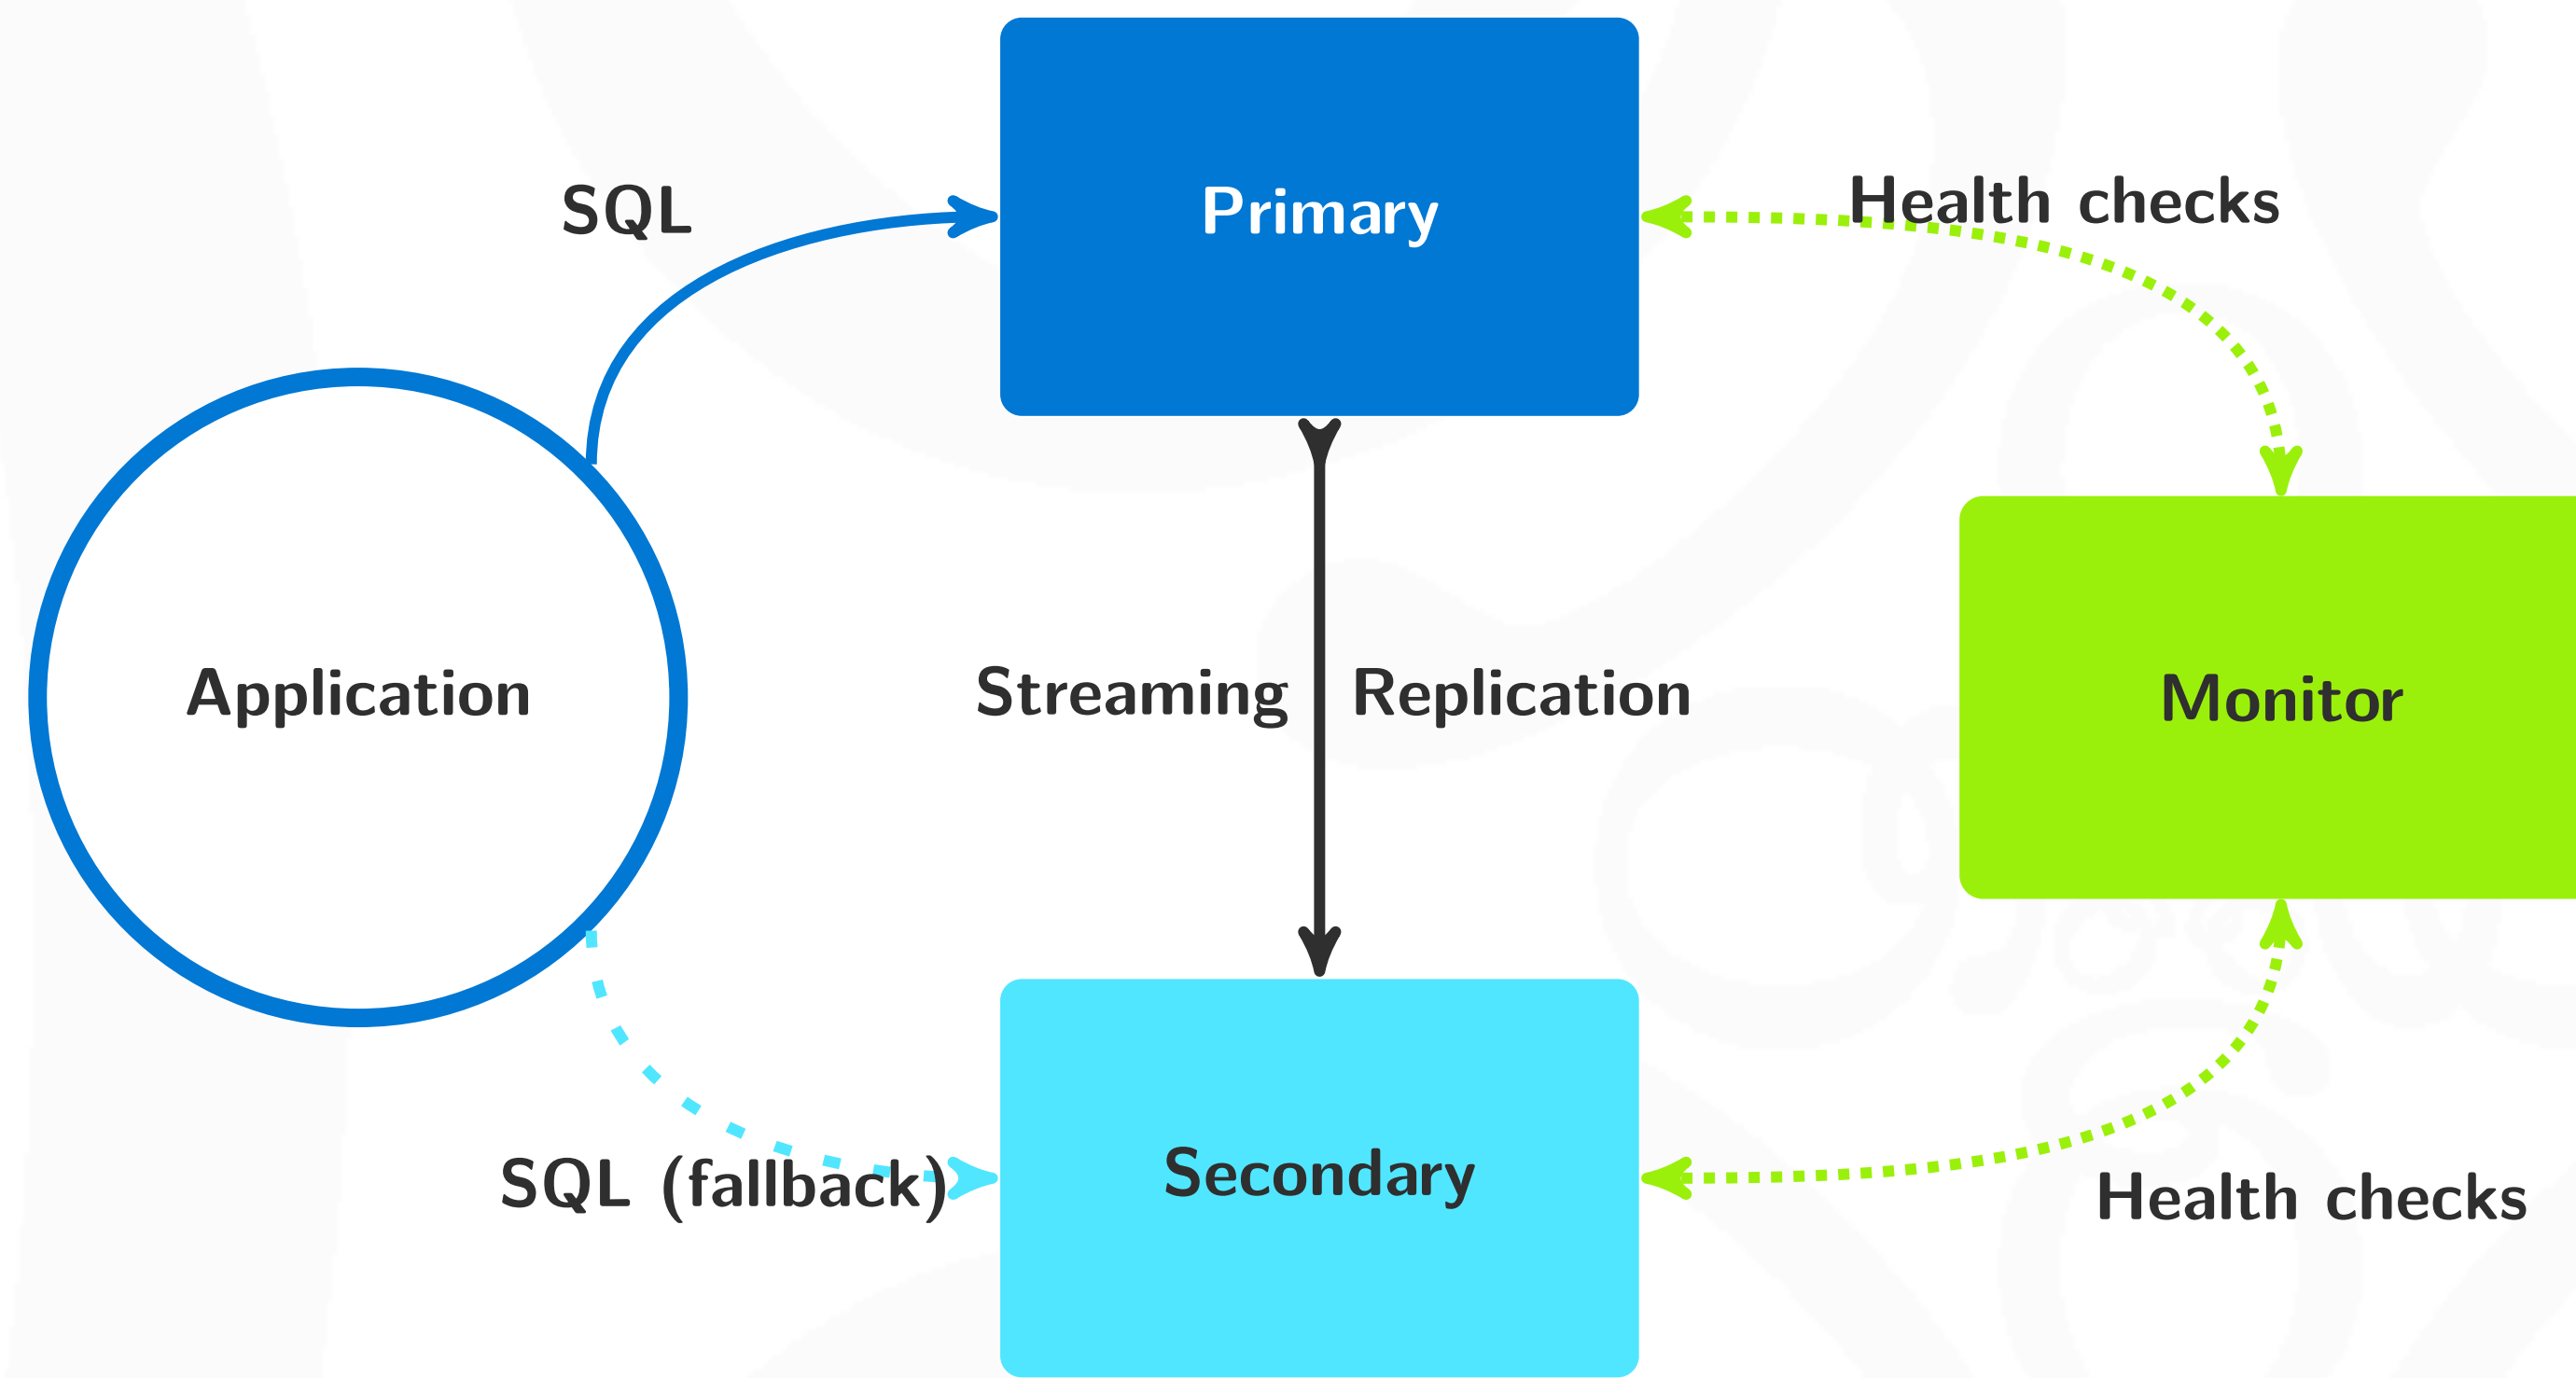


Automated Failover





pg_auto_failover



Single Standby

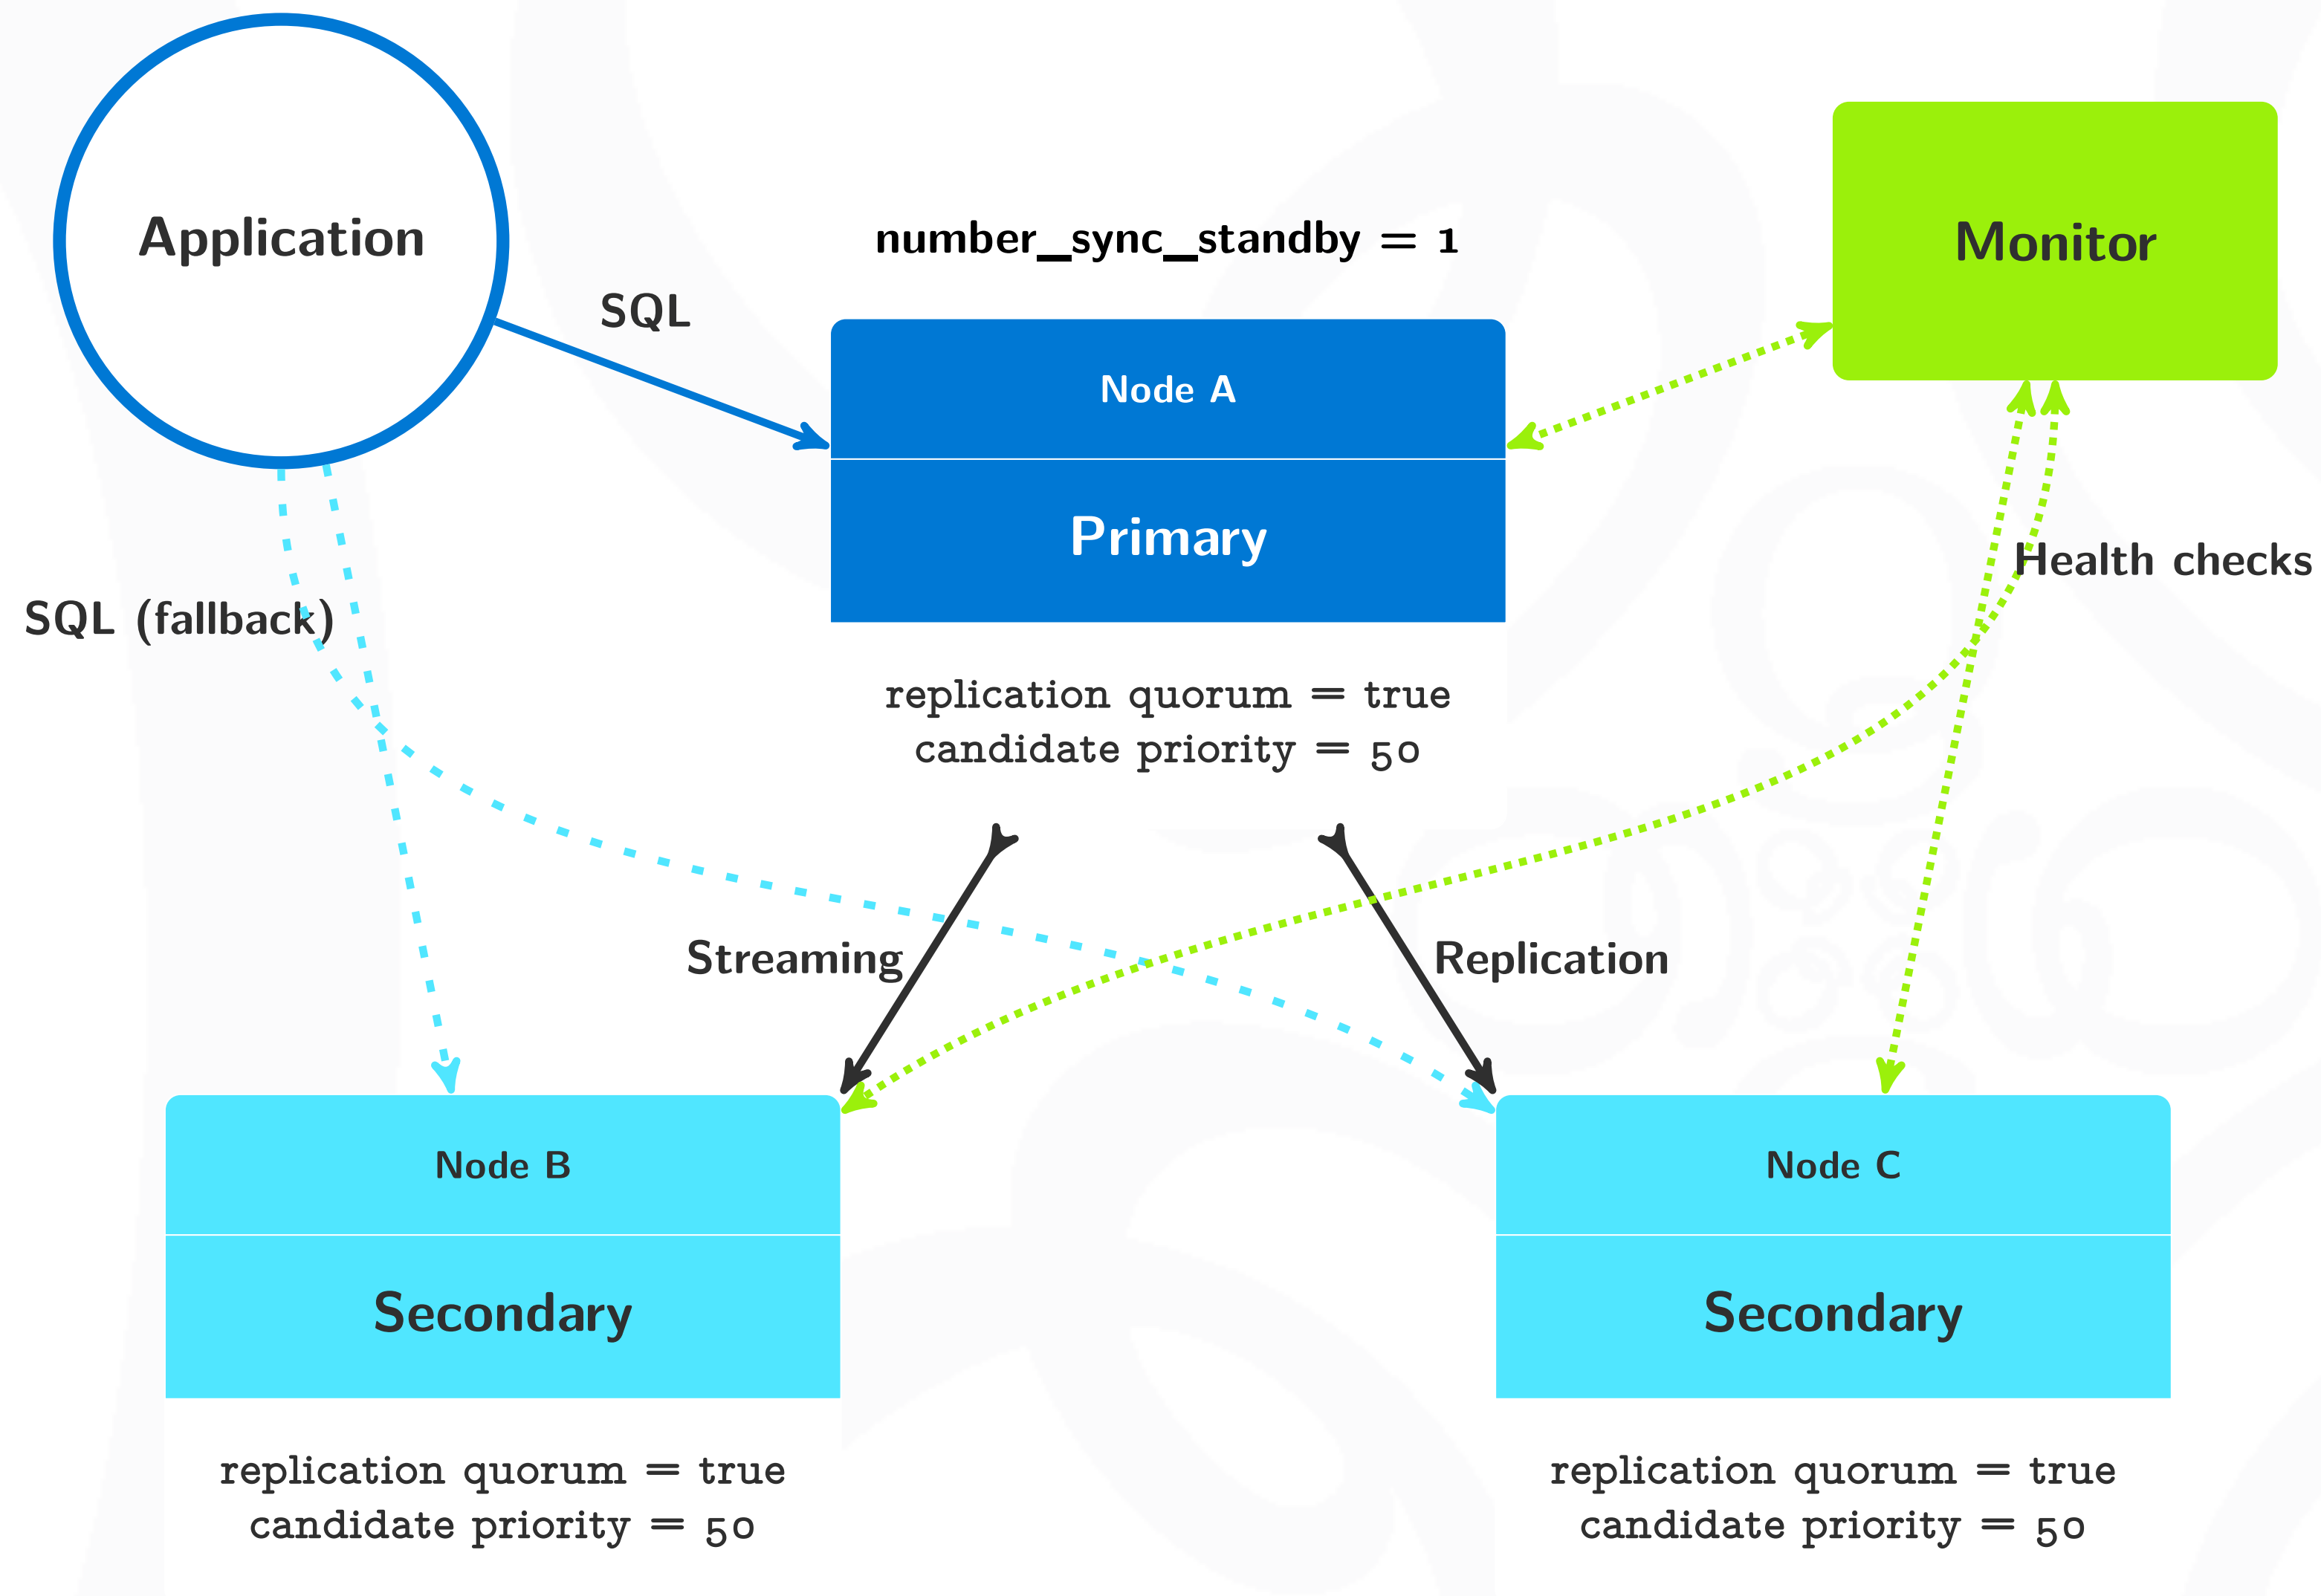
Only provides failover, no HA

Single Standby

```
$ pg_autoctl create monitor -auth trust -self-signed-ssl -run
```

```
$ pg_autoctl create postgres \
    -auth trust \
    -self-signed-ssl. \
    -monitor postgres://host/pg_auto_failover \
    -run
```

```
$ pg_autoctl create postgres \
    -auth trust \
    -self-signed-ssl. \
    -monitor postgres://host/pg_auto_failover \
    -run
```

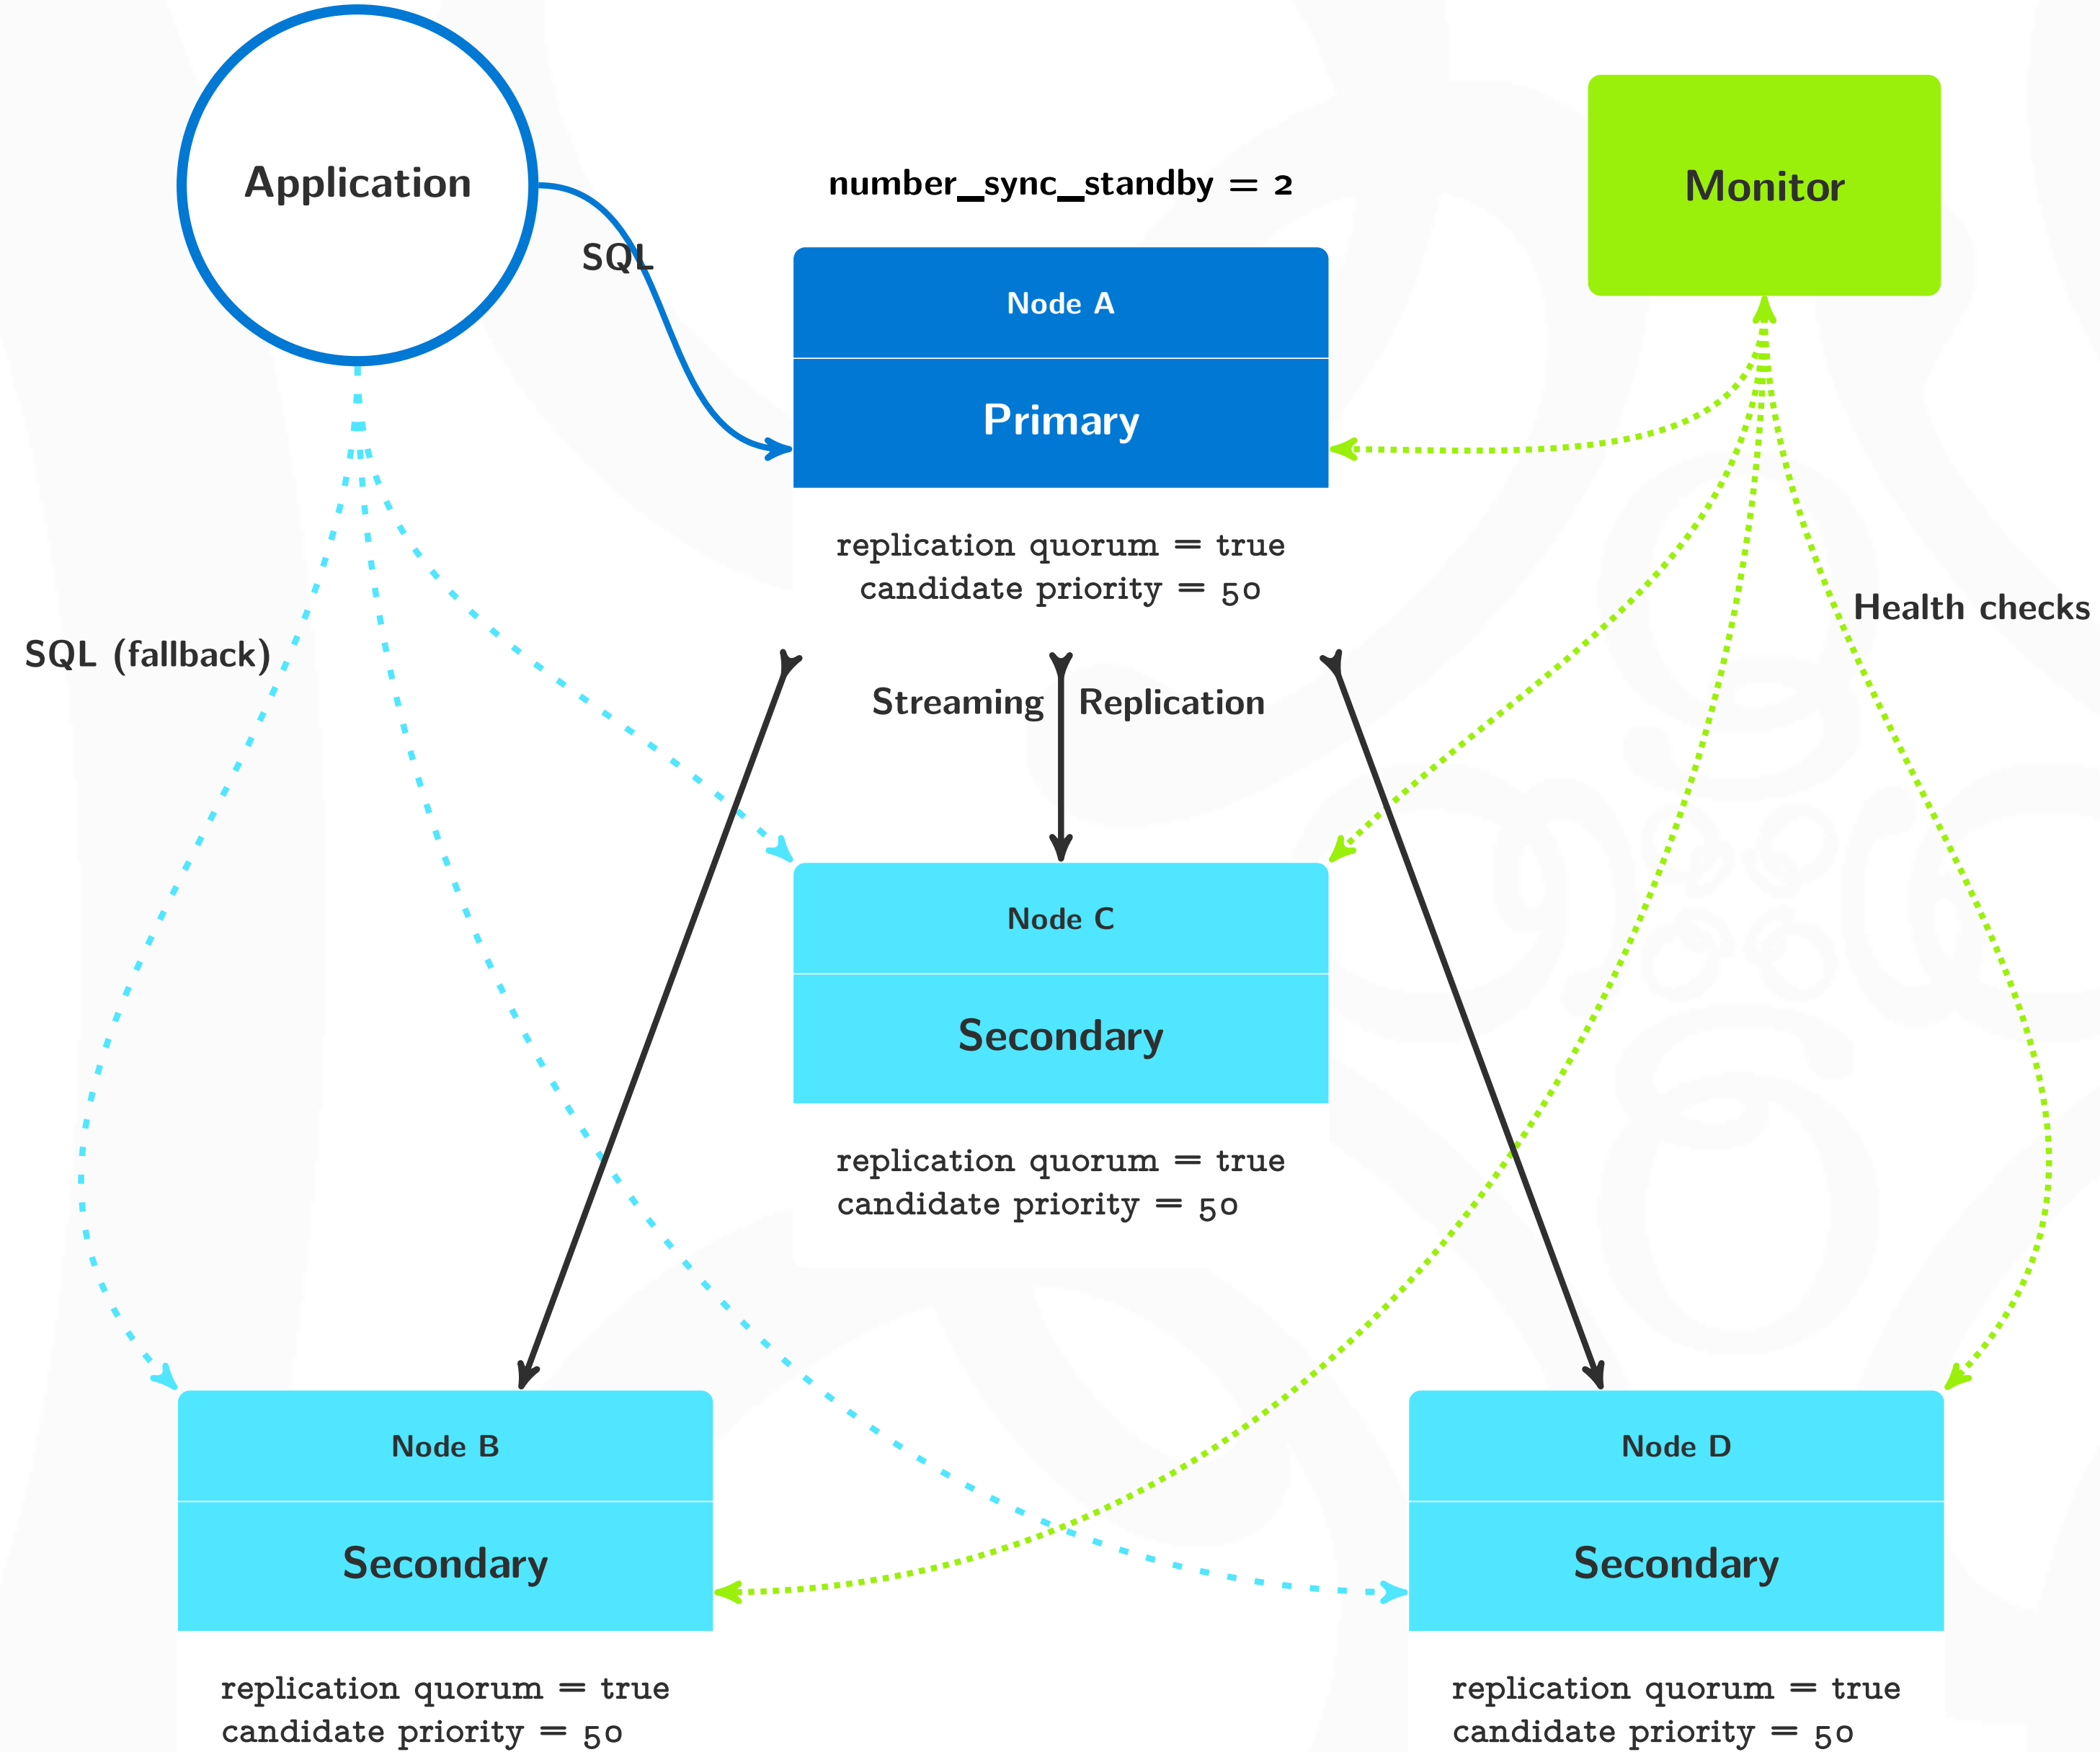



Two Standby Nodes

Allows to implement HA

Two Standby Nodes

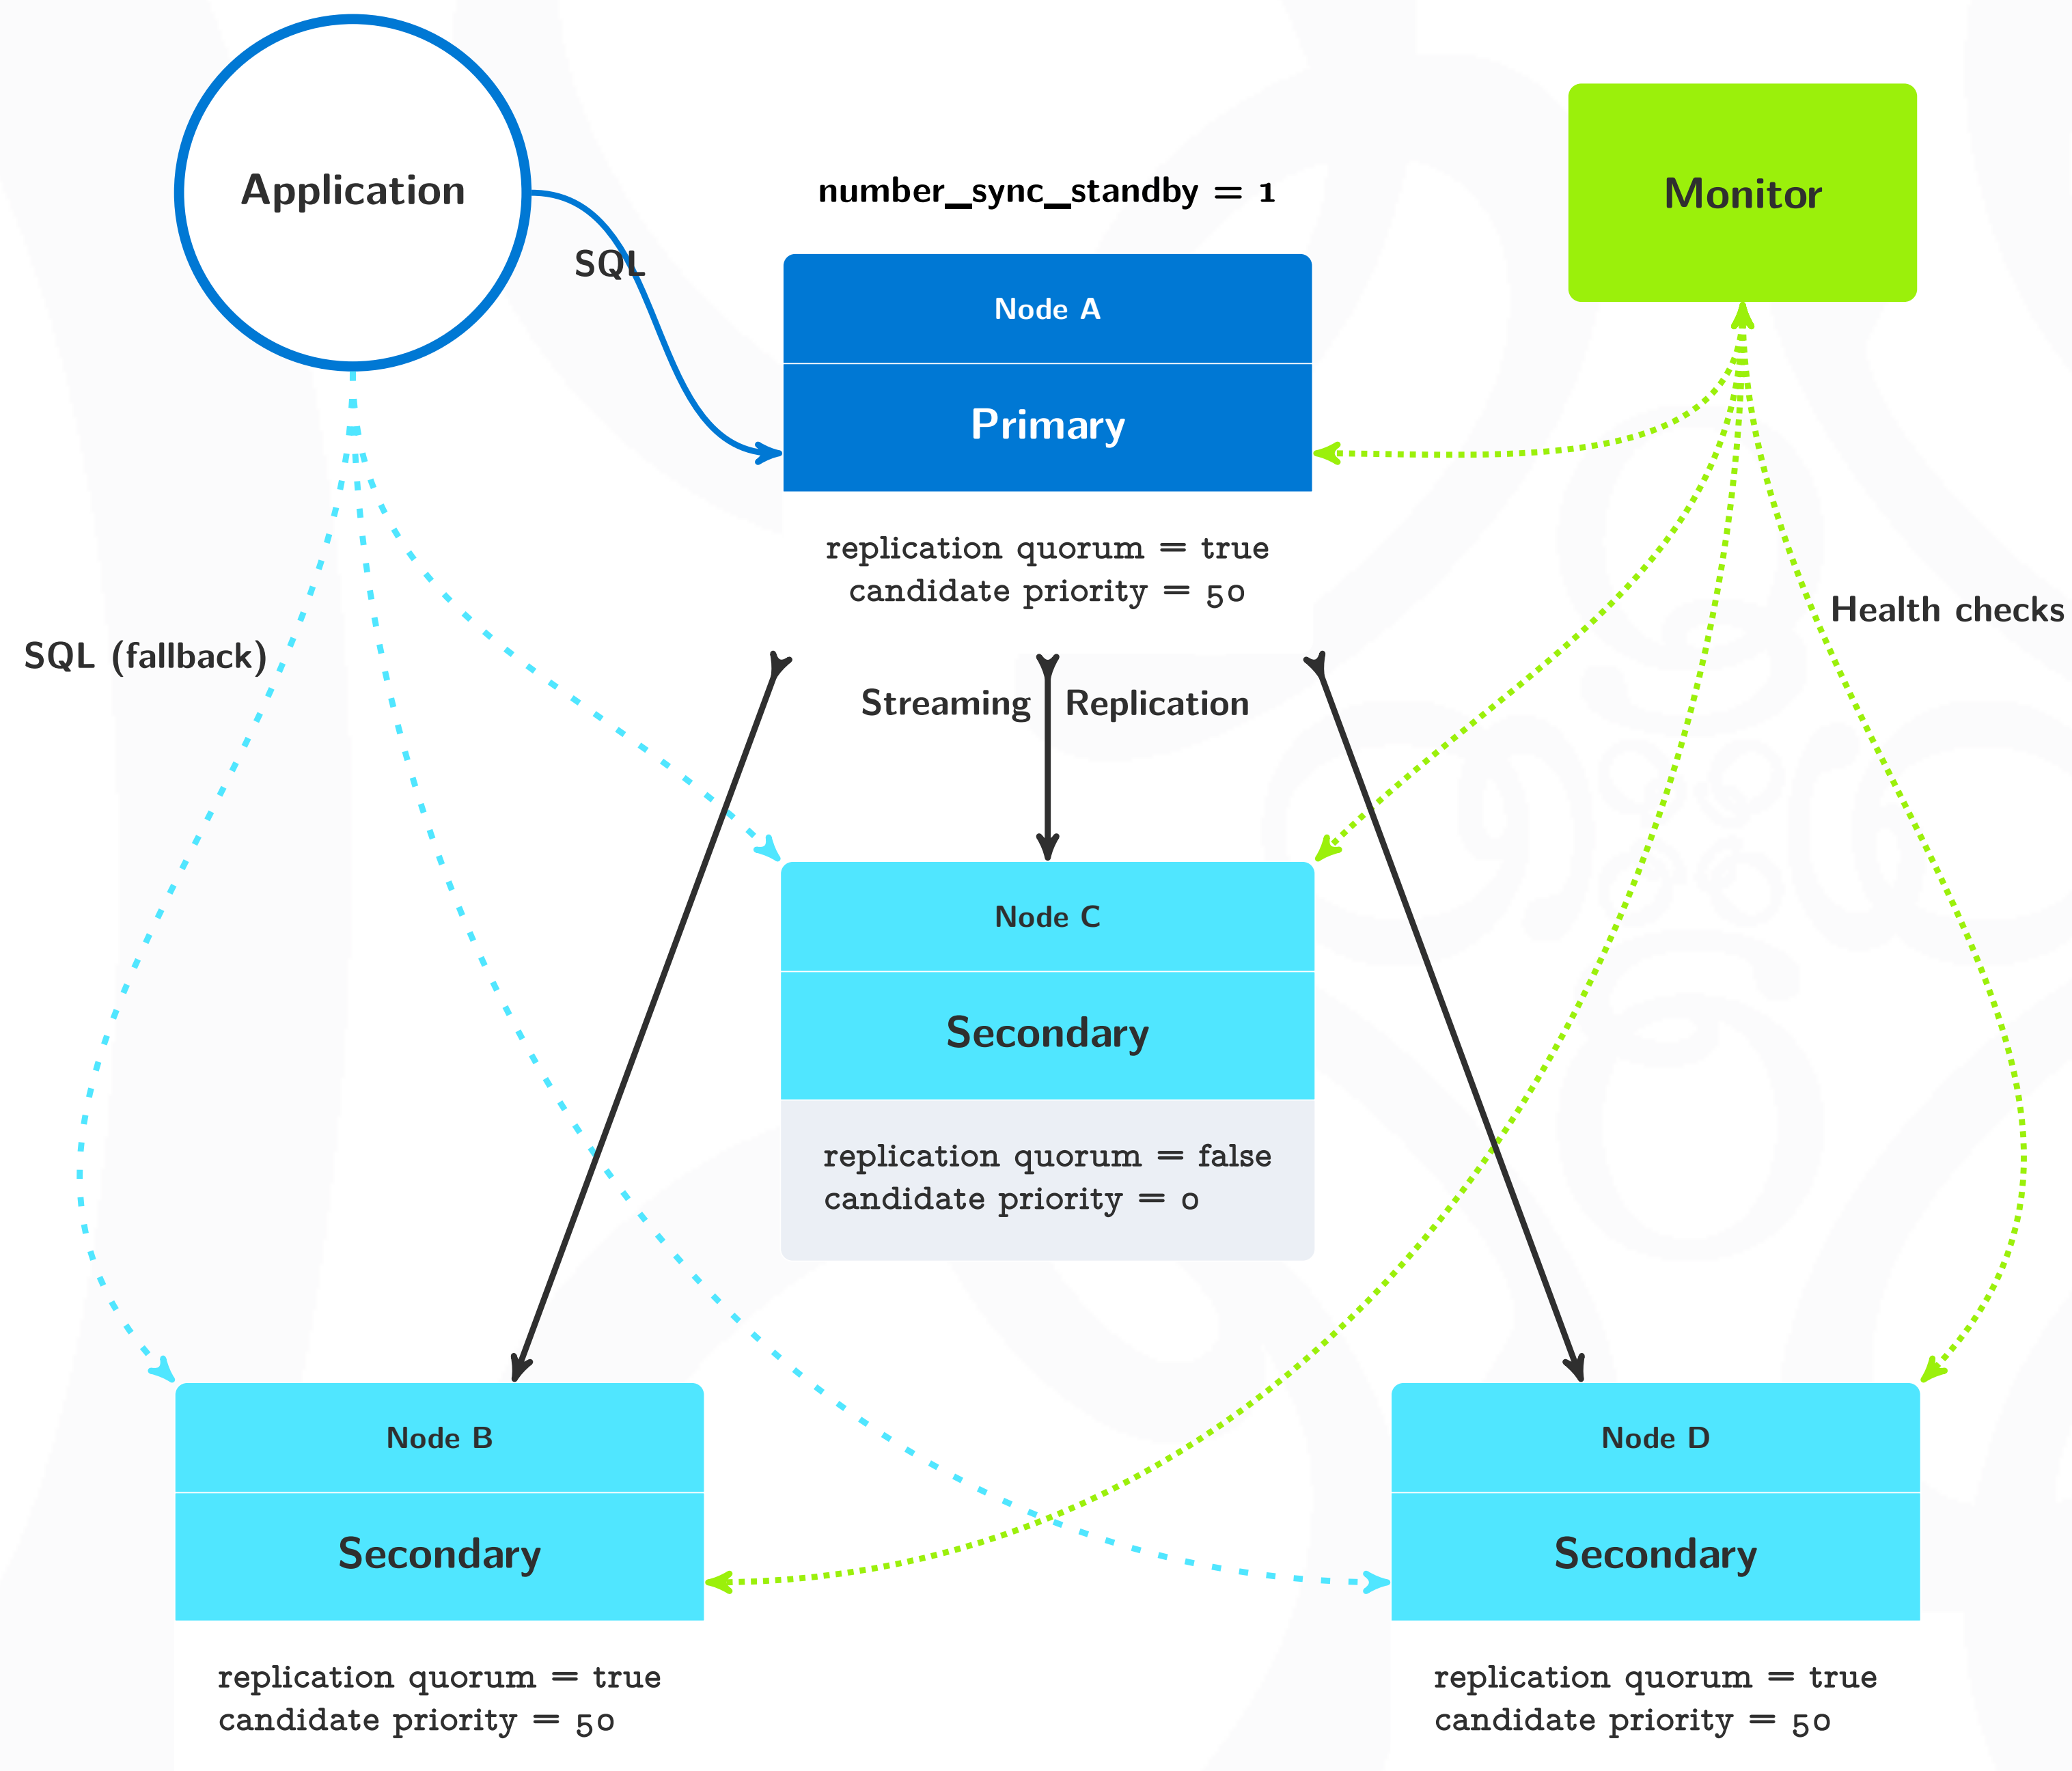
```
$ pg_autoctl create postgres \
    -auth trust \
    -self-signed-ssl. \
    -monitor postgres://host/pg_auto_failover \
    -run
```



Three Standby Nodes

Three Standby Nodes

```
$ pg_autoctl create postgres \
    -auth trust \
    -self-signed-ssl. \
    -monitor postgres://host/pg_auto_failover \
    -run
```

Three Standby Nodes, one async

Three Standby Nodes, One Async

```
$ pg_autoctl create postgres \
    -auth trust \
    -self-signed-ssl. \
    -monitor postgres://host/pg_auto_failover \
    -replication-quorum false \
    -candidate-priority 0 \
    -run
```

```
$ pg_autoctl set node candidate-priority 0
$ pg_autoctl set node replication-quorum false
```

Included in pg_auto_failover

- *Streaming Replication*
 - *Replication slots*
 - *Maintained on standby nodes*
- *Easy setup including SSL certificates*
- *Dynamic settings, online changes*

```
$ pg_autoctl create postgres
```



Included in pg_auto_failover

- Online “**Membership**” changes
- Idea that **roles are DYNAMIC**
 - Ensure retention during failover
 - For WALs
 - And for base backups
 - Automated HBA editing

```
$ pg_autoctl show uri  
$ pg_autoctl show state  
$ pg_autoctl get formation settings
```



Included in pg_auto_failover

- *Online membership changes*
- *Online formation setting changes*
- *Maintenance operations*

```
$ pg_autoctl set node candidate-priority 0  
$ pg_autoctl set node replication-quorum false  
$ pg_autoctl set formation number-sync-standbys 2
```

```
$ pg_autoctl enable maintenance  
$ pg_autoctl disable maintenance
```



Included in pg_auto_failover

- *Network split detection / protection*
- *Active monitoring and role changes*
 - *Automated and manual Failover*
 - *Including manual Switchover*
 - *Including manual choice of target node*

```
$ pg_autoctl perform failover  
$ pg_autoctl perform switchover  
$ pg_autoctl perform promotion --name node_d
```



Not included in pg_auto_failover ***yet***

- *Archiving and Disaster Recovery*
- *Configuration Management*
 - postgresql.conf
 - pg_hba.conf
 - *Some setting change require a restart*
- *Extension Management*
- *Postgres Upgrades*
 - Minor upgrade (scheduled restarts)
 - Major upgrades



POSTGRES BUILD 2020, DEC 8-9, ONLINE

PostgreSQL Architectures in Production

Dimitri Fontaine
PostgreSQL Major Contributor

THE ART OF POSTGRES QL

Ask Me Two Questions!

Dimitri Fontaine
Citus Data, Microsoft

@tapoueh