2ndQuadrant
Professional PostgreSQL

Deploying BDR

Simon Riggs
CTO, 2ndQuadrant &
Major Developer, PostgreSQL

February 2015

What is BDR?

- BDR is Bi Directional Replication
- BDR is the newest replication technology for core PostgreSQL
- BDR is open source, using PostgreSQL Licence
- BDR is a code submission to the PostgreSQL project
- BDR is a working, production quality solution

Trigger Based Replication

2006

- It works!
- Selective replication
- Allows online upgrade

Trigger Based Replication

- It works!
- Selective replication
- Allows online upgrade

- Not Integrated in Core
- Complex design & code
- Difficult to Understand
- High Maintenance
- Slow, high latency
- High overhead
- Serialization problems

Physical Streaming Replication

2010

- Integrated in core
- Simple design & code
- Ease of Use
- Low Maintenance (DDL)
- Fast, low latency
- Low overhead
- No serialization problems

Physical Streaming Replication

- Integrated in core
- Simple design & code
- Ease of Use
- Low Maintenance (DDL)
- Fast, low latency
- Low overhead
- No serialization problems

- Read-only access
- Same schema
- No temp tables
- Same users/security
- Physical effects
- Whole database only
- No Multi-Master
- Complex code
- No online upgrade

Ideal Replication

- Integrated in core
- Simple design & code
- Ease of Use
- Low Maintenance (DDL)
- Fast, low latency
- Low overhead
- No serialization problems
- Read-write access
- Allows different schema
- Allows temp tables
- Allows different security
- Limited Physical effects
- Selective replication
- Allows Multi-Master
- Backwards compatible
- Online Upgrade

2ndQuadrant designed
Logical Replication in 2009

2ndQuadrant started building
BDR in 2011

….we went production in 2014

1. Basic pgbench test: the overhead of BDR is very low and is very close to Hot Standby.

The small number for Londiste and Burcardo on Standby are the result of setting wal_level=minimal.

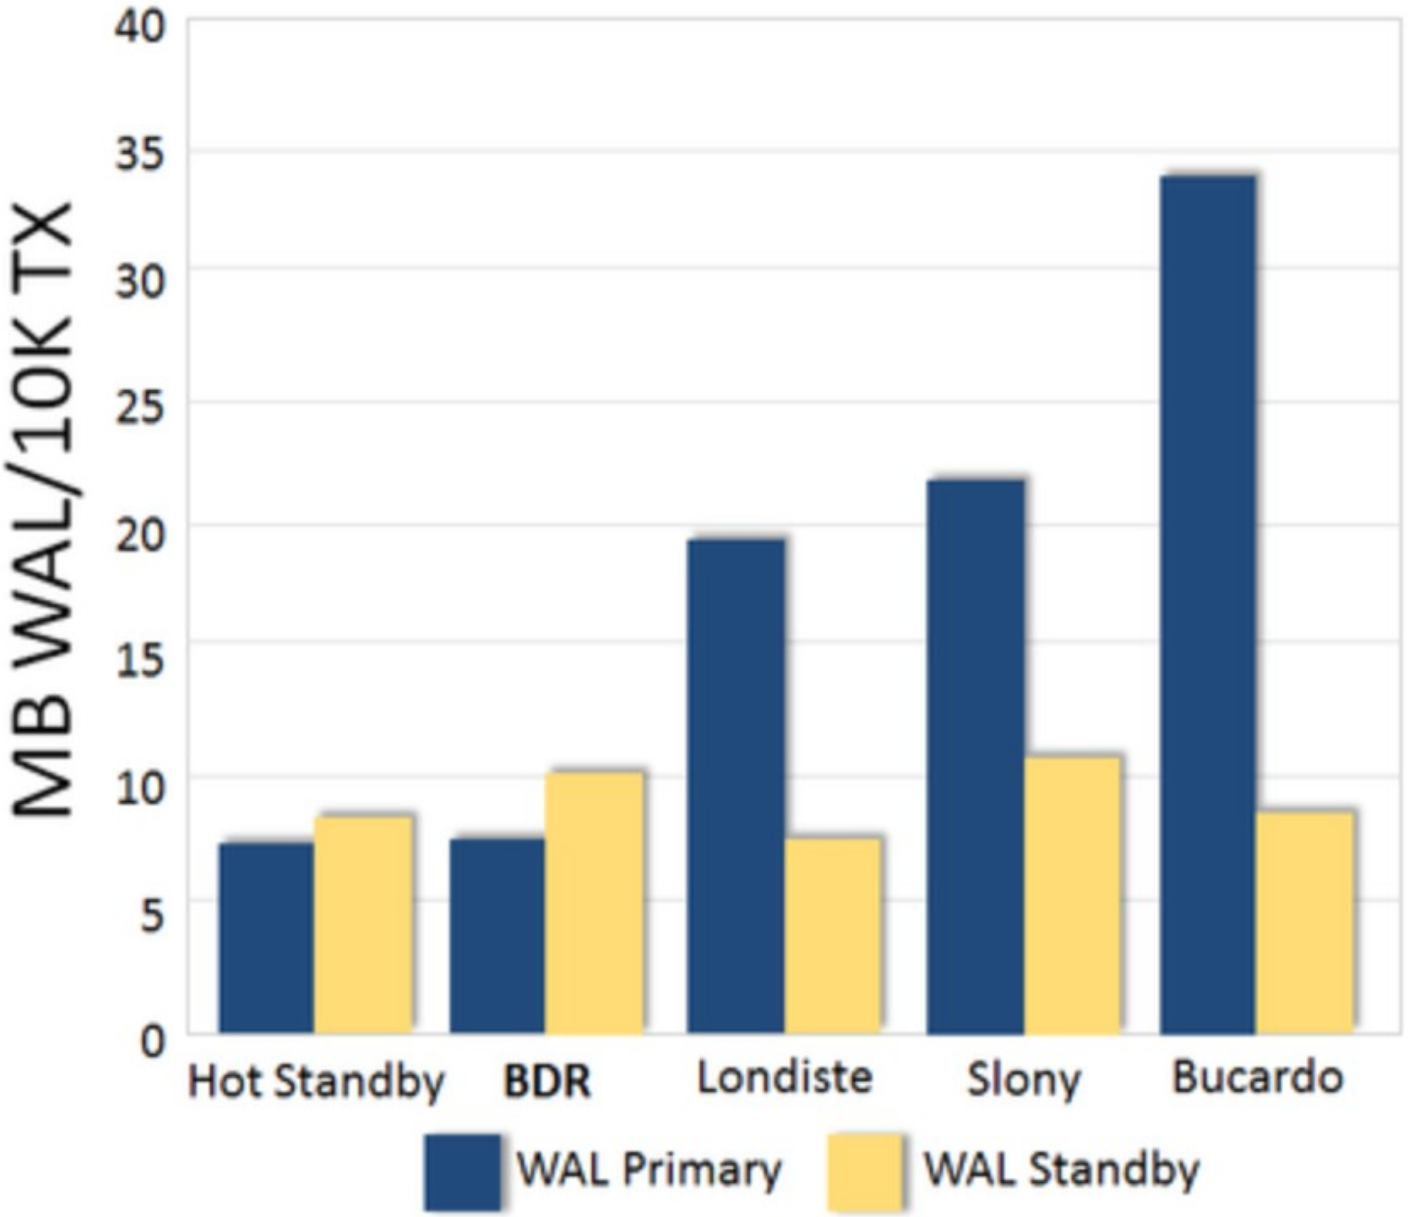

BDR

2. Transaction performance: again BDR's performance is very close to Hot Standby

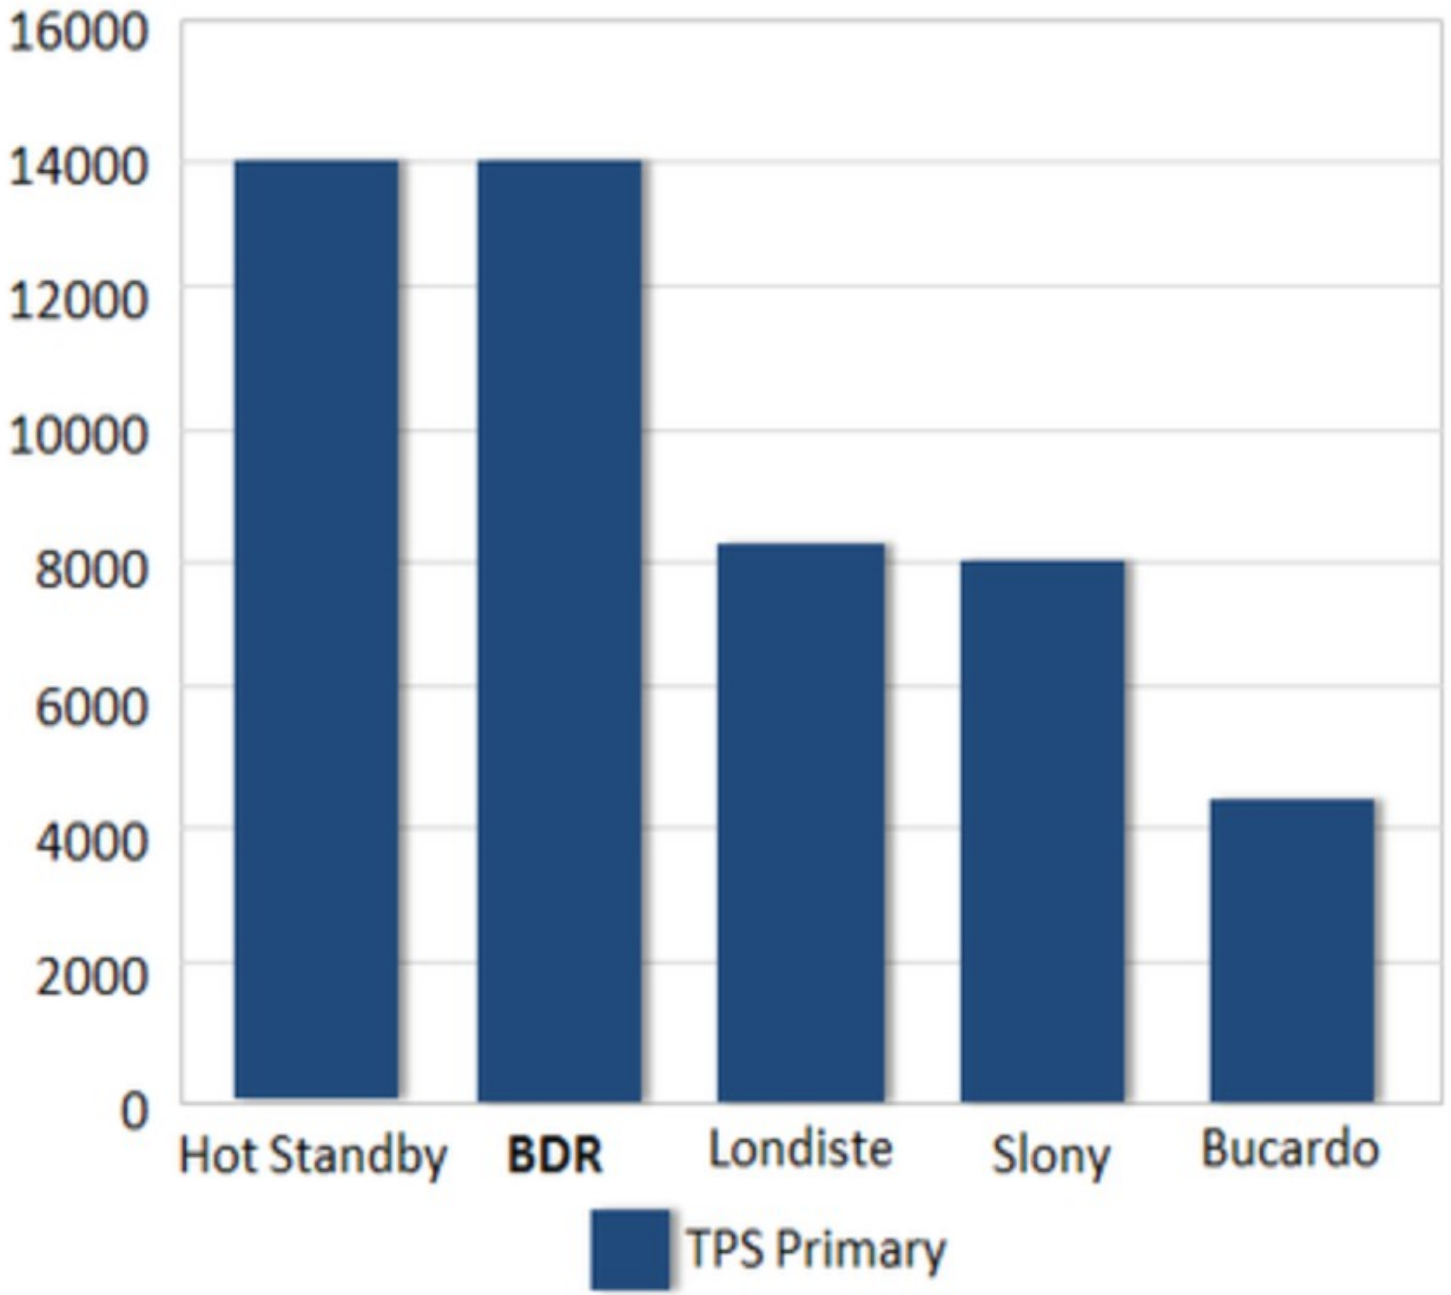

BDR

3. Latency: BDR avoids trigger processing so achieves a low latency per transaction.

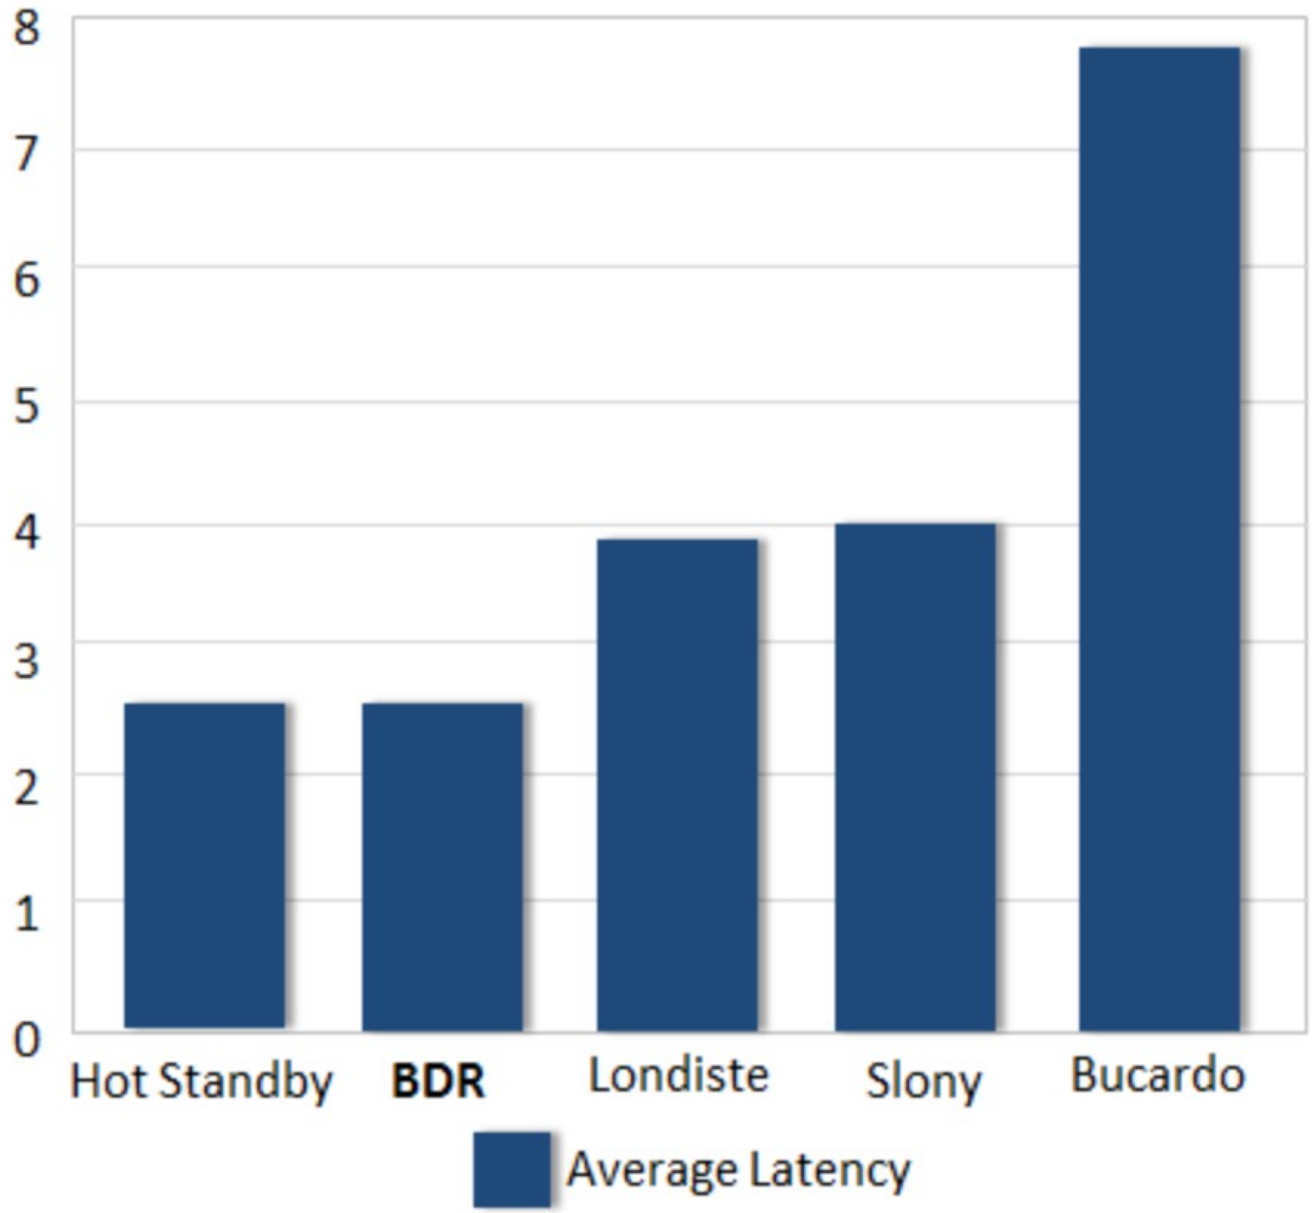

BDR

Latency measured in milliseconds

4. Custom workload throughput with latency < 2 seconds.

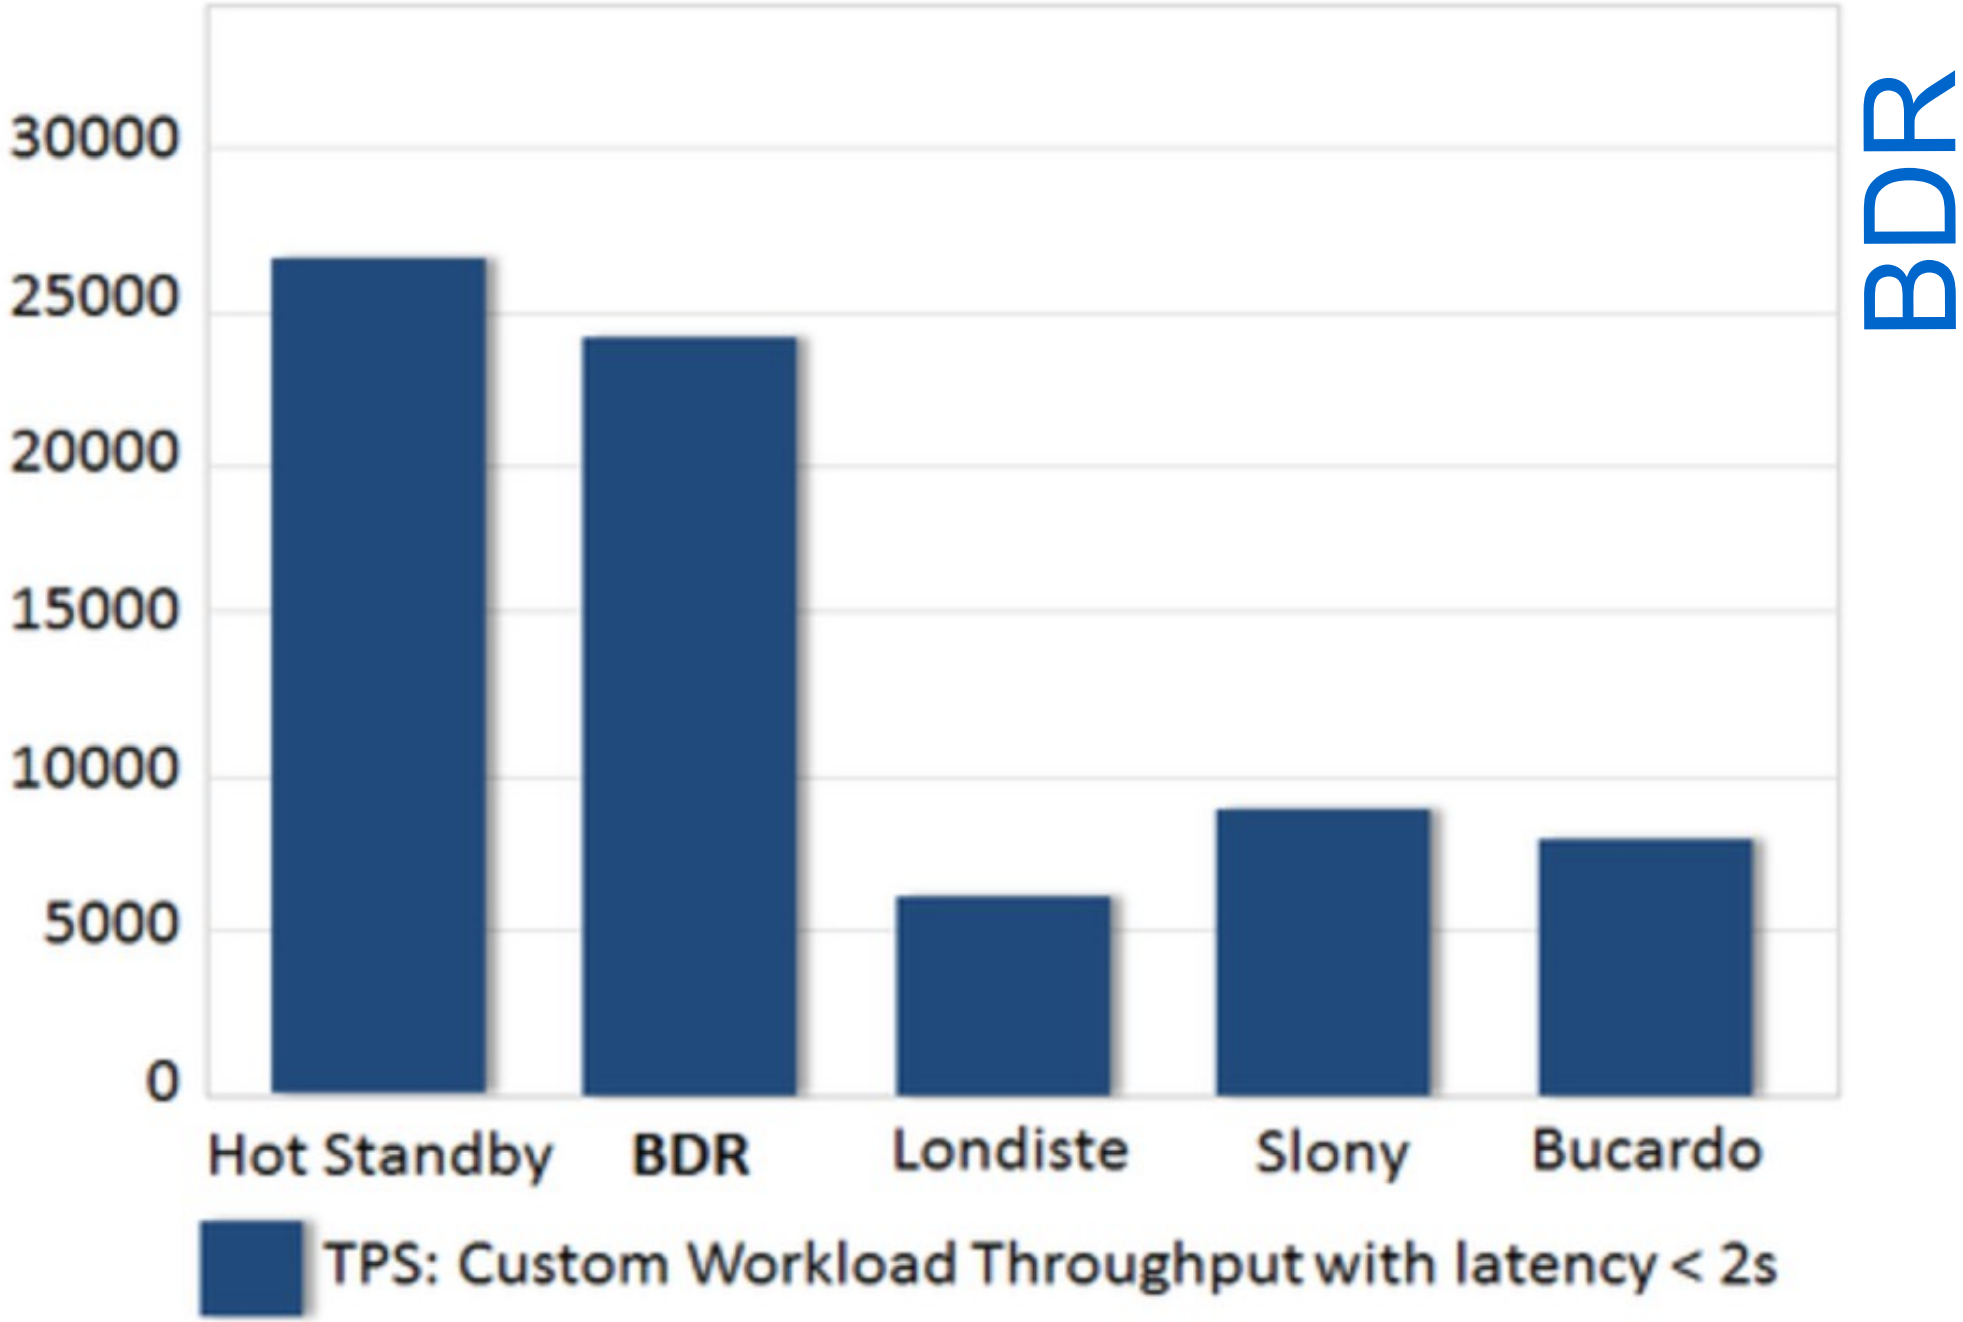

Database Access

USE CASE

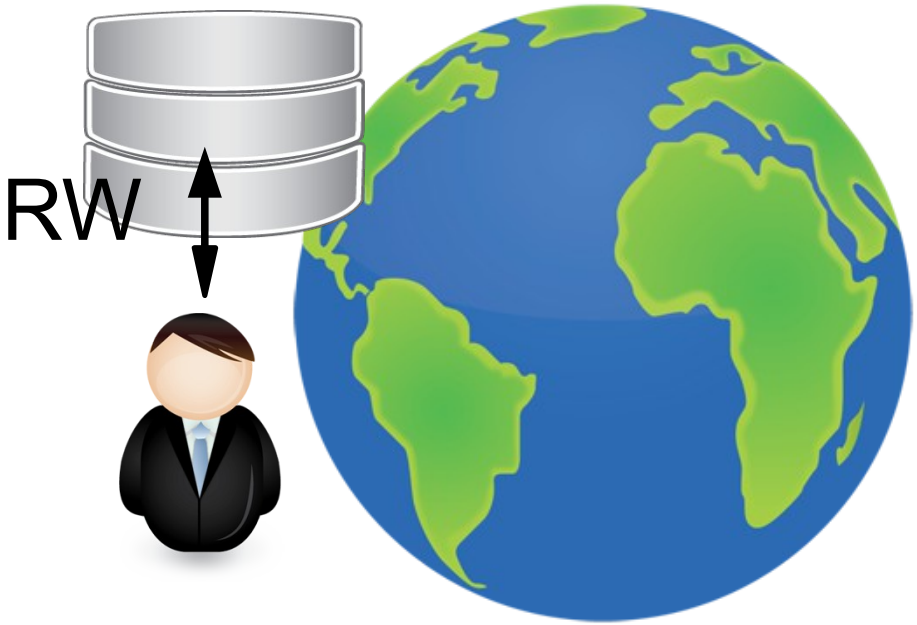

Database Access

USE CASE

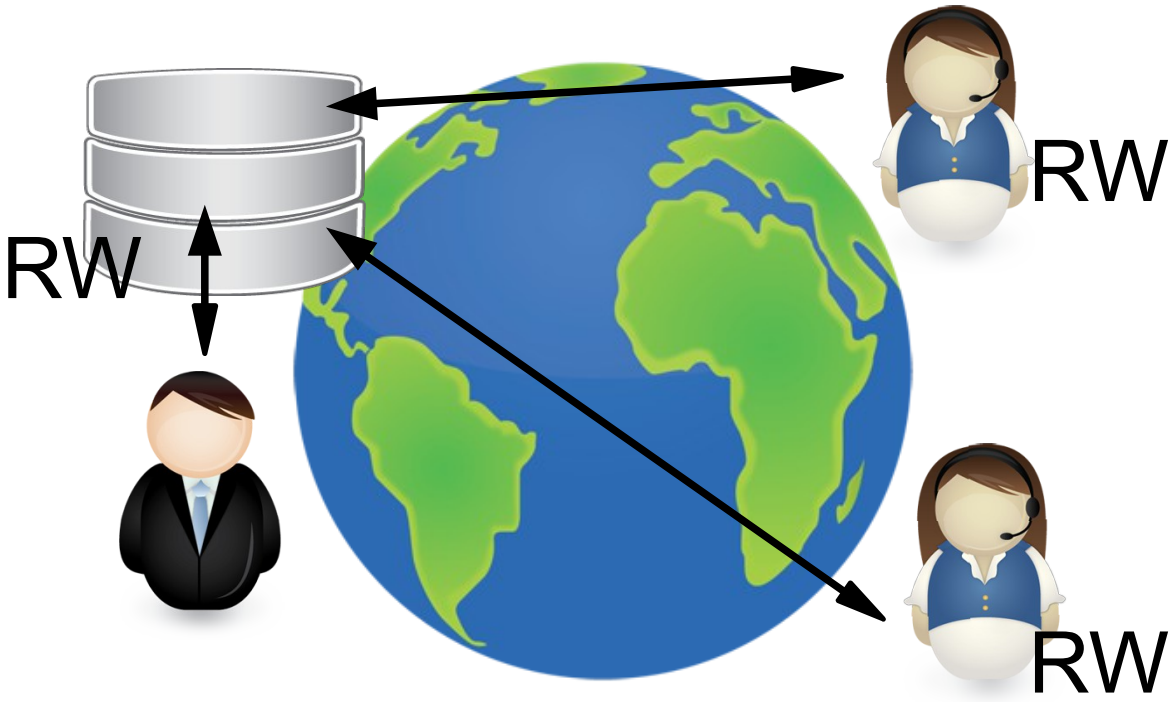

Database Copies

USE CASE

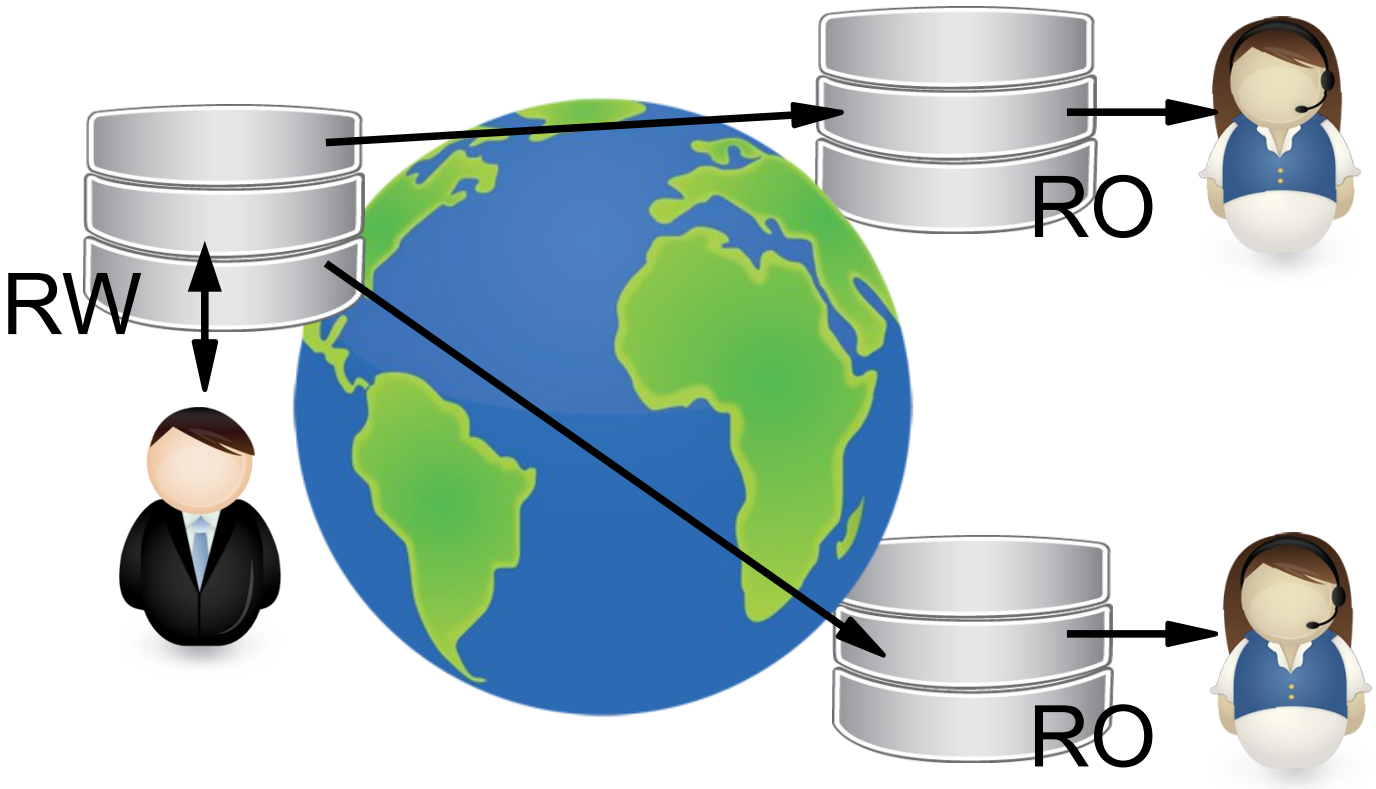

Multi-Master Database

USE CASE

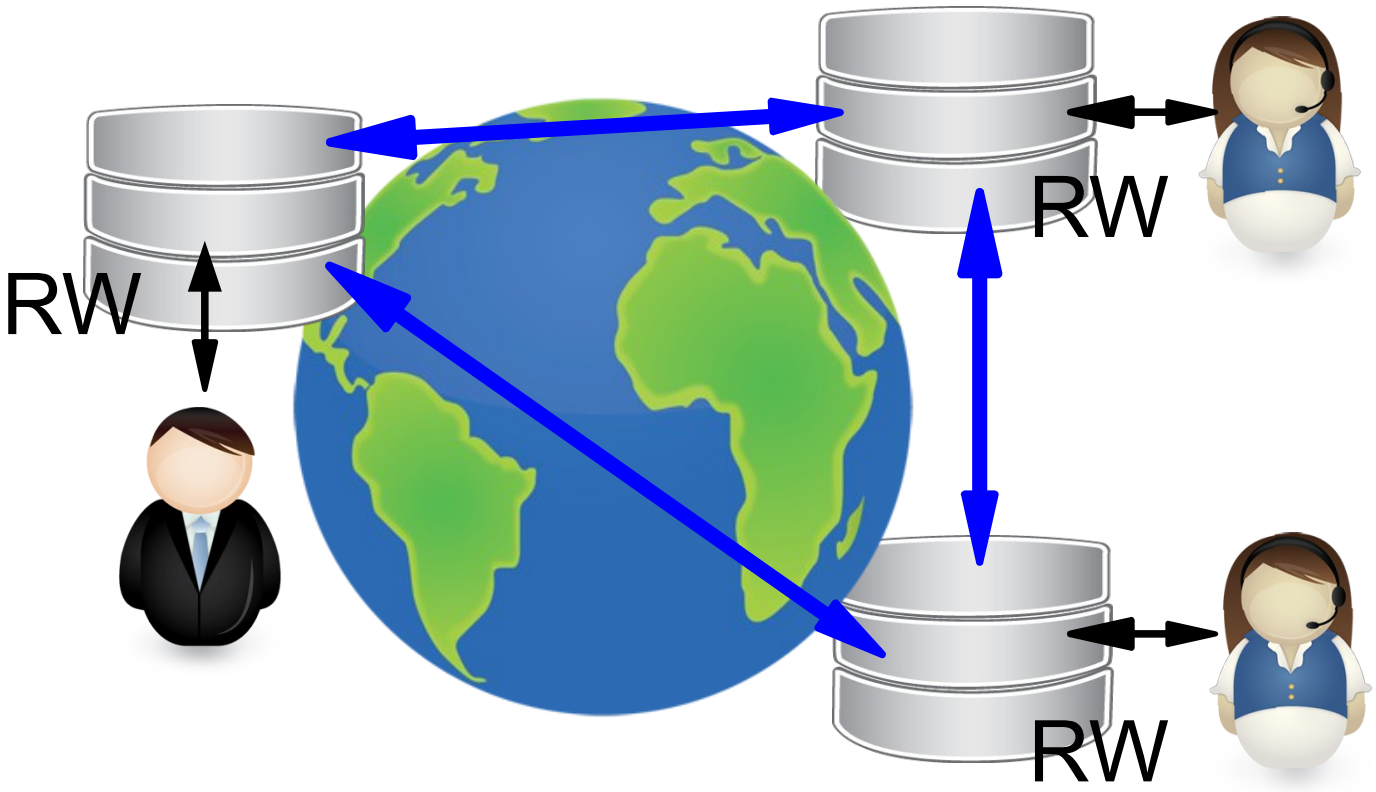

Force Majeure

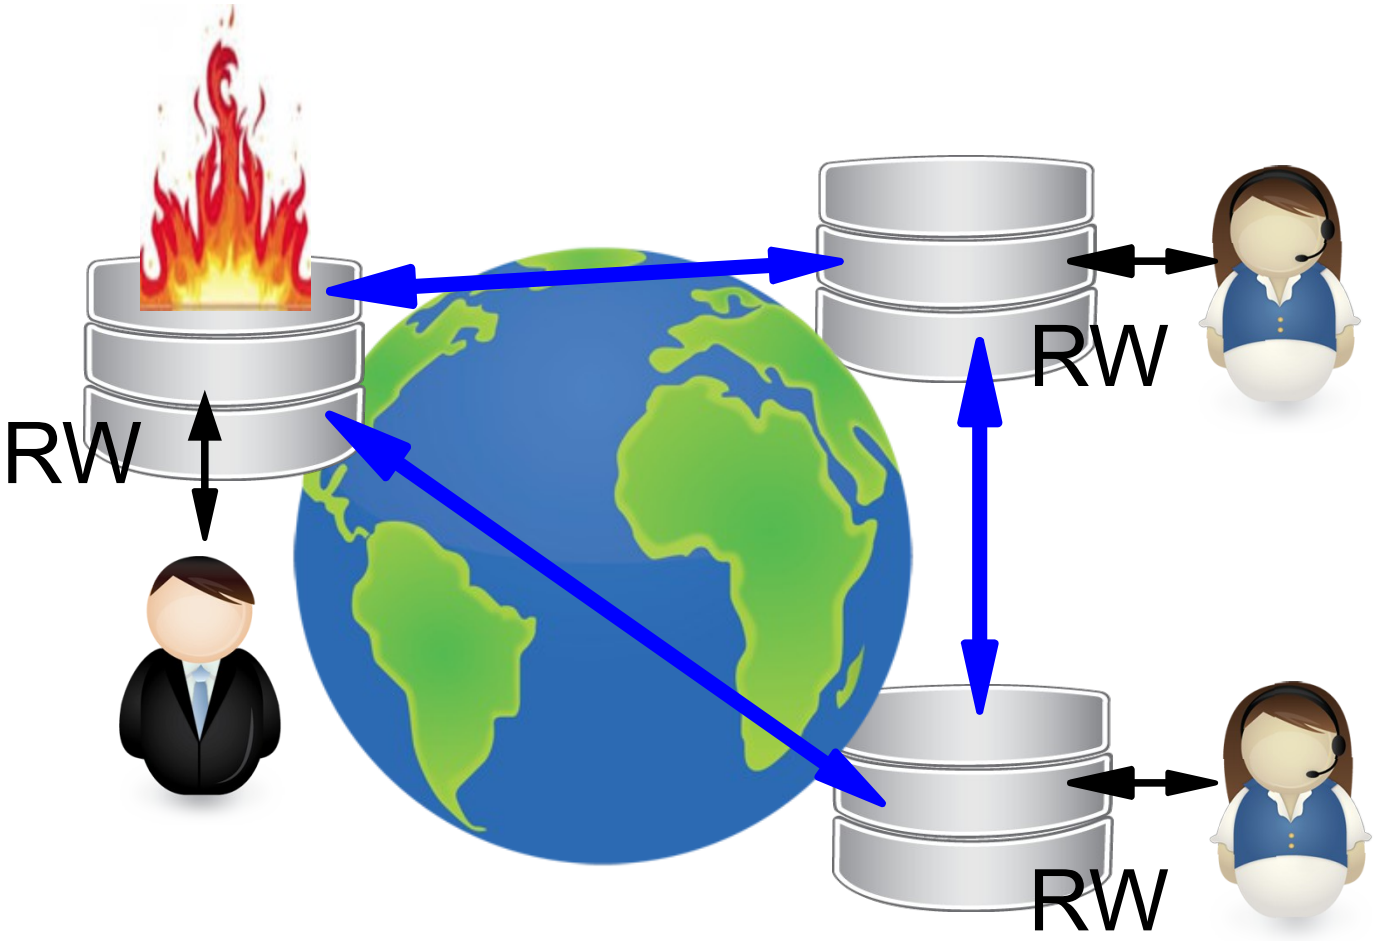

BAD NEWS

Logical & Physical

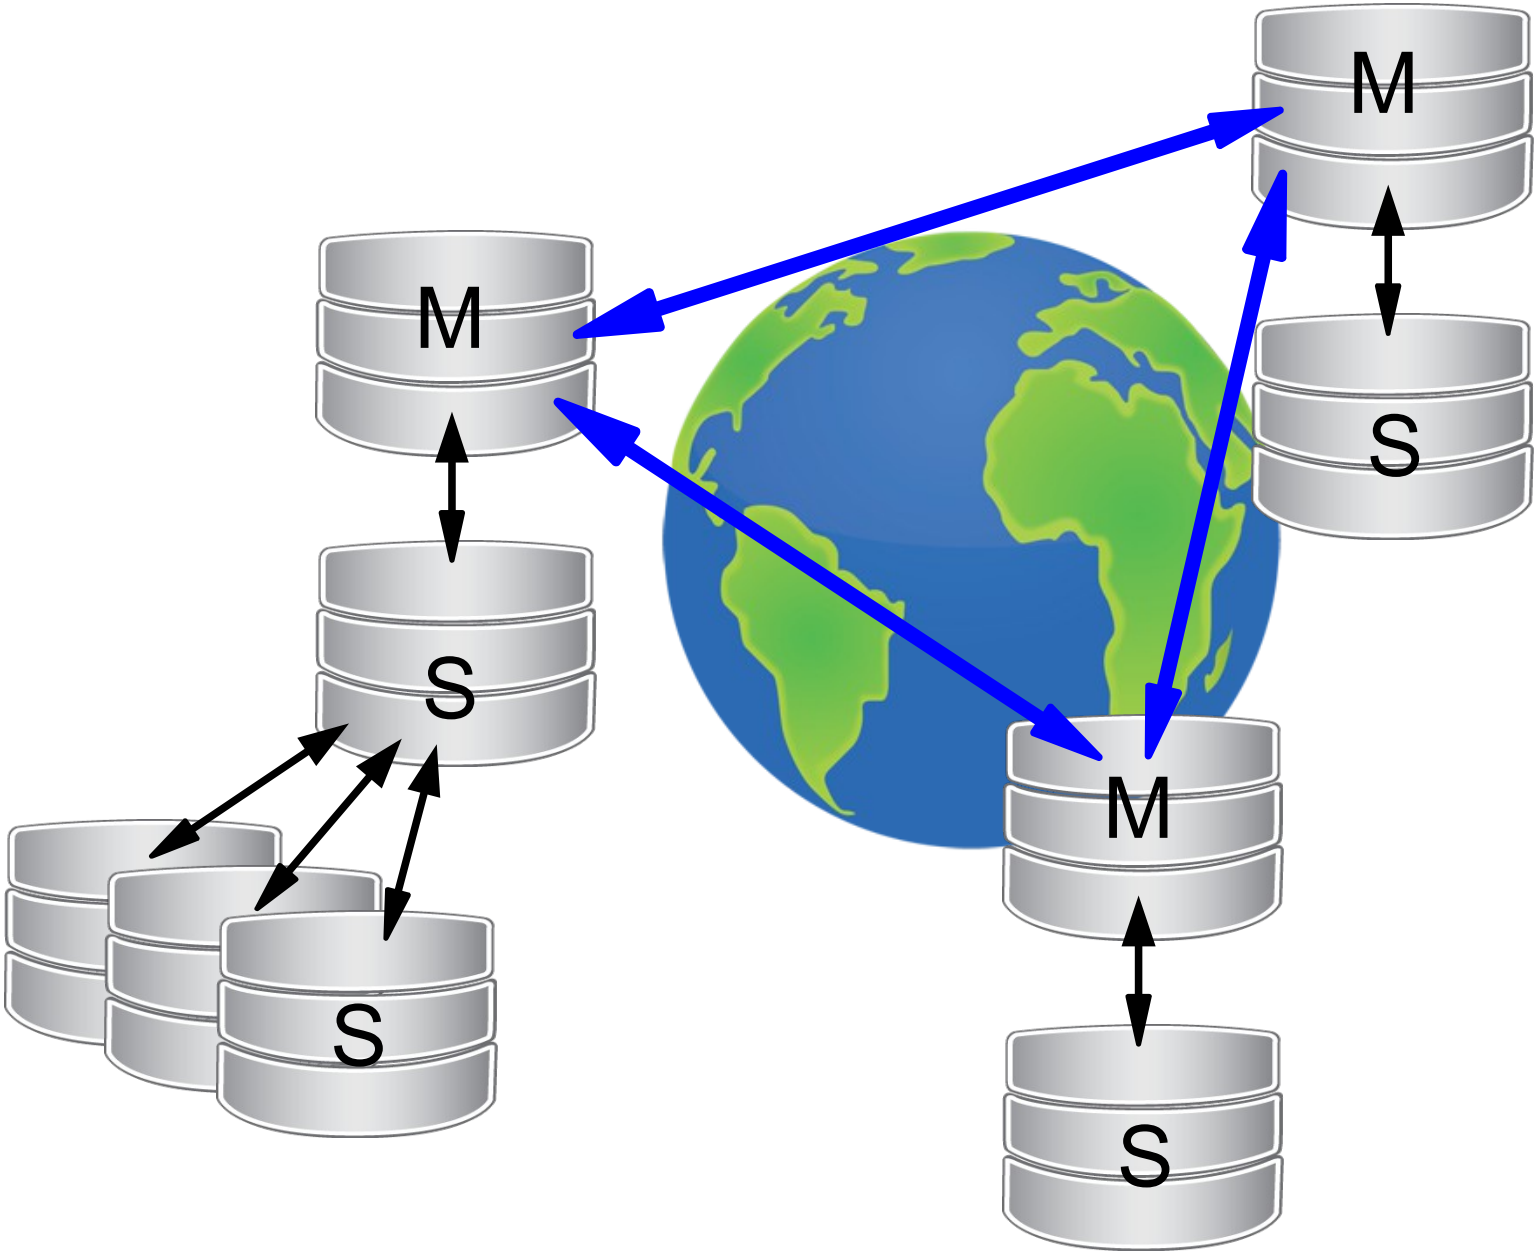

TECH ARCHITECTURE

Physical Streaming Replication

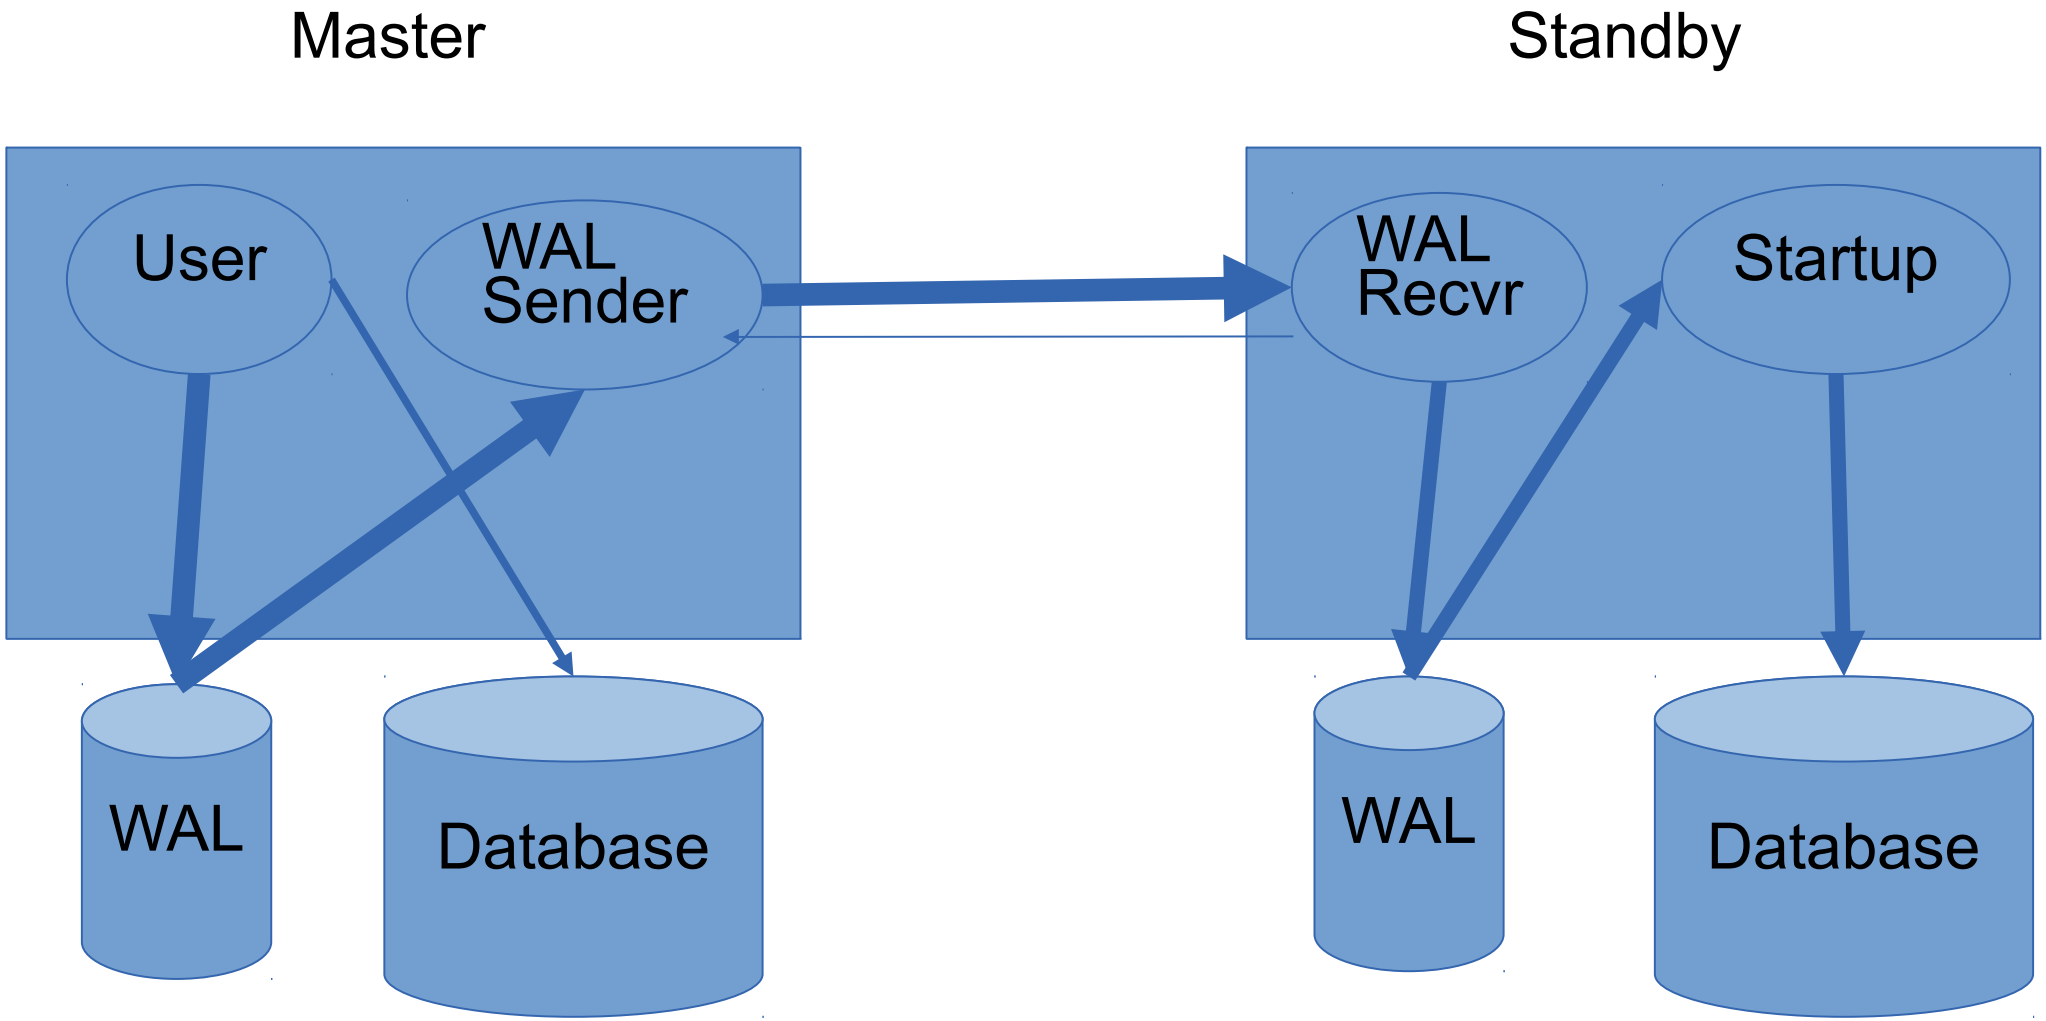

Logical Streaming Replication

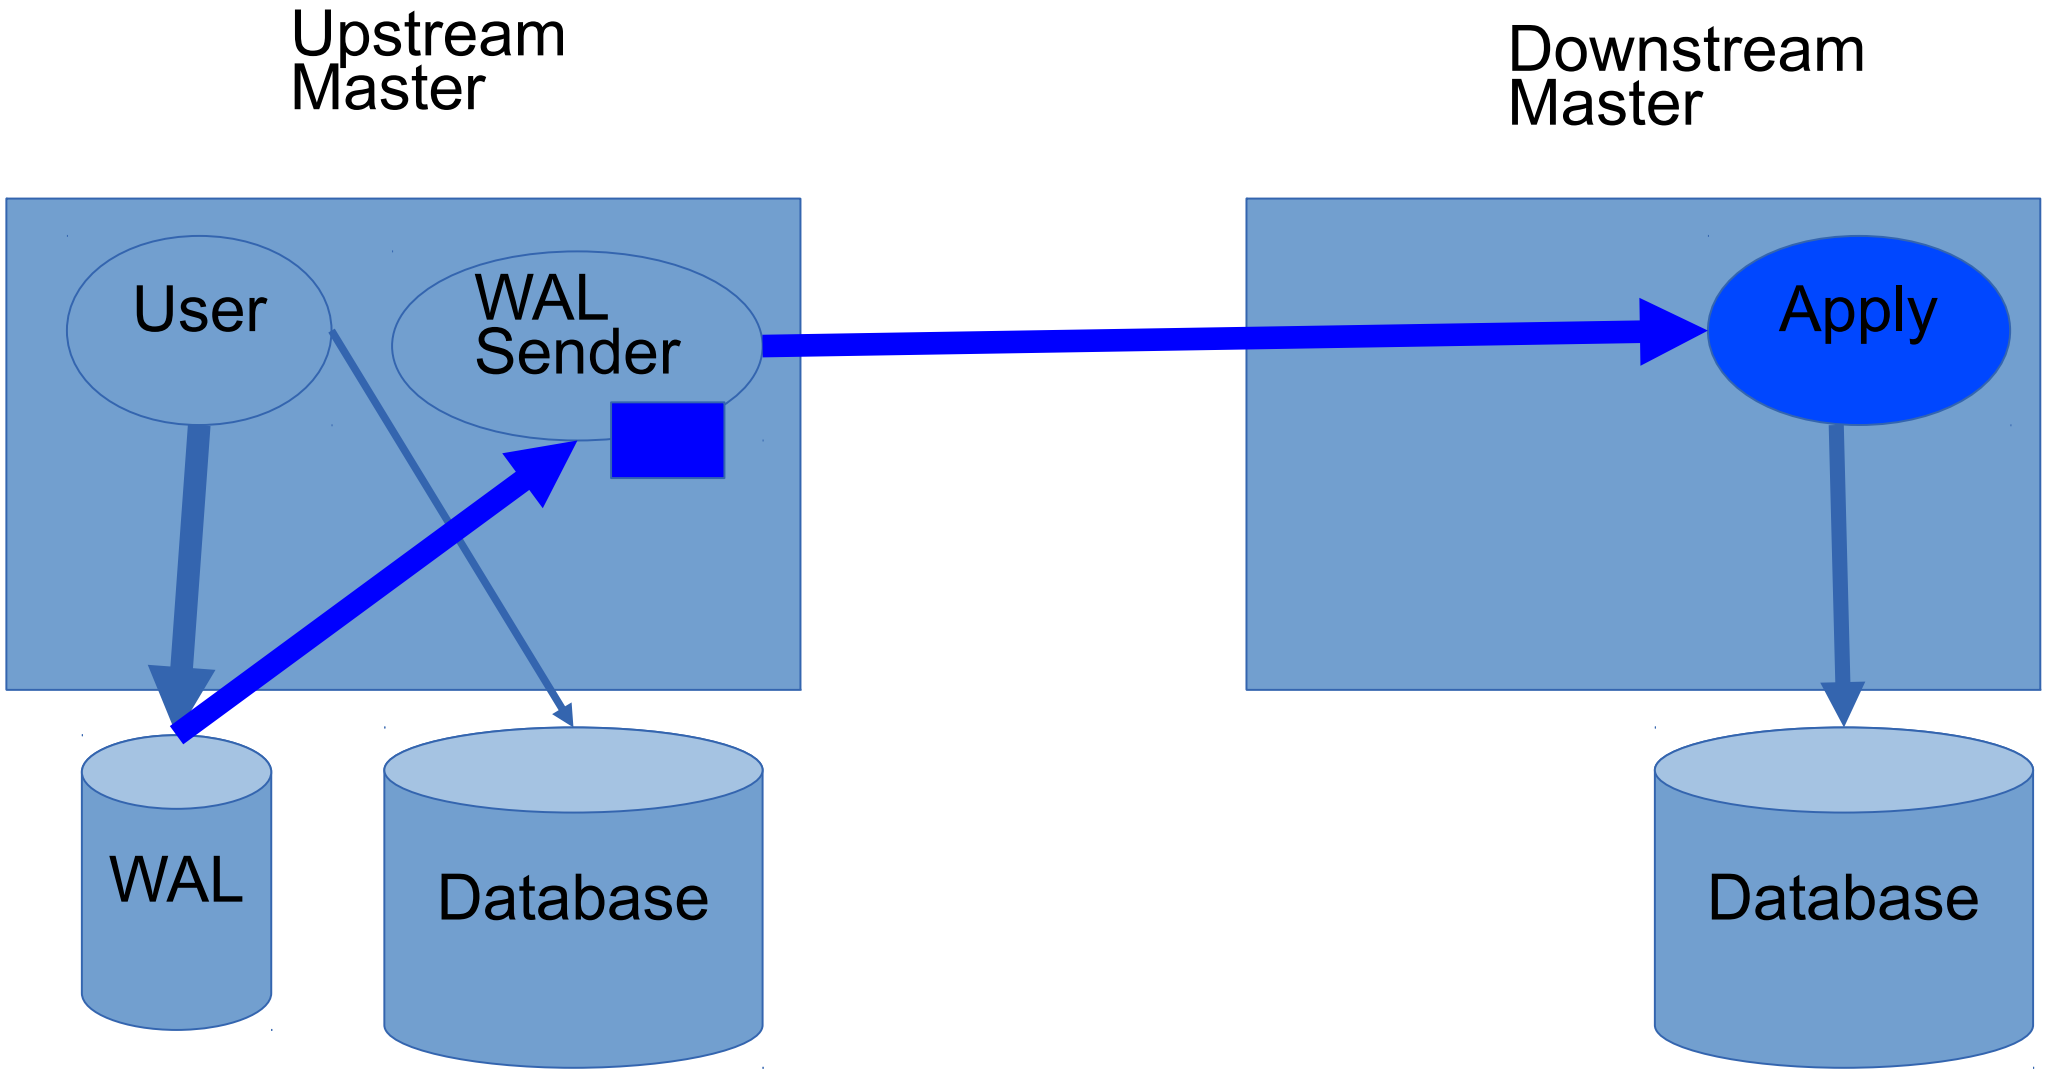

PostgreSQL 9.4 Infrastructure

- 9.3 Background Workers
- 9.4 Dynamic Shared Memory/Locks
- 9.4 Logical Decoding
- 9.4 pg_xlogdump
- 9.4 Replication Slots
- 9.4 REPLICA IDENTITY
- 9.4 user_catalog_table

BDR 0.8.0.7

Production NOW!

- Ease of Use
- Low Maintenance (DDL)
- Fast, low latency
- Low overhead
- No serialization problems
- Global Sequences

- Read-write access
- Allows different schema
- Allows temp tables
- Allows different security
- Limited Physical effects

- Allows Multi-Master
- Co-exists with Physical Rep

BDR 0.9

- Ease of Use
- Low Maintenance (DDL)
- Fast, low latency
- Low overhead
- No serialization problems
- Global Sequences

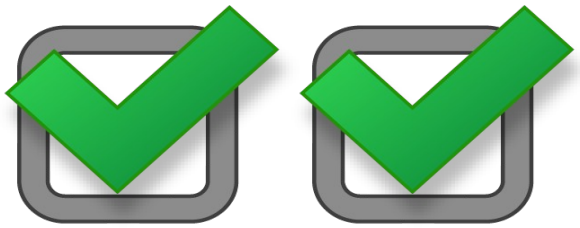

1Q2015

- Read-write access
- Allows different schema
- Allows temp tables
- Allows different security
- Limited Physical effects
- Selective replication
- Allows Multi-Master
- Co-exists with Physical Rep
- Online Upgrade

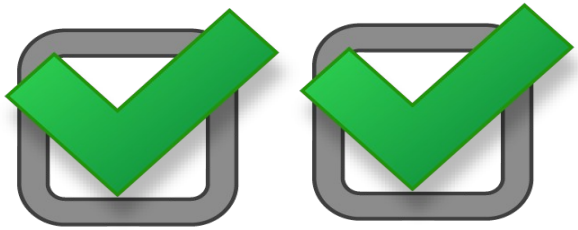

Deliverables

- **BDR 0.8.0.7**
  - Full function, already in use in production
  - Variant of PostgreSQL, open source
  - All changes submitted to PostgreSQL core
  - Production Now, merged later 9.6

- **BDR 0.9**
  - Extension for 9.4+
  - One-way replication

- **Integration into PostgreSQL (patches)**
  - 9.5: Online Upgrade from 9.4 to 9.5+
  - 9.6: More features

Integration Schedule

	Postgres Core	Extension	Variant Distro
9.4 – now!	+Logical Decoding +Dynamic Shm +Replication Slots +REPLICA IDENTITY	UDR Extension	Full BDR – up to 48 nodes
9.5	+Commit Timestamps +Sequence AM API +Replication Identifiers +DDL Event Triggers	UDR Extension (Faster!)	Slim BDR
9.6	+Cluster Metadata +UDR Plugin +Replication Sets +Global Sequences	(Gone!)	Multi-master BDR
9.7	+Multi-Master +New conflict models		(Gone!)

BDR

- 10 years ago, **Slony** delivered practical solution for PostgreSQL replication that works outside of core
  - but it had problems
- 10 years later, **BDR** delivers a practical solution for PostgreSQL replication that works outside of core
  - nearly all problems solved
  - *and* submitted changes for core PostgreSQL
- BDR will become “PostgreSQL replication”
  - Slony no longer being developed full-time
  - Londiste no longer being developed
  - It needs your support

Futures

- Many more features – we ain't done yet!
- More Practical Solutions
- Quorum Sync Rep
- Replication Transforms
- Auto Sharding
- Massively Parallel Query
- Distributed Optimisation