Ask the Expert:

How to Minimise Downtime and Increase Productivity with High Availability

Kevin Li, Borys Neselovskyi

23 February 2022

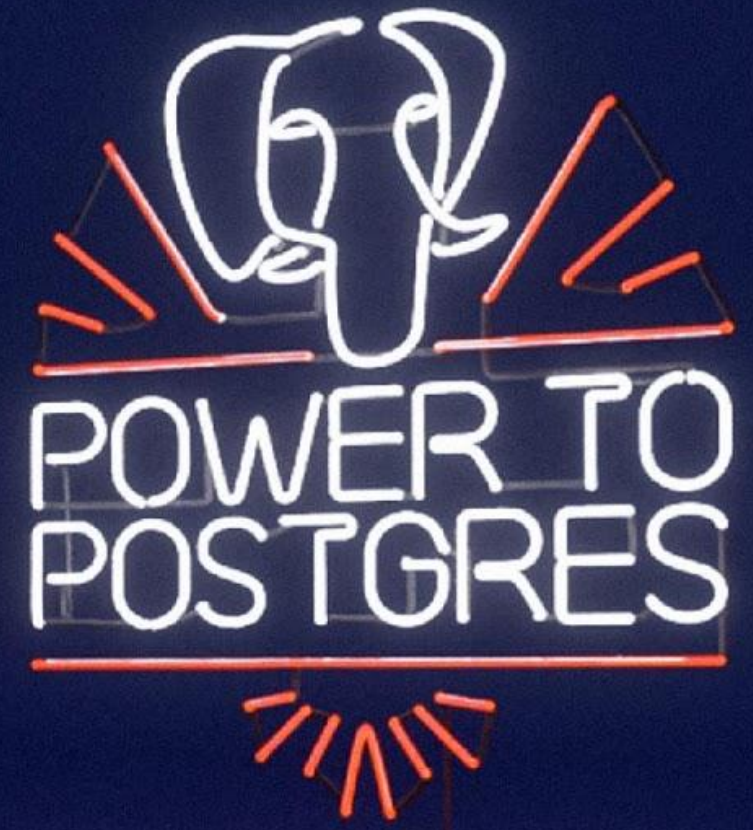

EDB

Our speakers

Host

Kevin Li
Director Sales Engineer, EMEA

Speaker

Borys Neselovskyi
Regional Sales Engineer, DACH

Agenda

- High Availability Concepts
- What are the RPO, RTO, GRO?
- Single Point of Failure
- How to Choose the right solution

High Availability Concepts

EDB

High Availability

High availability (HA) is a characteristic of a system, which aims to ensure an agreed level of operational performance, usually uptime, for a higher than normal period.

Key principles:

- Eliminate single point of failure
- Reliable crossover
- Detection of failures

Ref: https://en.wikipedia.org/wiki/High_availability

EDB

Scheduled/Unscheduled downtime

- **Scheduled/planned downtime** is a result of maintenance that is disruptive to system operation and usually cannot be avoided with a currently installed system design.
  - It include patches to system software that require a reboot or system configuration changes that only take effect upon a reboot.
- **Unscheduled/Unplanned downtime** is the result of downtime events due to some physical failures/events, such as hardware or software failure or environmental anomaly.
  - For example, power outages, failed CPU or RAM components (or possibly other hardware components failure), network failure, security breaches, or various applications, middleware, and operating system failures result in Unplanned outage/Unscheduled downtime.

EDB

Availability calculation

Calculated/expressed as a percentage of uptime in a given year based on the service level agreements. Some companies exclude the planned outage/scheduled downtime based on their agreements with customers on the availability of their services.

Availability %	Downtime per year	Downtime per month	Downtime per week	Downtime per day
99.99% ("four nines")	52.60 minutes	4.38 minutes	1.01 minutes	8.64 seconds
99.995% ("four and a half nines")	26.30 minutes	2.19 minutes	30.24 seconds	4.32 seconds
99.999% ("five nines")	5.26 minutes	26.30 seconds	6.05 seconds	864.00 milliseconds

EDB

POLL:

What is your definition of High Availability?

1) Elimination of single points of failure
2) Reliable crossover
3) Detection of failures as they occur
4) All of the above

What are the RPO/RTO/GRO?

Recovery Point Objective (RPO)

RPO is a measurement of time from the failure, disaster or comparable loss-causing event.

RPO can be used to measure:

- How far back must go, stretching back in time from the disaster to the last point where data is in a usable format
- How frequently you need to back-up your data, although an RPO doesn't represent additional needs like restore time and recovery time.
- How much data is lost following a disaster or loss-causing event
- Ex: RPO = 2 hours

  * In case of a crash I may forget everything that I did in the last 2 hours!

EDB

Recovery Time Objective (RTO)

The amount of time an application can be down and not result in significant damage to a business and the time that it takes for the system to go from loss to recovery

Recovery process includes

- The steps that IT must take to return the application
- And its data to its pre-disaster state.

RPO vs. RTO

RPOs and RTOs are key concepts for maintaining business continuity and function as business metrics for calculating how often your business needs to perform data backups.

- RTOs coincide with recovery point objectives (RPOs), a measurement of time from the failure, disaster or similar loss-causing event.
- RPOs calculate back in time to when your data was last usable, probably the most recent backup.

Geography Recovery Objectives (GRO)

If datacenter becomes unavailable, how long it takes for the service to become available again.

- It covers RPO/RTO for making services available across the geography.

Eliminate Single Point of Failure

Eliminate Single Point of failure

- Physical Replication
  - WAL shipping based replication: Replication based on the archived WAL
  - Streaming replication (SR)Streaming WAL files to one or more standbys
- Logical replication
  - Streaming logical data modifications from the WAL.

Physical Replication

Eliminate Single Point of failure

Hot Standby

- Identical to primary system
  - Data is still mirrored in real time
  - Allows READ
  - On failure, can replace primary
- Approaches
  - WAL shipping based
  - Streaming WAL (widely used after 9.0)

Eliminate Single Point of failure

Hot Standby: WAL shipping

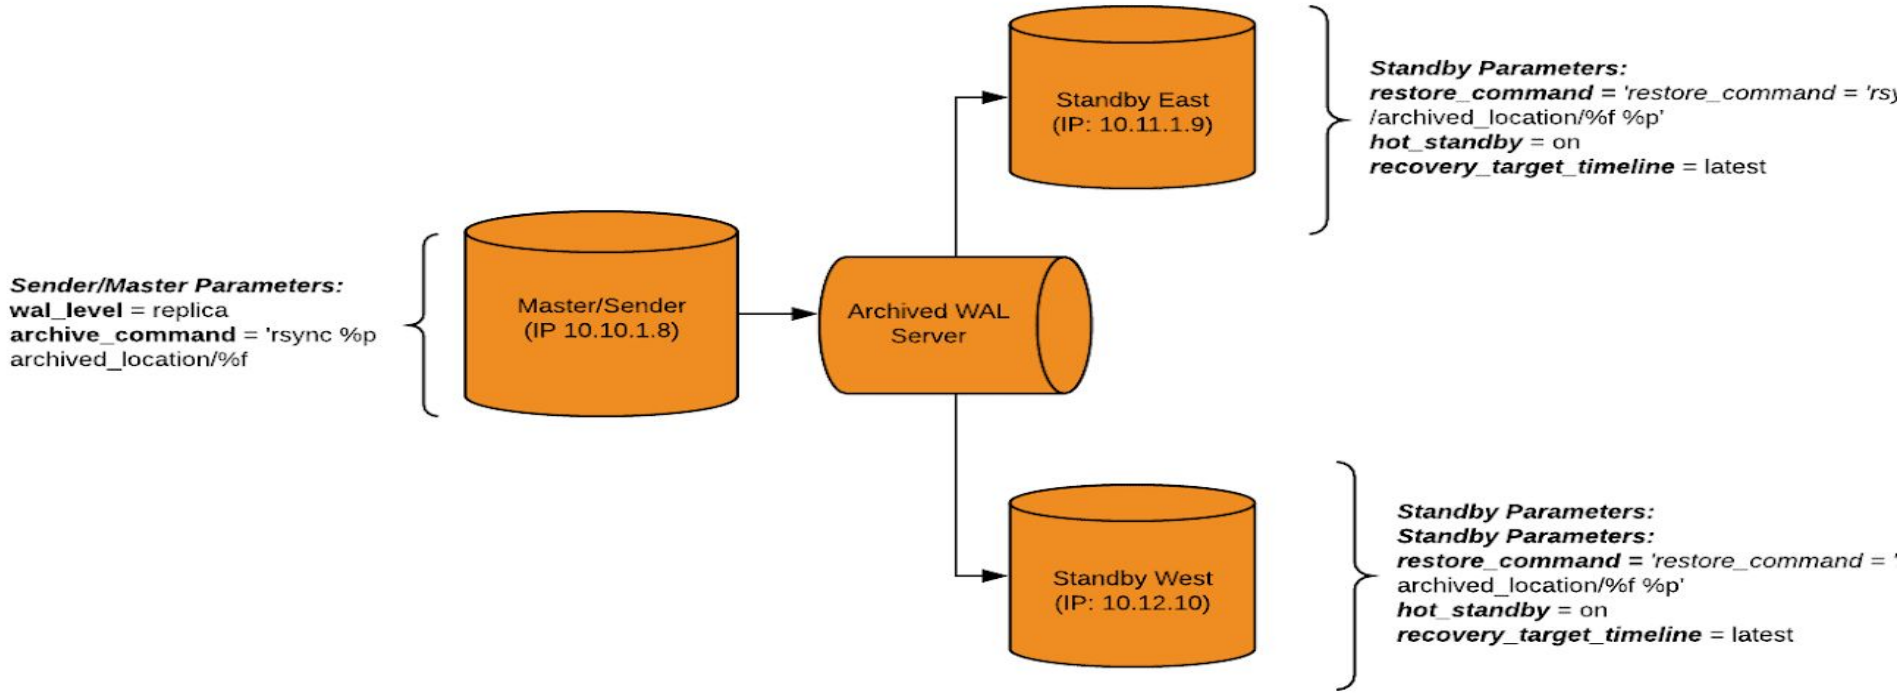

Eliminate Single Point of failure

Hot Standby: Streaming Replication

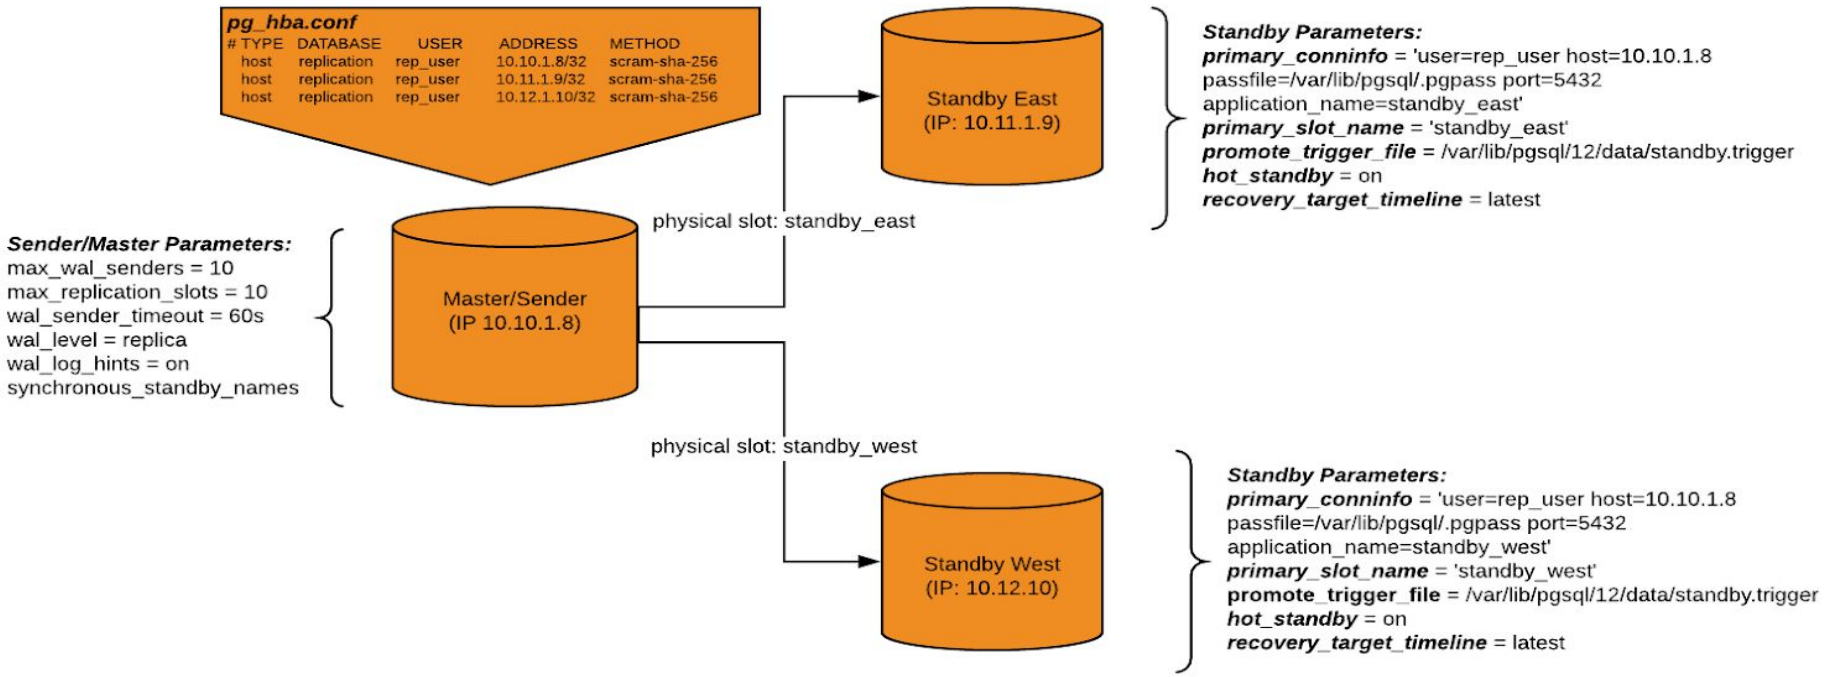

Eliminate Single Point of failure

Streaming Replication

- Asynchronous Streaming Replication
- Synchronous Streaming Replication
  - `synchronous_standby_names`

    E.g.
    - **`FIRST 1 (standby_east, standby_west)`**
    - **`ANY 3 (standby_east, standby_west, eu_standby_east, eu_standby_west)`**
  - `synchronous_commit`
    - **`synchronous_commit - off/local/remote_write/on/remote_apply`**

Reduce data loss → synchronous replication

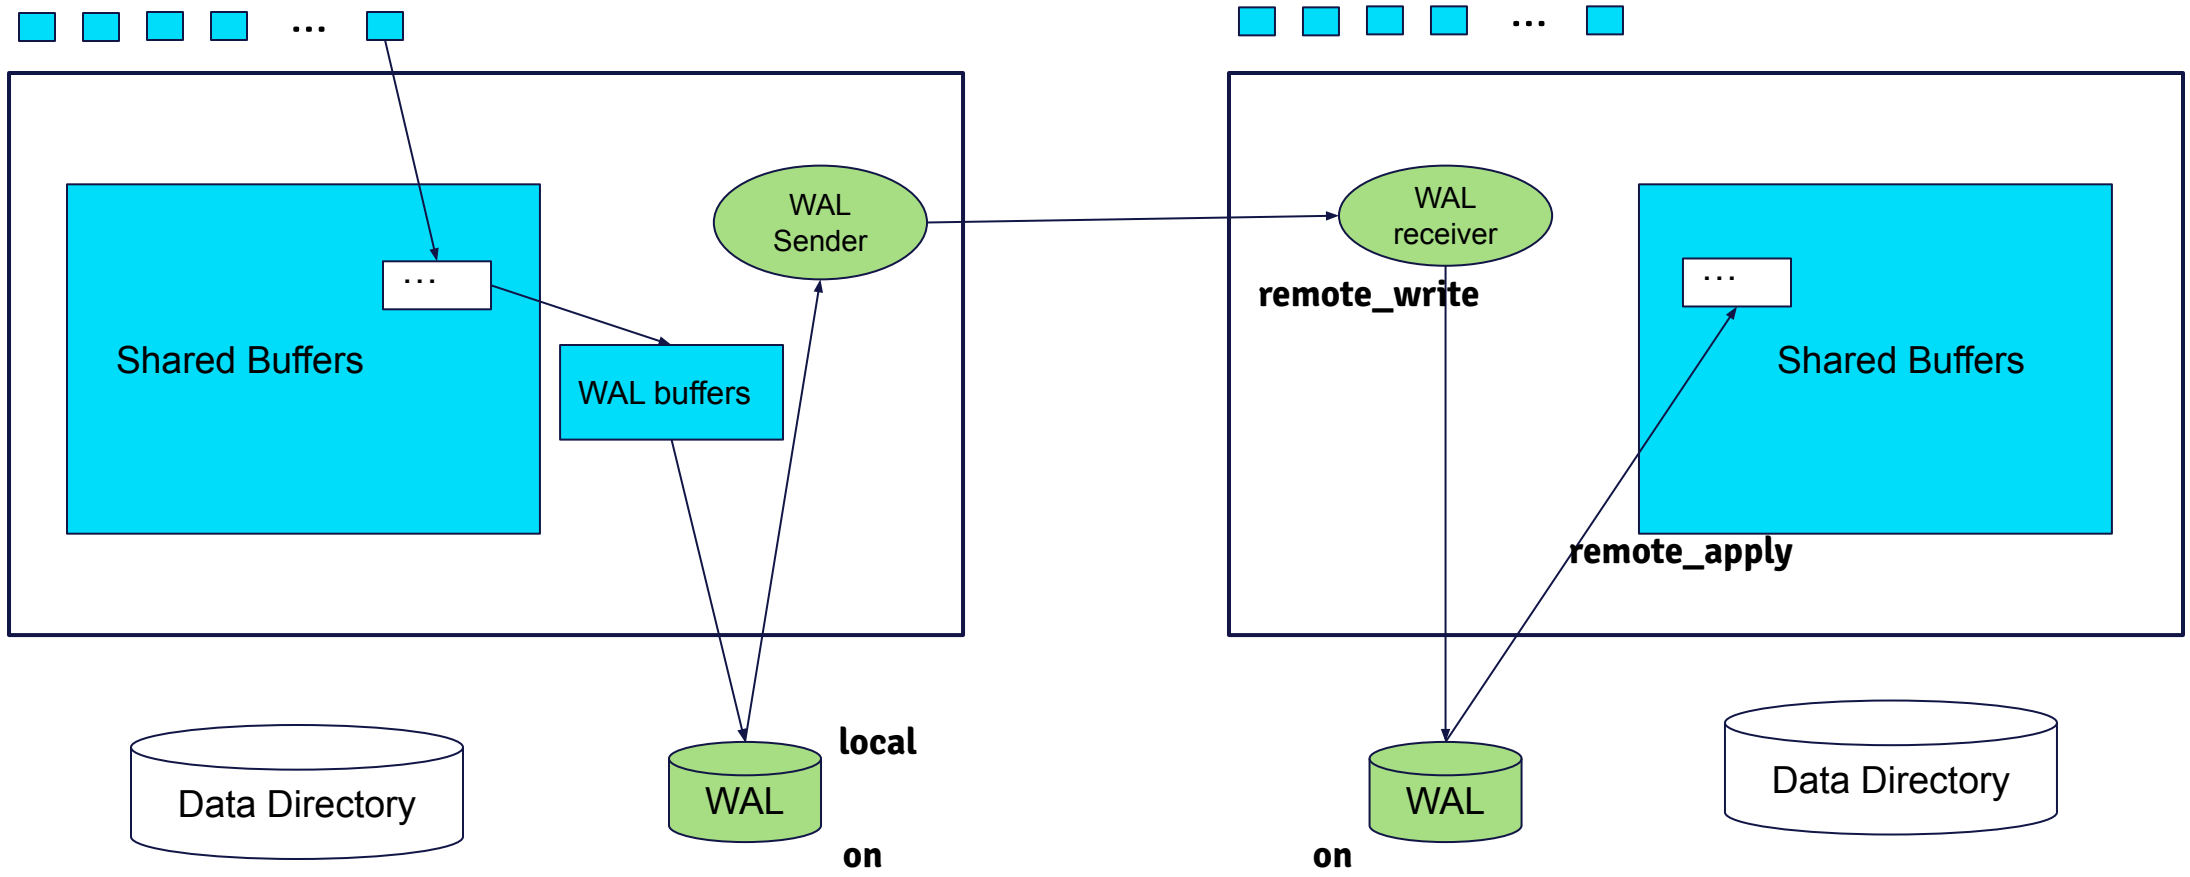

EDB

Reliable crossover and detection

Reliable Crossover & Detection

- In a redundant system, the crossover point itself becomes a single point of failure.
- Fault-tolerant systems must provide a reliable crossover or automatic switchover mechanism to avoid failure.
- Detection of failures:
  - If the above two principles are proactively monitored, then a user may never see a system failure.

Reliable Crossover & Detection

EDB Postgres Failover Manager:

- Continuously monitors your PostgreSQL service to automatically detect failures.
- After an outage is confirmed, Failover Manager automatically promotes the most up-to-date standby database as the new master.

Continuously monitors system health

Detects failures and takes action to maintain SLAs. Sends email alerts based on events.

Automatically fails over to the most current standby

Reconfigures other standbys to point to the new master.

Reconfigures load balancers on failover

Easily integrates with pgPool, F5 Networks, and other load balancers.

Avoids “split brain” scenarios

Prevents two nodes from thinking that each is the master, minimizing data loss.

Reliable Crossover & Detection (DEMO)

EDB Postgres Failover Manager:

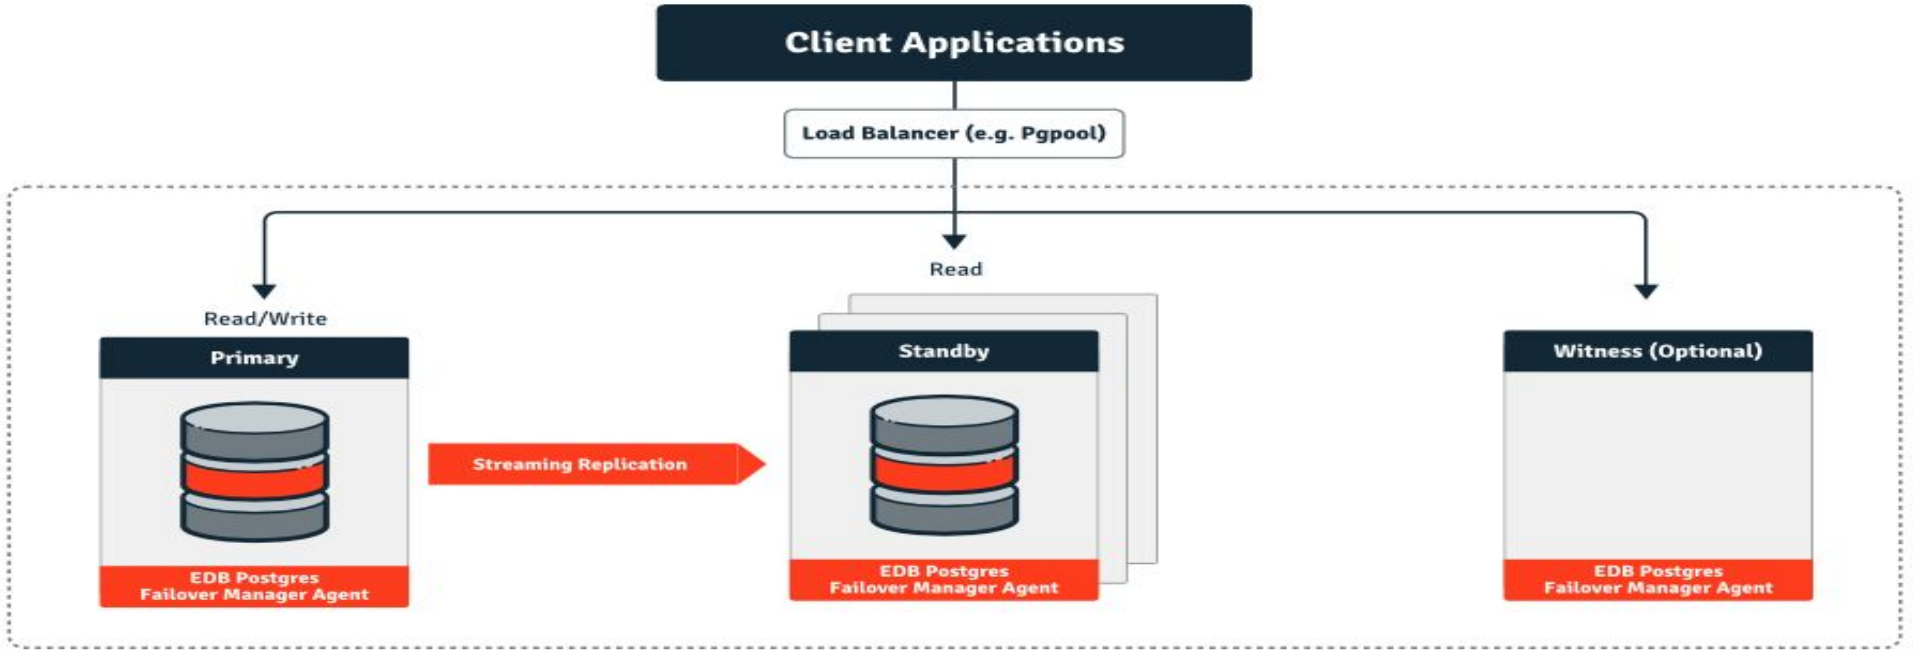

Logical replication

Logical replication

Use Cases

- Sharing a subset of the database between multiple databases
- Replicating between different major versions of PostgreSQL
- Replicating between PostgreSQL instances on different platforms (for example Linux to Windows)
- Consolidating multiple databases into a single one (for example for analytical purposes)
- Major Upgrade with the Minimal Downtime Approach
- Multi-Master Replication*
- Bi-Directional Replication**

* Multi-Master Replication is possible with native logical Replication, pglogical, BDR

** Bi-Directional Replication is possible with BDR

BDR is more than bi-directional replication

Multi-master replication enabling highly available and geographically distributed Postgres clusters

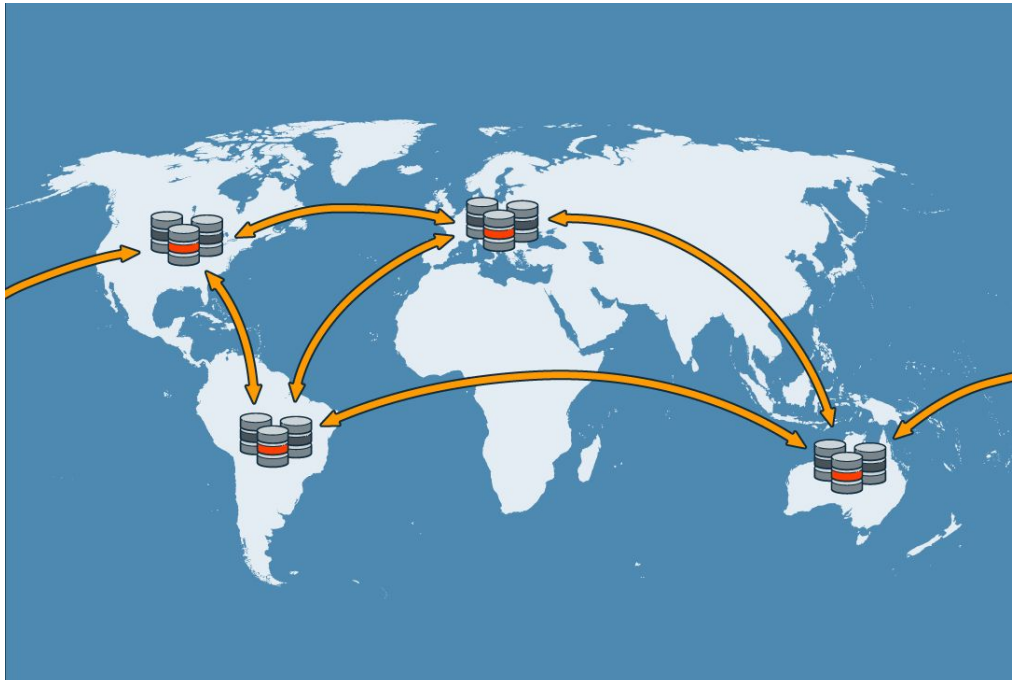

- Logical replication of data and schema enabled via standard PostgreSQL extension
- Data consistency options that span from immediate to eventual consistency
- Robust tooling to manage conflicts, monitor performance, and validate consistency
- Deploy natively to cloud, virtual, or bare metal environments

EDB

BDR Feature Overview

A full-featured multi-master replication solution for PostgreSQL clusters

Essentials

Provides the essential multi-master capabilities for PostgreSQL clusters

- Enables application and database upgrades without requiring downtime
- Provides clusters row level eventual consistency by default
- Tools to monitor operation and verify data consistency
- Extends PostgreSQL logical replication beyond unidirectional, standby use cases

Advanced

Includes advanced conflict management, data-loss protection, and up to 5X throughput

- Guards applications from committing transactions more than once
- Conflict-free synchronous replication with two phase commit
- Concurrent updates using conflict-free replicated data types (CRDTs)
- Configurable column level conflict resolution along with customizable conflict handling and transformation

BDR in AlwaysOn architecture (Demo...)

- Multi-master cluster
- Mesh architecture to minimize latency between nodes
- Raft consensus layer
- Integrated with other services
  - Pooling, backup, proxy
- Multiple possible architectures
  - Logical standbys
  - Subscriber-only nodes
  - Witness nodes
- Cloud, on-premises, or hybrid

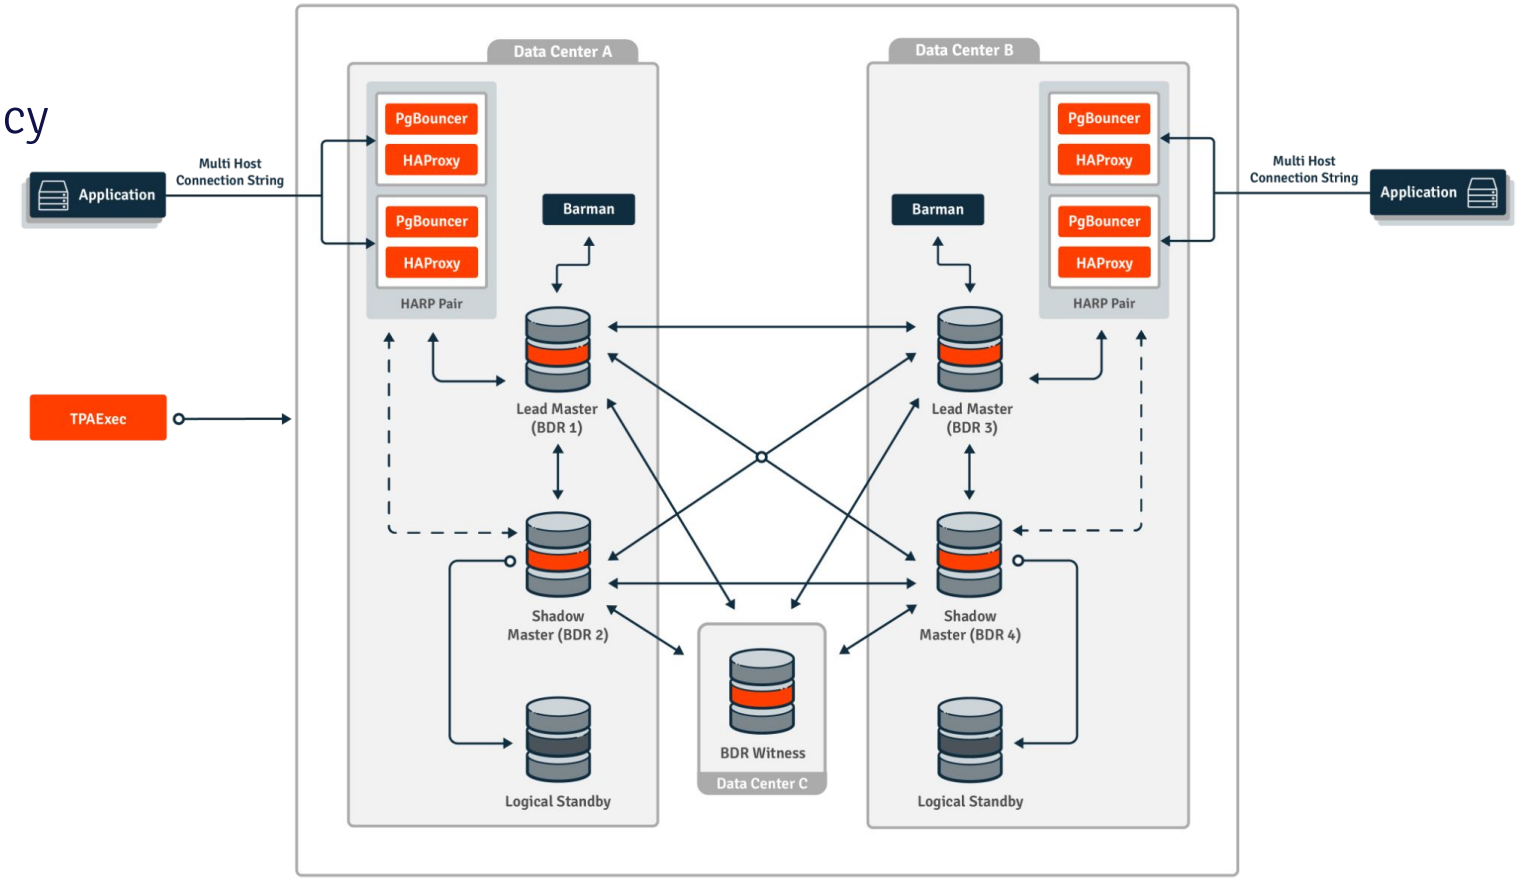

EDB

POLL:

What is your Service Level Agreement (SLA) requirement?

1) 99%
2) 99.9%
3) 99.99%
4) 99.999%

Choose the right solution

Which solution fits your requirements?

	Physical Replication	Logical Replication
High Availability	up to 9999	up to 99999
Scale Read Workloads	yes	yes
Scale Write Workloads	no	yes*
Bi-Directional Replication	no	yes**
Replication between different OS	no	yes
Replication between major versions	no	yes
Minor Upgrade with minimal Downtime	yes	yes***
Major Upgrade with minimal Downtime	no	yes***
Complexity of solution	moderate	complex

* Multi-Master Solution is a feature of logical replication (e.g. pglogical, BDR)

** Bi-Directional Replication is possible with BDR

*** Zero Downtime Patching is possible with BDR

If you have any further questions, please get in touch at:

AskTheExpert@enterprisedb.com

Thank you for attending **Ask the Expert: How to Minimise Downtime and Increase Productivity with High Availability**

Join our next Ask the Expert webinar on 9th March 2022:
How to Master PostgreSQL Databases in the Cloud