Master Class

Running 24x7 Business with Postgres

Ajit Gadge | Field CTO - APJ

September 22, 2022

EDB

Agenda

- Understanding the impact of outages.
- What are DR, HA, and FT and availability?
- EDB Postgres for the Enterprises
- Replication and Always On features.
- Solutions to run 24x7 business with EDB Postgres
- When to use what?

Understanding the impact of Outage

What do customers want today?

Reasons of Outages

- Disasters like storms, floods, and earthquakes.
- Human Errors.
- Network Failure.
- Server instability.
- Software / Services down.
- Storage / Hardware Failure
- Power Outages.
- Security Flaw.
- Usage Spike.
- Planned and Unplanned Maintenance.
- Don’t Know.

What is DR, HA and FT and availability?

Let’s understand DR/HA/FT.

	Disaster Recovery	High Availability	Fault Tolerance
Definition	**DR** refers to the set of policies and procedures in place to ensure the continuity and recovery of mission-critical systems during a disruptive event.	**HA** is to design the system in such a way that it will ensuring your critical systems are always functioning.	A **FT** system is extremely similar to HA but it ensure zero or nearly zero downtime.
Example/Reason	Flood , earthquake , Power Outage at DC	Server/Software failed, Failover, Switch over, Patching, etc.	Server / Software Failed, Network Failure etc.
Expected Availability	99% >	< 99.5	99.999 % - 100 %
RTO	Hrs	< 60 sec	~ 0 sec
RPO	Expected Data loss till last good backup	Near Zero or Zero Data loss	Zero Data loss
Method of Recovery	Backup , Disk mirroring , System mirroring	Replication, Data mirroring, Clustering	Replication, Master – Master
Cost	Low-Medium	Medium/High	High

Why is this so hard?

No downtime for maintenance!

HA Rating	Downtime/month (days:hours:min:secs)	Downtime/year (days:hours:min:secs)
99%	00:07:18:00	03:15:36:00
99.5%	00:03:39:00	01:19:48:00
99.9%	00:00:43:48	00:08:45:36
99.99%	00:00:04:23	00:00:52:34
99.999%	00:00:00:26	00:00:05:15

Patches, updates, and security fixes: at least four maintenance operations per year!

‘High Availability’ or ‘Always On’

Always On - the next step for high availability

- Historic perspective
  - Protects against hardware, network, and software failures
  - Assumes maintenance windows
  - Postgres tools: EDB Failover Manager, Repmgr, Patroni
- Always On - the new business imperative
  - Near-zero downtime, or “Always On” - a must-have for successful digital transformation in a global economy.
  - No more maintenance windows - but patches and upgrades have to be applied
  - 99.999% is the target level (ITIC 2020 Global Server Hardware, OS Reliability Survey)

How important is Always On?

- 99.99% is the standard today
- Downtime costs money and customers
  - Credit cards
  - Websites
  - Payment gateways
  - Single-sign on sites
  - ....

Eighty-eight Percent of Firms Say Hourly Downtime Costs Exceed $300K in 2020

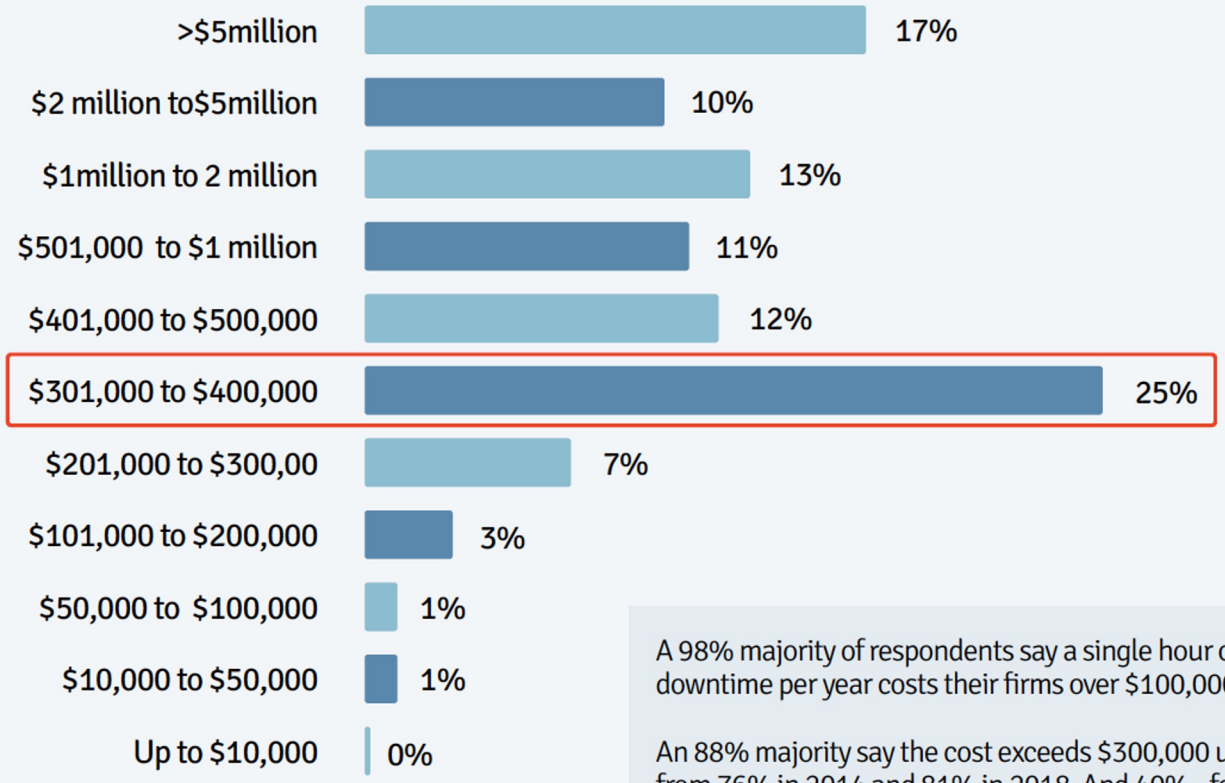

A 98% majority of respondents say a single hour of downtime per year costs their firms over $100,000

An 88% majority say the cost exceeds $300,000 up from 76% in 2014 and 81% in 2018. And 40% - four in 10 enterprises - say hourly downtime costs their firms $1M to >5M

Source: ITIC 2020

EDB - Postgres for the enterprise

Postgres as the clear winner in the database game

Recent data from multiple industry leaders

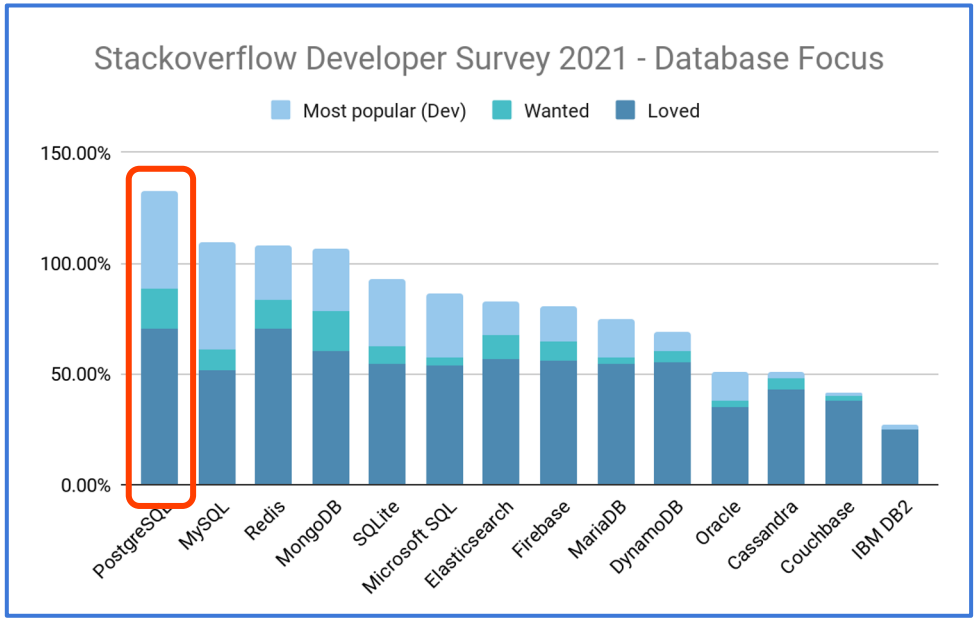

The CNCF End User Community was asked to describe what their companies recommend for different solutions: Adopt, Trial, Assess or Hold. This table shows how the End User companies rated each technology.

Level	Name	Votes
Adopt	Redis	
	Elasticsearch	
	PostgreSQL	
	MySQL	
	Kafka	
	Memcached	
Trial	MongoDB	
	Cassandra	
	Aurora	
	Consul	
	BigQuery	
	CloudSQL	
	AWS DynamoDB	
Assess	MariaDB	
	CockroachDB	
	Vitess	

26 companies contributed 217 data points

Cloud Native Computing Foundation Technology Radar

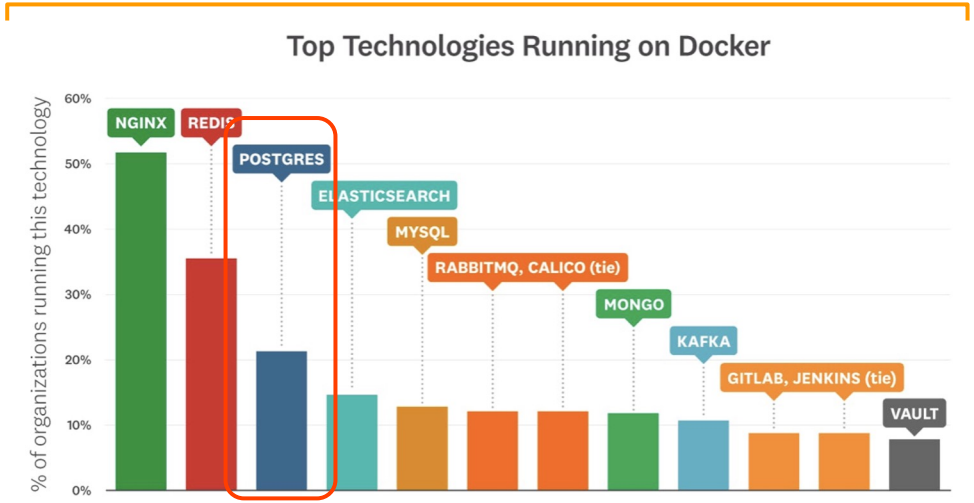

Source: Datadog

EDB Postgres - technology for the enterprise

Technology

PostgreSQL
ACID compliant; MVCC; object relational; extensible; scalable; reliable

EDB Postgres Advanced Server
Oracle compatibility: PL/SQL, 26 native packages, hints, password profiles, OCI/JDBC/ODBC/.NET drivers; Security; Manageability

High Availability
Streaming replication, logical replication, near-zero downtime operation

Management
Database monitoring, alerting, management by exception

Integration
Replication from/to 3rd party; FDW/SQL MED; geo-distributed replication

Migration
Schemas, stored procedures, data, snapshot migration, change data capture

Platforms

Self-managed

Virtualized
aws
Microsoft Azure
vmware

Containerized
RED HAT OPENSHIFT
docker
Amazon EKS
Azure Kubernetes Service (AKS)
Google Kubernetes Engine

Bare Metal

Managed
BIGANIMAL Powered by EDB

BigAnimal (EDB DbaaS)

IBM Cloud

POLARDB

Alibaba DBaaS

Replications and Always On

Choices for HA and AO in Postgres

- Shared disk solutions require 3rd party add-ons such as RedHat Cluster Server. They are rarely used with Postgres today
- We will focus on replication solutions
- Streaming replication and logical replication are key Postgres capabilities

EDB Failover Manager - Building on streaming replication

- Monitors the health of a Postgres Streaming Replication setup
- Automates failover process in the event of a failure (prevents split brain scenarios)
- Use with load balancer to support read scalability separating writes from reads
- Supports Disaster Recovery with offsite replicas
- Near zero downtime maintenance with controlled switchover
- Used in support of various “9s” based high availability requirements

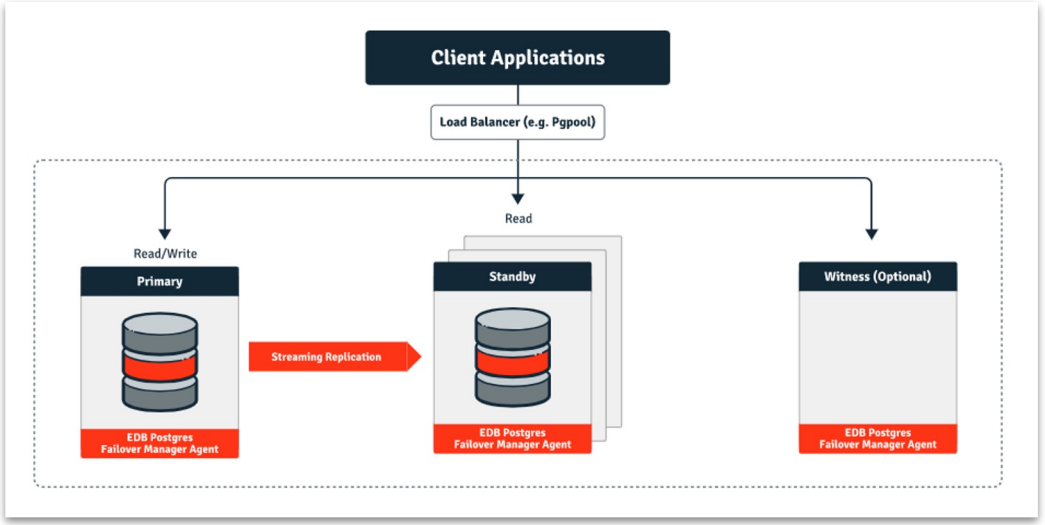

Streaming Example: Departmental Mission Critical

Single Primary and Multiple Standbys with Monitoring and Backup Capabilities

- Postgres database nodes with asynchronous or quorum based synchronous streaming replication
- EFM for management of failover and switchover
- Backup: Separate Barman Server

Property	Description
Recovery Time Objective	< 30 seconds
Recovery Point Objective	Typically a few seconds or less during failover. Dependent on backup strategy for full recovery.
Geographic Redundancy Objective	Dependent on node placement. Typically multi-availability zone.
Target Availability	99.99%

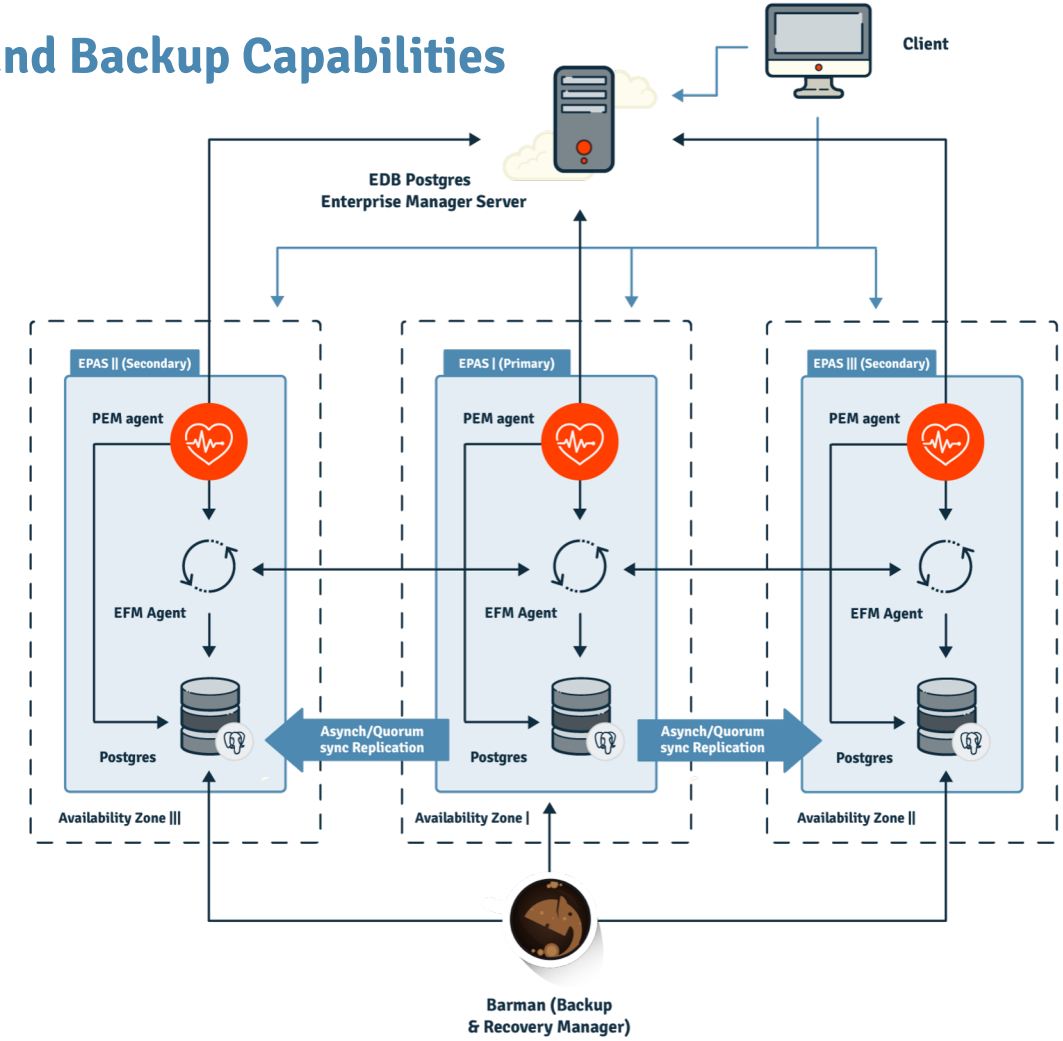

Limitations of Postgres Streaming Replication

- Read/write traffic goes to a single server
- Failover is lengthy process (< 30 sec)
  - Detecting failure
  - Verifying failure
  - Shutdown of current primary and promotion of replica
- Requires binary equivalence of primary and replica
  - Can only be used for minor version upgrades
  - All indexes, tables, etc. have to be identical on primary and replica
- All or nothing! No filters

But: It is easy to set up, super robust, and easy to manage!

Logical Replication Explained

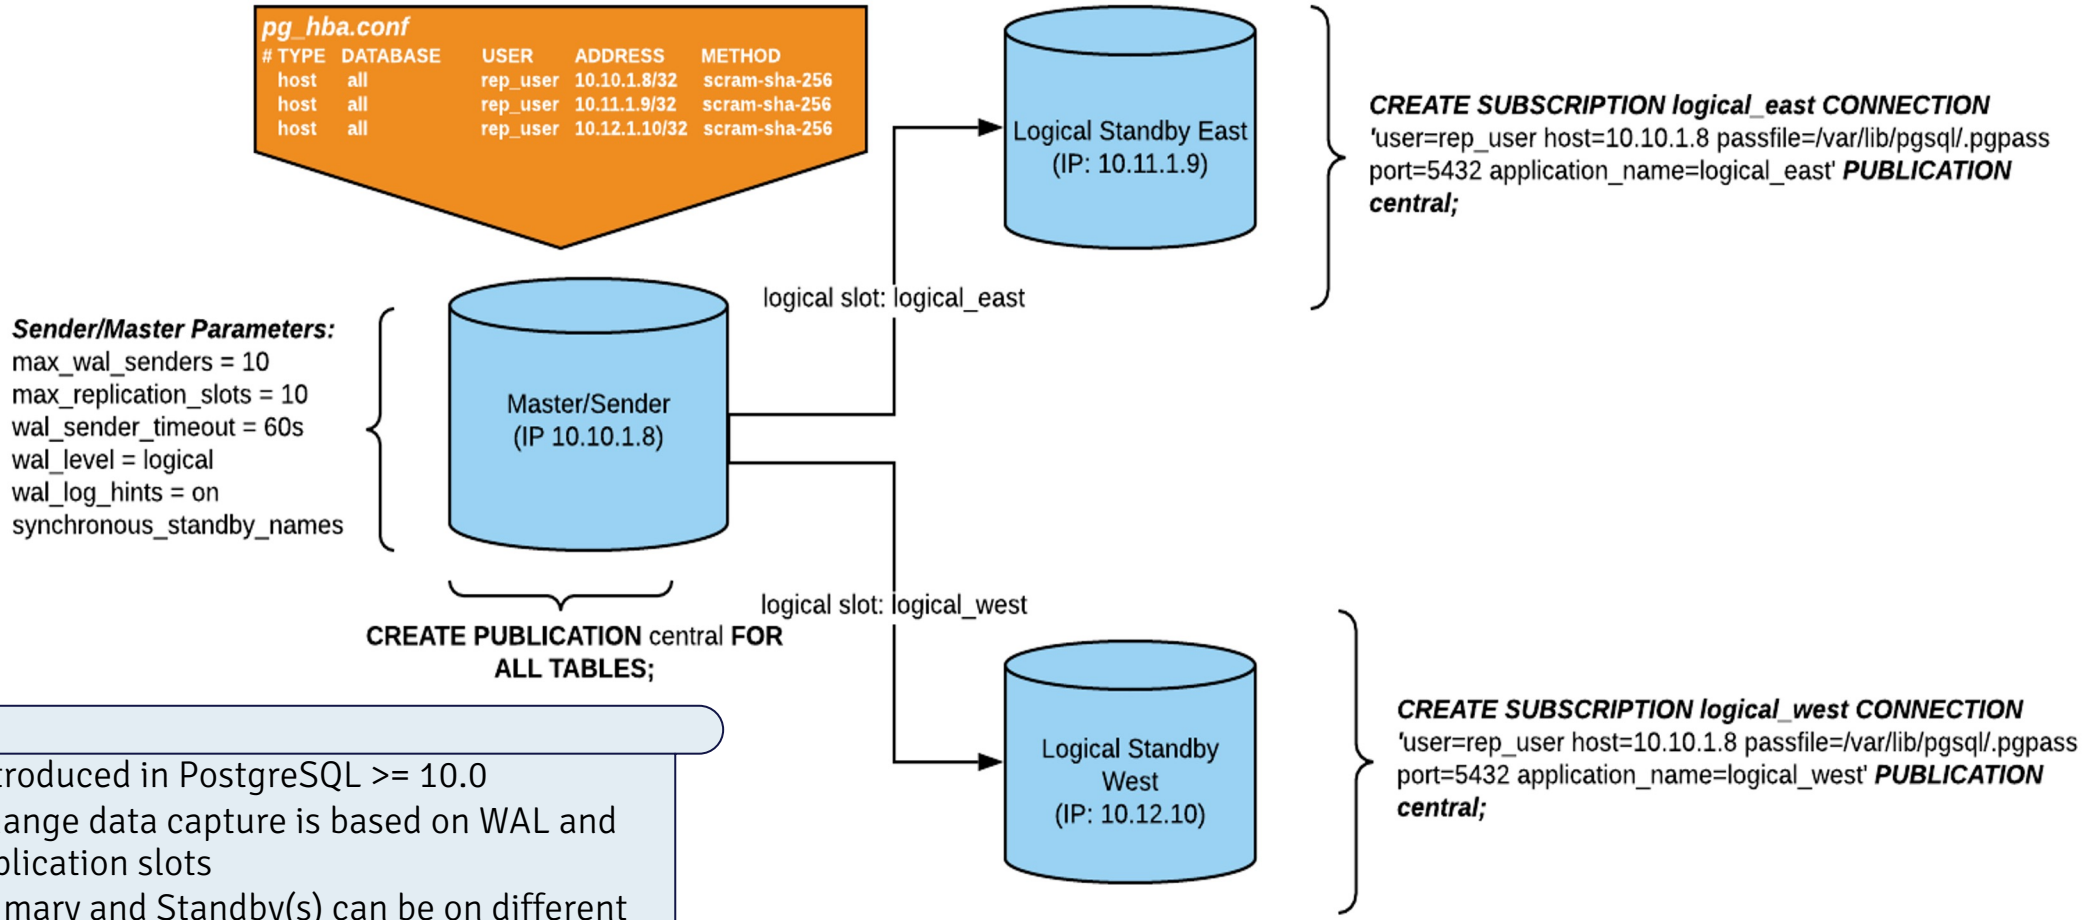

- Introduced in PostgreSQL >= 10.0
- Change data capture is based on WAL and replication slots
- Primary and Standby(s) can be on different minor and major versions.
- Selective table level replication is available.
- Doesn't support Large Object

Limitations of Postgres Native Logical Replication

A step towards future high availability

- No DDL replication
- No failover
- No integration with backup and recovery
- No built-in procedures for maintenance, updates and upgrades
- Unidirectional - no easy way to ‘fall back’

Introducing Postgres Distributed-BDR

Multi-master logical replication in a mesh network

- Automatic DDL and DML replication in an active-active mesh network
- Failover and switchover infrastructure to re-route traffic in case of failures or during the maintenance operation
- Advanced conflict detection and conflict management
- Management of distributed sequences
- Differentiated replication sets to control which data gets replicated and to which downstream databases.
- Cluster expansion/consolidation
- Rolling database software upgrades
- Rolling schema change/migration using cross-schema replication

Solutions to run 24x7 business with EDB

Postgres-BDR - Always On Platinum

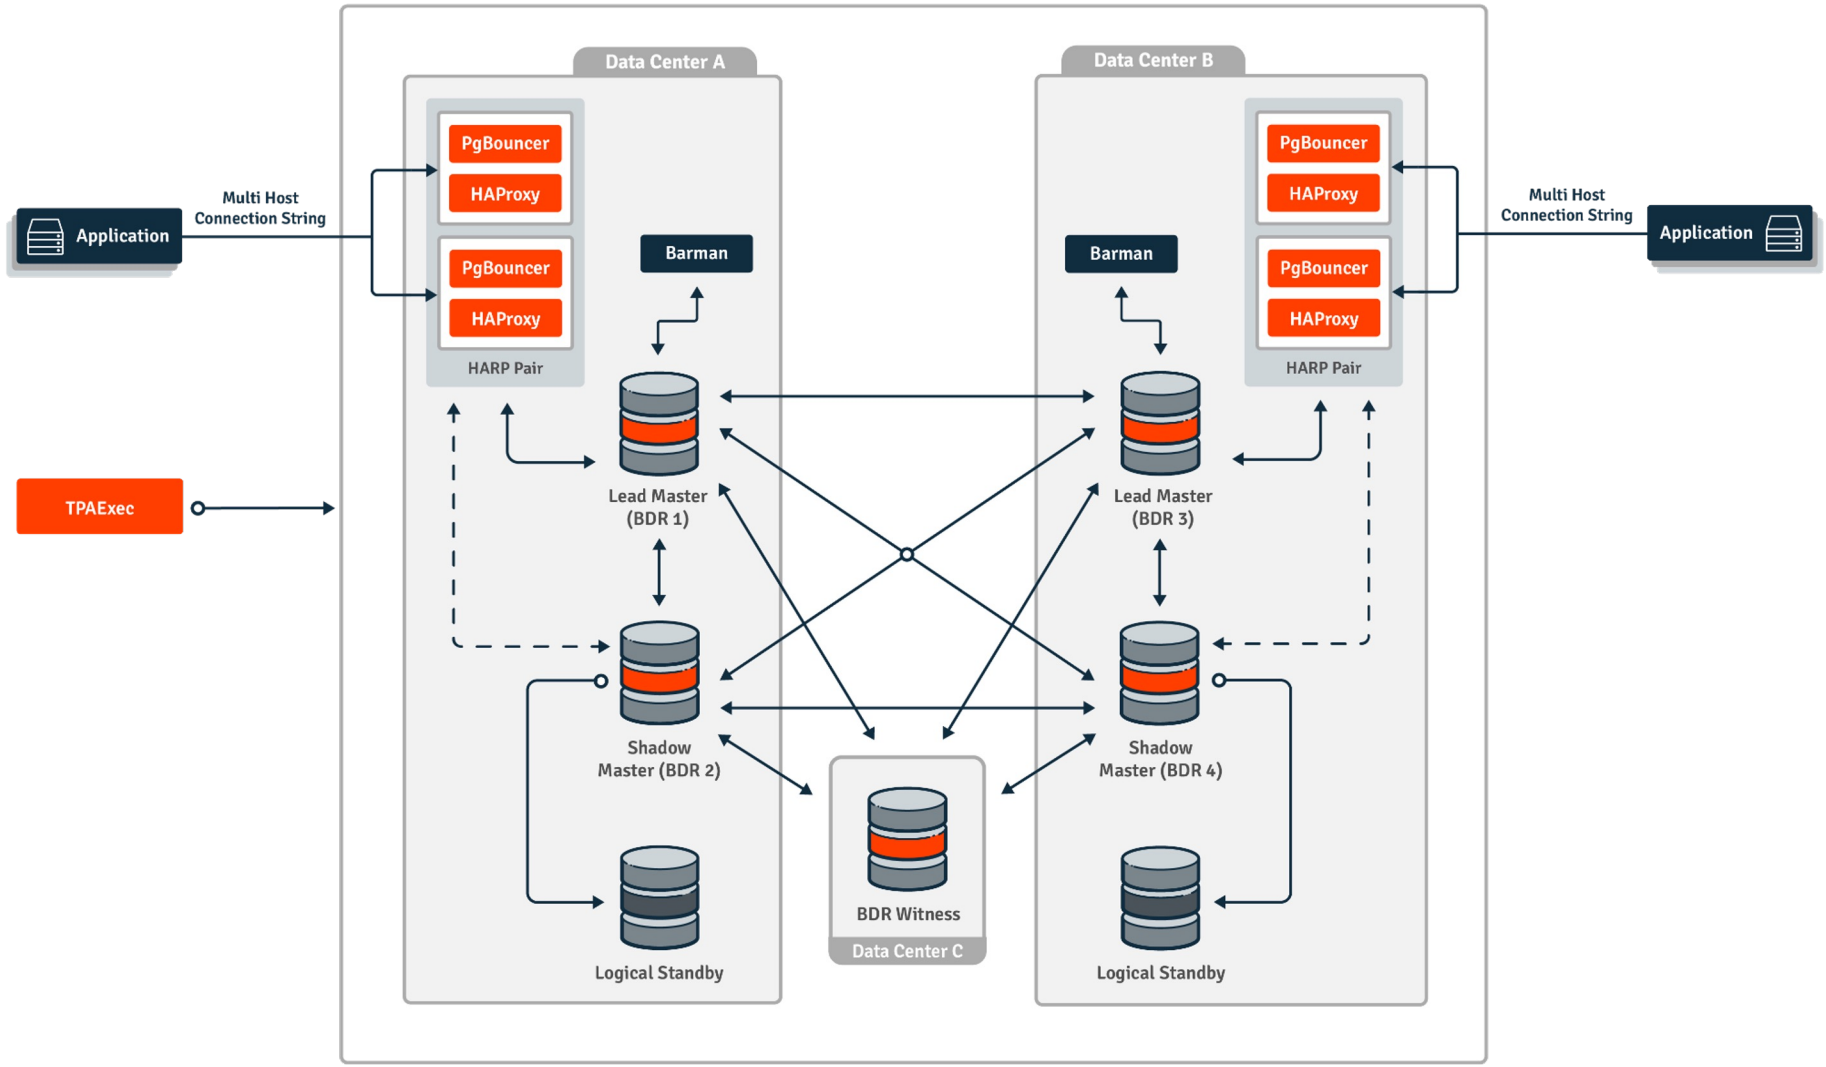

Always On Architectures

Always On Bronze (single active location)

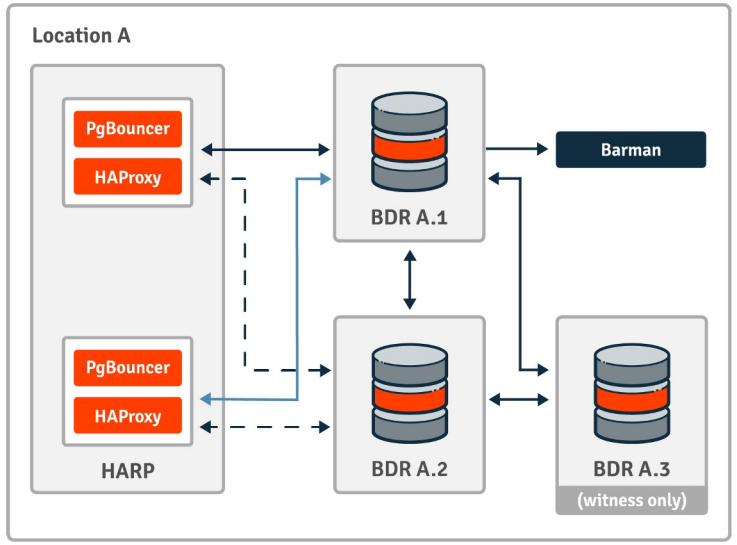

Always On Silver (single active location, backup in DR location)

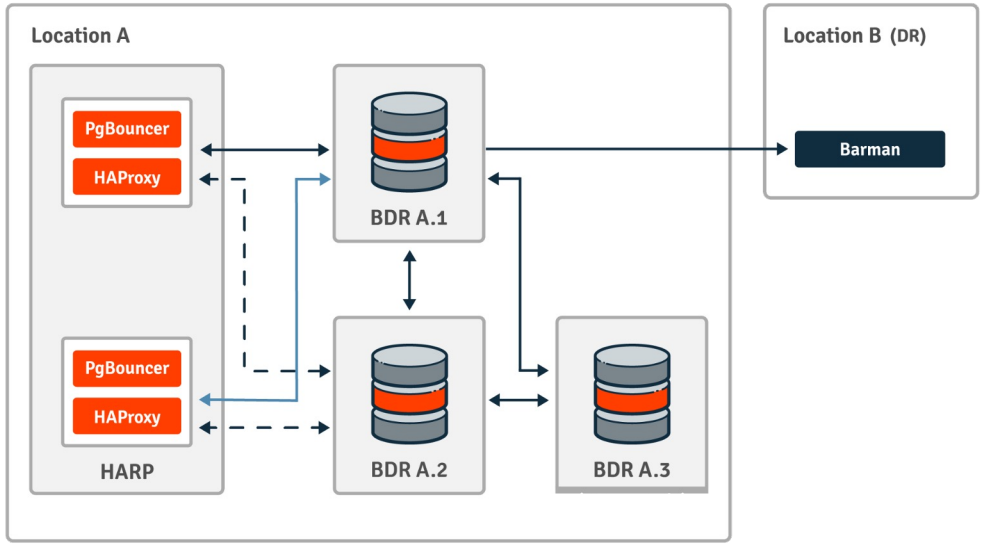

Always On Gold (two active location)

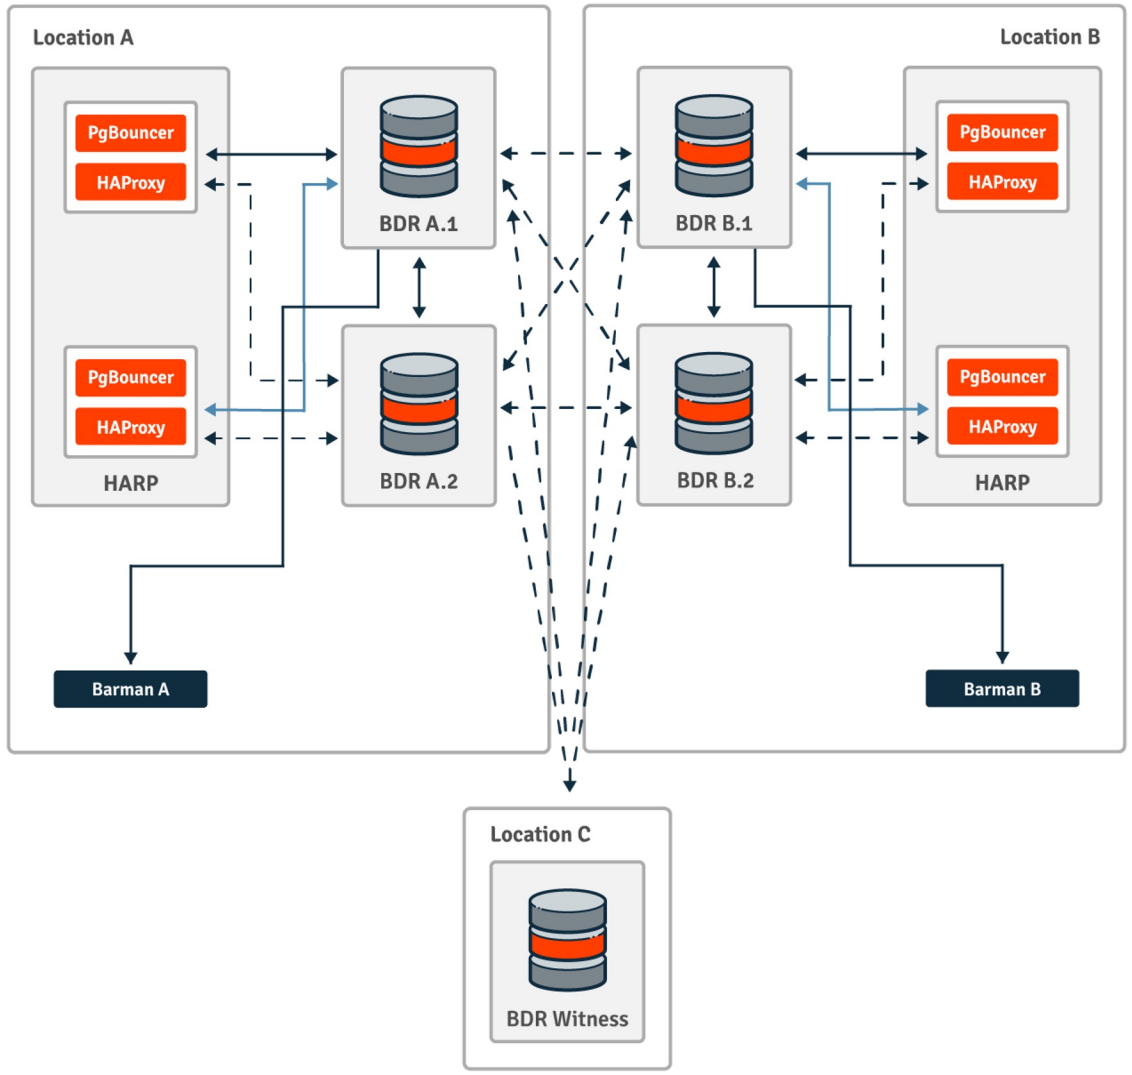

How to Choose the right BDR architecture

	Always On - Bronze	Always On - Silver	Always On - Gold	Always On - Platinum
Hardware failure protection	Yes	Yes	Yes	Yes
Location failure protection	No (unless Barman is moved offsite)	Yes - Recovery from backup	Yes - instant failover to fully functional site	Yes - instant failover to fully functional site
Failover to DR or Full second location	DR (if Barman is located offsite)	DR (if Barman is located offsite)	Full second location	Full second location
Zero downtime upgrade	Yes	Yes	Yes	Yes
Support of AZs in public/ private cloud	Yes	Yes	Yes	Yes
Fast local restoration of high availability after device failure	No; time to restore HA: (1) VM prov + (2) approx 60 min/500GB	Yes; three local BDR nodes allow to maintain HA after device failure	No; time to restore HA: (1) VM prov + (2) approx 60 min/500GB	Yes; logical standbys can quickly be promoted to full BDR nodes
Cross data center network traffic	No	Backup traffic only	Full replication traffic	Full replication traffic
BDR license cost	2 BDR nodes	3 BDR nodes	4 BDR nodes	4 BDR nodes 2 logical standbys

When to use what ?

Failover & Switchover benchmarking

How long does it take before the Postgres cluster is responding again?

BDR Always-On Silver

Scenario	Average	Minimum	Maximum
Switchover	296 ms	217 ms	453 ms
Postgres Crash	1,746 ms	1,546 ms	2,262 ms
System Crash	1,701 ms	1,041 ms	2,097 ms

EFM w. Streaming Replication

Scenario	Average	Minimum	Maximum
Switchover	12,000 ms	12,000 ms	12,000 ms
Postgres Crash	11,000 ms	11,000 ms	11,000 ms
System Crash	14,000 ms	13,000 ms	15,000 ms

Streaming replication or logical replication?

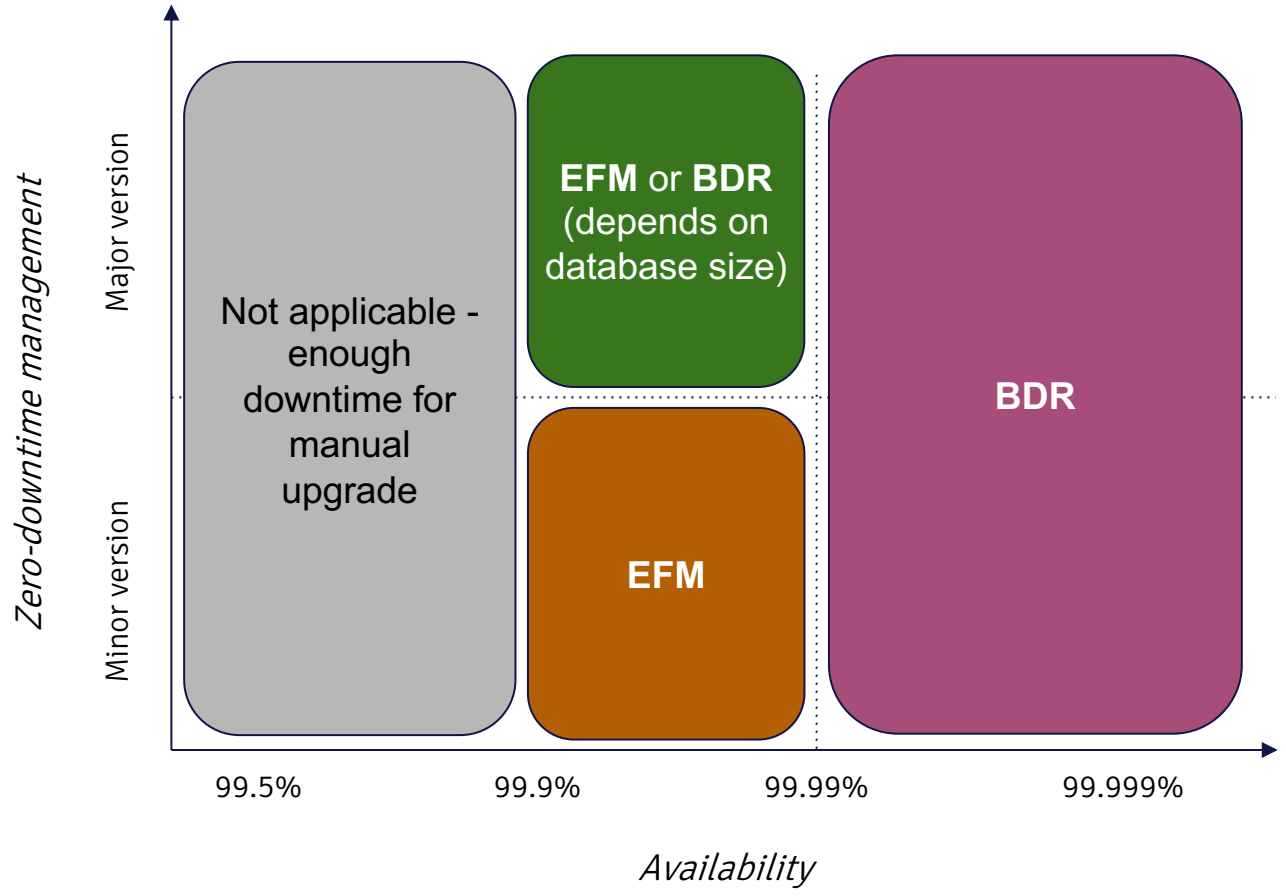

- Greater than 99.999% availability: Always use BDR
- Between 99.9% and 99.99%
  - EFM - if downtime for major version upgrade or reindex/ vacuum full does not exceed SLA
  - BDR - if downtime for major version upgrade or reindex/ vacuum full exceeds SLA

Learn more about Postgres-BDR

Whitepapers, webinars and conference presentations

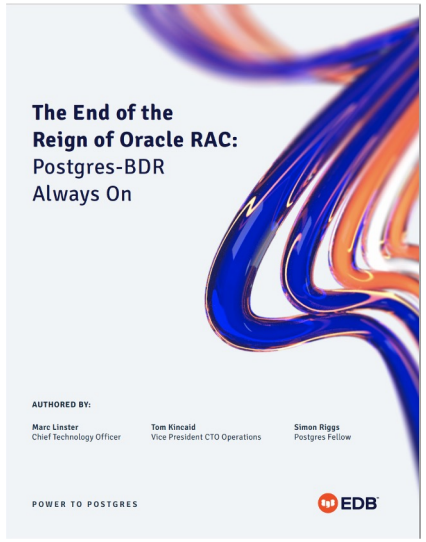

EDB Webinar:

- Upgrade easily with BDR; Simon Riggs; August 18 2021

Conference presentations:

- Building a business that never sleeps; Simon Riggs; Postgres Vision 2021
- A Case Study in “What If” PostgreSQL-BDR High-Availability/Disaster Recovery Analysis; Dominic Mortimer, Principal Software Engineer, ACI; 2Q Postgres Conference 2019

Blogs and Customer Success Stories

- How to Achieve Five Nines with Database Extreme High Availability: An Integral Part of Any Oracle Replacement Strategy, Jan Karremans · Sep 10, 2021
- 4 Reasons Why You Probably Don't Need RAC, Jan Karremans · Aug 24, 2021
- Project Management SaaS company uses Postgres-BDR to achieve consistent performance globally EDB Team · Aug 19, 2021
- Application High Availability and Resiliency: Steps to Improve Transaction Retry, Tom Kincaid Jul 30, 2021
- ACI Worldwide Modernizes Software Architecture While Reducing Risk and Lowering Costs with PostgreSQL; EDB Team · Mar 2, 2021

THANK YOU

(We can’t wait to see what you create)