



**EDB**

Postgres for the AI Generation

# EDB Postgres® Distributed

Ton Machielsen

Sales Engineer EMEA



Just solve it with Postgres






**EDB**  
Postgres for the AI Generation

# Agenda

- Introduction to EDB Postgres Distributed
- Deployment architectures
- Other considerations
- Demo



## EDB POSTGRES AI PLATFORM

### UNIFIED WORKLOAD MANAGEMENT

TRANSACTIONAL

ANALYTICAL

ARTIFICIAL INTELLIGENCE

### SINGLE-PANE-OF-GLASS ADMINISTRATION

HYBRID DATA ESTATE

INTELLIGENT OBSERVABILITY

ENTERPRISE SECURITY

### HYBRID AND MULTI-CLOUD DEPLOYMENT

PUBLIC CLOUD  
(MANAGED)

PRIVATE CLOUD  
(SOFTWARE)

ON PREMISES  
(APPLIANCE)

### PLATFORM TOOLS AND SERVICES

MIGRATION  
PORTAL

CONTINUOUS HIGH  
AVAILABILITY

BACKUP AND  
RECOVERY

### EXTENSIBILITY

CSP INTEGRATIONS

DEVOPS TOOLING

KUBERNETES TOOLING

GENAI & LLM INTEGRATIONS

LAKEHOUSE INTEGRATIONS

You're scaling to meet the needs of a global user base. Is your database prepared to scale with you?

Deploy multi-region clusters with five-nines availability to guarantee that data is consistent, timely, and complete – even during disruptions. This means an enhanced user experience and extended system capacity to deliver what you do best, whenever and wherever you operate.



# Highly available and geographically distributed

## Multi-master replication enables

- Logical replication of data and schema enabled via standard Postgres extension
- Robust tooling to manage conflicts, monitor performance, and validate consistency
- Deploy natively to cloud, virtual, or bare metal environments
- Geo-fencing, selectively replicating data for security compliance and jurisdiction control



# Benefits of multi-master replication for Postgres clusters

## Essential multi-master capabilities for Postgres clusters

- Application and database upgrades without requiring downtime
- Clusters row level eventual consistency by default
- Tools to monitor operation and verify data consistency
- 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 maintenance operation
- Dynamic cluster expansion/consolidation
- Improved security model with geo-fencing, selectively replicate data for security compliance and jurisdiction control

## Advanced conflict management, data-loss protection

- 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
- Replication lag control with asynchronous replication
- Faster, fully online major version upgrades
- Differentiated replication sets to control which data gets replicated and to which downstream databases



# Always on single-location

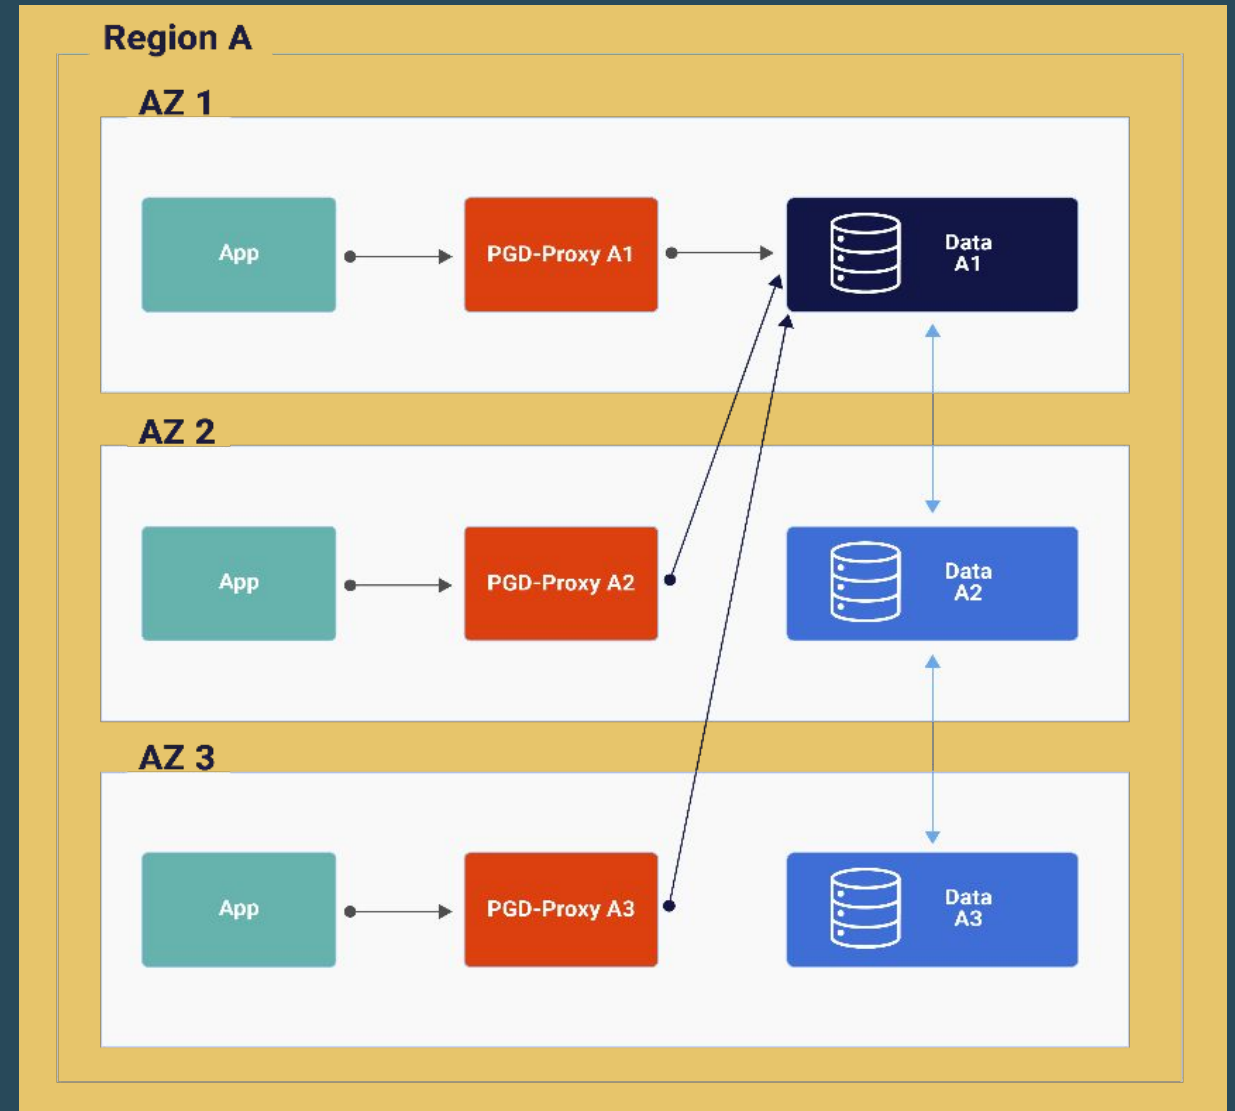
Redundant hardware to quickly restore from local failures

3 PGD nodes - could be 3 data nodes (recommended), or 2 data nodes and 1 witness which does not hold data (depicted)

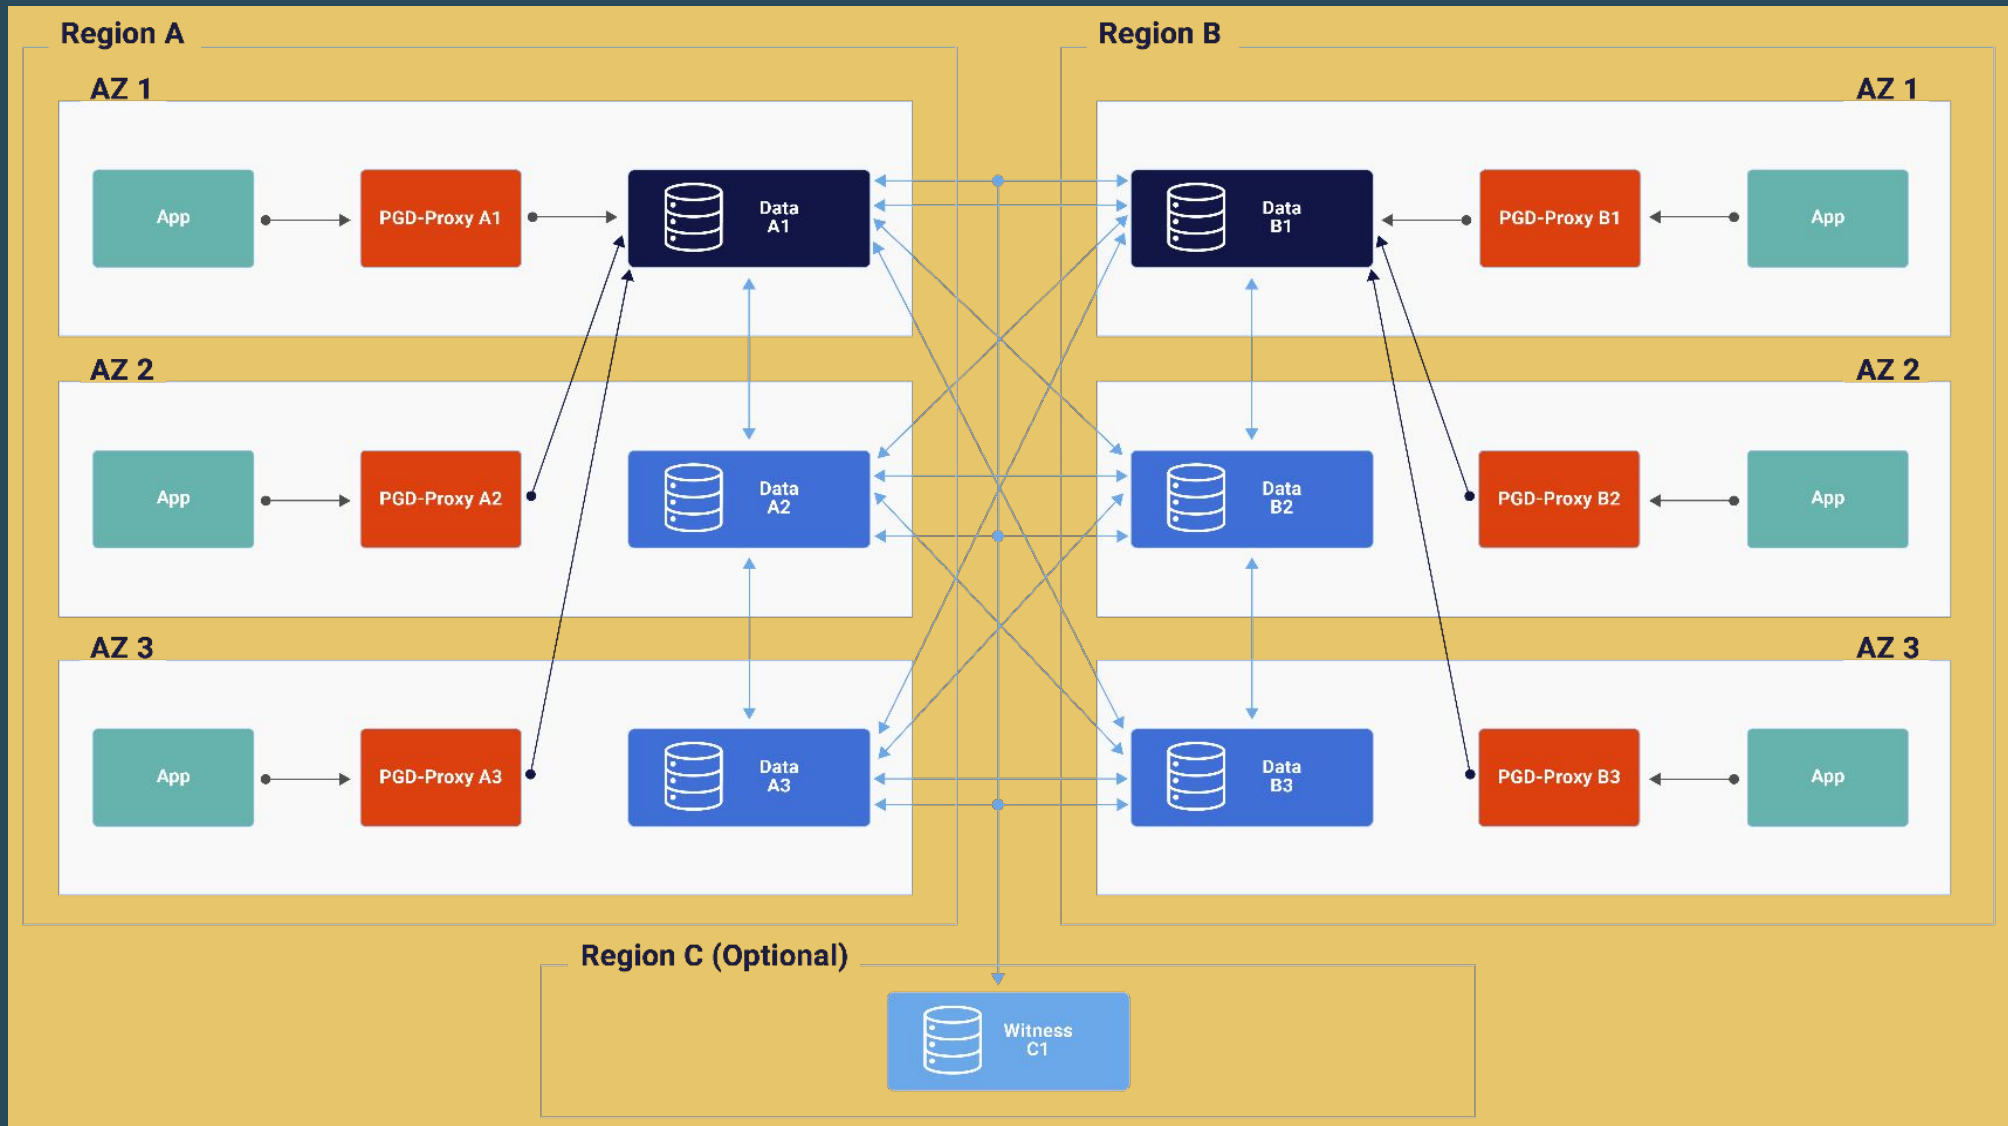
A PGD-Proxy for each data node with affinity to the applications - can be co-located with data node

Barman for backup and recovery (not depicted)

Postgres Enterprise Manager (PEM) for monitoring (not depicted)



# Always on multi-location (active/active or active/DR)





# Choosing your architecture

	Single Data Location	Two Data Locations	Two Data Locations + Witness	Three or More Data Locations
<b>Locations needed</b>	1	2	3	3
<b>Fast restoration of local HA in case of data node failure</b>	Yes - if 3 PGD data nodes No - if 2 PGD data nodes	Yes - if 3 PGD data nodes No - if 2 PGD data nodes	Yes - if 3 PGD data nodes No - if 2 PGD data nodes	Yes - if 3 PGD data nodes No - if 2 PGD data nodes
<b>Location failure protection for data</b>	No - unless offsite backup	Yes	Yes	Yes
<b>Global consensus in case of location failure</b>	N/A	No	Yes	Yes
<b>Data restore required after location failure</b>	Yes	No	No	No
<b>Immediate failover in case of location failure</b>	No - requires data restore from backup	Yes - alternate Location	Yes - alternate Location	Yes - alternate Location
<b>Cross location network traffic</b>	Only if offsite backup	Full replication traffic	Full replication traffic	Full replication traffic
<b>License cost</b>	2 or 3 PGD data nodes	4 or 6 PGD data nodes	4 or 6 PGD data nodes	6+ PGD data nodes



# Considerations

Postgres data nodes, Barman nodes, and PGD proxy nodes can be deployed on virtual machines or in a bare metal deployment mode. However, anti-affinity between data nodes and PGD Proxy nodes needs to be maintained:

- Multiple data nodes should not be on VMs that are co-located on the same physical hardware, as that reduces resiliency.
- Multiple PGD proxy nodes should not be on VMs that are co-located on the same physical hardware, as that also reduces resiliency.
- For better performance, users should ensure proper CPU and memory resource assignment when co-locating PG proxy nodes with data nodes on the same hardware.



# Demo time!!

## Automatic rolling system upgrades

- Works across major software releases of OS, database server, and EDB Postgres Distributed
- Major upgrades in place, not needing extra hardware
- Can be manually controlled, if desired
- Inter-node protocols are renegotiated for each link, all protocols backwards compatible

## Resolves the largest source of downtime

- Security/maintenance releases – 4 times per year
- Unscheduled bugs ~1 per year
- Parameter changes, hardware changes also

