EDB Postgres Distributed

The most advanced replication solution available for Postgres

Gianni Ciolli - VP, Solutions Architecture

Milano, 7 October 2022

EDB™

Content

- Do you need to be always on?
- Introduction to EDB Postgres Distributed
- Next steps

2022 Copyright © EnterpriseDB Corporation All Rights Reserved

Do you need to
be “always on”?

What is “always on?”

Delivering mission-critical applications/services 24x7

Finance

Payments

Bank account access

Telecom

Video conferencing

Texting/alerting

Transportation

Travel

Rideshare

GPS

2022 Copyright © EnterpriseDB Corporation All Rights Reserved

Why "always on?"

Why do some applications need to be "always on?"

Top-tier enterprise applications are critical to an organization's success in all regions where business is conducted, whether a single region or globally

The application represents a promise to its customers

The availability of the application directly ties to revenue generation

The application must perform well for a good user experience

The application data must always be current and available or the user loses trust

2022 Copyright © EnterpriseDB Corporation All Rights Reserved

How “always on” impacts your business

Immediate impact

- How much does downtime cost your business per second/minute/hour?
- Will users go elsewhere as soon as you’re down?
- Will there be regulatory consequences?

Aftermath

- Lost customers: can you win them back?
- Will competitors go after your remaining users?
- Reputation: will bad word-of-mouth slow down new user acquisition?

2022 Copyright © EnterpriseDB Corporation All Rights Reserved

Do you need to be “always on”?

Things to consider

What’s the reputation cost of downtime to your business?

Are there times when it’s okay for an application to be inaccessible?

Is access to data tied directly to revenue?

Are your customers globally distributed?

2022 Copyright © EnterpriseDB Corporation All Rights Reserved

A real world example

Imagine a brand that may have 1,000s of locations and eCommerce with payment processing

Processing 100 business transactions per second

And the average transaction is $85

If there is an outage/unplanned downtime of 5 minutes...

The business will miss out on 30,000 transactions ...

And be out **$2.5M** in revenue.

This disappoints 30,000 customers directly, plus the friends and family they talk to.

2022 Copyright © EnterpriseDB Corporation All Rights Reserved

EDB Postgres Distributed

Customer Success Story

Use case	SaaS project management and collaboration company experiencing massive growth with customers spanning individuals to large enterprises supporting over 100,000 teams globally.
The need	Provide consistent performance to their international customer base while scaling up their business in multiple regions.
Why EDB Postgres Distributed	ClickUp selected EDB Postgres Distributed to address geo-distributed needs of their customer base, and later upgraded to gain improved performance, rolling upgrades, and high availability.
Value	EDB Postgres Distributed has contributed to the 99.99% uptime for 12 consecutive months of ClickUp's 7 node global deployment.

ClickUp

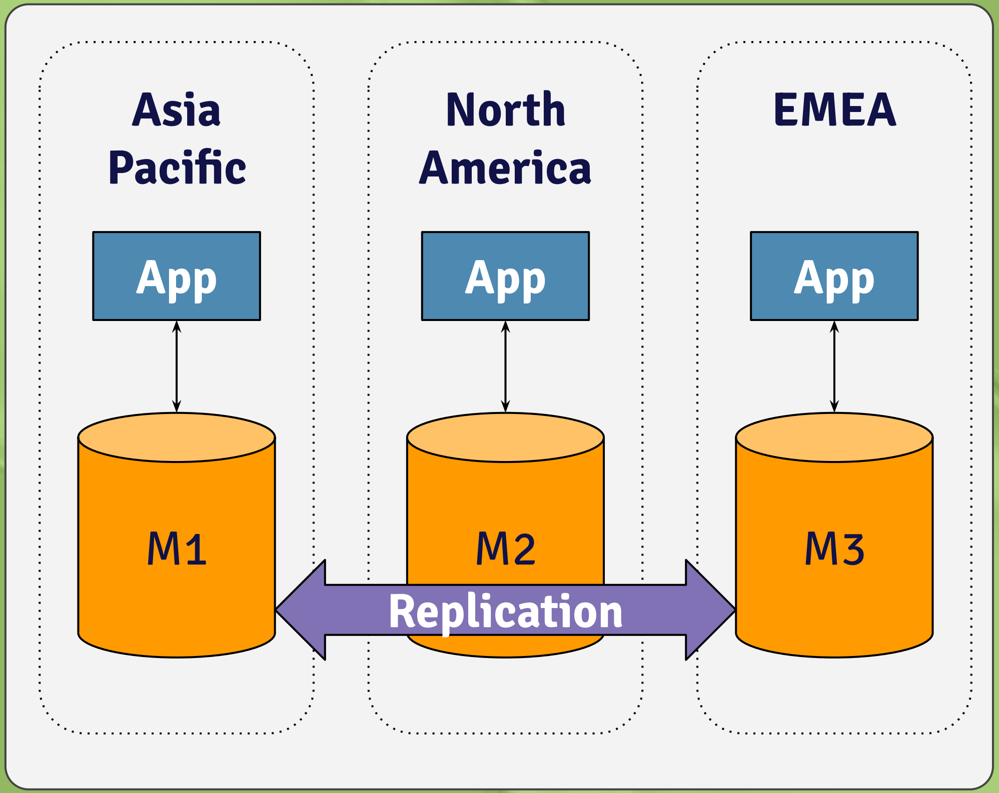

2022 Copyright © EnterpriseDB Corporation All Rights Reserved

Introduction to
EDB Postgres Distributed

- You’ve chosen Postgres
- Your application needs to be always on
- Any downtime = Loss of business continuity = Revenue loss

You need:

- A **replication solution**
- **Geographically distributed** database clusters
- **Extreme high availability**
- **Optimal performance** at all times
- **Advanced monitoring**

Enter EDB Postgres Distributed

The most advanced replication solution for Postgres

Maintain extreme high availability

Postgres clusters deployed with EDB Postgres Distributed keep top tier enterprise applications running

Upgrade with near zero downtime

Rolling upgrades of application and database software eliminate the largest source of downtime

Choose the level of consistency

Robust capabilities provide flexibility to meet application data loss requirements

2022 Copyright © EnterpriseDB Corporation All Rights Reserved

EDB Postgres Distributed is more than bi-directional replication

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

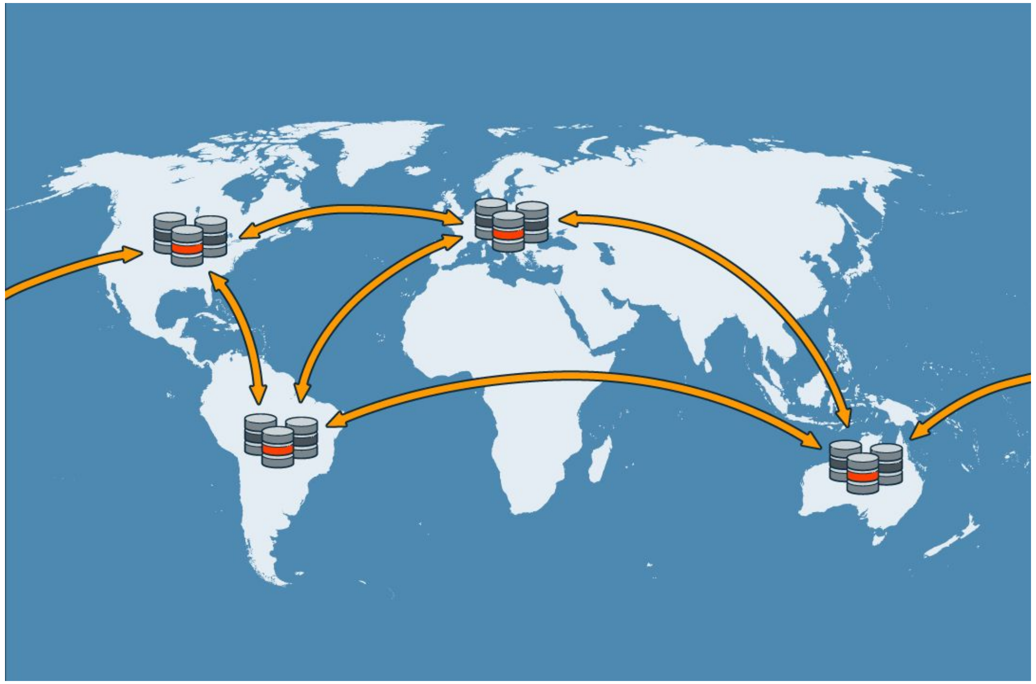

- Logical replication of data and schema enabled via standard Postgres extension
- Data consistency options that span from immediate to eventual consistency
- Robust tooling to manage conflicts, monitor performance, and validate consistency
- Deploy natively to cloud, virtual, or bare metal environments
- Geo-fencing, allowing selectively replicate data for security compliance and jurisdiction control.

2022 Copyright © EnterpriseDB Corporation All Rights Reserved

EDB Postgres Distributed feature overview

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

Essentials

Provides the essential multi-master capabilities for Postgres clusters

- Enables application and database upgrades without requiring downtime
- Provides clusters row level eventual consistency by default
- Concurrent streaming of large transactions smooths replication and reduces lag
- Tools to monitor operation and verify data consistency

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

2022 Copyright © EnterpriseDB Corporation All Rights Reserved

New in EDB Postgres Distributed 4.1

Reliability and Operability

Faster, fully online major version upgrades*

Performance and Scalability

Replication lag control with asynchronous replication*

High Availability and Recovery

More flexibility for durable replication with Group Commit

User Experience

Greater control with command line interface

User Experience

Simplified synchronous replication configuration with SQL level interface

User Experience

Automatic management of distributed sequences

2022 Copyright © EnterpriseDB Corporation All Rights Reserved

***Key differentiators and enhanced competitive advantages with v4.1**

Additional features in EDB Postgres Distributed

- 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
- Advanced conflict detection and conflict management
- Differentiated replication sets to control which data gets replicated and to which downstream databases
- Cluster expansion/consolidation
- Rolling schema change/migration using cross-schema replication
- Recovery from user error through solid integration with backup and recovery tools
- Improved security model with Geo-fencing, allowing selectively replicate data for security compliance and jurisdiction control

2022 Copyright © EnterpriseDB Corporation All Rights Reserved

Best fit

Where does EDB Postgres Distributed fit in EDB's portfolio?

Requirement / Characteristic	Good fit for EDB Postgres Distributed	Fit for another EDB solution
Availability Requirements	99.99 - 99.999% Four to Five 9s	99.9 - 99.99% Three to Four 9s
Data Centers & Geographic Regions	Globally distributed	Not globally distributed
Upgrade downtime	Rolling upgrades with near-zero downtime	Appetite for maintenance windows
Application tier	"Tier 1"	"Tier 2" & "Tier 3"
Application types	Payment gateways, telecommunications call routing, global collaboration	HR, Expense Reporting, CRM

2022 Copyright © EnterpriseDB Corporation All Rights Reserved

Two primary use cases

The most advanced replication solution for Postgres

1. Extreme High Availability

Postgres clusters deployed with EDB Postgres Distributed keep top tier enterprise applications running

2. Geographically Distributed Workloads

Global enterprise applications can address data sovereignty and stay in sync with EDB Postgres Distributed

2022 Copyright © EnterpriseDB Corporation All Rights Reserved

1. Extreme High Availability

Challenge: uptime

Your application needs **five 9s** availability to maintain and grow revenue.

Those **9s** can be stolen by:

1. upgrades
2. hardware failures
3. software failures
4. data center outages
5. operational errors

Solution: EDB Postgres Distributed

- EDB Postgres Distributed protects against the number one enemy of five 9s: **upgrades**
- Multi-master architecture provides **~2 second failover**
- Designed for **conflict handling**
- **Distributed transaction handling** to protect against data loss

2022 Copyright © EnterpriseDB Corporation All Rights Reserved

2. Geographically Distributed Workloads

Challenge: conflicts

Read/write transactions for a wide area user base means challenges with latency, performance, and replication

- Need to resolve multi-location data conflicts
- Requires design considerations
- Likely requires application changes
- Different paradigm than many existing applications

Solution: EDB Postgres Distributed

- Advanced column-level conflict handling and conflict-free replicated types (CRDTs)
- Keep regional data local for data sovereignty
- Designed for conflict handling

2022 Copyright © EnterpriseDB Corporation All Rights Reserved

EDB Postgres Distributed Summary

Providing what your application needs to be “always on”

Provides the essential multi-master capabilities for Postgres clusters

Advanced data-loss protection ensures transaction idempotence

Advanced conflict-handling scenarios

Up to 5X faster replication throughput

2022 Copyright © EnterpriseDB Corporation All Rights Reserved

Grazie!

2022 Copyright © EnterpriseDB Corporation All Rights Reserved