



Securing Postgres with Streaming Replication

PROUD CONTRIBUTOR TO



PostgreSQL
the world's most advanced open source database

Jan Karremans

Head of Sales

- Techie in Sales
- 35 years of database expertise
- 8 years of databases in Kubernetes



TRAINED



Why CYBERTEC?



Pure-play
PostgreSQL
Company



Specialized,
International and
Global Team



25 years as a
Key Contributor to
the PostgreSQL
Community




Product &
Service Around
PostgreSQL



Fundamentally
Independently
Operated



PostgreSQL is at the heart of everything we do, and we are proud to actively contribute to its community and projects.



Agenda

1. What Is Streaming Replication?

And why does it even matter?

2. Single Elephant To A Herd

Scaling PostgreSQL safely (nearly) to infinity

3. And This Logical Replication Thing?

What is it and why does that matter?

4. What Does This Mean?

Even the basic functions make a difference



Agenda

1. What Is Streaming Replication?

And why does it even matter?

2. Single Elephant To A Herd

Scaling PostgreSQL safely (nearly) to infinity

3. And This Logical Replication Thing?

What is it and why does that matter?

4. What Does This Mean?

Even the basic functions make a difference





**The world
runs on data**





**and you shouldn't
miss a speck**



Data Lives In Transactions

Change Is The Default

- Secure the changes
- “Replay” the changes
- Internal and external
- Dead and alive



Hans-jürgen Schöning & Simon Riggs

- Protocolled foundations for streaming replication
- Implemented streaming replication



Transaction Log

PRIMARY

Transaction



Engine



REPLICA



Reciever



WAL Log



Data Files



WAL Log



Data Files



Characteristics

- Binary compatible
- Works out-of-the-box
- Rock solid and super secure



Agenda

1. What Is Streaming Replication?

And why does it even matter?

2. Single Elephant To A Herd

Scaling PostgreSQL safely (nearly) to infinity

3. And This Logical Replication Thing?

What is it and why does that matter?


4. What Does This Mean?

Even the basic functions make a difference



Single Lonely Cluster



- Crash recovery
- Archive recovery




Single Lonely Cluster

With Some Support

- WAL file copy
- Transaction streaming
`pg_receivewal`



Primary And Replica



● Primary node



Async & Sync




Async & Sync - Dangers

synchronous commit setting	local durable commit	standby durable commit after PG crash	standby durable commit after OS crash	standby query consistency
remote_apply	X	X	X	X
on	X	X	X	
remote_write	X	X		
local	X			
off				

**CAUTION:
BIAS!**



Primary And Replicas




● Primary node




Primary And Replicas


`synchronous_standby_names = 'FIRST 2 (S1, S2, S3)'`



● Primary node



Primary And Cascading Replicas



● Primary node



More Tools



PGBOUNCER



Agenda

1. What Is Streaming Replication?

And why does it even matter?

2. Single Elephant To A Herd

Scaling PostgreSQL safely (nearly) to infinity

3. And This Logical Replication Thing?

What is it and why does that matter?

4. What Does This Mean?

Even the basic functions make a difference



What?

Physical


- Based in bits and bytes
- Rock solid and super secure
- Meant for technical resilience
- Seamless

Logical

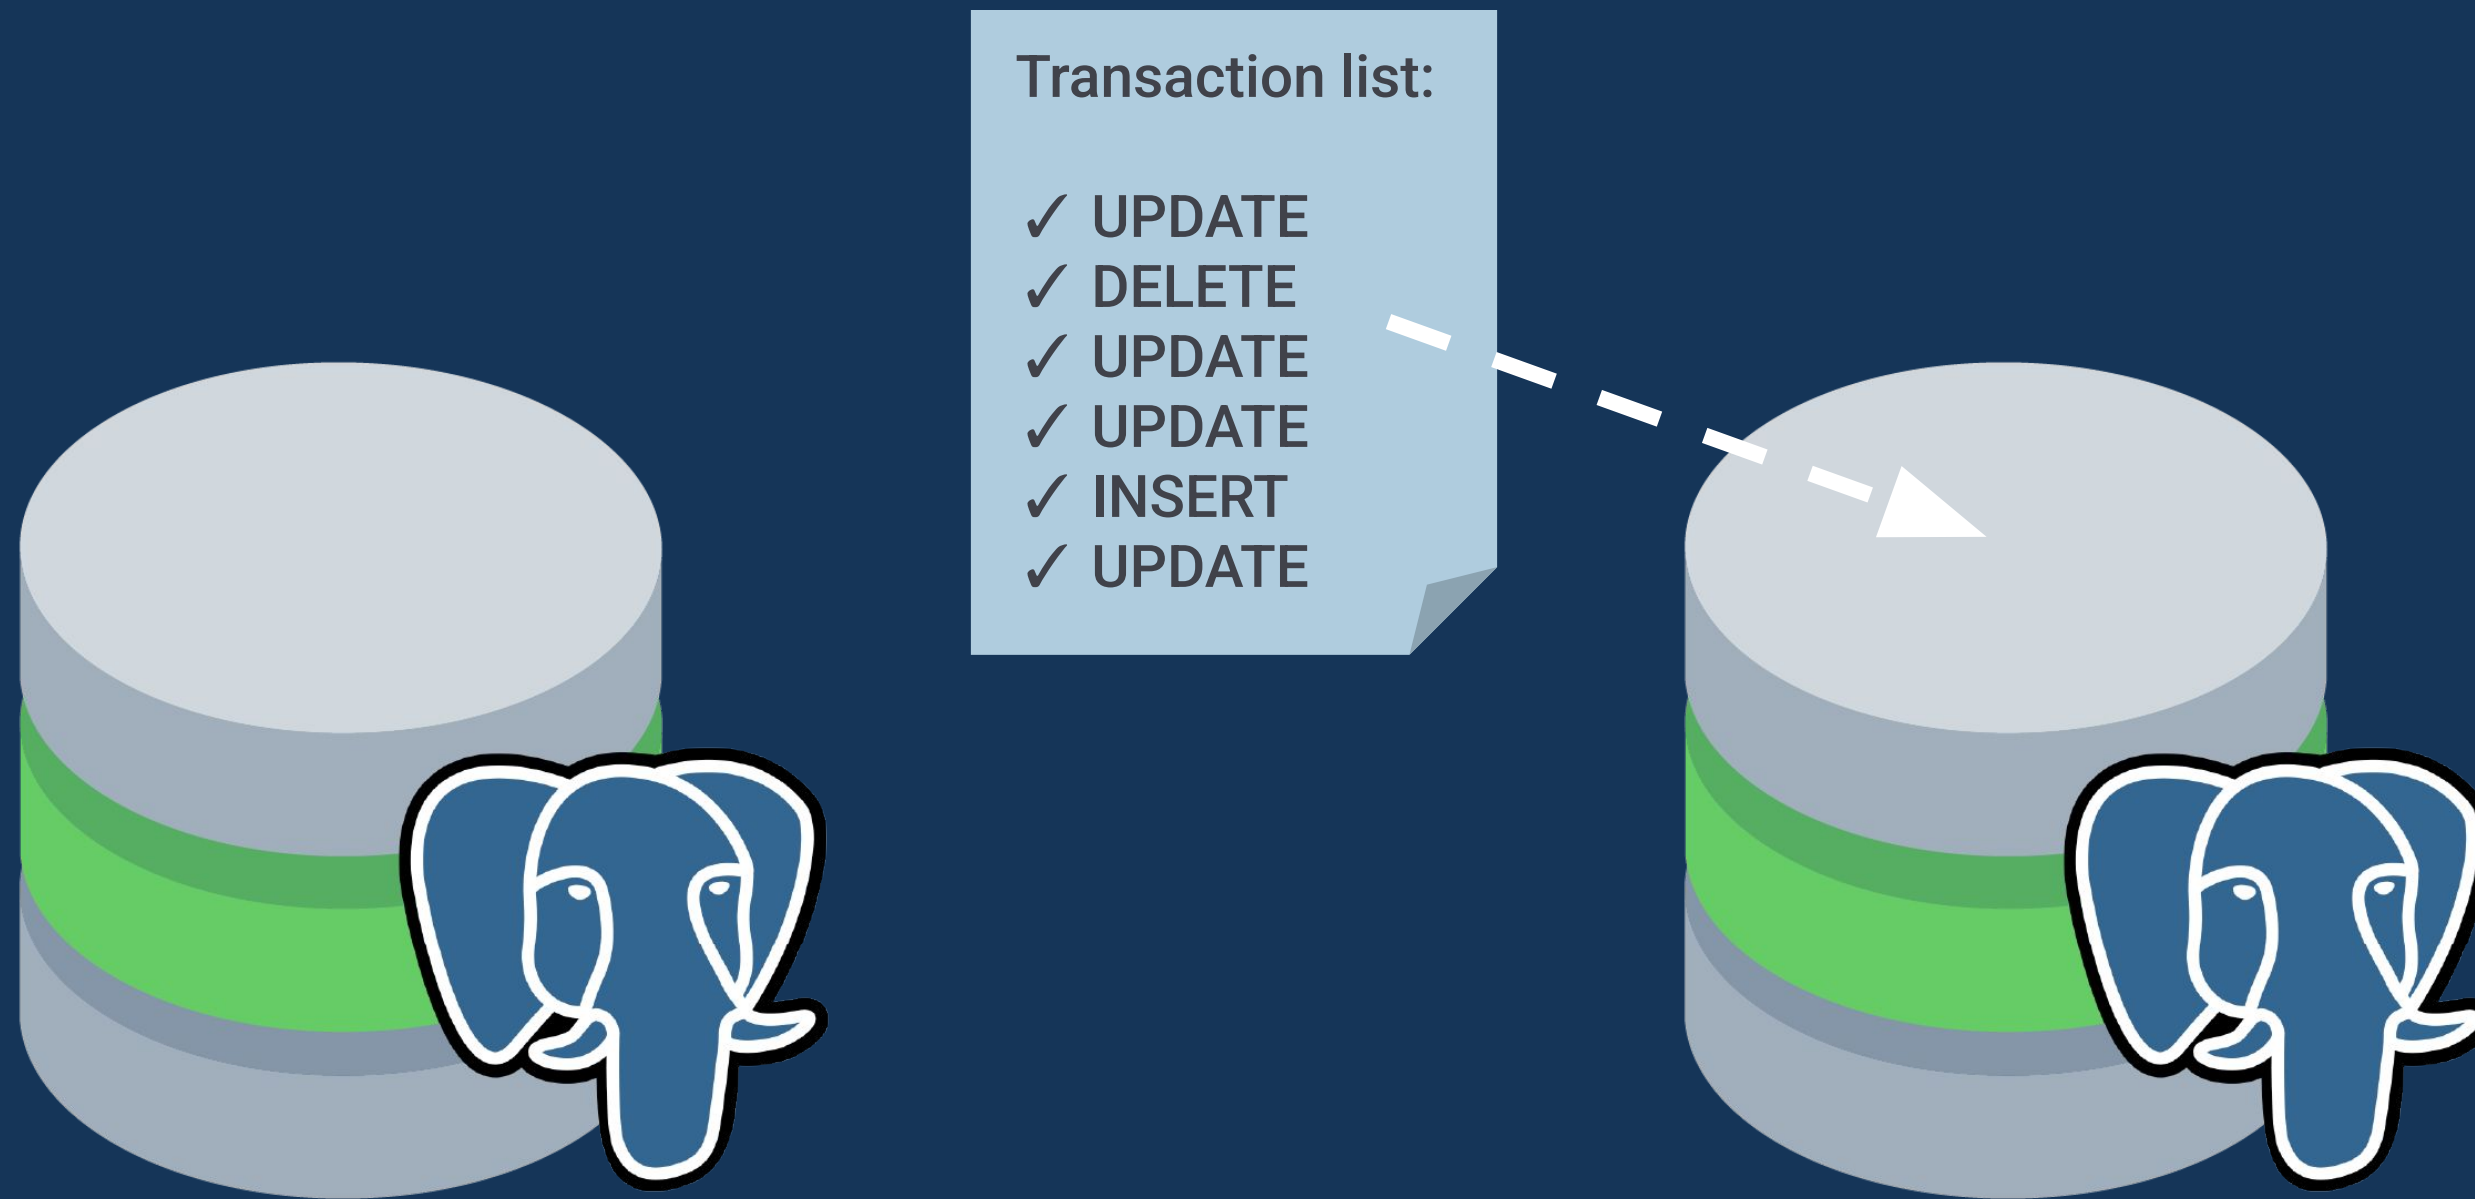
- Based on transactions
- Rock solid, but with a twist
- Meant for data "manipulation"
- Has requirements



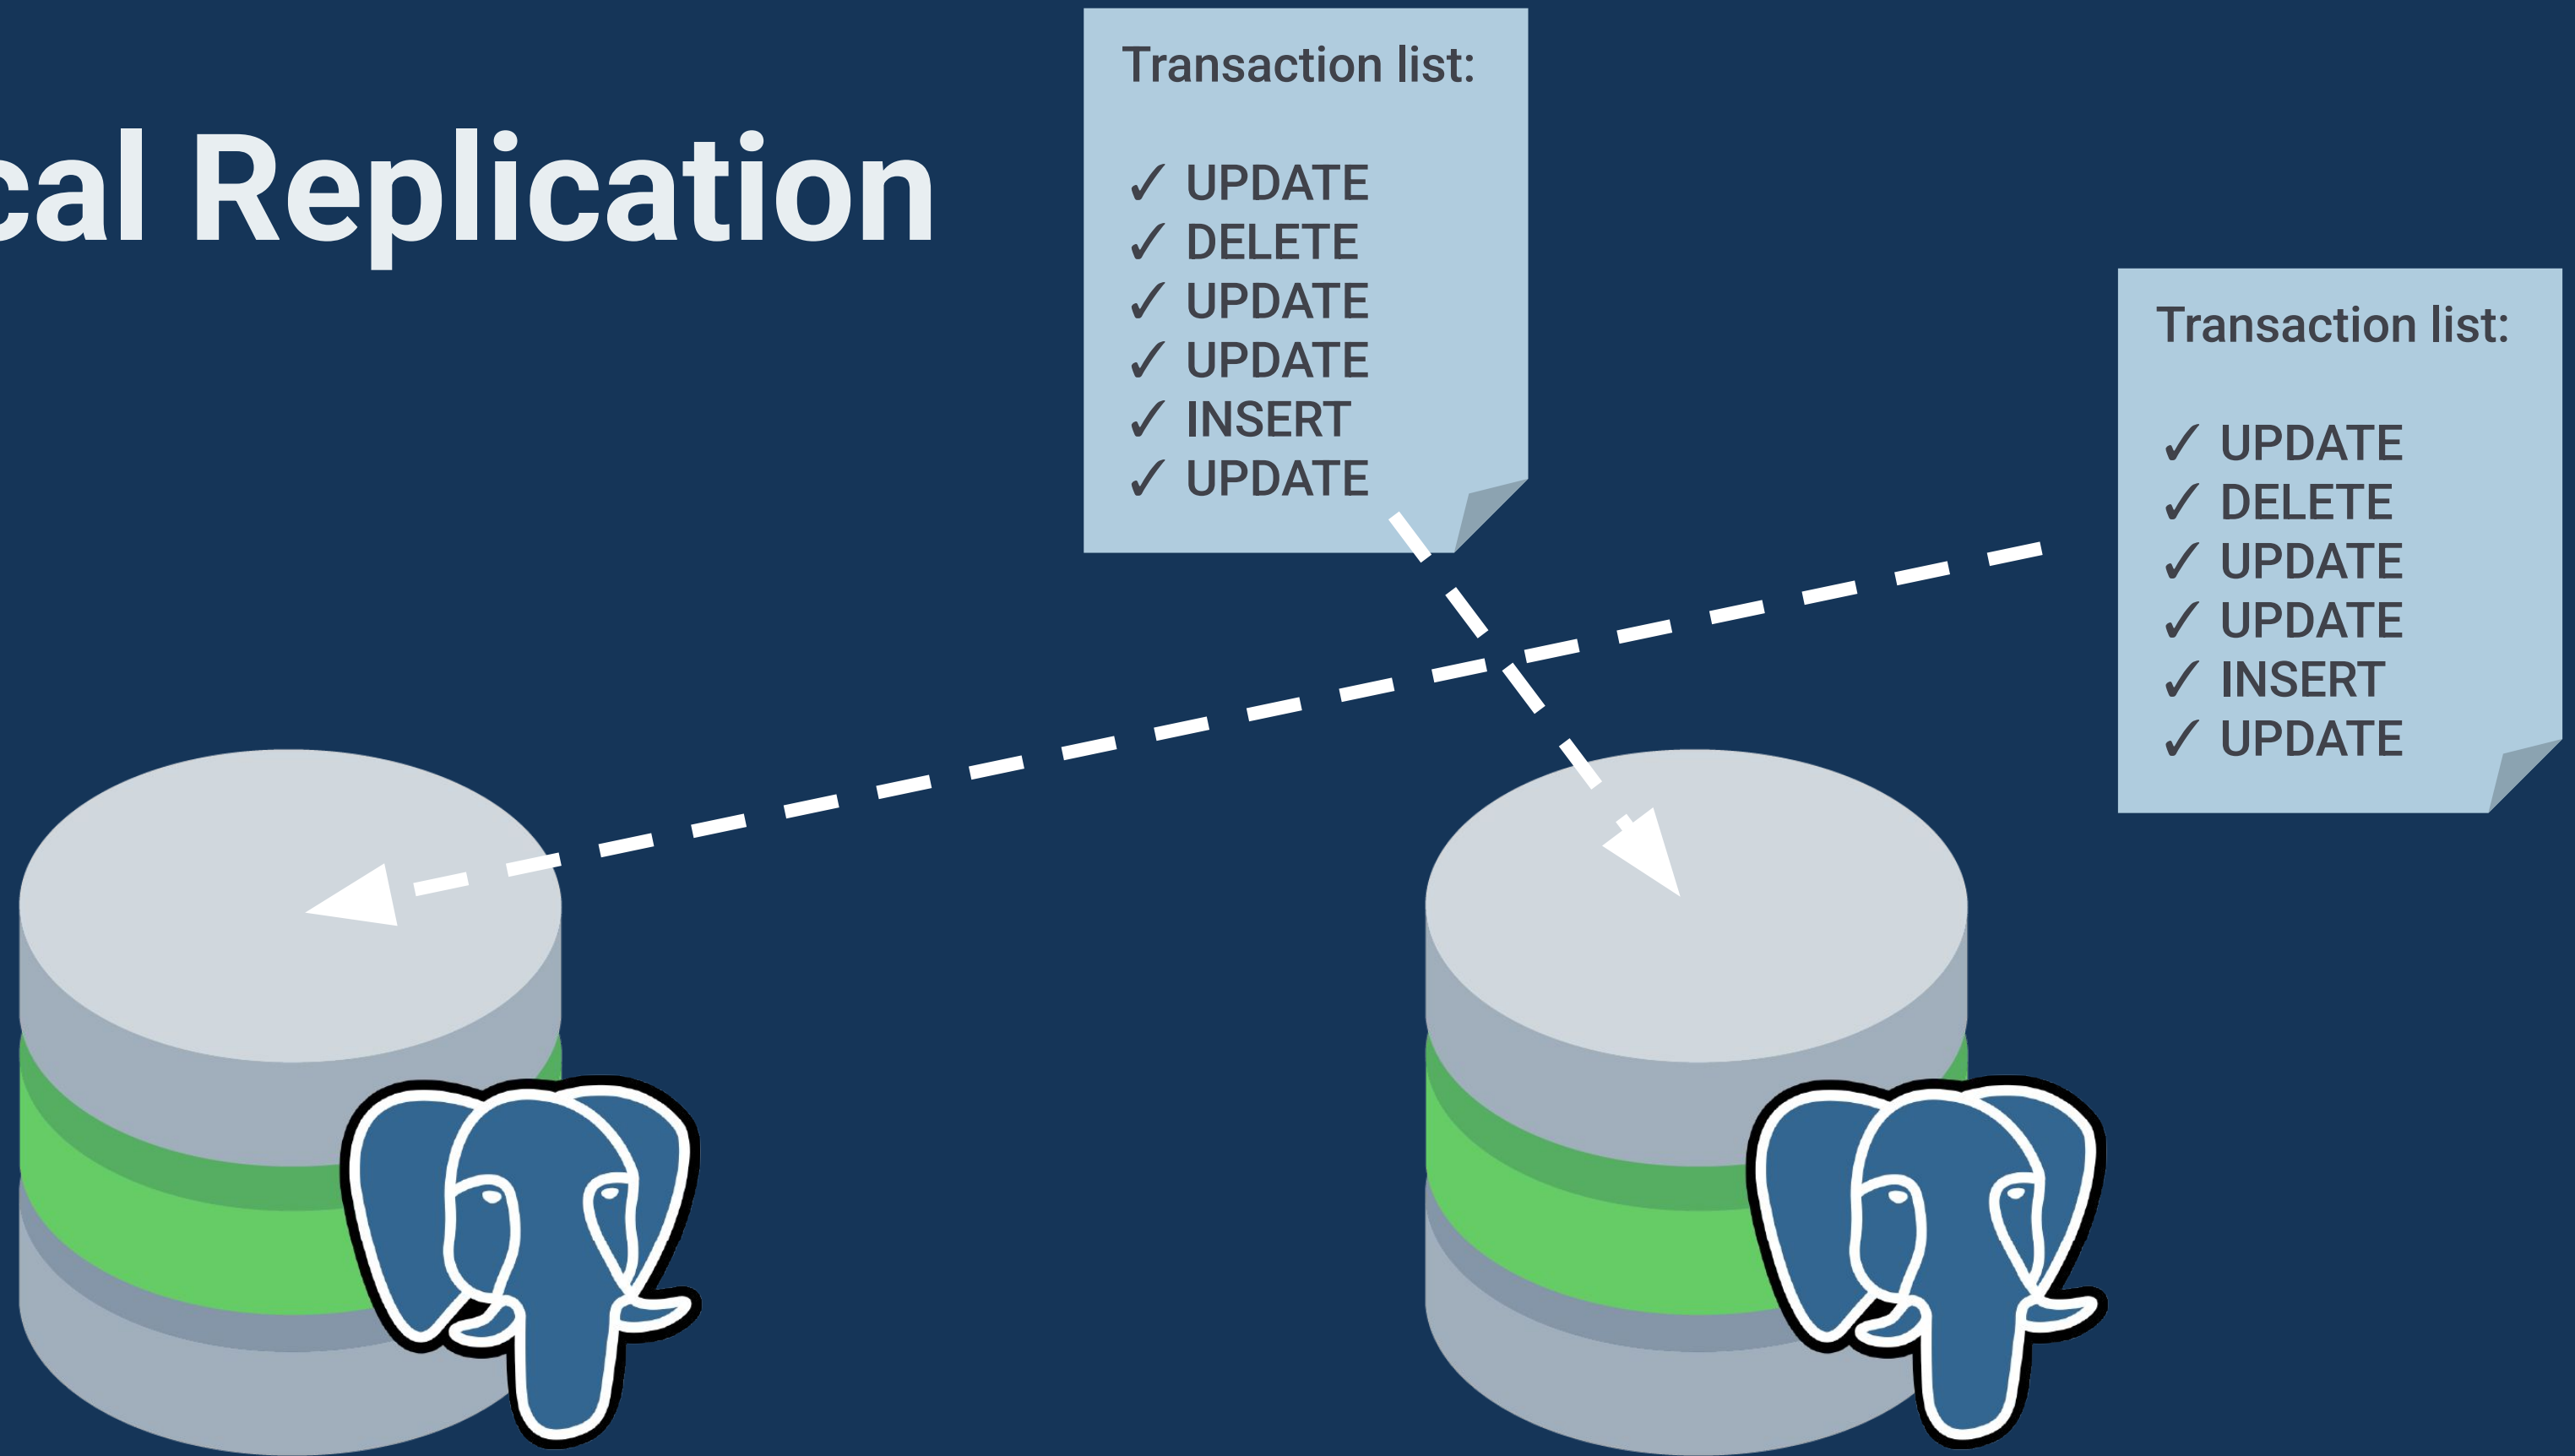
Streaming Replication



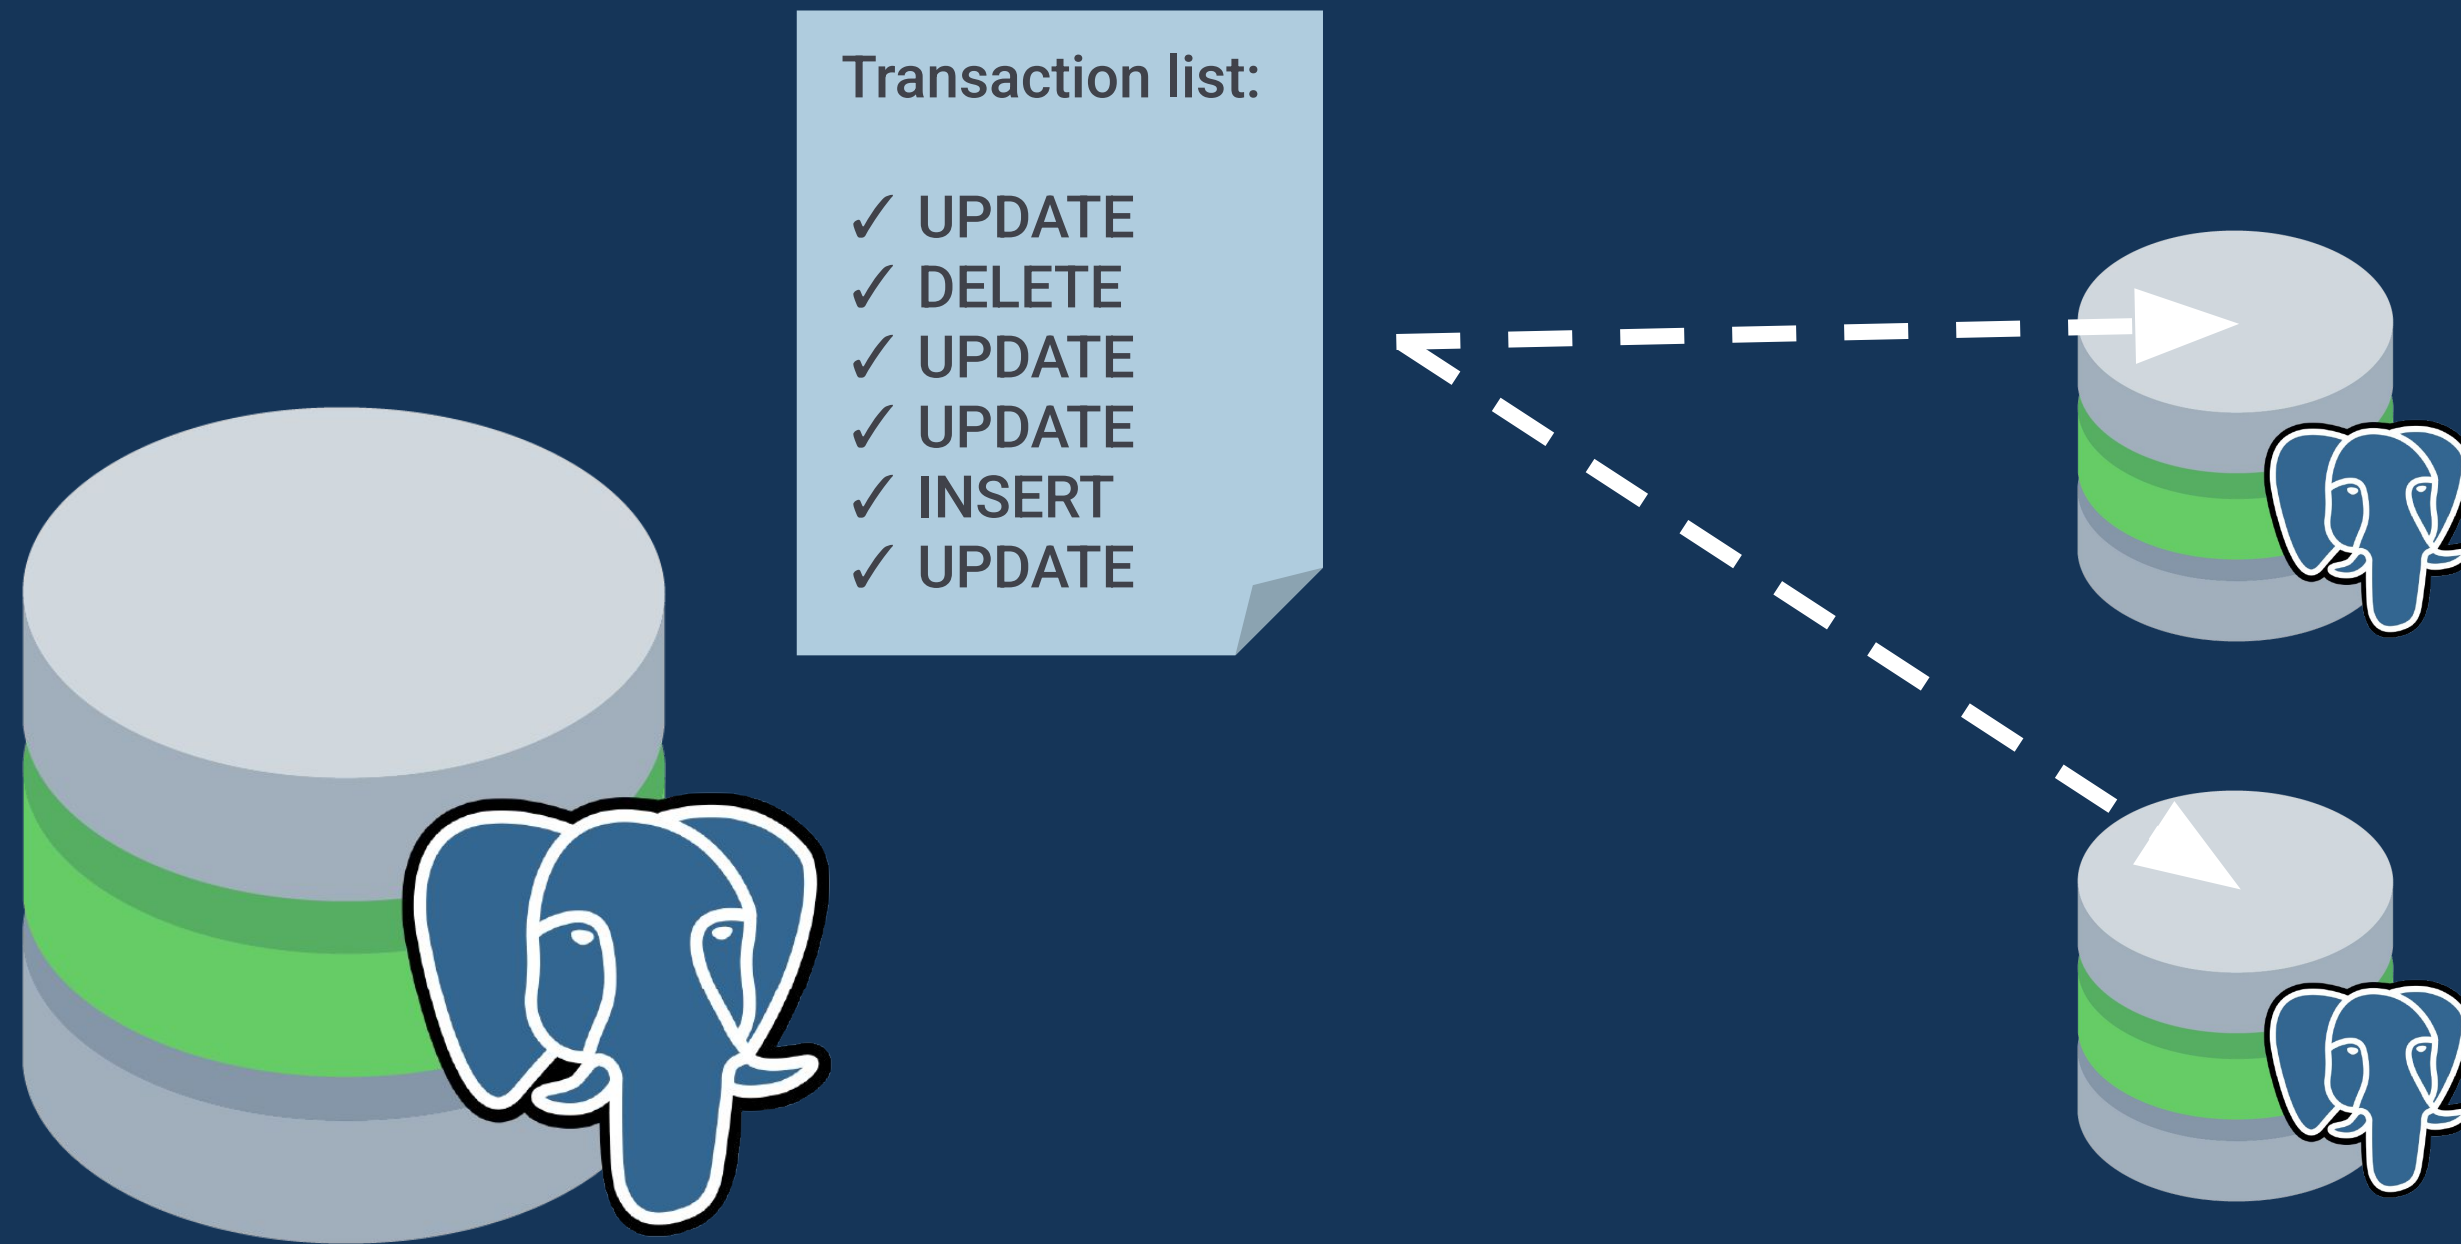
Logical Replication



Logical Replication



Logical Replication

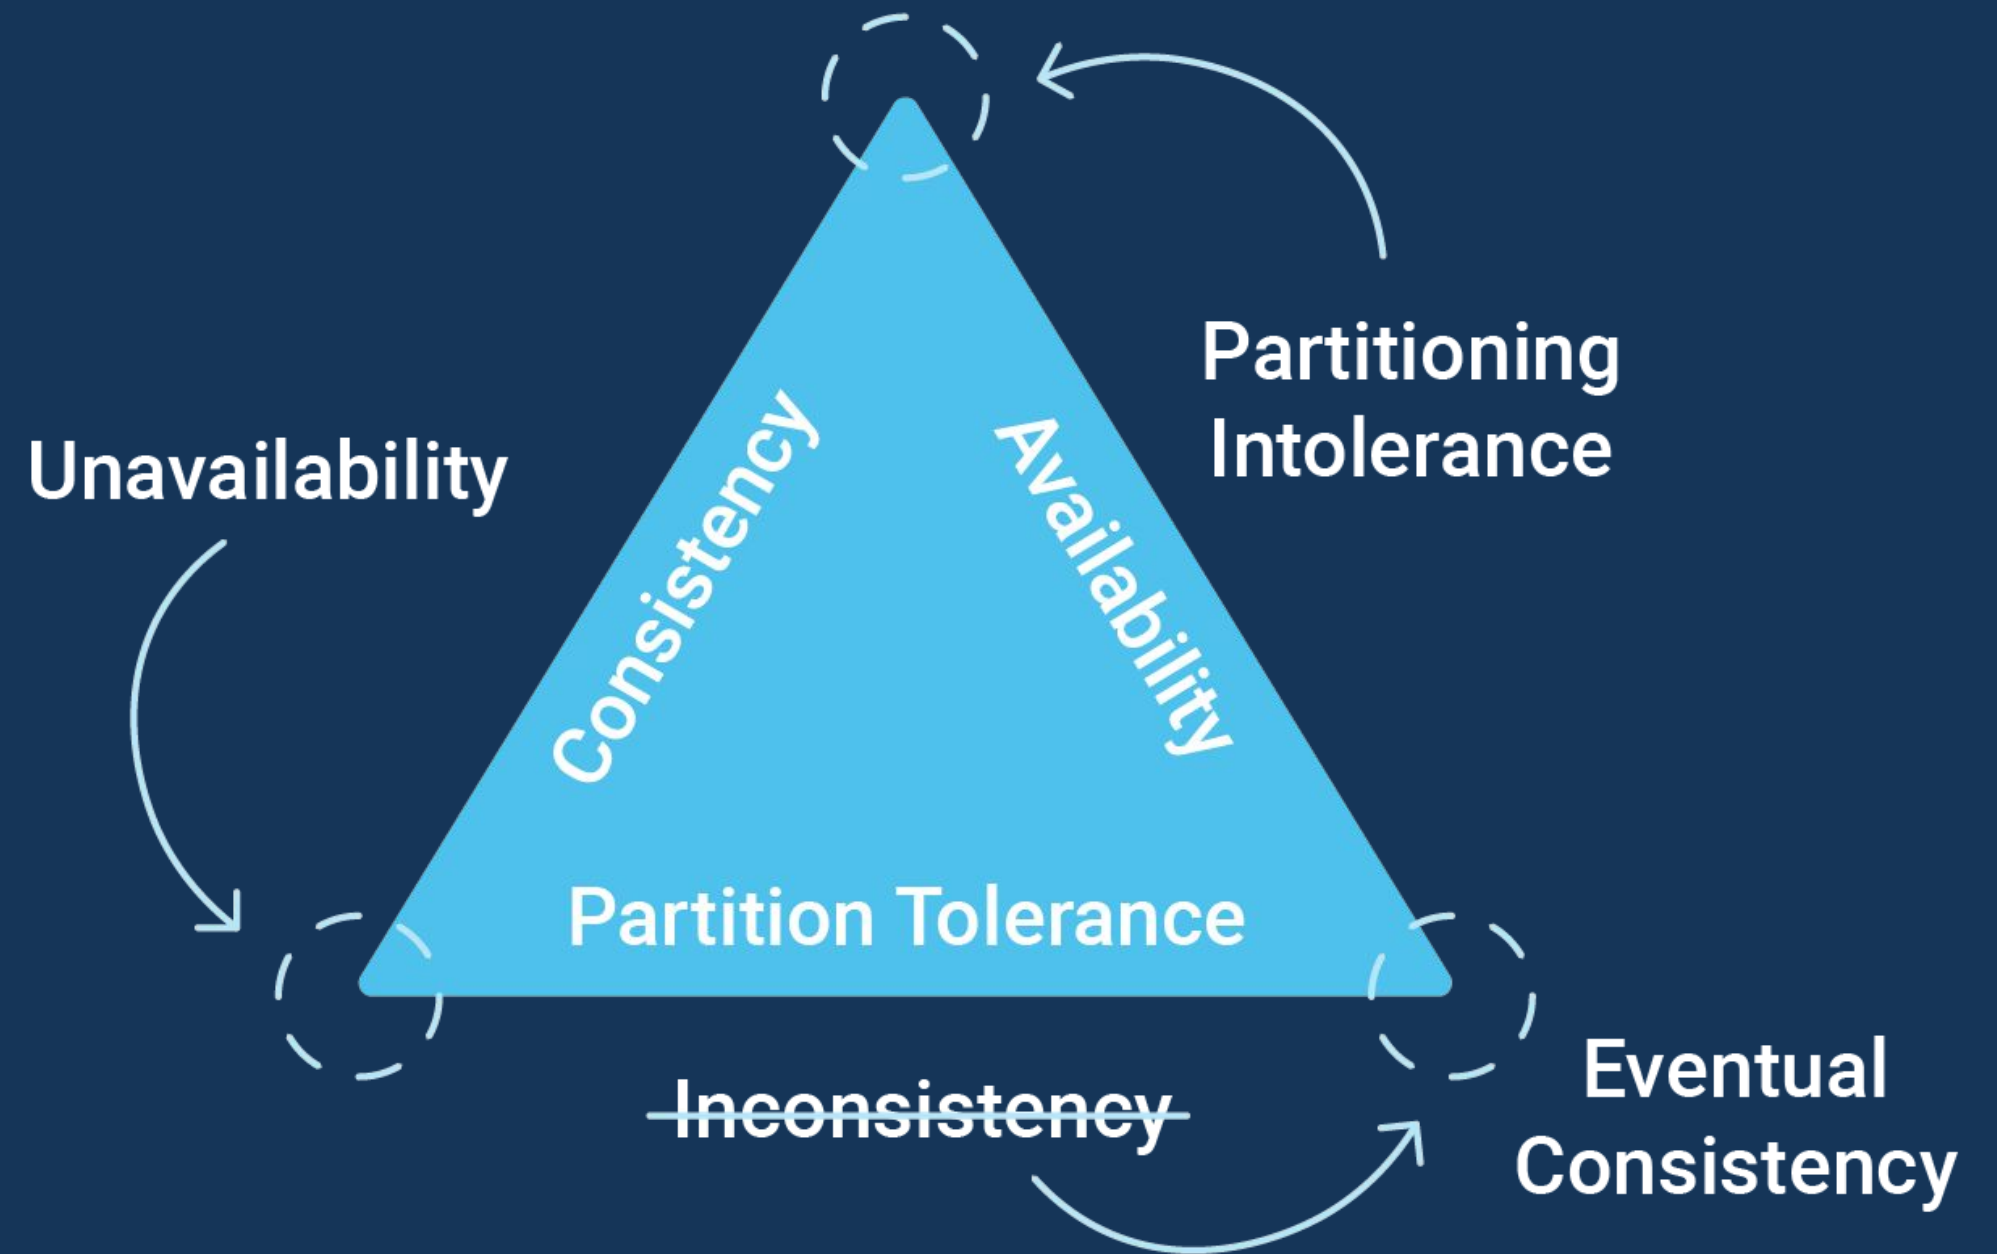


Risks

- Conflicts
- Logical corruption
- Conflicts
- Conflicts
- Conflicts

Requirements

- Is your application "aware"?
- How do you deal with sequences?
- Can be extremely useful
- If your use case needs it (1-2 %)



Agenda

1. What Is Streaming Replication?

And why does it even matter?

2. Single Elephant To A Herd

Scaling PostgreSQL safely (nearly) to infinity

3. And This Logical Replication Thing?

What is it and why does that matter?


4. What Does This Mean?

Even the basic functions make a difference



Solid Foundation

- Toolbox for creativity
- Rock solid and super secure
- Data protection knows no bounds
- From simple DR to global data protection strategies



Corollaries

Quoting The Legendary Ants Aasma

- PostgreSQL ecosystem has a solution for almost any problem you need to solve.
- For many specific tasks, there is something else out there that can do it better.
- However PostgreSQL will still get it done well enough.
- Very likely somebody has already made it do it and posted about it on the internet.



The Winning Formula

Quoting The Legendary Ants Aasma

1. Default to “Just use PostgreSQL”.
2. Use it until it no longer works.
3. Optimize until it works again.
4. Move the part that is hard to a special purpose tool.

If you are lucky, step 4 is never needed.

If it is, you arrive with knowledge and resources to tackle the problem correctly.





“ PostgreSQL is the best because
it is good enough for the task
you didn't know you had

Ants Aasma



Affiliations & Recognitions



Our Partners at PGDay Austria





Open Alliance

For PostgreSQL Education

Your Pathway to Verified PostgreSQL Skills

Scan for Updates



oapg-edu.org