



EDBTM

The Bones of
High Availability



Setting Expectations



RPO

Recovery Point Objective



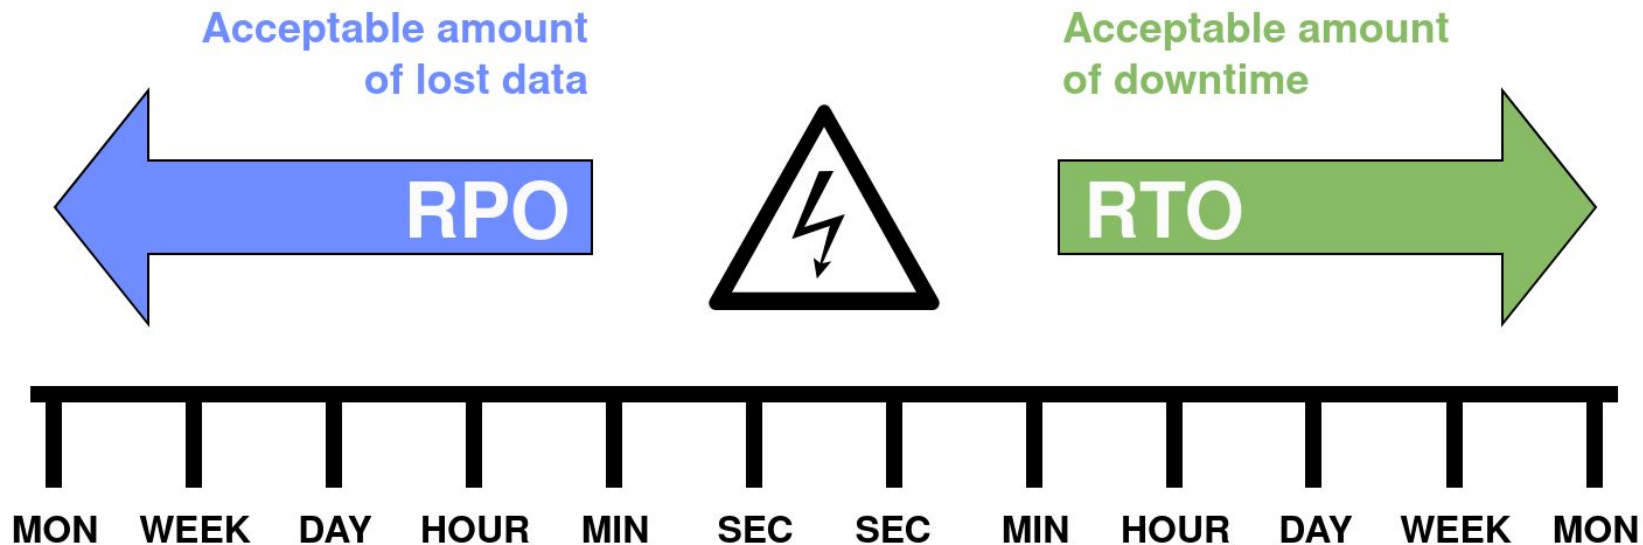


RTO

Recovery Time Objective

Percent	Per Year	Per Day
99%	3.7 days	14 minutes
99.9%	8.8 hours	86 seconds
99.99%	53 minutes	8.6 seconds
99.999%	5.3 minutes	0.86 seconds

RPO and RTO at a Glance



Glossary

What is High Availability?

\$

\$

\$

\$

\$



What is HA?

What is it, really?





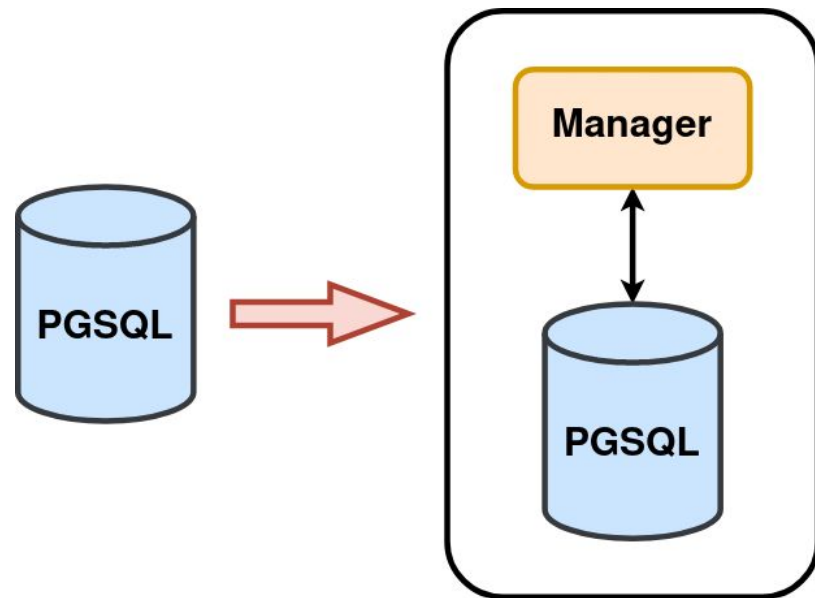
What is Quorum?

A bit odd





What is a Node?





What are Bones?

Fun Fact:

You contain a spooky skeleton!





Giving Postgres Bones


One is the Loneliest Number

Zone 1



Poor Postgres

Just look at him!




Eggs in One Basket

1. What happens if the node fails?
2. Where are backups in this scenario?
3. These actually exist “in the wild”

Two Can Be as Bad as One

Zone 1



Preventing Utter Disaster

3 – 2 – 1 Backup rule



**3 copies of
your data**

–



**2 different
media**

–



**1 copy
off site**

What About Backup Software?


Barman

- WAL archival
- WAL streaming
- Compression
- Incrementals
- Cloud options

pgBackRest

- WAL archival
- Encryption
- Compression
- Incrementals
- Cloud options

Three Isn't a Crowd

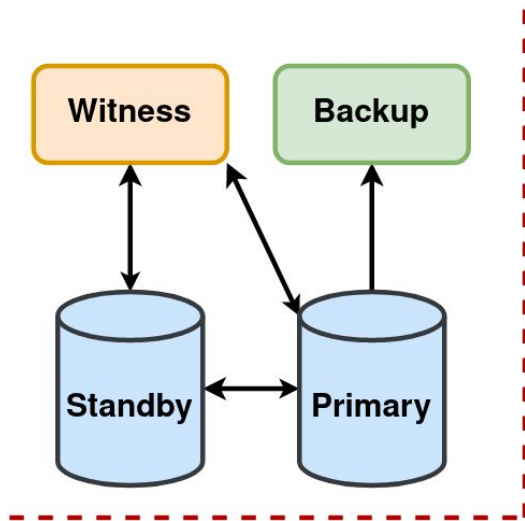


Three Amigos?

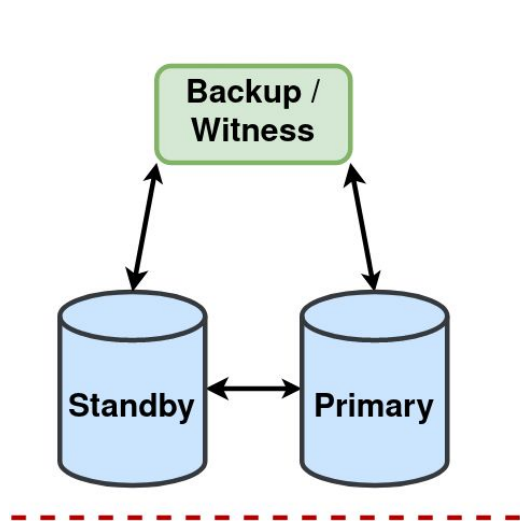
- Should we try to match hardware?
- Why isn't it safe to automate this?
- What does the RTO look like now?

Remember our Quorum

Zone 1



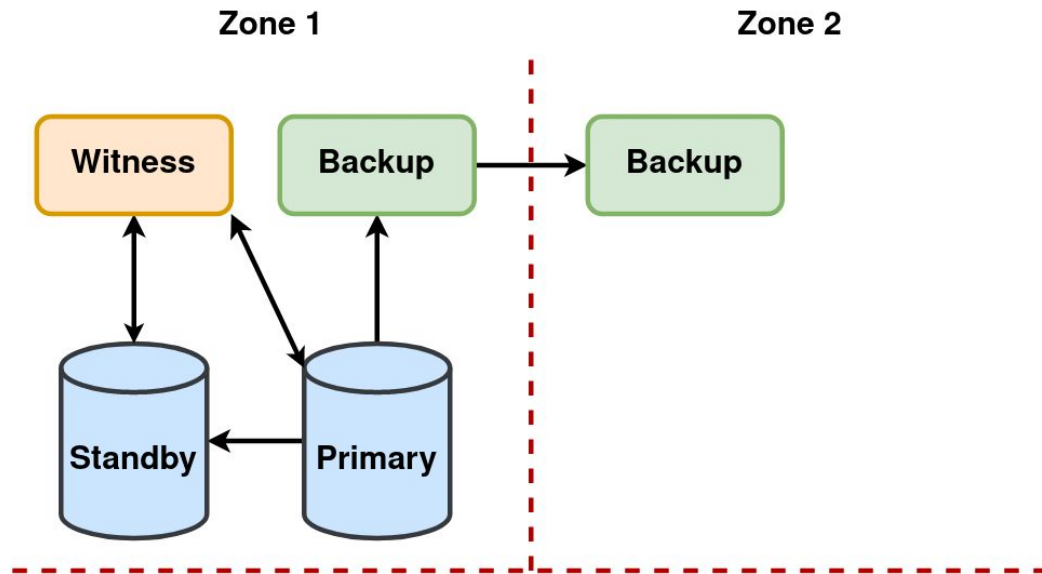
Zone 1




I Need a Witness

- Our first real HA cluster
- Remember the Management software
- What makes this better than the last example?

One for the Road



Two for the Road?



But Why?

*“Anything that can go wrong,
will go wrong”*

- **Murphy’s Law**

*“... and at the worst possible
time.”*

- **Finagle’s Law**

Floods!



Earthquakes!



Tornados!




Backhoes!




What's Missing?



Five... Golden... Servers!



Skipping Six



Maintaining Quorum

- Why did we add the witness again?
- What happens if Zone 1 fails?




Something is Missing

What can we add?




Round and Round we Go



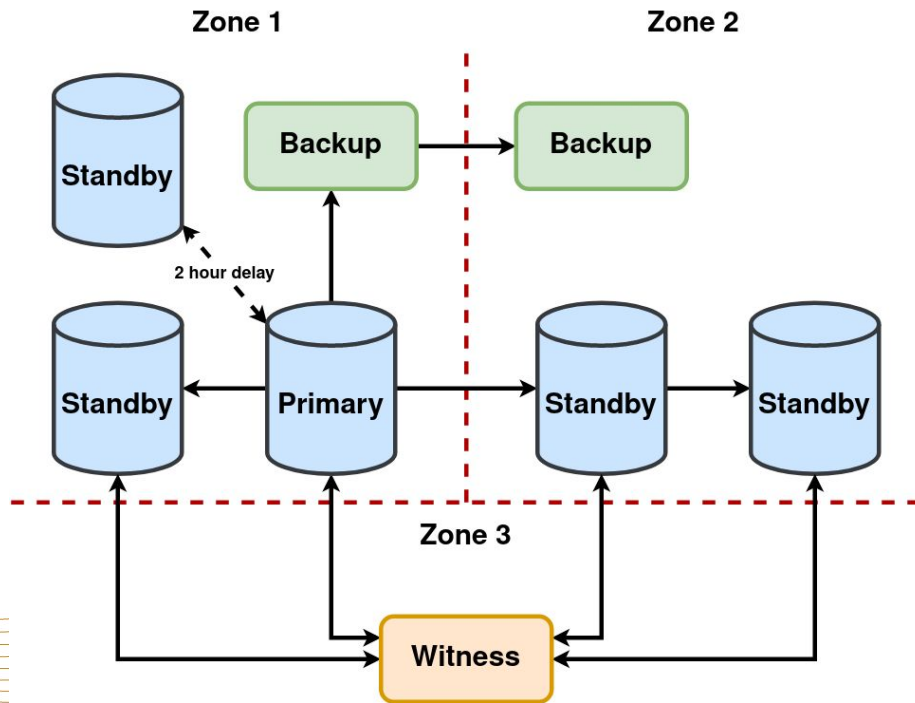
All's Well that Ends Well



Sudden Synchronization



When the Worst Happens



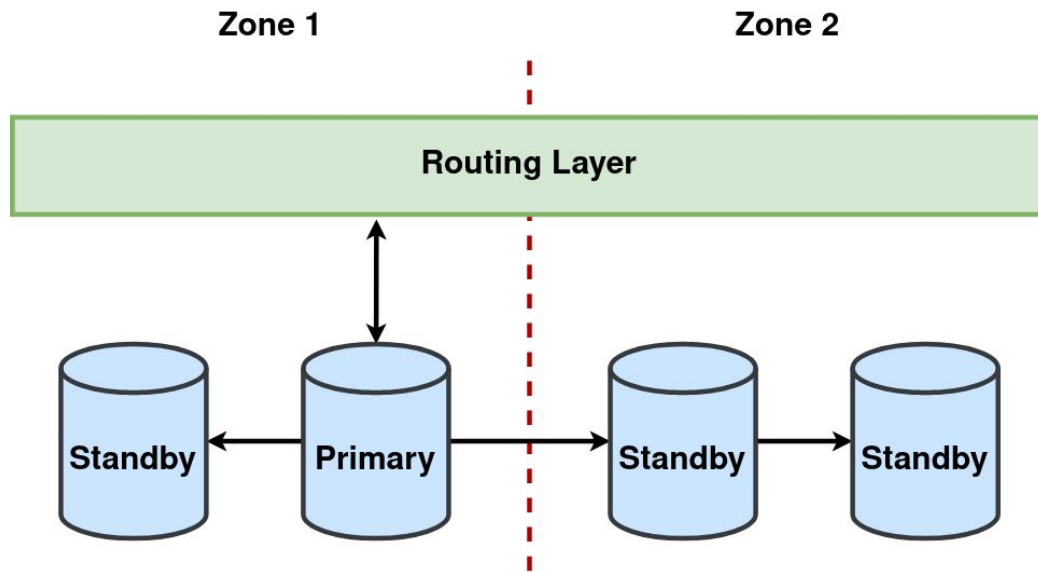
Delayed Gratification

- `recovery_min_apply_delay` parameter
- Goes well with:
 - `pg_wal_replay_pause()`
 - `pg_wal_replay_resume()`

Routing



Indirection!



Traffic Controller

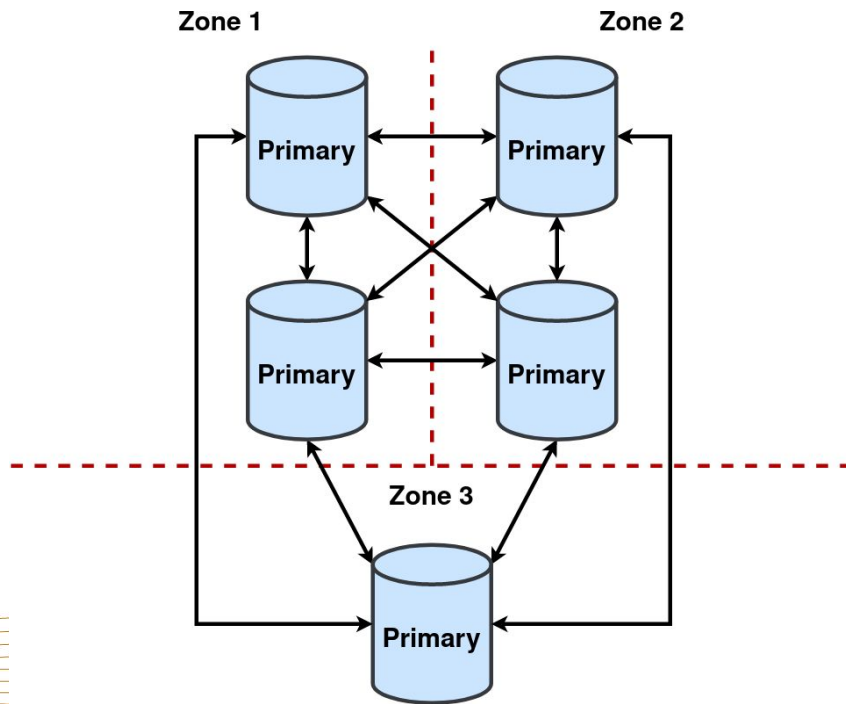
- Controlled by HA management service
- PgBouncer
- Pgpool-II
- HAProxy (Patroni)
- Load balancer



Hive Mind



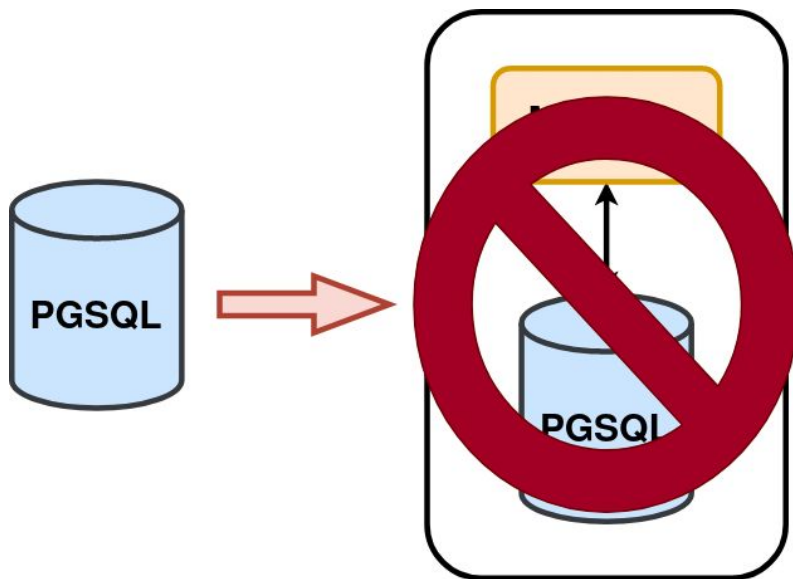
Multi-Master Logical Replication



The Final Cog

- Near instant failover (no promotion)
- Zero downtime major version upgrades
- Simplified routing
- And...

No more of this!



Alternatives

Why Bother?

“All of that looks really complicated. I think I’ll just use Firebase instead.”

- **THAT Conference Attendee**

BigAnimal: Faster, safer, smarter, better



Postgres Expertise

Expertise beyond the generalist cloud provider; we help steer the database roadmap and patch its bugs



Oracle Compatibility

Leave Oracle and further your cloud journey with a fully managed Postgres service



Continuous availability

High availability of your PostgreSQL clusters so you're always on, always available

Curious? Signup today for 14 day self-service FREE trial!
<https://portal.biganimal.com>

Managed Services

- High Availability
- Routing
- Backup



Amazon RDS




Google
Cloud SQL



Kubernetes Operators

- Configure with Yaml
- Manage whole cluster
- Many to choose from


EDB CloudNativePG



Crunchy Data PGO



Zalando postgres-operator



More Kubernetes Operators

Percona Operator for PostgreSQL



AppsCode KubeDB



OnGres StackGres



Join us for Part 2

[Register Today](#)



How to Achieve Zero Database Downtime

August 31st @ 11am ET

How to achieve up to five 9s availability
on your Postgres implementation

Gianni Cioli
VP of Solutions Architecture
& Field CTO



END