

patroni vs stolon vs repmgr

Как мы выбирали среди patroni, stolon, repmgr для нашего отказоустойчивого Постгреса

Андрей Фефелов
mastery.pro

Содержание

- High Availability применительно к Postgres
- Постановка задачи
- Краткий обзор и сравнение решений
 - Amazon Aurora/RDS
 - repmgr
 - stolon
 - patroni
- Решение задачи в нашем случае

High Availability применительно к Postgres

- NAS/SAN
- DRBD
- Logical (trigger based old fashion, pglogical)
- WAL shipping

Постановка задачи

- Hetzner
- 3 DC (Nuremberg, Falkenstein, Helsinki)
- RTO = 10min
- RPO --> 0 =)
- Automatic db failover
- Automatic app failover
- Fencing

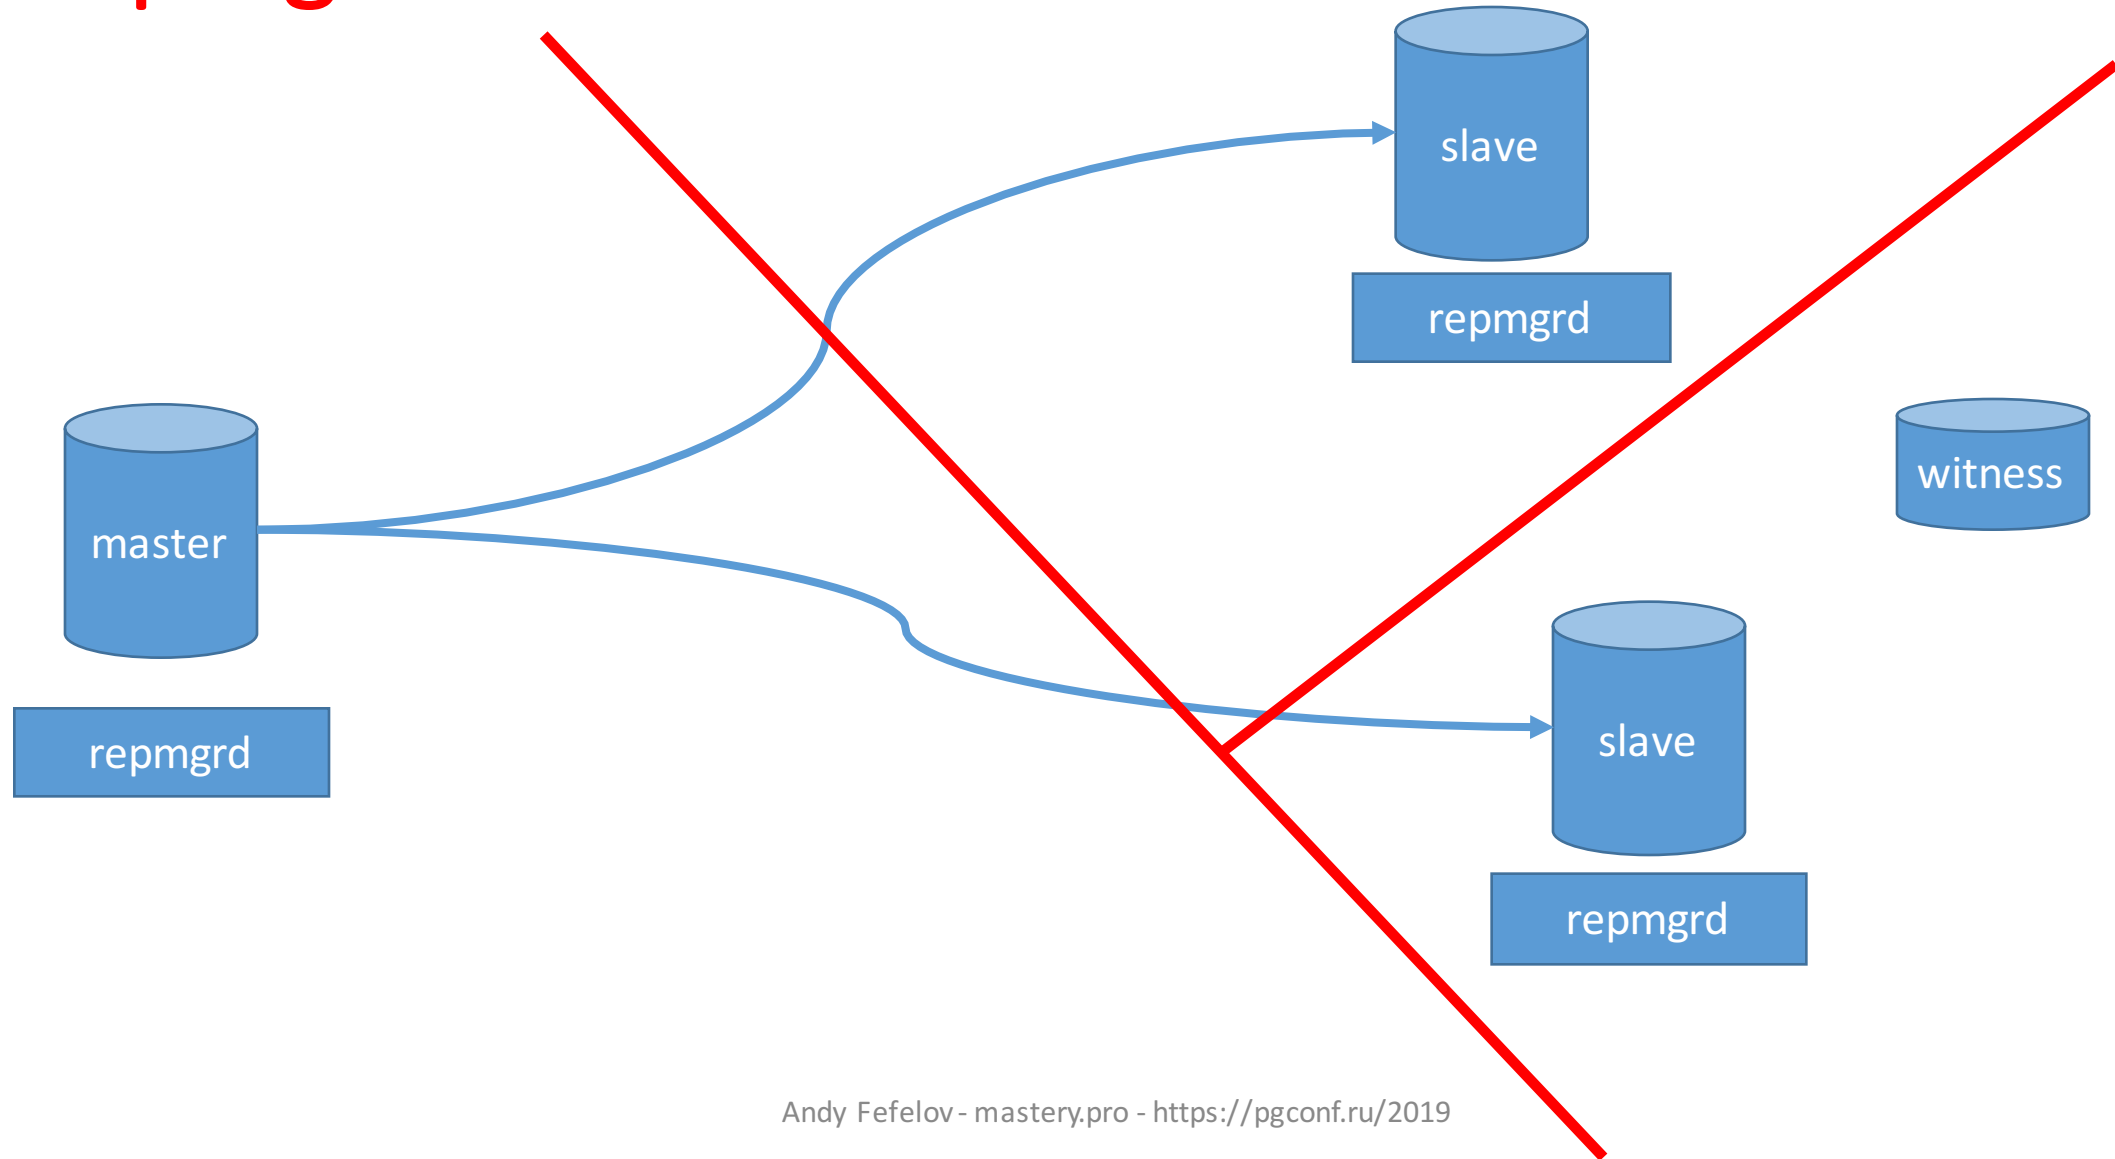
Amazon Aurora/RDS

- Aurora - быстро, очень быстро, очень дорого (aurora)
- Есть почти все extensions, нет pgq (pglogical since Sep 14, 2018)
- RDS - медленно (за те же деньги)
- Сложности «внешнего» бэкапа

repmgr

- 2ndQuadrant, C
- Any postgres from 9.3 -> 11
- Witness server concept
- No need for dcs
- Things to learn: 2
- No fencing
 - Workaround
<https://github.com/2ndQuadrant/repmgr/blob/master/doc/repmgrd-node-fencing.md>

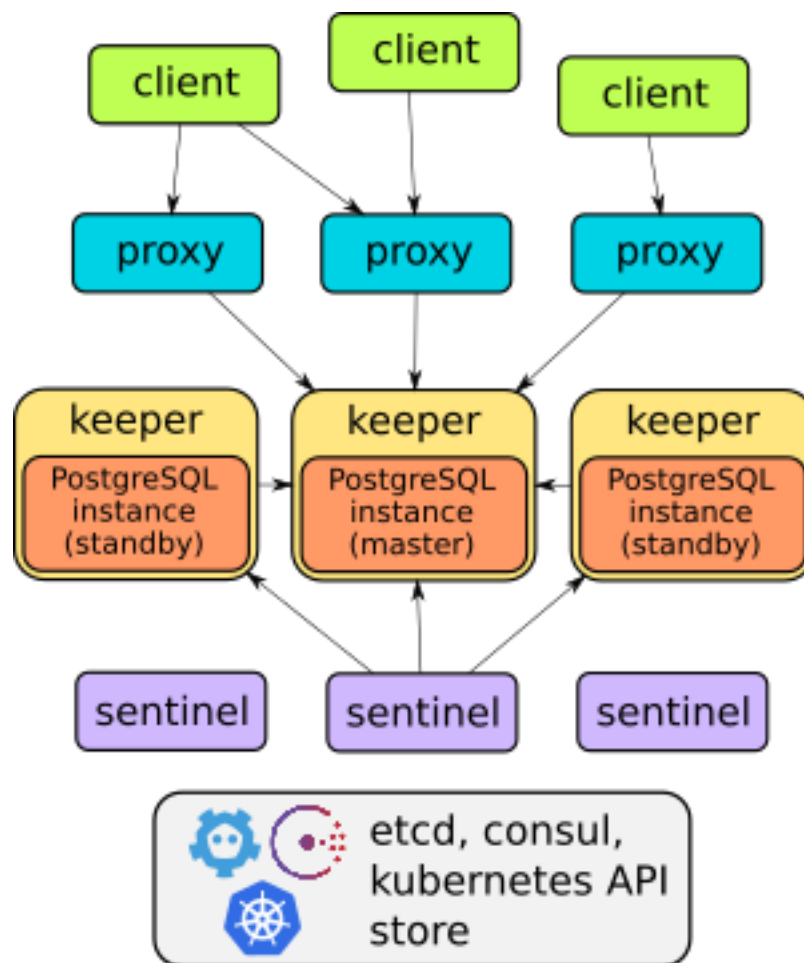
repmgr



stolon

- Sorintlab, go
- Use DCS
- k8s ready
- Sync/async replication
- Custom init
- Standby cluster
- RW+RO -> master only
- One way switchover
- pg_rewind
- Things to learn: 4

stolon



patroni + consul + wal-g + s3 + docker

- Zalando, python
- Use DCS
- k8s ready
- Sync/async replication
- Custom init
- Standby cluster
- pg_rewind
- Things to learn: 2

patroni + **consul** + wal-g + s3 + docker

- UI
- dns
- service checks
- patroni умеет анонсировать master/replica ❤️❤️❤️❤️❤️

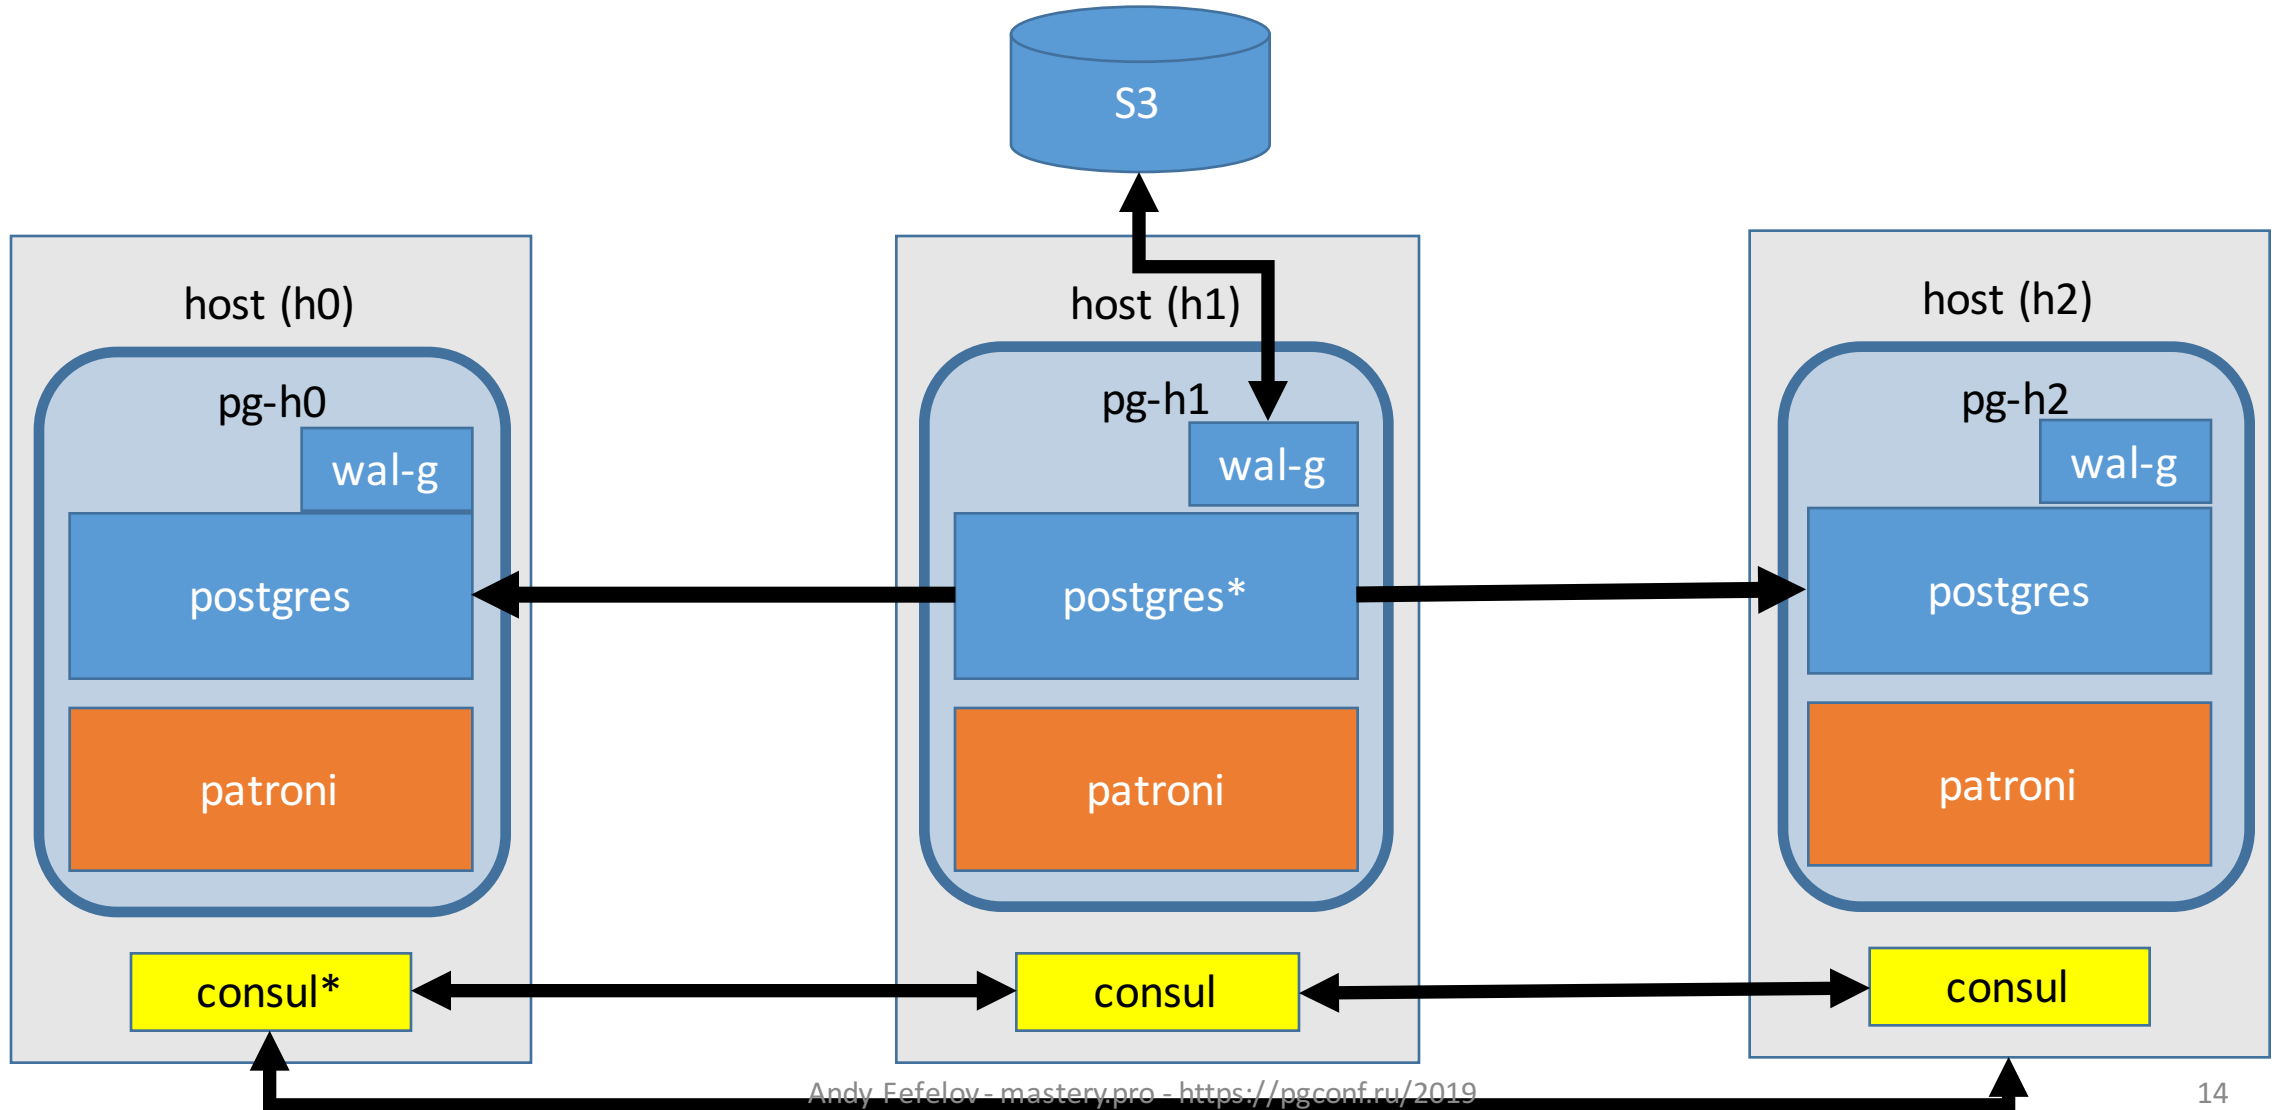
patroni + consul+ **wal-g** + s3 + docker

- fast (parallelized)
- compression
- quite simple
- production ready

patroni + consul+ wal-g + **s3** + **docker**

- we just use them

patroni + consul + wal-g + s3 + docker



patroni + consul+ wal-g + s3 + docker

- Pros:

- Simple setup
- Fast master init
- Consistent backup from any node (thanks to REST API)
- Failover works
- Hooks, one love!

- Cons:

- DNS should be tweaked
- Slow slave init
- Application should be ready to loose database

thanks to and links

- Александр Кукушкин (patroni)
- Андрей Бородин (wal-g)
- patroni.readthedocs.io/
- github.com/wal-g/wal-g
- github.com/afefelov/patroni-class