

vers le PetaByte

avec PostgreSQL

Dimitri Fontaine dimitri@2ndQuadrant.fr

13 Juin 2013

2ndQuadrant France PostgreSQL Major Contributor

- pgloader, prefix, skytools, debian, ...
- CREATE EXTENSION
- CREATE EVENT TRIGGER
- *Bi-Directional Réplication*
- *Partitionnement*




C'est quoi un *PetaByte*?

Beaucoup de données

1 PB =

- 1 024 **TB**
- 1 048 576 **GB**
- 1 073 741 824 **MB**
- 1 099 511 627 776 **KB**
- 1 125 899 906 842 624 **Bytes**





One Petabyte of one Gigabyte IBM 3380 (from 1980)

One million drives, weighs 250,000 Tonnes

Costs \$40,000,000






One Petabyte of 2TB Western Digital Drives

500 drives, weighs 365KG

Costs \$125,000

Big Data?



Data



Information



Presentation



Knowledge



EpicGraphic.com



The Horizontal Struggle


- 0.1 GB schema, backup: seconds
- 1 GB schema: seconds, backup: minutes (a few)
- 10 GB schema: depends, backup: minutes (many)
- 100 GB schema, backup: hours
- 1 TB schema: difficult, backup: days




VACUUM, AUTO VACUUM



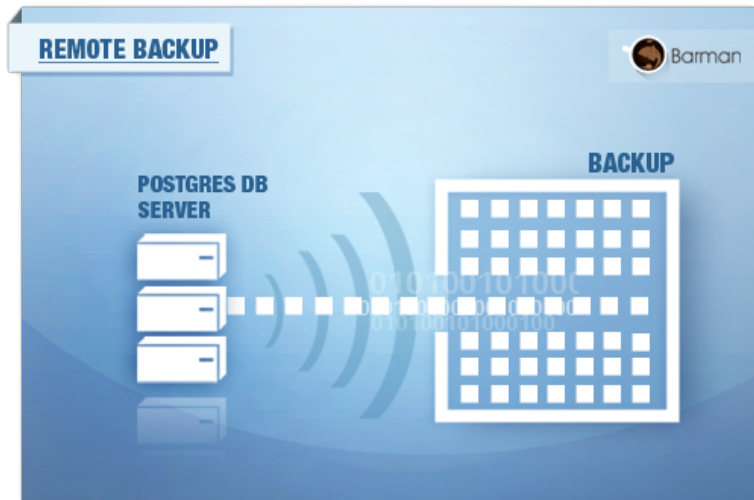
INDEXES



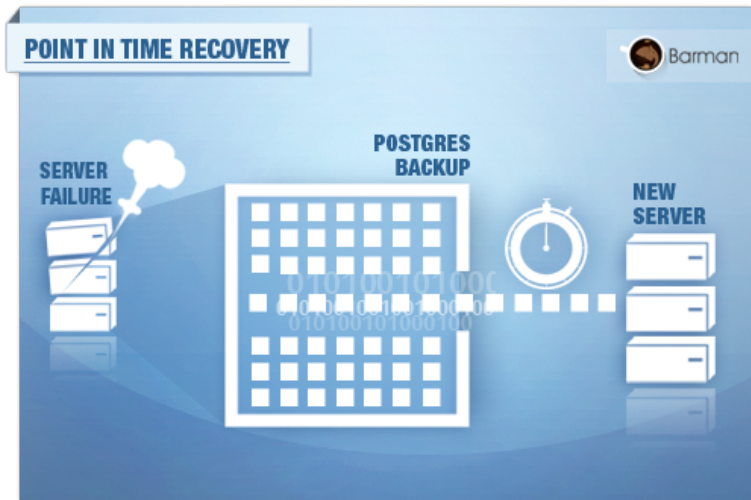
CHECKPOINT



Sauvegardes, pg_dump, basebackup



Archivage, PITR



Replication

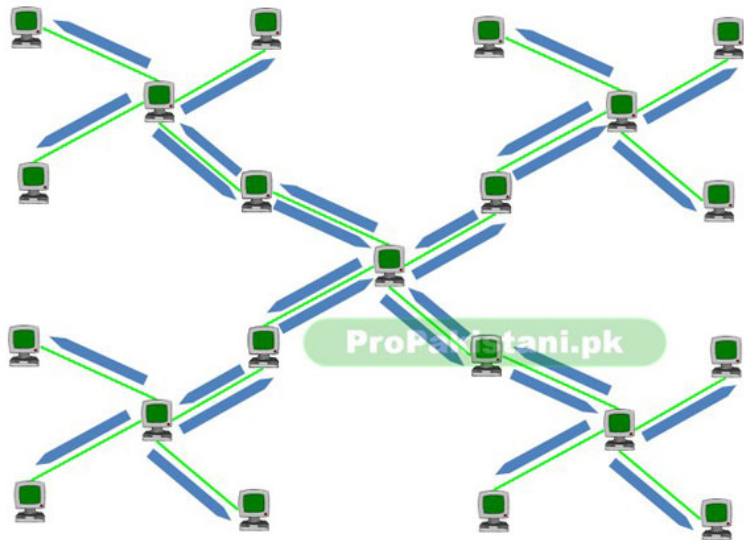


Développements en cours




- CHECKSUM
- COPY FREEZE
- `work_mem`
- REINDEX CONCURRENTLY
- VACUUM FREEZE, XID EPOCH
- `work_mem` > 1 GB

Architectures distribuées



Bi Directional Replication



- *Read/Write Foreign Data Wrappers*

- *PL/Proxy*

Conclusion

