



**Fast, cheap,
good ...
CHECK!**

*PGcon
2012
Jan Wieck*

AfiliasSM
GLOBAL REGISTRY SERVICES

Introduction

- **Jan Wieck**
PostgreSQL User and Contributor
Slony-I Author
- **Afilias Global Registry Services**
PostgreSQL User and Contributor
PostgreSQL Community Member
Slony-I Sponsor

Introduction

- **PostgreSQL has many options for replication.**
 - **They are all reasonably fast.**
 - **They are all cheap (how cheap is free?)**
 - **They are all good (for something)**
-
- **So what's missing?**

Introduction

**You need to tell us
what is missing!**

Use case Disaster Recovery

- **Disasters happen, are you prepared?**



Credit: <http://www.flickr.com/photos/noaaphotolib/>

Use case Disaster Recovery

- Backup and Recovery



Credit: <http://www.flickr.com/photos/kylemacdonald/>



Credit: <http://www.flickr.com/photos/mirjana/>

So how did we get here?

- **How did we get to using replication for backup and recovery?**
- **Backup strategies evolved to offer ever faster recovery, to the point of instant failover to a standby.**
- **Most backup strategies focus on “don't lose a single committed transaction!”**
- **Is that all that's required?**

What if the error is the user?

- You wanted clones?



Credit: <http://www.flickr.com/photos/evelynishere/>

What if the error is the user?

- But what good are perfect copies of the wrong data?



Credit: <http://www.flickr.com/photos/rogdavies/>

Desired features for Backup

- **Synchronous replication**
- **Delayed replication and PITR**
- **Taking other forms of backup from a replica.**

More about PITR

- **Rolling forward can be dangerous!**



Credit: <http://www.flickr.com/photos/edgoodwin/>

Desired features for Backup

- **Synchronous replication**
- **Delayed replication and PITR**
- **Taking other forms of backup from a replica**
- **Tools to analyze the replication stream**
- **Ability to stop PITR at the identified point in time**

Use case Scaling



Credit: <http://www.flickr.com/photos/neaaphotolib/>

Scaling by offloading

- **Move some of the heavy load to a replica**



Credit: <http://www.flickr.com/photos/bobolink/>

Offload Reporting

- **Take a careful look! That guy is WRITING!**



Credit: <http://www.flickr.com/photos/wwwworks/>

Offload special functionality

- **Pre-aggregating data**
- **Maintaining expensive indexing (full text search)**
- **Using new PostgreSQL features**

Desired features for Offloading

- **Selective replication (by table, by row, by column)**
- **Replica side triggers**
- **Replica side additional indexes**
- **Replica side additional tables**
- **Replica side additional columns**
- **Cross version**

Use case Road Warriors



Credit: <http://www.flickr.com/photos/thekog/>

What is a Road Warrior?

- **The sales representative, that needs to be able to work offline and synchronize with the central office later.**
- **The system will download central data (products, price lists, etc.) and upload the sales orders when connected.**

Features for the Road Warrior

- **Prolonged offline times**
- **Multiple origins**
- **Merge replication**
- **Selective replication**
- **Cross version and architecture**

Use case Consolidation



Credit: http://www.flickr.com/photos/stevenm_61/

What is Consolidation?

- **A user has multiple legacy systems that have eventually conflicting database requirements.**
- **Combine the data of all systems into one database to do cross application analysis and reporting.**

Features for Consolidation

- **Cross version and architecture**
- **Selective replication**
- **Renaming of objects during replication**
- **Marshalling of data**

Use case Migration



Credit: <http://www.flickr.com/photos/audreyjm529/>

What is Migration?

- **Moving data to new hardware/OS**
- **Upgrading to a new DB version**
- **Time consuming schema adjustments**

Features for Migration

- **Cross version and architecture**
- **Online installation, initial copy and then catch up without system downtime**
- **Switching master/slave (consider a connection pool)**
- **Marshalling of data**

Use case Untrusted Destination



Credit: <http://www.flickr.com/photos/humphreyking/>

What is an Untrusted Destination ?

- **A replication target that should not have all the data**
- **A replication target that if tampered with could affect your production systems**
- **A replication target that cannot establish access from the outside**

Features for Untrusted Destination

- **Selective replication**
- **Push only data transfer**
- **Zero impact if target becomes unavailable**

Questions and Answers

● ?!?