

PostgreSQL Replication Solutions

BRUCE MOMJIAN



Replication is a complex feature. POSTGRES SQL supports a variety of replication options.

Creative Commons Attribution License

<http://momjian.us/presentations>




Last updated: April, 2017

Uses for Replication





<https://www.flickr.com/photos/eugenius/>





Fail Over







Data Warehousing




Load Balancing



Remote Servers



Mobile Servers




Replication Solutions








<https://www.flickr.com/photos/paulbence/>

Shared Storage




- ▶ No overhead
- ▶ No data loss on fail-over
- ▶ Slave cannot execute queries

Storage Mirroring




DRBD






- ▶ No overhead on master
- ▶ Synchronous or asynchronous
- ▶ Possible data loss on fail-over when using asynchronous
- ▶ Slave cannot execute queries

Streaming Replication



- ▶ No overhead on master
- ▶ Slaves can execute queries
- ▶ Possible data loss on fail-over when using asynchronous mode
- ▶ Synchronous option available (Postgres 9.1)

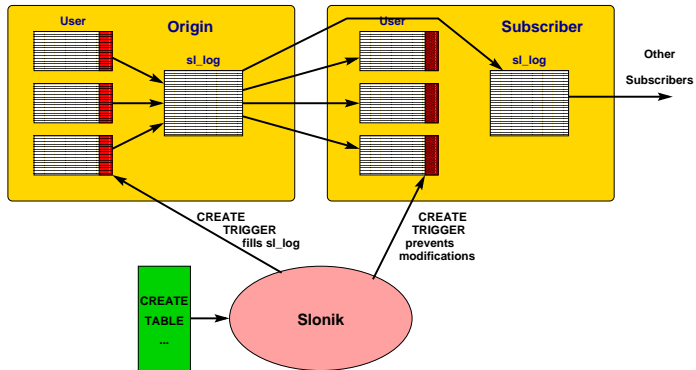
Slony




Asynchronous →

- ▶ Triggers add overhead to the master
- ▶ Possible data loss on fail-over
- ▶ Replication possible even over slow links
- ▶ Slave can execute read-only queries
- ▶ Table-level granularity allows complex data partitioning configurations



Slony Internals





Slony Master Switching



Bucardo




- ▶ Similar to Slony, except multi-master with conflict resolution
- ▶ Conflict resolution rules are user-configurable




Asynchronous
↔
with Conflict Resolution

Pgpool II



- ▶ Automatically load-balances read queries
- ▶ Queries with non-deterministic behavior can cause inconsistency
- ▶ Allows parallel query execution on all nodes
- ▶ Also does connection pooling and query caching

Pgpool II With Streaming Replication



Streaming replication avoids the problem of non-deterministic queries producing different results on different hosts.

Summary

Feature	Shared Disk Fail-over	File System Replic.	Transaction WAL Log Shipping	Trigger-based Replic.	Statement-Based Replication Middleware	Asynchronous Multi-Master Replic.	Synchronous Multi-Master Replic.
Most Popular Implementation	NAS	DRBD	Log shipping	Slony table rows	pgpool-II	Bucardo	
Communication Method	shared disk	disk blocks	WAL	table rows	SQL	table rows	table rows & row locks
No Special hardware required		•	•	•	•	•	•
Allows multiple master servers					•	•	•
No master server overhead	•		•		•		
No waiting for multiple servers	•		•	•		•	
Master failure will never lose data	•	•			•		•
Slaves accept read-only queries			•	•	•	•	•
Per-table granularity				•		•	•
No conflict resolution necessary	•	•	•	•			•

<http://momjian.us/presentations>