

PostgreSQL Replication Solutions

BRUCE MOMJIAN



Replication is a complex feature. POSTGRESQL supports a variety of replication options.

<https://momjian.us/presentations>



Creative Commons Attribution License




Last updated: June 2024

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






- 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

Slony

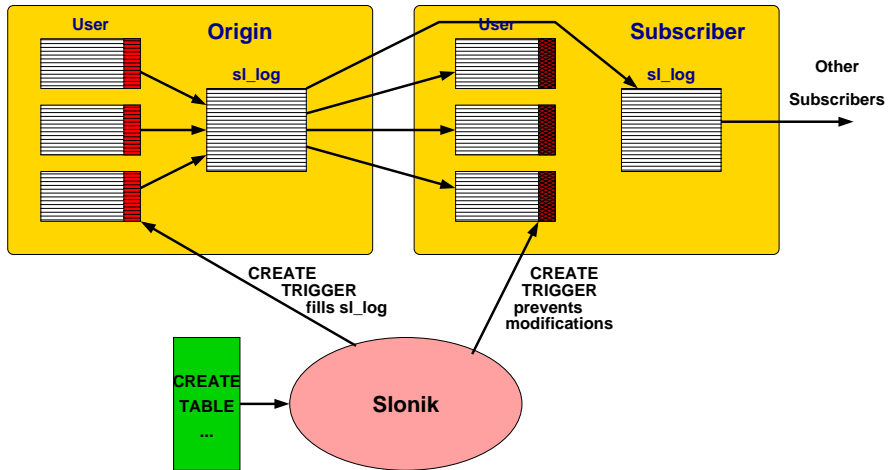


Asynchronous

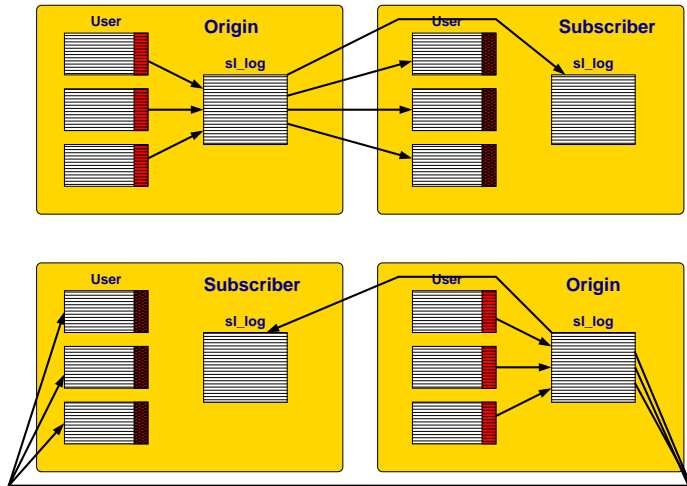


- Triggers add overhead to the master
- Possible data loss on fail-over
- Replication possible even over slow links
- Slave can execute read-only queries
- Table-level granularity allows complex data partitioning configurations



Slony Internals



Slony Master Switching



Bucardo




Asynchronous
with Conflict Resolution

A red dashed double-headed arrow points horizontally between the two server racks, indicating the replication direction. Below the arrow, the text "with Conflict Resolution" is written in red.


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

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	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	•	•	•	•			•



<https://momjian.us/presentations>