

Using streaming replication of PostgreSQL with pgpool-II

Tatsuo Ishii

President/PostgreSQL committer

SRA OSS, Inc. Japan

About SRA OSS

- Part of SRA group
 - SRA group is an independent software house in Japan
 - Employees: 1,700
 - Sales: 3.4 billion yen
- SRA OSS
 - Established in 2005
 - Specialized in OSS support and consultations
 - Has over 1,000 support customers
- Closely working with PostgreSQL community

About me


- PostgreSQL & pgpool-II developer
- Working as a PostgreSQL evangelist(16 talks in 2010)
- Working for SRA OSS, Inc. Japan




Why need replication?

- Because it:
 - Saves your life when one of your data is lost due to hard disk crash
 - Prevents the database system from accidental stopping due to hardware or network troubles. Thus allow you to continue business
 - Increases database performance by adding database servers

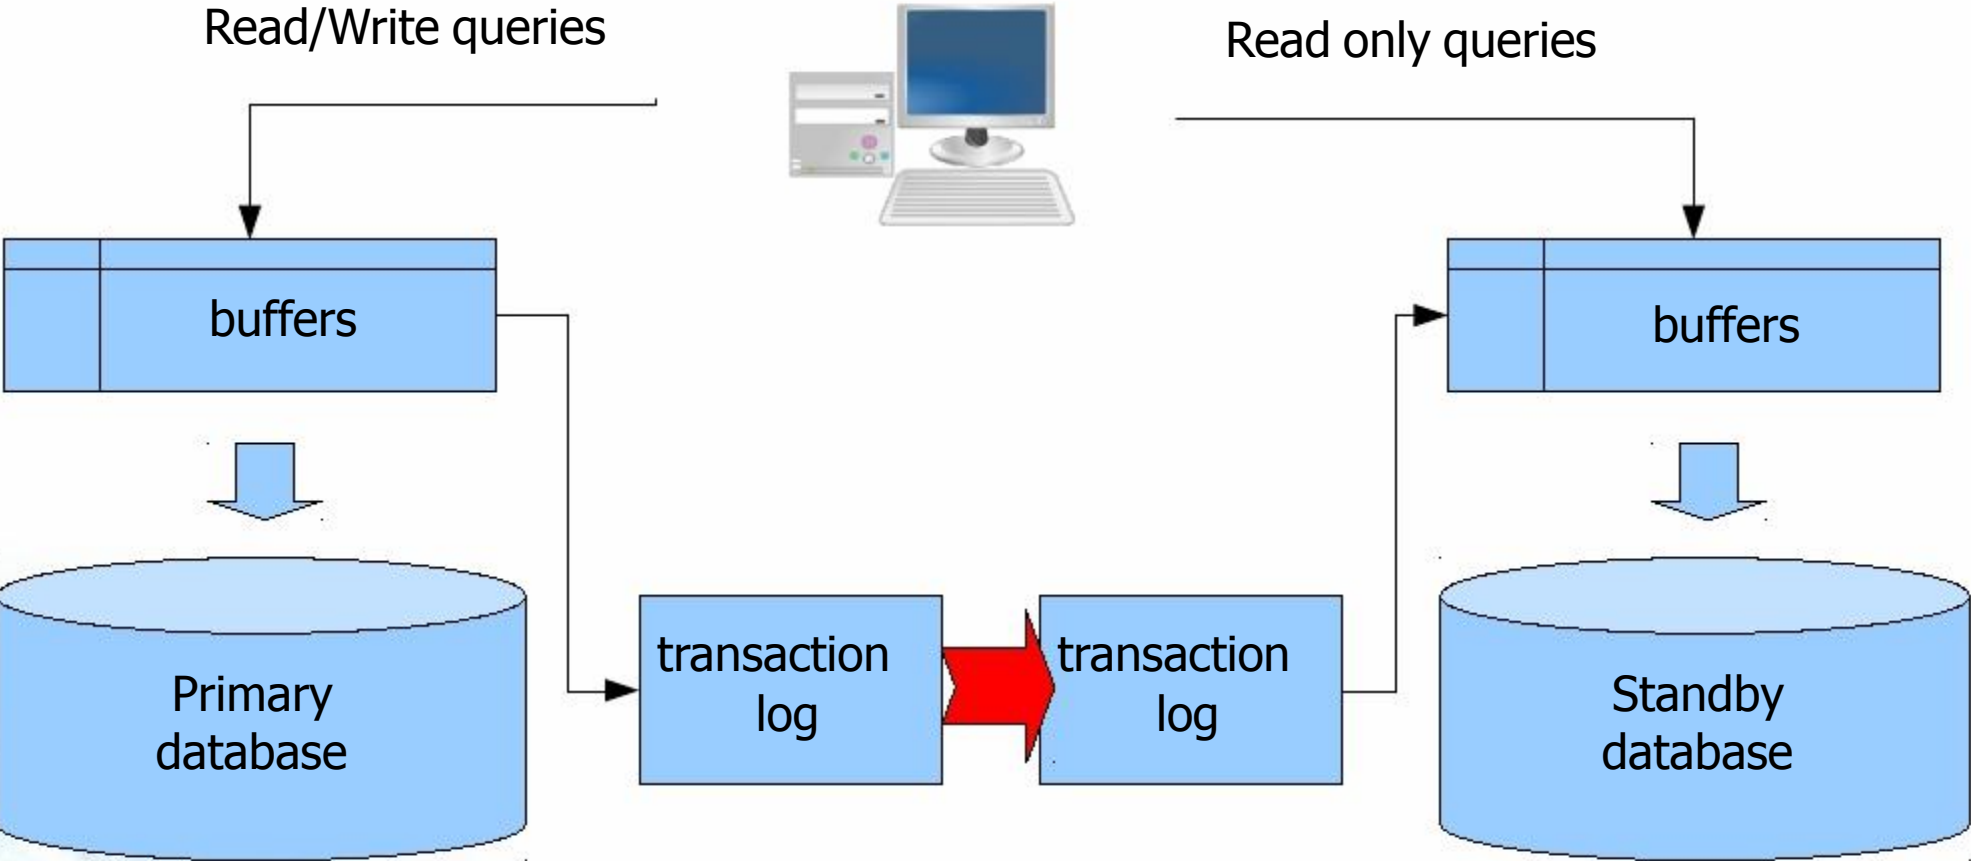
How to make database replicas?



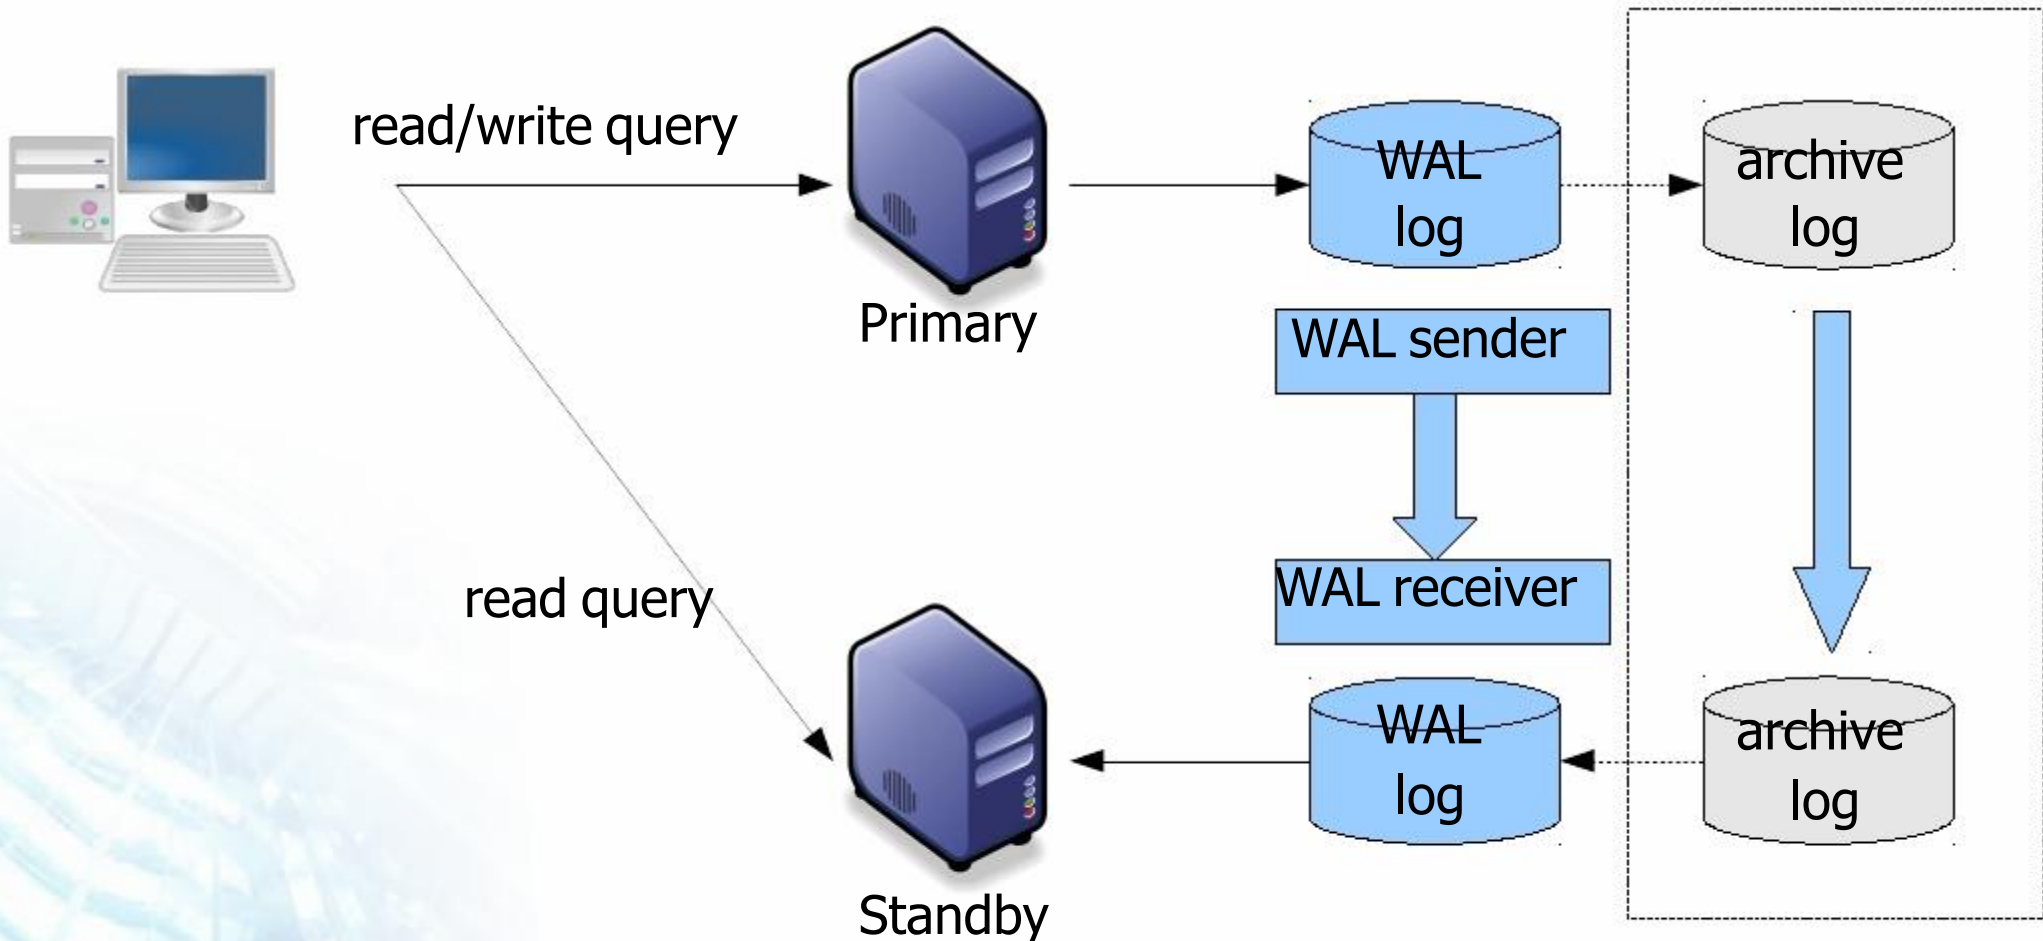
Transaction log



Idea of streaming replication

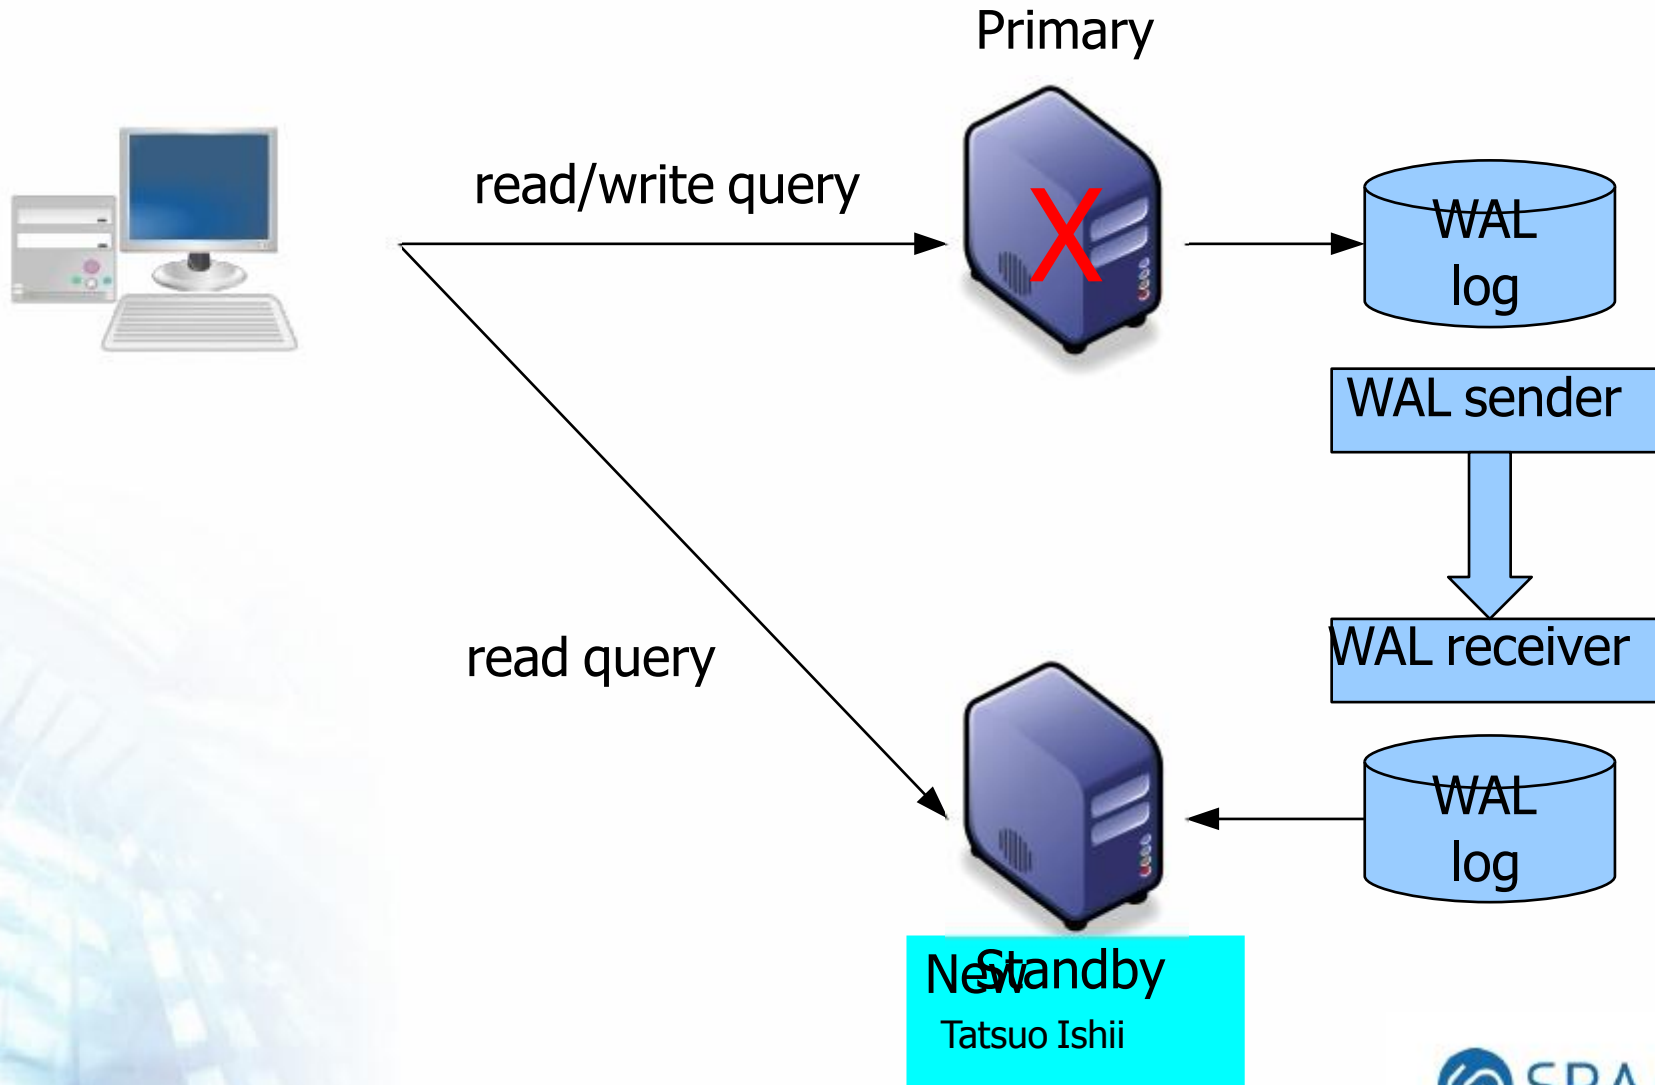


Architecture of streaming replication



Copyright(c)2011 Tatsuo Ishii

Promoting standby



New Standby
Tatsuo Ishii

Benefits of streaming replicaton


- PostgreSQL built-in replication available since 9.0
- Easy to setup
- Low overhead
- Everything can be replicated
 - Slony-I does not replicate DDLs, DCLs and large objects
- No application changes necessary as long as accessing primary

Things you might care about

- Replication lag(so called “asynchronous replication”)
- No automatic failover
- No built-in load balancing
- Can't send write queries to standby
 - DDLs, DCLs, DMLs
 - SELECTs including functions which write to databases
 - Some locking commands etc.

Why pgpool-II?

- Some of problems can be solved or at least moderated by using pgpool-II

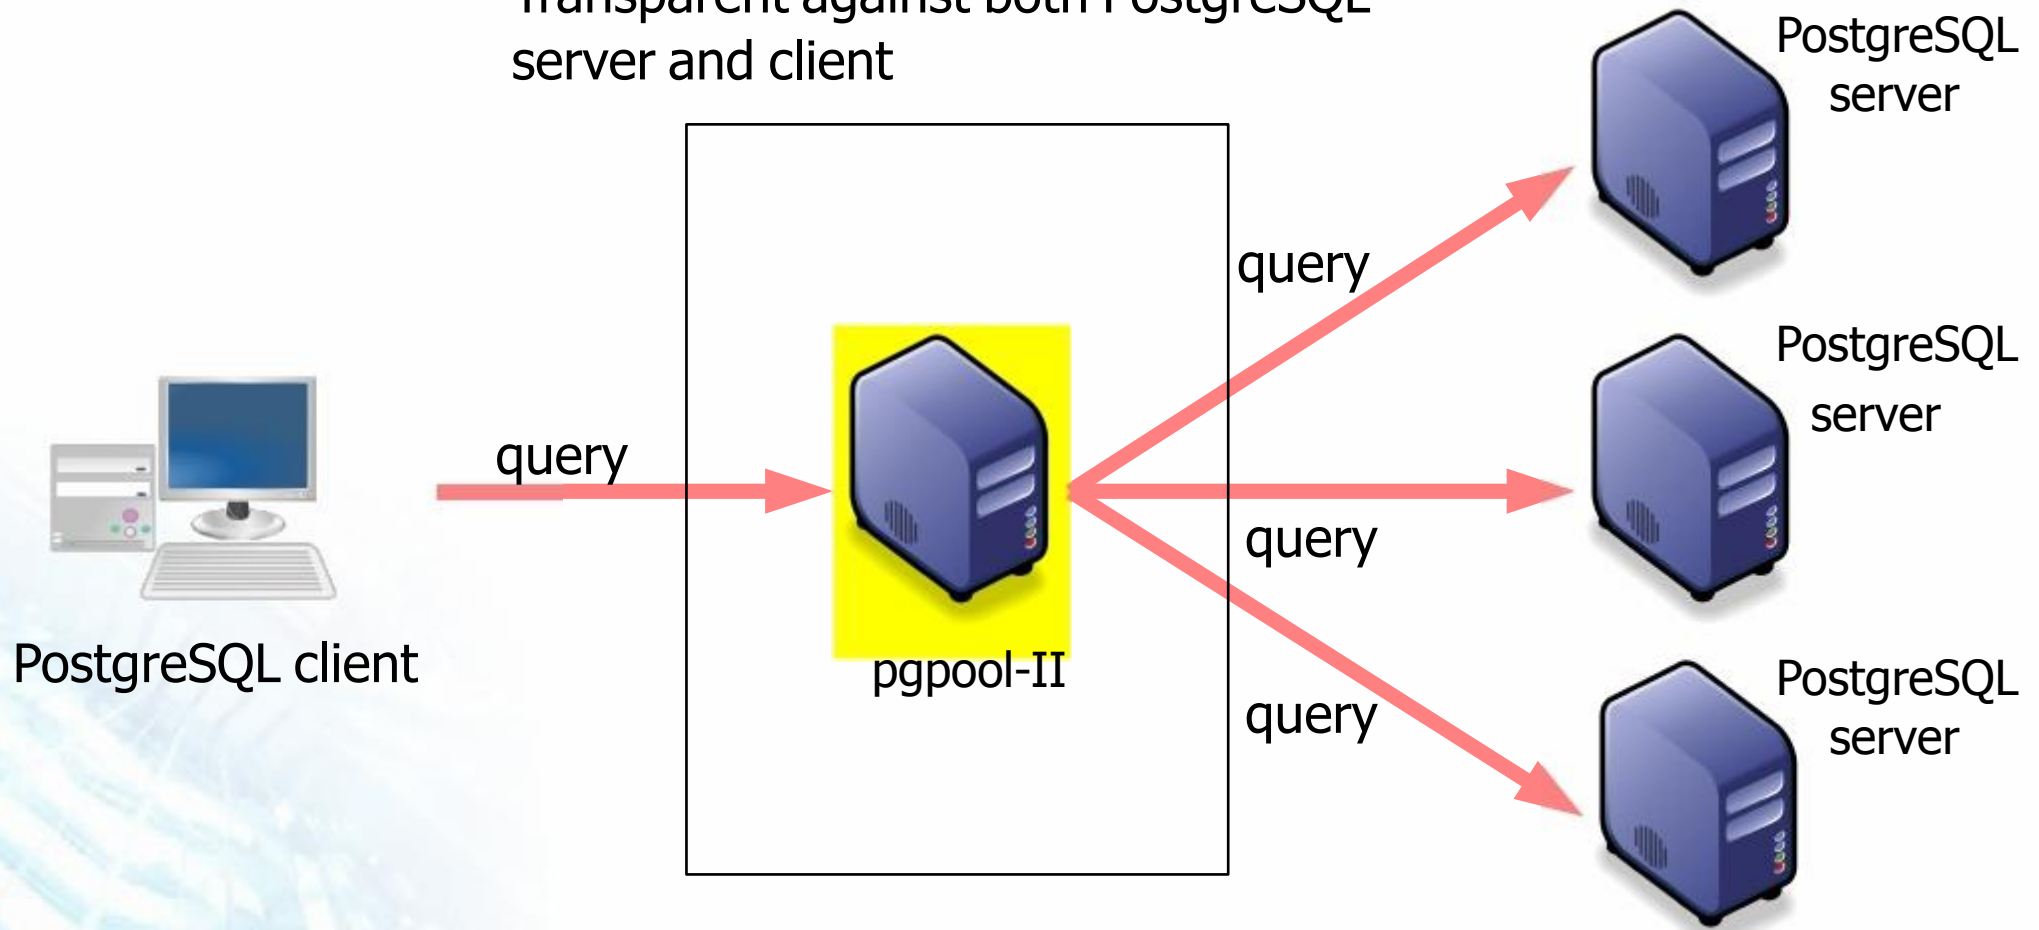


What is pgpool-II?

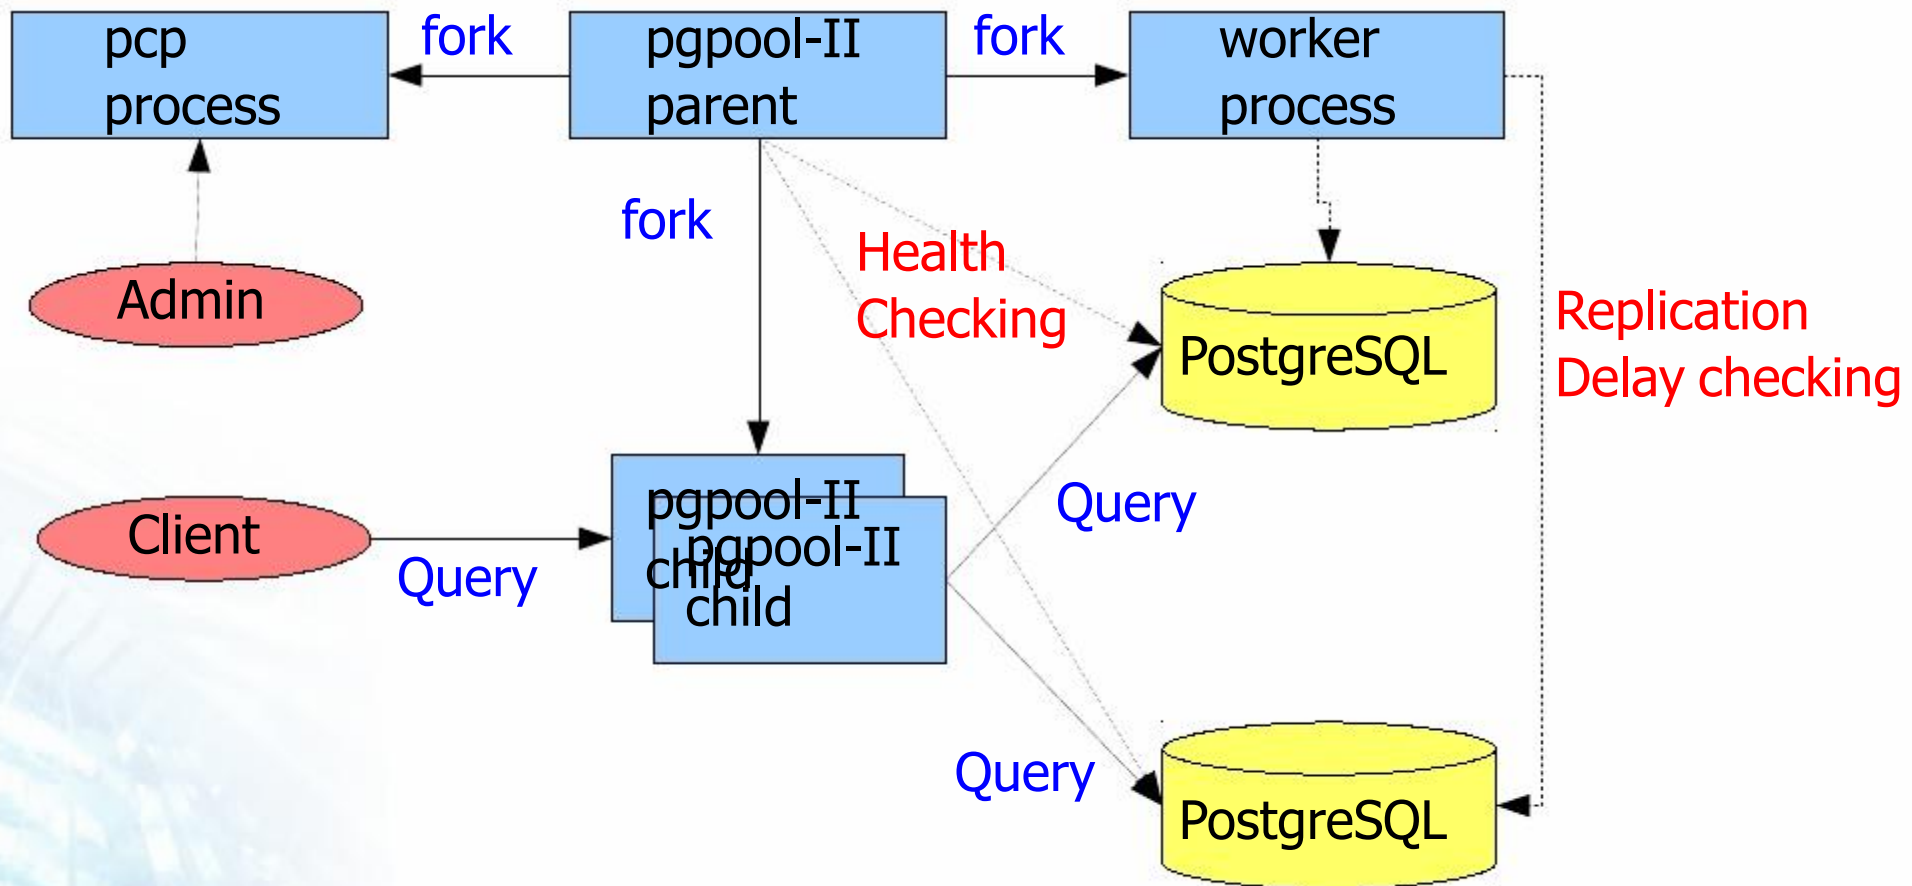
- A cluster management tool dedicated for PostgreSQL
- OSS project(pgfoundry, BSD License)
- Rich features
 - Synchronous replication
 - Load balancing, automatic failover, connection pooling etc.
 - Collaborating with other replication tools
 - Streaming replication, Slony-I
- Working as proxy between database clients and PostgreSQL

Basic idea of pgpool-II

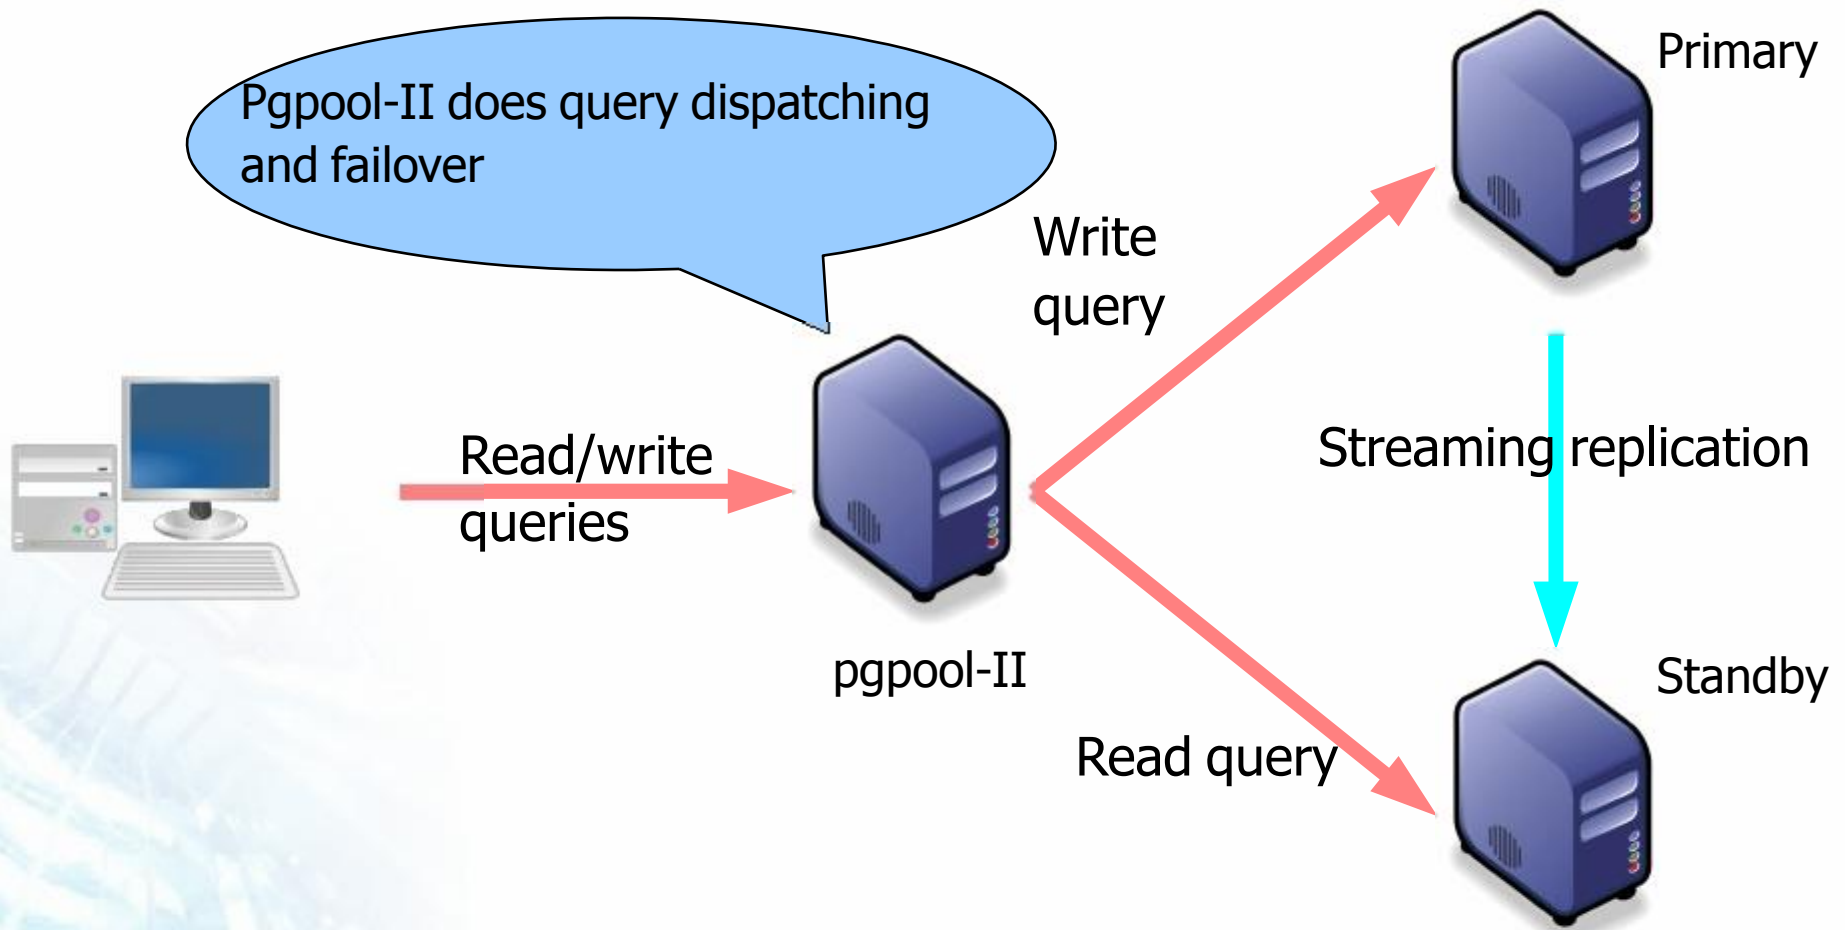
Transparent against both PostgreSQL server and client



The architecture of pgpool-II



Pgpool-II helps Streaming replication

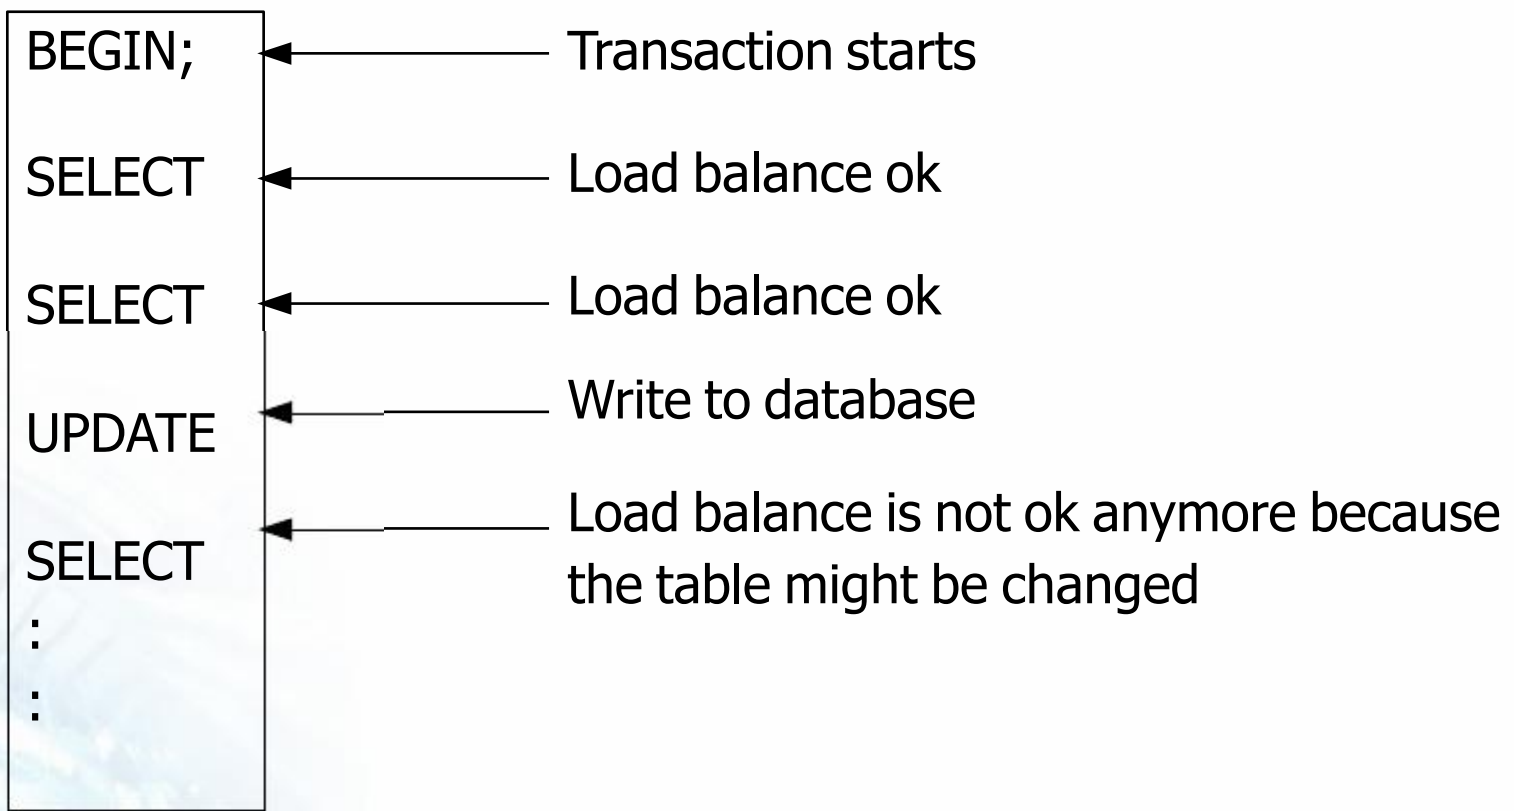


Pgpool-II helps users who use streaming replication

- Automatic query dispatching
 - Redirect DDLs, DCLs, DMLs and some commands(e.g. LOCK) to primary
 - Allow to specify functions which are writing to DB(they need to be executed on the primary)
 - Allow to load balance SELECTs
- Boost performance by enabling connection cache
- Automatic failover if one of database servers goes down
- Recover down node without stopping pgpool-II(online recovery)
 - Adding standby servers without stopping pgpool-II

Load balance SELECTs in an explicit transaction

transaction



Monitoring replication delay

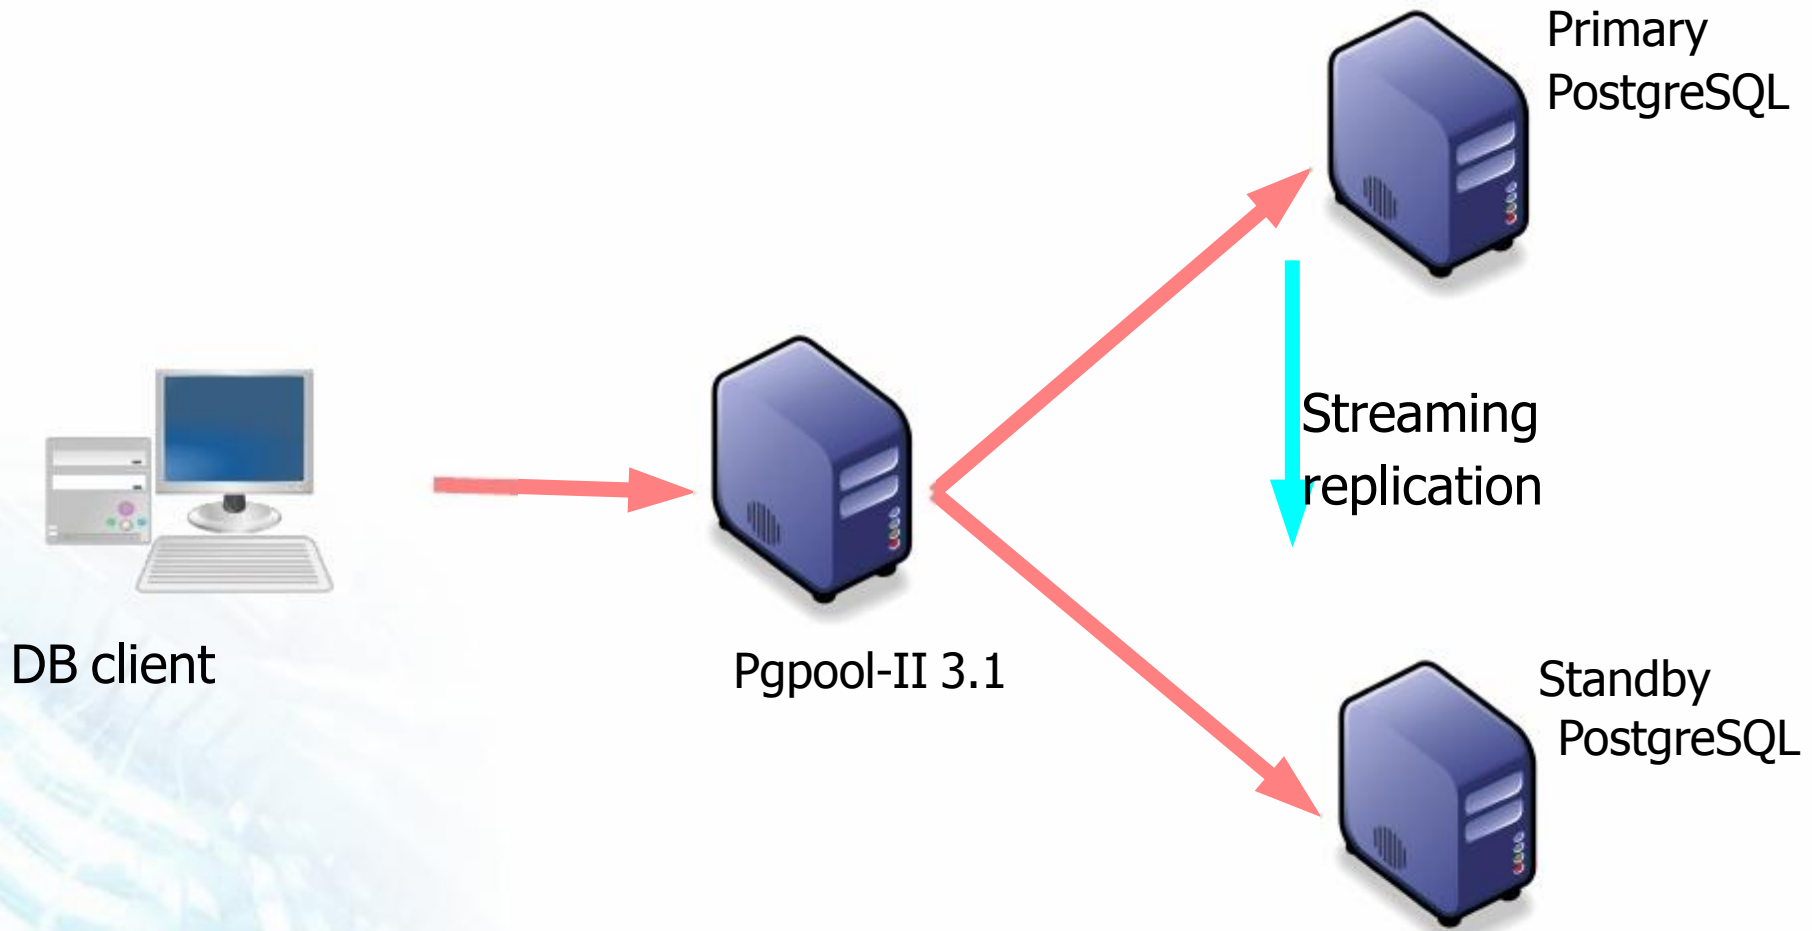
- log_standby_delay
 - 'none': no logging delay
 - 'if_over_threshold': log only when delay exceeds delay_threshold
 - 'always': always log delay
- If replication delay is too much, do not send SELECTs to standby

```
2010-06-28 15:51:32 LOG: pid 13223: Replication of node:1 is behind 1228800 bytes from the primary server (node:0)
2010-06-28 15:51:42 LOG: pid 13223: Replication of node:1 is behind 3325952 bytes from the primary server
(node:0)
2010-06-28 15:51:52 LOG: pid 13223: Replication of node:1 is behind 974848 bytes from the primary server (node:0)
2010-06-28 15:52:02 LOG: pid 13223: Replication of node:1 is behind 2990080 bytes from the primary server
(node:0)
2010-06-28 15:52:12 LOG: pid 13223: Replication of node:1 is behind 901120 bytes from the primary server (node:0)
2010-06-28 15:52:22 LOG: pid 13223: Replication of node:1 is behind 2433024 bytes from the primary server
(node:0)
```

What's new in pgpool-II next version (3.1)?

- Allow to use regular expression in black/white_function_list
- Allow to use syslog
- Better handling for more than 1 standby case
- After finishing online recovery, pgpool-II does not restart child process any more (thus client sessions are not disturbed)
- Import PostgreSQL 9.0 SQL parser

Demonstration



URL etc.

- Pgpool-II can be downloaded here:
 - <http://pgfoundry.org/projects/pgpool/>
- Twitter
 - @pgpool2
- SRA OSS's website
 - http://www.srasoss.co.jp/index_en.php