



# EDB & PGPOOL Relationship and PGPOOL II 3.4 Benchmarking results on AWS

May, 2015



- **Ahsan Hadi**
- **Senior Director of Product Development with EnterpriseDB**
- **Managing PGPOOL II development from EDB**

# EnterpriseDB Product Overview

## Cloud Enablement

Cloud Database for  
Amazon Web Services

Open Stack Support\*

APIs to enable the  
Platform as a Service\*

## Tools

xDB Replication  
Server

Back and Recovery  
Tool

Failover Manager

Migration Tool Kit

## Management and Monitoring

Postgres Enterprise  
Manager (PEM)

Stack Builder &  
Update Monitor

## The Database

### PostgreSQL / Postgres Plus Advanced Server

#### Core

PostgreSQL  
OR

Postgres Plus  
Advanced Server

#### Extension

Components (PostGIS, pgPool, pgBouncer, SQL/  
Protect, PL/Perl Python TCL, FDW)

Connectors (.Net, JDBC, ODBC, OCL)

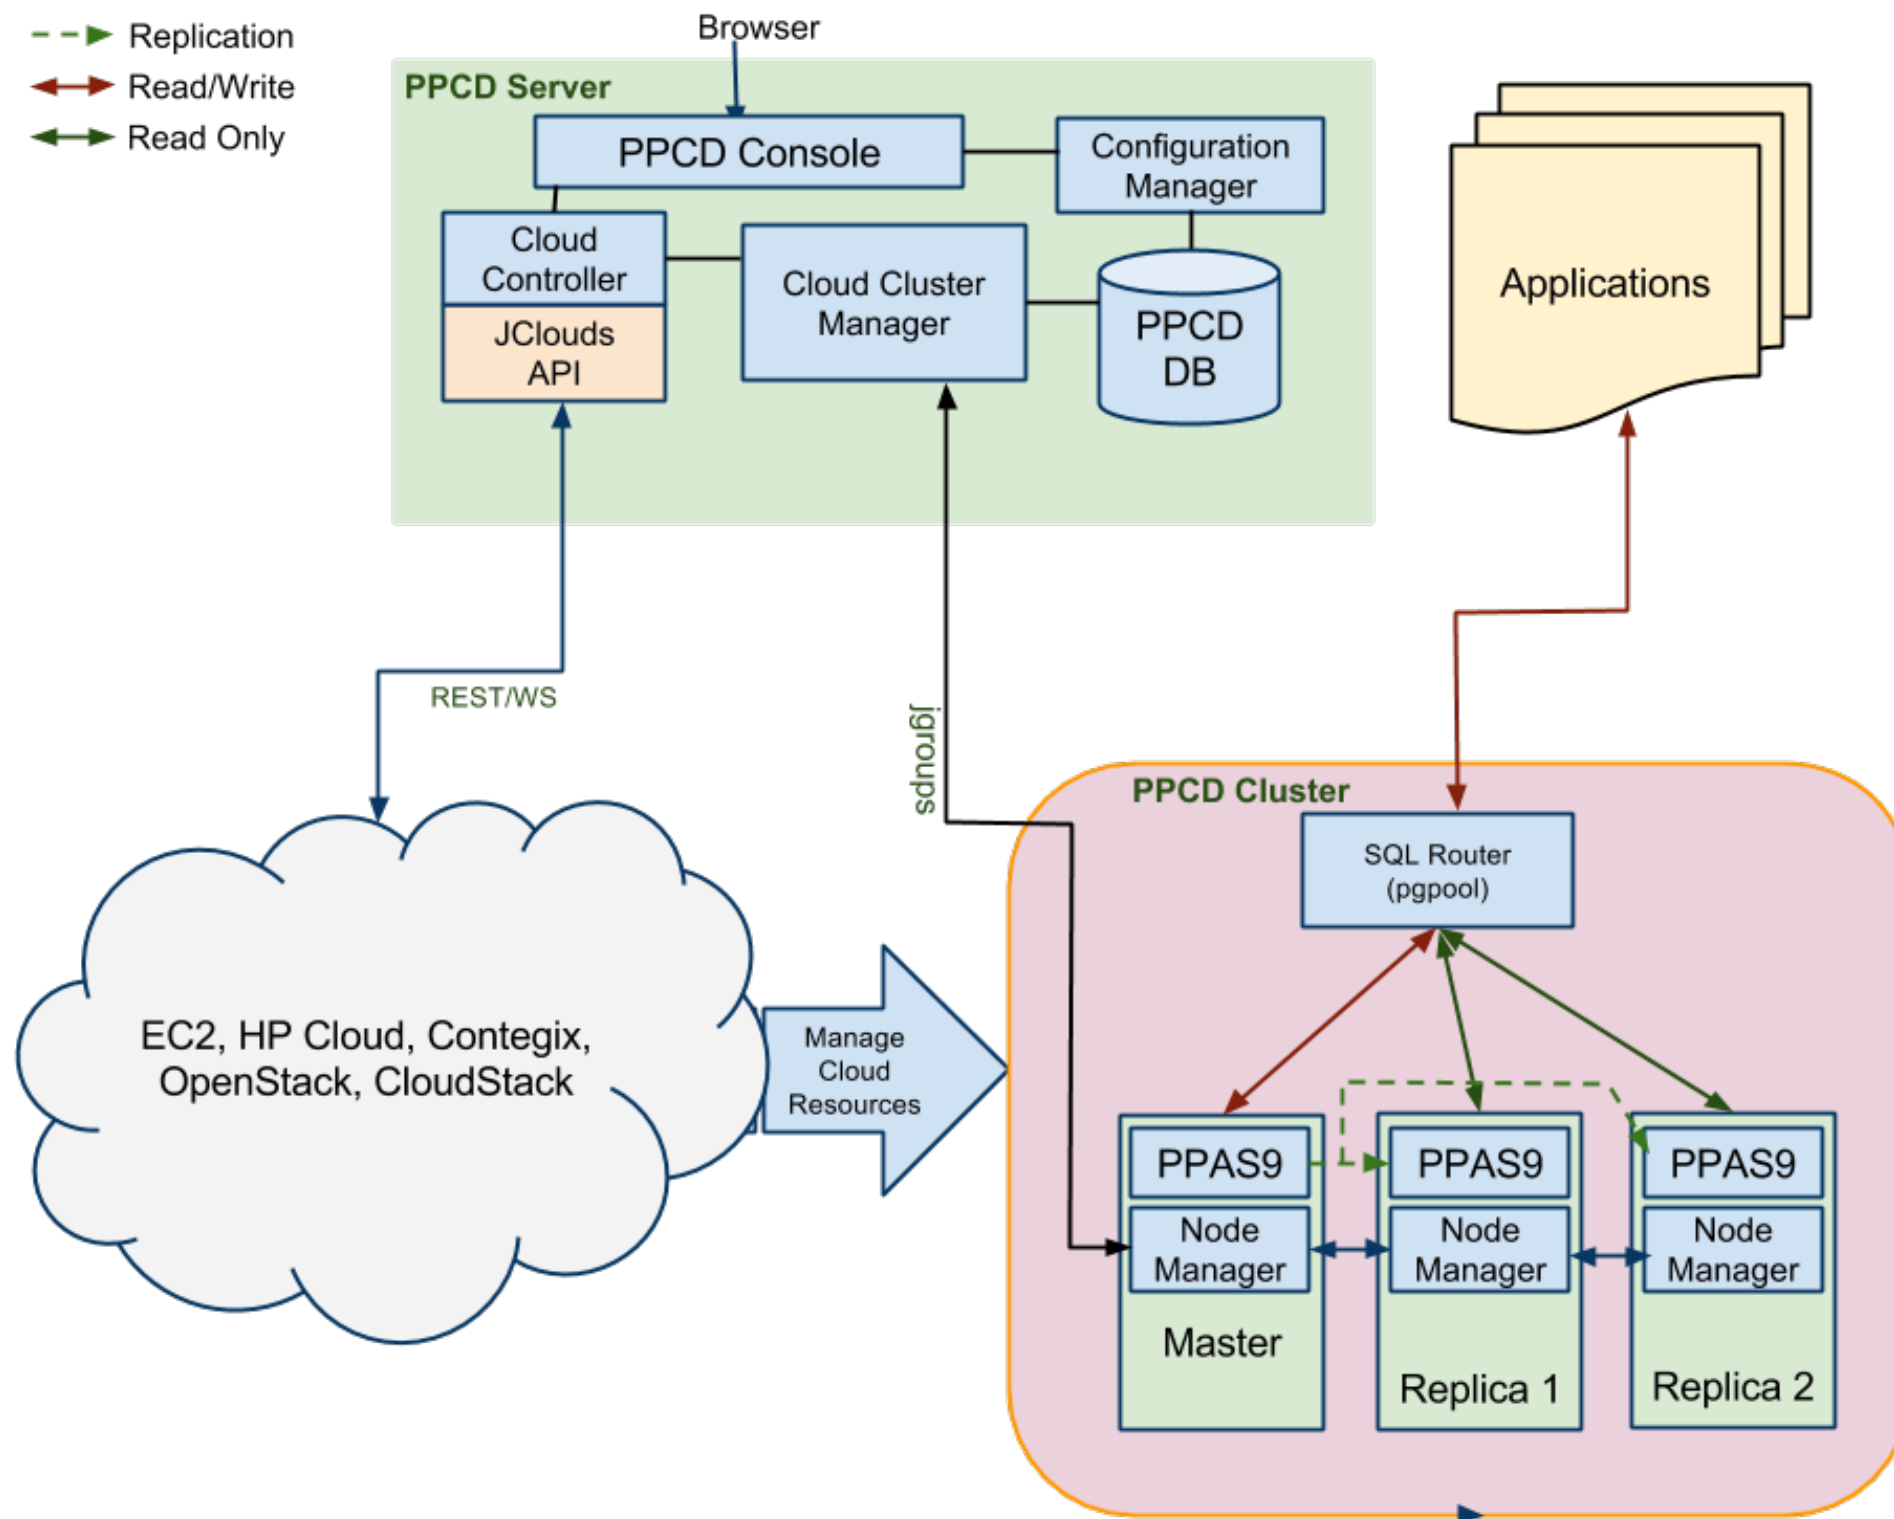
Utilities (EDBLoader, EDBPlus, ECPGPlus)

\* Roadmap

About my 

# How EDB uses PGPOOL II

# PPCD - pgpool-II in the cloud



# PPCD/pgpool II...

- **PPCD allows the user to host an instance of PG and/or PPAS in the cloud.**
- **PGPOOL II runs on the master node**
- **Client apps connect to PGPOOL.**
- **Performs load balancing by sending the write queries to master and read queries across read replicas**
- **Performs connection pooling**

# Other services

- **EDB offering support and services for pgpool II**
  - **Some BIG companies using pgpool II**
    - **US Based Credit card company**
    - **Indian Government Organisations**
    - **Healthcare**
    - **Telecom**
    - **IT solution providers and consultancy services**
- and More....**



# Most commonly used features

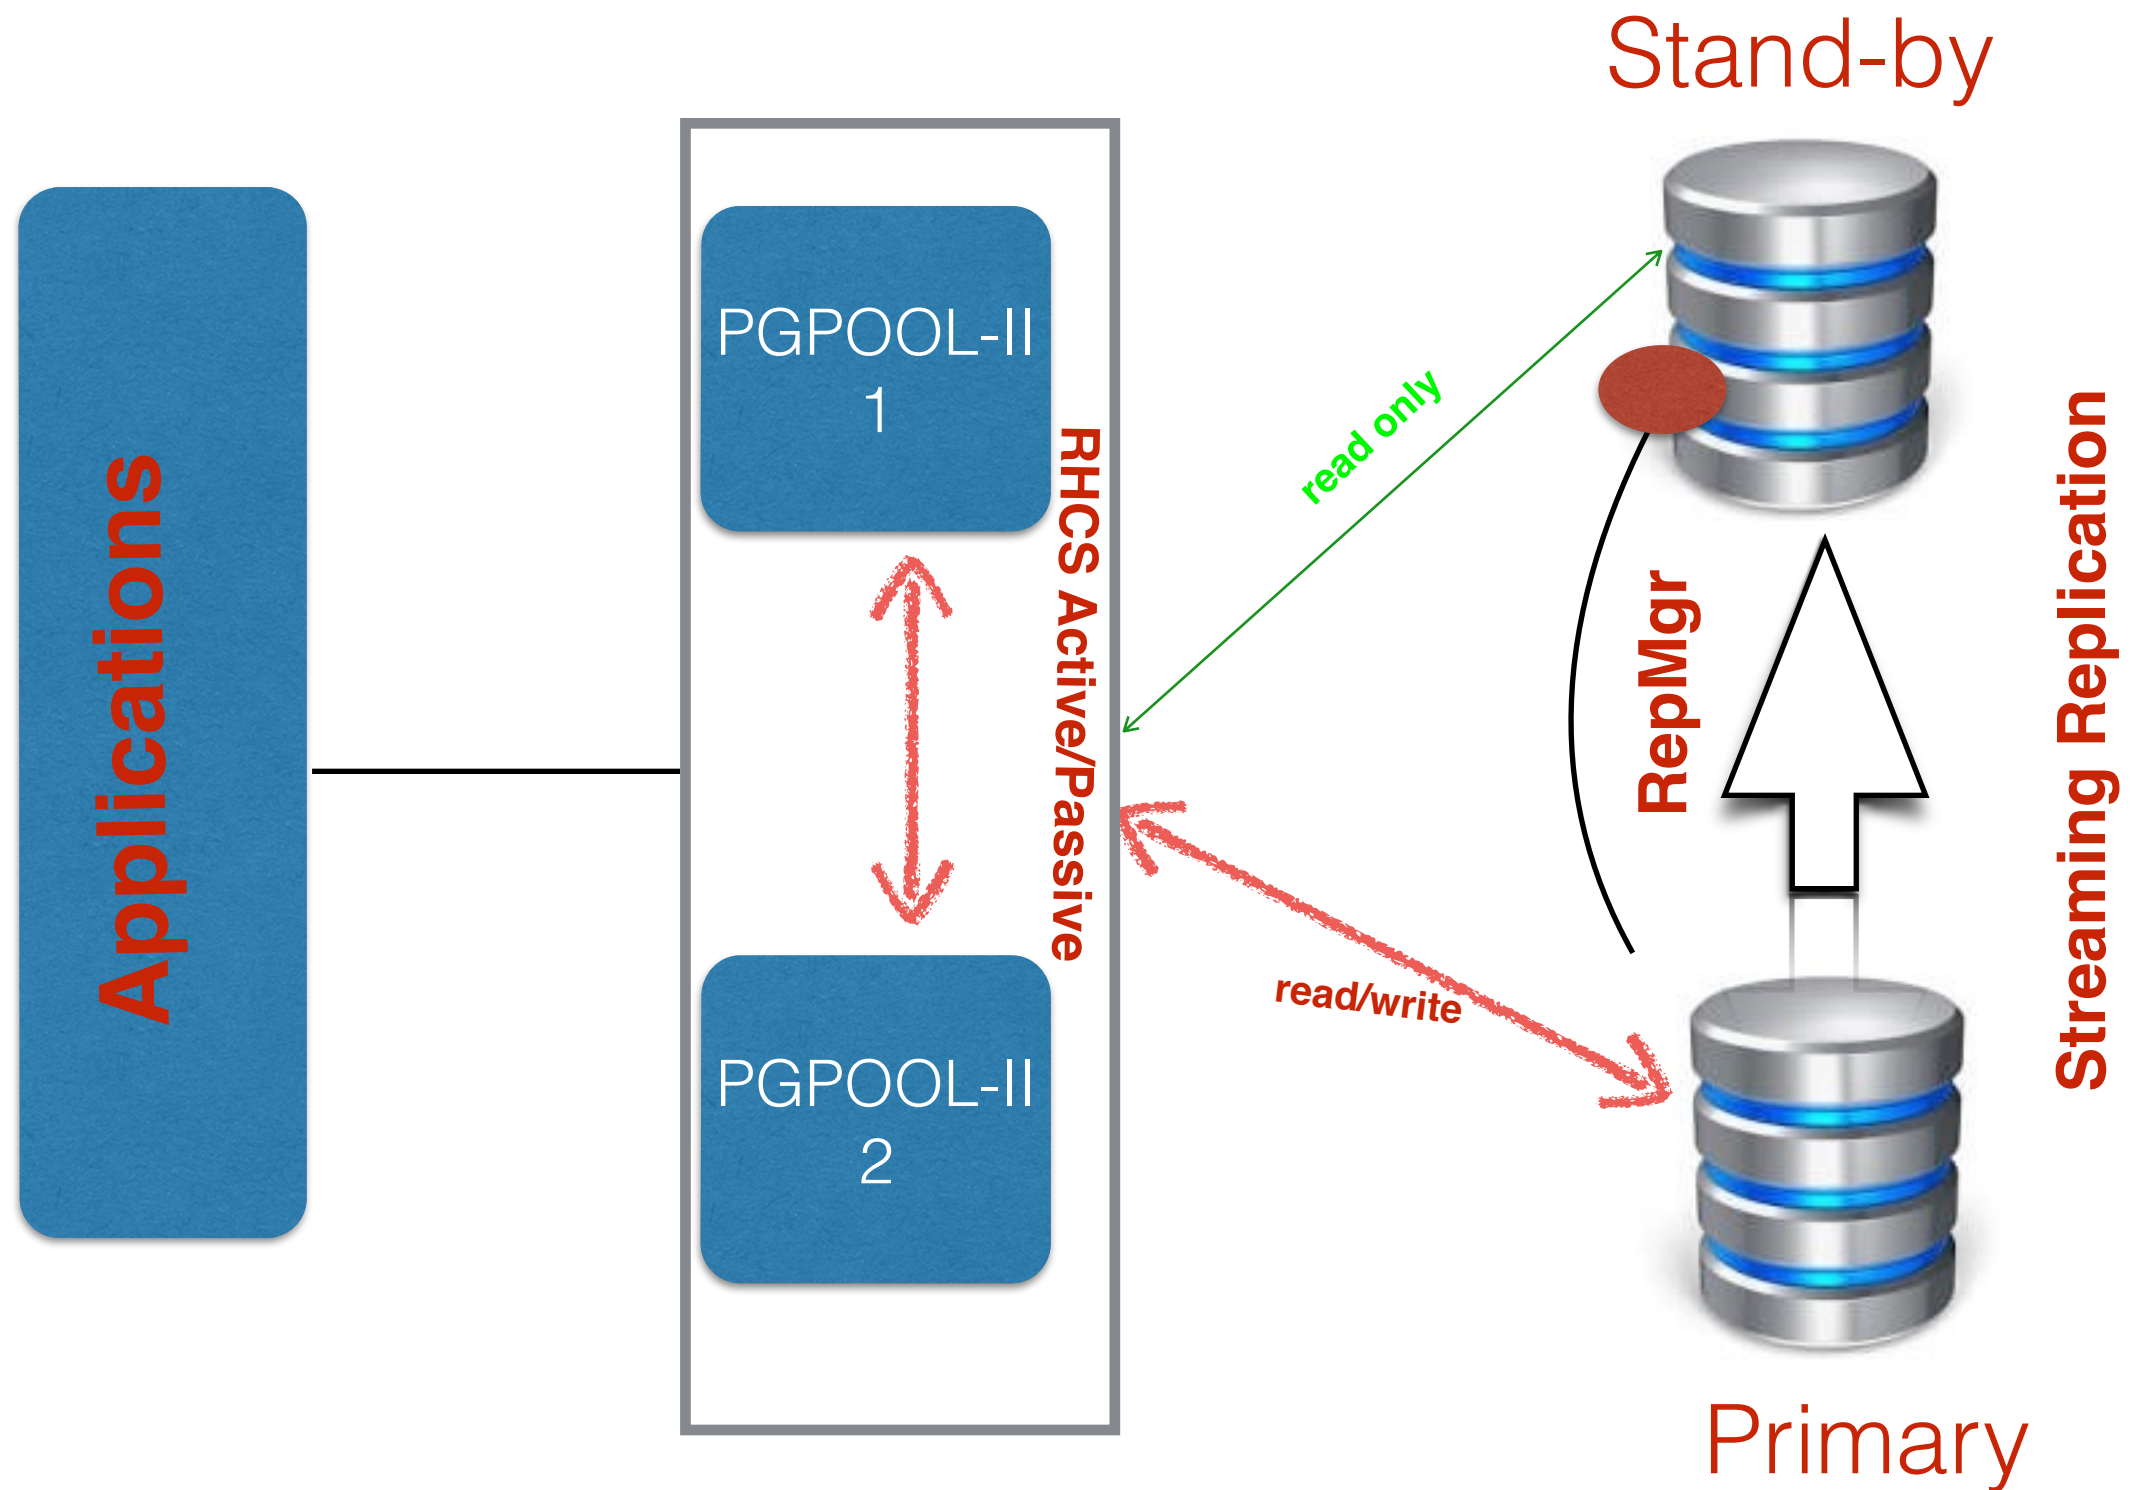
In order of priority...


- **Load Balancing (sql routing)**
- **Failover with streaming replication**
- **Connection Pooling**
- **Watchdog PGPOOL II HA**

# Features not used

- **Parallel Query**
- **PGPOOL Replication**

# A production customer scenario







- Overall it has gotten very stable and resilient in last few major releases.
- Types of customer issues reported in last two years (most of these issues have been resolved)
  - Handling SSL based authentication
  - Performance issues with extended query protocol (performance not as good as simple query protocol)
  - Issues with failover (not able to connect to stand-by immediately when failover happens)
  - Sequence functions don't work with pgpool
  - Performance issues with load balancing, the TPS is not very high
  - Configuration issues
  - Memory leaks (memory issues addressed by plugging in PG memory and exception management)

# WISHLIST !!

- **Improve performance with extended query protocol**
- **Watchdog enhancements (making it more resilient and handle scenarios like split-brain, automatic failover)**
- **Being able to install individual features of pgpool...I am only interested in connection pooling...**
- **Enhancements with Failover functionality..seamless failover**
- **Support more platforms, HPUX, solaris, windows...**
- **GUI Installer and GUI monitoring tool**
- **Better documentation**

# Benchmarking PGPOOL-II 3.4 on AWS - Environment

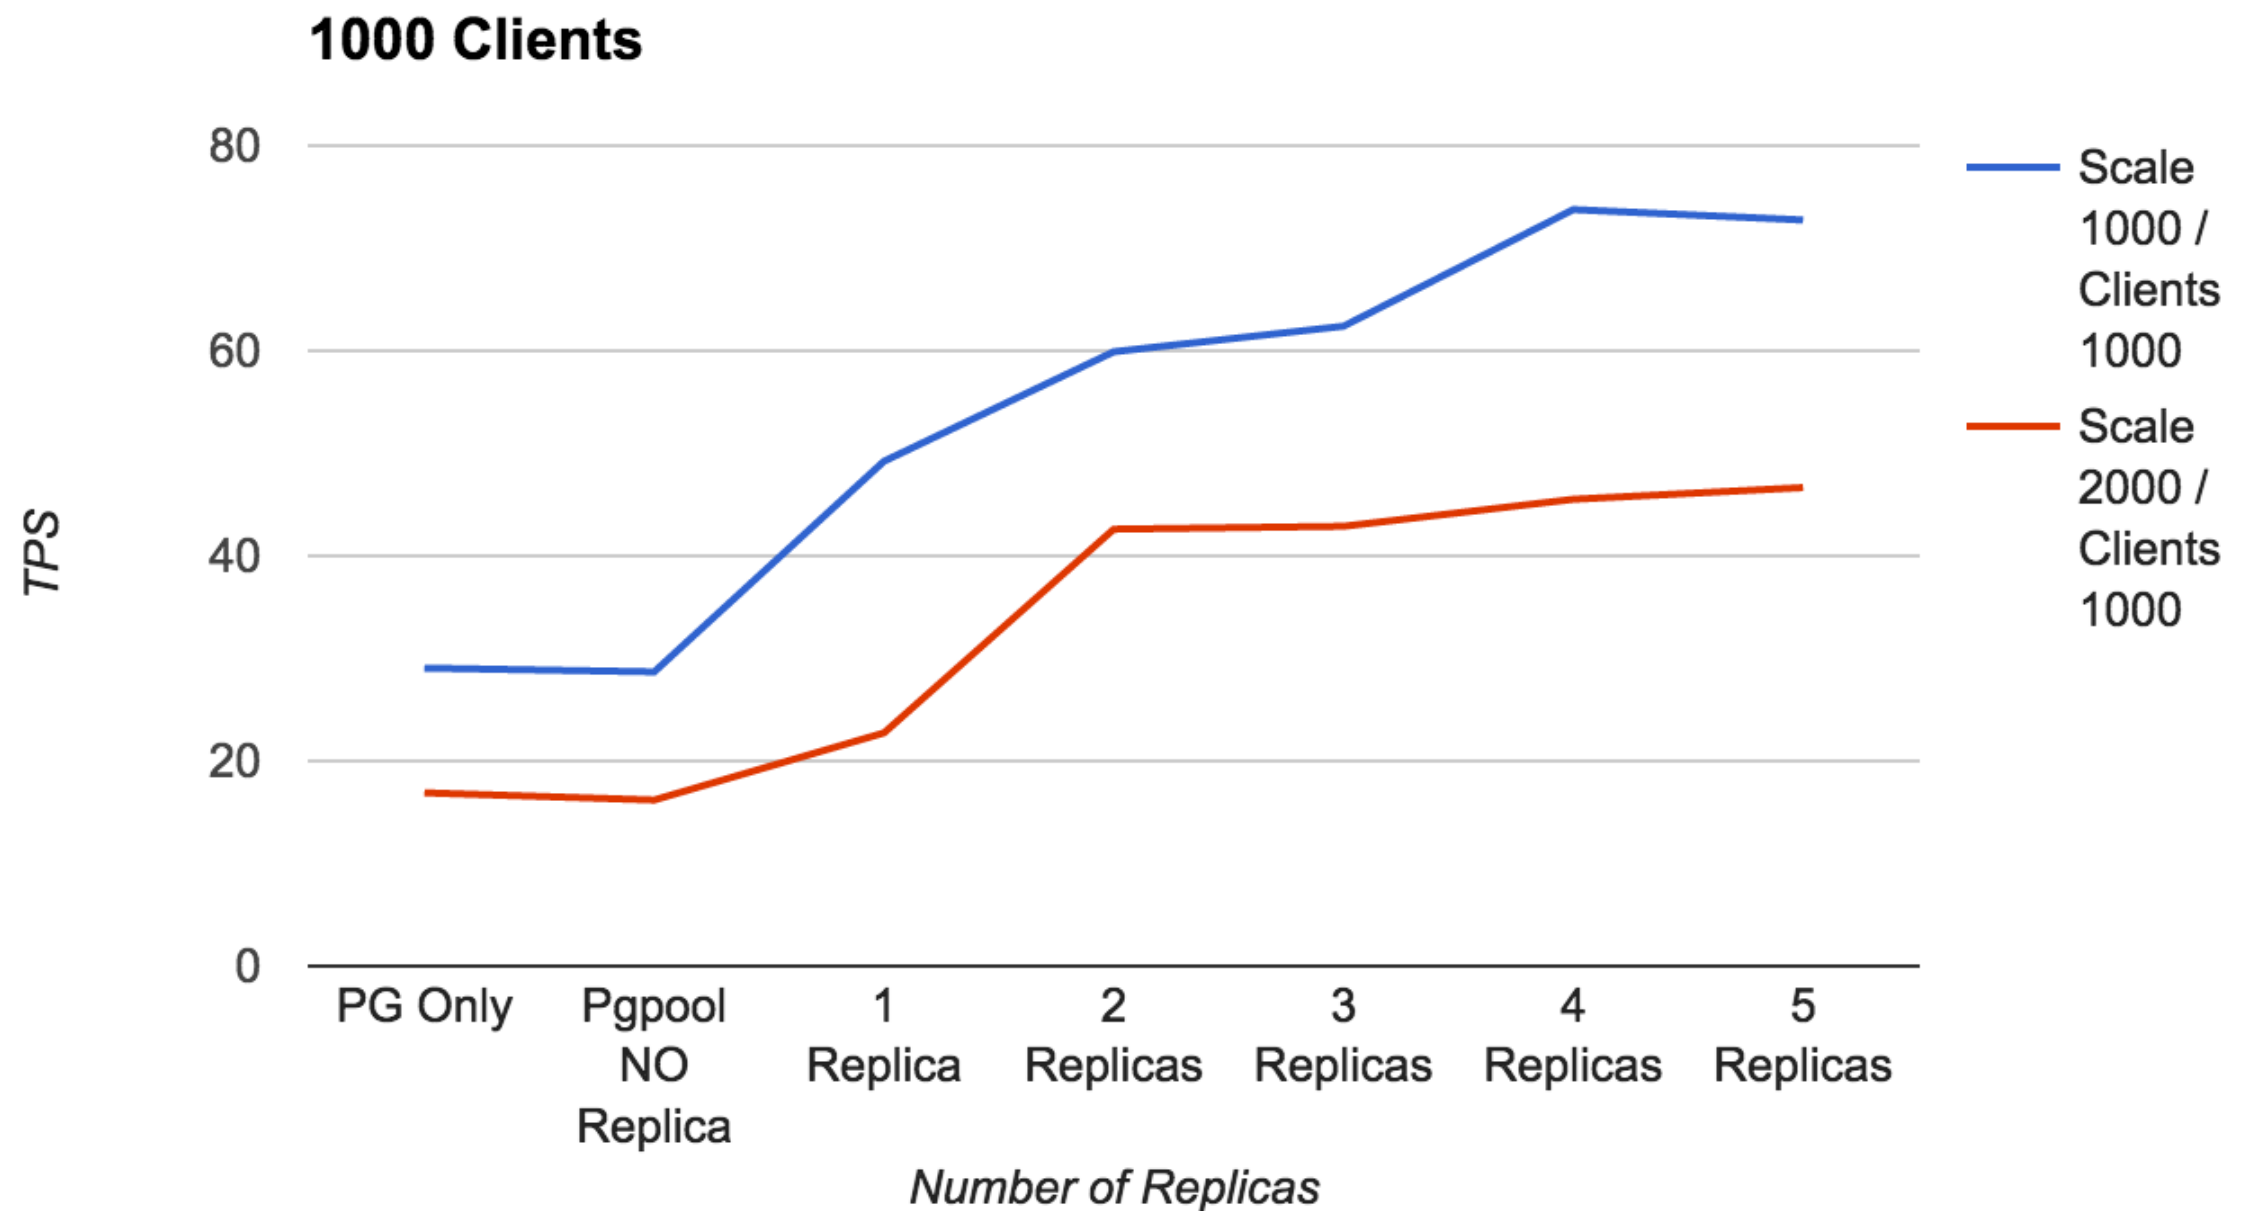
- PGPOOL-II version 3.4
- PG version 9.3.5
- Benchmarking tool pgbench
- Benchmarking done on AWS instances
- One pgpool-II instance with 5 read replicas
- Compared with PG running stand-alone
- PGPOOL-II configures with streaming replication, master/slave mode.
- Apples-to-Apples Comparison (Same optimal real-world PG configuration changes across standalone PG and pgpool-II nodes)

# Benchmarking PGPOOL-II 3.4 on AWS - Conclusions

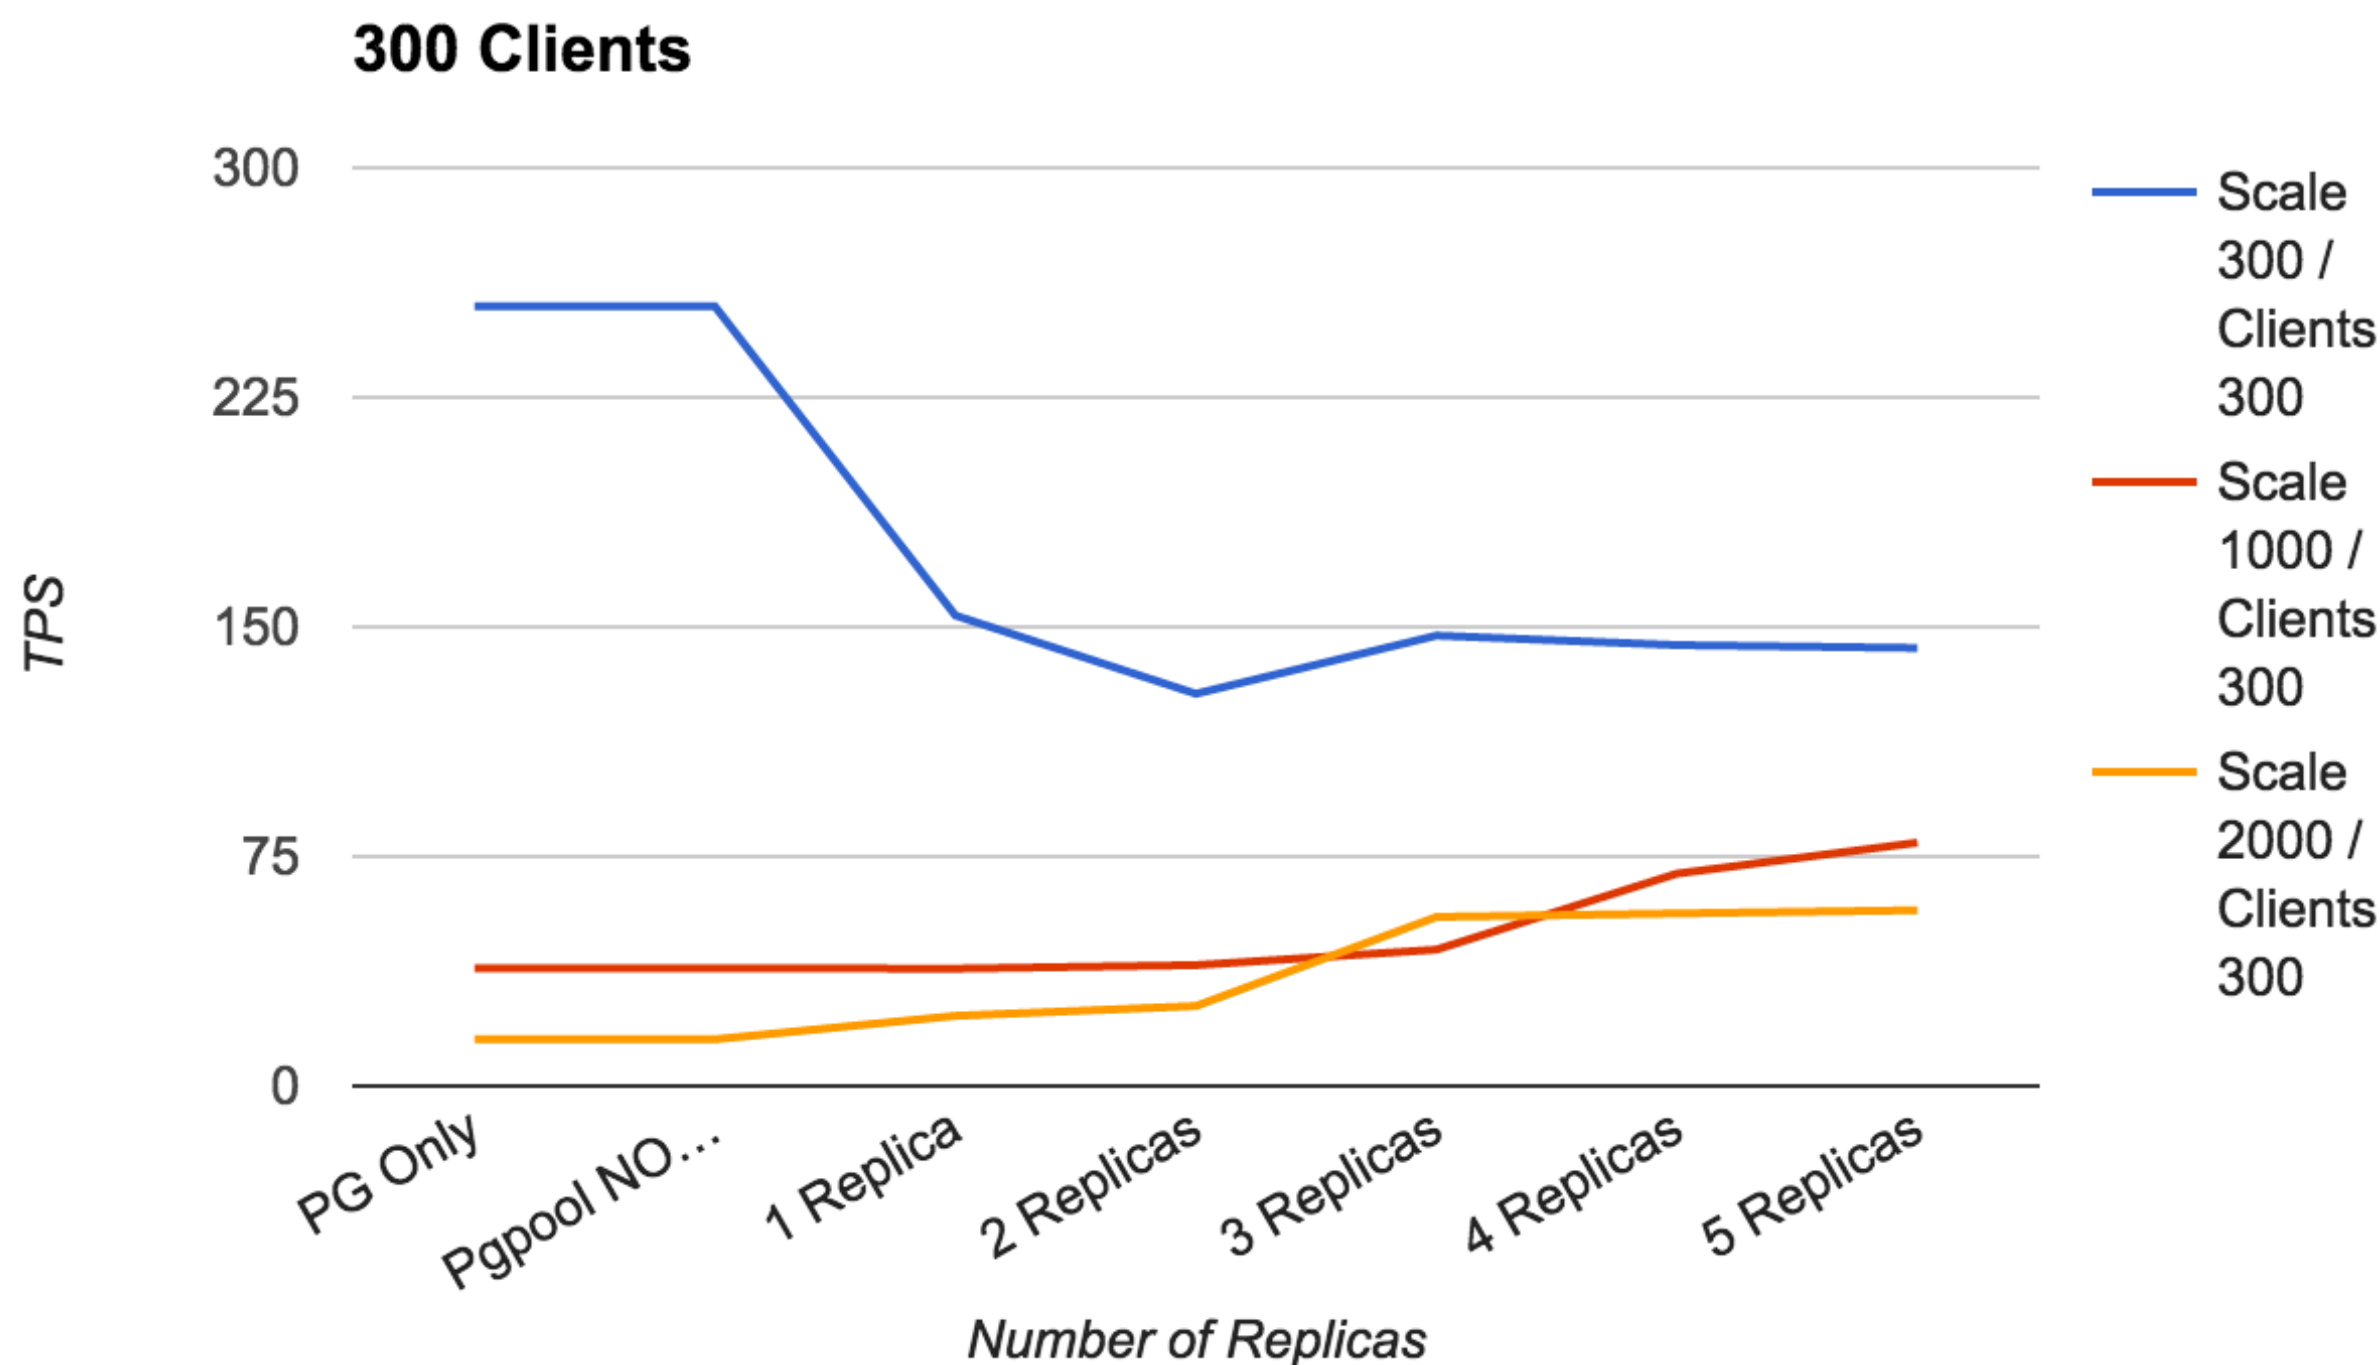
- Horizontal scalability when the database grows out of single server memory
- With big data, concurrent clients and 5 read replicas the TPS is 2/3 times then community PG
- After 4 read replica's the performance bar comes down with the given workload.
- Bad Performance with extended query protocol



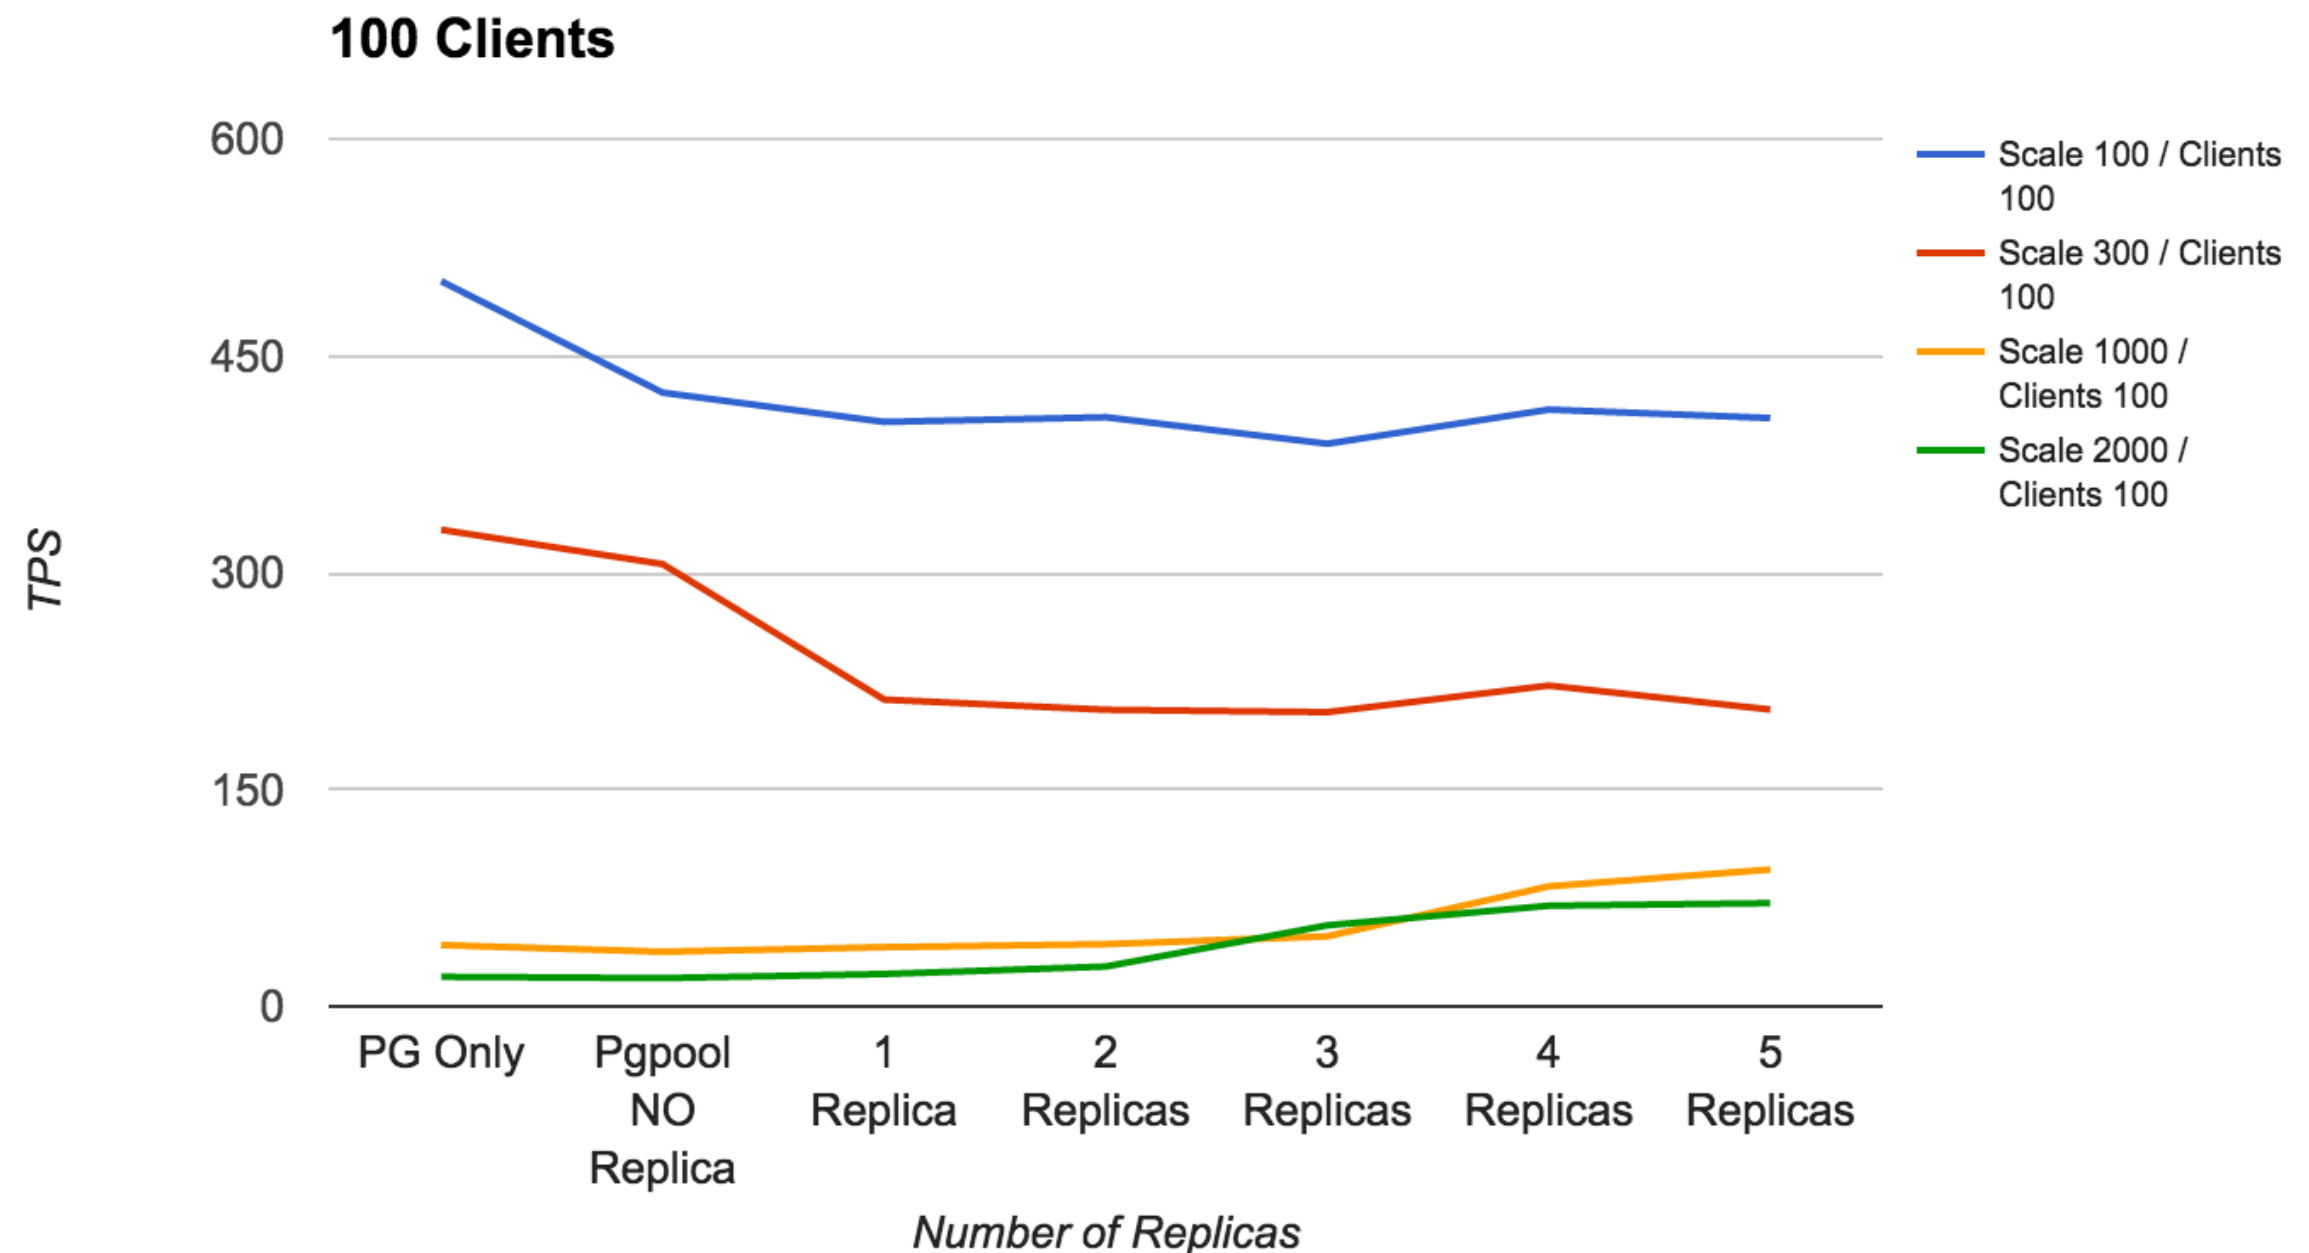
# Benchmarking PGPOOL-II 3.4 on AWS - Actual Results



# Benchmarking PGPOOL-II 3.4 on AWS - Actual Results



# Benchmarking PGPOOL-II 3.4 on AWS - Actual Results



# PGPOOL II Future

- Very Bright
- More remaining to be done for performance and failover management. Need to invest in tooling.
- Will become the middleware product of choice for Load Balancing, Failover and HA for PG.
- Need more funding and volunteers

THANK YOU

merci, grazie, spasiba, kam ouen, gratzias, manana, mahalo, cheers, toda, hvala, gracias, grassie, thank you, danki, kitos, welalin, tak, mahalo, danki, thanks, takk, domo arrigato, gratitude, danke, kitos, takk, dziekuje, miigwetch, talofa, modupe, mesi, na gode, dankon, gracias, merci, thanks, mahalo