

PostgreSQL Performance in Linux-containers and on top of Parallels Cloud Storage



Profit from the Cloud

Alexander Kirov

Program Manager, Parallels Cloud Storage



PGCONF.RUSSIA 2015

Agenda

1. About Parallels and Parallels products
2. Test stand description
3. PostgreSQL performance on native host
4. PostgreSQL performance in containers
5. About Parallels Cloud Storage
6. PostgreSQL performance on top Parallels Cloud Storage
7. How to try

> About Parallels

Что такое Parallels

Мощная основа

- Более 1000 сотрудников в мире
- Продажи в 125 странах из 16 офисов по всему миру
- Более 150 полученных патентов
- 500 партнеров, включая Microsoft, Apple, Intel, AMD, Dell, HP и IBM



Доказательства успеха

- В числе 100 крупнейших софтверных компаний мира
- Прибыльна, самофинансируема, 50% рост CAGR в течение 6 лет с ускорением темпов

Made in Russia

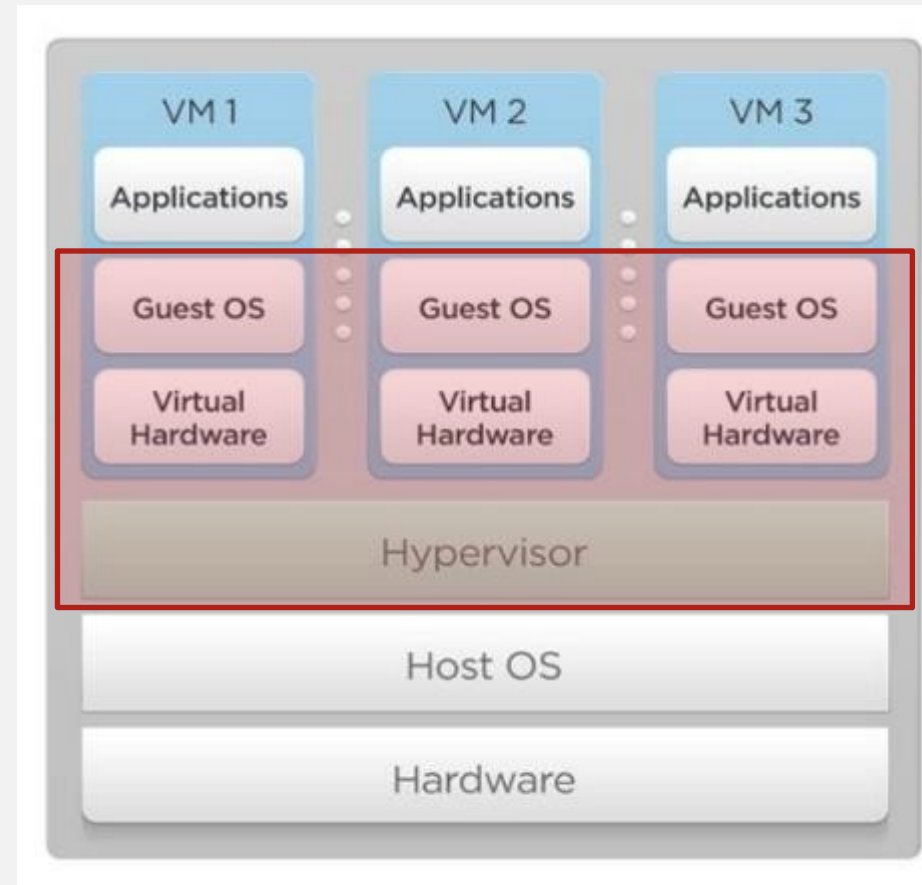
- Конкуренентоспособное ПО на мировом рынке
- Топ-4 крупнейших компаний с центром разработки в России (Kaspersky, Acronis, Parallels, ABBYY)
- Единственный в России разработчик законченных коммерческих решений по виртуализации
- R&D центры в Москве, Новосибирске и Санкт-Петербурге
- Parallels – участник проекта «Сколково»
- Основана выпускниками лучших ВУЗов страны: МФТИ, НГУ, МГУ, академии ФАПСИ и др.



> Parallels Products

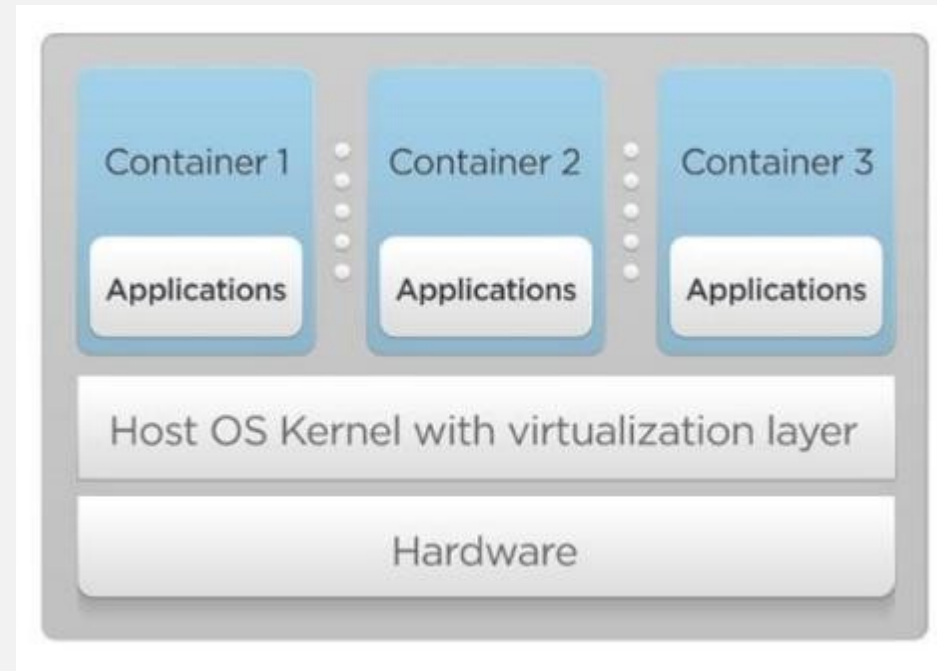
Parallels Hypervisor Virtualization

- OS flexibility
- HW emulation
- Bare metal installation



Parallels Containers

- More efficient memory management
- More efficient caching reduces I/O
- CT resource management
- Easy migration
- Easy backups&snapshots



Benefits for PostgreSQL in Containers

Many features are out of the box:

- Security: environment isolation
- Backups
- Snapshots
- Live-migration
- Limits & resources management
- Online resource management:
 - RAM hot-plug
 - CPU hot-plug
 - Online disk expand
- And **High performance**

> Test Stand Description

Test environments


1. PostgreSQL on **native hosts**
 - Synchronous Streaming Replication
2. PostgreSQL in **containers** on the host
 - Synchronous Streaming Replication in Container
3. PostgreSQL in **containers** on top of **Parallels Cloud Storage**
 - HA by Parallels Cloud Storage

Test Stand Hardware Description



CPU	1xQuadCore Intel Core i5 2400
Memory	16 GB (4x4GB), DDR3, 1333MHz
Disk	3x1TB 7200 RPM SATA drives
SSD	Intel S3700 100G SATA
Network Card	Intel 82599 10Gbit Network Connection
Network Switch	DELL PowerConnect 8132 10 Gb

Test Stand Software Description




PgBench



Operating system	Parallels Cloud Server 6 Update 8 (REHL 6 compatible)
Linux kernel version	2.6.32-042stab102.9
PostgreSQL version	9.4.0
Pgbench version	9.5devel

Pgbench configuration




- Databases initialization scale factor - 1000
- All tests were run with 50 clients
- Test time 600 seconds
- Each test was running 5 time

Test client PgBench command line for your homework:


```
./pgbench -h 172.0.14.10 -p 5432 -U postgres -c 50 -P 1 -T 600 bench 2>/dev/null
```

> PostgreSQL performance on native hosts and in containers


Test Stand #1 Description - “Native hosts”




Test Stand #2 Description – “Containers”



Single Test Node Structure



Results: PostgreSQL on native host vs containers



PostgreSQL in containers shows almost the same results as PostgreSQL shows on host

> Parallels Cloud Storage

Parallels Cloud Storage

The benefits of a SAN without the cost

Cost-effective

- Reuses existing hardware
- No vendor lock-in
- No need for dedicated, expensive hardware

High performance


- Automatic data and load balancing
- Out-of-the-box SSD caching improves read and write performance
- Improve overall throughput by utilizing idle and underutilized hardware

Scalable


- Expands easily by automatically identifying new additional nodes and devices and rebalancing the cluster
- Server storage capacity and bandwidth no longer limited by individual nodes

Fault tolerant and redundant

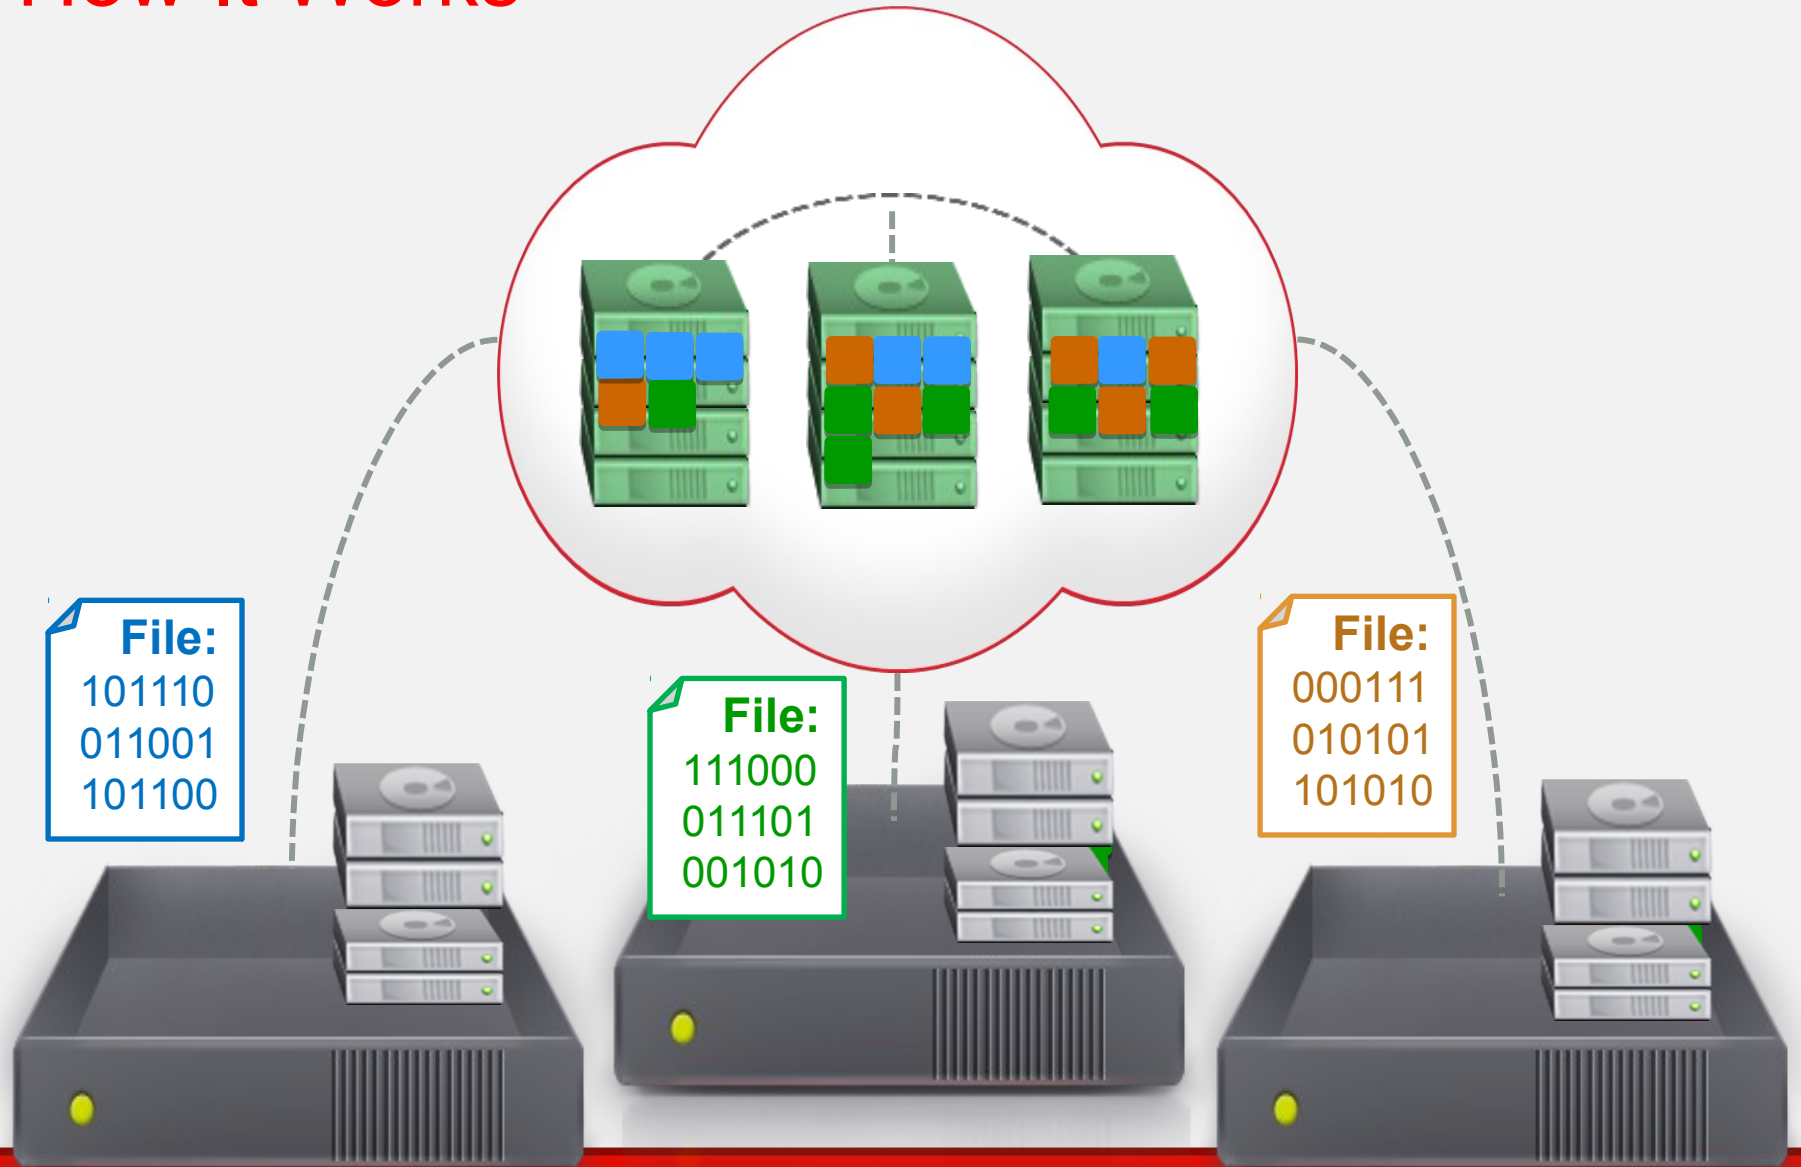
- Automatically manages redundancy and recovery
- Transparent data replication across storage cluster to maintain the required redundancy level




How It Works



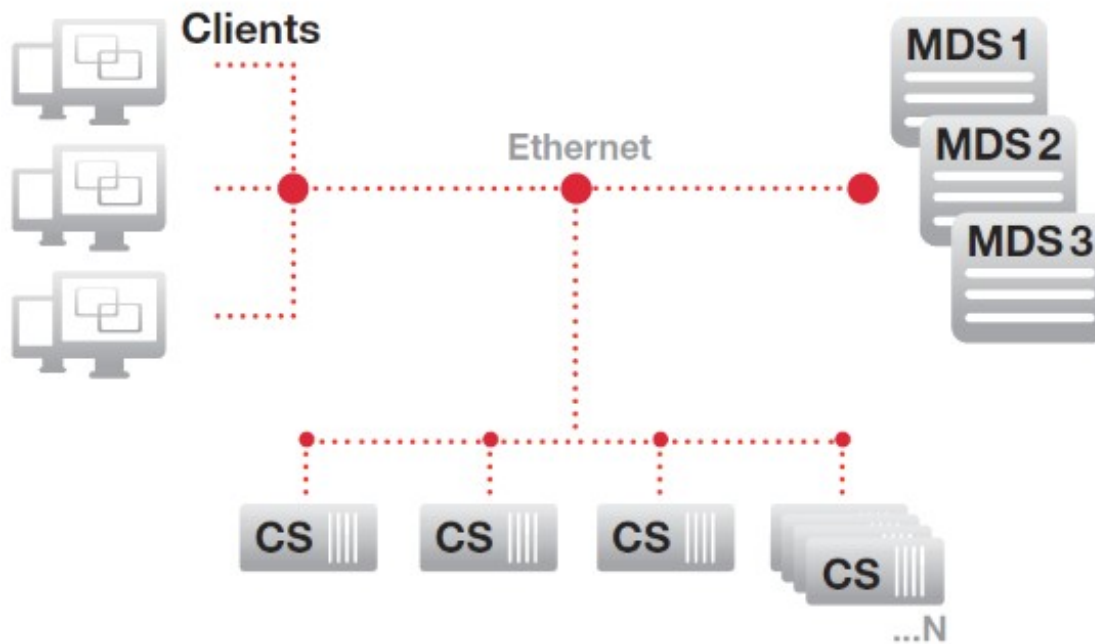
How It Works



How It Works



PStorage architecture




Meta Data Server (MDS)

- Stores metadata in memory
- Tracks data chunks and their versions
- Is highly available
- Can run on the same server as the chunk server and client

Chunk Server (CS)

- Stores data chunks
- Manages data chunks
- Performs read/write operations on data chunks
- Can run on the same server as the client

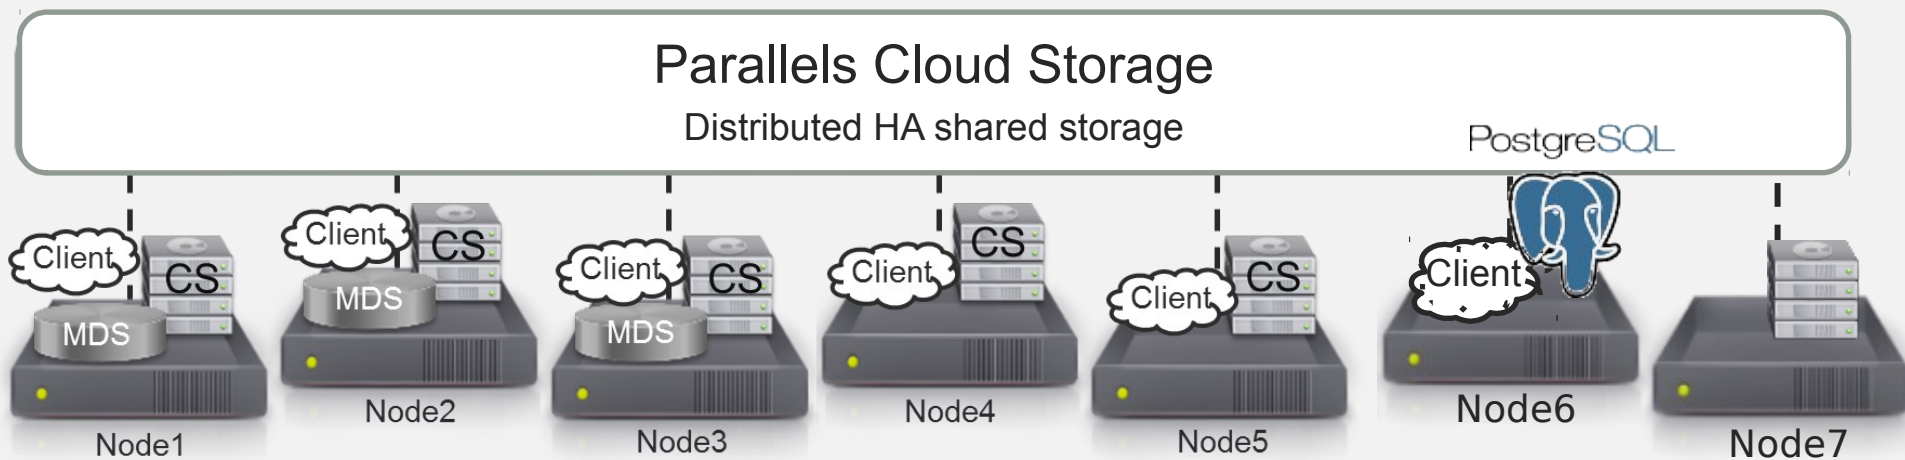
System Requirements



System Requirements:

- MDS: 1 GB free memory for each 100 TB stored data
- CS/client: 256 MB free memory
- Dedicated storage network (1GigE or faster)

Flexibility & Scalability



- Any mix of roles in one node
- Any mix of HDDs in a cluster
- Easy to add new node
- Scalable up to hundreds nodes
- Scalable up to PBs
- Automatic data and performance balancing
- The bigger cluster – the better performance
- The bigger cluster – the better reliability

High Availability (HA)


Parallels Cloud Server enables HA of VMs and Containers. When a node fails, all CTs/VMs are distributed to working nodes

Relocation algorithms:

- DRS (based on resource usage)
- To spare node
- Round-robin

Fencing:

- Node reboot
- Network disconnect



> PStorage Performance



Performance Comparison :: HW

iSCSI SAN Storage DELL EqualLogic PS6510E



x1 HW SAN EQL PS6510E

48 SATA Disks: 1TB 7200rpm
(Seagate ST31000524NS)

Network: 10Gbit
(Dell Force10 S4810)

VS.

Parallels Cloud Storage




x10 compute nodes

30 SATA Disks: 2TB 7200rpm
(Seagate ST2000DM001)
+ 10 SSD for caching
(Intel SSD 520)

Network: 10Gbit
(Brocade FastIron SuperX SX-F42XG)

PCS FASTER than HW SAN


Just 10 nodes
PCS cluster
faster than DELL
EQL SAN
(\$97000) in most
workloads



Parallels Cloud Storage Reliability Performance

High availability and fast recovery

- **Replication performance** defines time of reduced redundancy level. Critical!
- The shorter period of replication leads to the less probability to **lose data**
- In theory **Mean Time To Data Loss** $\sim 1/T^2$, where **T** is time to recover



The bigger the cluster, the faster recovery process is!

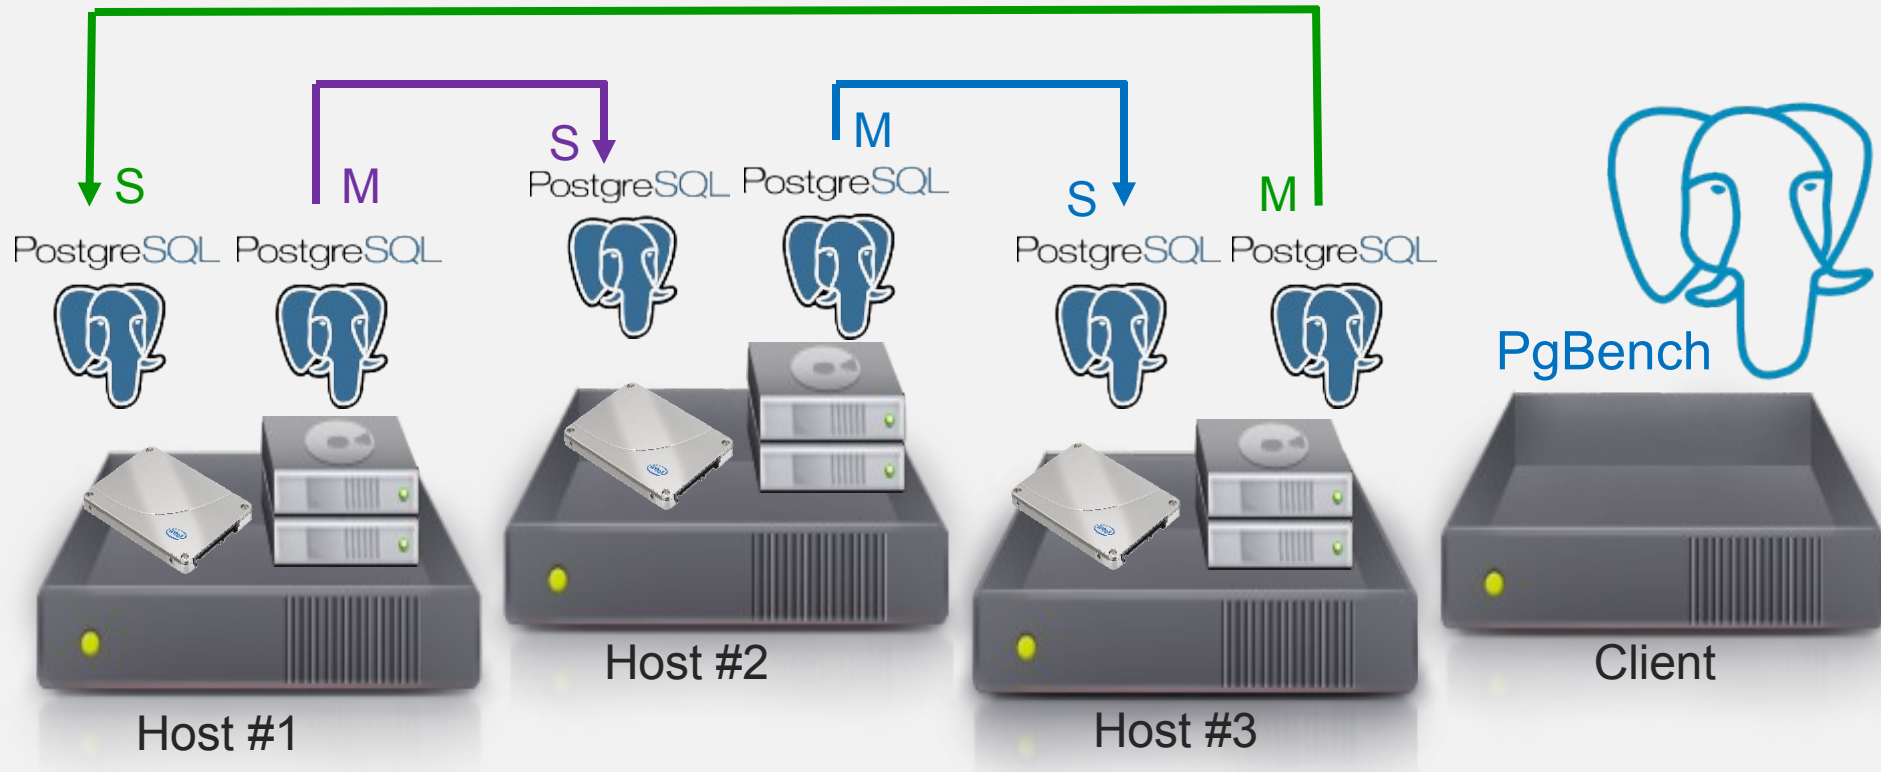
Cloud Storage performs replication in parallel.

Benefits for PostgreSQL


- Replication on storage layer
- No special PostgreSQL configuration
- Automatic cluster recovery
- Fast replication/recovery
- Instant live-migration
- Access whole cluster space from single server
- Access whole cluster IOPS
- Easy add new disks
- Built-in SSD-caching/journaling
- And **High performance**

> PostgreSQL performance in **containers**
on top of Parallels Cloud Storage


Test Stand Description



Test Stand Hardware Description



Results and Conclusion



PostgreSQL in container on top of Parallels Cloud Storage shows almost the same results as PostgreSQL shows on host

Conclusion

- PostgreSQL in Parallels containers shows almost the same results as PostgreSQL on host
- PostgreSQL on top of Parallels Cloud Storage shows almost the same results as PostgreSQL on host
- Containers provide security, resource management, migration, backups
- Parallels Cloud Storage provide scalability, redundancy, HA

> How to try



How to use it for FREE

Cloud Storage:

Parallels Cloud Storage: **FREE up to 100 GB**

<http://sp.parallels.com/products/pcs/cloud-storage/>

Virtualization:

OpenVZ: **FREE**

<http://openvz.org>

Parallels Cloud Server: License is required

<http://parallels.com/products/pcs/>

How to install Parallels Cloud Storage

Parallels Cloud Storage is available on the following distributives:

OpenVZ:

https://openvz.org/Parallels_Cloud_Storage

Parallels Cloud Server:

<http://sp.parallels.com/products/pcs/cloud-storage/>

Red Hat 7, CentOS 7: **Coming soon...**

http://download.pstorage.parallels.com/standalone/packages/rhel/7/x86_64/

Mail to: pstorage@parallels.com

Ubuntu 14.04: **Coming soon...**

<http://download.pstorage.parallels.com/standalone/packages/ubuntu/pool/main/p/pcs/>

Mail to: pstorage@parallels.com

> Q&A