Equnix Business Solutions, PT
Sampoerna Strategic Square, North Tower 25th Floor
Jl Jend Sudirman Kav. 45, Jakarta - Indonesia 12930
T: +6221 22866662 I info@equnix.asia
https://equnix.asia

Benchmark Report of PostgreSQL Tests on AMD EPYC Rome and Milan

Version : 1.1

Author : Equnix Research Team

No. : 20223342/RPT/II/EQUNIX

Date : 06 April 2022

Copyrights

Document Modification

Version	Date	Comments / Changes	Name
1.0	19-08-2021	Initial Version	Equnix Research Team
1.1	06-04-2022	Revise Benchmark comparison results on dual sockets	Equnix Research Team

Table Of Contents

Executive Summary

AMD EPYC has proven its ability in computing data intensive applications such as RDBMS. PostgreSQL as the most advanced and fastest Database Management System. It delivers tremendous results, compared to any dual socket servers in TPC-B and TPC-C Benchmark. TPC (Transaction Performance Council) is the world council dedicated as the ultimate reference for RDBMS benchmark (https://tpc.org)

Benchmarking PostgreSQL on top of AMD Zen 3 7543 has resulted in more than 3,021 TPS per Core for TPC-B, equivalent to 193,344 TPS for 64 core AMD EPYC processors. For TPC-C, it resulted in 45,952 TPM per Core, equivalent to 2,940,928 TPM for 64 core AMD EPYC processors.

On TPC-B, compared with AMD Zen 2 7742 performance, for a single processor, 7543 has 28% improvement. This improvement can be achieved as easily as replacing the socket to Zen 3 one. Meanwhile on TPC-C, 7543 has 48% improvement.

We do not recommend using dual sockets both on Zen 2 and Zen 3 in PostgreSQL, since it decreases the performance per core basis by around 64% for Zen 2 and 40% for Zen 3. In the other words, the performance using dual socket does not increase any significant improvement: 2% on Zen 2 and 18% on Zen 3.

AMD EPYC Zen 3 is the most recommended processor for High Performance Transaction System. This benchmark has already proven its ability the magic of Infinity Fabric, which is able to synchronize all cores to run in synergy. It answers almost all business needs for a High Performance Database System. For OLTP requirement, a great number of concurrency processes is a must feature to cater high loads of increasing user or customer accessing the system at the same time.

Methodology

Benchmarking a system has to be done by using proper mechanism or methodology. Improper usage of testing clearly results in nothing useful. The expected result should allow us to extrapolate in some manner to the real world condition and give us proper understanding or help us visualize the situation comparable to the production environment.

a. Construction

In performing the test, there are 2 configurations used on the AMD 7742 and AMD 7543 configuration as a Server. Below is the server configuration with 1 processor on each server.

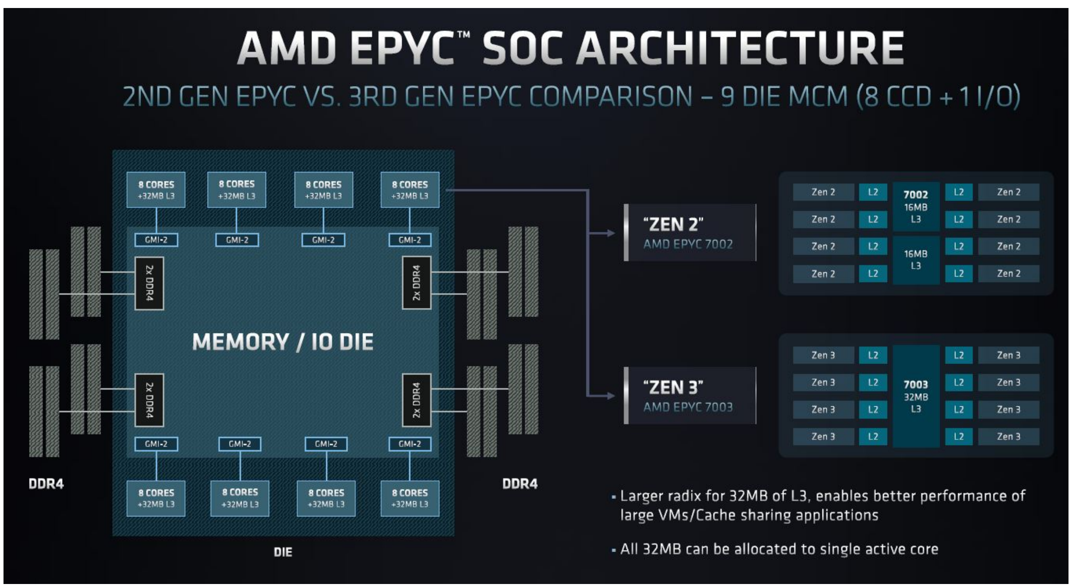

Hardware Specification:

No	Type	Specification	
1.	**AMD 7742**	Machine Type	GigaByte R272-Z30
2.		Processor	AMD EPYC 7742 64-Core Processor (128HT)
3.		Memory	512 GB RAM
4.		Disk 1	2 TB NVME

5.		Operating System	Debian GNU/Linux 10 (buster)
1.	**AMD 7543**	Machine Type	GigaByte R272-Z30
2.		Processor	AMD EPYC 7543 32-Core Processor (64HT)
3.		Memory	512 GB RAM
4.		Disk 1	2 TB NVME
5.		Operating System	Debian GNU/Linux 10 (buster)

Table 1: Hardware Specifications Single Socket

Below is the server configuration with 2 processor socket on each server.

No	Type	Specification	
1.	**AMD 7742**	Machine Type	GigaByte R282-Z90
2.		Processor	2x AMD EPYC 7742 64-Core Processor (256HT)
3.		Memory	512 GB RAM
4.		Disk 1	2 TB NVME
5.		Operating System	Debian GNU/Linux 10 (buster)
1.	**AMD 7543**	Machine Type	GigaByte R282-Z90
2.		Processor	2x AMD EPYC 7543 32-Core Processor (128HT)
3.		Memory	512 GB RAM
4.		Disk 1	2 TB NVME
5.		Operating System	Debian GNU/Linux 10 (buster)

Table 2: Hardware Specifications Dual Socket

Software involved in the test:

1. Debian GNU/Linux 10 (buster) (64-bit)
2. Database Server 11DB

3. pgBench for TPC-B tests.
4. HammerDB for TPC-C tests.

b. Software Installation

On Client Side, installs:

1. pgBench compile from source code
2. HammerDB, requires X Windows System.
3. some other utility tools such as: ntop, htop, iotop, etc.

On Server Side, installs:

1. 11DB from Source Code Equnix and compiled with GCC.

c. Benchmarking

There are 2 tools use for benchmark:

1. pgBench

   Use specifically for PostgreSQL benchmark with TPC-B scenario, TPC-B methodology is explained in the Terminology Section below. There are 3 parameters that determine the load of the test. This parameter should be set initially before every test, unless the next test is using the same parameter.

   1. **Scale** is the parameter for setting numbers of records on the database to be tested.
   2. **Endurance** is the number of second that test will passed. Higher number of the endurance affects the growing of CPU and memory usages, and at some point, it will consume all of the resources when the load is great enough. More usage of utilization of the CPU shows more effective the result comparing to the real world condition.
   3. **Concurrency**, determined number of simultaneous processes at the same time. It has linear relationship between concurrency and number of the Hardware Thread, by a factor of a constant number.

   pgBench initialization with parameters:

```
#pgBench -i -s 128 -h [host_server] -p [port_number] -U [username] [dbname]
```

-i is a parameter to perform initialization

-s is a parameter for determining the scale. 1 unit of scaling is 10,000 table record.

-h is a parameter to indicate the host address of the server.

-p is a parameter to indicate the port number of the server.

-U is the parameter to specify a user that will be used.

Running pgBench with parameter:

```
#pgBench -T [duration] -c [concurrency] -p [port_number] \
-h [host_address] -U [username] [dbname]
```

-T is a parameter to specify the duration in the second. In this case use the standard 600 second.

-c is a parameter to specify the value of concurrency to be used.

-p is a parameter to indicate the port number of the server.

-h is a parameter to indicate the host address of the server.

-U is a parameter to specify a user that will be used.

2. HammerDB

Use generally for benchmark with TPC-C scenario, TPC-C methodology is explained in the Terminology Section below. HammerDB was developed as Open Source and support many databases to be tested.

1. **Endurance** is the number of how long the benchmark will last, a similar parameter to the point 1. This number can be set by some number in seconds to determine capability of the database crunching high load at peak time.
2. **Concurrency**, the number of the simultaneous process of the database to cater high load which comes at the same time, that mimic the real world situation.

Parameter usage scope is:

- Number of Virtual User : 64 - 256
- Number of Warehouse : 64 - 256

- Duration of test : 30 minutes

Benchmark Result

1. TPC-B Benchmark comparison with 1 processor socket

TPC-B Result AMD Zen 3 7543 (64)

No	Scale	Concurrency	Thread	Duration (s)	fsync	Type	QPS
1	64	64	64	600	on	RW	91,132
2	64	64	64	600	on	R/O	**889,504**
3	64	64	64	600	off	RW	88,626
4	64	64	64	600	off	R/O	870,751
5	128	128	128	600	on	RW	**96,681**
6	128	128	128	600	on	R/O	854,046
7	128	128	128	600	off	RW	87,655
8	128	128	128	600	off	R/O	847,835

TPC-B Result AMD Zen 2 7742 (128)

No	Scale	Concurrency	Thread	Duration (s)	fsync	Type	QPS
1	64	64	64	600	on	RW	93,158
2	64	64	64	600	on	R/O	1,133,981
3	64	64	64	600	off	RW	29,038
4	64	64	64	600	off	R/O	1,134,455
5	128	128	128	600	on	RW	103,756
6	128	128	128	600	on	R/O	1,572,873
7	128	128	128	600	off	RW	95,771
8	128	128	128	600	off	R/O	1,573,100

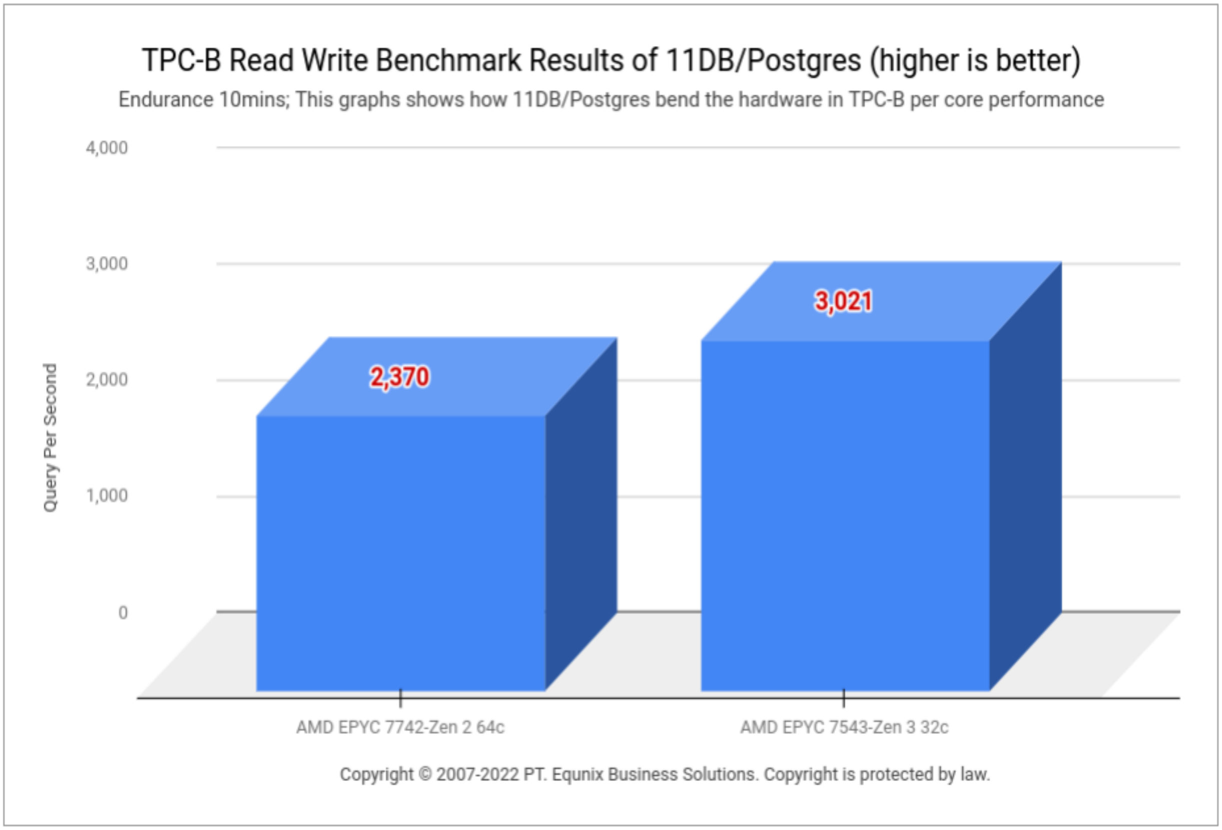

Illustration 1: Comparison TPC-B AMD Zen 2 and Zen 3 Per Core (Read Write)

Illustration 1 shows that PostgreSQL which is run on AMD Zen 3 7543 configuration has an improvement compared to AMD Zen 2 7742. On AMD Zen 2 7742, the optimum value is **2.370** QPS Per Core (151684 QPS / 64 Core) with concurrency 128. While on AMD Zen 3 7543, the optimum value is **3,021** QPS Per Core (96680 QPS / 32 Core) with concurrency 128.

$$1 + \frac{3{,}021 - 2{,}370}{2{,}370} = 1,2x$$

It shows about a **28%** increase in performance.

2. TPC-C Benchmark comparison with 1 processor socket

TPC-C Result AMD Zen 3 7543 (64)

No	Warehouses	User	Duration (s)	NOPM	TPM	fsync
1	64	32	30	615,320	1,415,076	off
2	64	64	30	604,972	1,426,019	off
3	64	32	30	608,224	1,403,205	on
4	64	64	30	602,867	1,394,219	on
5	128	64	30	631,345	**1,452,364**	off
6	128	128	30	628,794	1,448,582	off
7	128	64	30	639,469	**1,470,477**	on
8	128	128	30	596,069	1,373,997	on

TPC-C Result AMD Zen 2 7742 (128)

No	Warehouses	User	Duration (s)	NOPM	TPM	fsync
1	128	128	30	817,616	1,879,224	on
2	128	128	30	832,505	1,908,743	on
3	128	128	30	863,906	**1,979,870**	on
4	128	128	30	760,141	1,746,491	on

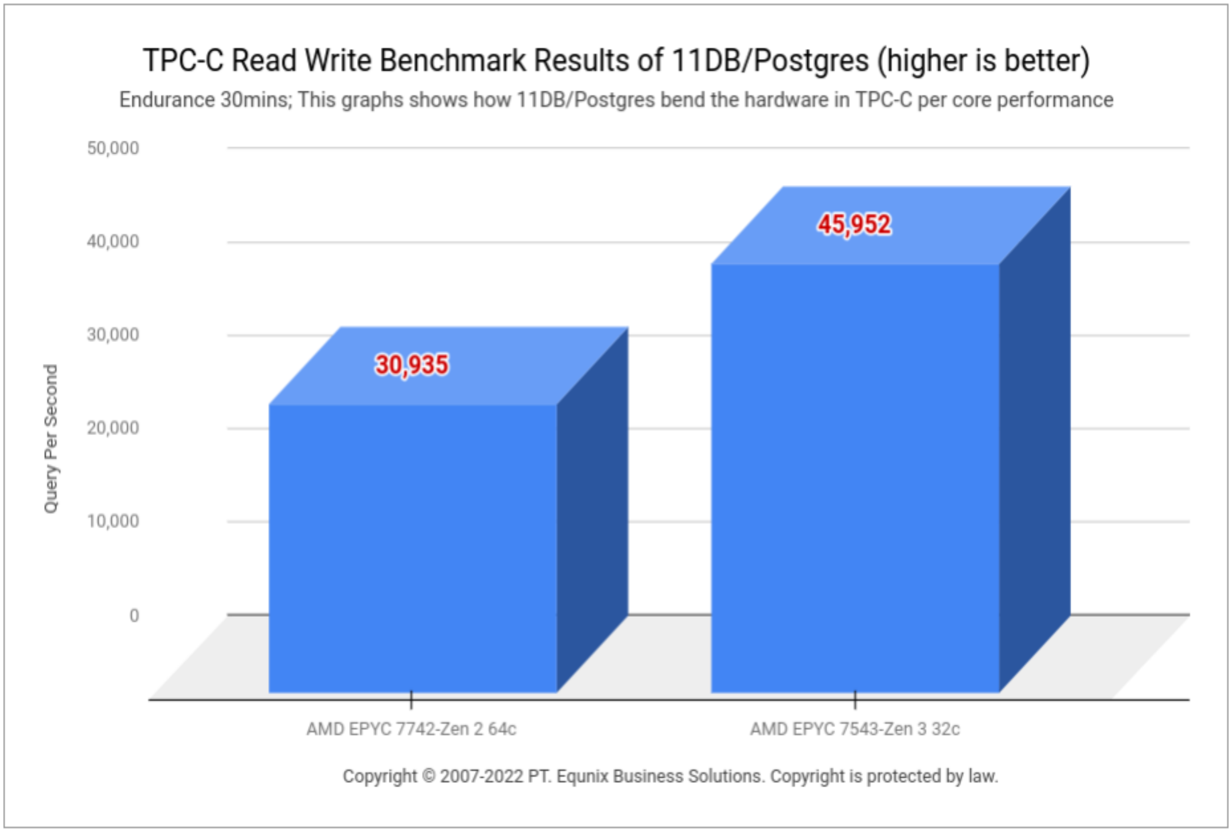

Illustration 2: Comparison TPC-C AMD Zen 3 32 Core and Zen 3 64 Core (Read Write)

Illustration 2 shows that PostgreSQL which is run on AMD Zen 3 7543 configuration has an improvement compared to AMD Zen 2 7742. On AMD Zen 2 7742, the optimum value is **30,935** QPS Per Core (1,979,870 QPS / 64 Core) with concurrency 128. While on AMD Zen 3 7543, the optimum value is **45,952** QPS Per Core (1,470,477 QPS / 32 Core) with concurrency 128.

$$1 + \frac{45{,}952 - 30{,}935}{30{,}935} = 1,48x$$

It shows about a **48%** increase in performance.

3. TPC-B Benchmark comparison with 2 processor sockets

TPC-B Result AMD Zen 3 7543 (128)

No	Scale	Concurrency	Thread	Duration (s)	fsync	Type	QPS
1	128	128	128	600	on	RW	102,958
2	128	128	128	600	on	R/O	1,396,416
3	128	128	128	600	off	RW	96,774
4	128	128	128	600	off	R/O	1,331,819
5	256	256	256	600	on	RW	**114,607**
6	256	256	256	600	on	R/O	1,376,968
7	256	256	256	600	off	RW	91,589
8	256	256	256	600	off	R/O	1,465,377

TPC-B Result AMD Zen 2 7742 (256)

No	Scale	Concurrency	Thread	Duration (s)	fsync	Type	QPS
1	128	128	128	600	on	RW	86,639
2	128	128	128	600	on	R/O	1,153,687
3	128	128	128	600	off	RW	107,938
4	128	128	128	600	off	R/O	1,145,422
5	256	256	256	600	on	RW	**105,752**
6	256	256	256	600	on	R/O	1,601,135
7	256	256	256	600	off	RW	124,977
8	256	256	256	600	off	R/O	1,576,313

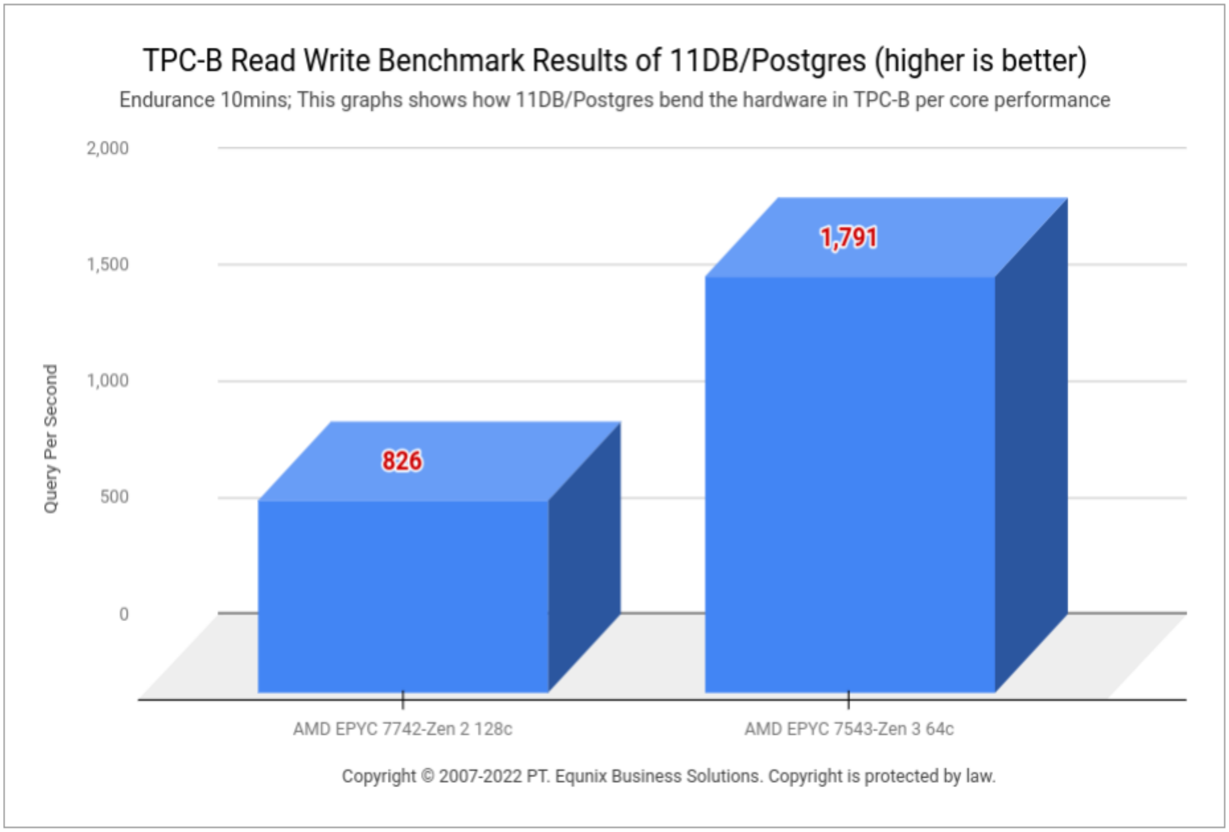

Illustration 3: Comparison TPC-C AMD Zen 3 32 Core and Zen 3 64 Core (Read Write)

Illustration 3 shows that PostgreSQL which is run on AMD Zen 3 7543 configuration has an improvement compared to AMD Zen 2 7742. On AMD Zen 2 7742, the optimum value is **826** QPS Per Core (105,752 QPS / 128 Core) with concurrency 256. While on AMD Zen 3 7543, the optimum value is **1791** QPS Per Core (114,607 QPS / 64 Core) with concurrency 256.

$$1 + \frac{1791-826}{826} = 2,16x$$

It shows about **116%** increase in performance.

Here is result comparison TPC-B AMD 7543 Zen 3 32 Core and 64 Core:

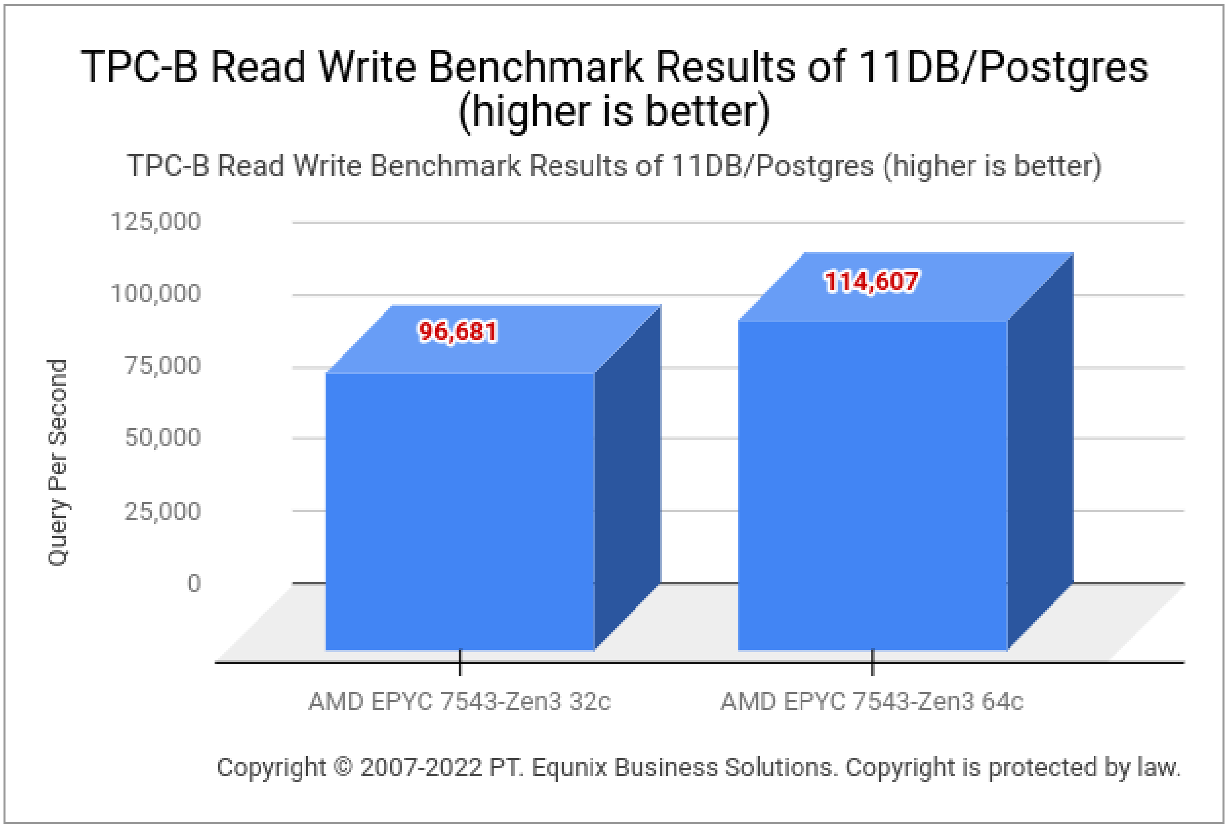

Illustration 4: Comparison TPC-B AMD Zen 3 32 Core and Zen 3 64 Core (Read Write)

Illustration 4 shows that PostgreSQL which is run on AMD Zen 3 7543 32 Core configuration has an improvement compared to AMD Zen 3 7543 64 Core. On AMD Zen 3 7543 64 Core, the optimum value is **114,607** QPS with concurrency 128. While on AMD Zen 3 7543 32 Core, the optimum value is **96,681** QPS with concurrency 128.

$$1 + \frac{114{,}607 - 96{,}681}{96{,}681} = 1,18x$$

It shows about an **18%** increase in performance.

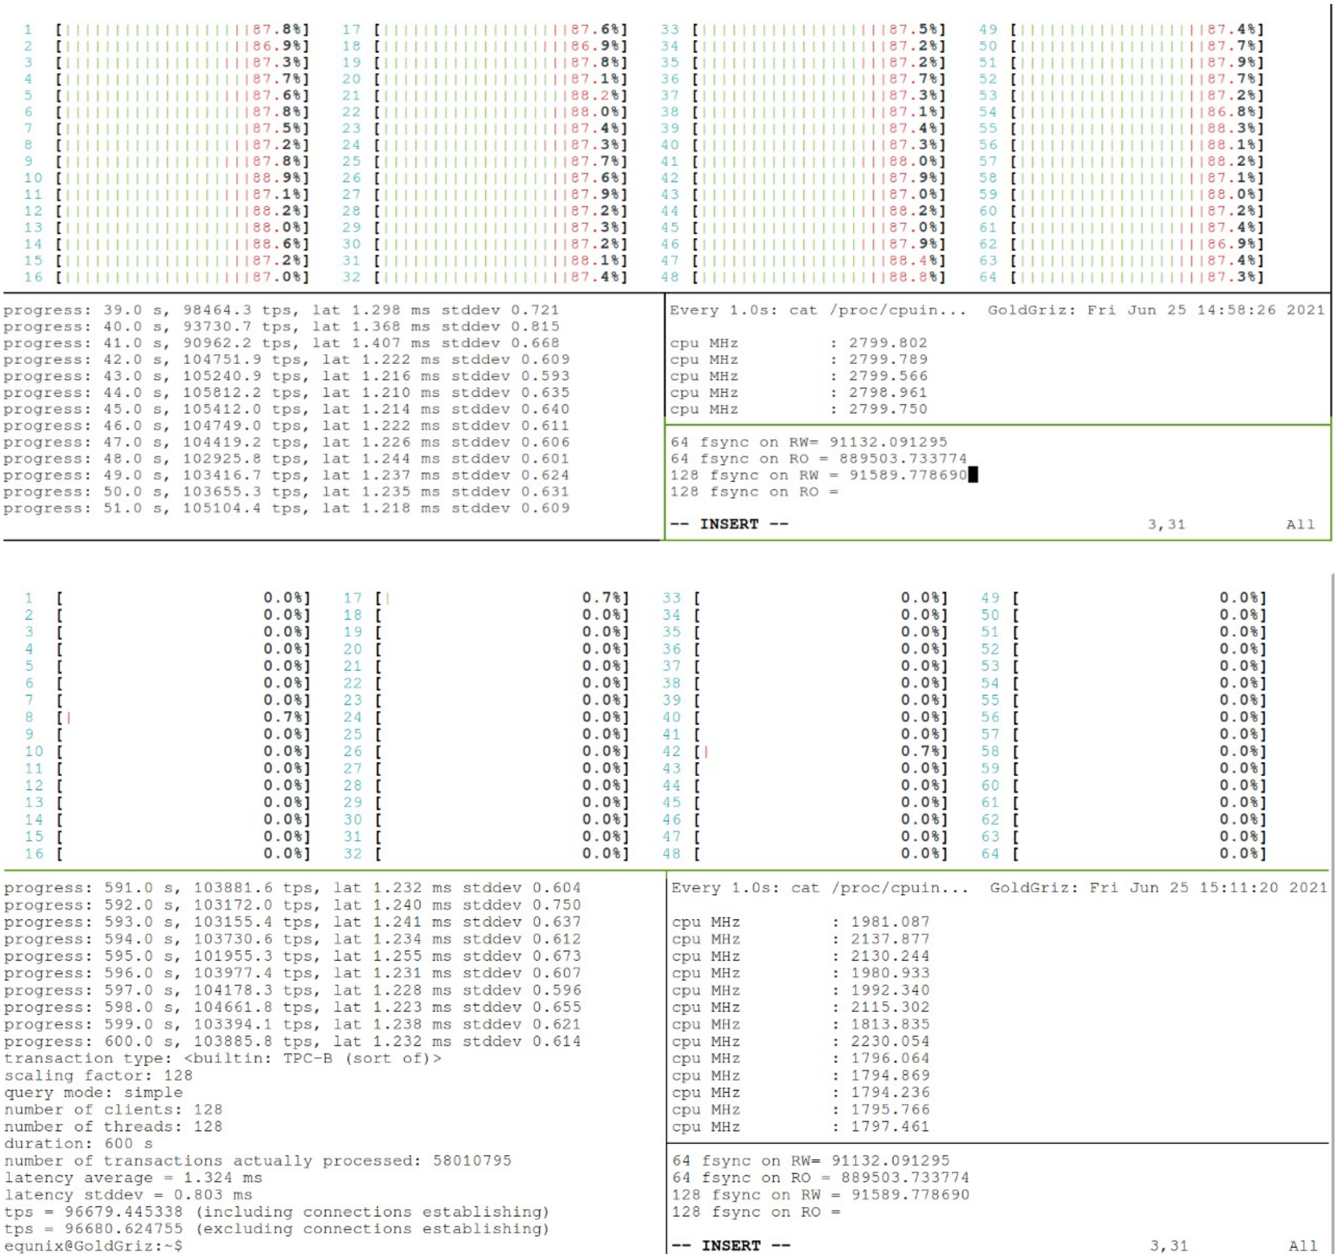

Illustration 2: CPU Usage of AMD Zen 3 75432 on TPC-B Bench

On AMD Zen 3 7543 configuration, when PostgreSQL is benchmarked, all CPU have more than 87% load. This describes that all CPU is used maximally, so PostgreSQL can achieve a huge performance.

Terminology

This chapter explains the definition and explanation of the term used in this document.

A. Software Terminology

i. RDBMS

RDBMS (Relational Database Management System) is an application for managing data in a structured and relational manner. The white-paper of Relational Data was invented by IBM Researcher Edgar F. Codd in 1970 as the foundation architecture of RDBMS today.

As RDBMS, there are 4 basic rules that should be complied with, Atomicity, Confidentiality, Isolation and Durable. These rules ensure the database can be fully functional on the transaction system. RDBMS has a Database Container which can contain: (1) Schemas, (2) Tables, (3) Triggers, (4) Stored Procedures, (5) Sequences, (6) Records, etc.

ii. PostgreSQL

PostgreSQL was derived from Ingres and Informix which was developed by UCB at 1970's up to 1990's led by Prof. Michael Stonebraker. In 1996, PostgreSQL Global Development Group has founded and became the major Developer of PostgreSQL until now. It is formed from many PostgreSQL Hacker all around the world. PostgreSQL World Conference is held every year in Ottawa, Canada in May. Equnix Business Solutions participated in the Conference as the Gold Sponsor 2014 as the contribution for the community of the Developers.

PostgreSQL has compiled 160 of 179 Conformance of SQL: 2011. The most conforming SQL database ever, with the new features of JSONB support in 2014, now PostgreSQL has the most features and joining the best of the both worlds, NOSQL and SQL.

Support for NoSQL or unstructured data meanwhile supporting Transactional and Structured data is the most awaited feature so far for many application which want flexibility as good as performance of financial related transactional.