CloudLinux OS

DIAWAY
Any Data Solutions

CPU PERFORMANCE IN SHARED HOSTING ENVIRONMENT

Analysis of AMD EPYC™ and 2nd Generation Intel® Xeon® Scalable Processors-based servers

- *Which PHP handler/web server is better to use?*
- *What is the comparative performance of mod_lsapi vs PHP-FPM?*
- *Which processor is a better choice today?*

TABLE OF CONTENTS

TESTING ENVIRONMENT & METHODOLOGY

1 Servers' hardware specs

HIGHER DENSITY GROUP- more cores with lower base frequencies (48 total cores per server at 2.2 or 2.3 GHz base frequency). Comparing 2 x 24-core Intel Xeon® vs 2 x 24 and 1 x 48 AMD EPYC servers.

HIGHER FREQUENCY GROUP - fewer cores at higher base frequency (32 total cores per server at 2.9 or 3.0 GH based frequency). Comparing 2 x 16-core Intel Xeon® vs 2 x 16 and 1 x 48 AMD EPYC™ servers.

Select configurations reflect today's gradual CSP transition to higher CPU core counts to provide higher multi-threaded performance, or faster response rates, at moderate price points. We've also included a new category of fully-featured single-socket AMD servers to see if they can provide competitive levels of performance vs traditional dual-socket counterparts.

This comparison included two configurations, based on recently-announced third-generation AMD EPYC™ processors, codenamed "Milan". Upcoming next-generation Xeon Scalable processors, codenamed "Ice-Lake SP", were unavailable at the time of testing.

All servers were equipped with identical DIMM and drive subsystems. We chose not to use the latest DDR3200 memory, supported by AMD EPYC™ processors, to fully equalize compared configurations, at the expense of losing certain points of performance on AMD-based platforms.

TESTING ENVIRONMENT & METHODOLOGY

1 Servers' hardware specs

Common efforts made by DIAWAY and CloudLinux were to identify the potential of AMD EPYC-based systems, running typical Web Hosting workloads, and objectively compare them with 2nd Generation Intel® Xeon® Scalable processors, that still hold the larger share in data centers. 32 and 48-core servers in DIAWAY's PoC lab in Estonia were configured for the remote tests by CloudLinux.

Specific Intel Xeon-based server configurations were selected to be a counterpart test subject, to provide a baseline performance/TCO reference in each core count category.

- For AMD single-socket configuration there was a DIAWAY Viimsi 1U Cloud Server for 1x AMD CPU, 12 NVMe U.2;
- For AMD dual-socket there was a DIAWAY Allika 1U Cloud Server for 2x AMD CPU, 12 NVMe U.2;
- For Intel configurations, it was DIAWAY Kalev All-Flash Dual Socket Storage Server.

> Typically, some applications require higher frequency, while other applications depend on higher core count. We tested 32 cores for high-frequency cases and 48 cores for high density as the foundation we used to define the two groups of processors. This sizing was selected for CSPs as it is considered optimal from the performance perspective (high frequency typically follows less core count), as well as financial (higher frequency and higher core density typically have a significant price penalty).
>
> The resulting figures and configurations are meant to reflect CSP's real-life workloads, transparently and objectively.
>
> DIAWAY Any Data Solutions

TESTING ENVIRONMENT & METHODOLOGY

DIAWAY Any Data Solutions | CloudLinux OS

HIGHER DENSITY GROUP

48 cores per server

1 (intel)

- **2 x Intel® Xeon® Gold 5220R**
- CPU count: 2
- Core count: 2x24
- Base frequency: 2.2 GHz
- TDP: 150W x 2
- Price: $1,505 x 2 (Intel RCP)
- RAM: 512GB Reg DDR2933
- (16 x 32GB 2933Mhz)

2 (AMD)

- **1 x AMD EPYC™ Rome 7552**
- CPU count: 1 (Single socket)
- Core count: 1x48
- Base frequency: 2.2 GHz
- TDP: 200W
- Price: $4,025 x 1 (AMD 1ku)
- RAM: 512GB x Reg DDR2933 (16 x 32GB 2933Mhz)

3 (AMD)

- **2 x AMD EPYC™ Rome 7352**
- CPU count: 2
- Core count: 2x24
- CPU base frequency: 2.3 GHz
- TDP: 155W x 2
- Price: $1,350 x 2 (AMD 1ku)
- RAM: 512GB Reg DDR2933 (16 x 32GB 2933Mhz)

4 (AMD)

- **1 x AMD EPYC™ Milan 7643**
- CPU count: 1 (single socket)
- Core count: 1x48
- CPU base frequency: 2.3 GHz
- TDP: 225W
- Price: $4,995 x 1 (AMD 1ku)
- RAM: 512GB x Reg DDR2933
- (16 x 32GB 2933Mhz)

HIGHER FREQUENCY GROUP

32 cores per server

1 (intel)

- **2 x Intel® Xeon® Gold 6226R**
- CPU count: 2
- Core count: 2x16
- CPU base frequency: 2.9 GHz
- TDP:150W x 2
- Price:$1,300 x 2 (Intel RCP)
- RAM: 512GB, Reg DDR2933
- (16 x 32GB 2933Mhz)

2 (AMD)

- **1 x AMD EPYC™ Rome 7542**
- CPU count: 1 (Single socket)
- Core count: 1x32
- CPU base frequency: 2.9 GHz
- TDP:225W
- Price:$3,400 x 1 (AMD 1ku)
- RAM: 512GB Reg DDR2933
- (16 x 32GB 2933Mhz)

3 (AMD)

- **2 x AMD EPYC™ Rome 7302**
- CPU Count: 2
- Core count: 2x16
- CPU base frequency: 3.0 GHz
- TDP:155W x 2
- Price: $978 x 2 (AMD 1ku)
- RAM: 512GB, Reg DDR2933
- (16 x 32GB 2933Mhz)

4 (AMD)

- **2 x AMD EPYC™ Milan 7313**
- CPU count: 1 (single socket)
- Core count: 2x16
- CPU base frequency: 3.0 GHz
- TDP: 155W x 2
- Price: $1,083 x 2 (AMD 1ku)
- RAM: 512GB x Reg DDR2933
- (16 x 32GB 2933Mhz)

TESTING ENVIRONMENT & METHODOLOGY

2 Software

WEB SERVERS

- Apache 2.4.6 for CloudLinux 7, mpm Worker
- Apache 2.4.37 for CloudLinux 8, mpm Worker
- Litespeed Web Server 5.4
- NGINX 1.16.1 + PHP-FPM

PHP

PHP version tested: 7.3 with enabled opcache extension.

OS AND RDBMS

- CloudLinux OS release 7.9 with MySQL 5.7.33
- CloudLinux OS release 8.2 with MariaDB 10.3.25

APACHE PHP HANDLERS

- mod_lsapi - PHP handler provided by CloudLinux
- PHP-FPM
- mod_suphp

TESTING ENVIRONMENT & METHODOLOGY

2 Software

Summary

As a result, using the software listed above, we have established the following set of configurations for PHP 7.3.

CloudLinux OS 7.9 with MySQL 5.7.33

1. Apache 2.4.6* + mod_lsapi**
2. Apache 2.4.6* + PHP-FPM
3. Apache 2.4.6* + mod_suphp
4. Litespeed Web Server 5.4
5. NGINX 1.16.1 + PHP-FPM

CloudLinux OS 8.2 with MariaDB 10.3.25

1. Apache 2.4.37* + mod_lsapi**
2. Apache 2.4.37* + PHP-FPM
3. Apache 2.4.37* + mod_suphp
4. Litespeed Web Server 5.4
5. NGINX 1.16.1 + PHP-FPM

* Worker

** With connection pool feature enabled

METHODOLOGY OF TESTING

Preconditions

We generated 200 user accounts for each server

We created a website on WordPress CMS for each user

We created 12 unique pages for each website

As a result, we arrived at 200 domains and 2400 unique URLs which roughly corresponds to the average shared hosting server

Since we wanted to reach the best server performance possible, we decided to use unlimited LVE limits for each user, and CageFS was enabled.

While testing PHP-FPM, we used the following opcache settings:

- opcache.memory_consumption=8192
- opcache.max_accelerated_files=800000

While testing other configurations (non PHP-FPM), we used the following settings:

- opcache.memory_consumption=128
- opcache.max_accelerated_files=4000

We will explain our choice of settings further in the conclusion. In other cases, the default configurations were used whenever possible.

METHODOLOGY OF TESTING

Testing Process

We have chosen the following testing procedure to mimic typical shared hosting environment:

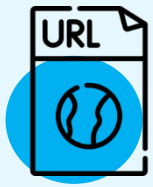

Requests were sent to 2400 pre-generated URLs randomly

One testing iteration lasted 5 minutes

Each iteration was with 150 concurrency level using non-keep-alive connections

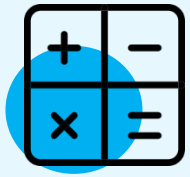

At the end of each iteration, the **total number of processed requests** and the **number of processed requests per second** were calculated

Each configuration listed in the table, went through these test four iterations. The final value of requests per second was calculated as the average over all iterations.

CPU PERFORMANCE RESULTS – HIGHER DENSITY GROUP

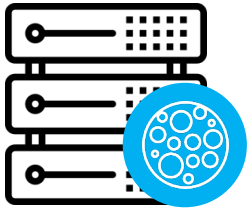

The Higher Density Group research has shown that the most powerful CPU is AMD EPYC™ Milan 1x7643 (48 cores, 2.3GHz): it can process 1258 queries per second. In terms of price/performance ratio (10M queries in our case) the best option is AMD EPYC™ 2x7352 (2x24 cores, 2.3GHz): 10M queries processing will cost $0.89.

At the same time, AMD EPYC™ Milan 1x7643 (48 cores, 2.3GHz), which is the most powerful one, is slightly more expensive, and it is third with $0.98 for 10M queries.

Aggregated Performance for Higher Density Group

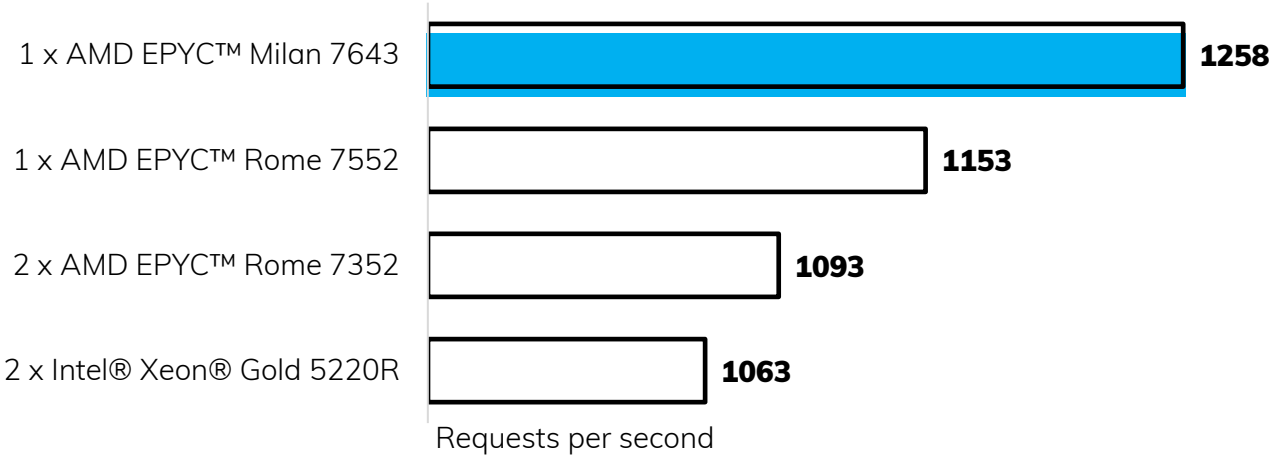

Price for 10M requests for Higher Density Group

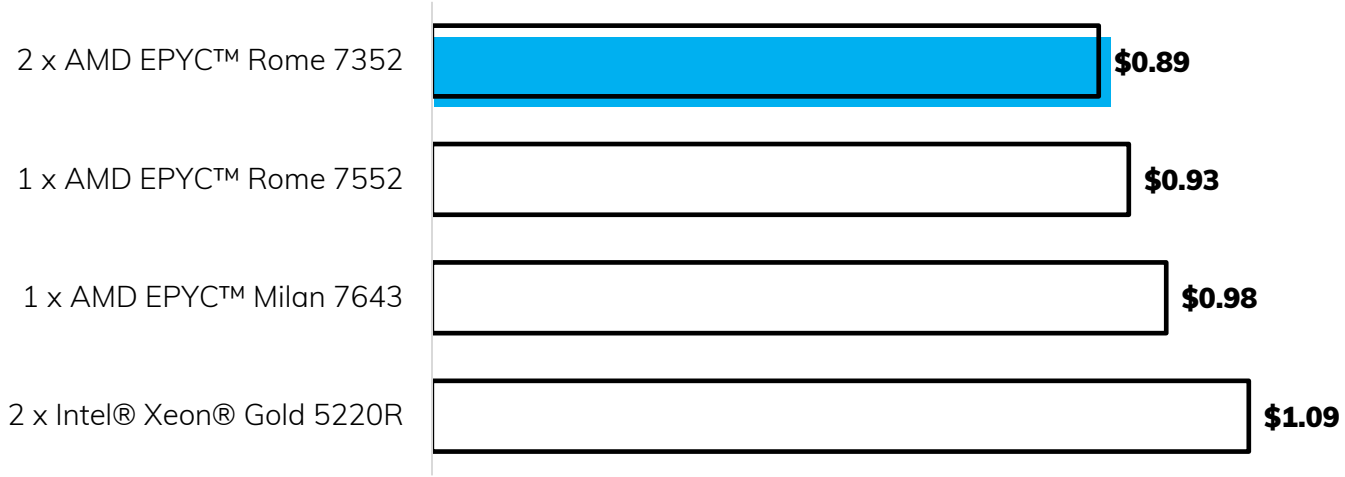

CPU PERFORMANCE RESULTS – HIGHER FREQUENCY GROUP

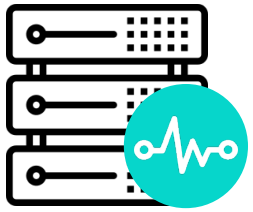

On the Higher Frequency Group, the best performance level is by AMD EPYC™ Milan 2x7313 (2x16 cores, 3GHz), it can process 1043 queries per second.

The same configuration was also the best one in terms of price/performance ratio: 10M queries processing will cost $0.96.

Aggregated Performance for Higher Frequency Group

Price for 10M requests for Higher Frequency Group

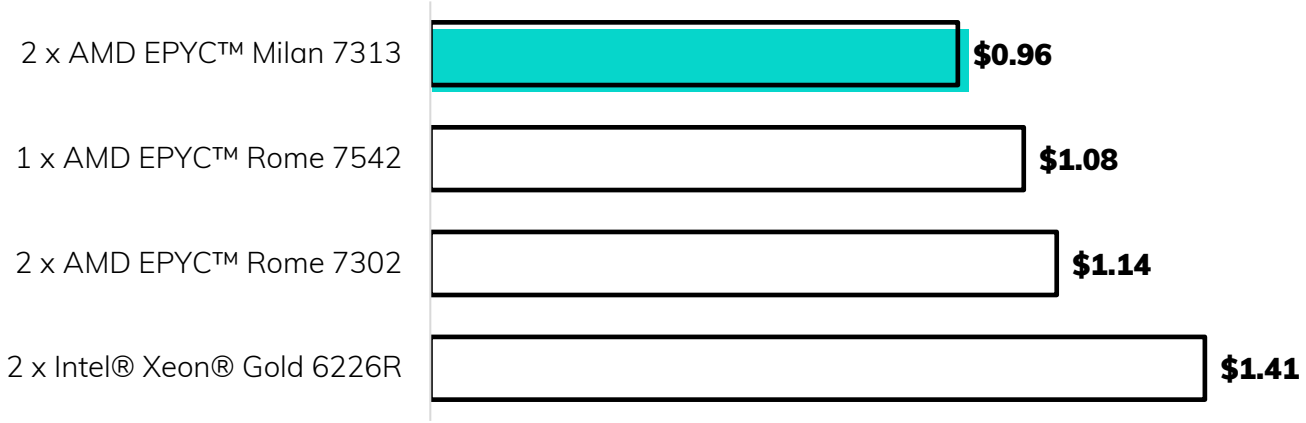

TOTAL COST OF OWNERSHIP FOR 4 YEARS

1. The AMD single-socket server platform proved to provide impressive performance and scalability. While it can replace dual-socket systems, it does come at a cost, unless you opt for the AMD EPYC 7xxxP single-socket CPUs, that are sold at lower price points vs their dual-socket counterparts with the same model number.
2. Both single- and dual-socket servers, based on the 2nd generation EPYC CPUs have demonstrated higher performance and performance per watt than the current 2nd Generation Intel® Xeon® Scalable servers that we have tested.
3. The newest 3rd generation EPYC Milan CPUs were the undisputed performance leaders in both 32- and 48-core server categories. However, their leadership have a somewhat higher power consumption vs the same core count 2nd generation EPYC models. It still translates into a leading performance-per-watt ratio in the test.
4. Dual-socket setups continue to play a vital role in CSP datacenters. Two lower core CPUs with faster clocks are beneficial to low-thread or latency-sensitive apps. Compared to single-socket servers, dual-socket systems can also provide higher peak memory bandwidth per core. However, not all applications can take advantage of that, so it is important to know the specifics of the application.

TCO 48m, monthly for Higher Density Group

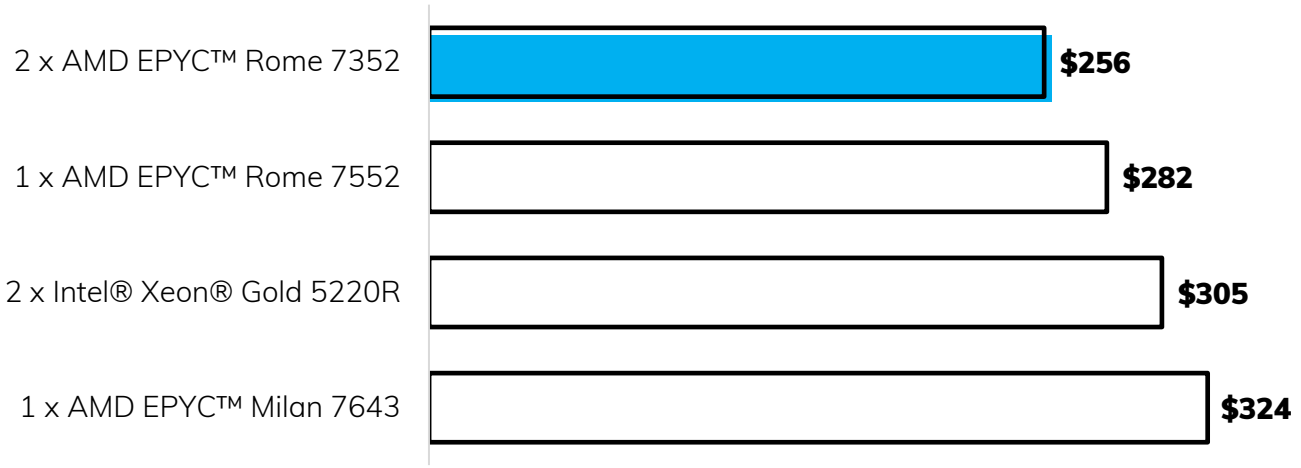

TCO 48m, monthly for Higher Frequency Group

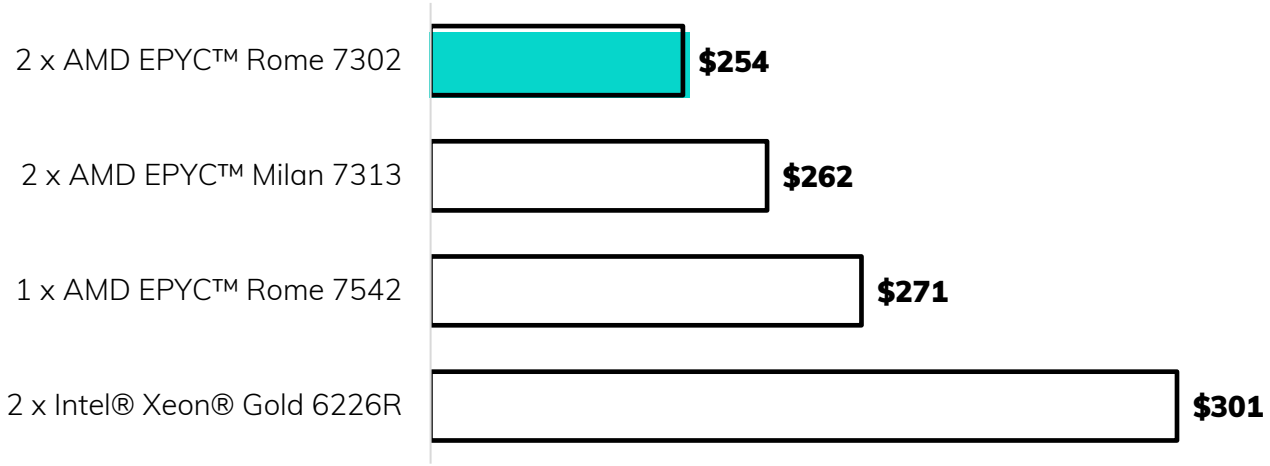

COST CALCULATION

Power consumption was measured during testing, and the results are tabulated and presented in Watts. Also added to the table, is the server cost and calculation of cost per kWatt. Monthly costs (TCO 48m, monthly) are calculated as follows: Server Cost / 48 month + 284 $ (Colocation cost per kWatt) * Power Consumption, Watts / 1000. Accordingly, TCO, 4 years is considered as TCO monthly * 48, values are also entered in the table.

The next step is to calculate the maximum performance in 4 years. Given the value of the number of processed requests per second, you can calculate how many requests will be processed by the server in 4 years. Aggregated Performance is the requests per seconds * 31536000 (number of seconds in a year) * 4 (years)*.

Further, it is easy to find out the price for one request Price per one request = TCO, 4-year / Aggregated Performance, requests for 4 years. For ease of comparison, in the last column of the table, we have provided the price for processing 10M requests.

*Real-life average server utilization normally stays below 100% so the calculation must be adjusted accordingly.

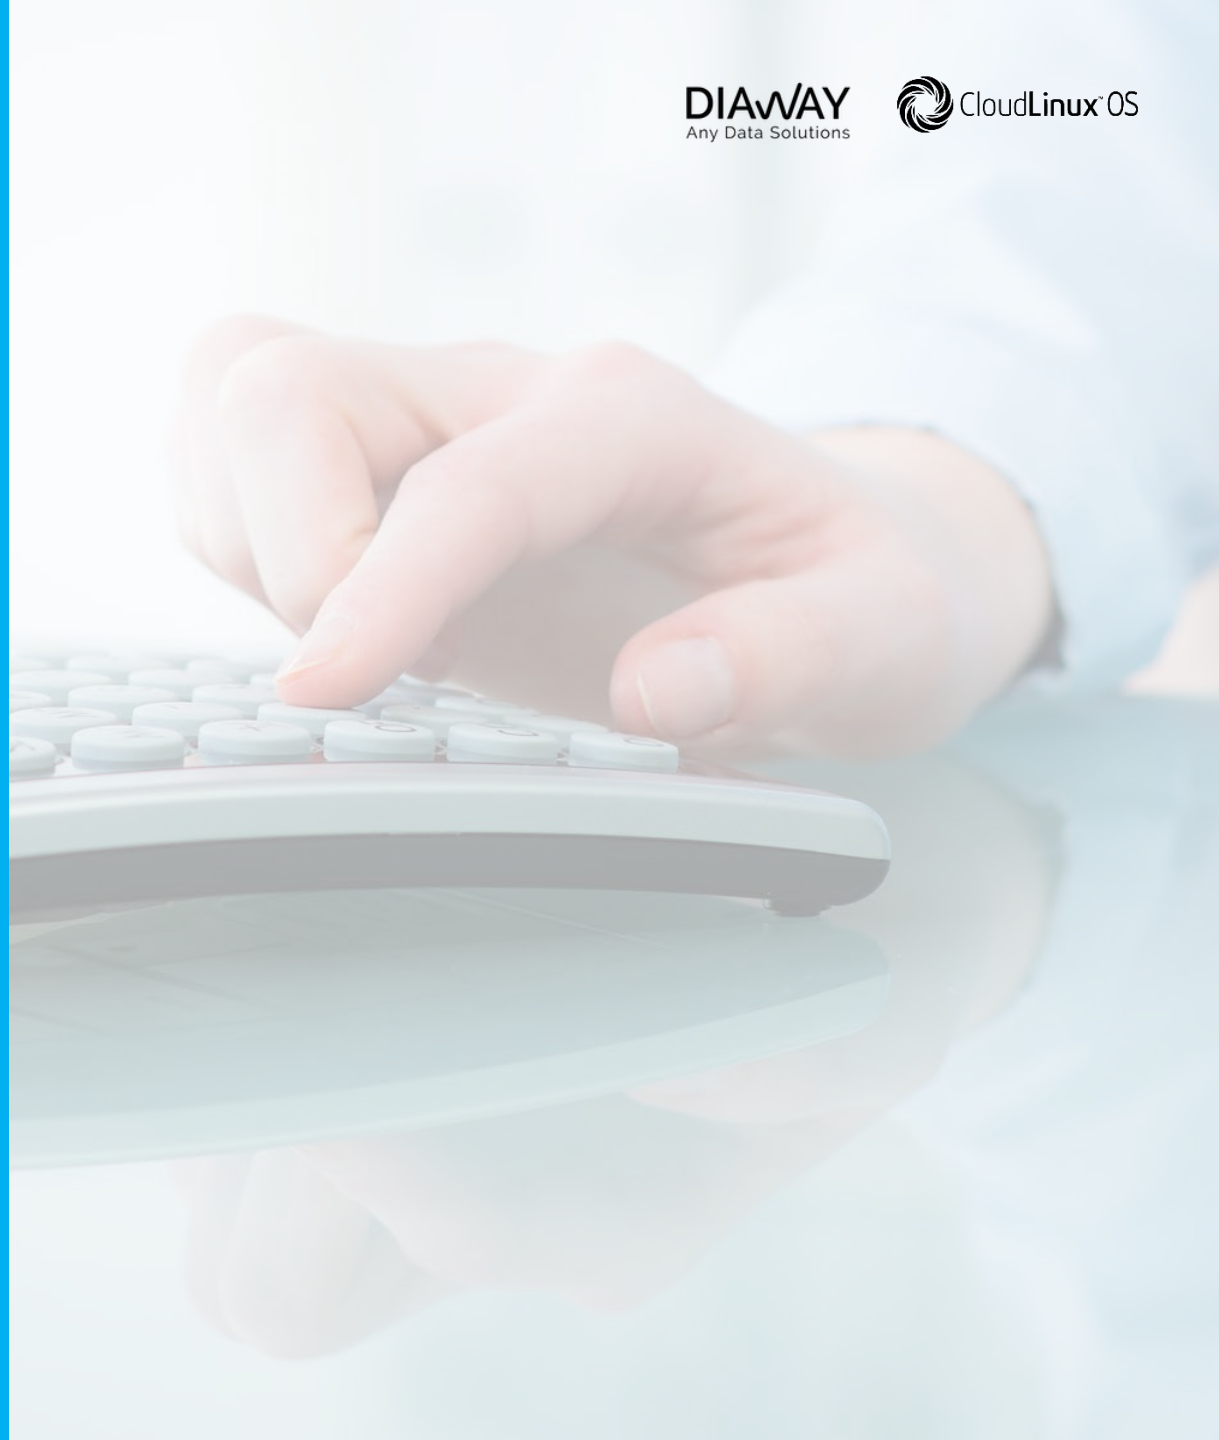

COST CALCULATION

Group	Configuration	Power Consumption, Watts	Server Cost	Colocation cost per kWatt	TCO 48m, monthly	TCO, 4 year	Aggregate Performance, requests per seconds (the greater the better)	Aggregate Performance, requests for 4 year	Price per one request	Price per 10M requests (the lower the better)
Higher frequency 32 cores	2 x Intel®Xeon®Gold 6226R	524	$7,304	$284	$301	$14,448	815	102,807,360,000.00	1.41E-07	$ 1.41
	2 x AMD EPYC™ Rome 7302	437	$6,255	$284	$254	$12,192	846	106,717,824,000.00	1.14E-07	$ 1.14
	1 x AMD EPYC™ Rome 7542	405	$7,499	$284	$271	$13,008	959	120,972,096,000.00	1.08E-07	$ 1.08
	2 x AMD EPYC™ Milan 7313	450	$6,465	$284	$262	$12,576	1043	131,568,192,000.00	9.56E-08	$ 0.96
Higher density 48 cores	2 x Intel®Xeon®Gold 5220R	502	$7,800	$284	$305	$14,640	1063	134,091,072,000.00	1.09E-07	$ 1.09
	2 x AMD EPYC™ Rome 7352	389	$6,999	$284	$256	$12,288	1093	137,875,392,000.00	8.91E-08	$ 0.89
	1 x AMD EPYC™ Rome 7552	396	$8,124	$284	$282	$13,536	1153	145,444,032,000.00	9.31E-08	$ 0.93
	1 x AMD EPYC™ Milan 7643	472	$9,094	$284	$324	$15,552	1258	158,689,152,000.00	9.80E-08	$ 0.98

WEB SERVERS AND PHP HANDLERS PERFORMANCE

mod_suphp

mod_suphp showed the worst results compared to other configurations. These are our expected results due to several disadvantages of its implementation, namely:

- There is no process manager - this means that for each HTTP request, a new PHP process is spawned using the exec function, which is a very cumbersome operation.
- There is no way to take advantage of opcache. After processing the PHP request, the process ends, therefore the memory that was allocated by opcache for further reuse is cleared.

Therefore, mod_suphp is currently deprecated, we do not recommend using it on shared hosting.

mod_lsapi

Despite a little lagging behind PHP-FPM in performance, mod_lsapi, due to the architectural limitations of PHP-FPM, is much better suited to shared hosting requirements:

- mod_lsapi provides memory savings in shared hosting environments.
  mod_lsapi provides independent PHP configuration for various accounts, making it easy to configure multi versions per domain.
- mod_lsapi provides a dedicated opcache pool for each account. This isolates the PHP processes of different accounts from each other, and eliminates the possibility of them competing for memory in the pool.
- The memory allocated for opcache pools of various accounts is counted in the memory they consume. Whereas in the case of PHP_FPM, the opcache pool is shared by different accounts.The memory allocated for it is not considered for any of the accounts, and the amount of memory consumed in the shared pool by different accounts cannot even be measured and, moreover, limited.

WEB SERVERS AND PHP HANDLERS PERFORMANCE

PHP-FPM

While testing PHP-FPM, we noticed that it uses a different opcache architecture. In this case, all domains share the same opcache buffer, while for the rest of the PHP handlers every domain has its own dedicated opcache.

Due to this, FPM handlers were tested with the different opcache settings:

- opcache.memory_consumption = 8192
- opcache.max_accelerated_files = 800000

If you have a large number of domains and you use opcache, you need to take these settings into account. For example, if you set the values between domains too small, the opcache memory race condition will occur and this will lead to a loss in performance. If the value is set too high, the memory will be allocated by the system for work, but will not be used. This is one of the downsides of PHP-FPM and the reason why it needs more memory to run than mod_lsapi.

NGINX + PHP-FPM

NGINX + PHP-FPM showed almost the same performance as Apache + PHP-FPM. We used WordPress sites with a small number of static files in testing, so there is no noticeable speed gain from using NGINX. In our case, performance depended more on PHP-FPM performance.

Litespeed Web Server

According to our results, Litespeed Web Server is the favorite of our entire list of configurations.

WEB SERVERS AND PHP HANDLERS PERFORMANCE RESULTS

If you are interested in the best free solution, we recommend using PHP-FPM + Apache or Nginx. However, you may face high-memory consumption and the complexity of having to configure the server yourself.

For a simple management solution with efficient memory consumption, we recommend mod_lsapi. If you are ready to pay for a web server, we recommend choosing Litespeed.

Regardless of what option you choose, you should keep in mind that you will need cache plugins, such as opcache, installed.