Comparative Performance Evaluation of Popular Virtual Private Servers

Josip Balen[1], Denis Vajak[1], Khaled Salah[2]

[1] Faculty of Electrical Engineering, Computer Science and Information Technology Osijek, J. J. Strossmayer University of Osijek, Croatia
[2] ECE Department, Khalifa University of Science, Technology and Research, UAE
josip.balen@ferit.hr, denis.vajak@ferit.hr, khaled.alah@kustar.ac.ae

Abstract

Virtual Private Server (VPS) enables user access to an operating system instance with a fine-grained control of private software and hardware resources. Many various factors can affect VPS performance and they primarily include physical server hardware specifications, installed operating system, virtualization layer, and the underlying network infrastructure. Therefore, it is very important to properly select a VPS host that meets users, applications and services resource and performance requirements. This paper presents a performance evaluation of three popular VPS hosts; namely Digital Ocean, Linode and VULTR. Performance measurement was conducted under the same controlled conditions for all three VPS hosts using a popular benchmark application for Unix operating systems - UnixBench. Three performance evaluation experiments with a focus on examining and studying key performance metrics which include CPU scheduling, memory management, hard disk drive management and Unix operating system task scheduling, were conducted. Performance measurement results show that VULTR achieves the best results under most of the tests in the first two experiments making it the best choice for low demand users, while DigitalOcean achieves the best results in the third experiment making it the best solution for high demand users who are looking for a high performance VPS.

Keywords: Benchmark, Performance evaluation, Virtual private server, UnixBench

1 Introduction

With a rapid explosion of information generated by users and needs for fast processing of huge amounts of data, personal computers have become inefficient in time and resources. Users are searching for other solutions and technologies that will allow fast and efficient processing of their tasks and data. Since the Internet is rapidly evolving into an always-on, always-connected, device-independent environment, a new trend inspired by cloud computing, can be observed, as described in [1-2] and [3]. In order to meet user demands for a high performance, new technologies and services emerged, such as grid, utility computing, virtualization or autonomic, mobile, and pervasive computing. The main advantage of these technologies is that users can use them according to their nees as they enable ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released with a minimal management effort or service provider interaction, as described in [4-6] and [7].

Virtual Private Servers (VPS) offer a cost-effective way of providing a powerful and flexible multipurpose server. From rendering videos, to hosting game servers and websites, VPSs have a wide array of possible usages and are used by many different users in order to meet the growing private and professional processing, computational, and storage requirements and needs. Although cloud computing is growing rapidly, server virtualization technology still holds a strong position on the market and VPS platforms has been a preferred choice for many users over cloud-hosted virtual machines due to several advantages it possesses. Although both platforms enable users to get powerful computing resources at low cost with a quick setup, the security and powerful admin control are the advantages of VPS hosting. As opposed to the cloud platforms, a VPS is completely owned by a user with clear isolation and boundaries for sharing the physical and logical resources, and with the ability of fine-grained control customization, and full administration of computational, storage, and network resources. Security level is also higher with a VPS since the underlying systems (software and hardware) as well as data have clearly defined physical and logical isolations [8] and [9]. Although cloud platform is scalable, scalability takes time, and is in many ways slower than the scalability of VPSs. Furthermore, Virtual Machine Mapping Problem (VMMP) could appear in the cloud infrastructure due to

*Corresponding Author: Josip Balen; E-mail: josip.balen@ferit.hr
DOI: 10.3966/160792642020032102003

a high number of reserved instances of virtual machines by cloud users. Therefore, some users' requests cannot be handled by using the available resources and they become rejected, like described in [10] and [11]. VPS hosts have the ability to quickly and dynamically create and terminate VPS instances, which makes them a reliable and suitable platform for on-demand scalable systems. It is also possible to establish a cost-efficient virtual MapReduce cluster by renting multiple VPSs from one provider [12].

Furthermore, one of the benefits of virtualization that often goes unnoticed is green computing. Compared to cloud computing, virtualization uses less equipment and resources, and consequently has a lower energy consumption, as explained in [13]. One of the most famous usages of VPSs is in the most successful and widely used Infrastructure-As-A-Service (IaaS) cloud platform - Amazon EC2. In its deployment it heavily relies on the Xen virtualization where each virtual machine, known as an instance, functions as a VPS, as described in [14].

Several superficial comparisons of different VPS hosts can be found online in [15-17] and [18]. However, a scientific research in the context of VPS host performance evaluation is still in its infancy. An exhaustive performance evaluation of the most widely used VPS hosts is necessary since it represents an important and common interest of multiple various user groups. In general, numerous private, professional and academic users need scientifically proven results and guidelines in order to properly select a VPS host that meets resource and performance requirements for their hosted applications or services. Therefore, the main motivation behind this paper is to obtain a systematic and comprehensive scientific research in the field of VPSs, to provide a high quality performance evaluation results for the most widely used VPS hosts, and consequently to bring new knowledge that has been missing in the earlier work. In this paper, an experimental research about VPS hosting features, configuration options, and pricing plans was conducted, as well as performance evaluation of the three most widely used VPS hosts: DigitalOcean [19], Linode [20] and VULTR [21]. DigitalOcean and Linode are selected since they are market leaders, while VULTR is newer to the VPS hosting scene but is often used for academic purposes. The main contributions of this paper can be summarized as follows:

- A review of related work and patents in the field of VPS performance measurement and evaluation is conducted,
- A review of the architectural design of VPSs, three different VPS hosts and a benchmark tool is presented,
- The performance of three popular VPS hosts in terms of key system metrics which include CPU scheduling, memory management, hard disk drive performance and operating system task performance, is experimentally measured, studied and evaluated,
- Important repeatable guidelines, scripts, and experimental are developed.

The remainder of this paper is organized as follows. Section 2 discusses the related work. Section 3 describes the architecture of VPSs and lists the VPS hosts used in performance evaluation. Section 4 describes the performance evaluation methodology and the benchmark application used in the experiments. Section 5 explains the experimental setup and the full evaluation procedure. Section 6 discusses performance measurement experimental results. Section 7 concludes our paper.

2 Related Work

Performance evaluation of different operating and computer systems is often complicated and complex to perform because numerous internal, constructional or external factors could have an impact on the system performance. To date, there exists a handful of research articles dealing with the performance of various systems. However, related scientific work in this specific filed of VPS host performance evaluation is limited. Therefore, we reviewed all available professional, scientific and patent references below.

2.1 Our Previous Work in The Field

In our previous work we have studied several different aspects of operating systems performance on desktop computers. In [22], we studied how different host operating systems influence virtual machine performance. In [23] we preformed performance evaluation in two different environments of three different versions of Windows operating systems similar as in [24] where we evaluated network performance. This paper, to some extent, continues our work described in [25] where we conducted performance evaluation of three cloud IaaS providers, namely Amazon EC2, ElasticHosts, and BlueLock. Three different benchmark applications, namely Simplex, STREAM and FIO with a basic set of performance measurements were used in three different VPS setups. The performance was evaluated in terms of execution time of CPU-bound processes, size of memory bandwidth, and speed of read write disk I/O. Performance measurement results showed that Bluelock has the faster execution speed for CPU-bound processes and the biggest memory bandwidth. However, Amazon EC2 outperforms ElasticHosts and Bluelock when it comes to disk read/write bandwidth. The main drawback of this paper is the inability to enable the same software and hardware configuration on all tested hosts due to the rigid configuration of the IaaS providers, which do not allow changing the amount of RAM or type and speed of CPU. Furthermore, different versions of Linux operating systems were used and in our other work in the field we showed the influence of the operating system on the performance of the whole computer system. Therefore, in this paper, we significantly expanded our previous work by conducting a much more detailed research, performing a comprehensive performance measurement and evaluation with developed model and methodology and by using three new and lately extremely popular VPSs based on the similar hosting plan and the same operating system.

2.2 Professional VPSs Performance Comparisons

Several qualitative and comparative studies of different VPS hosts performance can be found online mostly conducted by IT professionals. However, they are not based on a scientific approach with a systematic research and repeatable performance measurement methodology, so the quality of performance measurement results and evaluation is questionable. In [26], it is possible to compare the performance and pricing of different VPS hosts and plans using several different metrics. In terms of response time of DigitalOcean, Linode and VULTR, it can be concluded that VULTR has the fastest response and Linode the slowest. CPU utilization is lower for DigitalOcean, and Linode uses most of the CPU resources. Another comparison of all three VPSs was conducted in [27]. It is concluded that DigitalOcean has the highest performance to price ratio, the most geographic choices, and supports custom operating systems plus it provides redundancy on the disk drives. VULTR is also recommended but is lacking data integrity. Performance evaluation and comparison of four VPS hosts, namely Rackspace, DigitalOcean, Linode and VULTR, is presented in [28]. After extensive performance measurement with a set of different benchmark applications it is concluded that Linode offers the best performance for author's specific application. In [29] five different VPS providers, namely OVH, Linode, DigitalOcean, Scaleway and VULTR, are compared in detail. From the point of the features availability it can be concluded that VULTR provides the best ratio when compared to the pricing. When it comes to the performance evaluation the UnixBench tool was also used and in general it can be concluded that VULTR is again the winner since it provides the best overall performance when compared to the others.

Compared to the approach used in this paper, all online approaches are missing a scientific approach to this problem, that includes systematic research, detailed description and exhaustive experiments that will ensure performance measurement results accuracy and reliability and provide fair, complete and thorough performance evaluation.

2.3 Scientific Literature

In the scientific literature, there are many research papers dealing with the performance of various cloud infrastructures. However, they are not subject of this research. Furthermore, very few research papers study VPSs and present performance evaluation, and at the same time most of them are outdated and not dealing with the currently available and the most popular VPSs. In [30] a scalable, distributed database system that allows uniform access to concurrently distributed databases over VPSs by the Single Query Multiple Database (SQMD) mechanism was developed. Performance evaluation was performed in order to demonstrate the viability of developed system and several main problems occurred. The first one was a degradation of the performance of the system with an increasing number of responses from distributed database servers due to the global aggregation operation in the web service. Second one was found in extra hits that their approach generates. The third problem was found in unnecessary query processing. The performance evaluation approach in [30] is quite different when compared to the approach used in this paper since it evaluates the performance of the system developed and implemented on VSP and not the comparative performance of popular VSPs. A research on VPS architecture and performance was conducted in [31] where multiple different virtual machine platforms available at that time were tested to determine which performed the best for VPS hosting. Parallels, VMware, and Sun Virtualbox were tested by measuring the number of requests per second that a server was able to handle. The results show that Parallels has the best results in all experiments, both in unshared and shared situations making it a good choice for hosting a VPS at home or in a professional environment. Although their approach is in general very similar to ours, there are however, two huge differences. The first one is that in [31] authors have setup their own virtual environments which means that they are not using any of the most popular online VPSs. The second one is that no benchmark application was used for performance measurements and the authors measured their own metrics in the system which is prone to errors and can obtain uncertain measurement results. An interesting work was conducted in [32] and [33]. Three different widespread virtualization tools with different virtualization techniques, Xen, VMware, and OpenVZ, are presented and evaluated considering their use for router virtualization. Although it is concluded that Xen is the one that best fits virtual router requirements, the results show that OpenVZ which is the only virtualization tool that is based on the VPS, introduces less overhead over CPU, disk, and memory usage when compared to the other virtualization techniques. In their work authors used several benchmark applications to measure similar metrics to our (CPU scheduling, memory management and hard disk drive management). However, they also measure network performance that is not included in our work and on the other hand we measured Unix operating system task scheduling that is not included in [32] and [33] that we think is more important in VPS performance evaluation. Furthermore, similar to the paper above, authors have setup their own virtual environments which means that they are not using any of the most popular online VPSs. In [34] a performance comparison of two different server's architectures, an Apache (process-driven) and an Nginx (event-driven) architecture configured on two VPSs for hosting a website is presented. Two experiments were conducted where the response time, memory usage and efficiency were measured and compared. The results show that the Nginx provides better performance in terms of responsiveness and scalability, while Apache ensures the efficiency. However, similar to the papers above, this work was conducted on a personal computer and not on online available VPS hosts, and therefore it does not provide useful information for the VPS selection. In [35] VMs and containers were combined in order to enhance containers' isolation and extend VMs' functionality. The performance overhead that was introduced by running containers on various virtual machines was experimentally quantified. Compared with our research, at the container level there is no difference between a VPS and a VM assuming that they use the same virtualization technologies. Furthermore, similar performance metrics to our research was measured (CPU, memory and disk performance) but for each a different benchmark application was used. Compared to only UnixBench used in this paper there are no advantages and benefits in using more benchmark application. Furthermore, network performance was also measured but due to numerous factors that can influence it we decided not to include it in our research. The interesting fact is that authors also measured power consumption and this metric can be very useful to the service provider but since we did not have a possibility to measure it, it is not included in our research. Authors concluded that the security and isolation of VMs with containers was improved but the main drawback is the performance and energy overhead of the additional virtualization layer. In VPSs there is no need for the additional virtualization level and therefore, the performances are higher.

In [36] an automatic performance verification technology, which evaluates three types of server architectures (bare metals, containers and virtual machines) is proposed. It executes necessary performance tests automatically on provisioned user environments and based on the performance evaluation it recommends server architecture. Compared to our work, UnixBench was also used as the only benchmark tool but performance evaluation was conducted only based on the index score comparison. Furthermore, proposed solution also reduces users' efforts on selecting, designing and verifying servers' architectures that satisfy users' performance requirements. Moreover, research from the paper [36] resulted with a patent described in [37].

2.4 Patents in The Field

Apart from professional and scientific literature, several patents that are dealing with systems and methods for performance evaluation of various servers, clouds and virtualization systems are available in patents databases. In [37] a complex system consisting of automatic server selection device, method and software is described. Server selection device receives input with server configuration, performance requirements and operating system data. Based on that information it selects the best type of server that is able to satisfy users' performance requirements and to utilize the server resources. In [38] a public cloud evaluation system for evaluating various resources offered by multiple public cloud providers is proposed. It consists of a test application and performance evaluation software that is capable of instantiating and configuring one or more VMs in the public cloud, distributing and executing the performance evaluation software among the VMs, and saving and analyzing the performance evaluation results. In [39] a method for automatic optimization of available virtual machine resources by sizing of virtual machines is presented. The method first profiles applications in order to obtain the resource demand estimates. Based on those results it selects the best virtual machine configuration that can provide the most efficient results. System that enables selecting an appropriate computing equipment configuration for a virtualized computing environment is proposed in [40]. It is able to conduct performance evaluation and generate an overall performance efficiency for each virtualization computing equipment configuration based upon a combination of weighted measures of efficiency. Based on those results it compares, scores and ranks computing equipment configurations in order to determine which configuration is best suited for a particular set of virtualized computing environment requirements. A method for resource monitoring of a compute infrastructure interconnected with a network is proposed in [41]. The method comprises the steps of allotting a benchmark message and initiating and sending of the benchmark message to a reflecting entity. An advantage of the proposed technology is that it allows active and on-demand evaluation of e.g. a virtual machine/container performance. It provides an effective and efficient method for performing generalized tests of various performance and capacity issues related to virtual machines and containers, as well as other virtual resources, like storage from a centralized location.

All mentioned patents above are based on automation of performance evaluation process based on the users' and applications' needs and selection of the best and the most efficient system resources that can provide the best performance. Those types of an automatic performance evaluations could be very useful since it saves time and effort required for finding the best system and resources for specific needs. However, currently performance measurement, testing and evaluation need to be done manually for every separate system and the performance evaluation model presented in this paper serves as a guideline that can be repeatable and applied to the similar online and cloud systems.

3 Virtual Private Server

A VPS is a virtual server that users perceive as a dedicated/private server even though it is installed on a physical computer machine running multiple operating systems. A VPS can have a variety of features including a web server, a File Transfer Protocol (FTP) program, a mail server, and different software for blogging and e-commerce. Since VPSs can have their own copies of operating systems, users are generally given super-user permissions allowing them to fully manipulate the system and install/remove any software.

Due to the evolution of virtualization technology and software, numerous companies offer cheap and affordable VPS solutions for customers. However, since multiple users are using a single physical machine, there is a high risk of load balancing issues. Just a single user overloading a hardware instance can slow down other users' VPSs by taking up excess resources, as described in [42]. To counter this, VPS hosts often ensure fair allocations or limitations to resources and bandwidth.

3.1 VPS Architecture and Virtualization

VPSs, as described in [31] and [43], are built on a virtualization technology, where separate instances of virtual machines run on a single physical server but are virtually isolated from each other. This allows them to look and feel like a real server from a users' point of view. Isolation dramatically reduces security concerns, since there is no sharing among users. In addition, each VPS can be rebooted and has its own virtual objects like root, users, passwords, IP addresses, applications, folders and files.

The VPS architecture consists of the base virtual infrastructure, which is not visible for the users, that runs virtual machines and is able to dynamically create and destroy virtual instances. Each virtual machine runs a separate operating system with applications above them. Due to this design, multiple VPS instances can be run simultaneously, and resource management can be performed to enable an efficient performance of each instance. A general example of VPS architecture is shown in Figure 1. Since servers are also isolated at the network level, they are completely separated from each other. Each VPS can have a unique IP address (or multiple addresses), and users can additionally configure their servers to potentially filter unwanted traffic.

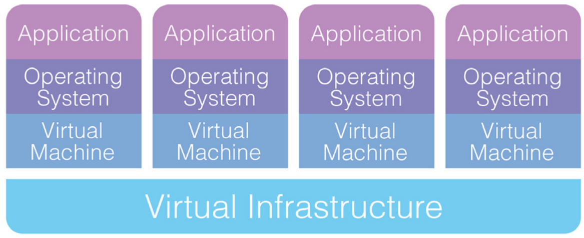

Figure 1. A typical architecture of VPS

Proper resource management is a highly critical component in the architectural design of a VPS, as it can play a major factor in providing the adequate performance and Quality of Service (QoS). Without proper resource management, a user that is overloading a server can potentially impact performance of other users hosted on the same physical server. Resource management can be dynamic or fixed, and it controls certain parameters like CPU power, disk space, memory, memory swap and others [44].

3.2 Description of Selected VPS Hosting Providers

An experimental research in this paper is conducted on the three popular and widely used VPS hosts. DigitalOcean and Linode are selected since they are market leaders while VULTR is commonly used for academic purposes. All three VPS hosts offer similar features. They have a variety of locations available for renting, support Linux operating systems, have high quality user interfaces for managing VPS instances, good customer support, newest technologies like Solid State Drives (SSDs) and newest processors, fast internet speed and other features.

DigitalOcean [19] provides developers with cloud services that help to deploy and scale applications that run simultaneously on multiple computers. It offers VPSs ("droplets") based on Kernel-based Virtual Machine (KVM) hypervisor, with various plans, options, GNU/Linux distributions, load balancers and available applications. It operates in 12 worldwide data centers and in January 2018 it was the third-largest cloud hosting company in the world in terms of web-facing computers.

Linode [20] offers KVM infrastructure based on Linux and robust set of tools which enables faster and easier applications development, deployment and scalability. It has ten data centers worldwide offering extensive peering relationships and Next Generation Network (NGN) for reducing latency and lowering the friction of scale. It also offers a Linode Backup service for automatic scheduled backups of servers, and Linode Manager and NodeBalancer that which can be used to control multiple server instances across a single system.

VULTR [21] offers the largest worldwide network spread across 16 data centers around the globe. Powerful networking features enable enterprises to deploy a highly customizable and scalable cloud infrastructure with high performance IP network, secure private networking, reserved IPs and redundant network design. It enables an easily scalable and low latency infrastructure solution with high processing power and fast I/O.

4 Benchmark Application

The performance measurement in this paper is conducted by using UnixBench, a system performance measurement and benchmarking tool [45]. It possesses comprehensive variety of individual tests targeted at specific areas such as CPU, RAM, graphics and hard drives and uses different algorithms and procedures. The following tests were used in this paper, as described in [46-47] and [48]:

(1) *Dhrystone*: This test represents CPU performance and focuses on handling strings without any floating-point calculations. It is heavily influenced by hardware, software, compiler, scripts, cache memory and integers.

(2) *Whetstone*: Similar to Dhrystone, it is focused on mathematical operations and it is designed to imitate the processor usage of several common set of programs. It uses a wide variety of C functions like *sin*, *cos*, *sqrt*, *log* and other integer and floating-point operations and measures the speed and efficiency of floating-point operations. The output metrics is defined with Millions of Whetstone Instructions Per Second (MWIPS) and the higher number means better performance.

(3) *Execl Throughput:* It tests execl throughput by measuring the number of execl calls that can be performed per second. The *exec()* family of functions replaces the current process image with a new one.

(4) *File Copy*: This test measures the rate at which files can be copied. It is executed with three different buffer sizes to determine the one best suited for a system.

(5) *Pipe Throughput*: It measures the number of times a process can write and read 512 bytes from a pipe.

(6) *Pipe-based Context Switching*: This test measures the number of times two separate processes can exchange an increasing integer through a pipe that often occurs in a realistic scenario.

(7) *Process Creation*: This test is designed to count the number of times a process can fork and reap the child process. It primarily tests the speed at which RAM can

(8) be allocated and then released, effectively measuring the RAM's bandwidth.

(9) *Shell Scripts:* It measures the number of times per minute that a process can start and reap concurrent copies of shell scripts. The test is performed with one and eight shell scripts.

(10) *System Call Overhead:* This final test estimates the cost of entering and leaving the operating system kernel. By repeatedly calling different system calls (such as *dup(), close(), getpid(), getuid(), and umask()*) it measures the number of returned calling processes *ids* per second and the highest number means better performance [49].

Each described test was run in three or ten iterations depending on the default test settings. In order to produce more reliable and consistent results we conducted an error analysis and found several peak measurement scores that were above or below the predetermined threshold of 100%. Those peak scores were discarded before calculating the average results since they are result of a system glitch (e.g. started operating system background processes) and could lead to misleading overall results. The obtained measurement results R_i are normalized to a set of baseline results R_{base} and the index values I_i are computed, as shown in (1) and described in [50]. Afterwards, these indexes from multiple test runs are added to the results and further computed and averaged by using the Perl code shown in Listing 1, based on the formula shown in (2). In the Perl code in the Listing 1 *$indexed* and *$numIndex* are kay-value lists that contain all categories and number of the tests in each category while *$indexed* will be empty if the test finishes without any result. The first two lines of the code add an object of those two lists so that they are available to the rest of the code. The code searches all keys in the *$indexed* key-value list and for all results calculates an average value for a certain category based on the *$sum* that is being calculated as the sum of natural logarithms of the index values of individual test runs - $\Sigma \ln I_i$. The '*index*' value is a newly calculated averaged index value for a certain category. Baseline results shown in Table 1 are benchmark scores obtained with Sun Microsystems workstation SPARCstation 20-61 that are used for calculating the index results. An overall index value of each test I_{over} is obtained by using a geometric mean of the individual normalized scores I_i, as shown in (3).

Listing 1. Calculating the index scores per category

```
$results->{'indexed'} = $indexed;
$results->{'numIndex'} = $numIndex;

foreach my $c (keys(%$indexed))
{
     if ($indexed->{$c}>0)
     {
          $results->{'index'}{$c}=exp($sum->
                 {$c}/$indexed-> {$c})*10;
     }
}
```

$$I_i = \frac{R_i}{R_{base}} \quad (1)$$

$$I_{cat} = e^{\frac{\sum \ln I_i}{n}} \times 10 \quad (2)$$

$$I_{over} = (\prod_{i=1}^{n} I_i)^{1/n} = \sqrt[n]{I_1 I_2 ... I_n} \quad (3)$$

Table 1. UnixBench baseline benchmark results

Metrics	Baseline Result R_{base}		Sample Period	Num. of Samples
dhry2reg	116700	lps	10.0 s	2
whetstone-dbl	55	MWIPS	10.0 s	2
execl	43	lps	20.0 s	1
fstime	3960	KBps	20.0 s	1
fsbuffer	1655	KBps	20.0 s	1
fsdisk	5800	KBps	20.0 s	1
pipe	12440	lps	10.0 s	2
context1	4000	lps	10.0 s	2
spawn	126	lps	20.0 s	1
shell1	42.4	lpm	60.0 s	1
shell8	6	lpm	60.0 s	1
syscall	15000	lps	10.0 s	2

5 Experimental Setup

5.1 Evaluated VPS Hosting Plans

In order to ensure efficient, reliable and fair performance comparison and evaluation, VPS instances are created on the three different VPS hosts with similar hosting plans, on the same location (London, UK), on the same Linux Debian distribution and the experiments were run at the same time of the day. Each VPS has several different plans with different resources available and during this research we measured performance of three different resource plans at each VPS by selecting the most similar ones in each experiment among all three different VPSs. Figure 2 shows the specifications of instances that are rented for the experimental analysis and performance measurement. In order to ensure thorough evaluation and cover all groups of users from less demanding to high demand users, different hosting plans were selected with the lowest, the medium and the highest resources available at each VPS. They are sorted by number of CPU cores and other specifications: 1 core, 2 cores and 4 cores with 1 GB, 2 GB and 8 GB of RAM, respectively. In addition, all hosts use SSD drives for a data storage. While the RAM specifications are similar across all instances of similar hosting plans, the third hosting plan on Linode deviates from this, as it has 4 GB of RAM instead of 8 GB that could impact its performance and measurement results. However, hosting plans are fixed at each provider and this one was the most similar one with the other two. The total number of different VPS instances evaluated is nine.

Figure 2. Specifications of evaluated VPS hosting plans

5.2 VPS Instances Setup Procedure

VPS instances setup procedure was very similar on all three VPS hosts and they were configured through available control panels. First, the Linux Debian distribution was installed and once the instances were ready, SSH protocol was used for terminal login and additional setup. Listing 2 shows the full procedure for all hosts that is identical in all experiments. However, DigitalOcean and VULTR first required manual installation of localization files, otherwise the system would keep showing a warning of missing localization on each reboot. This is the only difference between the three hosts. Afterwards, the operating system was upgraded to the newest version. This ensured that the systems are as similar as possible on all three VPS hosts that will minimize possible errors and differences in measurement process and results. After the setup procedure is finished, all three instances were rebooted. The final step before testing was to install UnixBench. It requires a build-essential package to be installed that was obtained with *apt-get* install command. Afterwards, UnixBench is downloaded using the *wget* command and extracted using *tar*. The systems are then again rebooted the one last time before running tests. The performance measurement procedure itself consists of simply running UnixBench from its directory and then system rebooting. In order to achieve more accurate and reliable results every test was repeated eight times. Once performance measurement was completed, the VPS instances were deleted from the user control panels and new ones for the next hosting plans were created.

Listing 2. UNIX commands executed during the VPS instance setup procedure

Only on DigitalOcean and VULTR:

Generate locale.
```
$ locale-gen UTF-8
```
Upgrade and restart the system.
```
$ apt-get update
$ apt-get upgrade
$ reboot
```

Prepare UnixBench.
```
$ apt-get install build-essential
$ wget https://github.com/kdlucas/byte-unixbench/archive/v5.1.3.tar.gz
$ tar xvf v5.1.3.tar.gz
$ reboot
```

Run UnixBench (repeat 8 times).
```
$ cd byte-unixbench-5.1.3/UnixBench/
$ ./Run
$ reboot
```

5.3 Performance Evaluation Model

Benchmark evaluation model requires performance metrics that is measurable and comparable among various VPS hosts. Therefore, we focused on examining and studying the key performance metrics which include CPU scheduling, memory management, hard disk drive management and Unix operating system task scheduling, similar to [51]. As described in Section 4, UnixBench uses a variety of tests to examine the performance of Unix based computer systems. Each of these tests uses different units with different sample running time and number of samples. The final values V_{fin} that are analyzed are the arithmetic means of results of eight separate test runs V_i. The formula for this is shown in (4).

$$V_{fin} = \frac{1}{N}\sum_{i=1}^{N} V_i \rightarrow N = 8 \quad \textbf{(4)}$$

The tests performed are identical across all VPS instances. Most of the tests measure the number of loops per second (lps) for specific operations. This is the number of times mathematical or system operations performed by UnixBench can be executed in a single second. The tests primarily cover the performance of CPU, RAM and SSD drives and their values are described and shown in Table 2. Three separate experiments were performed for different hosting plans and in order to minimize the difference among them, all specifications were set identically. In order to calculate the measurement percentage difference, DigitalOcean results were used as referent values to which other VPS hosts' performance measurement results are compared. DigitalOcean was chosen as a referent system since it has the highest number of users. The formula used to calculate the percentage difference is shown in (5).

$$Diff\,(\%) = \frac{Lin\,/\,VULTR_{value} - DigOc_{value}}{DigitalOcean_{value}} \times 100\% \quad \textbf{(5)}$$

Table 2. Performance metrics, sample period and units

Metrics	Unit	Sample Period	Number of Samples
Dhrystone 2 using register variables	lps	10.0 s	7 samples
Double-Precision Whetstone	MWIPS	9.8 s	7 samples
Execl Throughput	lps	30.0 s	2 samples
File Copy 1024 bufsize 2000 maxblocks	KBps	30.0 s	2 samples
File Copy 256 bufsize 500 maxblocks	KBps	30.0 s	2 samples
File Copy 4096 bufsize 8000 maxblocks	KBps	30.0 s	2 samples
Pipe Throughput	lps	10.0 s	7 samples
Pipe-based Context Switching	lps	10.0 s	7 samples
Process Creation	lps	30.0 s	2 samples
Shell Scripts (1 concurrent)	lpm	60.0 s	2 samples
Shell Scripts (8 concurrent)	lpm	60.0 s	2 samples
System Call Overhead	lps	10.0 s	7 samples

6 Performance Evaluation Results

Three separate experiments were conducted on different VPS instances sorted by the hosting plan, ordered primarily by the number of CPU cores, followed by the RAM amount, SSD capacity and network bandwidth. Results with significant differences are graphed and plotted, and other results with smaller differences are tabulated. In addition, performance measurement results percentage difference between the reference value and other VPSs is calculated and discussed in the text. All results were rounded to one decimal place since the second enters in the area of a measurement error.

6.1 The First Experiment

The first hosting plan specifications include 1 CPU core and 1 GB of RAM. The SSD drives capacity deviates between different hosts where DigitalOcean has 30 GB, Linode has 24 GB and VULTR has 20 GB of storage. In addition, all three instances have 2 TB of network bandwidth.

The first noticeable difference between VPSs in this experiment was obtained in Execl throughput test shown in Figure 3. DigitalOcean has the poorest performance at just 4,159.7 lps, while the best performance is achieved on VULTR at 4,833.5 lps with Linode achieving middle performance at 4,506.2 lps. Therefore, Linode outperforms DigitalOcean by 8.3% and VULTR by 16.2%. The second biggest difference can be seen during the Pipe-based Context Switching test, whereby DigitalOcean again exhibits the poorest performance at 309,798.7 lps, and VULTR the best performance at 370,769.0 lps. Similar to the previous test, Linode is situated in the middle with 341,024.1 lps. Under this test, Linode outperforms DigitalOcean by 10.08%, while VULTR by 19.68%. These results are plotted in Figure 3 (right).

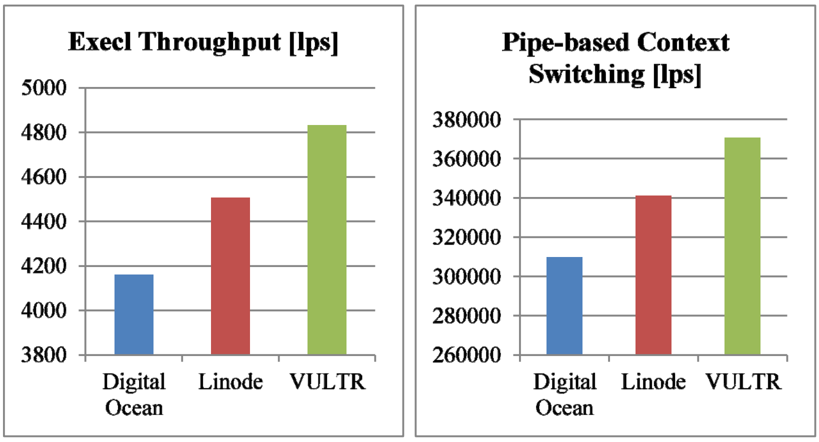

Figure 3. *Execl* throughput and pipe-based context switching results

The final noticeable difference was shown when using shell scripts testing. Unlike the previous two tests, Linode gives the worst performance, achieving a performance of 7,278.3 loops per minute (lpm) and 928.1 lpm for 1 and 8 concurrent scripts, respectively. VULTR again comes out at the top with 8,002.2 and 1,036.6 lpm, with DigitalOcean being close to Linode, but still slightly better at 7,366.9 and 960.9 lpm. These results make Linode slightly slower than DigitalOcean (1.2% and 3.4%, respectively), and VULTR noticeably better (8.6% and 7.8%, respectively). These results can be seen in the two graphs on Figure 4.

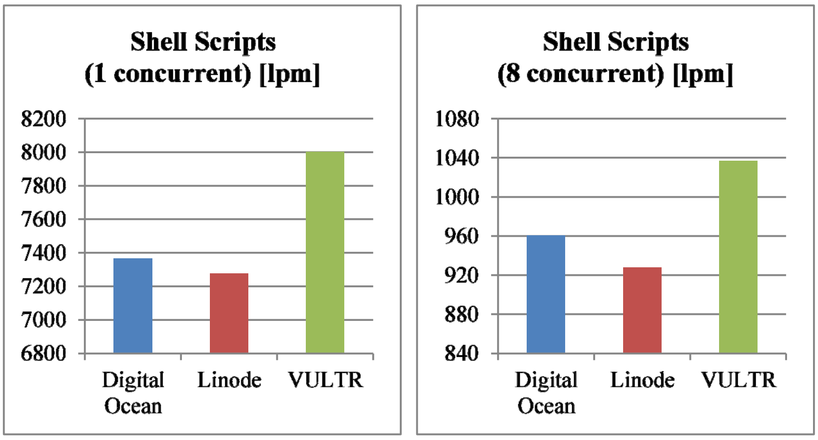

Figure 4. Results of the test involving shell scripts

The rest of the results are shown in Table 3 and are less noticeable, but still important to report. According to these results, VULTR achieves highest results in five different tests considering RAM and SSD handling while DigitalOcean in just one, the System Call Overhead test which means it is probably the best one at handling kernel processes. Although Linode shows the lowers performance in general in this test, it has the best CPU performance as it can been seen in Dhrystone and Whetstone tests results. Those results are aligned with the overall benchmark VPS performance score I_{over} that shows similar overall performance for DigitalOcean and Linode while VULTR having more than 8% better overall performance.

Table 3. Other performance measurement results in first experiment

Test	DigitalOcean	Linode	VULTR
Dhrystone 2 using register variables	27362726.4 lps	**27712759.7 lps**	26048713.0 lps
Double-Precision Whetstone	3484.9 MWIPS	**3592.5 MWIPS**	3426.2 MWIPS
File Copy 1024 bufsize 2000 maxblocks	880200.8 KBps	925459.3 KBps	**987083.8 KBps**
File Copy 256 bufsize 500 maxblocks	253026.4 KBps	273905.1 KBps	**280653.2 KBps**
File Copy 4096 bufsize 8000 maxblocks	1636312.7 KBps	1621821.3 KBps	**2028268.9 KBps**
Pipe Throughput	1862827.7 lps	1899487.1 lps	**1925244.2 lps**
Process Creation	10806.0 lps	10384.6 lps	**11393.2 lps**
System Call Overhead	**3772682.6 lps**	3565460.0 lps	3769993.8 lps

6.2 The Second Experiment

Second hosting plan includes 2 CPU cores and 2 GB of RAM across all hosts but again as in the first experiment, SSD sizes are slightly different. DigitalOcean has 40 GB, Linode has 48 GB and VULTR has 45 GB storage. All three hosts have 3 TB of network bandwidth.

Similar to the previous experiment, the first noticeable difference was obtained in the Execl throughput test (Figure 5 - left), where DigitalOcean achieves the lowest performance at 7,800.8 lps, Linode comes close in the second place with 7,889.3 lps and the best performance is achieved by VULTR at 8,518.4 lps. The difference between Linode and DigitalOcean is negligible (1.3%), but VULTR's score is 9.2% better when compared to the DigitalOcean's score. The second noticeable difference is in pipe-based context switching (Figure 5 - right). Here, DigitalOcean shows the lowest performance at 572,263.9 lps. Linode managed to run it at 638,849.4 lps, close to VULTR's score, which is 645,462.3 lps. The difference between DigitalOcean and the other two hosts is relatively similar since Linode shows 11.6% and DigitalOcean 12.8% better performance when compared to VULTR.

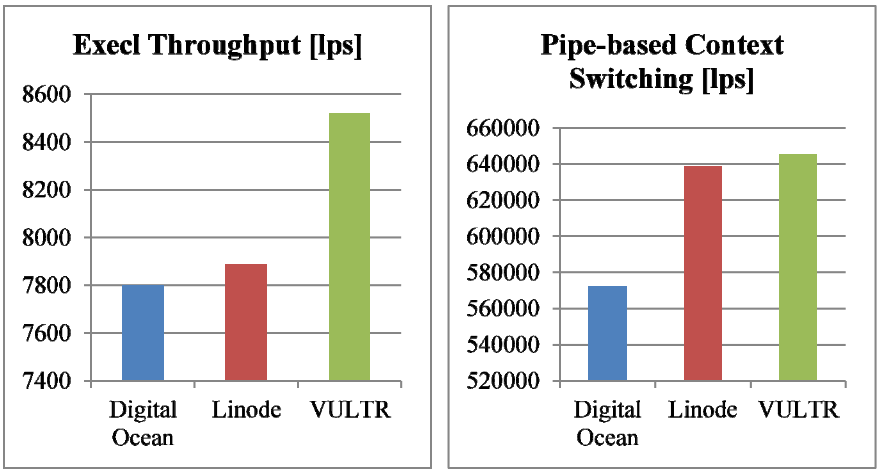

Figure 5. *Excel* throughput and pipe-based context swiching results

The third important difference is in the process creation test. DigitalOcean is the slowest at 16,582.8 lps. Linode is second at 17,237.0 lps, and VULTR is the fastest at 18,985.6 lps which means that Linode is 3.95% and VULTR is 14.49% faster than DigitalOcean. These results are exhibited in Figure 6 (left). The final bigger difference is visible in the results of the System call overhead test with the results shown in Figure 6 (right). VULTR and DigitalOcean's performance is similar (1.3% difference) but Linode is the slowest, having approximately 20% lower measurement results compared to other two VPSs. It means that it is the slowest at handling and managing kernel processes and consequently at executing the hardware-depended instructions that can influence applications performance when switching from user to kernel mode. Additional results are shown in Table 4. Similar to the previous experiment, VULTR obtained the highest scores in most of the tests, but this time, DigitalOcean achieved the best result under the pipe throughput test. Linode still shows the best CPU performance under the results of Dhrystone and Whetstone tests. VULTR is still the best in RAM and SSD management and operating system performance, but this time shows the lowest pipe throughput. Again, the results obtained in the Experiment 2 are aligned with the overall benchmark VPS performance score I_{over} that shows similar overall performance for DigitalOcean and Linode while VULTR shows 6.6% and 7.7% better overall performance.

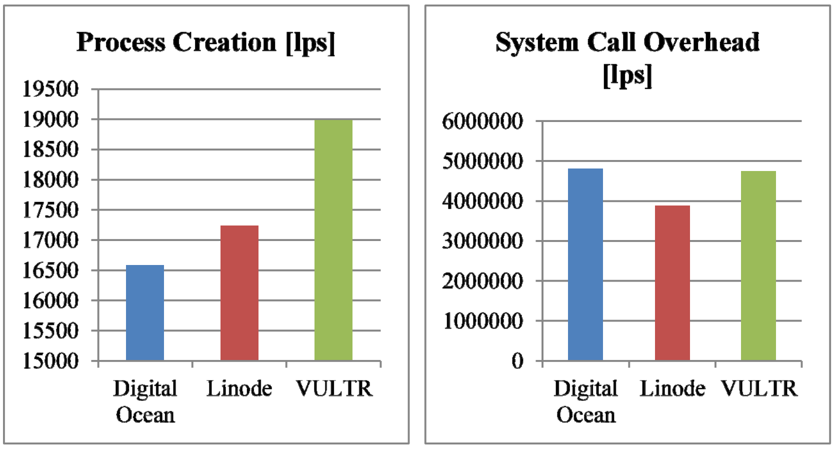

Figure 6. Process creation and system call overhead results

Table 4. Other performance measurement results in second experiment

Test	DigitalOcean	Linode	VULTR
Dhrystone 2 using register variables	55551085.9 lps	**56658679.3 lps**	54199197.0 lps
Double-Precision Whetstone	7124.8 MWIPS	**7251.3 MWIPS**	6963.6 MWIPS
File Copy 1024 bufsize 2000 maxblocks	922828.2 KBps	937483.5 KBps	**1051609.7 KBps**
File Copy 256 bufsize 500 maxblocks	255616.9 KBps	251309.6 KBps	**279408.0 KBps**
File Copy 4096 bufsize 8000 maxblocks	2403896.8 KBps	2762294.3 KBps	**3043364.7 KBps**
Pipe Throughput	**3636012.2 lps**	3511853.1 lps	3564136.5 lps
Shell Scripts (1 concurrent)	12960.8 lpm	12973.9 lpm	**14146.3 lpm**
Shell Scripts (8 concurrent)	1806.2 lpm	1897.1 lpm	**2000.4 lpm**

6.3 The Third Experiment

The third and the final hosting plan that was tested includes 4 CPU cores and 8 GB of RAM. However, Linode plan slightly differs, as it contains only 4 GB of RAM. There was no other plan available with 4 CPU cores, so this one was chosen for the test as the number of cores was determined as the most important factor. Furthermore, SSD drive capacities differ also in this test since DigitalOcean provides 80 GB, Linode 96 GB and VULTR 150 GB of storage. Both DigitalOcean and VULTR provide 5 TB of network bandwidth, while Linode provides only 4 TB.

The file copy test measurement results shown in Figure 7 notably deviate between the hosts. Linode appears to have the lowest results for all three different file copy tests: 758,753.5 KBps, 207,384.0 KBps and 2,324,683.7 KBps for buffer sizes of 1,024 bytes, 256 bytes and 4,096 bytes, respectively. DigitalOcean came in second place with 892,060.6 KBps, 238,145.4 KBps and 2,402,406.6 KBps. VULTR showed the best performance with the following results: 957,395.7 KBps, 255,166.3 KBps and 2,682,852.4 KBps. In other words, Linode obtains 14.9%, 12.9% and 3.2% lower results than DigitalOcean, while VULTR obtains 7.3%, 7.1% and 11.67% better results. The lower results for Linode VPS could be a consequence of double smaller available RAM size when compared to other two VPS plans since a Linux file transfer system uses RAM caching in order to make file transfer faster.

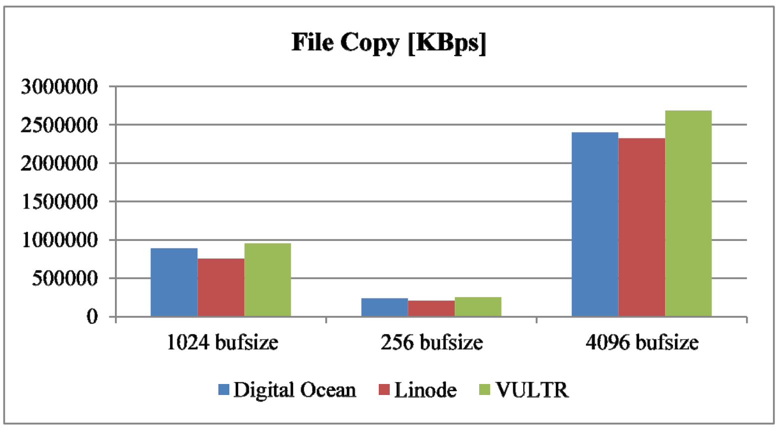

Figure 7. File copy results

The second obvious difference is in the process creation test, where Linode shows the lowest results at 22,573.1 lps. The second one is DigitalOcean which is able to perform at 24,600.7 lps, and finally VULTR shows the best results with 34,276.6 lps. Therefore, Linode achieves 8.2% lower results, while VULTR shows significantly better results, even 39.3% higher when compared to DigitalOcean. Results of this test are shown in Figure 8 (left). The final interesting results are obtained with system call overhead test, shown in Figure 8 (right). DigitalOcean shows the best results of 6,558,867.1 lps, while VULTR obtains 3.9% lower and Linode even 31.1% lower performance measurement results. This shows that DigitalOcean has the fastest handling and managing of kernel processes that can dramatically influence on user applications' performance.

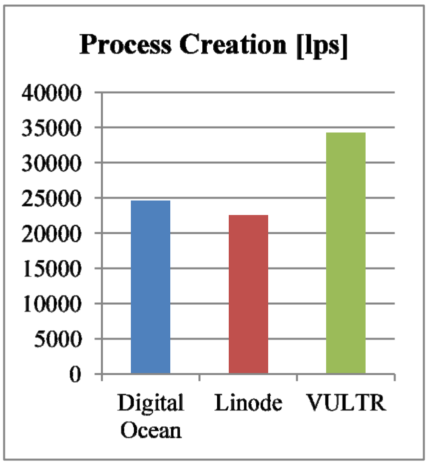

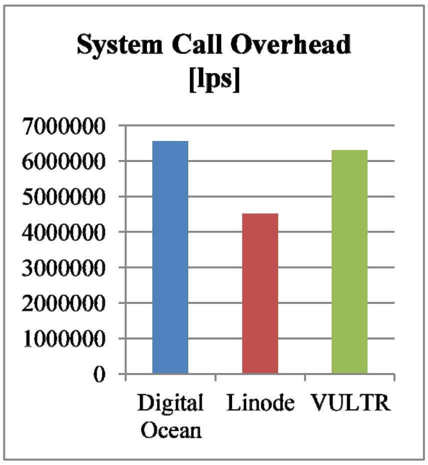

Figure 8. Process creation and system call overhead results

Results of other tests are shown in Table 5. It can be observed that DigitalOcean obtains the best results in almost all the tests which is interesting to observe, since in the past two experiments VULTR was achieving most of the best results and Digital Ocean had only several performance measurements that achieved the best results. Unlike the first two experiments, DigitalOcean shows the most efficient CPU handling and RAM management. In addition, it also shows the best performance in all operating system tests, making it the best at handling kernel processes. VULTR remains the best in terms of SSD performance, and Linode shows the lowest performance in all experiments except the pipe-based context switching. The overall score I_{over} shows that DigitalOcean and VULTR have obtained the best similar results in Experiment 3, much higher when compared to Linode, 8% and 10%, respectively.

Table 5. Other performance measurement results in third experiment

Test	DigitalOcean	Linode	VULTR
Dhrystone 2 using register variables	**116085140.6 lps**	112194227.8 lps	102696308.7 lps
Double-Precision Whetstone	**14530.9 MWIPS**	14476.0 MWIPS	13634.6 MWIPS
Execl Throughput	**16202.1 lps**	15292.7 lps	15336.2 lps
Pipe Throughput	**7364027.7 lps**	6984896.5 lps	6859096.4 lps
Pipe-based Context Switching	1179035.1 lps	**1278398.1 lps**	1245024.4 lps
Shell Scripts (1 concurrent)	**25813.0 lpm**	25331.7 lpm	25167.9 lpm
Shell Scripts (8 concurrent)	**3917.6 lpm**	3719.9 lpm	3672.2 lpm

7 Conclusion

Virtualization technology transforms traditional online services and has become an essential ingredient of the future information infrastructure. VPS service improves the cloud computing paradigm by offering a cost-effective way of providing a powerful and flexible multipurpose server with high availability, reliability, scalability and management flexibility. Performance evaluation of VPSs is crucial and beneficial to both service providers and service consumers and currently VPS performance is marginally represented in scientific research. Therefore, a comprehensive survey on related work, virtualization, VPSs architectural design and benchmark application is conducted in this paper in order to reflect the latest status in the field. Furthermore, a measurement-based approach is applied for detailed performance evaluation of the three popular VPS hosts, namely DigitalOcean, Linode and VULTR. In order to achieve accurate and reliable performance measurement results, three separate experiments were conducted on different hosting plans with multiple repeated performance measurements. Based on the performance evaluation results of all three experiments in terms of key system metrics which include CPU scheduling, memory management, hard disk drive performance and operating system task performance it can be concluded that VULTR provides the best performance for low and medium demanding users while DigitalOcean is the best solution for more demanding users who are looking for a high performance VPS host.

References

[1] S. Pal, Extending Mobile Cloud Platforms Using Opportunistic Networks: Survey, Classification and Open Issues, *Journal of Universal Computer Science*, Vol. 21, No. 12, pp. 1594-1634, December, 2015.

[2] K. Ma, R. Sun, A. Abraham, Toward a Module-centralized and Aspect-oriented Monitoring Framework in Clouds, *Journal of Universal Computer Science*, Vol. 19, No. 15, pp. 2241-2265, September, 2013.

[3] M. Perisa, G. Markovic, P. Kolarovszki, R. Madleňák, Proposal of a Conceptual Architecture System for Informing the User in the IoT Environment, *PROMET*, Vol. 31, No. 1, pp. 37-47, February, 2019.

[4] K. Venkatraman, K. Geetha, Dynamic virtual cluster cloud security using hybrid steganographic image authentication algorithm, *Automatika*, Vol. 60, No. 3, pp. 314-321, June, 2019.

[5] Z. Balogh, E. Gatial, L. Hluchy, R. Toegl, M. Pirker, D. Hein, Agent-Based Cloud Resource Management for Secure Cloud Infrastructures, *Computing and Informatics*, Vol. 33, No. 6, pp. 1333-1355, January, 2014.

[6] C. Li, L. Li, An Optimization Approach for Utilizing Cloud Services for Mobile Devices in Cloud Environment, *Informatica*, Vol. 26, No. 1, pp. 89-110, March, 2015.

[7] U. Rugwiro, C. Gu, W. Ding, Task Scheduling and Resource Allocation Based on Ant-Colony Optimization and Deep Reinforcement Learning, *Journal of Internet Technology*, Vol. 20, No. 5, pp. 1463-1475, September, 2019.

[8] M. Pearce, S. Zeadally, R. Hunt, Virtualization: Issues, Security Threats, and Solutions, *ACM Computing Surveys*, Vol. 45, No. 2, pp. 1-39, February, 2013.

[9] F. Lin, J. Xiahou, W. Zeng, A Study of Risk Evaluation Framework for Cloud Computing System, *Journal of Internet Technology*, Vol. 16, No. 7, pp. 1351-1366, December, 2015.

[10] S. Nesmachnow, S. Iturriaga, B. Dorronsoro, P. Bouvry, E. G. Talbi, Metaheuristics for the Virtual Machine Mapping Problem in Clouds, *Informatica*, Vol. 26, No. 1, pp. 111-134, March, 2015.

[11] N. Rajganesh, T. Ramkumar, A Review on Broker Based Cloud Service Model, *Journal of Computing and Information Technology*, Vol. 24, No. 3, pp. 283-292, September, 2016.

[12] M.-C. Lee, J.-C. Lin, R. Yahyapour, Hybrid Job-Driven Scheduling for Virtual MapReduce Clusters, *IEEE Transactions on Parallel and Distributed Systems*, Vol. 27, No. 6, pp. 1687-1699, June, 2016.

[13] N. Khan, A. Shah, K. Nusratullah, Adoption of Virtualization in Cloud Computing: A Foundation Step towards Green Computing, *International Journal of Green Computing (IJGC), IGI Global,* Vol. 6, No. 1, pp. 40-47, January, 2015.

[14] D. Wang, J. Liu, Optimizing Big Data Processing Performance in the Public Cloud: Opportunities and Approaches, *IEEE Network*, Vol. 29, No. 5, pp. 31-35, September-October, 2015.

[15] Cloud Spectator, *Cloud Vendor Benchmark 2015: Price-Performance Comparison Among 15 Top IaaS Providers*, Cloud Spectator, 2015.

[16] Cloud Spectator, *Cloud Spectator Releases 2017 Top Ten European Cloud Service Providers Report*, Cloud Spectator, 2017.

[17] J. L. Wilson, *The Best VPS Web Hosting Services for 2019*, http://www.pcmag.com/article2/0,2817,2455706,00.asp.

[18] A. Leslie, P. J. Fancher, R. Frankel, *2019's Best "VPS" Hosting Reviews*, http://hostingcurator.com/reviews/vps/index.html.

[19] *DigitalOcean*, https://www.digitalocean.com/.

[20] *Linode*, https://www.linode.com/.

[21] *Vultr Holdings: LLC.*, https://www.vultr.com/.

[22] G. Martinovic, J. Balen, S. Rimac-Drlje, Impact of the Host Operating Systems on Virtual Machine Performance, *33rd International Convention MIPRO*, Opatija, Croatia, 2010, pp. 613-618.

[23] G. Martinovic, J. Balen, B. Cukic, Performance Evaluation of Recent Windows Operating Systems, *Journal of Universal Computer Science*, Vol. 18, No. 2, pp. 218-263, January, 2012.

[24] G. Martinovic, J. Balen, Z. Hocenski, Network Performance Evaluation of Latest Windows Operating System, *20th International Conference on Software, Telecommunications and Computer Networks*, Split, Croatia, 2012, pp. 1-6.

[25] K. Salah, M. Al-Saba, M. Akhdhor, O. Shaaban, M. I. Buhari, Performance Evaluation of Popular Cloud IaaS Providers, *International Conference for Internet Technology and Secured Transactions*, Abu Dhabi, UAE, 2011, pp. 345-349.

[26] VPS Benchmarks: Compare VPS Plans, http://www.vpsbenchmarks.com/.

[27] R. Asher, *Qualitative Comparison of Linode, Digital Ocean and Vultr*, https://www.lowendtalk.com/discussion/41430/qualitative-comparison-of-linode-digital-ocean-and-vultr-updated.

[28] J. Wight, *Virtual Server Evaluation and Comparison Rackspace vs DigitalOcean vs Linode vs Vultr*, http://leadkrabi.com/2015/06/virtual-server-evaluation-and-comparison-rackspace-vs-digital-ocean-/vs-linode-vs-vultr/.

[29] J. D. Castro, VPS Comparison, https://github.com/joedicastro/vps-comparison.

[30] K. Kim, M. E. Pierce, R. Guha, SQMD: Architecture for Scalable, Distributed Database System Built on Virtual Private Servers, *IEEE Fourth International Conference on e-Science*, Indianapolis, USA, 2008, pp. 658-665.

[31] A. Almurayh, *Virtual Private Server*, UCCS-CS526, May,

2010.

[32] D. M. F. Mattos, L. H. Ferraz, L. H. Costa, C. Duarte, Evaluating Virtual Router Performance for a Pluralist Future Internet, *The 3rd International Conference on Information and Communication Systems*, New York, USA, 2012, pp. 1-7.

[33] D. M. F. Mattos, L. H. Ferraz, L. H. Costa, C. Duarte, Virtual Network Performance Evaluation for Future Internet Architectures, *Journal of Emerging Technologies in Web Intelligence*, Vol. 4, No. 4, pp. 304-314, November, 2012.

[34] P. Prakash, R. Biju, M. Kamath, Performance Analysis of Process Driven and Event Driven Web Servers, *The 9th IEEE International Conference on Intelligent Systems and Control*, Coimbatore, India, 2015, pp. 1-7.

[35] I. Mavridis, H. Karatza, Combining Containers and Virtual Machines to Enhance Isolation and Extend Functionality on Cloud Computing, *Future Generation Computer Systems*, Vol. 94, pp. 674-696, May, 2019.

[36] Y. Yamato, Performance-aware Server Architecture Recommendation and Automatic Performance Verification Technology on IaaS Cloud, *Service Oriented Computing and Applications*, Vol. 11, No. 2, pp. 121-135, June, 2017.

[37] Y. Yamato, *Server Selection Device, Server Selection Method, and Server Selection Program*, US20180113734A1, April, 2018.

[38] D. W. Chang, R. M. Pathak, *Performance-based Public Cloud Selection for a Hybrid Cloud Environment*, US201700 97841A1, September, 2018.

[39] R. Birke, Y. L. Chen, M. L. Schmatz, *Optimization of Virtual Machine Sizing and Consolidation*, US9569236B2, February, 2017.

[40] S. Kaiser, D. Chosnek, *Ranking of Computing Equipment Configurations for Satisfying Requirements of Virtualized Computing Environments Based on an Overall Performance Efficiency*, US9069890B2, June, 2015.

[41] A. Johnsson, C. Flinta, J. Christer, W. Granath, *Compute Infrastructure Resource Monitoring Method and Entities*, US20180107577A1, April, 2018.

[42] S. Furman, *Shared Web Hosting vs VPS Hosting vs Dedicated Servers*, http://www.codedwell.com/post/75/shared-web-hosting-vs-vps-vs-dedicated-servers.

[43] I. Jimenez, C. Maltzahn, J. Lofstead, A, Moody, K. Mohror, R. Arpaci-Dusseau, A. Arpaci-Dusseau, Characterizing and Reducing Cross-Platform Performance Variability Using OS-Level Virtualization, *IEEE International Parallel and Distributed Processing Symposium Workshops*, Chicago, USA, 2016, pp. 1077-1080.

[44] M. Niswar, A.-A. Sabri, E. Warni, M.-N. Musa, Memory Sharing Management on Virtual Private Server, *International Conference on ICT for Smart Society*, Jakarta, Indonesia, 2013, pp. 1-4.

[45] *Byte-unixbench*, http://code.google.com/p/byte-unixbench/.

[46] H. A. Ismail, M. Riasetiawan, CPU and Memory Performance Analysis on Dynamic and Dedicated Resource Allocation Using XenServer in Data Center Environment, *2nd International Conference on Science and Technology-Computer*, Yogyakarta, Indonesia, 2016, pp. 17-22.

[47] M. Raho, A. Spyridakis, M. Paolino, D. Raho, KVM, Xen and Docker: A Performance Analysis for ARM Based NFV and Cloud Computing, *3rd IEEE Workshop on Advances in Information, Electronic and Electrical Engineering*, Riga, Lativa, 2015, pp. 1-8.

[48] J. P. Mullerikkal, Y. Sastri, A Comparative Study of OpenStack and CloudStack, *Fifth International Conference on Advances in Computing and Communications*, Kochi, India, 2015, pp. 81-84.

[49] T. Critchley, *High-Performance IT Services*, Auerbach Publications, 2016.

[50] K. Lucas, *UnixBench*, https://github.com/kdlucas/byte-unixbench/tree/master/UnixBench.

[51] J. Feng, L. Zhang, J. Lu, B. Xu, CPE: A Cloud Server Performance Evaluation Model, *The 11th International Conference on Ubiquitous Information Management and Communication*, Beppu, Japan, 2017, pp. 1-8.

Biographies

Josip Balen is an assistant professor at Faculty of Electrical Engineering, Computer Science and Information Technology in Osijek, Croatia, where he received his Ph.D. study in 2014. His main research field include efficient message dissemination in VANETs and performance evaluation of various systems. He has been to many short and long research visits across Europe and he participated in several EU and national scientific and professional projects. For his work within IEEE organization he has won a dozen awards.

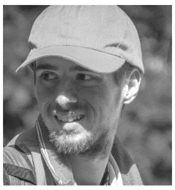

Denis Vajak is a Ph.D. student at the Faculty of Electrical Engineering, Computer Science and Information Technology in Osijek, Croatia. He has achieved B.Sc in informatics and M.Sc in the field of Computer Engineering. His main research field is automotive engineering. He spends most of his free time developing new applications and software, and exploring the possibilities and future of virtual reality.

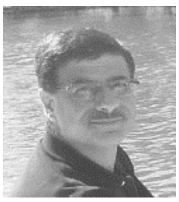

Khaled Salah is a full professor at the ECE Department, Khalifa University of Science, Technology and Research, UAE. Khaled has over 145 publications related to the areas of cloud computing, operating systems, and performance evaluation. He was also the recipient of Khalifa University Outstanding Research Award 2014, KFUPM University Excellence in Research Award 2009, and KFUPM Best Research Project Award 2010. He is a senior member of IEEE and has been serving on the editorial boards of a number of WoS-listed journals.