Cloud Computing in Amazon and Microsoft Azure platforms: performance and service comparison

Borislav S. Đorđević, Slobodan P. Jovanović and Valentina V. Timčenko, *Member, IEEE*

Abstract — **This paper addresses the performance and service comparison of two different Cloud Computing platforms: Amazon AWS and Microsoft Azure. Platforms are tested in similar virtual environments, namely for micro instances. Performance is measured by the collection benchmark program called Phoronix Test Suite 3, and the results are presented for Apache and Dbench benchmarks.**

Key words — **Cloud Computing, virtualization, Amazon AWS, Microsoft Azure, micro instance, performances, Phoronix test Suite.**

I. INTRODUCTION

C LOUD Computing (CC) stands for a technology that enables the provision of flexible access to available system computing resources. It operates independently of the resources physical location, allowing its fast and seamless allocation and reallocating based on the provided user demands. The available storage resources are usually virtualized and abstracted, whereas the CC is represented as a rising IT environment model that enables the development, deployment and real-time delivery of services, products and solutions over the Internet, while the charging for the use of such services is usually in proportion to the amount of the used resources.

As the CC at its core represents the concept of sites independence, provider resources are shared as needed to efficiently serve multiple users, by dynamically assigning and reallocating different physical and virtual resources according to the user's needs and subscriptions. The main advantage of this concept is that there is an initial investment in the information system, which is usually one of the main obstacles for the development of the information system of SMEs (Small and Medium Enterprises), as there is a need for high investments in the implementation of such systems.

Nevertheless, most of the legal issues related to this model arise from the fact that CC is a distributed technology in global market, whereas we still live in an era of territorially limited jurisdictions.

Therefore, it is highly needed to cope efficiently with the problems of data protection and confidentiality, intellectual property assurance, data access regulation, adequate law authority involvement, etc. CC technology is characterized by: high flexibility, low operating costs, devices/users independence of storage location, the possibility of sharing resources, all encompassed with guaranteed reliability, scalability and safety.

CC is a natural successor to virtualization which is the technology relying on service-basedarchitecture, providing to the user the abstraction and encapsulation of computer hardware and software elements, with an aim to use them efficiently for certain needs. Virtualization practically allows more efficient infrastructure use, as it enables its simultaneous utilization for different operation systems. When combined for CC needs, it provides more secure and reliable environment, mainly relying on virtual machine's (VM) isolation. The users of CC technology have no needs to be aware of system details, accessing the documents and resources by a simple web browser account to the interface of operator's cloud. The most important benefit of such approach is the convenience of having at the disposal a collection of data centres in one place while the price of using this service is several times lower than the cost of maintaining own data centre.

A data center is ideal for companies that need a customized, dedicated system that gives them full control over their data and equipment. A data center, however, has limited capacity - once company builds a data center, it will not be able to change the amount of storage and workload it can withstand without purchasing and installing more equipment.On the other hand, a cloud system is scalable to a company business needs. It has potentially unlimited capacity, based on the vendor's offerings and service plans. One disadvantage of the cloud is that user will not have as much control as it would a data center, since a third party is managing the system.

II. RELATED WORKS

In the last few years, there has been extensive research activity in the emerging area of Cloud. One of the most relevant topics is providing methodology and quantitative evaluation of QoS and different CC service performance levels [1]. Some authors have addressed the problem of possible interference between CCusers when sharing resources, thus have proposed different types of metrics for quantifying the performance isolation of cloud-based systems. They also specify new approaches for achieving

performance isolation in CC [2].The data centers are consuming a huge amount of energy, which is expected to grow noticeably under the actual technological trends. Thus, a number of studies target the arising problem of providing energy efficient CC solution [3]. To deal with the challenge of the CC resource scheduling optimization, some authors proposed specific cost-effective solution for service request scheduling in CC, with a goal to process more cost-effectively the dynamic user service requestswithout any Service Level Agreement (SLA) violation [4].The design enhancement, security policies and mechanisms enforcement are also explored in some studies [5-6]. This paper focuses on performance comparison of two CC platforms: Amazon AWS and MS Azure.

III. THE ARCHITECTURE AND CATEGORIES OF CLOUD COMPUTING SERVICES

The basis of CC is the IT infrastructure, which relies on a range of applied technologies enforcing the abstraction of physical resources through virtualization and their further sharing to different users. The CC architecture model is based on two entities:

Front end: represents the user controllable infrastructure characteristics, as well as the specifications for the access of the user to the cloud.

Back end: represents the cloud provider infrastructure. This model allows enterprises to upload and use applications much faster, with better control and less needs for maintenance, allowing IT sector companies to quickly and efficiently meet variable, unpredictable requirements.

Central server is responsible for system management, traffic and client requests monitoring this way providing stable functionality of the system. Cloud consists of a set of computers or server units, which intercommunicate based on specific middleware software. CC architecture is modular system encompassing for layers: hardware layer, infrastructure layer, platform layer and application layer.

Hardware layer is responsible for the management of physical cloud resources, including servers, routers, power and cooling systems management while it also takes care of system fault-tolerance management. It is typically represented by data centers, which can contain network of thousands of servers.

The infrastructure layer is often referred to as the virtualization layer since it is achieved by applying different virtualization techniques over the available physical resources. This layer is an essential component of CC as it encompasses computing and storage resources. Therefore, it is possible to provide dynamic resources allocation and their more efficient use.

Layer platform encompasses applied operating systems and application frameworks.

Application layer is the top of the CC architecture hierarchy layers, mainly consisting of needed up-to-date cloud applications. Unlike traditional applications, cloud applications can be automatically scaled in a way to achieve better performance, availability, and operate with lower operating costs.

Modular architecture enables CC to support a wide range of application requirements, thereby reducing the cost of management and maintenance.

CC provides three categories of services: Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), and Software-as-a-Service (SaaS).

A. IaaS : Infrastructure- as-a-Service

IaaS is a service provision model that implies the outsourcing of the organization equipment in a way to support other users operations, including storage, hardware, servers and networking components. The IaaS users typically pay the service on a per-use basis. Providers offer virtual machines, storage space, firewalls, and network load balancers. These resources can be rented on request, whereas the client does not control neither maintains the cloud infrastructure, but only manages on the site operating system, software updates and user-side applications [7].

B. PaaS : Platform- as-a-Service

PaaS is a service delivery model allows the user to rent virtualized servers and associated services for use of the existing applications and also for development, testing and use of new applications. However, unlike traditional systems, PaaS ensures low cost environment for the development of scalable applications [7].

C. SaaS : Software- as-a-Service

SaaS represents a software distribution model that relies on the applications hosted by service provider and their availability and global accessibility to the user via specific network or Internet. This approach allows for the optimal use of software resources, easier administration, costs decrease and needlessness of purchasing licenses, onsite installation, maintenance and ownership of the hardware resources needed for proper functioning of the applications. One of the most important benefits of SaaS over the traditional applications is the elasticity achieved by distributing the workload across a set of virtual machines. SaaS is comparable to the application service provider and on demand computing software delivery models. It can be delivered in a form of two slightly different models: (1) the hosted application management model, which is analogous to application service provider, as the provider hosts commercially available software for users and delivers it over the Internet; (2) the software on demand model is based on providing users a network-based access to a single copy of an application which is exclusively generated for SaaS distribution [7].

IV. TEST PLATFORMS

A. Microsoft Azure

Microsoft Azure is a flexible cloud platform that allows fast development, debugging and iteration of the applications, as well as their further management through a network of Microsoft data centers[8-9]. Applications can be developed with any tool, programming language, or existing framework, while there is possibility of integrating public cloud applications with existing IT environment.

To subscribe to the Azure, it is necessary to use some of the Microsoft Live accounts (Live, Hotmail, Outlook) and credit card. After completing the registration the user can make the purchase of needed resources in the cloud. VMs generation is performed from a management console, with the possibility of selecting different options among the list of those that are available. Fig. 1. provides and overview of creating VM environment under Linux, Ubuntu Server ver. 14.04LTS, 768MB RAM and one CPU core.

B. Amazon Web Services - AWS

Amazon Web Services is a CC platform offered by Amazon.The main features of the service include: low price, high speed, scalability, openness, adaptability, and guaranteed security.

EC2 and S3 are two most used options of this CC platform [10].EC2 (Elastic CC) is a central part of Amazon's platform. It allows users to rent virtual machines on which they can run their applications. Amazon S3 (Simple Storage Service) is an online file storage web service. It provides storage through web services interfaces using REST and SOAP protocols.

Amazon provides online services to other web sites or client applications, thus most of these services are not available to end users, but instead allow the developers the use and advantage of Amazon platform functionalities while developing their own applications. All services are charged according to the type of service and duration of its use. AWS is open, agile and flexible platform, thus it allows use of any programming language or development platform. The available amount of resources can be tailored to the user needs and dynamically increased or decreased by adding or removing processor cores or disk memory. AWS integrates a range of industrial security certificates: PCI DSS Level 1, ISO 27001, FISMA Moderate, HIPAA, and SAS 70 II.

Through the web interface user launches its VM instance, using the so-called Amazon Machine Image (AMI), a predefined template with the installation of the operating system. The process of creating a VM is similar to that in Azure, while AWS offers larger range of configuration options, such as choice of multiple versions of the Linux kernel and different Linux distributions.

What characterizes the AWS is their specially designed Linux distribution, Amazon Linux AMI, based on Fedora and Red Hat Enterprise Linux.

For the purposes of this paper, it will be created *t1.micro* instance of VM based on Ubuntu 14.04 LTS distribution of Linux (Table 1). EC2 AMI provides a user possibility to choose the number of instances, disk space for data storage, and network. In order to facilitate the access to the machine via SSH it is necessary to create a key pair (public and private).

V. MICROSOFT AZURE AND AWS PLATFORM SERVICES

Azure and AWS are offering top public cloud solutions, but when comparing the strongest instances that are on offer, it can be seen that Amazon offers a lot more for a certain amount of money. Extra Large (A4) Azure instance has 8 virtual CPU cores and 14GB for 0.72 dollars/hour,

while Amazon whose m3.2xlarge contains 8 CPU cores, 30GB of RAM and 2x80GB SSD disk space at price of 0.56 dollars/hour. Rates on a monthly basis are calculated on the basis of 750 working hours of virtual machines [16]. For the purpose of performance testing Azure and AWS, both VMs were installed are the same operating system, Ubuntu Linux Server 14.04 LTS 64-bit.

TABLE 1: REVIEW OF VIRTUAL MACHINES ON BOTH PLATFORMS

Azure	*Amazon*
System Information **Hardware:** Processor: Intel Xeon E5-2660 0 @ 2.19GHz (1 Core) Memory: 512 MB + 256 MB Disk: 31GB Virtual Disk + 21GB Vdisk **Software:** OS: Ubuntu 14.04 Kernel: 3.13.0-27-generic (x86_64) File-System: ext4 System Layer: Microsoft Hyper-V Server Core Count: 1 Thread Count: 1 Cache Size: 20480 KB Extensions: SSE 4.2 + AVX AES Encryption: YES Disk Scheduler: DEADLINE Disk Mount Opt: ordered,discard, relatime,rw	System Information **Hardware:** Processor: Intel Xeon E5-2650 0 @ 1.80GHz (1 Core) Memory: 588MB Disk: 30GB **Software:** OS: Ubuntu 14.04 Kernel: 3.13.0-24-generic (x86_64) File-System: ext4, System Layer: Xen 3.4.3.amazon Hypervisor Core Count: 1 Thread Count: 1 Cache Size: 20480 KB Extensions: SSE 4.2 + AVX AES Encryption: YES Disk Scheduler: DEADLINE Disk MountOptions: ordered,relatime,rw

TABLE 2: APACHE BENCHMARK RESULTS

Test	Azure (req/s)	AWS (req/s)
Test 1	1116.64	1022.43
Test 2	1119.92	1026.06
Test 3	1121.67	1027.38
Average	1119.41	1025.29

Fig. 1. Apache Benchmark

Phoronix Test Suite3 was used for testing, as it allows testing of Linux platforms and system performance in given conditions. The package is easy to install from Ubuntu Repository (Table 1).

Apache, Dbench and RAMspeed benchmark test procedures (integrated into the Phoronix Test Suite3) have been successfully carried out on both platforms.

Apache tests the performance of the overall system and refers to the number of requests that server can respond when having 1,000,000 requests and 100 of them are competitive. Comparative test results can be seen in Table 2 and Figure 1.

As it can be seen, test results give slight advantage to MS Azure platform. It is obvious that apparently stronger

MS Azure brings better test results. This could be an important fact for users that intend to use VM as a web or similar server with large number of requests per time unit.

Dbench test is application developed by Samba project, a free alternative to netbench tool. This software uses the requests to file system to test disk performance. The obtained testing results can be seen in Table 3. Dbench results indicate that average test score for AWS is a kind of better than MS Azure, but their results differ only in the second decimal place, which is negligible for the users.

TABLE 3: DBENCH BENCHMARK RESULTS

Test	Azure (MB/s)	AWS (MB/s)
Test 1	12.6516	12.6483
Test 2	12.6721	12.6642
Test 3	12.6319	12.6572
Average	12.65	12.66

RAMspeed is benchmark test that is used for measuring RAM performance. It consists of five memory test configurations: Copy, Scale, Add, Triad, and Average. The testing results can be seen in Table 4 and Figures2 and 3.
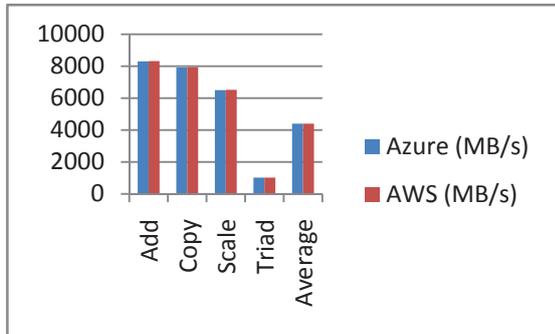
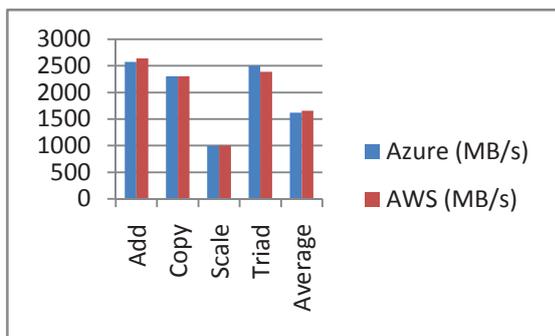

Fig. 2. RAM test – Integer

Fig. 3. RAM test – Floating point

The results of RamSpeed Benchmark tests show a little bit better performances for AWS in both analyzed variants, integer and floating point. This is consequence of a better optimisation of AWS for linux VM. This minor difference has no distinct influence on overal user expirience.

TABLE 4: RAMSPEED BENCHMARK

Test	Azure (MB/s)	AWS (MB/s)
Integer		
Add	8303.84	8318.52
Copy	7923.69	7934.63
Scale	6498.92	6522.35
Triad	1021.67	1027.38
Average	4399.17	4404.48
Float. point		
Add	2575.43	2637.03
Copy	2301.85	2306.91
Scale	998.71	1004.99
Triad	2503.42	2389.88
Average	1619.41	1657.77

VI. CONCLUSION

MS Azure have easy and intuitive user interface for managing virtual resources, but without possibility for specific VM adjustments. On the other hand, AWS offers more features for system fine tuning and gives more options oriented to managing Linux virtual machines.

When it comes to performances, the obtained results from testing micro instancesare very similar for both platforms, but test results give slight advantage to MS Azure platform when CPU and disk intensive operations are concerned.However, memory tests give one step up for AWS test system. It is understandable, considering that this MS Azureinstance of virtual hardware is basically stronger, which was confirmed by the obtained tests results. Nevertheless, benchmark tests of the more powerful virtual machine instances should be in accordance with the description of the virtual hardware that is made available to end users. AWS platform is more optimised for Linux virtual machines, because it gives more options that allow end user to adjust and this is substantiated by memory speed test results.

REFERENCES

[1] D. Armstrong, K. Djemame, "Performance Issues in Clouds: An Evaluation of Virtual Image Propagation and I/O Paravirtualization", *The Computer J.*, vol. 54, no. 6, 2011, pp. 836-849.
[2] R. Krebsa, C. Momma, Kounevb, "Metrics and techniques for quantifying performance isolation in cloud environments", Science of Computer Programming, vol. 90, 2014, pp. 116–134.
[3] Y. Jin, Y. Wen, Q. Chen, Z. Zhu, "An Empirical Investigation of the Impact of Server Virtualization on Energy Efficiency for GreenDataCenter", *The Computer J.*, vol. 56, no. 8, 2013, pp. 977-990.
[4] Z. Liu, S. Wang, Q. Sun, H. Zou and F. Yang, "Cost-Aware Cloud Service Request Scheduling for SaaS Providers", *The Computer J.*, vol. 57, no. 2, 2014, pp. 291-301.
[5] U. Tupakula, V. Varadharajan, "Trust Enhanced Security for Tenant Transactions in the Cloud Environment", *The Computer J.*, 2014, pp. 1-16.
[6] W. Zhao, P. Melliar-Smith, L. Moser, "Low Latency Fault Tolerance System", *The Computer J.*, vol. 56, 2013, pp. 716-740.
[7] T. Erl,Z. Mahmood and R. Puttini "Cloud Computing: Concepts, Technology & Architecture", Prentice Hall, JCS&T, vol. 13, no. 3, 2013, pp. 63-72
[8] Microsoft Azure at MS Azure Web site. Available: http://en.wikipedia.org/wiki/Microsoft_Azure
[9] Amazon Web Services. Available: http://aws.amazon.com
[10] Elastic Computer Cloud at AWS Amazon web site. Available: http://en.wikipedia.org/wiki/Amazon_Elastic_Compute_Cloud