2016

Performance Analysis of Public Cloud Computing Providers

Xiaolin Cheng
University of Paris-Sud, Sceaux, France, xiao-lin.cheng@u-psud.fr

Ahmed Bounfour
Universite Paris Sud (Paris XI), ahmed.bounfour@u-psud.fr

PERFORMANCE ANALYSIS OF PUBLIC CLOUD COMPUTING PROVIDERS

Research in Progress

Cheng, Xiaolin, University of Paris-Sud, Sceaux, France, xiao-lin.cheng@u-psud.fr

Bounfour, Ahmed, University of Paris-Sud, Sceaux, France, ahmed.bounfour@u-psud.fr

Abstract

The objective of this paper is to perform a comprehensive performance comparison of public cloud services for computing and to analyze the correlation between their prices and performance. Eight representative public cloud providers were divided into two groups using market share: small cloud providers and large cloud providers. Results revealed that these offered computing services vary widely in performance and price; most small cloud providers have more stable and better computing performance than large cloud providers; the performance of CPU impact price significantly.

Keywords: Public Cloud Computing, Performance Evaluation, Computing Service, Correlation

1 Introduction

Cloud computing technology is a virtual technology which distributes different services (infrastructure, platform, and software) based on different deployment models (public, private, hybrid and community). It is no longer a buzzword, it's a strategy, a business model, and a set of technologies. It has drawn significant attention from firms in recent years due to its agility, variety and ability to reduce cost. However, each company has different needs and constraints; cloud market is complex; more and more American and European companies are entering IT. These cloud computing providers offer different services which vary widely in performance and price. It is a big challenge to select appropriate cloud services which meet all the business strategies of the company.

This research in progress paper aims to provide a continuous comparison framework for public cloud services between small and large providers and a detailed analysis of the correlation between price and performance. Our research work has the following objectives:

- To compare the performance between small and large public providers
- To compare the prices of different public cloud providers
- To analyse the correlation between price and performance

The remainder of this article is organized as follows. Section 2 and Section 3 contribute to research background and literature review. Measurement methodology and selected cloud services are described in Section 4. Then, in Section 5 we focus on discussing benchmarking results and analyzing the correlation between the prices of public cloud services. Finally, we present our conclusions and introduce potential future research topics in Section 6.

2 Research background

Due to popularity of public cloud in different organizations, cloud performance evaluation is particularly important, and this evaluation can help users make right decisions. Public cloud computing is used by the general public and offer pay-as-you go charging model that enables customers to pay what they use. It is different from private cloud, internally used by some organizations. In contrast, public cloud infrastructure exists on the premises of cloud provider. The first public cloud Amazon Web Services was launched in 2006, and then more and more IT companies are riding their wave to offer a variety of public cloud computing services such as Google, Microsoft and IBM. Various public cloud providers offer different types of services with different pricing schemes raising big challenges on how to choose the best suited cloud services.

Ang Li identified common services of public cloud: elastic computing cluster, persistent storage, intra-cloud network and wide-area network (Li et al., 2010). Cluster runs application's codes using numerous virtual instances. Persistent storage is used to keep data of application and accessed through API calls. Intra-cloud network provides connection between application instances, wide-area network connects different data centers where the applications are hosted. This paper focuses on comparing the performance of elastic computing cluster between small and large public cloud providers. (Lenk et al., 2009) indicated that cloud storage is a major example of IaaS (Infrastructure as a Service). Computing service is another major example of IaaS.

3 Literature review

Simon L.Garfinkel measured the performance of Amazon's Grid Computing Services and details his experience working with these commodity computing services including analysis of Amazon's security model, implementation of the S3 client API and measurement of S3 performance from EC2 (Garfinkel, 2006). (Iosup et al., 2008) contributes to evaluate the performance of the Amazon Elastic Compute

Cloud (EC2) using micro-benchmarks, kernels, and e-Science workloads and compare the performance characteristics and pricing models of clouds with those of other scientific computing alternatives using long-term traces. (Ward, 2009) compared the performance of Amazon EC2 and Ubuntu Enterprise Cloud (UEC) using memory bandwidth, storage speed and application performance. Ward showed that for most computational tasks, UEC provides better performance than EC2, although EC2 provides the most mature IaaS cloud technology.

Yahoo! Cloud Serving Benchmark (YCSB) (Cooper et al., 2010) is a framework to benchmark cloud serving systems that provide online read/write access to data. Authors defined a set of benchmarks and presented comparison results of some widely used systems: Cassandra, HBase and PNUTS. CloudCmp (Li et al., 2010) is another framework to compare the performance and cost of cloud providers. This framework can be used to measure elastic computing, persistent storage, and networking services offered by a cloud service, however it only provides snapshot benchmarking results. Considering this research gap, we strive to compare elastic computing services and provide some more detailed continuous benchmarking results.

(Singh, 2014) emphasized that response time is a major factor that has the significant impact on cloud computing performance and it is reduced by selecting the appropriate type of broker service policies, i.e. closest data center, optimum response time and re-configure dynamically with load. Singh also indicated that response time is reaching towards constant value after 6 data centers. (Khanghahi & Ravanmehr, 2013) evaluated cloud computing performance in various scenarios considering different major factors in cloud computing performance. Their simulation and evaluation based on three categories: data centers, users and geographical region. Authors emphasized that distribution of data centers and use of the closest data center are better and more optimal than increasing its power and speed. It is also revealed that increasing the number of cloud users has increased the average response time, response time reduced drastically up to 10 data centers, so putting more than that only increases the cost.

(Iosup et al., 2011) aims to test whether the performance of clouds sufficient for Many-Task computing (MTC) based scientific computing. Authors performed an empirical evaluation of four public computing clouds using micro-benchmarks and suggested that computing performance of the tested cloud services is lower than traditional computing technologies grids and parallel infrastructures.

From literature review, it is inferred that majority of the research papers were focused on evaluating the performance of cloud providers and offered different comparison frameworks. The first worldwide public cloud service EC2 was the most popularly used service to make an analysis and response time was a major factor that contributed a lot to the performance. It is apparent that there is a need to compare the performance between small and large providers in order to help cloud users make right decisions.

4 Methodology

In this section, we strive to describe our measurement methodology and dataset. Our research work includes eight representative IaaS public cloud providers which were divided into two groups: large cloud providers and small cloud providers according to Wikibon Public Cloud Market Shares 2015 (Cloud & Shares, 2015)

4.1 Dataset and research model

Cloudscreener dataset provides information and standardized metrics related to various aspects of the performance of cloud computing technology. It provides a comprehensive set of indicators which helps to understand the variance of cloud performance. The dataset included 8 cloud providers in American and European countries for March and October of 2015. The extraction process yielded a total of 6 indicators, which described various aspects of cloud performance. Table 1 displays the selected indicators, their classification according to the framework proposed by Cloudscreener.

Service	Metric	Characteristic
Server	CPU	Events per seconds with 32 threads (numbers/s)
	Memory	RAM writing speed (MB/s)
Disk	IOPs	4k random write I/O speed (IOPs)
	Broadband	1M sequential write latency (Ko/s)
Network	Response time	Delay processing at server + Delay network(milliseconds)
	Availability	Interruption of data availability
Price	Linux/Windows	Dollars/Month

Table 1. Cloud Performance Metrices

Response time is the time taken by a cloud provider to respond to a request for cloud services, it is measured by subtracting start request from start response. Total response time is the delay of processing at server and network (Ristov, Gusev, & Kostoska, 2012)

H1: Response time is negatively related with the price of public cloud service.

IOPs is a common performance measurement used to benchmark computer disk devices. In the benchmark, this measure is computed as the average number of operations that go in and out per second obtained by using 4K random write operations and a standard block size.

H2: IOPs is positively related with the price of public cloud service

Availability is the proportion of time a system is in a functioning condition, it is measured by the ratio of a total time cloud service is capable of being used during a given interval to the length of the interval.

H3: Availability is positively related with the price of public cloud service.

CPU is measured by the average number of treated events per seconds with 32 threads, Memory is measured by the average throughput expressed in MB/s, and Broadband is measured by throughput (Ko/s) 100% 1M sequential write. And finally we should consider that whether cloud computing is cost effective before shifting to cloud computing. To analyze the correlation between price and other performance criteria is one of the important objects of our research work.

H4: CPU is positively related with the Price of public cloud service.

H5: Memory is positively related with the Price of public cloud service.

H6: Broadband is positively related with the Price of public cloud service.

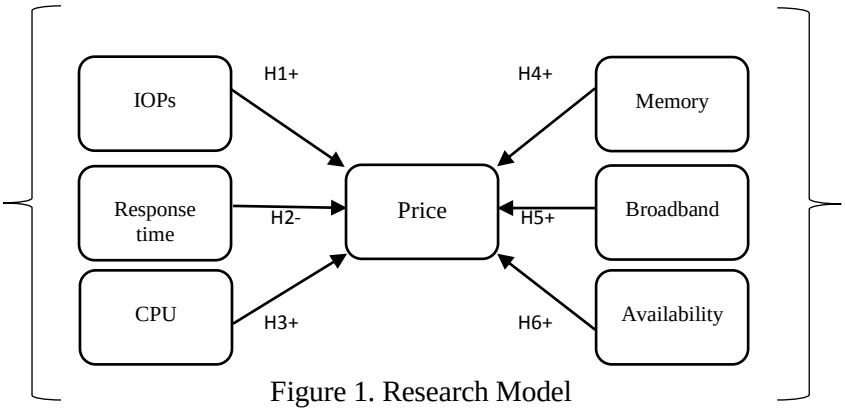

Figure 1. Research Model

Before analysing how these indicators may explain cloud performance, careful attention should be given to the different instances. In order to focus on understanding the performance variance of different cloud providers, medium instance was selected as a target.

Provider	Medium Instance	Large Instance	Extra-large Instance
AWS	m3.medium	M3.large	Extra large
Cloudwatt	n1.cw.standard-1	n1.cw.standard-2	n1.cw.standard-4
Google	n1.standard-1	n1.standard-2	n1-standard-4
Ikoula	m1.medium	Large	Extra
Windows Azure	standard A2	A3	A4
Numergy	small+	L+	XL+
Rackspace	4GB	8GB	15GB
Softlayer	Instance "Medium"	Instance "Large"	Instance "Extra"

Table 2. Type of instances

4.2 Classifying cloud providers

The second object of our paper is to analyze the public cloud service performance between small and large providers. Classifying selected cloud providers is the first step, we begin with some background and describe how to classify them by Wikibon Public Cloud Market Shares 1H 2015 (Cloud & Shares, 2015).

Provider	IaaS Market Share 1H 2015	
Amazon	3153	27.2%
Microsoft	1874	16.2%
IBM	1370	11.8%
Google Compute Engine	420	3.6%
Oracle	318	2.7%
Rackspace	282	2.4%
Other	4160	35.9%

Table 3. Wikibon cloud market shares

Figure 2 described that selected public cloud providers from CloudScreener database were divided into two groups: large providers (Google, Amazon, Microsoft, Rackspace) and small providers (Aruba, Cloudwatt, Numergy, Ikoula). Because the large cloud providers' services were popular and widely used by different types of firms, in this section, we just exhibit 4 selected small providers: Aruba, Numergy, Cloudwatt and Ikoula).

Ikoula is a French cloud provider and founded in 1998. It offers public cloud services from 2013 focusing on three different cloud services, more or less packaged. The first service Flex'Server offers dedicated virtual servers with processors, memory and different predefined storage spaces. Half of its clients are SMEs of websites or e-commerce. The second service FlexiCloud allows his clients to pick processors, memory and hard disk, in this case instances are often used for large architecture. The last one offers virtual machine at one euro, which offers the true automatic resource allocation without user validation. These virtual machines have also found an unexpected market in the financial world.

Aruba is a public cloud provider offering formally IaaS and cloud storage, it was created in 1994 in Italy. Aruba cloud would be similar with Amazon Web Services (AWS), but it is cheaper, more flexible

and better mastered. To succeed in the highly competitive French market, they decided to focus on innovation, ease of use and transparency. Aruba cloud settled especially on the reputation and strength of its parent that already has thousands of customers, and well established infrastructure. It also leverages its global strategy, in both local and global market. The implementation of Aruba in France fits into a broader strategy of extending its offer to European markets, including Germany, Spain and England. Aruba already presented in the Eastern European countries, such as the Czech Republic, Slovakia and Hungary.

Numergy and CloudWatt are two French cloud providers born from the will of the French government to establish a sovereign cloud services, they were launched in 2012. Four years later, the two firms are neck and neck. CloudWatt is managed by Orange and Thales, on the other side, Numergy is controlled by SFR and Dassault. Enjoying the data center and SFR expertise, Numergy entered to cloud market faster than CloudWatt and it offered servers, storage and network services, but there was no data centers abroad. While Numergy already had some distributors, CloudWatt chose the same indirect marketing model and hoped catch up. Compared to Numergy, CloudWatt positioned to target large organizations, public or private, with significant cloud projects, so it highlights concerns of hybrid cloud. Also, CloudWatt implemented OpenStack that introduced several differences with the strategy of Numergy. One of the main differences was that CloudWatt has not chosen the same network solutions as Numergy, however it deployed its own virtual private network infrastructure.

5 Benchmarking results

In this section, we present some preliminary benchmarking results of the common services offered by eight representative public cloud providers. The goal of cloud service benchmarking is to generate a comparison framework of performance. Our preliminary benchmarking results indicated that small cloud providers such as Ikoula, CloudWatt, Aruba and Numergy perform better than larger providers Amazon, Microsoft, Google and Rackspace in almost all the selected indicators except CPU. Also, in Table 5, we can find that the performance of public cloud services vary widely in different indicators.

The results inferred that conclude that market share is not positively related with the performance of public cloud services. It's an important point to be considered in the process of the selection of public cloud services. Considering legal concerns and keep focus on the comparison of performance for computing service, we anonymize the names of public cloud providers and refer to them as C1-C4 (large providers) and C5-C8 (Small providers). Table 4 and Table 5 display the selected indicators and their corresponding summary statistics.

Large	Availability		Broadband		CPU		IOPs		Memory		Response time	
Provider	March	October	March	October	March	October	March	October	March	October	March	October
C1	99.4	98.96	35857	28055	36	36	2721	2664	**196**	**265**	70	**67**
C2	**96.15**	99.32	73959	73976	80	81	1332	1498	1849	1876	49	47
C3	**96.91**	98.94	34501	30512	67	58	**235**	**362**	230	884	**84**	71
C4	99.45	99.31	187920	77980	**261**	**254**	1332	1301	360	**418**	53	51
Average	97.98	99.21	83059.25	52630.75	111	107,75	1405	1456.25	658.75	860.75	64	59

Small	Availability		Broadband		CPU		IOPs		Memory		Response time	
Provider	March	October	March	October	March	October	March	October	March	October	March	October
C5	99.34	99.32	79860	84514	52	51	5479	7671	2335	2224	45	46
C6	99.34	**98.73**	175088	193421	62	62	7063	1483	1389	1393	44	42
C7	99.36	99.33	170292	181007	130	**126**	6485	8988	2325	2287	45	45
C8	99.35	99.31	**692571**	**484560**	68	69	**22845**	**13806**	2707	2715	44	43
Average	99.3475	99.1725	279452,8	235875,5	78	77	10468	7987	2189	2154,75	44,5	44

Table 4. Benchmarking results

Variable	Mean	Std. Dev.	Min	Max
Price	52.11111	20.79561	27	97
IOPs	5041.167	5770.64	235	22845
Response time	52.61111	11.75805	42	84
CPU	89	66.25264	36	261
Memory	1571.389	944.8402	196	2715
Broadband	204696.1	220939	28055	785852
Availability	98.97278	.9146303	96.15	99.45

Table 5. Summary statistics

5.1 Comparison of the Service Price between Small and Large Providers

Performance and pricing are both key considerations of the public cloud services. A firm needing to use computing services must compare the alternatives of owning its computing infrastructure or leasing it from a cloud provider. Also, they should choose cost effective services with fewer resources on better performing services. In this subsection, we provide an overview of the cost items associated to Medium Instances. Table 6 indicated that small cloud providers have better performances than large providers in both of the systems (Windows and Linux) for March and October of 2015.

Service	Medium instance Linux		Medium instance Windows	
Price	March	October	March	October
C1	52	49	88	86
C2	33	27	58	53
C3	78	78	193	122
C4	95	70	121	89
Average	64.5	56	115	87.5

Service	Medium instance Linux		Medium instance Windows	
Price	March	October	March	October
C5	40	40	40	40
C6	41	36	72	55
C7	44	44	76	76
C8	40	40	56	56
Average	41.25	40	61	56.75

Table 6. Price of instance M (dollars/month)

5.2 Correlation between the Price of Public Cloud Service and Performance

Table 7 shows the correlations between identified factors. One of the important objects of our research is to find which factor influence the most the price of public cloud service. For the interpretation of this analysis, we look at the first column to identify which variable has the largest value. We found that there is a highlighted, positive correlation between price and CPU. Therefore we can conclude that CPU is the major factor impacting the price of public cloud service. Return to the hypotheses that we did, the results confirmed that H5: CPU is positively related with the price of public cloud service

Variable	Price	IOPs	Res.Time	CPU	Memory	Broadband	Availability
Price	1.0000						
IOPs	-0.2232	1.0000					
Response time	0.1902	-0.4132	1.0000				
CPU	**0.6415**	-0.1553	-0.2057	1.0000			
Memory	-0.3546	0.6658	-0.6882	-0.2805	1.0000		
Broadband	0.3103	0.5429	-0.4195	0.4624	0.1716	1.0000	
Availability	0.2333	0.3046	-0.3384	0.1101	0.1753	0.2467	1.0000

Table 7. Correlations between different indicators

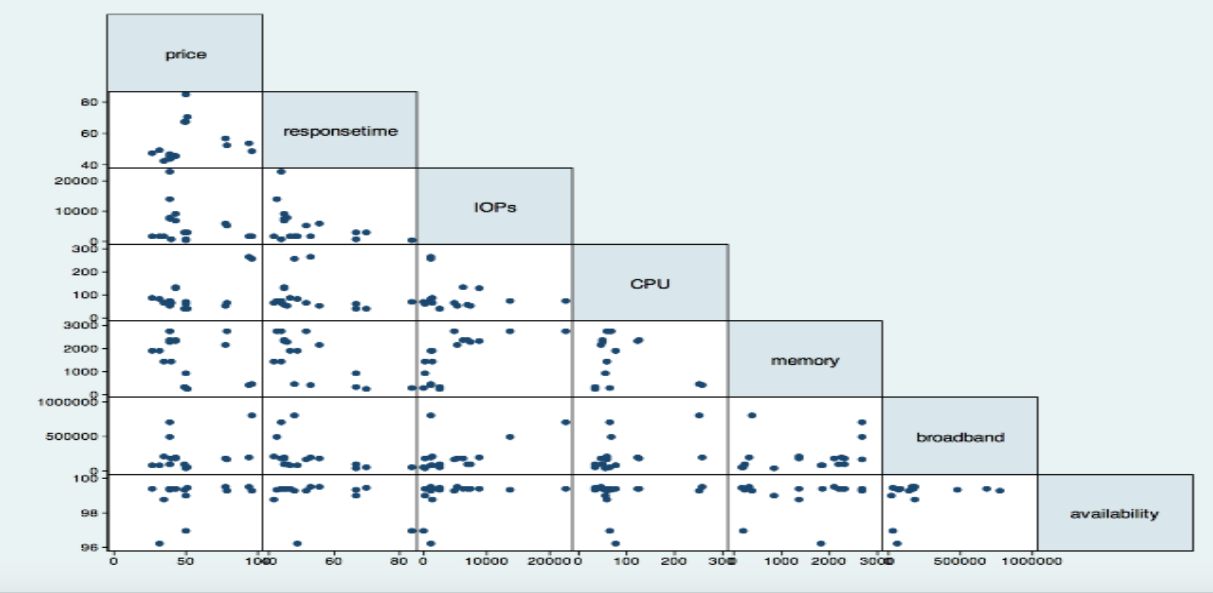

Figure 2. Correlations between different indicators

6 Conclusion and Future research

This section contributes to discuss contributions and limitations of our research work and also future research directions. This study not only examined the performance of different public cloud providers, but also tracked performance variability for two month periods. The methodology allowed us to capture performance variability over time. The current study complements previous work by analyzing the correlation between price and performance factors, comparing the performance between small public cloud providers and large provideers. From our premium results, we can find that CPU is the key factor of the performance that has significant impact on the price of public cloud services. Small cloud providers offer more stable services and pricing models than large providers. Such a systematic benchmarking research work to compare public cloud performance can make a significant impact and create healthy competition among cloud providers. We believe that our comparison framework is a significant step toward analyzing different public cloud performance.

As it stands, one of our current research limitations is that the hypothesis are not based on literature, and also, it lacks some technical depth. In our future research, we will focus on these issues, to proceed with a deep analysis statistically; to analyze more public cloud providers and offer toolboxes to evaluate applications' performance based on the results that we obtained.

References

Cloud, P., & Shares, M. (2015). Public Cloud Market Shares 2014 and 2015, (August), 1–14.

Cooper, B. F., Silberstein, A., Tam, E., Ramakrishnan, R., & Sears, R. (2010). Benchmarking cloud serving systems with YCSB. *Proceedings of the 1st ACM Symposium on Cloud Computing - SoCC '10*, 143–154.

Ec, A., Cloud, U. E., & Ward , J. S. (2009). A Performance Comparison of Clouds.

Garfinkel, S. L. (2006). Technical Report TR-08-07 : An Evaluation of Amazon ' s Grid Computing Services : EC2 , S3 and SQS. *Applied Sciences*, 1–15.

Iosup, a, Ostermann, S., Yigitbasi, M. N., Prodan, R., Fahringer, T., & Epema, D. H. J. (2011). Performance Analysis of Cloud Computing Services for Many-Tasks Scientific Computing. *IEEE Transactions on Parallel and Distributed Systems*, *22*(December), 931–945.

Iosup, A., Ostermann, S., Yigitbasi, M. N., Prodan, R., Fahringer, T., Epema, D. H. J., & Yigitbasi, N. (2011). Performance Analysis of Cloud Computing Services for Many-Tasks Scientific Computing. *IEEE Transactions on Parallel and Distributed Systems*, *22*(November), 931–945.

Khanghahi, N., & Ravanmehr, R. (2013). Cloud Computing Performance Evaluation: Issues and Challenges, *3*(5), 29–41.

Li, A., Yang, X., Kandula, S., & Zhang, M. (2010). CloudCmp: Comparing Public Cloud Providers. *Proceedings of the 10th Annual Conference on Internet Measurement - IMC '10*, 1.

Ristov, S., Gusev, M., & Kostoska, M. (2012). Cloud Computing Security in Business. *International Journal of Network Security and Its Applications*, *4*(2), 75–93.

Singh, J. (2014). Study of Response Time in Cloud Computing, (October), 36–43.