2015

An Empirical Performance Analysis Of IaaS Clouds With CloudStone Web 2.0 Benchmarking Tool

Neha Soni
University of North Florida, nehasony8@gmail.com

AN EMPIRICAL PERFORMANCE ANALYSIS OF IAAS CLOUDS WITH CLOUDSTONE WEB 2.0 BENCHMARKING TOOL

by

Neha Soni

A thesis submitted to the
School of Computing
in partial fulfillment of the requirement for the degree of

Master of Science in Computing and Information Sciences

UNIVERSITY OF NORTH FLORIDA
SCHOOL OF COMPUTING

August, 2015

The thesis "An Empirical Performance Analysis of IaaS Clouds with Cloudstone Web 2.0 Benchmarking Tool" submitted by Neha Soni in partial fulfillment of the requirements for the degree of Master of Science in Computing and Information Sciences has been

Approved by the thesis committee: Date

_____________________________________ ____________________

Dr. Sanjay P. Ahuja
Thesis Advisor and Committee Chairperson

_____________________________________ ____________________

Dr. Roger Eggen

_____________________________________ ____________________

Dr. Karthikeyan Umapathy

Accepted for the School of Computing:

_____________________________________ ____________________

Dr. Asai Asaithambi
Director of the School

Accepted for the College of Computing, Engineering, and Construction:

_____________________________________ ____________________

Dr. Mark A. Tumeo
Dean of the College

Accepted for the University:

_____________________________________ ____________________

Dr. John Kantner
Dean of the Graduate School

ACKNOWLEDGEMENT

I would like to take this opportunity to express my gratitude to the people without whose support this thesis would not have been possible. I would like to thank my thesis advisor, Dr. Sanjay P. Ahuja for his valuable guidance and support, and I would like to thank my committee members, Dr. Roger Eggen and Dr. Karthikeyan Umapathy, for their valuable suggestions. I would also like to thank Dr. Asai Asaithambi and Dr. Ching-Hua Chuan for attending my thesis defense and providing suggestions. I would also like to thank Mr. Jim Littleton for providing valuable editorial suggestions on my thesis write up. Finally, I would like to thank my husband and family for their continuous support, encouragement and love during the long process in achieving this important goal.

CONTENTS

FIGURES

TABLES

ABSTRACT

Web 2.0 applications have become ubiquitous over the past few years because they provide useful features such as a rich, responsive graphical user interface that supports interactive and dynamic content. Social networking websites, blogs, auctions, online banking, online shopping and video sharing websites are noteworthy examples of Web 2.0 applications. The market for public cloud service providers is growing rapidly, and cloud providers offer an ever-growing list of services. As a result, developers and researchers find it challenging when deciding which public cloud service to use for deploying, experimenting or testing Web 2.0 applications. This study compares the scalability and performance of a social-events calendar application on two Infrastructure as a Service (IaaS) cloud services – Amazon EC2 and HP Cloud. This study captures and compares metrics on three different instance configurations for each cloud service such as the number of concurrent users (load), as well as response time and throughput (performance). Additionally, the total price of the three different instance configurations for each cloud service is calculated and compared. This comparison of the scalability, performance and price metrics provides developers and researchers with an insight into the scalability and performance characteristics of the three instance configurations for each cloud service, which simplifies the process of determining which cloud service and instance configuration to use for deploying their Web 2.0 applications. This study uses CloudStone – an open-source, three-tier web application benchmarking tool that simulates Web 2.0 application activities – as a realistic workload generator and to capture

the intended metrics. The comparison of the collected metrics indicate that all of the tested Amazon EC2 instance configurations provide better scalability and lower latency at a lower cost than the respective HP Cloud instance configurations; however, the tested HP Cloud instance configurations provide a greater storage capacity than the Amazon EC2 instance configurations, which is an important consideration for data-intensive Web 2.0 applications.

Chapter 1

INTRODUCTION

Web applications have evolved over the past two decades from static content files to the dynamically-generated user-interactive web pages. Traditional Web 1.0 applications had several limitations including static, read-only files that only supported passive, one-way communication between a website and its clients. Web 1.0 applications also had limited scalability capabilities; therefore, as the number of clients requesting data from a website increased, the response time – the time required for the client to receive data from the website – increased; resulting in a decrease in the website's performance and causing communication delays. Conversely, Web 2.0 applications allow users to interact with the content of a web page rather than simply consuming the content. Such dynamic applications provide important features such as a rich user interface and active, two-way communication that supports collaboration amongst the application and its users. The social networking services Facebook and Twitter are two examples of popular Web 2.0 applications.

As Web 2.0 applications became an integral part of the daily activities of people throughout the world, providers migrated their applications to large-scale distributed computing platforms capable of supporting the increasing demand for online services. One such distributed platform is the Cloud, which is emerging as the dominant computing platform for Web 2.0 applications. Cloud computing provides a number of

benefits to providers, as well as their customers, that make it a better choice over other distributed computing platforms. One benefit of cloud computing is that the bulk of the data associated with an application and its customers resides in the cloud, which means customers can access their data regardless of their location or the device with which they connect. Another benefit of cloud computing is that policies can be established to maintain acceptable levels of load and latency as the demand on the system fluctuates. A final, but likely the most important, benefit of cloud computing is that it provides scalability, which allows the resources, assigned to an application to expand and contract as the number of concurrent users fluctuates.

As cloud computing has emerged as the dominant platform for Web 2.0 applications, researchers have begun to study public cloud services to help application developers choose the cloud service that best supports their applications. This study provides measurements such as the number of concurrent users (scalability), response time and throughput (performance), and total cost for various virtual machine instances on the Amazon EC2 and HP Cloud Infrastructure as a Service (IaaS) providers. The CloudStone web application benchmarking tool is used to capture and measure the scalability and performance of the tested cloud service providers and instances.

1.1 CloudStone Overview

CloudStone is an open-source, multi-platform tool – developed by the University of California, Berkeley and Sun Microsystems – for benchmarking Web 2.0 applications

operating on cloud computing platforms to generate a perception of the performance characteristics of cloud service providers [Sitaram11]. As the only modern Web 2.0 application benchmarking tool available today, CloudStone is extremely useful to researchers studying Web 2.0 applications on cloud services that provide on-demand virtual instances. CloudStone runs a Web 2.0 application called Olio that simulates a social-events application using three virtual machines – application server, database server, and a client server. CloudStone is comprised of three major components – the Olio application, a workload generator (Faban), and a set of measuring and automation tools for running large experiments on cloud computing platforms. CloudStone defines Service Level Agreements (SLAs), which are described in Chapter 4, that specify response time criteria for its operations.

1.1.1 Olio Application

Olio simulates a social-events calendar application that serves as reference architecture for testing and evaluating the characteristics of Web 2.0 applications. It supports functionality representative of Web 2.0 applications - user generated metadata, social networking functions such as posting, sharing, tagging, searching, and commenting on social events and a rich AJAX-based GUI [Beitch10]. Similar to a social media application, it also indicates the number of friendship requests a particular user has. Olio currently supports three web application framework implementations – PHP, J2EE and Ruby on Rails, and this study utilizes the PHP implementation.

The Olio workload component, which is responsible for generating a load on the Olio application, emulates a number of concurrent or active users during a test. The workload on the Olio application can be scaled up simply by increasing the number of concurrent users, which helps to identify the maximum number of concurrent users a particular cloud-based virtual machine instance supports. The maximum number of concurrent users is defined as the number of active users using the Olio application without violating the SLA set by CloudStone [Sitaram11]. Similar to any social media application, the Olio application creates 100 times concurrent users in the database that are referred to as registered users. A social media application has a large number of registered users however only few of them will be actively using the application. The Olio application responds to seven page operations that are essentially page requests that result in one or more HTTP request/response cycles [Subramanyam11]. All these operations are explained in detail in Appendix B.

1. HomePage – Landing page of the Olio application, which includes static content and thumbnail images, as well as the option to login or logout
2. Login – A registered user is randomly selected to log into the application using a valid username and password. If the selected user is already logged into the application, a log out operation is performed first.
3. TagSearch – Enables users to browse events by related tags. The results of a tag search are limited to 125 events.
4. EventDetail – Displays the details of a selected event. Events are randomly selected from the events listed on the user's home page.

5. PersonDetail – Displays the details of a selected user. Users are randomly selected from the list of registered users.
6. AddPerson – Enables a user to add a new registered user using randomly generated values. If the selected user is already logged into the application, a log out operation is performed first.
7. AddEvent – Enables a user to add a new event using randomly generated values for the Title, Summary, Description, Address, Event Date and Event Time fields.

Of these seven page operations, the HomePage, TagSearch, EventDetail and Login operations are performed more frequently than the AddEvent, AddPerson and PersonDetail operations. The frequency of each page operation performed by concurrent users is expressed as a percentage of all the page operations performed, known as Operation Mix percentage, and all of the operations equate to 100% [Subramanyam11].

1.1.2 Faban

Faban is a free, open source performance workload generator that runs on a client machine and generates a load on the application server machine by simulating a large number of concurrent users accessing the Olio application. Faban is comprised of two major components: the Faban Harness and the Faban Driver Framework. The Faban Harness acts as a container for hosting and automating the benchmark, and it provides a simple web interface to schedule, queue and compare the runs, collect statistics, and display results and graphs. The Faban Driver Framework is a high-level API-based

benchmark development framework, and a component model controls the life cycle of a benchmark run [Faban14].

1.1.3 CloudStone Architecture

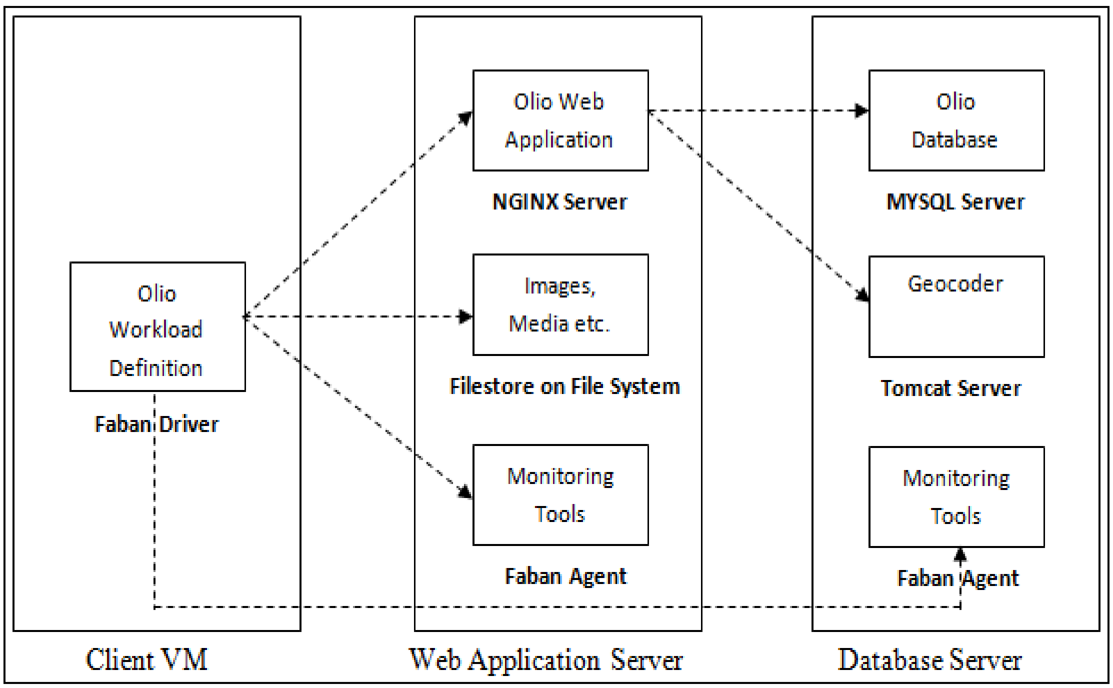

Figure 1: CloudStone architecture [Grozev14]

As shown in Figure 1, the Faban workload generator is installed and runs on the client machine, and it generates a workload by simulating a large number of users connecting to the application server machine. Faban is copied onto two other machines where it acts as an agent and monitors their performance throughout the benchmark execution process. The application server machine runs the Olio web application in the Nginx server, and it has file storage for users' images and other multimedia content. The database machine runs MySQL, and it provides access to application data. The database machine also hosts

a GeoCoder that implements a geocoding process, which helps with mapping geocoding services to the application server [Grozev14].

1.2 Cloud Architectures

"Infrastructure as a Service (IaaS) is a provision model in which an organization outsources the equipment used to support operations, including storage, hardware, servers and networking components. The service provider owns the equipment and is responsible for housing, running and maintaining it. The client typically pays on a per-use basis" [Rouse10]. Notable IaaS cloud service providers include Amazon AWS, HP Cloud, Windows Azure, Google Compute Engine, Rackspace Open Cloud, IBM SmartCloud Enterprise, AT&T and GoGrid. Although each of these providers offer cloud services to their customers, Amazon AWS is the current leader of the cloud computing market. Due to limited funding, this study conducts experiments only on two clouds – Amazon EC2 and HP Cloud.

1.2.1 Amazon EC2

Amazon Elastic Compute Cloud (EC2) is the cloud service offered by Amazon Web Services (AWS), which allows researchers and developers to pay for only the resources needed without any upfront investment. Amazon EC2 can be cost effective, and it does not require a long-term commitment from its customers [2ndwatch14]. Furthermore, it provides a simple interface through which virtual machine instances are easily added,

launched, deleted and maintained. The Amazon EC2 cloud service provides the Amazon Machine Image (AMI), which is a template of a virtual machine instance that contains an operating system, an application server and applications. An AMI must be specified prior to launching an instance of the virtual machine, but multiple instances can be launched from the same AMI. Amazon EC2 offers a variety of operating systems such as Linux, Sun Microsystems, Open Solaris, and Windows Server 2008. Amazon EC2 provides three data storage options. Amazon Elastic Block Storage acts like a persistent hard disk attached to an instance. An Instance Store is expensive, temporary storage that does not persist if an instance is terminated or stopped. Lastly, Amazon Simple Storage Service (Amazon S3) acts like a repository for Internet data that supports the storage, as well as the retrieval of data anytime and anywhere on the web [AWS15C].

Amazon EC2 offers a broad collection of instance types, which determine the hardware configuration of launched instances, and the hardware configurations are optimized for are general purpose, compute-intensive, memory-intensive, GPU-intensive, and storage-intensive operations. The hardware configurations for these instance types differ in the number of Elastic Compute Units (ECUs), memory, storage, and network performance. An ECU represents the unit amount of CPU allocated to a particular instance, and one ECU is equivalent to a CPU capacity of a 1.0-1.2 GHz 2007 Opteron® or 2007 Xeon® processor [Wikipedia15]. Table 1 lists the key features and limitations of Amazon EC2 [Lê-Quôc13].

Key features	Limitations
Free usage tier – 750 hours/month of T2.micro instance usage for first 12 months	No free credit unlike HP that can be used towards other services
Cost effective – pay for only what is consumed with affordable prices	Performance – Resources may not be running at desired performance levels due to multi-tenancy
Complete control over virtual machines	Web Console – Navigation becomes difficult when an account has more than 20 instances
Highly reliable and secured with good customer service and support	Multi-tenancy – Multiple accounts competing for same server, network and storage in over-subscription model

Table 1: Amazon EC2 key features and limitations

1.2.2 HP Cloud

HP Cloud is a public cloud infrastructure that provides cloud services to developers, software vendors, and businesses. It is built on OpenStack® technology and implies an on-demand, pay-as-you-go model [HPCloud12A]. HP Cloud provides a simple web console to manage cloud resources, and launch virtual instances effectively and efficiently. HP cloud offers four options for launching instances: boot from image, boot from snapshot, boot from volume, and boot from volume snapshot. The boot from image option allows for the launching of instances using a predefined software configuration, which include operating systems such as Windows Server 2008, CentOS, Debian, Fedora, SUSE and Ubuntu [HPcloud14B]. HP Cloud offers three data storage options that focus on performance, durability and availability: Block Storage, Object Storage and Content Delivery Network (CDN). Block storage provides persistent storage that can be attached to an instance. Object storage ensures ultra-high durability and unlimited

storage capacity, which supports the retrieval of large amounts of data immediately and securely. CDN enables access to data by storing the data on the server nearest virtual instances, which minimizes latency [HPCloud14D].

HP Cloud provides standard and high memory instance types to its customers. Each instance type varies in terms of the number of HP Compute Units, memory, and storage. An HP Compute Unit (CCU) is a unit of CPU capacity that represents the computational power of a virtual core. According to HP Cloud, 6.5 CCUs are equivalent to the minimum power of one logical core of an Intel® 2012 Xeon® 2.60 GHz CPU [HPCloud14C]. Table 2 lists the key features and limitations of HP Cloud [Sullivan14].

Key features	Limitations
Larger and powerful virtual instances	Relatively new in IaaS cloud market and limited track record
Better costs compared to Rackspace cloud	Pricing and billing higher than few other peers
Excellent 24*7 customer service and support	Launching virtual instances is little slower
Free $100 credit for first three months	Low limits on the number of virtual instances and amount of RAM

Table 2: HP Cloud key features and limitations

Chapter 2

LITERATURE REVIEW

Despite being one of the most useful and modern tools for benchmarking Web 2.0 applications, few peer-reviewed papers on CloudStone exist. Will Sobel *et al*. discusses the differences between Web 1.0 and Web 2.0 applications and workloads, as well as how CloudStone addresses Web 2.0 application requirements [Sobel08]. Furthermore, challenges related to benchmarking Web 2.0 applications such as database tuning, database performance, and server deployment are discussed. Experiments conducted on Amazon EC2 and Sun's Niagara 2 enterprise server using CloudStone showed that Amazon's EC2 had better concurrent-user support and lower response times than Sun's Niagara 2 [Sobel08].

Emmanuel Cecchet *et al*. proposed BenchLab – an open testbed that computes web application performance using existing web browsers, which is important when benchmarking Web 2.0 applications that utilize JavaScript or AJAX technologies that allow for complex interactions between the application and a web browser. As a result, BenchLab addresses the importance of measuring a Web 2.0 application's performance while emulating complex interactions that most traditional benchmarks fail to address. BenchLab uses CloudStone and Wikibooks as realistic Web 2.0 application backends and allows developers and researchers to measure the performance of their Web 2.0 applications in existing WAN environments. BenchLab focuses on three key dimensions

required in modern tools for benchmarking Web 2.0 applications – realistic server-side application, realistic workload generator, and realistic workload injector emulating the browser experience [Cecchet11].

William Voorsluys *et al.* performed experiments on the migration of virtual machines using CloudStone to calculate the reduction of responsiveness and availability experienced by applications during the migration of virtual machines. A case study, beneficial to environments in which Service Level Agreements (SLAs) drive system availability and responsiveness, was identified using Xen virtual machines running Ubuntu Linux and Olio – a Web 2.0 application – and measuring the cost of migrating virtual machines while varying the workload (number of concurrent users) on the application. The number of concurrent users the application could handle during the virtual machine migration was measured using the Service Level Agreement (SLA) metric defined in CloudStone [Voorsluys09].

Deepal Jayasinghe *et al*. performed three experiments related to the performance and scalability analysis of IaaS clouds using six clouds (three public and three private) – the private clouds were each built using a different commercial hypervisor. The public cloud experiments focused on the Emulab, Open Cirrus and Amazon EC2 platforms, and the RUBBoS benchmarking tool was used to compare the performance and scalability of the public clouds. To validate the results of the public cloud experiments, the same experiments were performed on three private clouds created using the commercial hypervisors (CVM), Xen and KVM. The results indicated that the hardware and software

configuration that performed best in Emulab was the worst-performing configuration in Amazon EC2 whose performance was limited by a combination of network sending buffers, low resource utilization and high response times. The three private clouds indicated a high variation in performance in which Xen performance was 75 percent better than CVM – using the read-write RUBBoS workload, and CVM performance was ten percent better than Xen – using the CloudStone workload [Jayasinghe14].

As a consequence of the highly-dynamic and interactive nature of Web 2.0 applications, a dynamic storage backend is necessary to support the workloads produced by these applications. To meet this need, VMware introduced Virtual SAN – a robust, distributed, and scalable virtualized storage system, which is comprised of solid-state (SSD) and traditional magnetic drives. A study on Web 2.0 applications using the CloudStone benchmarking tool and VMware's Virtual SAN storage system indicated that the Olio (Web 2.0 – social-events calendar) application performed well with the Virtual SAN storage system compared to traditional storage systems due to its low latency over time [Singaravelu14].

Although numerous studies on CloudStone exist in the literature, none have compared price, performance, scalability, and throughput of varied instances of IaaS clouds using the CloudStone benchmark. Furthermore, no study has focused on benchmarking Web 2.0 applications using the HP Cloud, which offers 'pay-as-you-go' computing with lower prices and better performance than other IaaS clouds. The purpose of this study is to provide a reliable, vendor-neutral source of information for architects, developers, and

researchers to compare the scalability, performance characteristics, and pricing models using varied instance configurations on two public IaaS cloud providers – Amazon EC2 and HP Cloud – hosting Web 2.0 applications.

Chapter 3

RESEARCH METHODOLOGY

CloudStone, an open source, three-tier web application benchmarking tool, is used to compare the performance characteristics of Amazon EC2 and HP Cloud using similar virtual instance configurations. The following research methodology is used:

1. Create and launch three t2.medium instances in Amazon EC2 running 64-bit Ubuntu 14.04 LTS – name the instances Client VM, Web Server VM and Database VM.
2. Install the CloudStone framework on each of the three virtual machines.
3. Run CloudStone and schedule a run.
4. Configure benchmark parameters such as the number of concurrent users and the addresses of the three virtual machines (an explanation of parameters is provided in Appendix A).
5. View the Summary Results to observe metrics such as response time and throughput, and view the Detailed Results to observe graphs for the metrics.
6. Increase the concurrent users parameter until the threshold value is reached (i.e. performance degrades and response time increases).
7. Record the maximum number of concurrent users, as well as the response time and throughput of the scenario.

8. Record the duration (in hours) of the experiment, and calculate the Price of the Experiment = Duration (in hours) * Hourly Price of the Instance.
9. Repeat steps 2 through 8 using three m3.xLarge instances in Amazon EC2.
10. Repeat steps 2 through 8 using three m3.2xLarge instances in Amazon EC2.
11. Repeat steps 2 through 8 using three Standard medium instances in HP Cloud.
12. Repeat steps 2 through 8 using three Standard XL instances in HP Cloud.
13. Repeat steps 2 through 8 using three Standard 2XL instances in HP Cloud.
14. Compare the performance characteristics of Amazon EC2 and HP Cloud.

Chapter 4

METRICS

The following metrics are examined when comparing Amazon EC2 to HP Cloud:

1. Load – The maximum number of concurrent users supported by the virtual instance. This is determined using the response time and throughput metrics. The load is the threshold at which throughput does not change or begins to decrease and the response time increases until it exceeds the 90^{th} percentile SLA requirement (see Table 3).
2. Response Time – The duration from when the user request is submitted and the application receives a response. CloudStone reports Response Time in terms of average (mean), maximum, standard deviation, 90^{th} percentile, required 90^{th} percentile, and pass or fail. The mean response time for each of the seven operations performed during the experiments were calculated.
3. Throughput – The number of operations carried out per second calculated as the Total Operations performed during the steady-state interval divided by the Total Seconds in the steady-state interval (a steady-state of 300 seconds was used during the experiments).
4. Price – The cost of the experiment, based on the time (in hours) required to complete the experiment, calculated using the formula:

   Price = Hourly Rate of an instance * Hours spent * Number of virtual machines

Operation	Required 90th Percentile Response Time (sec)
HomePage	1
Login	1
TagSearch	2
EventDetail	2
PersonDetail	2
AddPerson	3
AddEvent	4

Table 3: Olio operations and their required 90th percentile response time in seconds [Subramanyam11]

Chapter 5

HARDWARE AND SOFTWARE REQUIREMENTS

5.1 Hardware Requirements

The CloudStone framework requires a physical or virtual machine capable of running Linux and executing the programs bash and ssh.

5.1.1 Amazon EC2

Table 4 lists the configuration parameters for each type of Amazon EC2 instance (three virtual machine instances are required for each configuration).

	Instance Type		
	T2.Medium	M3.xLarge	M3.2xLarge
vCPUs	2	4	8
ECUs	Variable	13	26
RAM	4 GB	15 GB	30 GB
Storage	EBS	2*40 SSD (GB)	2*80 SSD (GB)
Hourly Rate	$ 0.052	$ 0.28	$ 0.56
Clock Speed	2.5 GHz	2.5 GHz	2.5 GHz

Table 4: Configuration of Amazon EC2 instances

5.1.2 HP Cloud

Table 5 lists the configuration parameters for each type of HP Cloud instance (three virtual machine instances are required for each configuration).

	Instance Type		
	Standard Medium	Standard xLarge	Standard 2xLarge
vCPUs	2	4	8
HP CUs	4	15	30
RAM	4 GB	15 GB	30 GB
Storage	50 GB	270 GB	470 GB
Hourly Rate	$ 0.12	$ 0.45	$ 0.90

Table 5: Configuration of HP Cloud instances

5.2 Software Requirements

The 64-bit version of the Ubuntu 14.04 LTS operating system, as well as the CloudStone benchmark framework, needs to be installed and configured on each of the three virtual machine instances prior to conducting each experiment (Appendix A provides detailed information on installing and running the CloudStone benchmark).

Chapter 6

RESULTS AND ANALYSIS

This study evaluates and compares the performance of Web 2.0 applications on the Amazon EC2 and HP Cloud IaaS services using the web application benchmarking tool CloudStone.. The Microsoft Excel 2010 T-TEST function is used to perform statistical analysis on the collected data to obtain p-values for the response time and throughput metrics. The T-TEST function is used to determine whether the mean values for two different data sets are statistically the same. The T-TEST function is performed using the one-tailed distribution (tails = 1) and two-sample unequal variance (type = 3) options to calculate the p-values since the data were collected independent of each other during the experiments on the two cloud services, and the data was distributed in one direction with unequal variance. A p-value of 0.05 is used to determine whether the data is statistically significant (p-value <= 0.05). IBM's SPSS tool was not used to calculate the p-values of the data because it does not support one-tailed T-TESTs.

6.1 Load

6.1.1 T2.Medium vs. Standard Medium

Table 6 and Figure 2 show how the average response time (seconds) and Table 7 and Figure 3 show how throughput (operations per second) varies as the number of concurrent users is varied on the Amazon EC2 T2.Medium and HP Cloud Standard medium instances. Table 6 and Figure 2 indicate a gradual increase in average response time on the HP Cloud Standard medium instance up to 275 concurrent users, and beyond 275 concurrent users, average response time increases exponentially; however, on the Amazon EC2 T2.Medium instance, the average response time continues to gradually increase up to 1000 concurrent users where the average response time begins to increase exponentially. Table 7 and Figure 3 indicate a linear increase in average throughput on the HP Cloud Standard medium instance up to 275 concurrent users, and beyond 275 concurrent users, average throughput begins to decrease; however, on the Amazon EC2 T2.Medium instance, the average throughput continues to linearly increase up to 1000 concurrent users where the average throughput begins to decrease. The respective exponential increases in average response time and decreases in average throughput indicates that the HP Cloud Standard medium instance supports a maximum of 275 concurrent users while the Amazon EC2 T2.Medium instance supports a maximum of 1000 concurrent users.

T2.Medium (Amazon EC2)		Standard Medium (HP Cloud)	
No. of concurrent users	Average Response Time (sec)	No. of concurrent users	Average Response Time (sec)
50	0.0315	50	0.0651
100	0.0300	100	0.0910
150	0.0288	150	0.1310
200	0.0297	200	0.2155
275	0.0299	275	0.4540
500	0.0341	285	1.9050
750	0.0472		
1000	0.1160		
1050	1.9650		

Table 6: Response Time Vs Number of concurrent users – T2.Medium Vs Standard Medium

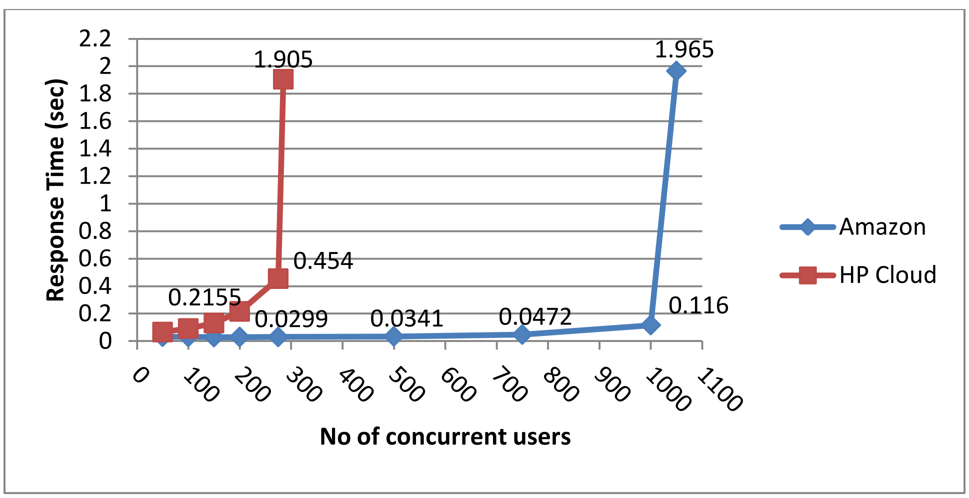

Figure 2: Response Time Vs Number of concurrent users – T2.Medium Vs Standard Medium

T2.Medium (Amazon EC2)		Standard Medium (HP Cloud)	
No. of concurrent users	Throughput (ops/sec)	No. of concurrent users	Throughput (ops/sec)
50	10.1185	50	10.2450
100	19.9350	100	20.2565
150	30.3200	150	30.0515
200	40.2150	200	40.3065
275	55.3635	275	54.6020
500	100.2700	285	43.5850
750	151.7030		
1000	201.5650		
1050	188.1450		

Table 7: Throughput Vs Number of concurrent users – T2.Medium Vs Standard Medium

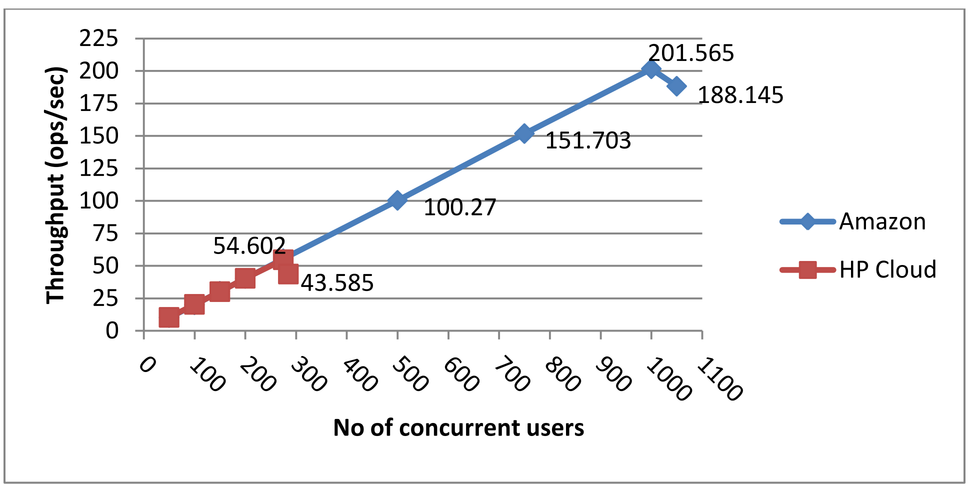

Figure 3: Throughput Vs Number of concurrent users – T2.Medium Vs Standard Medium

6.1.2 M3.xLarge vs. Standard xLarge

Table 8 and Figure 4 show how the average response time (seconds) and Table 9 and Figure 5 show how throughput (operations per second) varies as the number of concurrent users is varied on the Amazon EC2 M3.xLarge and HP Cloud Standard xLarge instances. Table 8 and Figure 4 indicate a gradual increase in average response time on the HP Cloud Standard xLarge instance up to 800 concurrent users, and beyond 800 concurrent users, average response time increases exponentially; however, on the Amazon EC2 M3.xLarge instance, the average response time continues to gradually increase up to 1300 concurrent users where the average response time begins to increase exponentially. Table 9 and Figure 5 indicate a linear increase in average throughput on the HP Cloud Standard xLarge instance up to 800 concurrent users, and beyond 800 concurrent users, average throughput begins to decrease; however, on the Amazon EC2 M3.xLarge instance, the average throughput continues to linearly increase up to 1300 concurrent users where the average throughput begins to decrease. The respective exponential increases in average response time and decreases in average throughput indicates that the HP Cloud Standard xLarge instance supports a maximum of 800 concurrent users while the Amazon EC2 M3.xLarge instance supports a maximum of 1300 concurrent users.

M3.xLarge (Amazon EC2)		Standard xLarge (HP Cloud)	
No. of concurrent users	Average Response Time (sec)	No. of concurrent users	Average Response Time (sec)
100	0.0250	100	0.0365
200	0.0260	200	0.0550
400	0.0280	400	0.0495
600	0.0300	600	0.0810
800	0.0325	800	0.3600
1000	0.0445	850	1.8855
1100	0.0595		
1200	0.0800		
1300	0.4050		
1350	3.1300		

Table 8: Response Time Vs Number of concurrent users – M3.xLarge Vs Standard xLarge

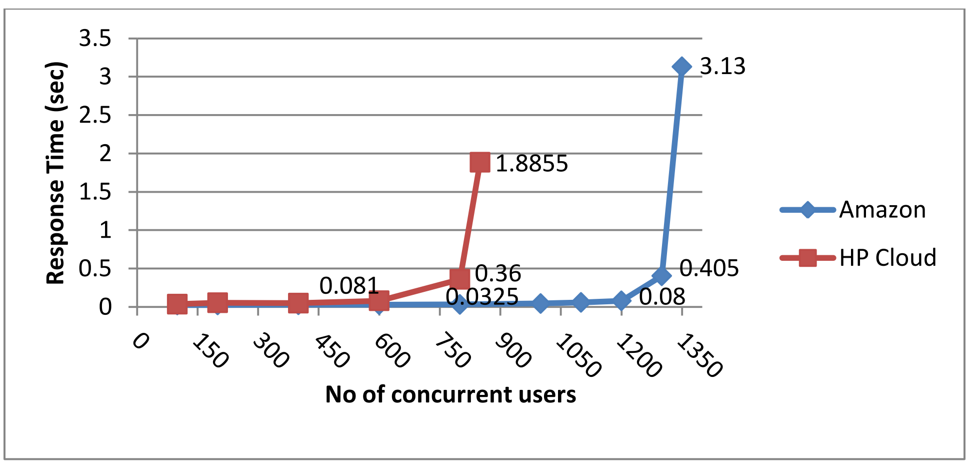

Figure 4: Response Time Vs Number of concurrent users – M3.xLarge Vs Standard xLarge

M3.xLarge (Amazon EC2)		Standard xLarge (HP Cloud)	
No. of concurrent users	Throughput (ops/sec)	No. of concurrent users	Throughput (ops/sec)
100	20.3865	100	19.9600
200	40.3570	200	40.2585
400	80.3265	400	80.7235
600	121.0600	600	120.5330
800	160.9450	800	159.7070
1000	201.4415	850	141.1315
1100	220.5400		
1200	242.2670		
1300	258.1720		
1350	193.5785		

Table 9: Throughput Vs Number of concurrent users – M3.xLarge Vs Standard xLarge

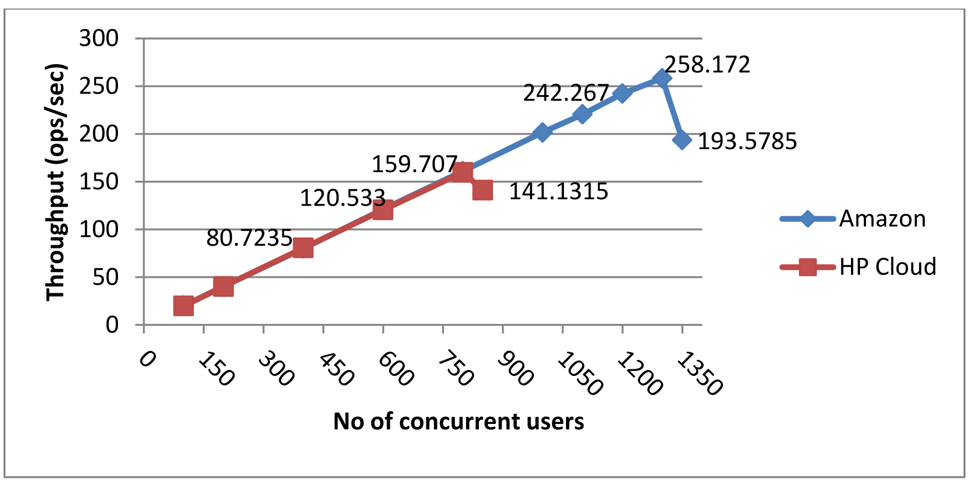

Figure 5: Throughput Vs Number of concurrent users – M3.xLarge Vs Standard xLarge

6.1.3 M3.2xLarge vs. Standard 2xLarge

Table 10 and Figure 6 show how the average response time (seconds) and Table 11 and Figure 7 show how throughput (operations per second) varies as the number of concurrent users is varied on the Amazon EC2 M3.2xLarge and HP Cloud Standard 2xLarge instances. Table 10 and Figure 6 indicate a gradual increase in average response time on the HP Cloud Standard 2xLarge instance up to 1030 concurrent users, and beyond 1030 concurrent users, average response time increases exponentially; however, on the Amazon EC2 M3.2xLarge instance, the average response time continues to gradually increase up to 1900 concurrent users where the average response time begins to increase exponentially. Table 11 and Figure 7 indicate a linear increase in average throughput on the HP Cloud Standard 2xLarge instance up to 1030 concurrent users, and beyond 1030 concurrent users, average throughput begins to decrease; however, on the Amazon EC2 M3.2xLarge instance, the average throughput continues to linearly increase up to 1900 concurrent users where the average throughput begins to decrease. The respective exponential increases in average response time and decreases in average throughput indicates that the HP Cloud Standard 2xLarge instance supports a maximum of 1030 concurrent users while the Amazon EC2 M3.2xLarge instance supports a maximum of 1900 concurrent users.

M3.2xLarge (Amazon EC2)		Standard 2xLarge (HP Cloud)	
No. of concurrent users	Average Response Time (sec)	No. of concurrent users	Average Response Time (sec)
200	0.0250	200	0.0465
400	0.0242	400	0.0635
600	0.0245	600	0.0860
800	0.0257	800	0.1470
1030	0.0270	1030	0.2740
1200	0.0305	1050	2.5895
1500	0.0375		
1800	0.0860		
1900	0.3450		
1950	3.6250		

Table 10: Response Time Vs Number of concurrent users –
M3.2xLarge Vs Standard 2xLarge

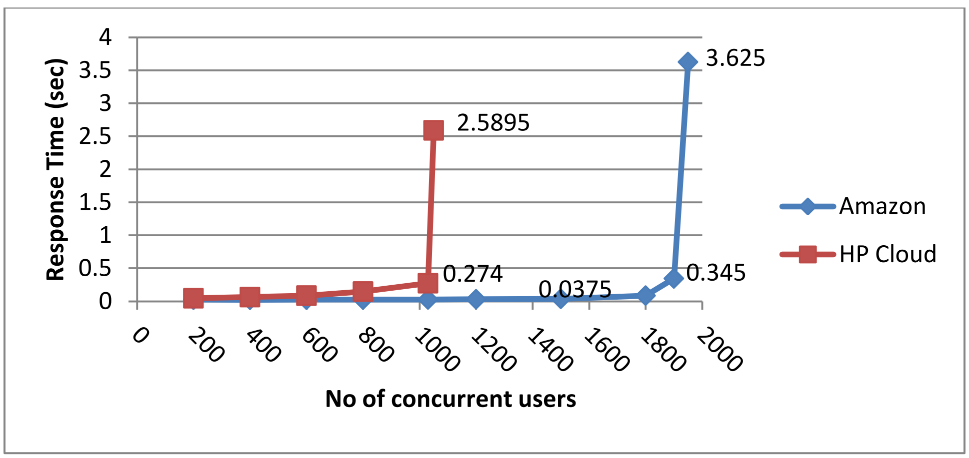

Figure 6: Response Time Vs Number of concurrent users –
M3.2xLarge Vs Standard 2xLarge

M3.2xLarge (Amazon EC2)		Standard 2xLarge (HP Cloud)	
No. of concurrent users	Throughput (ops/sec)	No. of concurrent users	Throughput (ops/sec)
200	40.1600	200	40.2550
400	80.2050	400	80.0135
600	120.7770	600	119.4100
800	160.1950	800	160.5050
1030	208.0230	1030	205.9610
1200	241.3535	1050	169.4215
1500	302.0420		
1800	361.3830		
1900	380.0715		
1950	185.0185		

Table 11: Throughput Vs Number of concurrent users – M3.2xLarge Vs Standard 2xLarge

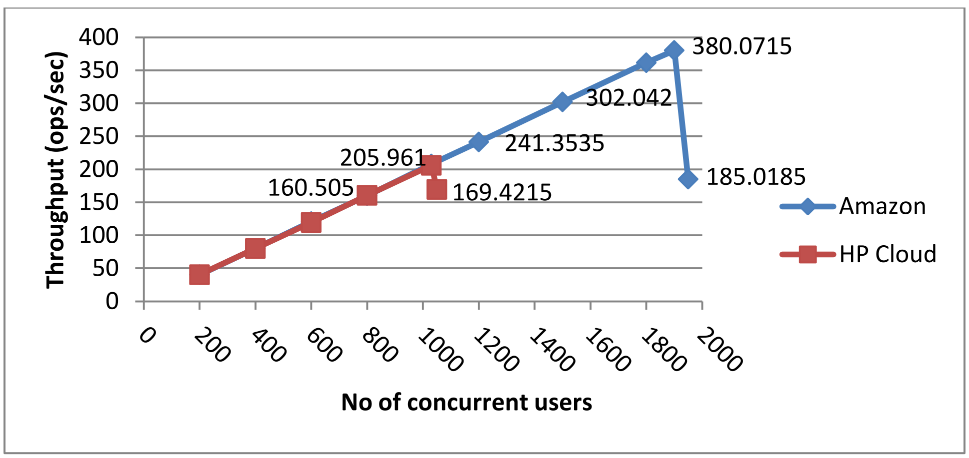

Figure 7: Throughput Vs Number of concurrent users – M3.2xLarge Vs Standard 2xLarge

Figure 8 and Table 12 summarize the number of concurrent users supported by the three respective instance types on Amazon EC2 and HP Cloud, and the data indicates that the Amazon EC2 instances are more scalable than the HP Cloud instances.

	HP Cloud		Amazon EC2	
Instance Type	No of concurrent users	Scale up Factor	No of concurrent users	Scale up Factor
Medium	275	x	1000	x
xLarge	800	2.91x	1300	1.3x
2xLarge	1030	3.74x	1900	1.9x

Table 12: Scale up factor for Amazon EC2 and HP Cloud instances

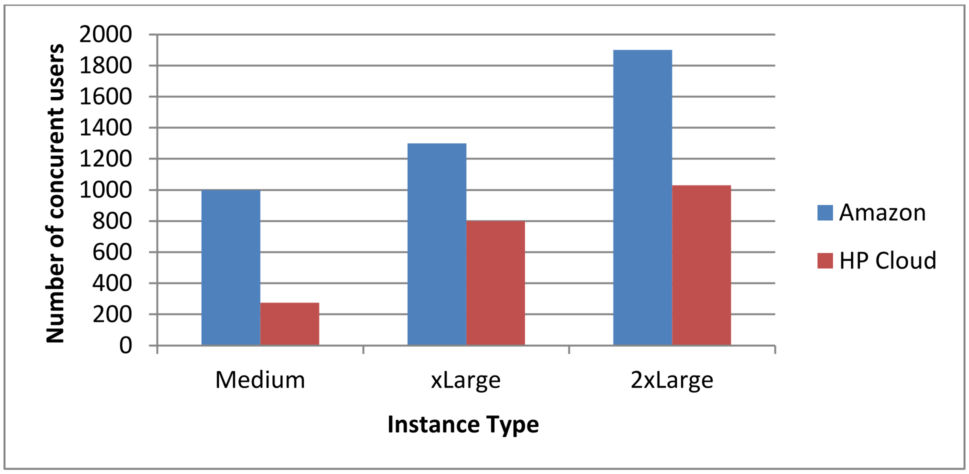

Figure 8: Instance Type Vs Number of concurrent users for Amazon EC2 and HP Cloud

6.2 Response Time

Since the HP Cloud instances tested consistently supported a smaller number of concurrent users than the Amazon EC2 instances, the average response time of each set of virtual instances were studied and compared using 275, 800 and 1030 concurrent users, respectively (T2.Medium vs. Standard medium, M3.xLarge vs. Standard xLarge and M3.2xLarge vs. Standard 2xLarge).

6.2.1 T2.Medium vs. Standard Medium

The results of the T-TEST statistical analysis performed on the average response time values collected on the medium instance of both cloud services indicates that the Amazon EC2 T2.Medium instance performed better than the HP Cloud Standard medium instance with respect to average response time (refer to Table 6 and Figure 2). The p-value of the T-TEST analysis was less than 0.05, which indicates that the average response time results are statistically significant.

6.2.2 M3.xLarge vs. Standard xLarge

The results of the T-TEST statistical analysis performed on the average response time values collected on the xLarge instance of both cloud services indicates that the Amazon EC2 M3.xLarge instance performed better than the HP Cloud Standard xLarge instance with respect to average response time (refer to Table 8 and Figure 4). It is important to note that though there is a significant difference in the average response time reported by both the clouds for 800 concurrent users, this difference is not significant up to 600 concurrent users. The p-value of the T-TEST analysis was greater than 0.05, which indicates that the average response time results are statistically insignificant.

6.2.3 M3.2xLarge vs. Standard 2xLarge

The results of the T-TEST statistical analysis performed on the average response time values collected on the 2xLarge instance of both cloud services indicates that the Amazon EC2 M3.2xLarge instance performed better than the HP Cloud Standard 2xLarge instance with respect to average response time (refer to Table 10 and Figure 6). The p-value of the T-TEST analysis was less than 0.05, which indicates that the average response time results are statistically significant.

6.3 Throughput

Since the HP Cloud instances tested consistently supported a smaller number of concurrent users than the Amazon EC2 instances, the average throughput of each set of virtual instances were studied and compared using 275, 800 and 1030 concurrent users, respectively (T2.Medium vs. Standard medium, M3.xLarge vs. Standard xLarge and M3.2xLarge vs. Standard 2xLarge).

6.3.1 T2.Medium vs. Standard Medium

The results of the T-TEST statistical analysis performed on the average throughput values collected on the medium instance of both cloud services indicates that the Amazon EC2 T2. Medium and HP Cloud Standard medium instances were statistically equal with respect to average throughput (refer to Table 7 and Figure 3). The p-value of the T-TEST

analysis was greater than 0.05, which indicates that the average throughput results are statistically insignificant.

6.3.2 M3.xLarge vs. Standard xLarge

The results of the T-TEST statistical analysis performed on the average throughput values collected on the xLarge instance of both cloud services indicates that the Amazon EC2 M3.xLarge and HP Cloud Standard xLarge instances were statistically equal with respect to average throughput (refer to Table 9 and Figure 5). The p-value of the T-TEST analysis was greater than 0.05, which indicates that the average throughput results are statistically insignificant.

6.3.3 M3.2xLarge vs. Standard 2xLarge

The results of the T-TEST statistical analysis performed on the average throughput values collected on the 2xLarge instance of both cloud services indicates that the Amazon EC2 M3.2xLarge and HP Cloud Standard 2xLarge instances were statistically equal with respect to average throughput (refer to Table 11 and Figure 7). The p-value of the T-TEST analysis was greater than 0.05, which indicates that the average throughput results are statistically insignificant.

Instance Type	Metric	p-value
T2.Medium vs. Standard Medium	Response Time	0.0419
	Throughput	0.9931
M3.xLarge vs. Standard xLarge	Response Time	0.1120
	Throughput	0.9920
M3.2xLarge vs. Standard 2xLarge	Response Time	0.0380
	Throughput	0.9880

Table 13: T-TEST Results

6.4 Price

Since the Amazon EC2 service is relatively less expensive than the HP Cloud service, with respect to the virtual instances tested (refer to the Hardware Requirements section), the three Amazon EC2 instances tested outperformed the three HP Cloud instances, respectively, with regards to Price. Table 14 and Figure 9 show the resulting price for the maximum number of concurrent users supported on each Medium, xLarge and 2xLarge instance of Amazon EC2and HP Cloud, respectively. The price for each instance is calculated using the formula:

Price = Hourly Rate of an instance * Hours spent * Number of virtual machines

In each experiment, three virtual machines were used – web application server, database server and client server. Table 14 and Figure 9 indicate that the Amazon EC2 instances are less expensive than the HP Cloud instances.

	HP Cloud		Amazon EC2	
Instance Type	No of concurrent users	Price	No of concurrent users	Price
Medium	275	$ 2.52	1000	$ 1.09
xLarge	800	$ 13.50	1300	$ 9.24
2xLarge	1030	$ 29.70	1900	$ 23.52

Table 14: Price vs. Performance in Amazon EC2 and HP Cloud instances

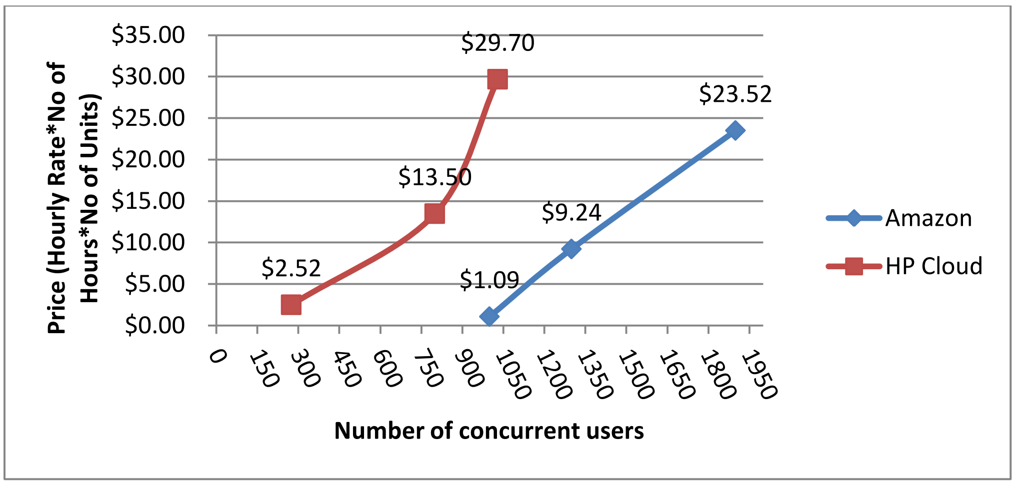

Figure 9: Price Vs Performance (number of concurrent users) in Amazon EC2 and HP Cloud

6.5 Assessment of Clouds

6.5.1 Amazon EC2

Amazon EC2 provides a very easy to use web interface for adding or removing virtual machines within seconds. Users can view their account's billing information and the detailed description of the resource usage unlike the HP Cloud. Amazon EC2 provides

powerful instances at a low cost when compared to other IaaS cloud services. Unlike HP Cloud, users do not have to activate services such as Compute, Object Storage, and Monitoring before using them. It was noticed that loading the database for more than 100,000 users in the database servers in Amazon EC2 took significantly more time in the experiments. It was also observed that ‘wa’ parameter in the vmstat output was high which was then resolved by comparing and changing the MySQL configuration parameters of my.cnf in the database server with my.cnf under Olio application. Despite the fact that Amazon EC2 offers 750 hours per month of T2.micro instances and other features to its users, it does not offer any free credit unlike HP Cloud that can be used towards other cloud services or products. The cost to setup CloudStone, get the Olio application configured and running, perform trial runs, and fix any issues was approximately $170. Although the setup costs were approximately the same for both services (with the HP Cloud $100 credit), once the cloud service environments were setup, Amazon EC2 outperforms HP Cloud with regards to the per-hour rates charged per instance. Finally, the limits placed on the number of running instances, storage and networking are higher on Amazon EC2 (compared to HP Cloud).

6.5.2 HP Cloud

Although HP Cloud is relatively new to the IaaS market, it offers an impressive assortment of products and services. HP Cloud provides some wonderful features such as simple and easy to us web interface, free trial credit, excellent 24*7 customer support, and virtual instances with greater storage capacity etc. Since HP Cloud virtual machines

have the greater storage capacity, loading the database for more than 100,000 users on HP Cloud was faster than Amazon EC2 virtual machines. However, a few issues had to be fixed while configuring the CloudStone on HP Cloud virtual machines. Before using any service, the user first needs to activate the service and manually assign a floating IP for that instance. Additionally, CloudStone cannot resolve the host names of the HP Cloud virtual machines which can be fixed by resolving the host name with machine's private IP address in the hosts file under the /etc folder in all the HP Cloud virtual machines. The cost to setup CloudStone, get the Olio application configured and running, perform trial runs, and fix any issues was approximately $280, but the $100 credit reduced the setup cost to approximately $180. Finally, the limits placed on the number of running instances, storage and networking are lower on HP Cloud (compared to Amazon EC2), but these limits can be increased; however, doing so requires the user to contact customer support.

Chapter 7

CONCLUSIONS

Chapter 6 presented the results of experiments conducted using the CloudStone benchmarking tool on three virtual machine instances of Amazon EC2 (T2.Medium, M3.xLarge and M3.2xLarge) and three virtual machine instances of HP Cloud (Standard medium, Standard xLarge and Standard 2xLarge), and the maximum number of concurrent users supported by each virtual machine instance type, as well as the average response time and throughput, was studied and compared. Furthermore, statistical analysis of the results was performed using Microsoft Excel's T-TEST function, and the p-values for average response time and throughput were examined. Figures 2 and 3 indicate that the Amazon EC2 T2.Medium virtual machine instance is more scalable than the HP Cloud Standard medium virtual machine instance. Although the vCPU (2), RAM (4 GB), and storage (50 GB) configuration was identical for both instance types, the computational power associated with each instance type was different. The Amazon EC2 T2.medium instance provides more computational power than the HP Cloud Standard medium instance. The Amazon EC2 T2.medium instance provides two vCPUs, each with a clock speed of 2.5 GHz, which provides a total of 5 GHz of computational power. Furthermore, the T2.medium instance is bursty, which means that the instance can automatically and transparently scale up to another full core when additional computational power is needed [AWS14B]. The HP Cloud Standard medium instance provides four HP Compute Units (6.5 HP Compute Units is equivalent to one logical core

of an Intel® 2012 Xeon® 2.60 GHz CPU), which provides a total of 1.6 GHz of computational power. Consequently, the average response times reported by the CloudStone benchmarking tool were significantly higher for the HP Cloud Standard medium instance than those of the Amazon EC2 T2.Medium instance. Furthermore, the throughput results reported by the CloudStone benchmarking tool were equivalent for the two instances up to 275 concurrent users, but beyond 275 concurrent users, the HP Cloud Standard medium instance became saturated (over utilized).

Figures 4 and 5 indicate that the Amazon EC2 M3.xLarge virtual machine instance is more scalable than the HP Cloud Standard xLarge virtual machine instance. Although the vCPU (4) and RAM (15 GB) configuration was identical for both instance types, the computational power and storage associated with each instance type was different. The amount of storage used for the Amazon EC2 M3.xLarge instance (50 GB) was approximately one-fifth that of the storage used for the HP Cloud Standard xLarge instance (270 GB); however, this lack of storage is insignificant when compared to the difference in computational power because the Amazon EC2 M3.xLarge instance provides more computational power than the HP Cloud Standard xLarge instance. The Amazon EC M3.xLarge instance provides four vCPUs, each with a clock speed of 2.5 GHz, which provides a total of 10 GHz of computational power. The HP Cloud Standard xLarge instance provides 15 HP Compute Units (6.5 HP Compute Units is equivalent to one logical core of an Intel® 2012 Xeon® 2.60 GHz CPU), which provides a total of 6 GHz of computational power. Consequently, the average response times reported by the CloudStone benchmarking tool were significantly higher for the HP Cloud Standard

xLarge instance than those of the Amazon EC2 M3.xLarge instance. Furthermore, the throughput results reported by the CloudStone benchmarking tool were equivalent for the two instances up to 800 concurrent users, but beyond 800 concurrent users, the HP Cloud Standard xLarge instance became saturated (over utilized).

Figures 6 and 7 indicate that the Amazon EC2 M3.2xLarge virtual machine instance is more scalable than the HP Cloud Standard 2xLarge virtual machine instance. Although the vCPU (8) and RAM (30 GB) configuration was identical for both instance types, the computational power and storage associated with each instance type was different. The amount of storage used for the Amazon EC2 M3.2xLarge instance (50 GB) was approximately one-tenth that of the storage used for the HP Cloud Standard 2xLarge instance (470 GB); however, this lack of storage is insignificant when compared to the difference in computational power because the Amazon EC2 M3.2xLarge instance provides more computational power than the HP Cloud Standard 2xLarge instance. The Amazon EC M3.2xLarge instance provides eight vCPUs, each with a clock speed of 2.5 GHz, which provides a total of 20 GHz of computational power. The HP Cloud Standard 2xLarge instance provides 30 HP Compute Units (6.5 HP Compute Units is equivalent to one logical core of an Intel® 2012 Xeon® 2.60 GHz CPU), which provides approximately 12 GHz (11.5 GHz) of computational power. Consequently, the average response times reported by the CloudStone benchmarking tool were significantly higher for the HP Cloud Standard 2xLarge instance than those of the Amazon EC2 M3.2xLarge instance. Furthermore, the throughput results reported by the CloudStone benchmarking tool were

equivalent for the two instances up to 1030 concurrent users, but beyond 1030 concurrent users, the HP Cloud Standard 2xLarge instance became saturated (over utilized). According to the data available, each of the Amazon EC2 instances (T2.Medium, M3.xLarge and M3.2xLarge) are less expensive (Price per Hour) and more scalable than each of the respective HP Cloud instances (Standard medium, Standard xLarge and Standard 2xLarge); therefore, each of the Amazon EC2 instances used in the experiments outperformed their respective HP Cloud instance counterpart in terms of price.

A review of all the test results collected for the various metrics (scalability, response time, throughput and price), the results indicate that it is cheaper to deploy Web 2.0 applications on Amazon EC2 instances rather than on HP Cloud instances because Amazon EC2 instances provide greater computational power scalability and less latency. If storage capacity is a significant factor, however, it may be cheaper to deploy Web 2.0 applications on HP Cloud instances rather than on Amazon EC2 instances because the default storage capacity for the HP Cloud Standard xLarge and Standard 2xLarge instances are much greater than the Amazon EC2 M3.xLarge and M3.2xLarge instances (5x and 10x, respectively), but this is not true for the medium instances of the Amazon EC2 (T2.Medium) and HP Cloud (Standard medium) since the storage capacity of these instances are identical.

The results of this study will be helpful to researchers and developers planning to deploy Web 2.0 applications on one of the reviewed cloud services. Furthermore, the results of this study will help researchers and developers choose the best compromise between

metrics such as scalability, response time, throughput, and price when deploying Web 2.0 applications on one of the reviewed cloud services.

Chapter 8

FUTURE WORK

This study is limited to evaluating the performance of Web 2.0 applications on the Amazon EC2 and HP Cloud services using the CloudStone benchmarking tool with respect to scalability, response time, throughput and price; however, since little research focused on evaluating the performance of Web 2.0 applications using the CloudStone benchmarking tool on public and private clouds exists, the study can serve as a reference to future studies related to Web 2.0 application performance on other IaaS clouds.

One such study might evaluate the performance of Web 2.0 applications in cloud environments where additional virtual machine instances are used to balance the load as the number of concurrent users increases. Another study might evaluate the performance of Web 2.0 applications on other IaaS cloud services such as Rackspace, IBM SmartCloud, and Google Compute Engine. Another study might evaluate the performance of Web 2.0 applications on PaaS cloud services such as Google App Engine, Microsoft Azure, and Salesforce.com. Finally, a study might explore whether the Ruby on Rails or J2EE implementations of the Olio application produce a set of results that are different from the PHP implementation of the application used in this study.

REFERENCES

Print Publications:

[Cecchet11]
Emmanuel Cecchet, Veena Udayabhanu, Timothy Wood, Prashant Shenoy. "BenchLab: An Open Testbed for Realistic Benchmarking ." *WebApps '11: 2nd USENIX Conference on Web Application Development* (USENIX Association), 2011: 37-48.

[Jayasinghe14]
Deepal Jayasinghe, Simon Malkowski, Jack Li, Qingyang Wang, Zhikui Wang, and Calton Pu. "Variations in Performance and Scalability: An Experimental Study in IaaS Clouds Using Multi-Tier Workloads." *IEEE TRANSACTIONS ON SERVICES COMPUTING* 2 (April 2014): 293-306.

[Pallis10]
Pallis, George. "Cloud Computing." *IEEE Internet Computing* (IEEE Computer Society), September 2010: 70-73.

[Voorsluys09]
William Voorsluys, James Broberg, Srikumar Venugopal, Rajkumar Buyya. *Cost of Virtual Machine Live Migration in Clouds: A Performance Evaluation.* Vol. 5931, in *Cloud Computing*, by Gansen Zhao, Chunming Rong Martin Gilje Jaatun, 254-265. Beijing: Springer Berlin Heidelberg, 2009.

Electronic Sources:

[2ndwatch14]
2ndwatch. "Benefits of Amazon Web Services (AWS)." *2ndwatch.* January 28, 2014. http://2ndwatch.com/benefits-of-amazon-web-services-aws/ (accessed NOvember 4, 2014).

[AWS14A]
services, amazon web. "Amazon EC2." *aws.* November 4, 2014. http://aws.amazon.com/ec2/ (accessed November 4, 2014).

[AWS14B]
services, amazon web. "New Low Cost EC2 Instances with Burstable Performance." *aws.* July 1, 2014. https://aws.amazon.com/blogs/aws/low-cost-burstable-ec2-instances/ (accessed April 8, 2015).

[AWS15C]

"Products & Services." *aws.amazon.com.* April 8, 2015.
http://aws.amazon.com/products/?nc2=h_ql_sf_ls (accessed April 8, 2015).

[Beitch10]
Aaron Beitch, Brandon Liu, Timothy Yung, Rean Griffith, Armando Fox, David A. Patterson. *Rain: A Workload Generation Toolkit for Cloud Computing Applications.* Berkeley: Electrical Engineering and Computer Sciences University of California at Berkeley, 2010.

[Cormode08]
Graham Cormode, Balachander Krishnamurthy. "Key differences between Web 1.0 and Web 2.0." *First Monday* 13 (2008).

[Faban14]
Faban. "Faban Driver Framework Developers Guide - Introduction." *faban.org.* October 25, 2014. http://faban.org/1.1/docs/guide/driver/intro.html (accessed November 4, 2014).

[Grozev14]
Grozev, Nikolay. *Automated CloudStone Setup in Ubuntu VMs.* May 10, 2014. http://nikolaygrozev.wordpress.com/2014/05/10/automated-cloudstone-setup-in-ubuntu-vms/ (accessed November 2, 2014).

[HPCloud12A]
"HP Cloud Services." *hpcloud.* April 30, 2012. https://www.hpcloud.com/sites/default/files/HPCS%20Overview.PDF (accessed November 4, 2014).

[HPcloud14B]
Cloud, HP. "Getting Started with HP Helion Public Cloud Console." *HP Cloud.* December 31, 2014. http://docs.hpcloud.com/publiccloud/hpcloudconsole/ (accessed April 8, 2015).

[HPCloud14C]
"HP Cloud Pricing." *HP Cloud.* December 31, 2014. http://www.hpcloud.com/pricing (accessed April 8, 2015).

[HPCloud14D]
"HP Cloud Storage." *HP Cloud.* December 31, 2014. http://www.hpcloud.com/products-services/storage-cdn (accessed April 8, 2015).

[Le-Quoc13]

Alexis Lê-Quôc, Mike Fiedler,Carlo Cabanilla. "The Top 5 AWS EC2 Performance Problems." *datadoghq.* July 31, 2013. http://www.datadoghq.com/wp-content/uploads/2013/07/top_5_aws_ec2_performance_problems_ebook.pdf (accessed November 4, 2014).

[Rouse10]
Rouse, Margaret. "Infrastructure as a Service (IaaS)." *SearchCloudComputing.* August 31, 2010. http://searchcloudcomputing.techtarget.com/definition/Infrastructure-as-a-Service-IaaS (accessed November 4, 2014).

[Singaravelu14]
Lenin Singaravelu, Zach Shen. *Web 2.0 Applications on VMware Virtual SAN.* September 4, 2014. Web 2.0 Applications on VMware Virtual SAN (accessed November 4, 2014).

[Sitaram11]
Dinkar Sitaram, Geetha Manjunath. *Moving to the Cloud: Developing Apps in the New World of Cloud Computing.* Syngress, 2011.

[Sobel08]
Will Sobel, Shanti Subramanyam, Akara Sucharitakul, Jimmy Nguyen, Hubert Wong, Sheetal Patil, Armando Fox, David Patterson. *Cloudstone: Multi-Platform, Multi-Language Benchmark and.* October 08, 2008. http://cyberaide.googlecode.com/svn/trunk/misc/cloud-papers/cca08/33.pdf (accessed November 3, 2014).

[Subramanyam11]
Subramanyam, Shanti. "Olio Workload Description." *github.* July 25, 2011. https://github.com/shanti/olio/wiki/Olio-Workload-Description (accessed April 9, 2015).

[Sullivan14]
Sullivan, Dan. "IaaS Providers List: 2014 Comparison And Guide." *tomsitpro.* February 14, 2014. http://www.tomsitpro.com/articles/iaas-providers,1-1560.html (accessed November 4, 2014).

[Wikipedia15]
Wikipedia. "Amazon Elastic Compute Cloud." *wikipedia.* April 6, 2015. http://en.wikipedia.org/wiki/Amazon_Elastic_Compute_Cloud (accessed April 8, 2015).

Appendix A

CLOUDSTONE SETUP

Launch virtual machines on Amazon EC2 or HP Cloud. CloudStone implicitly requires that all the machines can talk to and ping each other at random ports, due to which we need to configure the security group while launching an instance as below:

Inbound rules for sg-e63db383 (Selected security groups: sg-e63db383)

Type	Protocol	Port Range	Source
All TCP	TCP	0 - 65535	0.0.0.0/0
Custom ICMP Rule	Echo Request	N/A	0.0.0.0/0

Figure 1: Security Group configurations

After installing the CloudStone following installation steps available at http://parsa.epfl.ch/cloudsuite/web.html, point browser to http://[client machine address]:9980 in order to run the CloudStone benchmark. The following screen gets displayed:

Figure 2: Cloudstone benchmark Home Screen

To schedule a run, click ‘Schedule Run’, enter a profile name and click ‘Select’. Configure the CloudStone setup for ‘Java’, ‘Driver’, ‘Web Server’ and ‘Data Servers’ tabs as below.

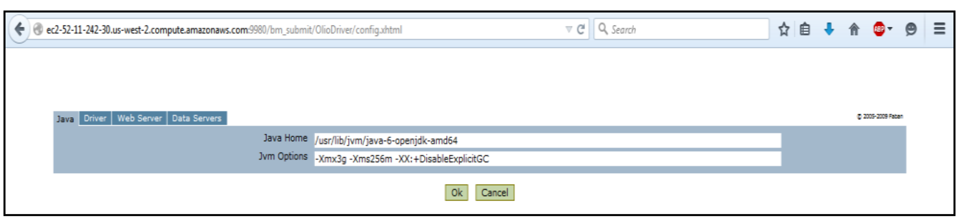

Figure 3: Cloudstone benchmark configuration - Java tab

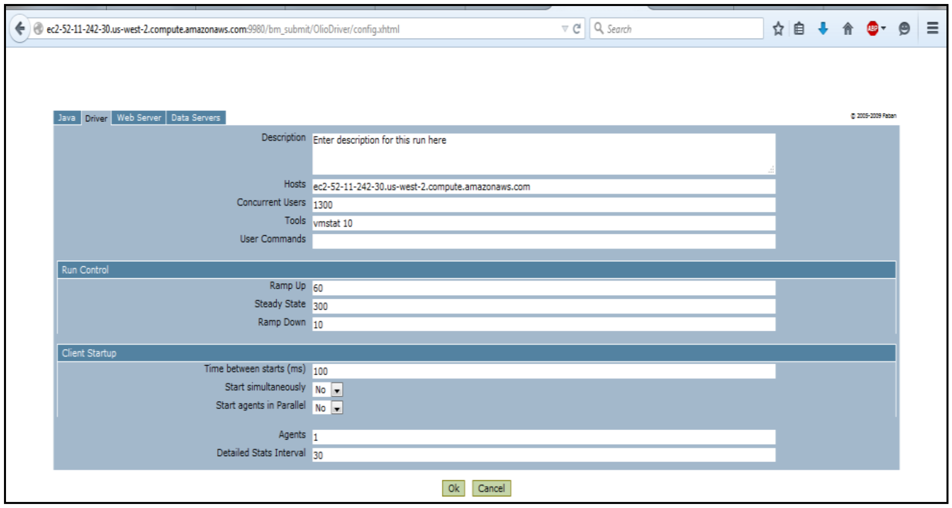

Figure 4: Cloudstone benchmark configuration - Driver tab

As shown in Figure 4, 'Hosts' has the address of the client machine.

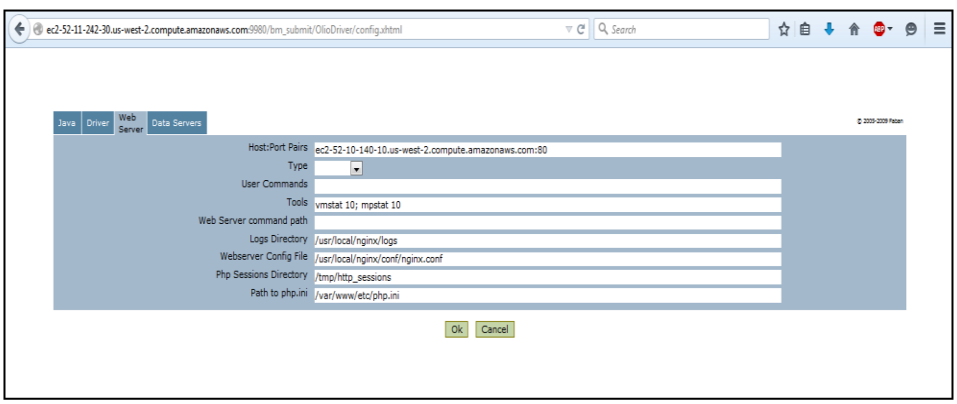

Figure 5: Cloudstone benchmark configuration – Web Server tab

As shown in Figure 5, 'Host:Port Pairs' has the address of the web application server machine.

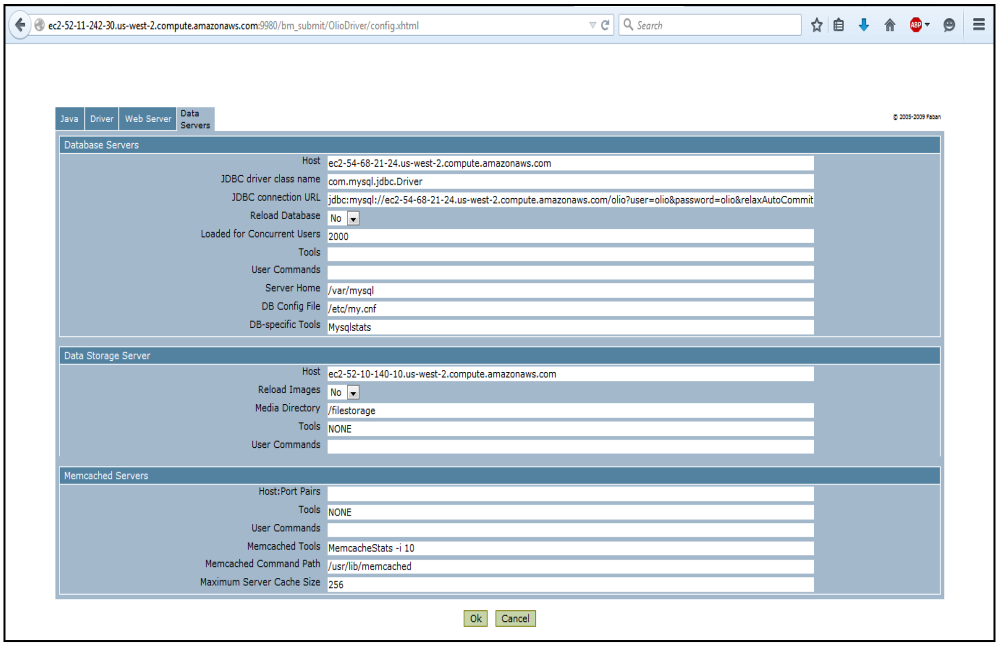

Figure 6: Cloudstone benchmark configuration – Data Servers tab

As shown in Figure 6, 'Host' under 'Database Servers' has the address of the database server machine and 'JDBC connection URL' points to the MySQL of the database server machine with correct username and password to connect to the database. 'Host' under 'Data Storage Server' has the address of the web application server machine. Once the benchmark is configured as mentioned in the screens above, the benchmark progress and the results can be viewed using 'View Results' link as shown in Figure 7.

Figure 7: Cloudstone benchmark configuration – View Results

To view the summary result, detailed results and run log click on the RunID (Refer Figures 8 to 14).

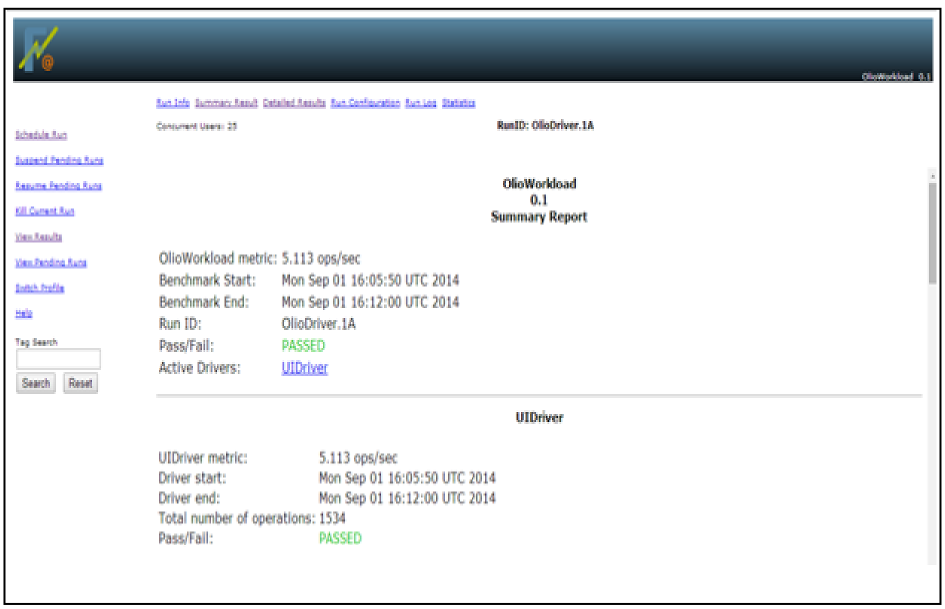

Figure 8: Cloudstone benchmark Summary Result - UIDriver

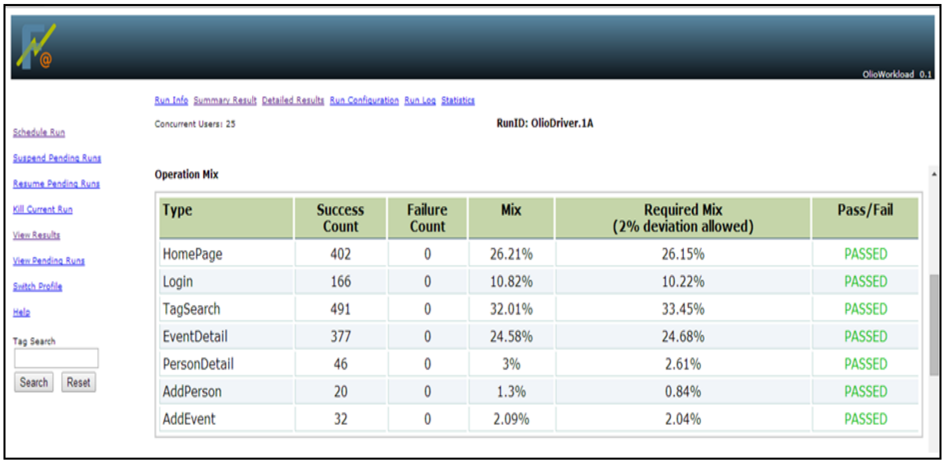

Type	Success Count	Failure Count	Mix	Required Mix (2% deviation allowed)	Pass/Fail
HomePage	402	0	26.21%	26.15%	PASSED
Login	166	0	10.82%	10.22%	PASSED
TagSearch	491	0	32.01%	33.45%	PASSED
EventDetail	377	0	24.58%	24.68%	PASSED
PersonDetail	46	0	3%	2.61%	PASSED
AddPerson	20	0	1.3%	0.84%	PASSED
AddEvent	32	0	2.09%	2.04%	PASSED

Figure 9: Cloudstone benchmark Summary Result – Operation Mix

Type	Avg	Max	SD	90th%	Reqd. 90th%	Pass/Fail
HomePage	0.015	0.040	0.007	0.040	1.000	PASSED
Login	0.011	0.016	0.001	0.020	1.000	PASSED
TagSearch	0.030	0.051	0.010	0.060	2.000	PASSED
EventDetail	0.015	0.045	0.005	0.040	2.000	PASSED
PersonDetail	0.020	0.032	0.006	0.040	2.000	PASSED
AddPerson	0.065	0.073	0.004	0.080	3.000	PASSED
AddEvent	0.085	0.137	0.016	0.100	4.000	PASSED

Figure 10: Cloudstone benchmark Summary Result – Response Times (seconds)

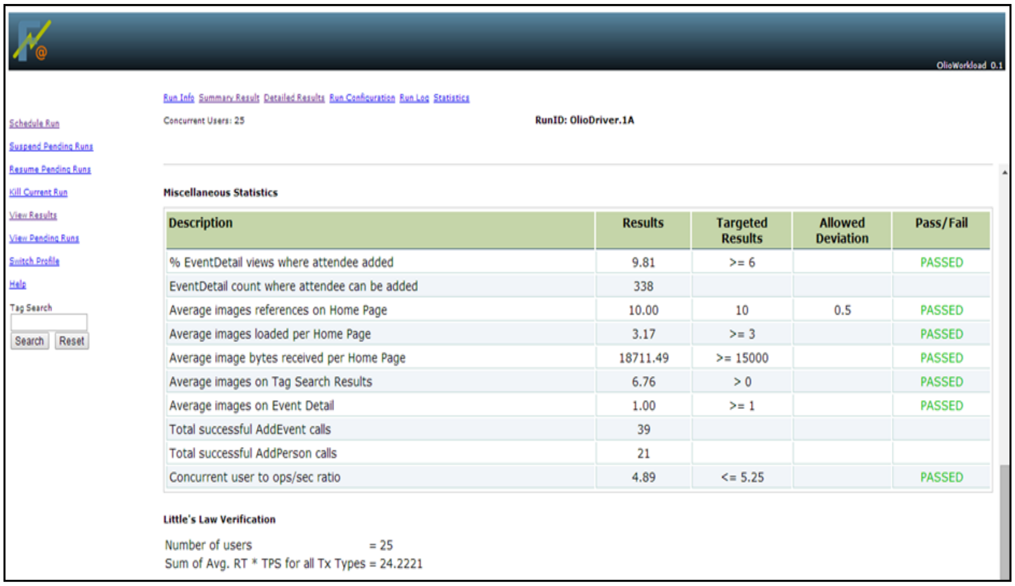

Figure 11: Cloudstone benchmark Summary Result – Miscellaneous Statistics

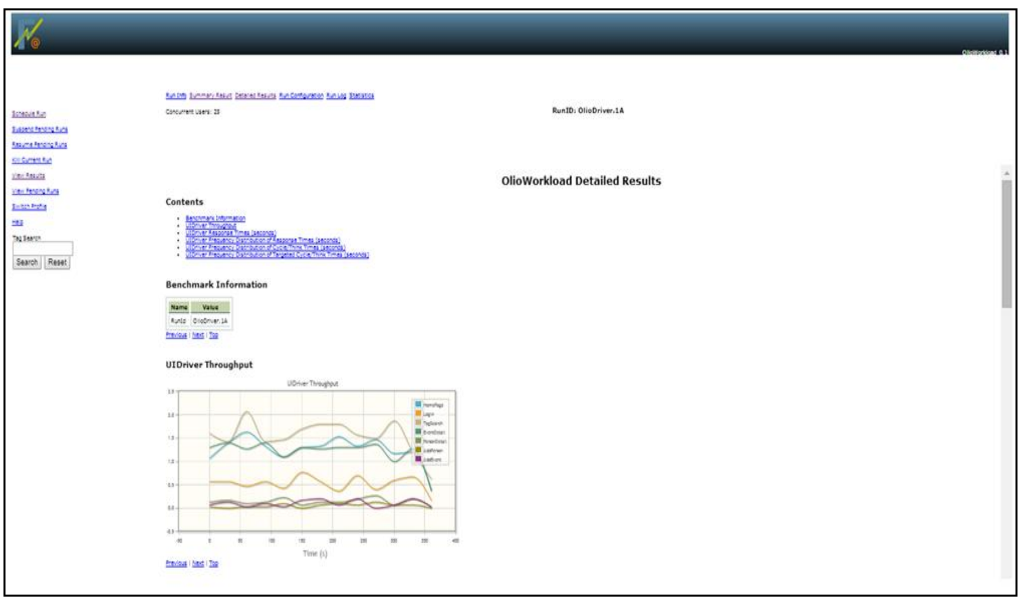

Figure 12: Cloudstone benchmark Detailed Results

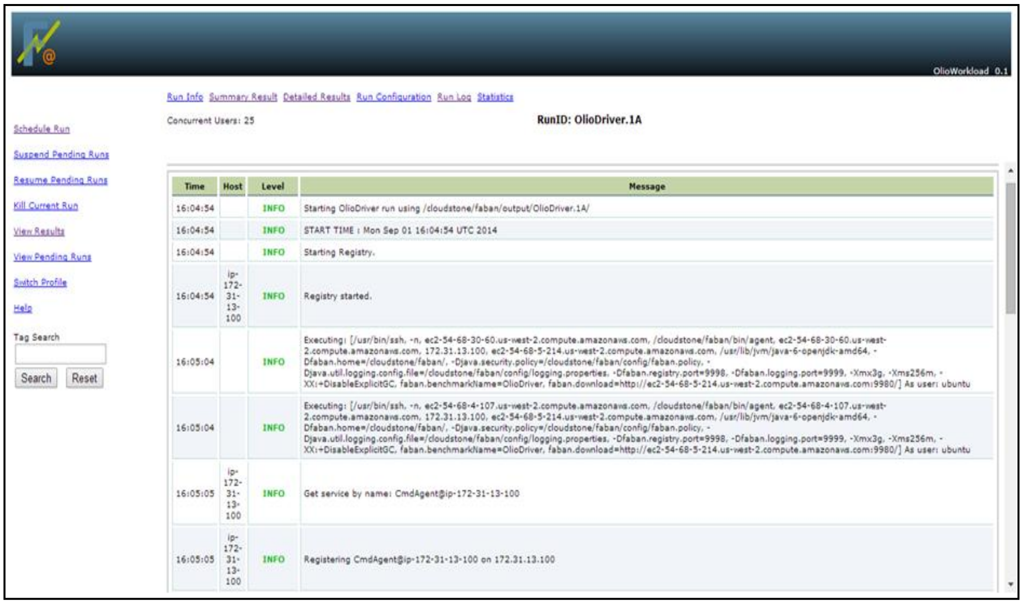

Figure 13: Cloudstone benchmark - Run Log

To compare two or more runs, select runs and click ‘Compare’.

Figure 14: Cloudstone benchmark – Compare Runs

Appendix B

OPERATIONS IN THE OLIO APPLICATION

The Olio application is comprised of seven operations as described earlier in the thesis. The HomePage displays all the available events with the thumbnails and a link to navigate to their details (Refer Figure 1). These events can be filtered according to their zip code, created date and event date as shown in Figure 1. The user can log in to the application using a valid username and password pair. Depending on the login/logout status of the user, an appropriate message is displayed on the HomePage (Refer Figures 1 and 2). Once the user is logged in to the application, he can add an event, search for users, update or reset his own profile. In addition, links to logout, friendship requests and upcoming events are displayed on the HomePage as shown in the Figure 2. The user can navigate to an upcoming event and view its details along with the number and details of the attendees for that event (Refer Figure 3). To update or reset all the details such as username, password, email, telephone, image and address, the user can use the Edit Profile tab as shown in Figure 4. The user can search for users using the Users tab available on the top and browse for events using a particular tag using Search Tags available on the left (Refer Figure 5). The user can use the Add Event tab available on the top to add a new event with the values for Title, Summary, Description, Address, Event Date and Time for an event chosen at random (Refer Figure 6). Once the user opts to logout of the application using logout link, the user is navigated back to the HomePage (Refer Figure 7).

Figure 1: Olio Application – Home Page

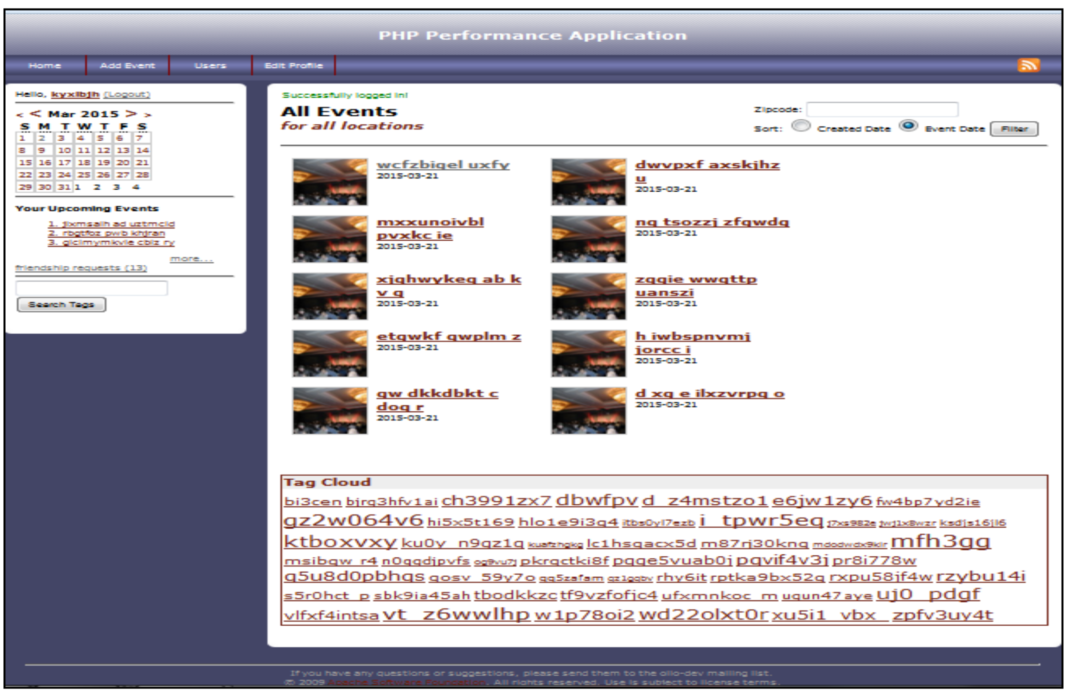

Figure 2: Olio Application – Successful Login

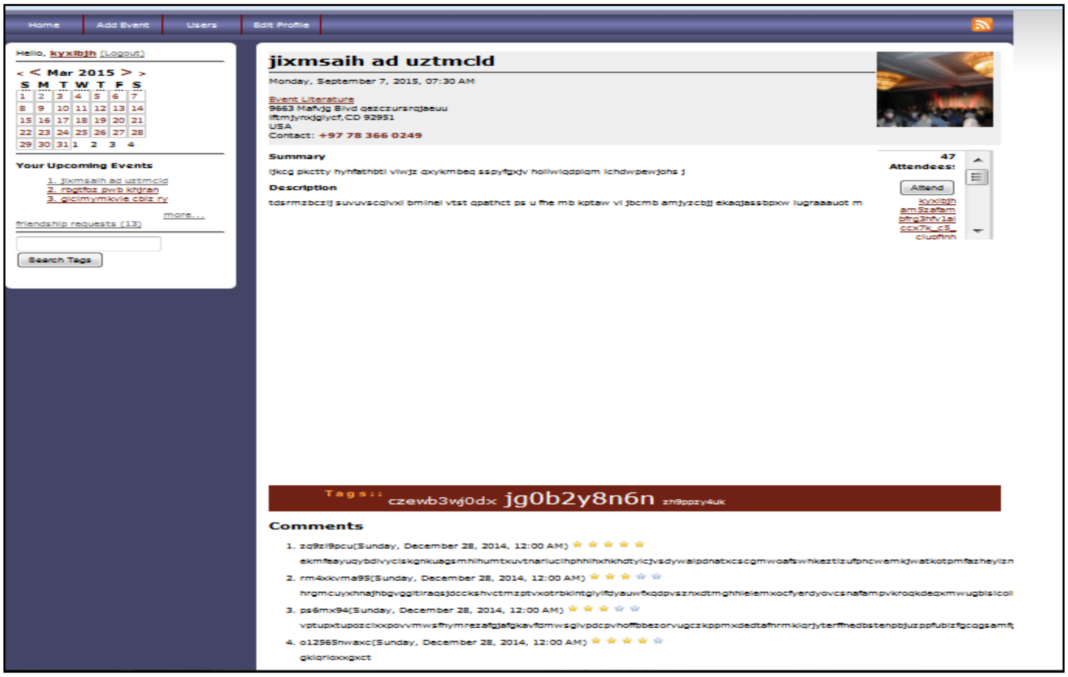

Figure 3: Olio Application – Event Details

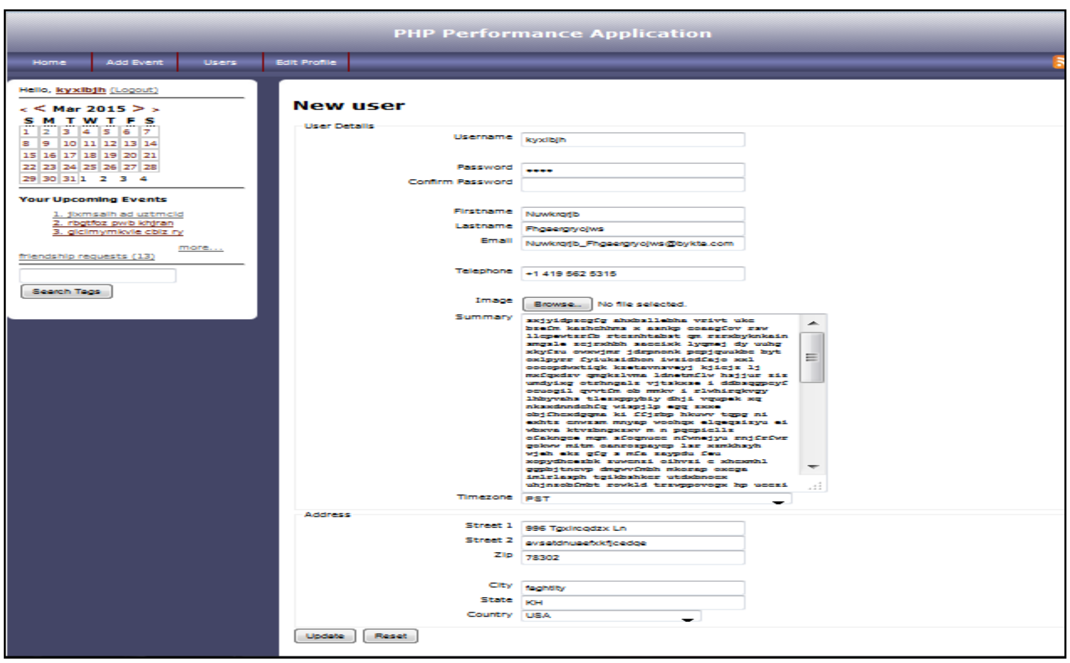

Figure 4: Olio Application – Update/Reset User Details

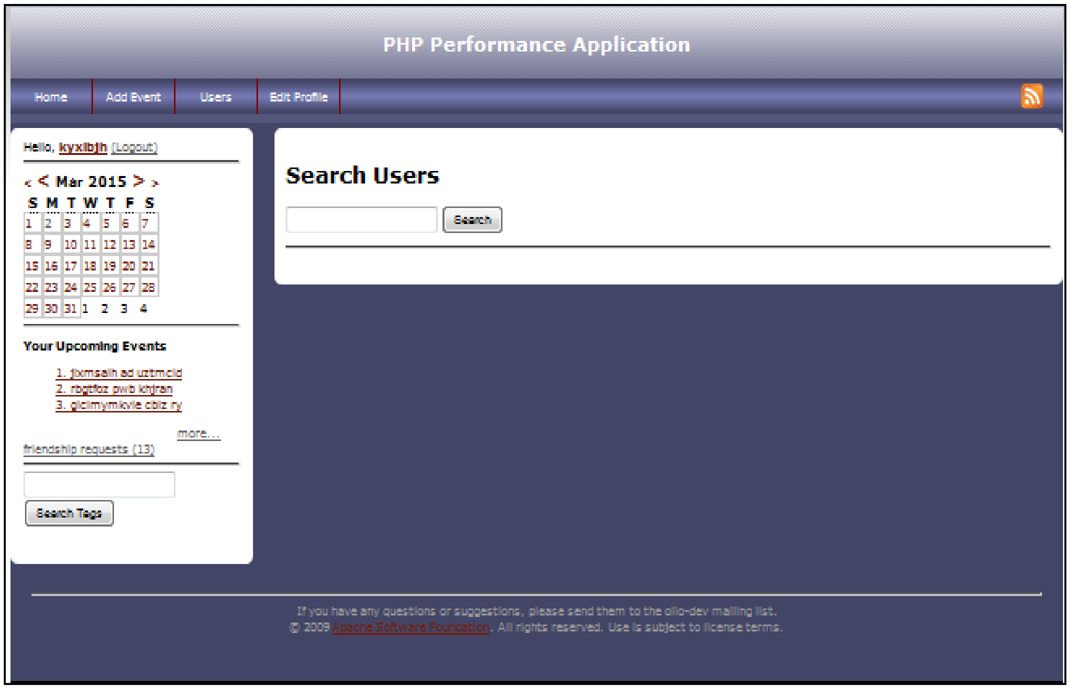

Figure 5: Olio Application – Search Users and Events using tags

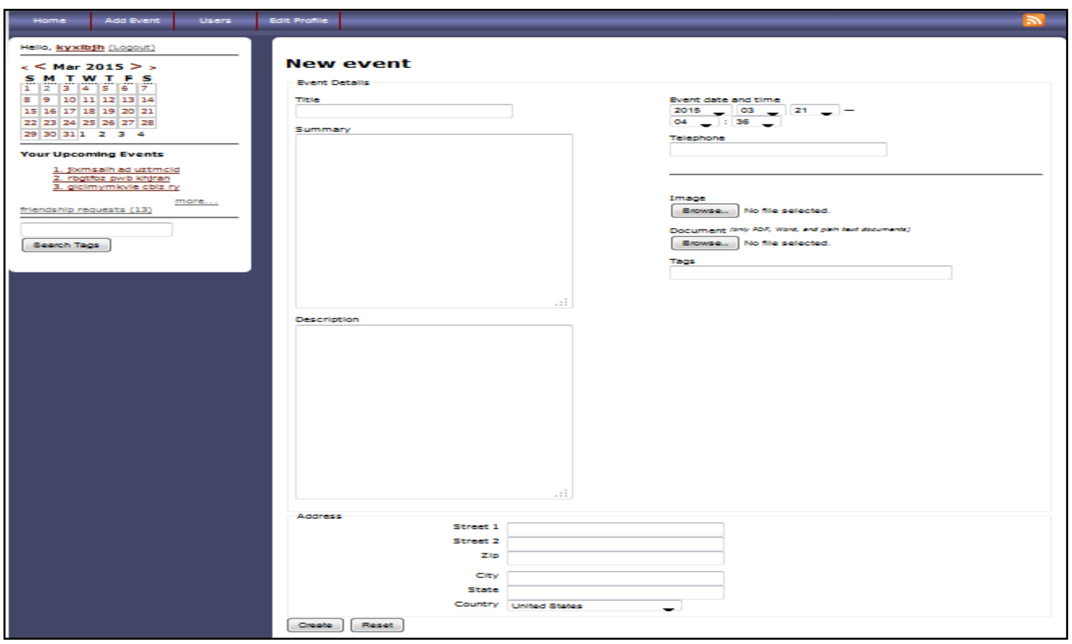

Figure 6: Olio Application – Add Event

Figure 7: Olio Application – Logout

VITA

Neha Soni earned her undergraduate degree in Information Technology from Maharashtra Institute of Technology (MIT), Pune, India. She has worked as a Senior Systems Engineer with a top-notch service-based company in India for four years where she has developed innovative web solutions for clients like Boeing, Microsoft and Toyota. She is currently pursuing a master’s degree in Computing and Information Sciences with a Software Engineering concentration at University of North Florida, Jacksonville. She plans to work once she graduates and dreams to start her own Information Technology firm in near future.