

Technical Report
Cloud Computing Platforms
Author: Konstantinos Domdouzis
Date: 14 September 2012

1. Definition

Cloud computing is the use of hardware and software resources over the Internet. These resources can be offered from different data centers which all together consist of a cloud. Cloud computing is based on the use of virtual servers which offer software resources, infrastructure, and hosting to users. The users are not involved to the maintenance of any of these servers and they are only focused to the services they provide.

The architecture of a Cloud computing platform includes two elements: the “front-end” and the “back-end”. The “front-end” includes the client’s device and the required applications for accessing the “back-end”. The “back-end” is the cloud itself. This includes various computer machines, data storage systems, and servers. A central server is used to monitor the whole system. A special type of software called “Middleware” is used to allow computers to communicate with each other.

There are four major deployment models for Cloud computing:

- Public Cloud**: This is the most common model in use. It is a public or external cloud where services are provided over the Internet.
- Community Cloud**: If several organizations seek to share resources, then a community cloud can be established.
- Hybrid Cloud**: In this model, two separate clouds are joined together.
- Private Cloud**: These are clouds on a private host.

Depending on the services offered by the Clouds, Cloud computing can be classified as following:

- Infrastructure as a Service (IaaS)
- Platform as a Service (PaaS)
- Software as a Service (SaaS)
- Storage as a Service (STaaS)
- Security as a Service (SECaaS)

- Data as a Service (DaaS)
- Test Environment as a Service (TEaaS)
- Desktop as a Service (DaaS)
- API as a Service (APIaaS)

The most common Cloud computing platforms are Microsoft Azure, Amazon Elastic Cloud (EC2), Google AppEngine, IBM Smart Cloud Services, Heroku, Skytap, Rackspace Cloud, Go Grid, and Enlight Cloud.

2. Cloud Computing Platforms

This report focuses on three main Cloud platforms: Microsoft Azure, Amazon Elastic Cloud (EC2) and the Google AppEngine.

•Amazon Elastic Cloud (EC2) (also known as Amazon Web Services EC2 or AWS EC2)

This is a web service that provides to the users resizable computing capacity in the cloud. Amazon Elastic EC2 allows the provision of a virtual server. The users are billed only for the use of resources. Amazon EC2 works in conjunction with Amazon Simple Storage Service (Amazon S3), Amazon Relational Database Service (RDS), Amazon SimpleDB, and Amazon Simple Queue Service (Amazon SQS).

-Amazon S3: This is a simple web services interface that can be used to store and retrieve any amount of data, at any time, from anywhere on the web.

-Amazon SQS: This service allows the moving of data between distributed components of applications that perform different tasks.

-Amazon SimpleDB: This is a web service for running queries on structured data in real-time.

The most important element though for the interaction of the users with the clouds is the Amazon Machine Instance (AMI). An AMI is a template that allows the users to launch instances which are running copies of the AMI. The AMI can be accessed by the Amazon Web Services (AWS) Console which is at the following address:

<https://console.aws.amazon.com/ec2/>

There are different types of instances that a user can specify. These are Standard instances, Micro, High-Memory, High-CPU, Cluster Compute, Cluster GPU, and High I/O. Standard Instances include the following:

- Small Instance (Default) 1.7 GB of memory, 1 EC2 Compute Unit (1 virtual core with 1 EC2 Compute Unit), 160 GB of local instance storage, 32-bit or 64-bit platform
- Medium Instance 3.75 GB of memory, 2 EC2 Compute Units (1 virtual core with 2 EC2 Compute Units each), 410 GB of local instance storage, 32-bit or 64-bit platform
- Large Instance 7.5 GB of memory, 4 EC2 Compute Units (2 virtual cores with 2 EC2 Compute Units each), 850 GB of local instance storage, 64-bit platform

•Extra Large Instance 15 GB of memory, 8 EC2 Compute Units (4 virtual cores with 2 EC2 Compute Units each), 1690 GB of local instance storage, 64-bit platform

The cost of using Amazon EC2 Platform can be calculate monthly through the online calculator: <http://calculator.s3.amazonaws.com/calc5.html> New customers of Amazon EC2 can use the AWS Free Usage Tier which allows users to run for free an Amazon EC2 Micro-Instance for one year. The region however where the Amazon EC2 Platform instances run can also specify the cost. For example, for the European Union, the instances are located in Ireland and the prices are defined as follows:

Region: EU (Ireland) ▾		
	Linux/UNIX Usage	Windows Usage
Standard On-Demand Instances		
Small (Default)	\$0.085 per Hour	\$0.115 per Hour
Medium	\$0.170 per Hour	\$0.230 per Hour
Large	\$0.340 per Hour	\$0.460 per Hour
Extra Large	\$0.680 per Hour	\$0.920 per Hour
Micro On-Demand Instances		
Micro	\$0.020 per Hour	\$0.020 per Hour
High-Memory On-Demand Instances		
Extra Large	\$0.506 per Hour	\$0.570 per Hour
Double Extra Large	\$1.012 per Hour	\$1.140 per Hour
Quadruple Extra Large	\$2.024 per Hour	\$2.280 per Hour
High-CPU On-Demand Instances		
Medium	\$0.186 per Hour	\$0.285 per Hour
Extra Large	\$0.744 per Hour	\$1.140 per Hour
Cluster Compute Instances		
Quadruple Extra Large	N/A*	N/A*

Additional instances are offered by Amazon when Microsoft Windows run in combination with Microsoft SQL Server.

	Windows Usage	Windows with SQL Standard Usage	Windows with SQL Web Usage
Standard On-Demand Instances			
Small (Default)	\$0.115 per Hour	\$0.629 per Hour	\$0.155 per Hour
Medium	\$0.230 per Hour	\$0.744 per Hour	\$0.270 per Hour
Large	\$0.460 per Hour	\$0.974 per Hour	\$0.500 per Hour
Extra Large	\$0.920 per Hour	\$1.434 per Hour	\$0.960 per Hour
Micro On-Demand Instances			
Micro	\$0.020 per Hour	N/A	\$0.075 per Hour
High-Memory On-Demand Instances			
Extra Large	\$0.570 per Hour	\$1.084 per Hour	\$0.610 per Hour
Double Extra Large	\$1.140 per Hour	\$1.654 per Hour	\$1.180 per Hour
Quadruple Extra Large	\$2.280 per Hour	\$3.307 per Hour	\$2.360 per Hour
High-CPU On-Demand Instances			
Medium	\$0.285 per Hour	N/A	\$0.325 per Hour
Extra Large	\$1.140 per Hour	\$2.167 per Hour	\$1.220 per Hour
Cluster Compute Instances			
Quadruple Extra Large	N/A*	N/A	N/A

The customers can also rent reserved instances. In this case, they have to make an upfront payment with a significant discount. There are three types of reserved instances (Light, Medium, Heavy Utilization).

• Microsoft Windows Azure

Windows Azure Platform is Microsoft's platform for running applications in the cloud. The user gets on-demand computing and storage to host, scale and manage web applications through Microsoft data centers.

The components of the Microsoft Windows Azure Platform are the following:


-Windows Azure: provides a Windows-based environment for running applications and storing data on servers in Microsoft data centers. Windows Azure allows the running of many applications by simultaneous users. This is allowed through the provision of multiple instances, each executing in its own virtual machine. The user can access Windows Azure through the Windows Azure portal and can select between two different accounts: hosting or storage.

-SQL Azure: provides data services in the cloud based on SQL server. Storage in Windows Azure Platform is achieved through the use of blobs, tables, and queues. A blob contains binary data and when a user selects a storage account, then blob containers which include more than one blobs are included. In order to store data in a better format, tables are used. These are not relational tables, however the data are stored in relational entities which include attributes.

-Windows Azure Platform AppFabric: provides cloud services for connecting applications running in the cloud. The lifecycle of AppFabric includes three stages. These are the Developer Tools, the AppFabric Container, and the Application Manager. Developer Tools allow users to build cloud-based applications using the .NET Framework and Visual Studio. The AppFabric Container includes a number

of cluster which can be used to host custom or Windows Communication Foundation (WCF) code. The Application Manager is used for the management of the whole AppFabric Application.

The functionality of the Microsoft Windows Azure Platform can be concisely presented in the following diagram:



Windows Azure offers a 90-day free trial, at the end of which the user can elect to continue using it and this comes with 10 shared instances for an additional 12-month at no charge. The trial also includes a Windows Azure SQL database. The following items are offered per month at no charge:

- **Web Sites:** Up to 10 web sites
- **Data Transfers:** 165 MB of outbound data transfers per day, per subscription; unlimited inbound data transfers
- **Storage:** 1 GB
- **Relational Database:** 20 MB of a third-party, MySQL database

The users can select to switch to the Reserved Instances but in this case a Pay-as-you-Go rate applies. The user however are able to use 100 instances and 10GB of storage space. Based on the Reserved Instance size, the following costs apply:

Reserved Instance Size	CPU Cores	Memory	Preview* price/hour (33% Discount)	General Availability (GA)* price/hour
Small	1	1.75 GB	\$0.08	\$0.12
Medium	2	3.5 GB	\$0.16	\$0.24
Large	4	7 GB	\$0.32	\$0.48

• Google App Engine

Google App Engine allows users to run web applications on Google’s infrastructure. It supports applications written in different programming languages, such as Java, JavaScript, or Ruby. App Engine uses two Python runtime environments and provides a Go runtime environment that runs natively Go code. The storage of Google App Engine is achieved through App Engine Datastore, Google Cloud SQL, and Google Cloud Storage. The App Engine can serve both traditional website content as well as dynamic.

The function of Google App Engine is based on three parts: the runtime environment, the datastore, and the scalable services.

• Runtime Environment: Google App Engine provides two runtime environments: Python and Java. The environment which is chosen depends on the application that must be developed.

• Datastore: Data are written to the datastore as entities. Each entity has its own key that uniquely identifies it. An entity has one or more properties, each of which has a name, and a value that is one of several primitive data types.

•**Scalable Services:** Google App Engine allows the development of scalable applications by constraining how these applications access data.

According to Google [<https://developers.google.com/appengine/docs/billing>], the following costs apply to Google App Engine:

Resource	Unit	Unit cost
Outgoing Bandwidth	Gigabytes	\$0.12
Frontend Instances (F1)	Instance hours	\$0.08
Frontend Instances (F2)	Instance hours	\$0.16
Frontend Instances (F4)	Instance hours	\$0.32
Discounted Instances	Instance hours	\$0.05
Backend Instances (B1 class)	Hourly per instance	\$0.08
Backend Instances (B2 class)	Hourly per instance	\$0.16
Backend Instances (B4 class)	Hourly per instance	\$0.32
Backend Instances (B8 class)	Hourly per instance	\$0.64
Stored Data (Blobstore)	Gigabytes per month	\$0.13
Stored Data (Datastore)	Gigabytes per month	\$0.24
Stored Data (Logs Data)	Gigabytes per month	\$0.24
Stored Data (Task Queue)	Gigabytes per month	\$0.24
Channel	Channel opened	\$0.00001 (\$0.01/100 channels)
Recipients Emailed	Email	\$0.0001
XMPP	XMPP stanzas	\$0.000001 (\$0.10/100,000 stanzas)
Logs API	Gigabytes	\$0.12
SNI SSL certificates	Sets of five SNI certificate slots per month	\$9.00
SSL Virtual IPs (VIPs)	Virtual IP per month	\$99.00
PageSpeed bandwidth	Gigabytes (in addition to outgoing bandwidth charges)	\$0.39

There are three types of low-level Datastore operations:

Operation	Cost
Write	\$0.10 per 100k operations
Read	\$0.07 per 100k operations
Small	\$0.01 per 100k operations

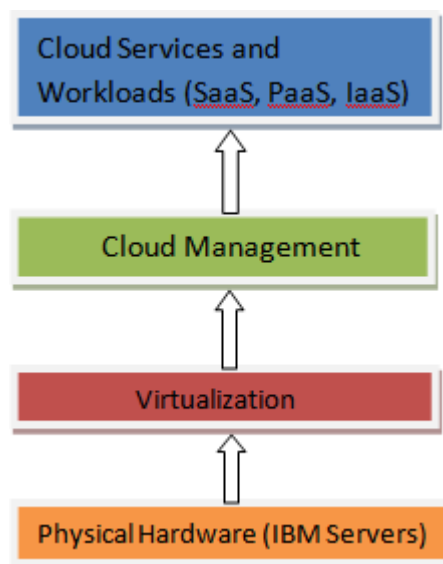
These low-level operations map to API-level operations as follows:

High-Level Operation	Low-Level Operations Required
-----------------------------	--------------------------------------

Entity Get (per entity)	1 read
New Entity Put (per entity, regardless of entity size)	2 writes + 2 writes per indexed property value + 1 write per composite index value
Existing Entity Put (per entity)	1 write + 4 writes per modified indexed property value + 2 writes per modified composite index value
Entity Delete (per entity)	2 writes + 2 writes per indexed property value + 1 write per composite index value
Query	1 read + 1 small per entity retrieved
Query (keys only)	1 read + 1 small per key retrieved
Key allocation (per key)	1 small

• **IBM SmartCloud**

IBM SmartCloud is a branded ecosystem of cloud computing products and solutions for IBM. It includes characteristics of IaaS, PaaS, and SaaS. It combines these characteristics under three umbrellas: SmartCloud Foundation, SmartCloud Services, and SmartCloud Solutions. The IBM Cloud Platform includes a web portal to administer cloud resources and APIs that enable users to control cloud resources. The functionality of IBM SmartCloud can be presented in the following diagram:



The basic characteristics of the IBM SmartCloud Platform are the following:

- The **overview** (main page, ibm.com/cloud/developer/) is your access point to the Cloud environment.
- The **control panel** lets you add and manage instances, images, and storage. You have three view options:
 1. All of your provisioned instances
 2. The images you created from your customized instances
 3. Persistent cloud storage units that you provisioned

- The **asset catalog** powered by Rational Asset Manager software, is a robust interface you use to manage the details of your created custom software images. You can add permissions, user guides, and general content, as well as customize the details of your image offerings. The asset catalog also contains some key developer forums and other cloud materials and artifacts. (More discussion on the catalog after this list.)
- The **account page** holds users' private keys (more about keys in the next section), reserved IP addresses, an aggregation of your account and instance/image notifications, some global statistics, and basic user information.
- The **support page** gives you access to the IBM Cloud communities and a diverse variety of other cloud materials. You can bring an issue to the attention of others and view feedback. There is also a document and video library here for all kinds of detailed resources.

Sheffield Hallam University

Technical Report - Cloud Computing Platforms

DOMDOUZIS, Konstantinos <<http://orcid.org/0000-0003-3679-3527>>

Available from the Sheffield Hallam University Research Archive (SHURA) at:

<http://shura.shu.ac.uk/33234/>

Copyright and re-use policy

Please visit <http://shura.shu.ac.uk/33234/> and <http://shura.shu.ac.uk/information.html> for further details about copyright and re-use permissions.