Cloud Computing Solutions

Felician ALECU
Academy of Economic Studies, Bucharest, Romania
alecu.felician@ie.ase.ro

The cloud means the Internet. The term is derived from the way in which the Internet is represented into the network diagrams. Cloud computing is a paradigm that incorporates the concept of software as a service. This means the software and data are stored on servers that can be accessed over the Internet. The current cloud computing architecture involves the existence of data centers that are able to provide services to the clients located all over the world. The major players in field of the cloud computing are Google, Microsoft, Amazon, Yahoo, and some traditional hardware producers like HP, IBM, Intel.
Keywords: *cloud computing, Internet, software services, applications*

1 About the Cloud Computing

Cloud computing allows to move the processing effort from the local devices to the data center facilities. In such a way, any phone, for example, could be able to solve complex differential equation systems by simply passing the specific arguments to a data center service that will be capable to give back the results in a very short time. In these conditions, the security of data and applications becomes a major issue.

Simply speaking, the cloud means the Internet. This term is derived from the way in which the Internet is often represented into the network diagrams.

A storage cloud provides storage services (block or file based services); a data cloud provides data management services (record-based, column-based or object-based services); and a compute cloud provides computational services. Often these are layered (compute services over data services over storage service) to create a stack of cloud services that serves as a computing platform for developing cloud-based applications.

Cloud computing represents a new paradigm of the Internet computing in which the software is seen as a service and the applications and data are stored on multiple servers that can be accessed from the Internet.

Cloud computing is totally different than grid computing that tries to create a virtual computer by joining together a cluster of computers. The aim of a grid computing architecture is to solve large tasks by using the advantage of concurrency and parallelism.

The current cloud computing architecture (figure 1) involves the existence of data centers that are able to provide services to the clients located all over the world. In this context, the cloud can be seen as a unique access point for all the requests coming from the customers/clients.

It is important to notice a client could be a hardware device and/or a software application, like a browser, for example.

A mobile phone or PDA can successfully play the role of the cloud client. For this reason, the mobile device should run on the Android or iPhone platforms. Also, a web browser, like Google Chrome, can be a cloud client without any problem. Basically, the client is renting or simply accessing the processing capacity needed from the data center. The quality of the service becomes a crucial factor of the cloud computing success.

Cloud computing allows to move the processing effort from the local devices to the data center facilities. In such a way, any phone, for example, could be able to solve complex differential equation systems by simply passing the specific arguments to a data center service that will be capable to give back the results in a very short time. In these conditions, the security of data and applications becomes a very major issue.

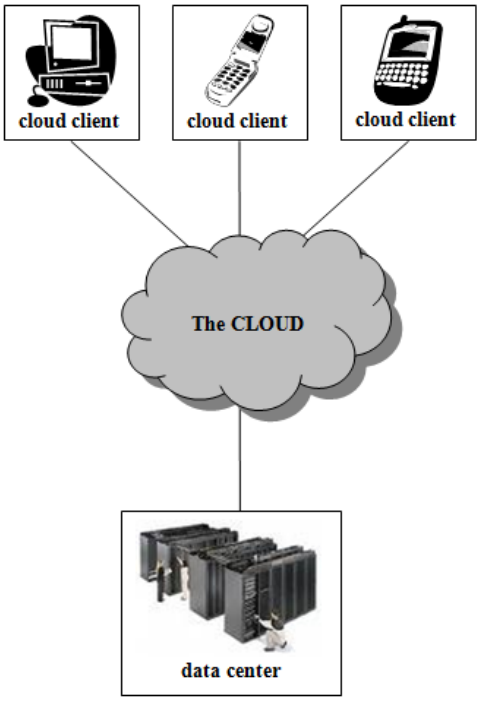

Fig. 1. Cloud Computing Architecture

The main advantages of the cloud computing are the following:

- there is no need to download or install a specific software;
- the cost is low or even free, in some cases;
- if the client computer crashes, there is almost nothing lost because everything is stored into the cloud;
- there is no need to update the local system when some new fix packs are released;
- cloud computing can be used on clients having minimal hardware requirements, like mobile phones or PDAs;
- the problem of licensing different software packages is moved to the data center level;
- no costs (or very small ones) for hardware upgrades;
- the users are not dependent by their personal computer because they can use any other device having an Internet connection and minimum software requirements.

Of course there are some disadvantages as well, like:

- an Internet connection is required in order to be able to access and use the cloud and this Internet dependence the offline mode impossible. On the other hand, some applications require a high speed Internet connection so the traffic speed may affect the overall performances;
- on a long term basis, the subscription fee may be more expensive than buying the hardware, for example;
- a very big concern is the data security because the data and the software are located on remote servers that can crash or disappear without any additional warnings. In this context, the service quality becomes crucial and the need of the backups is vital.

2 Security Benefits of Cloud Computing

As stated before, a very big concern regarding the cloud computing is the data security because the data and the software are located on remote servers that can crash or disappear without any additional warnings.

Even if it seems improbable, the cloud computing provides some major security

benefits for individuals or companies, like the following:

- centralized data storage – this goes to reduced effects of losing some hardware items, like a laptop, for example. While the main part of the applications and data is stored into the cloud, losing a client is not a big issue anymore – there are no sensitive data lost and a new client can be connected to the cloud very fast;
- monitoring of data access becomes easier because it is enough to monitor only one place, not thousands of computers belonging to a major company, for example;
- increased uncertainty – it is almost impossible for a thief to determine which physical component to steal in order to get a digital asset;
- virtualization allows a rapid replacement of a compromised server located into the cloud without major costs or damages. Also, the downtime for computers in the cloud could be substantially reduced because it is very easy to create a clone by using an image;
- logging – extended logs can be activated because the cloud is big enough to store large collections of data;
- the security changes can be easily tested and implemented

On the other hand, table 1 shows the major incidents that impacted the cloud computing providers during the last years [5].

Table 1. The most significant cloud computing incidents
(taken from *Cloud Computing Incidents Database*)

Date	***Product***	***Incident Type***	***Comments***
2009-01-31	Google	Outage	Lasted up to 1 hour
2009-01-30	Ma.gnolia	Data Loss	Both online and backup databases affected.
2008-10-18	AWS Services	Security	Issue present since service launch
2008-10-15	Gmail	Outage	Lasted more than 24 hours
2008-09-18	Google Docs	Security	Limited to ISP(s) in Thailand
2008-09-15	App Engine	Outage	Datastore writes experienced elevated latencies and error-rates.
2008-09-02	Google Apps	Security	Malicious service provider could impersonate a user at other service providers.
2008-08-26	FlexiScale	Outage	Full extended outage
2008-08-12	Gmail	Outage	Users unable to use webmail due to issues with loading contacts between 14:00 and 16:00
2008-08-08	The Linkup	Data Loss	Data claimed to be safe but inaccessible
2008-07-20	Amazon S3	Outage	Full outage for 8 (weekend) hours
2008-07-10	MobileMe	Outage	Scheduled outage window exceeded during upgrade to MobileMe
2008-07-09	.Mac	Outage	Full outage (except mail) during upgrade to MobileMe 18:00-00:00
2008-04-28	EC2	Outage	Result of a customer creating a large number of firewall rules and instances.
2008-02-15	Amazon S3	Outage	Early morning outage (04:31-06:48) caused by authentication service overload
2008-01-07	Salesforce.com	Outage	Affected all instances and supporting infrastructure
2007-09-29	EC2	Outage	Result of a customer creating a large number of firewall rules and instances.

3 Applications of Cloud Computing

Due to the fact cloud computing becomes a very interesting subject, Microsoft announced it is developing a new Windows platform, called Windows Azure [1], which will be able to run cloud based applications.

The Azure Services Platform uses a specialized operating system, Windows Azure, to run its "fabric layer" — a cluster hosted at Microsoft's datacenters that manages computing and storage resources of the computers and provisions the resources (or a subset of them) to applications running on top of Windows Azure.

Windows Azure, which was known as "Red Dog" during its development, has been described as a "cloud layer" on top of a number of Windows Server systems, which use Windows Server 2008 and Hyper-V to provide virtualization of services.

The platform includes five services — Live Services, SQL Services, .NET Services, SharePoint Services and Dynamics CRM Services — which the developers can use to build the applications that will run in the cloud. A client library, in managed code, and associated tools are also provided for developing cloud applications in Visual Studio.

Scaling and reliability are controlled by the Windows Azure Fabric Controller so the services and environment don't crash if one of the servers crash within the Microsoft datacenter and provides the management of the user's web application like memory resources and load balancing.

The main components of the Azure Services Platform are: Windows Azure itself, providing a Windows-based environment for running applications and storing data on servers in Microsoft datacenters; .NET Services, linking the distributed infrastructure to applications; SQL Services, for database use; and Live Services, which links Microsoft's Live applications and provides data synchronization, search and downloading, among other features; SharePoint services and finally the Dynamics CRM ones.

The Azure Services Platform can currently run .NET Framework applications written in C#, while supporting the ASP.NET application framework and associated deployment methods to deploy the applications onto the cloud platform. Two SDKs have been made available for interoperability with the Azure Services Platform: The Java SDK for .NET Services and the Ruby SDK for .NET Services. These enable Java and Ruby developers to integrate with .NET Services.

Fig. 2. Windows Azure Services

One of the first application suite belongs to the Amazon [2] and it is called AWS (Amazon Web Services), launched in 2002. AWS is a collection of remote services

intended for client applications or web sites. According to the Amazon news, there are almost 500.000 developers that subscribed to the AWS.

Amazon Web Services' offerings are accessed over HTTP, using REST and SOAP protocols. All are billed on usage, with the exact form of usage varying from service to service.

The Amazon Web Services suite includes a component called Amazon Elastic Compute Cloud (or EC2), that allows to the users to rent from Amazon processing power to be used to run their own applications. Basically, the users are paying some fees and they receive virtual machines. The cloud is an elastic one just because the user can start, stop and create the virtual machines through the web service. There are three predefines sizes for the virtual machines that can be rented: small, medium and large, depending on the physical hardware performances.

Amazon Elastic Compute Cloud is a commercial web service that allows customers to rent computers on which to run their own computer applications. EC2 allows scalable deployment of applications by providing a web services interface through which a customer can create virtual machines, i.e. server instances, on which the customer can load any software of his choice. A customer can create, launch, and terminate server instances as needed, paying by the hour for active servers, hence the term "elastic". A customer can set up server instances in zones insulated from each other for most failure causes so that one may be a backup for the other and minimize down time. Amazon.com provides EC2 as one of several web services marketed under the blanket term Amazon Web Services (AWS).

The main advantage of the AWC EC2 web service suite is the user doesn't need to install or run applications on the local computer, so there is no need of hardware support and maintenance.

Today, EC2 is able to host Windows Server and SQL Server database. From October 2008, also Oracle is running on the EC2 platform.

Amazon charges customers in two primary ways: *hourly charge per virtual machine* and *data transfer charge*. The hourly virtual machine rate is fixed, based on the capacity and features of the virtual machine. Amazon advertising describes the pricing scheme as "you pay for resources you consume," but defines resources such that an idle virtual machine is consuming resources, as opposed to other pricing schemes where one would pay for basic resources such as CPU time.

Customers can easily start and stop virtual machines to control charges, with Amazon measuring with one hour granularity. Some are thus able to keep each virtual machine running near capacity and effectively pay only for CPU time actually used.

As of March 2009, Amazon's time charge is about $73/month for the smallest virtual machine without Windows and twelve times that for the largest one running Windows. The data transfer charge ranges from $.10 to $.17 per gigabyte, depending on the direction and monthly volume. Amazon does not have monthly minimums or account maintenance charges.

Sun network.com (Sun Grid) enables the user to run Solaris OS, Java, C, C++, and FORTRAN based applications. First, the user has to build and debug his applications and runtime scripts in a local development environment that is configured to be similar to that on the Sun Grid. Then, he needs to create a bundled zip archive (containing all the related scripts, libraries, executable binaries and input data) and upload it to Sun Grid. Finally, he can execute and monitor the application using the Sun Grid web portal or API. After the completion of the application, the user will need to download the execution results to his local development environment for viewing.

Aneka, which is being commercialized through Manjrasoft, is a .NET-based service-oriented resource management platform. It is designed to support multiple application models, persistence and security solutions, and communication protocols such that the preferred selection can be changed at anytime without affecting an existing Aneka

ecosystem. To create an Aneka Cloud, the service provider only needs to start an instance of the configurable Aneka container hosting required services on each selected desktop computer. The purpose of the Aneka container is to initialize services and acts as a single point for interaction with the rest of the Aneka Cloud. Aneka provides SLA support such that the user can specify QoS requirements such as deadline (maximum time period which the application needs to be completed in) and budget (maximum cost that the user is willing to pay for meeting the deadline). The user can access the Aneka Cloud remotely through the Gridbus broker. The Gridbus broker also enables the user to negotiate and agree upon the QoS requirements to be provided by the service provider.

One of the biggest promoters of the cloud computing is Google [3]. They already have a massive computer infrastructure (the cloud) where millions of people are connecting to. Moreover, Google joined with IBM [4] trying to lead the cloud computing to a new stage.

The Google cloud can be accessed by using the Google Apps, intended to be a software as a service suite dedicated to information sharing and security.

Google App Engine is a platform for building and hosting web applications on Google Web Servers. It was first released as a beta version in April 2008. Google App Engine has both free and paid accounts. According to Google, a free account can use up to 1 GB of persistent storage and enough CPU and bandwidth for about 5 million page views a month. no hardware or software licenses are needed and the data can be accessed in a secured manner from anywhere in the world by simply using a device connected to the Internet.

Currently, the only supported programming language is Python. A limited version of the Django web framework is available, as well as a custom Google-written web app framework similar to JSP or ASP.NET. Google has said that it plans to support more languages in the future, and that the Google App Engine has been written to be language independent. Any Python framework that supports the WSGI using the CGI adapter can be used to create an application; the framework can be uploaded with the developed application. Third-party libraries written in pure Python may also be uploaded.

Fig. 3. Google Apps

Compared to other scalable hosting services such as Amazon EC2, App Engine provides more infrastructures to make it easy to write scalable applications, but can only run a limited range of applications designed for that infrastructure. App Engine's

infrastructure removes many of the system administration and development challenges of building applications to scale to millions of hits. Google handles deploying code to a cluster, database shading, monitoring, failover, and launching application instances as necessary.

While other services let users install and configure nearly any compatible software, AppEngine requires developers to use Python as the programming language and a limited set of APIs. Current APIs allow storing and retrieving data from a BigTable non-relational database; making HTTP requests; sending e-mail; manipulating images; and caching. Most existing Web applications can't run on App Engine without modification, because they require a relational database.

Per-day and per-minute quotas restrict bandwidth and CPU use, number of requests served, number of concurrent requests, and calls to the various APIs, and individual requests are terminated if they take more than 30 seconds or return more than 10MB of data.

Google App Engine's datastore has a SQL-like syntax called "GQL". Select statements in GQL can be performed on one table only. GQL intentionally does not support the *join* statement, because it is seen to be inefficient when queries span more than one machine. Instead, one-to-many and many-to-many relationships can be accomplished using *ReferenceProperty*(). This shared-nothing approach allows disks to fail without the system failing.

The where clause of select statements can perform >, >=, <, <= operations on one column only. Therefore, only simple were clauses can be constructed. Switching from a relational database to the Datastore requires a paradigm shift for developers when modeling their data.

App Engine limits the maximum rows returned from an entity get to 1000 rows per Datastore call. Most web database applications use paging and caching, and hence do not require this much data at once, so this is a non-issue in most scenarios. If an application needs more than 1,000 records per operation, it can use its own client-side software or an Ajax page to perform an operation on an unlimited number of rows.

Task scheduling is a heavily requested feature by App Engine users for future releases. Batch updates that take longer than 30 seconds cannot be executed. Data export and bulk operations using MapReduce can be managed using software running outside of App Engine.

Unlike a relational database such as Oracle, Microsoft SQL Server, MySQL, or PostgreSQL, the Datastore API is not relational in the SQL sense.

There are three main categories of services, as the following:

1. Messaging – includes Gmail, Calendar and Google Talk;
2. Collaboration – Google Docs, Video and Sites;
3. Security – email security, encryption and archiving.

But Microsoft, Amazon and Google are far from the only companies to understand and embrace cloud computing. Several reports have suggested that the criminal underground is also using cloud computing in what people are starting to call fraud-as-a-service.

4 Conclusions

Cloud computing represents a very interesting concept of our days. There are a lot of applications able to exploit the cloud and the list is expanding very fast. Many devices are already cloud compatible – the traditional computers, PDAs, mobile phones and even browsers, like the Google Chrome.

Today, the information infrastructure is moving faster to a simple but very innovative concept called cloud computing. In this context, cloud computing is potentially able to offer major security benefits for the cloud enabled solutions and applications.

References

[1] Windows Azure Services Platform - www.microsoft.com/azure
[2] Amazon Web Services - aws.amazon.com
[3] Google Apps – www.google.com/a

[4] IBM Cloud Computing - www.ibm.com/ibm/cloud/

[5] Cloud Computing Incidents Database - wiki.cloudcommunity.org/wiki

Felician ALECU has graduated the Faculty of Cybernetics, Statistics and Economic Informatics in 2000 and he holds a PhD diploma in Economics from 2006. Currently he is lecturer of Economic Informatics within the Department of Economic Informatics at Faculty of Cybernetics, Statistics and Economic Informatics from the Academy of Economic Studies. He is the author of more than 20 journal articles in the field of parallel computers, grid computing and distributed processing.