Cloud Service and Deployment Models

Cloud Strategy Partners, LLC

Sponsored by: IEEE Educational Activities and IEEE Cloud Computing

Course Presenter's Biography

This IEEE Cloud Computing tutorial has been developed by Cloud Strategy Partners, LLC. Cloud Strategy Partners, LLC is an expert consultancy firm that specializes in Technology and Strategy relating to Cloud Computing.

Course Summary

In this tutorial we will review the Cloud Computing definition including Cloud properties, service models, deployment models. We will also discuss Cloud services models IaaS, PaaS, SaaS and corresponding general use cases as well as High Performance Computing (HPC) in clouds. Computer Grids and Cloud Computing will also be reviewed. We will present Cloud deployment models and corresponding use cases including Basic cloud deployment models: Private, Public, Hybrid, Community as well as emerging cloud deployment model such as Federated clouds and Interclouds. Finally, we will provide general enterprise oriented use cases including Virtual Private Cloud (TBA).

Course Outline

This Lesson will cover the following topics:

- Background: Cloud Computing definition including Cloud properties, service models, deployment models
- Cloud services models IaaS, PaaS, SaaS and corresponding general use cases
- High Performance Computing (HPC) in clouds
- Computer Grids and Cloud Computing
- Cloud deployment models and corresponding use cases including Basic cloud deployment models: Private, Public, Hybrid, Community And also
- Emerging cloud deployment models: Federated clouds, Interclouds
- General Enterprise oriented use cases including Virtual Private Cloud (TBA)
- Summary and take away

Cloud Computing Definition: Components

This slide is a visual representation of the NIST Cloud Computing Definition. We have grouped the NIST concepts into three areas-Deployment Models, Service Models, and Essential Characteristics. As you can see the key concepts in each area are Cloud characteristics On-demand self-service Broad network access Resource pooling Rapid elasticity Measured Service. Basic service models Software as a Service (SaaS) Platform as a Service (PaaS) Infrastructure as a Service (IaaS). Deployment models Private clouds Public clouds Hybrid clouds Community clouds Federated clouds, Interclouds

Cloud Service Models

There are actually more service models than the three (IaaS, PaaS, SaaS) widely in use today. Service models like Data Analytics as a Service and HPC/Grid as a Service are emerging as useful models. How one Selects the appropriate service model depends on factors such as availability of suitable application software Need for development and test environment, need for effective computing infrastructure control and management required distribution of data, services, and infrastructure, existence and complexity of enterprise IT infrastructure and datacenter/warehouse.

Use cases: Infrastructure as a Service (IaaS)

Now let us look at each of these Service Models individually.

First consider Infrastructure as a Service (IaaS). IaaS is easily used for
Website hosting, where a web server and operating system stack are put on VM's, where they can easily take advantage of cloud features such as easy scaling, global availability, managed environment, geographical load balancing, special content delivery front-end or infrastructure
Disaster Recovery preparedness, where a running server snapshot is made (including in memory image) which can be reconstituted in the event of a disaster
Test and Development where easy to create VM's can be used eliminating the need for creating test environment which may require complex infrastructure for testing new products under stress load. Actually any kind of Short-term collaborative projects, requiring a lab on demand
And of course a way to solve the needs of Cyclical and/or seasonal capacity, where the software architecture allows use of scale-out sa a way to handle increased capacity

Use cases: Platform as a Service (PaaS)

Now let us look at Platform as a Service (PaaS). Here, users have control over applications configuration, deployment and management. Instead of using a server metaphor, because they are given an Applications hosting and deployment platform abstracted from the server Providers usually provide all the details for the computing platform, connectivity, elasticity/scalability, backup etc.

PaaS is good for deploying applications which came from an "applications container" world before like J2EE or .NET. In fact PaaS systems such as RedHat OpenShift or Cloud Foundry from IBM or Pivotal are very much like J2EE, and Windows Azure from Microsoft is very much like .NET.

SaaS applications like Customer Relations Management (CRM) and Supply Chain Management (SCM) often come with a corresponding PaaS environment which allows for users to create their own extensions or variations on the SaaS application and host them there right "next" to the SaaS system. An example of this is the PaaS environment Force.com which complements the SaaS application Salesforce.com

Use cases: Cloud Application/Software as a Service (Apps/SaaS)

Now let's look at the Service Model of Cloud Application/Software as a Service (Apps/SaaS)

These are complete applications on the cloud ready for the user to use directly, usually through a browser, but in general through any “finished” user interface.
Email is a perfect example of a complete application one accesses through a browser. As an alternative example, the “cloud storage” applications are also SaaS with a very simple user interface –a “folder”.

In the Mobile world SaaS is known as “an app”, because the front end User Interface sits on the phone, while the back end sits in the cloud. In fact while many Mobile apps don’t look like it, they are built with the reconfigurable Mobile version of a browser called WebKit.
The SaaS providers can ensure good levels of security first by obtaining the special certification and compliance level needed in that industry and also ensuring that the underlying cloud provider obtains the special certification and compliance level needed in that industry

Cloud and High Performance Computing (HPC)

Scientists and Researchers traditionally build specialized computing systems such as Grids and Supercomputers to satisfy their needs for massive storage and computational capability. These specialized systems are very expensive.

Cloud systems have become very cost effective, comparatively, and also have been able to scale to much larger sizes than any of the biggest grids or supercomputers. As a result, Scientists and Researchers have been experimenting with using Clouds as a platform for HPC.

While the loosely coupled nature of Cloud Computing lends itself to Big Data problems (after all Internet Search was the “original” big data problem) the typical HPC problem requires closely coupled and even large coherent memory space capability. On the one hand Cloud engineers are finding ways to create specialized Cloud platforms which provide some of the capabilities HPC and Grid platforms. And on the other hand HPC and Grid scientists are finding algorithms more suited to the Cloud architecture. Together this is catalyzing an explosion of the scientific use of Clouds where previously only Grids and Supercomputers would suffice.

Computer Grids and Computer Clouds

Let’s look more deeply at the concepts of Grids and how they relate to Cloud Computing.
The Computer Grid concept was developed by the Scientific community to address needs for data-driven science where the size of the data and/or the size of the computation required is larger by orders of magnitude than the largest single “scale-up” computer. Grids have

become standardized and interoperable and have been a very successful platform for science, as the slide details.

In fact, many would say that Computer Grids can be treated as a predecessor of Cloud Computing. However Computational grids focus on computationally intensive operations, which is a different architectural objective than Clouds. This is why Clouds have made adjustments to hopefully lure Grid scientists to try to use Clouds.

High Performance Computing (HPC) and Clouds

HPC libraries and applications are optimized for specific hardware/supercomputer. HPC user application needs sometimes bypass OS kernel and communicate directly with remote user processes and hardware/storage. Most HPC systems employ parallel file system to increase I/O bandwidth of computing nodes. Standard Cloud is not designed for HPC but rather for Highly Scalability Computing (HSC). The slide details some of these scenarios.

HPC and Cloud

This illustration shows the Relationship between HPC, Cloud HSC, Cloud HPC, and Grid computing models.

The programming model many Scientific problems use on Grids is called Message Passing Interface (MPI). Also, there is an assumption that the platform is implementing a clustering technique which in turn requires a special low latency network and special high performance I/O processes.

Here you can see which technique is suited for which capability. It is still an art for to be able to run a HPC problem on a cloud and part of the battle is to find the right cloud platform for the type of problem at hand

---

HPC Cloud Use Cases – Vendors, Operators

Here are some examples of specialized Cloud platforms which are designed to serve the HPC/Grid community. Amazon has an offering which is physically separate from their main cloud. There are also add-on software products which further enhance the HPC capability on Amazon. There are also HPC clouds offered by and to research community Examples for XSEDE, EGI, and DEISA are detailed.

Amazon EC2 HPC

This slide has extremely detailed information about Amazon EC2 HPC, as Currently Amazon offers dedicated computation optimized VMs and "cluster instances" that deliver better performance to HPC users. They also offer access to GPU (Graphics Processing Units –used as Vector Machines with a special API). Please note the examples given.

Univa HPC with RightScale and Amazon AWS

This slide details the additional capability one can get on Amazon with Univa HPC and Rightscale.

Cloud Deployment Models

Now we turn to study the different Cloud Deployment Models

- Private Clouds
- Public Clouds
- Hybrid Clouds
- Community Clouds
- Federated Clouds
- Multi-clouds and Inter-clouds

Private Cloud

First we consider the Private Cloud. The NIST definition is: the cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises.

Private clouds are a choice for companies that already own datacenter and developed IT infrastructure and have particular needs around security or performance. They are a better choice for the company datacenter than Legacy servers in so many ways, bringing many benefits derived from virtualization and automation. However they also provide Challenges and disadvantages, mostly in that the enterprise needs to migrate or re-factor applications to take advantage of the Cloud automation.

Public Cloud

Now let us consider the Public Cloud. The NIST definition is: the cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider and is a form of providing public cloud services and a Cloud Service Providers business model Brings economy of scale in pooling datacenter resources, virtualization and on-demand provisioning Allows outsourcing enterprise IT infrastructure Solves/addresses disaster recovery problem Suitable for SME and agile companies.

However, the Data and processing environment are not under control of enterprises, which for applications or data with a security requirement may raise concerns. Furthermore, Service Providers can't be perfect, power outages, network issues, and so on can disrupt service. Although the security and reliability of the public cloud will almost certainly exceed that of the private cloud due to the skill and staff size of the public cloud operator versus the enterprise itself, it is not 100% fail safe. Their SLA will certainly be problematic for many enterprises.

Hybrid Cloud

The NIST definition for Hybrid Cloud is: the cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds).
Hybrid cloud, while the most complicated configuration to manage, is also the most economical model for modern companies. It combines core cloud based enterprise infrastructure and high load tasks outsourcing to public clouds. It also Combines benefits of the controlled environment in private clouds and rapid elasticity of public clouds
However, it requires deeper enterprise IT/cloud modernization. Processes/workflow require re-engineering and re-architecting. And the challenges of getting seamless integration between the Private and the Public cloud can be solved but by custom work, and the issues of Compatibility, standardization and compliance are not turnkey yet.

Community Cloud

The Community Cloud, while used by several constituencies, is something slightly different from a Public Cloud, The NIST definition is – The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises. The Community Cloud then involves cooperation and integration of IT infrastructure and resources from multiple organizations. It May serve large inter-organizational projects. For example, the scientists form many organizations at the CERN Large Hadron Collider (HC) share a Community Cloud. It requires interoperability and compliance between member organizations and their resources, including Identity management.

Emerging models: Federated and Intercloud

There are some interesting emerging models. Cloud Researchers have suggested that, like the Internet, a mechanism should exist for users to be able to utilize multiple clouds from multiple providers or companies, and not have the details visible. In the Internet, the user sees a uniform and global topology. The user is unaware of which Internet Service Provider is hosting the web site which is he viewing, and that user is also very unaware how packets have traveled in between their browser and that web site.

We will now look at the models of Cloud federation, which allows inter-cloud resources sharing and combined provisioning.

There are two types of cloud federation; a Provider side federation for resources sharing and provisioning, and a User/customer side federation that allows creation of multi-provider heterogeneous cloud infrastructures.

The provider side federation is called “Intercloud”. The Intercloud deployment model provides a general framework for multi-provider heterogeneous cloud based services and infrastructures building and operation.

Use cases for Private clouds

Now we turn to use cases for each type of Cloud. First we consider Use cases for Private clouds. Private clouds will be chosen if there is a requirement for High data I/O and low network latency: disk intensive processes, wide sensor network, or process control. There are Legacy applications and in some cases special equipment requirements. There are Specialty hardware or configuration requirements, e.g. VM with 32+GB for in-memory data processing and governance or regulatory requirements.

Use cases for Public and Hybrid Clouds

Use cases for Public and Hybrid Clouds include, situations where there are Unpredictable growth: game or social websites, marketing campaigns. Where Cyclical: applications with regular daily or seasonal traffic fluctuation such as financial markets or eCommerce. Or for easily parallelized: applications using batch processing, data analytics, media encoding Example use cases Web/social web and mobile applications Development and test, proof of concept Big Data analytics, business reporting.

Use cases for Hybrid clouds

Use cases for Hybrid clouds covers where a company combines both private and public clouds. For example, the Private cloud hosts regular workload and master processes, also security and compliance critical applications/tasks. The Public cloud hosts non-critical and not regular but computing intensive workload. This scenario requires compatibility between private and public cloud platforms.

Hybrid cloud supports the notion of Cloudbursting – term widely used by businesses to describe a situation when workload is temporarily migrated to cloud, extending and replicating the private cloud resources and VMs (using formula “buy the base, rent a spike”)

Use cases for Intercloud and Multi-cloud

As to Use cases for Intercloud and Multi-cloud, the Intercloud deployment model provides a basis for provisioning heterogeneous multi-provider cloud based project oriented infrastructures on-demand

The first and most powerful idea is for the providers to connect, and provide a seamless integration of their public clouds, just as the internet is. Just as the Internet connects to nearly everything these days, so would the Intercloud.

The other use cases have to do with uniting various combinations of clouds (enteprrise/campus, groups of Community clouds, etc) in a transparent way.
It is clear that today Clouds are somewhat “walled gardens” and interoperability and federation will emerge in some fashion or another, eventually.

Use case: Intercloud infrastructure provisioning: Workflow => Logical (Cloud) Infrastructure (1)

Let us examine a use case, where I have a large application which can be mapped out in functionality by its Master Workflow. This is illustrated in the Slide on the top. I have several cloud service providers, as illustrated by the four entities on the bottom of the slide. Actually, a participant in this can be non-cloud, as long as it participates in the same federation protocols, as indicated by the terms “Resource Service Provider”.

Through Federation, a “logical Intercloud” is formed, and the 4 different datacenters let’s say, appear to part of one virtual cloud. This is indicated as the “cloud in the middle”. To further complicate our example, the problem is being looked after by two universities and so the User Groups A and B represent traditional diagrams of a control function.

Use case: Intercloud infrastructure provisioning: Workflow => Logical (Cloud) Infrastructure (2)

Enterprise/Scientific workflow Is mapped to heterogeneous cloud infrastructure containing IaaS, PaaS components.

This slide shows the mapping of the virtual Resources

Use case: Intercloud infrastructure provisioning: Logical Infrastructure => Network Infrastructure (1)

The application is mapped to run on the virtual VM's which are mapped to the actual underlying services. The two campuses in this example can monitor and manage the Workflow across VMs across the federated cloud. Here, the network of one for the cloud carriers is used to choreograph all that occurs.

Use case: Intercloud infrastructure provisioning: Logical Infrastructure => Network Infrastructure (2)

Or as in the previous example, but the orchestration occurs over the public internet. Either way the composite application is choreographed or the elements of application code are set to run on the public cloud.

These sets of slides meant to illustrate scenarios where our definition of "an application" is not constrained to one cloud or to our native service provider either. Intercloud allows for a composite application to be deployed as it makes sense, putting some components with one Cloud Service Provider and some with another.

Enterprise related use cases

Now we are going to consider a number of Enterprise related use cases.

- Moving Enterprise IT infrastructure to cloud
- Private cloud, Hybrid cloud, Virtual Private Cloud Website hosting on cloud Application development and testing Lab on-demand in cloud Computer Cluster on-demand
- See also Univa, RightScale, Amazon use case above.

Use cases and business relationships

In considering Use cases and business relationships, there has been detailed analysis in looking at the important combinations possible, for example, considering classification Service models and deployment models, Stakeholders involvement and business relations, and Industry or community use cases.

What we get are scenarios like End users to Cloud, Enterprise to Cloud to End users, Enterprise to Cloud, Enterprise to Cloud to Enterprise, Private Cloud, Changing Cloud Providers, etc. The following slides examine several of these.

Website hosting in Cloud

The easiest and most used use case for Cloud is that of Website hosting. This use case shows how to deploy the infrastructure for this in the cloud.

In the Company's IT infrastructure, we need a Local database server, access to the various Department IT resources, and IT Administration capability. We link to the cloud using a Secure Virtual Private Network (VPN) or Virtual Private Cloud (VPC) connection, In the Cloud, we spin up the web site facilities (using a 3-tier model) including Database, servers, storage; Web servers, Apps servers; and a Load balancer.

The illustration shows the topology for this use case. This use case leverages Cloud for outsourcing web-based customer/user facing services while keeping in house IT infrastructure and a set of resources to support internal/office operation and store critical data required for local company's services. It allows the enterprise to: automatically scale up and down cloud based resources that have cyclical or random demand increase, paying only for what you need and when you need; simplify geographical expansion and geographical load distribution; avoid building expensive physical private webserver and application servers, also projected for peak demand periods; maintain enterprise control and security over cloud resources using VPN access to cloud resources

Test and Development, Proof of Concept

A second illustrative use case is the one for Test and Development, Proof of Concept. Using the cloud Eliminates need for creating test environment which may require complex infrastructure for testing new products under stress load It is a Cheap alternative for proof-of-concept services try out without capital investment

Application development and testing

Another excellent use case for the Cloud is to utilize it for Application development and testing. In the Company's IT infrastructure there is the Integrated Development Environment (IDevE), Quality Assurance (QA), Performance testing (PERF), and the Developers. In the Cloud based facilities one places a Replicated IDevE, QA, and PERF environment, along with Web servers, Apps servers, Database/Storage, and test/load generator. The company and the Cloud are connected via a Secure Virtual Private Network (VPN) or Virtual Private Cloud (VPC) connection.

Each time an expanded development and testing capability is needed, the cloud environment is spun up. Moving application development and testing to clouds allows for testing of close to real new applications infrastructure and load without creating expensive in-house test environment. Enterprise IT infrastructure needs only resources to support development with limited local test functionalities. Non-production test environment is moved from production facilities to clouds. Provide development and QA teams with on-demand close to real testing environment using full scale load. Use cloud platform based monitoring and performance measurement tools. Empower company to respond quickly and cost-effectively to market demands for test, trial and proof-of-concept activities.

Lab on-demand in cloud

As an extension to the Application development and testing Use Case, is to have a full Lab on-demand in the Cloud. It is nearly identical to the previous use case but in a more general way organizations can put a variety of test and measurement gear up in the cloud.
Suitable for both development lab and educational lab, fully functional virtual lab facilities, on-demand provisioning, adopted to cyclical educational process, preset configuration, easy user management.

Computer Cluster on-demand

A final example use case is that of a Computer Cluster on-demand. This use case addresses enterprises need for Big Data analytics capabilities, makes use of cloud based Data Analytics facilities affordable for SME, solves Big Data storage issues, moves computational workflow/jobs management.

In the Company's IT infrastructure there are the Data repositories, Data servers, Data management system, Analytics capabilities repository , and Data visualization facilities. In the Cloud based facilities there are Data servers (SQL, NoSQL, etc), Data Analytics tools, High

Performance Computing (HPC) cluster, and Scientific/Business Workflow Management. A High Performance Optical Network for fast data upload and download is used.
This allows one to:

- Build proprietary computer cluster using general cloud compute an storage resources or use cloud based Data Analytics services
- Makes HPC and Big Data analytics services accessible for SME
- Reduce capex for big companies by keeping limited data analytics capabilities in-house outsourcing peak computation job to cloud
- Move computational HPC workflow/jobs to cloud using the same scientific/business workflow management/monitoring system/tools Requires high-performance optical network, typically dedicated.

Wrap up and Take away

Clouds provide 3 basic service models and a number of deployment models that reflects the variety of business and service relations Selecting right service and deployment model is an important stage in cloud implementation by enterprise Enterprise may adopt gradual/staged cloud implementation from IaaS to PaaS/SaaS and from private cloud to hybrid and public cloud.

Demand from Big Data and data intensive applications facilitates demand for HPC services as a part of cloud infrastructure Clouds are not designed for HPC but rather as Highly Scalable Computing (HSC) environment Clouds can be effectively used for tasks parallelization (also called “embarrassingly parallel tasks”).