# **Research Agenda in Cloud Computing**

Amarnath Paswan [1], Neelam Gupta[2], Naresh Kumar Marwal [3]

[1] Amarnath Paswan [CSE], JECRC UDML College of Engineering, Kukas , Jaipur Rajasthan, India

[2] Neelam Gupta [CSE], JECRC UDML College of Engineering, Kukas, Jaipur Rajasthan, India

[3] Naresk Kumar Marwal [CSE], JECRC UDML College of Engineering, Kukas , Jaipur Rajasthan, India

Email : [1] <u>paswan.Amarnath@gmail.com</u> [2] <u>neelamgupta.cse@udml.ac.in</u> [3] nareshmarwal.cse@udml.ac.in

## ABSTRACT

Cloud computing is the latest effort in delivering computing resources as a service. It represents a shift away from computing as a product that is purchased, to computing as a service that is delivered to consumers over the internet from large-scale data centres - or "clouds". "Cloud" computing - a relatively recent term, builds on decades of research in virtualization, distributed computing, utility computing, and more recently networking, web and software services. It implies a service oriented architecture, reduced information technology overhead for the end-user, great flexibility, reduced total cost of ownership, ondemand services and many other things. Whilst cloud computing is gaining growing popularity in the IT industry, academia appeared to be lagging behind the rapid developments in this field. This paper is the first systematic review of peer-reviewed academic research published in this field, and aims to provide an overview of the swiftly developing advances in the technical foundations of cloud computing and their research efforts. Structured along the technical aspects on the cloud agenda, we discuss lessons from related technologies; advances in the introduction of protocols, interfaces, and standards; techniques for modelling and building clouds; and new use-cases arising through cloud computing.

#### Keywords:

Cloud computing, cloud technologies, utility computing, end-to-end quality of service.

## **1. INTRODUCTION**

Cloud computing has recently reached popularity and developed into a major trend in IT. While industry has been pushing the Cloud research agenda at high pace, academia has only recently joined, as can be seen through the sharp rise in workshops and conferences focussing on Cloud Computing. Lately, these have brought out many peer-reviewed papers on aspects of cloud computing, and made a systematic review necessary, which analyses the research done and explains the resulting research agenda. We performed such a systematic review of all peer-reviewed academic research on cloud computing, and explain the technical challenges facing in this paper. My paper aims to provide a comprehensive review of the academic research done in cloud computing and to highlight the research agenda academia is pursuing. We are well aware that a survey in such a fast moving field will soon be out of date, but feel such a survey would provide a good base for the 1st ACM Symposium on Cloud Computing to set new work in context with, and that it can act as a resource for researchers new in this area. Research in this field appeared to be split into two distinct viewpoints. One investigates the technical issues that arise when building and providing clouds, and the other looks at implications of cloud computing on enterprises and users. In this paper we discuss the advances and research questions in technical aspects of Cloud Computing, such as protocols, interoperability and techniques for building clouds, while we discuss the research challenges facing enterprise users, such as cost evaluations, legal issues, trust, privacy, security, and the effects of cloud computing on the work of IT departments, elsewhere .

## **2. DEFINITIONS**

There has been much discussion in industry as to what cloud computing actually means. The term cloud computing seems to originate from computer network diagrams that represent the internet as a cloud. Most of the major IT companies and market research firms such as IBM , Sun Microsystems , Gartner and Forrester Research have produced whitepapers that attempt to define the meaning of this term. These discussions are mostly coming to an end and a common definition is starting to emerge. The US National Institute of Standards and Technology (NIST) has developed a working definition that covers the commonly agreed aspects of cloud computing. The NIST working definition summarises cloud computing as:

**97 |** Page

ww.ijltemas.in

a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.

This definition describes cloud computing as having five essential characteristics, three service models, and four deployment models. The essential characteristics are:-

□ Ondemand self-service: computing resources can be acquired and used at anytime without the need for human interaction with cloud service providers. Computing resources include processing power, storage, virtual machines etc.

□ Broad network access: the previously mentioned resources can be accessed over a network using heterogeneous devices such as laptops or mobiles phones.

□ Resource pooling: cloud service providers pool their resources that are then shared by multiple users. This is referred to as multi-tenancy where for example a physical server may host several virtual machines belonging to different users.

□ Rapid elasticity: a user can quickly acquire more resources from the cloud by scaling out. They can scale back in by releasing those resources once they are no longer required.

□ Measured service: resource usage is metered using appropriate metrics such monitoring storage usage, CPU hours, bandwidth usage etc.

The above characteristics apply to all clouds but each cloud provides users with services at a different level of abstraction, which is referred to as a service model in the NIST definition. The three most common service models are:-

□ Software as a Service (SaaS): this is where users simply make use of a web-browser to access software that others have developed and offer as a service over the web. At the SaaS level, users do not have control or access to the underlying infrastructure being used to host the software. Salesforce's Customer Relationship Management software3 and Google Docs4 are popular examples that use the SaaS model of cloud computing.

□ Platform as a Service (PaaS): this is where applications are developed using a set of programming languages and tools that are supported by the PaaS provider. PaaS provides users with a high level of abstraction that allows them to focus on developing their applications and not worry about the underlying infrastructure. Just like the SaaS model, users do not have control or access to the underlying infrastructure being used to host their applications at the PaaS level. Google App Engine5 and Microsoft Azure are popular PaaS examples. □ Infrastructure as a Service (IaaS): this is where users acquire computing resources such as processing power, memory and storage from an IaaS provider and use the resources to deploy and run their applications. In contrast to the PaaS model, the IaaS model is a low level of abstraction that allows users to access the underlying infrastructure through the use of virtual machines. IaaS gives users more flexibility than PaaS as it allows the user to deploy any software stack on top of the operating system. However, flexibility comes with a cost and users are responsible for updating and patching the operating system at the IaaS level. Amazon Web Services' EC2 and S37 are popular IaaS examples.

The service models described in the NIST definition are deployed in clouds, but there are different types of clouds depending on who owns and uses them. This is referred to as a cloud deployment model in the NIST definition and the four common models are:

 $\Box$  Private cloud: a cloud that is used exclusively by one organisation. The cloud may be operated by the organisation itself or a third party.

■ Public cloud: a cloud that can be used (for a fee) by the general public. Public clouds require significant investment and are usually owned by large corporations such as Microsoft, Google or Amazon.

□ Community cloud: a cloud that is shared by several organisations and is usually setup for their specific requirements. The Open Cirrus cloud testbed could be regarded as a community cloud that aims to support research in cloud computing.

□ Hybrid cloud: a cloud that is setup using a mixture of the above three deployment models. Each cloud in a hybrid cloud could be independently managed but applications and data would be allowed to move across the hybrid cloud. Hybrid clouds allow cloud bursting to take place, which is where a private cloud can burst-out to a public cloud when it requires more resources.

Figure 1 provides an overview of the common deployment and service models in cloud computing, where the three service models could be deployed on top of any of the four deployment models.



ww.ijltemas.in

Figure 2: Cloud computing deployment and service models

#### **3. STANDARDS AND INTERFACES**

Cloud computing seeks to be a utility delivered in a similar as way electricity is delivered. Due to the higher complexity involved in delivering IT resources, open standards are necessary that enable an open market of providing and consuming resources. Currently, each vendor develops its own solution and avoids too much openness, to tie consumers in to their services and make it hard for them to switch to competitors. However, to new adopters the fear of vendor lock-in presents a barrier to cloud adoption and increases the required trust. There are three groups currently working on standards for cloud computing: The Cloud Computing Interoperability Forum, the Open Cloud Consortium10, and the DMTF Open Cloud Standards Incubator11. There is also a document called the open cloud manifesto12, in which various stakeholders express why open standards will benefit cloud computing. In literature, Grossman [2009] points out that the current state of standards and interoperability in cloud computing is similar to the early Internet era where each organization had its own network and data transfer was difficult.

Keahey et al. looked into the difficulties of developing standards and summarised the main goals of achieving interoperability between different IaaS providers as being machine-image compatibility, contextualization compatibility and API-level compatibility. Image compatibility is an issue as there are multiple incompatible virtualisation implementations such as the Xen, KVM, and VMWare hypervisors. When users want to move entire VMs between different IaaS providers, from the technological point of view this can only work when both providers use the same form of Contextualization virtualisation. compatibility problems exist because different IaaS providers use different methods of customizing the context of VMs, for example setting the operating system's username and password for access after deployment must be done in different ways. Finally, there are no widely agreed APIs between different IaaS providers that can be used to manage virtual infrastructures and access VMs. For machine image or VM compatibility there is an ongoing attempt to create an open standard called the Open Virtual Machine Format (OVF).

Cloud computing can benefit from standardised API interfaces as generic tools that manage cloud infrastructures can be developed for all offerings. For IaaS there are developments towards standards and Eucalyptus is looking to become the de-facto standard. For PaaS and SaaS stakeholders need to join the standardisation groups to work towards it. Achieving standardised APIs appears to be rather politically than technically challenging, hence there seems to be little space for academic involvement. However, standardised interfaces alone do not suffice to prevent vendor lock-in. For an open cloud, there is a need for protocols and software artefacts that allow interoperability to unlock more of the potential benefits from cloud computing.

#### **4. BUILDING CLOUDS**

In this section we describe work that helps building cloud offerings. This requires management software, hardware provision, simulators to evaluate the design, and evaluating management choices.

Sotomayor et al. presents two tools for managing cloud infrastructures: OpenNebula, a virtual infrastructure manager, and Haizea, a resource lease manager.. Haizea can act as a scheduling backend for OpenNebula, and together they advance other virtual infrastructure managers by giving the functionality to scale out to external clouds, and providing support for scheduling groups of VMs, such that either the entire group of VMs are provided resources or no member of the group. In combination they can provide resources by best-effort, as done by Amazon EC2, by immediate provision, as done by Eucalyptus, and in addition using advance reservations.

Sriram discusses some of the issues with scaling the size of data centres used to provide cloud computing services. He presents the development and initial results of a simulation tool for predicting the performance of cloud computing data centres which incorporates normal failures, failures that occur frequently due to the sheer number of components and the expected average lifecycle of each component and that are treated as the normal case rather than as an exception. Sriram shows that for small data centres and small failure rates the middleware protocol does not play a role, but for large data centres distributed middleware protocols scale better. CloudSim, another modelling and simulation toolkit has been proposed by Buyya et al. CloudSim simulates the performance of consumer applications executed in the cloud. The topology contains a resource broker and the data centres where the application is executed. The simulator can then estimate the performance overhead of the cloud solution.

In summary, several projects research into the way future clouds can be built. Given the methodology we chose earlier, the papers discussed in this section differ too much to conclude with a single research direction in which academia is heading when looking into building future clouds. In fact, it seems there are many more research directions we will be facing when it comes to building new cloud facilities.

## 5. NEW USE CASES IN CLOUD COMPUTING

In this paper we have so far presented work that seeks to advance the technology of cloud computing. We end this by presenting new technologies and use cases that become possible through the use of cloud computing. Chun and Maniatis describe one such use-case, where cloud computing enables a technology which otherwise would not be possible: to overcome hardware limitations and enable more powerful applications on smartphones, they use external resources. This is done by partially off-loading execution from the smartphone and using cloud resources. But, Chun and Maniatis also include laptops or desktops near the phone in their "cloud" because of the network latency for phones. Depending on the use case, their model offloads entire computations or parts thereof, and only has the remainder executed locally.

Another use-case that becomes feasible and affordable through the use of cloud computing is large-scale nonfunctional requirements testing, as described by Ganon and Zilbershtein. They tested Network Management Systems for systems where much of the functionality is in the endpoints, such as in voice over IP software. They discuss the advantages and disadvantages of cloud-based testing over testing against real elements or a simulator, and describe how a cloud based test setup can be created using agents that are deployed into the cloud and with the use of cloud elasticity. Further, implications of using the cloud for this setup are evaluated, such as security, safety of intellectual property or software export restrictions, and solutions to tasks such as creating setups that emulate problems including noisy or delayed network connections are presented. Ganon and Zilbershtein reach the conclusion that there are significant benefits of using cloud-based testing, although it cannot completely replace traditional testing against real managed endpoints. They round up their insightful paper with a use case of a test scenario that was carried out on Amazon's cloud and resulted in improvements to the software that could not have been highlighted with other feasible forms of testing, and with disclosing the costs occurred to carry out the cloud based test. Matthew and Spraetz also looked at testing in cloud computing. They explained an effort to automate testing for SaaS providers along the example of Salesforces' Apex test framework. Because consumers can use the Force.com PaaS offering to customise their business solutions into the CRM system using an entire Java-like programming language, it becomes unfeasible to test all possible states of the CRM beforehand. Instead, a test framework is provided, that allows users to specify regression tests, which can be carried out before every update to the SaaS offering. This is crucial because in a SaaS world there is no choice of version. Once an update is rolled out it is effective for all users.

In a cloud that offers IaaS, the number of VMs and thus instances of operating systems that need to be managed increases significantly. To avoid having to deploy software and updates into each virtual machine, and to avoid lengthy installation processes, entire so called "virtual appliances" will be managed. This means, in cloud computing the operating system will no longer be viewed separated from the applications deployed, but rather both will be deployed and maintained jointly. For service providers this means, they now have the ability to offer a virtual appliance, as functional disc image, instead of having to create lengthy installation procedures to guarantee compatibility with other applications in the VM. Wilson describes Coronary, a software configuration management tool for virtual appliances. Coronary takes the idea of incremental updates from configuration management software such as CVS or subversion, and uses this technology to manage virtual appliances over their lifecycle. Wilson discusses the new requirements of version control when used for virtual appliances, and how Coronary handles them.

# 6. CONCLUSION

This paper has presented the work published by the academic community advancing the technology of cloud computing. Much of the work has focussed on creating standards and allowing interoperability, and describes ways of designing and building clouds.

Various definitions of cloud computing were discussed and the NIST working definition by Mell and Grance was found to be the most useful as it described cloud computing using a number of characteristics, service models and deployment models. The socio-technical aspects of cloud computing that were reviewed included the costs of using and building clouds, the security, legal and privacy implications that cloud computing raises as well as the effects of cloud computing on the work of IT departments. The technological aspects that were reviewed included standards, cloud interoperability, lessons from related technologies, building clouds, and use-cases that presented new technological possibilities enabled by the cloud.

To conclude, this paper discussed the research academia has pursued to advance the technological aspects of cloud computing, and highlighted the resulting directions of research facing the academic community. This paper reviewed the technical aspects of research in cloud computing.

#### REFERENCES

[1] Sun Microsystems, Introduction to Cloud Computing Architecture, 2009

[2] KHAJEH-HOSSEINI, A., SOMMERVILLE, I. and SRIRAM, I. "Research Challenges for Enterprise Cloud Computing," (unpublished).( Submitted to 1st

ACM Symposium on Cloud Computing, Indianapolis, Indiana, USA, June 2010, under paper id 54)

[3] IBM, Staying aloft in tough times, 2009

[4] PLUMMER, D.C., BITTMAN, T.J., AUSTIN, T.,

CEARLEY, D.W., and SMITH D.M., Cloud Computing: Defining and Describing an Emerging Phenomenon, 2008

[5] STATEN, J., Is Cloud Computing Ready For The Enterprise?, 2008.

[6] MELL, P. and GRANCE, T. 2009. Draft NIST Working Definition of Cloud Computing.

[7] Open CirrusTM: the HP/Intel/Yahoo! Open Cloud Computing Research Testbed. Retrieved December 1, 2009, from https://opencirrus.org/

[8] KEAHEY, K., TSUGAWA, M., MATSUNAGA, A. and FORTES, J. 2009. Sky Computing. Internet Computing, IEEE 13, 5, 43-51.

[9] SOTOMAYOR, B., RUBé, A., LLORENTE, I. and FOSTER, I. 2009. Virtual Infrastructure Management in Private and Hybrid Clouds. Internet Computing, IEEE 13, 5, 14-22.

[10] SRIRAM, I. 2009. A Simulation Tool Exploring Cloud-Scale Data Centres. In 1st International Conference on Cloud Computing (CloudCom 2009), pp. 381-392.

[11] BUYYA, R., RANJAN, R. and CALHEIROS, R. N. 2009. Modelling and simulation of scalable Cloud computing environments and the Cloud-Sim toolkit: Challenges and opportunities. In High Performance Computing & Simulation, 2009. HPCS'09. International Conference on, 1-11.

[12] CHUN, B.-G. and MANIATIS, P. 2009. Augmented Smart Phone Applications Through Clone Cloud Execution. In Proceedings of the 12th Workshop on Hot Topics in Operating Systems (HotOS XII).

[13] GANON, Z. and ZILBERSHTEIN, I. E. 2009. Cloud-based Performance Testing of Network Management Systems. In Computer Aided Modelling and Design of Communication Links and Networks, 2009. CAMAD '09. IEEE 14th International Workshop on, 1-6.

[14] MATHEW, R. and SPRAETZ, R. 2009. Test Automation on a SaaS Platform. In Software Testing Verification and Validation, 2009. ICST '09. International Conference on, 317-325.

[15]WILSON, M. 2009. Constructing and Managing Appliances for Cloud Deployments from Repositories of Reusable Components.