International Journal of Engineering and Technical Research (IJETR)
ISSN: 2321-0869, Volume-2, Issue-5, May 2014

Agenda in Cloud Technologies

Ujwal Bhatia, Sukhwinder Singh

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". 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.**

Categories and Subject Descriptors

A.1 [General Literature]: Introductory and Survey C.2.4 [Computer Communication Networks]: Distributed Systems – ***Cloud Computing***

General Terms

Management, Measurement, Performance, Design, Economics, Reliability, Experimentation, Standardization

Index Terms— **Cloud computing, cloud technologies, review**

I. 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.

There were several whitepapers and general introductions to cloud computing, which provide an overview of the field, [e.g. 1, 2, 3, 4, 5], but yet there is no systematic review of the agenda academia has taken. Pastaki Rad *et al.* [6] presented a preliminary survey that included a short overview of storage systems and Infrastructure as a Service (IaaS), which, however, was not systematic and fell short of providing a good overview of the state-of-the-art and lacked a discussion of the research challenges. Our 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 [7]. This paper is structured as follows: the methodology used to carry out this review is shown in the Section 2; Section 3 discusses various definitions of cloud computing; Section 4 outlines the lessons to be learnt from related areas; Section 5 and Section 6 review the work on standardised interfaces and Cloud interoperability respectively; Section 7 summarises various other research done in support of building Cloud infrastructures; while use cases of Cloud computing are reviewed in Section 8; finally Section 9 concludes the review by summing up the research directions academia faces.

II. Methodology

This review surveyed the existing literature using a principled and systematic approach: we searched each of the major research databases for computer science, the ACM Digital Library, IEEE Xplore, SpringerLink, ScienceDirect and Google Scholar, for the following keywords: cloud computing, elastic computing, utility computing, Infrastructure as a Service, IaaS, Platform as a Service, PaaS, Software as a Service, SaaS, Everything as a Service, XaaS. The date range for this search was limited from 2005 until October 2009. This date range was chosen because this survey work was commenced in October 2009, and because all public clouds were launched after 2005. For example, Amazon first launched EC2 (Elastic Compute Cloud) in August 20061 and Google launched App Engine in April 20082. According to Google Trends, the term *cloudcomputing* started becoming popular in 2007 as shown in Figure 1.

Manuscript received May 12, 2014.

Ujwal Bhatia, Student, Electronics and Communication Engineering Department, PEC University of Technology, Chandigarh, India

Sukhwinder Singh, Mentor, Assistant Professor, Electronics and Communication Engineering Department, PEC University of Technology, Chandigarh, India

The searches from the five target databases returned over 150 papers. The titles and abstracts of these papers were read and for quality reasons we decided to use only peer-reviewed papers for the review; only a small number of non peer-reviewed publications were included, such as well quoted definitions or a summary of a workshop discussing research challenges academia is facing, as these were relevant and not matched by comparable peer-reviewed work. Furthermore, papers that had misleading titles or abstracts and those that were purely focused on High Performance Computing and e-Science were also left out of the review as these areas are not within the core focus of our review. The citation-references of the selected papers were checked but no additional papers were found to be necessary to add to this review based on the criteria mentioned above. This resulted in a total of 56 publications being selected for review. The papers were split into three categories based on their main focus; the categories were: general introductions, technological aspects of cloud computing and organizational aspects. The latter category is discussed elsewhere [7]. The papers that provided general introductions to cloud computing are referenced throughout this paper. The technological category was further broken down into papers that dealt with protocols, interfaces, standards, lessons from related technologies, techniques for modelling and building clouds, and new use-cases arising through cloud computing.. Table 1 provides an overview of the papers reviewed in this review and their categories. As it can be seen in the table, the majority of the

Figure 1: Searches for "cloud computing" on Google.com, taken from Google Trends3.

III. Lessons from Related Technologies

The remainder of this paper reviews the research that describes technological aspects of research in cloud computing. This starts with a look at lessons to be learnt from related fields of research. In the following, standards and interfaces in cloud computing as well as interoperability between different cloud systems are explained. Then, techniques for designing and building clouds are summarised, which include advances in management software, hardware provisioning, and simulators that have been developed to evaluate design decisions and cloud management choices. This is rounded up by presenting new use-cases that have become possible through cloud computing.

Voas and Zhang [20] identified cloud computing as the next computing paradigm that follows on from mainframes, PCs, networked computing, the internet and grid computing. These developments are likely to have similarly profound effects as the move from mainframes to PCs had on the ways in which software was developed and deployed. One of the reasons that prevented grid computing from being widely used was the lack of virtualization that resulted in jobs being dependant on the underlying infrastructure. This often resulted in unnecessary complexity that had an effect on wider adoption [21]. Ian Foster – who was one of the pioneers of grid computing – compared cloud computing with grid computing and concluded that although the details and technologies of the two are different, their vision is essentially the same [22]. This vision is to provide computing as a utility in the same way that other public utilities such as gas and electricity are provided. In fact the dream of utility computing has been around since the 1960s and advocated by the likes of John McCarthy and Douglas Parkhill. For example, the influential mainframe operating system *Multics* had a number of design goals that are remarkably similar to the aims of current cloud computing providers. These design goals included remote terminal access, continuous operational provision (inspired by electricity and telephone services), scalability, reliable file systems that users trust to store their only copy of files, information sharing controls, and an ability to support different programming environments [23]. Therefore it is unsurprising that many people compare cloud computing to mainframe computing. However, it should be noted that although many of the ideas are the same, the user experience of cloud computing is almost completely the opposite of mainframe computing. Mainframe computing limited people's freedom by restricting them to a very rigid environment; cloud computing expands their freedom by giving them access to a variety of resources and services in a self-service manner.
Foster *et al.* [22] compare and contrast cloud computing with grid computing. They believe cloud computing is an evolved version of grid computing, in such a way that it answers the new requirements of today's time, takes into account the expensiveness of running clusters, and the existence of low-cost virtualisation. IT has greatly evolved in the last 15 years since grid computing was invented, and at present it is on a much larger scale that enables fundamentally different approaches. Foster *et al.* see similarities between the two concepts in their vision and architecture, see a relation between the concepts in some fields as in the programming model ("MapReduce is only yet another parallel programming model") and application model (but clouds are not appropriate for HPC applications that require special interconnects for efficient multi-core scaling), and they explain fundamental differences in the business model, security, resource management, and abstractions. Foster *et al.* find that in many of these fields there is scope for both the cloud and grid research communities to learn from each other's findings, and highlight the need for open protocols in the cloud, something grid computing adopted in its early days. Finally, Foster *et al.* believe that neither the electric nor computing grid of the future will look like the traditional electric grid. Instead, for both grids they see a mix of

micro-productions (alternative energy or grid computing) and large utilities (large power plants or data centres).

In Market-Oriented Cloud Computing, a follow-on work from their Market-Oriented Grid Computing and Market-Oriented Utility Computing papers, Buyya *et al.* [24] describe their work on market oriented resource allocation and their Aneka resource broker: In the case of limited availability of resources, not all service requests will be of equal importance, and a resource broker will regulate the supply and demand of resources at market equilibrium. A batch job for example might be preferably processed when the resource value is low, while a critical live service request would need to be processed at any price. Aneka, commercialised through Manjrasoft, is a servicebroker that mediates between consumers and providers by buying capacities from the provider and subleasing them to the consumers. However, such resource trading requires the availability of ubiquitous cloud platforms with limited resources, and is in contrast to the desire for simple pricing models.

As cloud computing delivers IT as a service, cloud researchers can also learn from service oriented architecture (SOA). In fact, the first paper that introduced PaaS [25] described PaaS as an artefact of combining infrastructure provisioning with the principles of SaaS and SOA. Since then, no academic work has been published in the field of PaaS. We have to take our to-date understanding of PaaS from the current developments in industry, in particular from the two major vendors, Force.com and from Google App Engine. Sedayao [26] built a monitoring tool using SOA services and principles, and describe their experience from building a robust distributed application consisting of unreliable parts and the implication for cloud computing. As design goal for distributed computing scenarios such as cloud computing they propose, “like routers in a network, any service using other cloud services needs to validate input and have hold down periods before determining that a service is down”[26]. Zhang and Zhou [27] analyse convergence from SOA and virtualisation for cloud computing and present seven architectural principles and derive ten interconnected architectural modules. These build the foundation for their IBM cloud usage model, which is proposed as Cloud Computing Open Architecture (CCOA). Vouk [21] described cloud computing from a SOA perspective and talked about the Virtual Computing Laboratory (VCL) as an implementation of a cloud. VCL is an "open source implementation of a secure production-level on-demand utility and service oriented technology for wide-area access to solutions based on virtualised resources, including computational, storage and software resources" [21]. In this respect, VCL could be categorised as an IaaS layer service.

Napper and Bientinesi [28] ran an experiment to compare the potential performance of Amazon’s cloud computing with the performance of the most powerful, purpose build, high performance computers (HPC) in the Top500 list in terms of solving scientific calculations using the LINPACK benchmark. They found that the performance of individual nodes in the cloud is similar to those in HPC, but that there is a severe loss in performance when using multiple nodes, although the used benchmark was expected to scale linearly. The AMD instances scaled significantly better than the Intel instances, but the cost for the computations were equivalent with both types. As the performance achieved decreased exponentially in the cloud and only linearly in HPC systems, Napper and Bientinesi [28] conclude that despite the vast availability of resources in cloud computing, these offerings are not able to compete with the supercomputers in the Top500 list for scientific computations.

In a non peer-reviewed summary of keynote speeches for a workshop on distributed systems Birman *et al.* [29] express that the distributed systems research agenda is quite different to the cloud agenda. They argue that while technologies from distributed systems are relevant for cloud computing, they are no longer central aspects of research. As example they list strong synchronisation and consistency as ongoing research topics from distributed systems. In cloud computing they remain relevant, but as the overarching design goal in the cloud is scalability, the search is now for decoupling and thus avoiding synchronisation, rather than improving synchronisation technologies. Birman *et al.* [29] come to a cloud research agenda comprising four directions: managing the existing compute power and the loads present in the data centre; developing stabile large-scale event notification platforms and management technologies; improving virtualisation technology; and understanding how to work efficiently with a large number of low-end and faulty components.

Cloud computing has been compared to several related fields of research. This section has shown that the cloud computing research agenda differs from the agenda in related fields, but that there are several findings in related research communities the research community can benefit from. We have also seen, that practitioners in distributed computing, grid computing, and SOA have joined the cloud community and proposed goals for research based on the background of their field. In the following, we shall look at the research more from the point of view of the cloud agenda.

IV. 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 Forum9, 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. This changed with the introduction of TCP and other Internet standards. However, these standards were initially resisted by vendors just as standardisation attempts in cloud computing are being resisted by some vendors.

Keahey *et al.* [30] 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 virtualisation. Contextualization compatibility problems exist 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). At the API-level, for PaaS AppScale13, an open source effort to re-implement the interfaces of Google App Engine, is aiming to become a standard, and for IaaS management, Amazon EC2's APIs are quickly becoming a *de-facto* standard, popularised through their open source re-implementation Eucalyptus.

13 http://code.google.com/p/appscale

14 http://www.linux-kvm.org

15 http://www.flexiscale.com

16 http://www.newservers.com

Eucalyptus is an open-source software package that can be used to build IaaS clouds from computer clusters [31]. Eucalyptus emulates the proprietary Amazon EC2 SOAP and Query interface, and thus an IaaS infrastructure set up using Eucalyptus can be controlled with the same tools and software that is used for EC2. The open source nature of Eucalyptus gives the community a useful research tool to experiment with IaaS provisioning. The initial version of Eucalyptus used Xen as hypervisor for virtual machines, but since the publication of that version, support for further hypervisors has been added, in particular for the newly popular KVM hypervisor14. Eucalyptus has a hierarchical design that makes it reasonably easy to predict its performance. However, for very large data centres this centralised design might not scale particularly well, hence Nurmi *et al.* recommend it for typical settings in present in academia. Although Eucalyptus just re-implemented the Amazon EC2 interfaces, to date it is one of the most fundamental contributions by the research community towards standards in cloud computing, although only a few other providers use these interface APIs yet. But, for reasons such as fault tolerance or performance, or freedom from lock-in, consumers may wish to use multiple cloud providers. In the absence of open standards, or when attempts at providing open interface standards like Eucalyptus are not followed by some providers, there will be heterogeneous interfaces. Dodda *et al.* [32] address the problem of managing cloud resources with such heterogeneous access, by proposing a generic interface to the specific interface presented by individual cloud providers. They use their interface to an interface to compare the performance of Amazon EC2's Query and SOAP interface, and find that the average response time for the SOAP interface was nearly double that of the Query interface. These results emphasise the importance of selecting the interface through which resources from a given provider are managed. In a similar effort, Harmer *et al.* [33] present a cloud resource interface that hides the details of individual APIs to allow provider agnostic resource usage. They present the interface to create a new instance at Amazon EC2, at Flexiscale15, and at a provider of on-demand non-virtualised servers called NewServers16, and implemented an abstraction layer for these APIs. The solution from Harmer *et al.* goes beyond hiding API details and contains functionality to
compensate for loss of core infrastructure in scenarios where multiple providers are used.

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. This technically rich direction will be discussed in the following section.

V. CLOUD INTEROPERABILITY AND NOVEL PROTOCOLS

The next steps from compatible and standardised interfaces towards utility provisioning are universal open and standard protocols that allow interoperability between clouds and enable the use of different offerings for different use cases. Bernstein *et al.* [34] describe an in-depth overview of the technological research agenda and open questions for interoperability in the cloud. They are looking for ways of allowing cloud services to interoperate with other clouds and highlight many goals and challenges, such as that cloud services should be able to implicitly use others through some form of library without the need to explicitly reference them, e.g. with their domain name and port. The collection of protocols inside and in-between the clouds that solve interoperability in the cloud are termed *intercloud protocols*. The intercloud protocol research agenda is made up of several areas: addressing, naming identity and trust, presence and messaging, virtual machines, multicast, time synchronisation, and reliable application transport. For cloud computing, each of these areas contains several issues. In addressing for example, the research problem is that there is the limited address space in IPv4 and that its successor IPv6 might be an inappropriate approach in a large and highly virtualised environment, as the cloud, due to its static addressing scheme: Bernstein *et al* criticise that IP addresses traditionally embody network locations for routing purposes and identity information, but in the cloud context identifiers should allow the objects to move into different subnets dynamically. This problem of static addresses is addressed by Ohlman *et al.* [35]. They recommend the usage of Networking of Information (NetInf) for cloud computing systems. Unlike URLs which are location-dependent, NetInf uses a location-independent model of naming objects, and offers an API that hides the dynamics of object locations and network topologies. Ohlman *et al.* demonstrate how this can ease management in the cloud, where the design desires transparency of location.

VI. 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. We were surprised so far not to see significant contributions to the usage and scaling properties of Hadoop/MapReduce, which is a new programming paradigm in the cloud. Similarly, there was no work published yet on effective usage of PaaS offerings such as Google Apps.

Various definitions of cloud computing were discussed and the NIST working definition by Mell and Grance [11] 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.

A number of authors have discussed the new research challenges that are raised by cloud computing. Bernstein *et al.* [34] listed a research agenda and open questions to achieve interoperability, and Birman *et al.* [29] described a research agenda that seeks to facilitate industry in building successful clouds. Vouk [21] described the problems of managing virtual machine (VM) images. It would be difficult to manually update a large number of VM images and verify their integrity by checking their contents. Mei *et al.* [51] compared the input-output, storage and processing features of cloud computing with pervasive computing and service computing to highlight new research challenges. Cloud computing could benefit from the functionality modelling issues studied in service computing, and the context-sensitivity issues studied in pervasive computing [51]. However, it is difficult to talk about cloud computing without having a particular abstraction layer in mind. The comparisons done by Mei *et al.* are reasonable at an IaaS layer, but they are not very meaningful at the SaaS layer where storage and processing features might not be visible at all. Youseff *et al.* [16] briefly discussed the research challenges in IaaS clouds mentioning that system monitoring information could be used for application optimization in clouds. However, making such information available to users in a useful manner is a challenge [16]. Armbrust *et al.* [18] looked at other research challenges in cloud computing. They highlighted ten obstacles in cloud computing that included technical challenges relating to the adoption of cloud computing, such as availability of service and data lock-in. The lack of scalable storage, performance unpredictability and data transfer bottlenecks are also obstacles that could limit the growth of cloud computing. These obstacles present a number of new research opportunities in cloud computing and Armbrust *et al.* provided some ideas of how these obstacles could be tackled.

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. In this way the various projects were set in context, and the research agenda followed by and facing academia was presented. The review showed that there are several ways in which the cloud research community can learn from related communities, and has shown there is interest in academia for describing these similarities. Further, there have been attempts at building unified APIs to access clouds which seem to be more politically than technically challenging. Then, the perhaps clearest research agenda was presented towards interoperability in the cloud and the challenges that need to be overcome. Finally, both for building clouds and presenting use cases in the cloud, the research efforts were shown to be very diverse, making it hard to suggest in which way academia will be moving. This paper reviewed the technical aspects of research in cloud computing. Together with [7], which discussed the work on implications of cloud computing on enterprises and users, this forms a complete survey of all research published on Cloud Computing, providing a solid basis for the 1st ACM Symposium on Cloud Computing.

ACKNOWLEDGMENTS

I would like to make my acknowledgements to Professor Sukhwinder Singh, PEC University of Technology, Chandigarh, India. for his guidance and giving me a better understanding of modern cloud techniques.

REFERENCES

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

[2] FELLOWS, W.2008. Partly Cloudy, Blue-Sky Thinking About Cloud Computing. 451 Group.

[3] VARIA, J.2009. *Cloud Architectures*. Amazon Web Services.

[4] CHAPPELL, D.2009. *Introducing the Azure Services Platform*. David Chappell & Associates.

[5] RAYPORT, J. F. and HEYWARD, A.2009. *Envisioning the Cloud: The Next Computing Paradigm*. Marketspace.

[6] PASTAKI RAD, M., SAJEDI BADASHIAN, A., MEYDANIPOUR, G., ASHURZAD DELCHEH, M., ALIPOUR, M. and AFZALI, H. 2009. A Survey of Cloud Platforms and Their Future.

[7] 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)

[8] IBM, *Staying aloft in tough times*, 2009

[9] PLUMMER, D.C., BITTMAN , T.J., AUSTIN, T., CEARLEY, D.W., and SMITH D.M., *Cloud Computing: Defining and Describing an Emerging Phenomenon* , 2008

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

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

[12] ERDOGMUS, H. 2009. Cloud Computing: Does Nirvana Hide behind the Nebula? *Software, IEEE* 26, 2, 4-6.

[13] LEMOS, R. 2009. Inside One Firm's Private Cloud Journey. Retrieved December 1, 2009, from http://www.cio.com/article/506114/Inside_One_Firm_s_Private_Cloud_Journey

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

[15] VAQUERO, L., MERINO, L., CACERES, J. and LINDNER, M. 2009. A break in the clouds: towards a cloud definition. *SIGCOMM Comput. Commun. Rev*. 39, 1, 50-55.

[16] YOUSEFF, L., BUTRICO, M. and DA SILVA, D. 2008. Toward a Unified Ontology of Cloud Computing. In *Grid Computing Environments Workshop, 2008. GCE '08*, 1-10.

[17] J. Appavoo, V. Uhlig, and A. Waterland, "Project Kittyhawk: Building a Global-Scale Computer," vol. 42, 2008, pp. pp. 77-84.

[18] M. Armbrust, A. Fox, R. Griffith, A. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, and M. Zaharia, *Above the Clouds: A Berkeley View of Cloud Computing*, 2009.

[19] WEINMAN, J. 2008. 10 Reasons Why Telcos Will Dominate Enterprise Cloud Computing

[20] VOAS, J. and ZHANG, J. 2009. Cloud Computing: New Wine or Just a New Bottle? *IT Professional* 11, 2, 15-17.

[21] VOUK, M. A. 2008. Cloud computing — Issues, research and implementations. In *Information Technology Interfaces, 2008. ITI 2008. 30th International Conference on*, 31-40.

[22] FOSTER, I., ZHAO, Y., RAICU, I. and LU, S. 2008. Cloud Computing and Grid Computing 360-Degree Compared. In *Grid Computing Environments Workshop (*GCE '08), Austin, Texas, USA, November 2008, 1-10.

[23] Corbató, F. J., Saltzer, J. H., and Clingen, C. T. 1972. Multics: the first seven years. In *Proceedings of the May 16-18, 1972, Spring Joint Computer Conference*, Atlantic City, New Jersey, May 1972, 571-583.

[24] BUYYA, R., YEO, C. and VENUGOPAL, S. 2008. Market-Oriented Cloud Computing: Vision, Hype, and Reality for Delivering IT Services as Computing Utilities. In *High Performance Computing and Communications, 2008. HPCC '08. 10th IEEE International Conference on*, 5-13.

[25] CHANG, M., HE, J., and E. Leon, "Service-Orientation in the Computing Infrastructure," 2006, pp. 27-33.

[26] SEDAYAO, J. 2008. Implementing and operating an internet scale distributed application using service oriented architecture principles and cloud computing infrastructure. In *iiWAS '08: Proceedings of the 10th International Conference on Information Integration and Web-based Applications & Services*, 417-421.

[27] ZHANG, L. and ZHOU, Q. 2009. CCOA: Cloud Computing Open Architecture. In *Web Services, 2009. ICWS 2009. IEEE International Conference on*, 607-616.

[28] NAPPER, J. and BIENTINESI, P. 2009. Can cloud computing reach the top500? In UCHPC-MAW '09: Proceedings of the combined workshops on UnConventional high performance computing workshop plus memory access workshop, 17-20.

[29] Birman, K., Chockler, G., and van Renesse, R. 2009. Toward a cloud computing research agenda. *SIGACT News*, 40, 2, 68-80.

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