Cloud Computing

Nariman Mirzaei (nmirzaei@indiana.edu)

Fall 2008

Abstract

“Cloud” computing – a relatively recent term, defines the paths ahead in computer science world. Being built on decades of research it utilizes all recent achievements in virtualization, distributed computing, utility computing, and networking. It implies a service oriented architecture through offering softwares and platforms as services, reduced information technology overhead for the end-user, great flexibility, reduced total cost of ownership, on demand services and many other things. This paper is a brief survey based of readings on “cloud” computing and it tries to address, related research topics, challenges ahead and possible applications.

1 Introduction

Cloud computing is the next generation in computation. Maybe Clouds can save the world; possibly people can have everything they need on the cloud. Cloud computing is the next natural step in the evolution of on-demand information technology services and products. The Cloud is a metaphor for the Internet, based on how it is depicted in computer network diagrams, and is an abstraction for the complex infrastructure it conceals.

It is a style of computing in which IT-related capabilities are provided “as a service”, allowing users to access technology-enabled services from the Internet (i.e., the Cloud) without knowledge of, expertise with, or control over the technology infrastructure that supports them.

Email was probably the first service on the “cloud”. As the computing industry shifts toward providing Platform as a Service (PaaS) and Software as a Service (SaaS) for consumers and enterprises to access on demand regardless of time and location, there will be an increase in the number of Cloud platforms available.

But it seems that Cloud computing cannot save the universe. Cloud computing cannot run for President. Cloud computing is a very specific type of computing that has very specific benefits. But it has specific negatives as well. And it does not serve the needs of real businesses to hear only the hype about cloud computing – both positive and negative. One thing that is hoped to be accomplished with this paper is not only a clear picture of what the cloud does extremely well and a brief overview of them, but also a short survey on their criteria and challenges ahead of them.

2 Background

2.1 Cyberinfrastructure

“The comprehensive infrastructure needed to capitalize on dramatic advances in information technology has been termed cyberinfrastructure [1]”. The term "cyberinfrastructure" describes the new research environments that support advanced data acquisition, data storage, data management, data integration, data mining, data visualization and other computing and information processing services over the Internet. In scientific usage, cyberinfrastructure is a technological solution to the problem of efficiently connecting data, computers, and people with the goal of enabling derivation of novel scientific theories and knowledge [2].

“Cyberinfrastructure makes applications dramatically easier to develop and deploy, thus expanding the feasible scope of applications possible within budget and organizational constraints, and shifting the scientist’s and engineer’s effort away from information technology development and concentrating it on scientific and engineering research. Cyberinfrastructure also increases efficiency, quality, and reliability by capturing commonalities among application needs, and facilitates the efficient sharing of equipment and services. [3]”

Today, almost any business or major activity uses, or relies in some form, on IT and IT services. These services need to be enabling and appliance-like, and there must be an economy of- scale for the total-cost-of-ownership to be better than it would be without cyberinfrastructure. Technology needs to improve end-user productivity and reduce technology-driven overhead. For example, unless IT is the primary business of an organization, less than 20% of its efforts not directly connected to its primary business should have to do with IT overhead; even though 80% of its business might be conducted using electronic means [4].

2.2 Service Oriented Architecture

SOA is a way of reorganizing a portfolio of previously siloed software applications and support infrastructure into an interconnected set of services, each accessible through standard interfaces and messaging protocols. Once all the elements of an enterprise architecture are in place, existing and future applications can access these services as necessary without the need of convoluted point-to-point solutions based on inscrutable proprietary protocols. This architectural approach is particularly applicable when multiple applications running on varied technologies and platforms need to communicate with each other. In this way, enterprises can mix and match services to perform business transactions with minimal programming effort [5].

Service-oriented architecture offers a way of thinking about IT assets as service components, establishing a software architectural approach to building business applications. The service-oriented architecture approach is based on creating stand-alone, task-specific reusable software components that function and are made available as services.

A service-oriented architecture service exposes a clearly defined activity—like credit card validation—to consuming business applications that might need to perform that function (such as an order processing application). At the core of the service-oriented architecture philosophy is the modularization of business functions for greater flexibility, manageability, and reusability.

2.3 Workflows

A workflow is a depiction of a sequence of operations, declared as work of a person, work of a simple or complex mechanism, work of a group of persons, work of an organization of staff, or machines. Workflow may be seen as any abstraction of real work, segregated in workshare, work split or whatever types of ordering. For control purposes, workflow may be a view on real work under a chosen aspect, thus serving as a virtual representation of actual work. The flow being described often refers to a document that is being transferred from one step to another [6].

A workflow is a model to represent real work for further assessment, e.g., for describing a reliably repeatable sequence of operations. More abstractly, a workflow is a pattern of activity enabled by a systematic organization of resources, defined roles and mass, energy and information flows, into a work process that can be documented and learned. Workflows are designed to achieve processing intents of some sort, such as physical transformation, service provision, or information processing. A workflow can be represented by a directed graph that represents data-flows that connect loosely and tightly coupled (and often asynchronous) processing components. One such graph is shown in Figure 1. It illustrates a Workflow as a part of Experiment Builder in Lead Project [7].

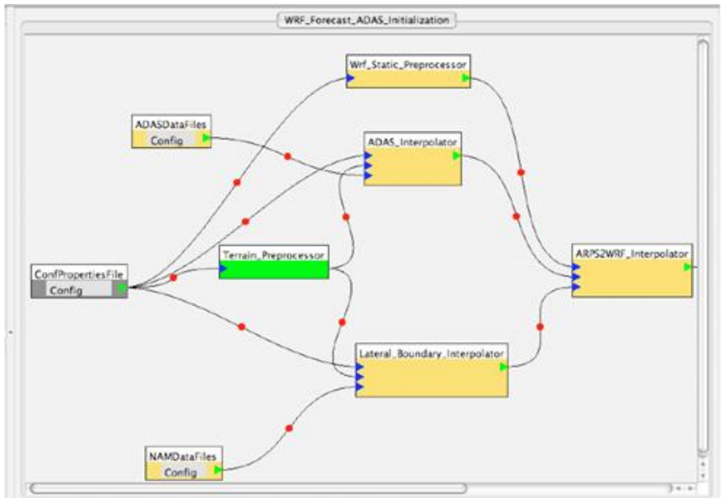

Figure 1

2.4 Virtualization

Virtualization is a framework or methodology of dividing the resources of a computer into multiple execution environments, by applying one or more concepts or technologies such as hardware and software partitioning, time-sharing, partial or complete machine simulation, emulation, quality of service, and many others. It allows abstraction and isolation of lower-level functionalities and underlying hardware. This enables portability of higher-level functions and sharing and/or aggregation of the physical resources [8].

There are lots of virtualization products, and a number of small and large companies that make them. For instance, in the operating systems and software applications space are VMware[1], Xen - an open source Linux-based product developed by XenSource[2], and Microsoft virtualization products, can be mentioned.

3 Cloud Computing

Cloud computing is a paradigm that focuses on sharing data and computations over a scalable network of nodes. Examples of such nodes include end user computers, data centers, and Web Services. We term such a network of nodes as a *cloud*. An application based on such clouds is taken as a *cloud application* [9]. Basically cloud is a metaphor for internet and is an abstraction for the complex infrastructure it conceals. The main idea is to use the existing infrastructure in order to bring all feasible services to the cloud and make it possible to access those services regardless of time and location.

Whether it's called *Cloud Computing* or *On-demand Computing*, *Software as a Service*, or *the Internet as Platform*, the common element is a shift in the geography of computation. When you create a spreadsheet with the Google Docs service, major components of the software reside on unseen computers, whereabouts unknown, possibly scattered across continents.

The shift from locally installed programs to cloud computing is just getting under way in earnest. Shrink-wrap software still dominates the market and is not about to disappear, but the focus of innovation indeed seems to be ascending into the clouds. Some substantial fraction of computing activity is migrating away from the desktop and the corporate server room. The change will affect all levels of the computational ecosystem, from casual user to software developer, IT manager, even hardware manufacturer.

Recently, a lot of vendors have started talking about "cloud computing" in their marketing materials. Merrill Lynch has estimated a $160- billion addressable market opportunity, including $95-billion in business and productivity applications, and another $65-billion in online advertising for Cloud Computing [10]. But the main question is whether the users are ready to give up using services on their local machines and shift to the Cloud since shifting to cloud computing has both advantages and disadvantages for all possible users; nevertheless, they may have different level of importance for different users.

3.1 Pros

- **Reduced Cost:** Cloud technology is paid incrementally (you pay only for what you need), saving organizations money in the short run. Money saved can be used for other important resources.
- **Increased Storage:** Organizations can store more data than on private computer systems.
- **Highly Automated:** IT personnel not needed to keep software up to date as maintenance is the job of the service provider on the cloud.
- **More Mobility:** Employees can access information wherever they are, rather than having to remain at their desks.
- **Allows IT to Shift Focus:** No longer having to worry about constant server updates and other computing issues, government organizations will be free to concentrate on innovation.

3.2 Cons

GNU founder Richard Stallman says that the interesting thing about cloud computing is that we've redefined cloud computing to include everything that we already do. One reason you should not use web applications to do your computing is that you lose control. It's just as bad as using a proprietary program [11]. But certainly shifting to cloud computing has other problems including:

- **Security:** Is there a security standard?
- **Reliance on 3rd Party:** Control over own data is lost in the hands of an "difficult-to-trust" provider
- **Cost of transition:** Is it feasible for me to move from the existing architecture of my data center to the architecture of the cloud?
- **Uncertainty of benefits:** Are there any long term benefits?

4 End Users/Providers

The main end users/providers can be divided to the following major groups.

4.1 Ordinary People

This group of users is just using services from the cloud. They do not care much about high performance and the main problem they may face is having an internet connection all the time and also information privacy. Cloud computing can help this group providing them hardware resources and accessibility through pervasive handhelds with limited resources.

4.2 Academia

Academia is building its own clouds upon the current cyberinfrastructure they have. They are building cloud systems upon on their grid resources (like Teragrid) to resolve Grids limitations. The availability of these large, virtualized pools of compute resources raises the possibility of a new compute paradigm for scientific research with many advantages. For research groups, cloud computing provides convenient access to reliable, high performance clusters and storage, without the need to purchase and maintain sophisticated hardware. For developers, virtualization allows scientific codes to be optimized and pre-installed on machine images, facilitating control over the computational environment. [12]

It is still difficult and time consuming to develop and deploy a grid application, and complexity issues span over programmatic, technology and management perspectives. This has kept many users from the utilization of grid computing, choosing instead simpler technologies like Web Services and traditional databases. The use of middleware applications and libraries has imposed a level of homogenization on top of the grid fabric, composed of heterogeneous hardware and software. This simplifies resource management inside a particular virtual organization. However, developing cross-grid applications that span across different virtual organizations has remained difficult.

Existing grid middleware can be deployed in a cloud environment, as grid services can run inside image instances and multiple agents performing the same functions can be spawned from a single image easily. If an index service is available for discovering these services, the network configuration of each VM may be unknown, adding flexibility to the design. In this fashion, the master-worker paradigm can be easily deployed in a cloud environment. [13]

4.3 Enterprises

The emergence of clouds has already caused an impact in the IT industry. Many enterprises are deciding to make use of virtual datacenters to facilitate infrastructure managing and sparing the need of hardware maintenance. This type of cyberinfrastructure reduces the complexity involved in deployment of services, at the cost of losing flexibility with a narrower interface, a cost that many users may be willing to pay to deploy applications in a distributed environment.

Leader IT companies are already building their own clouds. Starting and small size enterprises are also becoming users of cloud services like Salesforce [14] (ERP and accounting systems), Google Apps [15], QuickBooks [16] Online instead of using local softwares. But the question still remains about big and middle size non IT enterprises [17]. Where the information privacy is the most important issue and they have already spent lots of money for their local systems.

5 Current Works

Currently there are various cloud systems on both academic and industrial world are being built. Following is a brief review of what is undergoing presently.

5.1 Academia

5.1.1 Eucalyptus

Eucalyptus[18] is an opensource software framework developed by University of California – Santa Barbara for cloud computing that implements what is commonly referred to as Infrastructure as a Service (IaaS); systems that give users the ability to run and control entire virtual machine instances deployed across a variety physical resources. [19]

The current interface to Eucalyptus is compatible with Amazon's EC2 interface, but the infrastructure is designed to support multiple client-side interfaces. EUCALYPTUS is implemented using commonly available Linux tools and basic Web-service technologies making it easy to install and maintain.

The system is being used to experiment with HPC and cloud computing by trying to combine cloud computing systems like Eucalyptus and EC2 with the Teragrid, as a platform to compare cloud computing systems' performance.

5.1.2 Nimbus

The University of Chicago Science Cloud, codenamed "Nimbus"[20], provides compute capability in the form of Xen virtual machines (VMs) that are deployed on physical nodes of the University of Chicago TeraPort cluster (currently 16 nodes) using the Nimbus software. The Nimbus cloud is available to all members of scientific community wanting to run in the cloud. [21]

Nimbus supports both WSRF and EC2 interfaces and it can be configured to use familiar schedulers like PBS and SGE to manage VMs. To be mentioned, University of Florida has also are deployed virtual workspaces services in collaboration with the Nimbus. [22]

Other current scientific world cloud systems that can be named are Kupa by Masaryk University [23], Wispy by Purdue University [24], Virtual Computing Laboratory (VCL) by North Carolina State University [25] and CARMEN by 11 UK universities in collaboration [26].

5.2 Enterprise

These days enterprise Clouds consisting of hundreds of thousands of computing nodes are common (Amazon EC2 [27], Google App Engine [15], Microsoft Live Mesh [28]) and hence federating them together leads to a massive scale environment. It seems that all leading IT companies have understood the importance of cloud computing and its great future needs and they are moving toward it no matter what happens.

6 Challenges Ahead

One of the most important challenges ahead is that clouds will always be compared to local machine in the time of usage. It's important for the user to know what he gains of shifting to the cloud. Obviously using services on local machines, the user needs more resources but at least he knows that he has access to his data all the time and he has the data he owns on his local machine. But who is in charge of restoring his data if something happens to the cloud and the fact that the user is not aware of the physical place which his data is stored makes cloud more unreliable for him. Here is a list of issues that cloud computing is currently facing.

6.1 Information Policy

Cloud computing raises a range of important policy issues, which include issues of privacy, security, anonymity, telecommunications capacity, government surveillance, reliability, and liability, among others. At a minimum, users will likely expect that a cloud will provide:

- **Reliability and Liability.** Users will expect the cloud to be a reliable resource, especially if a cloud provider takes over the task of running “mission-critical” applications and will expect clear delineation of liability if serious problems occur.
- **Security, privacy, and anonymity.** Users will expect that the cloud provider will prevent unauthorized access to both data and code, and that sensitive data will remain private. Users will also expect that the cloud provider, other third parties, and
- **Access and usage restrictions.** Users will expect to be able to access and use the cloud where and when they wish without hindrance from the cloud provider or third parties, while their intellectual property rights are upheld. [29]

Here the most important issue is security and the way that the provider has to assure the user of providing it. Also one of the most important aspects of cloud in which academia is more interested is high performance and adding securing will always reduce performance. Thus there is a need to find a way of implementing security with the least effect on performance.

Here are seven of the specific security issues Gartner says customers should raise with vendors before selecting a cloud vendor.

1. **Privileged user access**. Sensitive data processed outside the enterprise brings with it an inherent level of risk, because outsourced services bypass the "physical, logical and personnel controls" IT shops exert over in-house programs. Get as much information as you can about the people who manage your data. "Ask providers to supply specific information on the hiring and oversight of privileged administrators, and the controls over their access," Gartner says.
2. **Regulatory compliance**. Customers are ultimately responsible for the security and integrity of their own data, even when it is held by a service provider. Traditional service providers are subjected to external audits and security certifications. Cloud computing providers who refuse to undergo this scrutiny are "signaling that customers can only use them for the most trivial functions," according to Gartner.
3. **Data location**. When you use the cloud, you probably won't know exactly where your data is hosted. In fact, you might not even know what country it will be stored in. Ask providers if they will commit to storing and processing data in specific jurisdictions, and whether they will make a

contractual commitment to obey local privacy requirements on behalf of their customers, Gartner advises.

4. **Data segregation**. Data in the cloud is typically in a shared environment alongside data from other customers. Encryption is effective but isn't a cure-all. "Find out what is done to segregate data at rest," Gartner advises. The cloud provider should provide evidence that encryption schemes were designed and tested by experienced specialists. "Encryption accidents can make data totally unusable, and even normal encryption can complicate availability," Gartner says.
5. **Recovery**. Even if you don't know where your data is, a cloud provider should tell you what will happen to your data and service in case of a disaster. "Any offering that does not replicate the data and application infrastructure across multiple sites is vulnerable to a total failure," Gartner says. Ask your provider if it has "the ability to do a complete restoration, and how long it will take."
6. **Investigative support**. Investigating inappropriate or illegal activity may be impossible in cloud computing, Gartner warns. "Cloud services are especially difficult to investigate, because logging and data for multiple customers may be co-located and may also be spread across an ever-changing set of hosts and data centers. If you cannot get a contractual commitment to support specific forms of investigation, along with evidence that the vendor has already successfully supported such activities, then your only safe assumption is that investigation and discovery requests will be impossible."
7. **Long-term viability**. Ideally, your cloud computing provider will never go broke or get acquired and swallowed up by a larger company. But you must be sure your data will remain available even after such an event. "Ask potential providers how you would get your data back and if it would be in a format that you could import into a replacement application," Gartner says. [30]

6.2 Provenance Data

Cloud provenance data, and in general meta-data management, is an open issue. Open challenges include: How to collect provenance information in a standardized and seamless way and with minimal overhead – modularized design and integrated provenance recording; How to store this information in a permanent way so that one can come back to it at anytime, - Standardized schema; and How to present this information to the user in a logical manner – an intuitive user web interface. [31]

There are also issues like the scalability, portability of services, cloud interactions, interoperability, fault tolerance, energy cost and the cost of building clouds versus keeping the current systems. Given the proliferation of different virtualization environments, and the variety in the hardware, standardization of image formats is of considerable interest. Some open solutions exist or are under consideration, and a number of more proprietary solutions are here already.

7 Conclusion

Cloud computing is an emerging computing paradigm that is increasingly popular. Leaders in the industry, such as Microsoft, Google, and IBM, have provided their initiatives in promoting cloud computing. However, the public literature that discusses the research issues in cloud computing are still inadequate.

In a study of the research literature surrounding cloud computing, I found that there is a distinct focus on the needs of the scientific computing community. Big IT companies are also building their own version of cloud. But still there are many question have left without an answer and indeed the most important one is security.

One of the other aspects of the cloud which is left is the social aspect of it. The Cloud is going to happen but which services should be offered on the cloud and for whom. What happens if smaller IT companies start to offer their services on the cloud and no one uses them?! I believe that everything eventually can move to the Cloud. The question is if users are ready for that and if it's the right move and this need must be addressed.

8 References

[1] From "NSF'S Cyberinfrastructure Vision for 21st Century Discovery," NSF Cyberinfrastructure Council, September 26th, 2005, Ver.4.0, pg 4.

[2] Wikipedia, "Cyberinfrastructure", http://en.wikipedia.org/wiki/Cyberinfrastructure

[3] Ditto, Appendix A (http://www.nsf.gov/od/oci/reports/APXA.pdf)

[4] M.A. Vouk, "Virtualization of Information Technology Resources", in Electronic Commerce: A Managerial Perspective 2008, 5th Edition y Turban, Prentice-Hall Business Publishing, to appear.

[5] Mike P. Papazoglou, "Service -Oriented Computing: Concepts, Characteristics and Directions", Tilburg University, INFOLAB,

[6] Wikipedia, "Workflow", http://en.wikipedia.org/wiki/Workflow

[7] Lead Project, https://portal.leadproject.org/

[8] An Introduction to Virtualization, http://www.kernelthread.com/publications/virtualization/

[9] Lijun Mei, W.K. Chan, T.H. Tse, "A Tale of Clouds: Paradigm Comparisons and Some Thoughts on Research Issues", To appear in Proceedings of the 2008 IEEE Asia-Pacific Services Computing Conference (APSCC 2008), IEEE Computer Society Press, Los Alamitos, CA

[10] R. Buyya, C. S. Yeo, and S. Venugopa, "Marketoriented cloud computing: Vision, hype, and reality for delivering it services as computing utilities" In Proceedings of the 10th IEEE International Conference on High Performance Computing and Communications (HPCC-08, IEEE CS Press, Los Alamitos,CA, USA) 2008.

[11] Mike Ricciuti, "Stallman: Cloud computing is 'stupidity'", http://news.cnet.com/8301-1001_3-10054253-92.html

[12] J. J. Rehr, J. P. Gardner, M. Prange, L. Svec and F. Vila, "Scientific Computing in the Cloud", Department of Physics, University of Washington, Seattle

[13] Laboratório Nacional de Computação Científica, "Using Clouds to address Grid Limitations", Av. Getúlio Vargas, 333 – Quitandinha

[14] Salesforce, http://www.salesforce.com/

[15] Google App Engine, http://appengine.google.com

[16] QuickBooks Online, http://oe.quickbooks.com/_bb/index.cfm

[17] Mathew Schwartz, "Running Your Business In The Cloud", http://www.bmighty.com/security/showArticle.jhtml;jsessionid=21DDR5RZ44AGGQSNDLPCKHSCJUNN2JVN?articleID=210604071&pgno=1

[18] Eucalyptus, http://eucalyptus.cs.ucsb.edu/

[19] Daniel Nurmi, Rich Wolski, Chris Grzegorczyk, Graziano Obertelli, Sunil Soman, Lamia Youseff, Dmitrii Zagorodnov, "The Eucalyptus Open-source Cloud-computing System", Computer Science Department, University of California - Santa Barbara, Santa Barbara, California 93106

[20] Nimbus, http://workspace.globus.org/

[21] Introduction to Nimbus @ UC, http://workspace.globus.org/clouds/nimbus.html

[22] STRUTS, http://www.acis.ufl.edu/vws/

[23] Kupa, http://meta.cesnet.cz/cms/opencms/en/docs/clouds

[24] Wiapy, http://www.rcac.purdue.edu/teragrid/resources/#wispy

[25] Virtual Computing Laboratory (VCL), http://vcl.ncsu.edu

[26] CARMEN, http://www.carmen.org.uk/

[27] Amazon Elastic Compute Cloud (EC2), http://www.amazon.com/ec2/

[28] Microsoft Live Mesh, http://www.mesh.com

[29] Paul T. Jaeger , Jimmy Lin, Justin M. Grimes, Cloud Computing and Information Policy: Computing in a Policy Cloud? , Forthcoming in the Journal of Information Technology and Politics, 5(3).

[30] Jon Brodkin, "Gartner: Seven cloud-computing security risks", InfoWord, http://www.infoworld.com/article/08/07/02/Gartner_Seven_cloudcomputing_security_risks_1.html, July 2008

[31] Mladen A. Vouk, "Cloud Computing – Issues, Research and Implementations", Proceedings of the ITI 2008 30th Int. Conf. on Information Technology Interfaces, June 23-26, 2008, Cavtat, Croatia

9 Other Resources

[a] Andy Greenberg, "FORBES: Computing In The Cloud (Layered Tech)", Forbes Magazine, March 2008, http://layeredtech.wordpress.com/2008/03/27/forbes-computing-in-the-cloud-layered-tech

[b] Barry Lynn," Are Enterprises Ready for Cloud Computing?", Web2.0 Journal, October 208, http://web2.sys-con.com/node/643387

[c] Marianne C. Murphy, Marty McClelland, "Computer Lab to Go: A "Cloud" Computing Implementation", Proc ISECON 2008, v25

[d] T. V. Raman, "Cloud Computing and Equal Access for All", Proceedings of the 2008 international cross-disciplinary conference on Web accessibility (W4A), Beijing, China

[e] Ed Sperling, "Cloud Computing isn't For Everyone", Forbes, November 2008, http://www.forbes.com/technology/2008/11/07/cio-cloud-computing-tech-cio-cx_es_1110cloud.html?feed=rss_technology

[f] Wikipedia, "Cloud Computing", http://en.wikipedia.org/wiki/Cloud_computing

[g] E. Naone, "Computer in the Cloud," Technology, Review, MIT, Sept 18, 2007, http://www.technologyreview.com/infotech/19397/?a=f

[h] Paul Watson, Phillip Lord, Frank Gibson, Panayiotis Periorellis, Georgios Pitsilis, "Cloud Computing for e-Science with CARMEN", School of Computing Science, Newcastle University, Newcastle-upon-Tyne, UK

[i] Ian Foster, Yong Zhao, Ioan Raicu, Shiyong Lu, "Cloud Computing and Grid Computing 360-Degree Compared"

[j] Rajiv Ranjan, Rajkumar Buyya, "Decentralized Overlay for Federation of Enterprise Clouds", Grid Computing and Distributed Systems (GRIDS) Laboratory, Department of Computer Science and Software Engineering, The University of Melbourne, Australia

[k] Hiroshi Maruyama, "Challenges and Opportunities for Computer Science in Services Science", IBM Research, Tokyo Research Laboratory