

Collaborative learning in the clouds

Hajar Mousannif · Ismail Khalil · Gabriele Kotsis

© Springer Science+Business Media, LLC 2012

Abstract Cloud computing has emerged as one of the most highly discussed topics both in the academic community and in the computing industry. While most of the work that has been conducted to explore this field focuses either on establishing the basis for cloud computing or almost exclusively on the issues surrounding security and data privacy, this paper takes the first exploratory step into exploring the actual internal working of cloud computing and demonstrates its viability for organizations, more specifically educational establishments. The paper starts by introducing the most important key clouds computing concepts, including virtualization technologies, Web services and Service Oriented Architectures (SOA), and distributed computing. Light will be then shed on the impact and potential benefits of cloud computing on teaching and learning in educational institutions. The paper closes by describing building a private cloud inside educational institution and highlights its offerings for students, staff and lecturers.

Keywords Cloud computing · Virtualization · Private clouds

1 Introduction

With the significant advances in Information and Communications Technology over the last half century, cloud computing is rapidly emerging as the new information technology platform. Cloud computing is a model for enabling ubiquitous,

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 (Mell and Grance 2011). The major idea behind cloud computing is to deliver Infrastructure, Platform and Software as services in a “pay-as-you-go” model, hence rendering the Internet a large repository where resources are available to everyone as services. With such paradigm, users can access services based on their requirements regardless of where the services are hosted or how they are delivered. Moreover, clients only pay for the quantity of the rented resources (data storage, computation, etc.) they consume.

Three service models are currently being differentiated: 1° Software-as-a-Service (SaaS), i.e., online applications, such as web-based email. GoogleApps and Salesforce.com are examples of this model, 2° Platform-as-a-Service (PaaS): The Google App Engine, the Microsoft Windows Azure and Amazon Web Services are PaaS examples which allow customers to deploy their own applications, and 3° Infrastructure-as-a-Service (IaaS), which provides processing power or storage. Amazon Elastic Compute Cloud (EC2), Rackspace, Nimbus, and Eucalyptus (Nurmi et al 2009) are examples of this model. A well-detailed cloud computing services taxonomy can be found in (Höfer and Karagiannis 2011).

Most of the work that has been conducted to explore this field focuses either on establishing the basis for cloud computing (Zhang et al. 2010; Yoo 2011; Marston et al. 2010; Kossmann and Kraska 2010; Wang et al. 2010; Weinhardt et al. 2009) or almost exclusively on the issues surrounding security and data privacy (Zissis and Lekkas 2012; Paquette et al. 2010; Subashini and Kavitha 2011; Oh et al. 2010; Kaur and Kaushal 2011), with very few final answers available. In this paper, we attempt to demystify the current trends in the space of cloud computing by providing an

H. Mousannif (✉)
Faculty of Sciences and Techniques, Cadi Ayyad University,
Marrakesh, Morocco
e-mail: hajar.mousannif@gmail.com

I. Khalil · G. Kotsis
Department of Telecooperation, Johannes Kepler University Linz,
Linz, Austria

overview of its core aspects and assessing some current cloud architectures implementations.

As cloud computing is becoming more and more an adoptable technology for many organizations, we demonstrate its viability for educational institutions in particular and highlight the potential benefits it provides in the educational arena. We also introduce building a private cloud inside educational institutions. More specifically, we show how this cloud is built and highlight its offerings for students, staff and lecturers.

The remainder of this paper is organized as follows. Section 2 provides an overview of the fundamental building blocks of cloud computing such as virtualization, service-oriented architectures and distributed computing. Section 3 overviews and assesses several existing cloud architectures. In section 4, we harness the potentials of cloud computing in the educational sector. In section 5, we introduce the private cloud philosophy. Section 6 presents an effort in building a private cloud inside educational institutions and highlight its offerings.

2 Cloud computing building blocks

Projecting as an evolutionary step, following the transition from mainframe computers to client/server deployment models, much of the basis for Cloud Computing is empowered by an overhyped repackaging of already existing technologies and few new concepts in several research fields like virtualization, Service-Oriented Architectures (SOA), and distributed and grid computing.

Virtualization (Barham et al. 2003) is the technology that hides the physical characteristics of a computing platform from the users, instead presenting an abstract, emulated computing platform. Virtualization renders a single physical machine able to function as a set of multiple logical Virtual Machines (VMs); it actually allows hosting and running multiple operating system environments, totally isolated from one another (Smith and Nair 2005), hence providing maximum flexibility to configure various partitions of resources on the same piece of hardware to different specific requirements of service requests. With respect to this, a cloud can be seen as a pool of virtualized resources across the Internet that follows a pay-per-use model and that can be dynamically reconfigured to satisfy user requests via on-the-fly provisioning/deprovisioning of virtual machines (Vaquero et al. 2009). Since the cloud provider can dynamically offer different hardware and software configurations to the cloud user, the user can select a machine configuration and manage this machine with full privileges without interfering in cloud machines available to the others. Quality of virtualization determines the robustness of a cloud infrastructure. Good virtualization can effectively assist sharing of cloud facilities,

managing of complex systems, and isolation of data/application (Xu 2012). Some market-leading virtualization technologies are VMware, which we will be discussing in more details in section 5 and 6, and Xen (Abels et al. 2005).

In addition to virtualization, cloud computing platforms possess characteristics of both clusters and grids. Cloud computing and grid computing differ, however, in target, in focus, and in the implementation technologies (Foster et al. 2008). On the one hand, with respect to a grid or cluster, a cloud strongly supports failover and content replication since VMs can be easily migrated from one node to another. The distributed nature of clouds gives, however, the cloud user less control over the location of data and computation. On the other hand, cluster and grid computing were initially designed and developed for facilitating High Performance Computing (HPC) applications, such as simulations, while cloud computing aims also at supporting highly collaborative applications across organizations and users of different expertise and varying levels of authorization, hence making management costs much lower than those of clusters and grids and management less cumbersome. More comparisons highlighting strengths and weaknesses of the three above-mentioned major computing paradigms are available in (Mateescua et al. 2011).

Given the current trend of increasing complexity of collaboration among users and systems and the automation of business processes using IT services, other building blocks of cloud computing include Service-Oriented Architectures (SOA) and Web Services. In systems following the SOA approach, functionalities are delivered and consumed as services. The SOA facilitates interoperable services between distributed systems to communicate and exchange data with one another, thus providing a uniform means for service users and providers to discover and offer services respectively. The interfaces cloud infrastructures provide to the users are either command-line or in the form of Web service, using standards such as Simple Object Access Protocol (SOAP), REpresentational State Transfer (REST) and Web Services Description Language (WSDL) (Muehlen et al. 2005). Web services and SOA are important to manage cloud services mainly because they enable service reusability and interoperability and reduce development and maintenance costs (Radhakrishnan and Sriraman 2007).

3 Cloud architectures overview

Several cloud platforms have been investigated and developed. They basically fall into two major categories: Industrial and academic. Industrial cloud platforms include Amazon Elastic Compute Cloud (EC2), Google App Engine, Microsoft Azure, Sun network.com (Sun Grid) and Aneka among many others. These cloud platforms have been heavily introduced in the

literature. Readers can refer to (Rimal et al. 2009) for more details about these cloud systems, as well as their strengths, weaknesses, and challenges. Academic efforts include the RESERVOIR project (Rochwerger et al. 2009), Virtual Workspaces (Keahey et al. 2005), and OpenNebula (OpenNebula 2011).

Most cloud computing architectures are structured in a modular fashion in a sense that on the highest level of their abstraction, the main roles and the corresponding architectural elements are defined allowing drilling down for each of these elements as needed. Most cloud computing architectures define three fundamental roles: cloud service consumer, cloud service provider, and cloud service creator. Each role can be fulfilled by a single person, a group of people, an organization or an IT system.

People or IT systems acting in the role of Cloud Service Consumers browse the service offering catalog delivered by a particular cloud service, consume the offered service instances, and get billed for all or part of them.

A Cloud Service Provider is in charge of providing cloud services to Cloud Service Consumers. Several cloud computing providers make use of open-source software and platforms. Amazon uses the open-source Xen technologies and Google's PaaS offering is built around the open-source Python programming language, but their core cloud computing service and additional services are kept closed-source.

A Cloud Service Creator is responsible for creating a cloud service: Infrastructure, Platform, Software or Business Process Services, which can be run by a Cloud Service Provider and by that exposed to Cloud Service Consumers.

The IBM Cloud Computing Reference Architecture (CCRA) (Behrendt et al. 2011), as an example, defines with respect to the three high level roles identified above some other functional cross cutting layers such as integration and governance. The integration functionality is specifically relevant in the context of hybrid clouds, where seamless integrated management, usage and interoperability of cloud services in integration with on-premise IT is critical. Governance in a cloud (Wohl 2006) consists in exercising control over services, imposing policies and monitoring services in order to deliver value to the business. It is widely argued that governance can effectively manage increasing risk and multiple regulations, and ensure continuity of critical business processes in the cloud.

From an infrastructural point of view, a cloud can be seen as a distributed system that consists of many servers, storage devices, and networking devices. Its simplest form is a 3-tier client-server architecture that consists of three main layers: 1°) the physical hardware at the bottom, 2°) a virtualization environment at the next higher layer, and 3°) the software for supporting the PaaS or the SaaS at the top. It is important to note that the physical existence of a virtualized

infrastructure on the cloud service provider side is not mandatory, since it is possible for a cloud service provider to consume infrastructure as a service from a different cloud service provider and put higher value cloud services on top.

4 Cloud computing and education

In this section we will be discussing the benefits of cloud computing to higher education institutions and different educational uses.

Any educational institution should be able to (Sultan 2010):

- provide students and staff with software (e.g., email accounts, operating systems, etc.) and hardware (e.g., PCs, Servers, etc.);
- provide researchers and postgraduate students with the required special software and hardware to run experiments that are likely to involve a great deal of processing and computation;
- provide Web developers with the development tools needed to write and host Web applications.

With educational institutions being more and more under budget shortage and wanting to better reallocate their limited internal resources to better support their corporate priorities, many schools and universities have already started taking advantage of available cloud-based applications offered by service providers. Students, administrative staff and lecturers can be made to use the services of providers of SaaS and IaaS clouds, developers can use all the software they need for their development online and all the hardware for hosting their applications through a PaaS cloud provider, while researchers whose projects require a great deal of processing power and/or additional server capacity can do so in minutes and at the click of a button through an IaaS cloud provider. Amazon's Elastic Compute Cloud (EC2), as an example, provides a computing environment that allows users to rent virtual computers on which they run their own computer applications and use Web service interfaces in order to manage (create, launch, stop .etc.) server instances. EC2 relies on Amazon Machine Images (AMIs) which contain a variety of operating systems, applications, libraries, data and associated configuration settings. Amazon EC2 charges the user for the use of the instances.

Many other cloud providers are offering a whole set of free services to staff and students at educational institutions which include email, contact lists, calendars, document storage/creation/sharing, and the ability to create websites. (Google Apps for Education 2011) is one of them and is already being used in 61 of the Top 100 schools according to (US News & World Report 2011). It is a collection of Web-based messaging (e.g., Gmail, Google Talk, and Google

Calendar) and productivity and collaboration tools (Google Docs: text files, spreadsheets, and presentations).

Cloud computing offers the educational sector a way to increase IT capacity and add on the fly capabilities without investing in new infrastructure, new training, or licensing new software. Many educational establishments across the world found cloud computing to be very attractive to use and has already integrated it for educational usage. The University of California (UC) at Berkeley, as an example, was able to move a course that focused on developing and deploying software-as-a-service (SaaS) applications from its locally-owned infrastructure to the cloud (Fox 2009). A virtual machine image containing the whole software stack was created, and each student or team could deploy that image on an EC2 server instance and instantly have the same experience as if they themselves were administering the application server in a datacenter. While on a shared Berkeley server, root access would never be granted, with EC2 students can have root access on their own image, and any damage they do can be undone by simply re-instantiating the image. Another complete case study on how The University of Westminster, which is a UK educational establishment with more than 22,000 students, has embraced cloud computing is fully detailed in (Sultan 2010).

All examples mentioned above demonstrate that not only the equipment setup can be outsourced, but the ongoing IT administration of the resources as well.

5 The private cloud philosophy

In terms of resources availability and deployment strategy, clouds can be classified into three different types:

- **Public clouds:** Resource providers offer computing resources as services to a wide range of interested customers in a pay-per-use basis. The use of machines is leased to the user during the requested time. As such, public clouds come the closest to the vision of utility computing that advances that computing will be one day commoditized and delivered in a manner similar to traditional utilities such as water, electricity, gas, and telephony (Buyya et al. 2009).
- **Private clouds or internal clouds:** Clouds with resources that can be accessed and used by individuals inside an organization for its internal purposes, having similarities with data farms or private grids. A private cloud is typically hosted on the company's own servers, within their own network infrastructure.
- **Hybrid clouds:** Bring together public and private clouds. Some cloud providers make it possible to connect the virtual instances to the company's network via VPN (Virtual Private Network), to make the company

network seem like one big scalable IT infrastructure, resulting in a combination of control over performance and security with elasticity. An example of this is BlueLock's BlueLock Box (BlueLock Box 2011), which is a hardware device that works as a miniature cloud server but is managed by BlueLock's support team and can also connect to BlueLock's own cloud infrastructure for adding capacity or backup/restore functionality.

As security is becoming more and more a serious issue with public cloud computing in any form, private clouds are more appealing to organizations that need or want more control over their data than they can get by using a third-party hosted service such as Amazon's EC2 or Simple Storage Service (S3).

Many private cloud technologies have already attracted many organizations. Such solutions include Eucalyptus (Nurmi et al 2009), Vmware (Vmware Vsphere 2011), and Microsoft (Microsoft 2011).

Eucalyptus is a major vendor of private clouds creation technology. The platform interfaces directly with Amazon EC2 and provides a REST and SOAP interface compatible with Amazon Web Services (AWS). It also exposes administrative functionalities (e.g., user management, storage configuration, network management, hypervisor configuration, etc.) for managing and maintaining the cloud.

No mention of virtualization in any form would be complete without some discussion of VMware. Their own private-cloud architecture is built on top of VMware vSphere which we used for building private cloud infrastructure and which we will be discussing in the next section. They also offer an IaaS version via vCloud Express which is an IaaS offering that allows developers to address various infrastructure and programming needs such as experimentation, prototyping and testing.

Microsoft came late to virtualization, but managed lately to offer private cloud solutions based on Windows Server with Hyper-V and System Center.

Private clouds have already penetrated many organizations in the educational sector. The Hochschule Furtwangen University (HFU), as an example, is running their own private cloud infrastructure (Doelitzscher et al. 2011), called Cloud Infrastructure and Application (CloudIA), which is mainly dedicated to HFU staff and students running e-Learning applications, and external people for collaboration purposes. Other examples of private clouds use in educational institutions include The Virtual Computing Laboratory (VCL) (Vouk et al 2008), developed by North Carolina State University (USA), which enables students to reserve and access virtual machines (VMs) with a basic image or specific applications environments, such as Matlab and Autodesk along with physical computing hardware management and the BlueSky cloud framework (Dong et al. 2009), developed by Xi'an Jiaotong University (China),

which enables physical machines to be virtualized and allocated on-demand for e-Learning purposes.

6 Building a private cloud

Motivated by all the aforementioned benefits of integrating cloud computing in educational institutions, and especially those related to private clouds, we will be describing the effort in building a private cloud inside educational institution, from the architectural and functional point of views, and we will be highlighting its offering.

As a first step towards implementing a more comprehensive private cloud platform, we opted for a simple 3-tier client-server architecture that involves 3 main layers: The physical hardware at the bottom, a virtualization environment at the next higher level, and the service layer at the top, as depicted in Fig. 1. A description of each layer is provided bellow.

6.1 The virtualization layer

In order to allow multiple operating systems to share a single hardware host, a hypervisor is needed. A hypervisor is actually a program that controls the host processor and resources by allocating what is needed to each operating system in turn and making sure that the virtual machines cannot disrupt each other.

There are two types of hypervisors: native (also called bare-metal or type 1) and hosted (also called type 2). The former runs directly on the physical hardware while the latter runs on a host operating system (Linux, Windows, etc.), which runs on the physical hardware. Native hypervisors include IBM's original CP/CMS and z/VM, VMware's ESXi, Citrix's Citrix XenServer and Microsoft's Hyper-V, etc. Oracle/Sun's VirtualBox and VMware's Workstation are examples of hosted hypervisors.

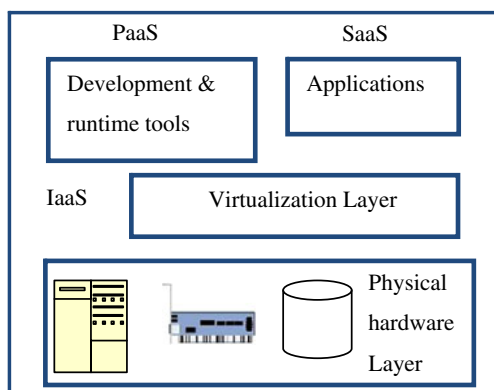


Fig. 1 Private cloud architecture overview

For this cloud architecture, we deployed a native hypervisor using VMware's Vsphere ESXi (Vmware Vsphere 2011).

6.2 The physical layer

The private cloud architecture is built on top of an existing hardware infrastructure that consists of two computer pools, a PC pool and a server pool, which are both used for teaching purposes, e.g. running labs exercises and practical projects, during the day. The servers have VMware vSphere Hypervisor ESXi version 5.0 directly installed on top of the hardware while the computers in the PC pool have a thin client installed and can simply access the VMs on the server pool. As a thin client we used ThinStation version 2.2.2 (ThinStation 2011) which is a basic and small Open Source thin client operating system. The main advantage of using a thin client is that an old Pentium 100 MHz with 32 MB RAM can be a perfectly useful workstation. Moreover, a hard disk may not even be needed since a direct boot off the network is possible using a network boot protocol (PXE). Both pools are managed using VMware's Vcenter (Vmware Vsphere 2011) which has the important feature of ensuring that the most important applications get the computing, network and storage resources they need by prioritizing workloads. All hosts and services can be visualized and managed through a web browser, displaying real-time data about the virtual infrastructure including performance, configuration, storage, alarms and access permissions.

6.3 The service layer

This layer provides software for supporting the PaaS and the SaaS the cloud users need.

All third, fourth and fifth year students in universities are required to take at least one practical project in various topics such as designing and implementing new software architectures, and evaluating new networking or programming techniques, we created ready-to-use and customized VMs that contain most of what students need to develop, deploy and test their own applications. The VMs include predefined development tools such as MySQL, Tomcat, PHP, and Apache web server and students simply access the required VM using the thin client.

This private cloud is also an IaaS framework that allows creating virtual machines by choosing base images and post installing software packages selected by the user on the fly the same way Amazon Machine Images (AMIs) are used. However, this feature is currently available to course coordinators only.

7 Benefits of educational private cloud

Universities acknowledge the potential benefits of using the deployed cloud. It actually allows the university to avoid over-

provisioning IT infrastructure and reduce its maintenance costs. Indeed, instead of providing physical machines to students/staff, this cloud provides VMs. As an example, VMs use is particularly suitable for networking exercises like routing labs or firewall configuration where students usually need multiple physical machines to complete the exercise. Before using the cloud, PCs were assigned among students at the beginning of each semester and at the end of the semester, PCs need to be collected and maintained (i.e. components functionality check, wiping of installed software, new OS installation, etc.). This process consumes a considerable amount of time of IT resources management department depending on how many PCs were used. Using the cloud considerably optimizes this process since only image re-instantiation is required which is a matter of minutes instead of hours.

Moreover, before deploying the cloud, lecturers were most of the time complaining about errors due to OS misconfiguration, missing software components or missing programming libraries. Now with predefined VMs use, all attending students use the same development environment, perfectly known to the lecturer.

References

- Abels, T., Dhawan, P., & Chandrasekaran, B. (2005). An overview of xen virtualization, <http://www.dell.com/downloads/global/power/ps3q05-20050191-Abels.pdf>. Accessed 15th September 2011.
- Barham, P., Dragovic, B., Fraser, K., Hand, S., Harris, T., Ho, A., et al. (2003). Xen and the art of virtualization. *19th ACM Symposium on Operating Systems Principles*. doi:10.1145/1165389.945462.
- Behrendt, M., Glasner, B., Kopp, P., Dieckmann, R., Breiter, G., Papp, S. et al. (2011). <https://www.opengroup.org/cloudcomputing/uploads/40/23840/CCRA.IBMSubmission.02282011.doc>. Accessed 12th July 2011.
- BlueLock Box (2011). <http://www.bluelock.com/cloud-hosting/virtual-private-cloud/>. Accessed 27th July 2011.
- Buyya, R., Yeo, C. S., Venugopal, S., Broberg, J., & Brandic, I. (2009). Cloud computing and emerging IT platforms: vision, hype, and reality for delivering computing as the 5th utility. *Future Generation Computer Systems*, 25(6), 599–616.
- Doelitzscher, F., Sulistio, A., Reich, C., Kuijs, H., & Wolf, D. (2011). Private cloud for collaboration and e-Learning services: from IaaS to SaaS. *Computing*, 91(1), 23–42.
- Dong, B., Zheng, Q., Qiao, M., Shu, J., & Yang, J. (2009). BlueSky cloud framework: an e-learning framework embracing cloud computing. *Lecture Notes in Computer Science*. doi:10.1007/978-3-642-10665-1_55.
- Foster, I., Zhao, Y., Raicu, I., & Lu, S. (2008). Cloud computing and grid computing 360-degree compared. *Grid computing environments workshop*. doi:10.1109/GCE.2008.4738445.
- Fox, A. (2009). Cloud computing in education. Berkeley iNews, <https://inews.berkeley.edu/articles/Spring2009/cloud-computing>. Accessed 27th July 2011.
- Google Apps for Education (2011). <http://www.google.com/apps/intl/en/edu/index.html>. Accessed 15th July 2011.
- Höfer, C. N., & Karagiannis, G. (2011). Cloud computing services: taxonomy and comparison. *Journal of Internet Services and Applications*, 2(2), 81–94.
- Kaur, P. J., & Kaushal, S. (2011). *Security concerns in cloud computing, communications in computer and information science, high performance architecture and grid computing* (pp. 103–112). New York: Springer.
- Keahey, K., Foster, I., Freeman, T., & Zhang, X. (2005). Virtual workspaces: achieving quality of service and quality of life in the grid. *Scientific Programming*, 13(4), 265–275.
- Kossmann, D., & Kraska, T. (2010). Data Management in the cloud: promises, state-of-the-art, and open questions. *Datenbank-Spektrum: Zeitschrift für Datenbanktechnologien und Information Retrieval*, 10(3), 121–129. doi:10.1007/s13222-010-0033-3.
- Marston, S., Li, Z., Bandyopadhyay, S., Zhang, J., & Ghalsasi, A. (2010). Cloud computing-The business perspective. *Decision Support Systems*, 51(1), 176–189.
- Mateescu, G., Gentsch, W., & Ribbens, C. J. (2011). Hybrid computing—where HPC meets grid and cloud computing. *Future Generation Computer Systems*, 27(5), 440–453.
- Mell, P., & Grance, T. (2011). The NIST Definition of Cloud Computing, Recommendations of the National Institute of Standards and Technology. <http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf>. Accessed 27 October 2011.
- Microsoft (2011). <http://www.microsoft.com/en-us/server-cloud/private-cloud/default.aspx>. Accessed 27th July 2011.
- Muehlen, M. Z., Nickerson, J. V., & Swenson, K. D. (2005). Developing web services choreography standards—the case of REST vs. SOAP. *Decision Support Systems*, 40(1), 9–29.
- Nurmi, D., Wolski, R., Grzegorzczak, C., Obertelli, G., Soman, S., Youseff, L., et al. (2009). The Eucalyptus open-source cloud computing system. 9th IEEE/ACM international symposium on cluster computing and the grid (CCGRID), doi: 10.1109/CCGRID.2009.93.
- Oh, T. H., Lim, S., Choi, Y. B., Park, K., Lee, H., & Choi, H. (2010). State of the art of network security perspectives in cloud computing. *Communications in Computer and Information Science*. doi:10.1007/978-3-642-16444-6_7978.
- OpenNebula (2011). OpenNebula project. <http://www.opennebula.org/>. Accessed 20th July 2011.
- Paquette, S., Jaeger, P. T., & Wilson, S. C. (2010). Identifying the security risks associated with governmental use of cloud computing. *Government Information Quarterly*, 27(3), 245–253.
- Radhakrishnan, R., & Sriraman, B. (2007). Aligning Architectural Approaches towards an SOA-Based Enterprise Architecture. WICSA '07 Proceedings of the Sixth Working IEEE/IFIP Conference on Software Architecture, doi: 10.1109/WICSA.2007.5.
- Rimal, B. P., Choi, E., & Lumb, I. (2009). A taxonomy and survey of cloud computing systems. In: NCM'09: proceedings of the 2009 fifth international joint conference on INC, IMS and IDC, doi: 10.1109/NCM.2009.218.
- Rochwerger, B., Breitgand, D., Levy, E., Galis, A., Nagin, K., Llorente, I. M., et al. (2009). The reservoir model and architecture for open federated cloud computing. *IBM Journal of Research and Development*, 53(4).
- Smith, J. E., & Nair, R. (2005). The architecture of virtual machines. *Computer*, 38(5), 32–38.
- Subashini, S., & Kavitha, V. (2011). A survey on security issues in service delivery models of cloud computing. *Journal of Network and Computer Applications*, 34(1), 1–11.
- Sultan, N. (2010). Cloud computing for education: a new dawn? *International Journal of Information Management*. doi:10.1016/j.ijinfor.2009.09.004.
- ThinStation (2011). ThinStation, <http://thinstation.org/>. Accessed 27th July 2011.
- Vaquero, L. M., Roderer-Merino, L., Caceres, J., & Lindner, M. (2009). A break in the clouds: towards a cloud definition. *ACM SIGCOMM Computer Communication Review*, 39(1), 50–55.

- VMware Vsphere (2011). <http://www.vmware.com/products/vsphere/overview.html>. Accessed 27th July 2011.
- Vouk, M., Averritt, S., Bugaev, M., Kurth, A., Peeler, A., Schaffer, H., et al (2008). Powered by VCL—using Virtual Computing Laboratory (VCL) Technology to Power Cloud Computing. In: Proceedings of the 2nd international conference on the virtual computing initiative (ICVCI'08).
- Wang, L., Laszewski, G., Younge, A., He, X., Kunze, M., Tao, J., et al. (2010). Cloud computing: a perspective study. *New Generation Computers*. doi:10.1007/s00354-008-0081-5.
- Weinhardt, C., Anandasivam, A., Blau, B., Borissov, N., Meinel, T., Michalk, W., et al. (2009). Cloud computing—a classification business models, and research directions. *Business Information Systems Engineering*, 1(5), 391–399.
- Wohl Associates (2006). SOA Governance. An IBM White Paper. http://www-01.ibm.com/software/solutions/soa/Amy_Wohl_SOA_Governance_Analyst_White_Paper.pdf. Accessed on 5th July 2011.
- US News & World Report (2011). http://www.usnews.com/usnews/store/college_compass_google_deal.htm. Accessed 15th July 2011.
- Xu, X. (2012). From cloud computing to cloud manufacturing. *Robotics and Computer-Integrated Manufacturing*. doi:10.1016/j.rcim.2011.07.002.
- Yoo, C. S. (2011). Cloud computing: architectural and policy implications. *Review of Industrial Organization*, 38(4), 405–421.
- Zhang, Q., Cheng, L., & Boutaba, R. (2010). Cloud computing: state-of-the-art and research challenges. *Journal of Internet Services and Applications*, 1(1), 7–18.
- Zissis, D., & Lekkas, D. (2012). Addressing cloud computing security issues. *Future Generation Computer Systems*, 28(3), 583–592.

Hajar Mousannif holds a PhD degree in computer Sciences on her work on wireless sensor networks and vehicular networks. She received an engineering degree in Telecommunications and Computer Sciences from the National Institute of Posts and Telecommunications (INPT), Rabat (Morocco) in 2005. She joined, the same year, the Faculty of Sciences and Techniques Marrakech (Cadi Ayyad University, Morocco) as a telecom engineer where she was mainly responsible for managing and administrating the network. Since 2006, she is a teacher within the Department of Computer Sciences at the same Faculty. Her primary research interests include cloud computing, wireless sensor networks, vehicular adhoc/sensor networks, and next generation internet technologies. In addition to her academic experience, she was in the program committee chair of many conferences.

Ismail Khalil is a senior researcher and lecturer at the institute of telecooperation, Johannes Kepler University Linz, Austria since October 2002. He is the president of the international organization of Information Integration and Web-based Applications & Services @WAS (www.iiwas.org). He holds a PhD in computer engineering and received his habilitation degree in applied computer science on his work on agents' interaction in ubiquitous environments in May 2008. He

currently teaches, consults, and conducts research in Mobile Multimedia, Agent Technologies, and the Semantic Web and is also interested in the broader business, social, and policy implications associated with the emerging information technologies.

Before joining Johannes Kepler University of Linz, he was a research fellow at the Intelligent Systems Group at Utrecht University, Netherlands from 2001-2002 and the project manager of AgenCom project at the Software Competence Center Hagenberg - Austria from 2000-2001.

He has authored around 100 scientific publications, books, and book chapters. He is the editor of the Handbook of Research on Mobile Multimedia series, the book Mobile Multimedia: Communication Engineering Perspective, the book Multimedia Transcoding in Mobile and Wireless Networks and the book Innovations in Mobile Multimedia Communications and Applications: New Technologies. He serves as the editor-in-chief of the International Journal on Web Information Systems (IJWIS), International Journal on Pervasive Computing and Communication (IJGCC) both published by Emerald Group publishing, UK, Journal of Mobile Multimedia (JMM) published by Rinton Press, USA, International Journal of Mobile Computing and Multimedia Communication (IJMCMC) published by IGI Global, USA, Advances in Next Generation Mobile Multimedia book series published by IGI Global, USA, and Atlantis Ambient and Pervasive Intelligence book series published by Atlantis and Springer. He is on the editorial board of several international journals. His work has been published and presented at various conferences and workshops.

Gabriele Kotsis started her scientific career at the University of Vienna. She received her masters degree (1991, honored with the Award of the Austrian Computer Society), her PhD (1995, honored with the Heinz-Zemanek Preis) and the *venia docendi* in computer science (2000) from the University of Vienna. Before joining Johannes Kepler University of Linz (JKU Linz) in October 2002, she was working as a researcher and teacher at the University of Vienna (1991-2001), at the Vienna University for Economics and Business Administration (2001) and at the Copenhagen Business School (2002).

Since 2002, Gabriele Kotsis is holding a full professor position in computer science at JKU Linz. She is chairing the Department of Telecooperation with a research focus in mobile computing, multimedia and hypermedia systems as well as cooperative and collaborative systems. Research in those areas includes the investigation of methods, techniques and tools for system development as well as evaluation and analysis with focus on performance evaluation. The Department is participating in numerous national and international projects, including CRUISE, a European network of excellence in sensor networks, Euro-FGI, a network of excellence on Future Generation Internet, the AustrianGrid project, or ModelCVS a project on semantics in SW and system modelling, and actively involved in the organisation of international conferences, including for example iiWAS and MoMM.

She is author of numerous publications in international conferences and journals and is co-editor of several books. She is member of the OCG and the ACM. From April 2003 to April 2007 she was president of the Austrian Computer Society.