Efficient Management of Hybrid Clouds

Sofie Van Hoecke, Tom Waterbley, Jan Devos
Electronics and Information Technology Lab (ELIT)
University College West-Flanders (Howest)
Ghent University Association - Courtrai, Belgium
{*sofie.van.hoecke, tom.waterbley, jan.devos*}*@howest.be*

Tijl Deneut, Johan De Gelas
Sizing Servers Lab
University College West-Flanders (Howest)
Ghent University Association - Courtrai, Belgium
{*tijl.deneut, johan.de.gelas*}*@howest.be*

***Abstract*—Cloud computing has become a significant technology trend driven by big players that is transforming our current IT industry. Public cloud computing comes with advantages such as cost savings, high availability, and easy scalability. However, Small and Medium-sized Enterprises are driven by different reasons for not outsourcing their IT infrastructure entirely. By combining the benefits of the private and public cloud, hybrid cloud computing allows Small and Medium-sized Enterprises to optimize their infrastructure and run their virtual machines where they will be most effective and efficient. We present in this paper a virtual infrastructure management tool that allows to set-up and manage hybrid clouds efficiently in a user-friendly way. Our tool provides automatic load balancing between the private and public clouds at the virtual machine level, and dynamically upscales on-premise virtual machines to public cloud servers based on cost and performance information.**

***Keywords*-hybrid cloud, virtual infrastructure management, SME, load balancing.**

I. INTRODUCTION

Cloud computing has become a significant technology trend, driven by big players like Amazon, Microsoft, Google and VMware, and transforming our current IT industry. Cloud computing delivers large-scale utility computing services to a wide range of consumers. Within cloud computing, users on various types of devices access programs, storage, processing and applications over the Internet, offered by cloud computing providers, resulting in a previously unprecedented elasticity of resources. Through economies of scale, cloud computing comes with advantages such as cost savings, high availability, easy scalability [1], and the transformation of capital IT expenditures into operational IT expenditures. Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS) relieves cloud users from maintaining their infrastructure, development areas, and respectively software.

The increasing adoption rate of cloud computing is currently driving developers, integrators and hosting companies to take cloud computing into account. The last few years, especially the number of providers delivering IaaS has increased quickly [2]. Consequently, companies need to revise their current assets as cloud computing is becoming a strategic weapon. The expansion of IaaS providers increases the options available to companies when acquiring resources in a cost effective manner. Berkeley even predicts that the economy of scale and statistical multiplexing may ultimately lead to a handful of cloud computing providers and "data center-less" companies [3].

However, we do not believe that most Small and Medium-sized Enterprises (SMEs) will become "data center-less" in the near future [4][5]: SMEs cover a wide spectrum of industries [6] and the number of SMEs far exceeds the number of large and very large organizations in almost every country all over the world. Besides the large number of SMEs worldwide, it is also recognized that SMEs constitutes a growing importance as providers of employment opportunities and key players for the well-being of local and regional communities. SMEs are driven by different reasons to maintain their own data center, such as legislation of storing data in-house, investments in the current infrastructure, or the extra latency and performance requirements. This drive is supported by the fact that SMEs have already invested heavily in their own private server equipment and software.

Consequently, we feel that a hybrid approach makes more sense for SMEs. Through the creation of hybrid clouds [7], one can use the internal infrastructure combined with public cloud resources (see Figure 1). This way, on one hand, critical applications can run on the hardware in the private data center or collocated at an SME hosting provider, and, on the other hand, the public cloud can be used as a solution to manage peak demands (cloudbursting) or for disaster recovery. These hybrid clouds capitalize on investments made on the private IT infrastructure, and upscale to the public cloud for specific application requirements.

The architectural concept of a hybrid cloud is overpowering: being able to dynamically move virtualized servers between your data center and a public cloud provider. However, there are still many research challenges to be tackled before hybrid cloud computing can become a reality. One key question is how to enable virtual infrastructure management, meaning the dynamic orchestration of virtual machines (VMs). The scaling efficiency and elasticity of hybrid cloud computing all depend on the efficiency of the virtual infrastructure management [8]. As many cloud providers are incompatible and use proprietary cloud soft-

 ISBN: 978-1-61208-153-3

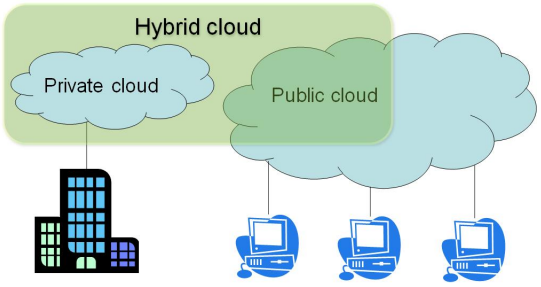

Figure 1. Hybrid cloud computing model

ware and APIs, it is hard to set-up hybrid clouds integrating different cloud solutions, resulting in vendor lock-ins. This is even more the case when it comes to SMEs as they depend strongly on external IT expertise [9][10]. There is little general support with respect to how to set up hybrid clouds and how to manage resources in hybrid environments where management has to act across different resource infrastructures [8]. Existing solutions for hybrid cloud computing and virtual infrastructure management require multiple tools to cooperate and a lot of manual configuration. To the author's knowledge, user-friendly tools that allow SMEs to manage the virtual infrastructure themselves are non-existing.

Therefore, the aim of this research is to design and implement an integrated virtual infrastructure management tool that allows SMEs to set-up and manage hybrid clouds efficiently in a user-friendly way. Our virtual infrastructure management tool provides automatic load balancing between the private and public clouds at the VM level, and dynamically upscales on-premise VMs to the public cloud servers based on cost and performance information. This way, SMEs can optimize their infrastructure and run their VMs where they will be most effective and efficient.

The remainder of this paper is as follows. The next section outlines the benefits and challenges of hybrid cloud computing for SMEs. Subsequently, related work in the field is presented in Section III, after which we define the objectives, design and implementation of our virtual infrastructure management tool in Section IV. Section V covers our experimental evaluation and results, after which we summarize the most important conclusions of our work in Section VI.

II. Benefits and Challenges of Hybrid Cloud Computing for SMEs

The hybrid cloud model extends the private cloud model by using both local and remote resources. These remote cloud resources are seamlessly integrated in the private infrastructure. Hybrid clouds are usually used to scale out when the local resources are exhausted, called cloudbursting. This way, companies can create highly elastic environments.

The benefits of hybrid clouds are amongst others [3][11]:

- **Optimal utilization**
  As peak loads can be up to ten times higher than the average load, traditional data centers need to be over-dimensioned, resulting in idle servers during average load and unnecessary costs. Hybrid clouds scale out to the public cloud to handle peak loads, so the private infrastructure can be dimensioned to handle the average case.
- **Risk transfer**
  The risk of downtime is reduced. Whenever there are problems with the private IT infrastructure, the load can be transferred to the public cloud who ensures high uptimes.
- **Availability**
  Hybrid clouds can upscale to the public cloud or even let the public cloud completely take over operations, this way providing high availability without requiring redundancy and geographic dissemination in the private infrastructure.

However, there are also many challenges and issues for hybrid cloud computing [3][11][12], especially when targetting SMEs:

- **Interoperability and vendor lock-in**
  Vendor lock-ins were already identified as a major risk factor in IT outsourcing [13]. Additionally, failures can also hit public cloud providers, even the big players. According to [3], the only solution is using multiple cloud providers, but, again here, this is only possible when vendor lock-ins are avoided.
- **Hybrid cost**
  Hybrid cloud infrastructures have on one hand a setup and operating cost for the private IT infrastructure (such as hardware, power, cooling, maintenance) and, on the other hand, a pay-per-use cost for the public part at the cloud provider. This hybrid cost model makes it hard to reveal and predict the total costs and benefits of an IT investment project.
- **Security**
  Hybrid cloud computing requires solid service level agreements with and trust in the public cloud providers. As the servers are no longer shielded by the company's firewall, other security measures have to be applied.

III. State of the Art

Today, large technology vendors as well as open-source software projects both address the hybrid cloud market and are developing virtual infrastructure management tools to set-up and manage hybrid clouds. VMware's vCloud offers live migration of virtual appliances and machines between data centers and allows service providers to offer IaaS while maintaining compatibility with internal VMware deployments. HP provides three offerings for hybrid cloud computing: HP Operations Orchestration for provisioning, HP Cloud Assure for cost control, and HP Communications

 ISBN: 978-1-61208-153-3

as a Service for service providers to offer small businesses on-demand solutions. Also Amazon reaches towards hybrid cloud models with its Virtual Private Cloud service.

In addition to the large technology vendors, also open-source software projects are providing on-premise and public cloud integration. Eucalyptus Enterprise provides a software infrastructure for on-premise cloud computing and enables to work within VMware environments and provision resources to Amazon Web Services. Ubuntu Enterprise Cloud (UEC) is shifting from Eucalyptus towards OpenStack and is also compatible with Amazons EC2. OpenNebula [7] is another open-source project that supports the dynamic execution of multi-tier services on a distributed infrastructure consisting of both data center resources and remote cloud resources. Nimbus also provides a virtualization framework to help manage cloud deployments for IaaS. Finally, openQRM extended its focus and now also supports public clouds, currently however limited to Amazon EC2.

Although many options are on the table, and many advertise they support hybrid cloud computing, the current initiatives of the large vendors either result in (i) a hypervisor and/or vendor lock-in, (ii) require a separate interface to manage the private and public cloud, or (iii) require additional tools to implement the load balancing. vCloud only supports ESX as hypervisor, UEC Xen and KVM, and openQRM ESX, Xen and KVM as hypervisor. The virtual infrastructure management capabilities of the open-source solutions provide more choice but require a lot of scripting and are difficult to configure and use. If a flexible hybrid cloud is the goal, the choice of the underlying virtualization platform is crucial, putting the open-source solutions afront. As Nimbus only supports a limited number of hypervisors, and Eucalyptus is more appropriate for private clouds [7][11], OpenNebula and openQRM are the best options today [14].

However, although OpenNebula and openQRM are the best options today, important features are missing like monitoring VM instances or retrieving the VM's IP addresses in order to implement advanced load balancing mechanisms. Also, a graphical interface is either missing or not user-friendly. As general speaking, SMEs are lacking behind in adoption of IT compared to larger companies [15], a good, user-friendly, vendor-independent virtual infrastructure management tool is needed to help SMEs efficiently set-up and manage hybrid clouds. To our knowledge, we are the first to provide such a tool that allows automatic load balancing between the private and public clouds at the VM level, and dynamically upscales on-premise VMs to the public cloud servers based on cost and performance information.

IV. Design of the virtual infrastructure management tool

Below the objectives, general concept and implementation of the virtual infrastructure management tool are described.

A. Objectives

Three requirements for hybrid clouds will be fulfilled: firstly, the hybrid cloud should be able to autonomously handle different load patterns, including peak demands, balance the load and upscale when needed to the public cloud. Secondly, the hybrid cloud should be transparent for the VMs in the infrastructure. This makes sure that the hybrid cloud logic needs only to be implemented in our tool, and standard virtualization software can be used on the VMs. And thirdly, the hybrid cloud architecture should be easy and user-friendly to set-up and configure. Human-platform interaction for configuring the hybrid cloud should be straightforward by using user-friendly user interfaces so that training and dependency on external IT expertise can be minimized. Current hybrid cloud initiatives do not fulfill these requirements and therefore a new tool is designed and presented in this paper.

B. General concept

Consider the use case of the SME Nieuws.be to illustrate the general concept and show how current SME data centers can be optimized in order to gain competitive assets to the public cloud. Nieuws.be is an internet company that aggregates and distributes national and international news on the web, collected by the redaction or contributed by one of their readers. By providing filtering and news-on-demand, they fulfill new market requirements and gain market share within the widely spread news business. As a result, they need an infrastructure to support their heavily visited news site.

In the traditional way, Nieuws.be buys and maintains an infrastructure of six load balanced web servers and one database. Using virtualization and by setting up a hybrid cloud, they can optimize their infrastructure by consolidating lightly used servers. Using virtualization, three web servers can be ran on a single machine. The database is heavily used, so is virtualized on a dedicated machine. As a result, the same performance can be achieved by only three (heavily used) physical servers. In case of peak demands, the heavily used servers cannot handle the load, and therefore cloudbursting is used to automatically allocate additional resources in the public cloud. Additionally, as Nieuws.be is a local news site, the request pattern also depends strongly on the time of the day. As can be seen in Figure 2, during night hours, requests to Nieuws.be fall back to a minimum. Therefore, during down times, the hybrid cloud can further optimize by migrating lightly used VMs to a single host, allowing to shutdown one of the servers.

One can immediatly see that this hybrid cloud architecture results in a two-fold cost reduction: on one hand, only three instead of six servers need to be bought; on the other hand, the infrastructure has a huge energy saving. Besides the cost benefit of consuming less energy, it additionally has the social benefit of reducing their energy footprint.

 ISBN: 978-1-61208-153-3

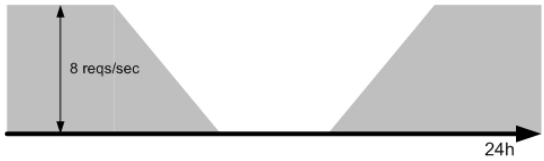

Figure 2. Request pattern of Nieuws.be

Other use cases that illustrate the benefit of hybrid cloud computing to handle peak loads are the Hallmark infrastructure around Valentine's day, or the Colruyt infrastructure in December due to holiday purchases and gifts (see Figure 3).

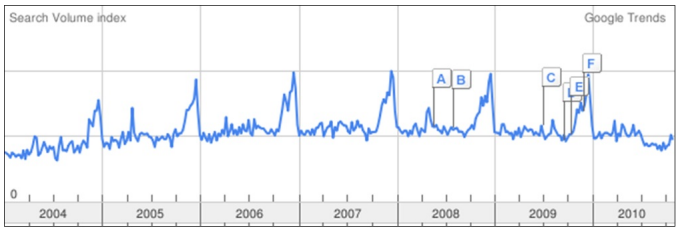

Figure 3. Request pattern of Colruyt

C. Implementation

The main functionality of the virtual infrastructure management tool provides efficient management of hybrid clouds through automatic scaling and load balancing. Figure 4 illustrates the architecture for the virtual infrastructure management tool.

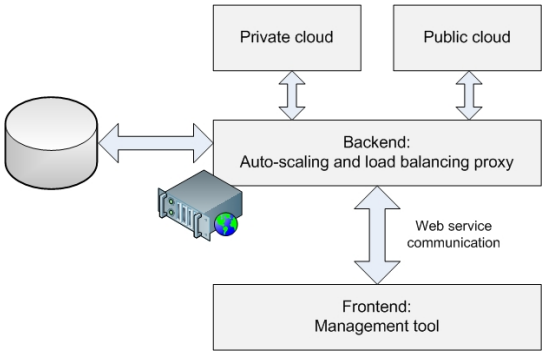

Figure 4. Architecture of the virtual infrastructure management tool

As can be seen in this figure, two parts make up the virtual infrastructure management tool: on one hand, a proxy that implements different load balancing algorithms and provides configurable thresholds for upscaling to the public cloud; on the other hand, a management interface that visualizes and manages the hybrid cloud, clusters VMs and remotely configures the proxy.

The major task of the proxy is forwarding the incoming requeststs to the appropriate VMs. The proxy supports different load balancing algorithms in order to do so. At the moment round robin and weighted round robin are supported, but more algorithms will follow in the future. Round robin is especially suited for load balancing when the different VMs have (almost) the same performance. If the VMs have different specifications, weighted round robin can be used to compensate for these differences. There, servers are presented requests in proportion to their weighting resulting in fairly distributing the requests amongst VMs, instead of equally distributing the requests. To support the weighted load balancing, the performance of all VMs needs to be monitored and the thresholds for up- and downscaling need to be configured. As each public cloud instance type differentiates itself from the others in terms of price, number of virtual cores, available memory and I/O bandwidth, these pricing and performance models are used to derive the weights for the weighted round robin load balancing and can also be used to implement more advanced load balancing and upscaling algorithms in the proxy.

The management interface, presented in Figure 5, provides a tab for visualizing the current VMs in the infrastructure, both private and public ones, as well as the functionality to start and stop these VMs. New VMs can also be added, and clusters can be generated. A cluster is a group of VMs providing the same service. Each cluster can use its own load balancing and scaling settings. The management interface also provides a tab to configure the proxy's load balancing algorithms and scaling thresholds.

Both the proxy and management tool are implemented using C# in combination with the .NET Framework. In order to fulfill the defined requirements above, it is important that the VMs in the private cloud can be addressed the same way as the VMs in the public cloud in order to simplify the hybrid cloud management. Therefore, the proxy is implemented on a dedicated VM in the cloud without an interface, and uses a plug-in system to communicate with different private and public cloud providers. In order to support the automatic upscaling, the VMs also need to be monitored. In order to do so, the APIs of the different cloud providers are used. All requests sent to the proxy are load balanced and forwarded to the according VM (private or public); the response however is directly sent to the client, skipping the proxy. The communication between the management tool and the proxy is implemented using Web service technology and WCF. In order to adjust the proxy settings, the user can use the graphical management interface which sends Web service requests to the proxy in order to configure the thresholds and load balancing algorithms. At the backend, a database is used to store the properties of the hosts, clusters, VMs, scaling thresholds and load balancing constraints. This database is updated by periodically monitoring the VMs in the background. The database is implemented using SQL Server and the Entity Framework. This way, the relational structure of the database is abstracted and one can directly

 ISBN: 978-1-61208-153-3

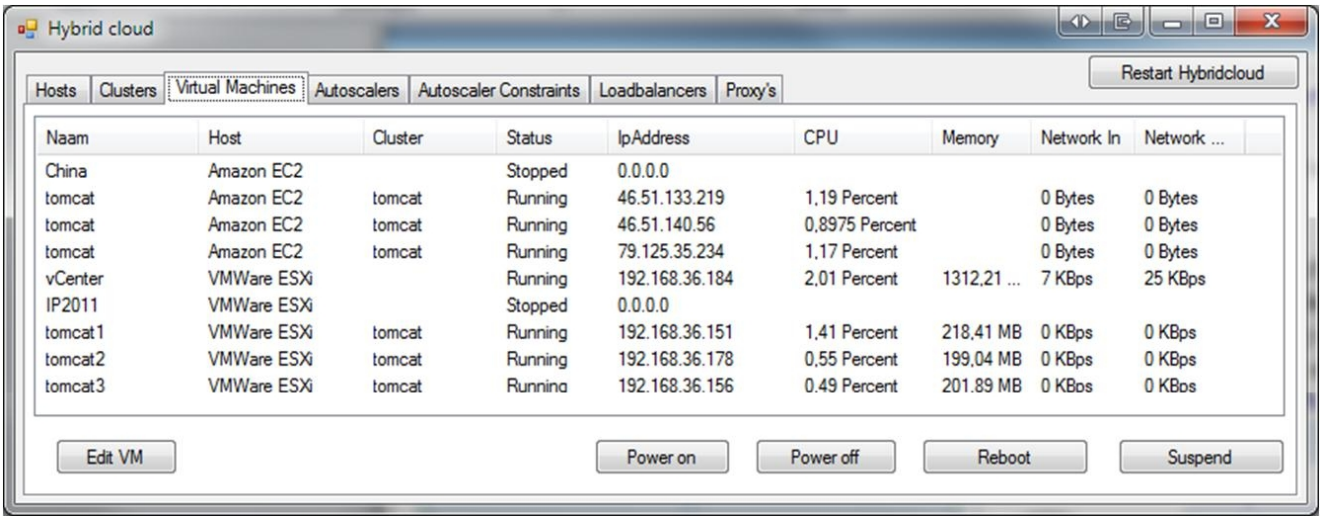

Figure 5. Screenshot of the management tool presenting the infrastructure overview

work with the created object classes.

Currently, VMware is supported for the private part of the hybrid cloud, and Amazon for the public part of the hybrid cloud. More cloud providers will be added in the future. For the public cloud part, the Amazon EC2 service is used to manage the public VMs, and the Amazon CloudWatch service is used to monitor the status of the VMs. The Amazon AWS API can be used starting from Microsoft .NET Framework v2.0. The VMware vSphere API is similar to the Amazon AWS API for .NET, but then applied to a VMware cloud. In order to set up the hybrid cloud, the VMware vSphere PowerCLI API is used. This PowerCLI offers in the first place an interface to Windows PowerShell (which is a new and advanced command-line shell for Windows), but the dlls can also be imported in .NET projects resulting in the same functionality being available using programming code.

V. Evaluation and Results

The components of the virtual infrastructure management tool have been implemented and are currently being evaluated.

Figure 5 presents a screenshot of the management interface where the current infrastructure is visualized, presenting all private and public VMs and their properties. Each of these VMs can be started, stopped, rebooted or suspended. The management tool also provides tabs to add additional VMs to the infrastructure or to create service clusters. The proxy configuration can also be done in the graphical tool: thresholds for up- and downscaling can be configured (see Figure 6) and a load balancing algorithm can be selected and tuned.

The operation of our virtual infrastructure management tool was verified through an experimental performance study. During the evaluation, we started with two web servers in the private cloud and upscaled to maximum six web servers in the public cloud, resulting in an infrastructure consisting of eight VMs. The number of requests per second

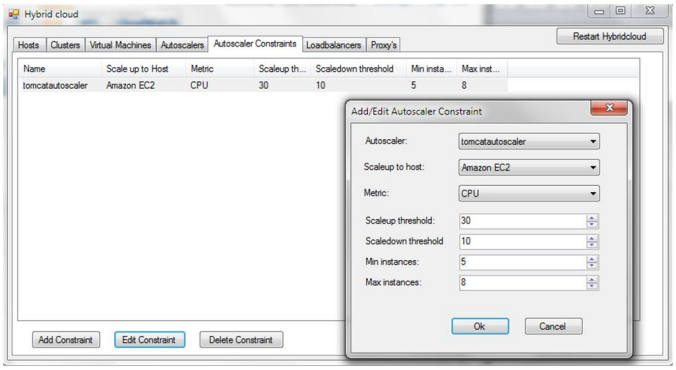

Figure 6. Screenshot of the management tool presenting the scaling tresholds configuration

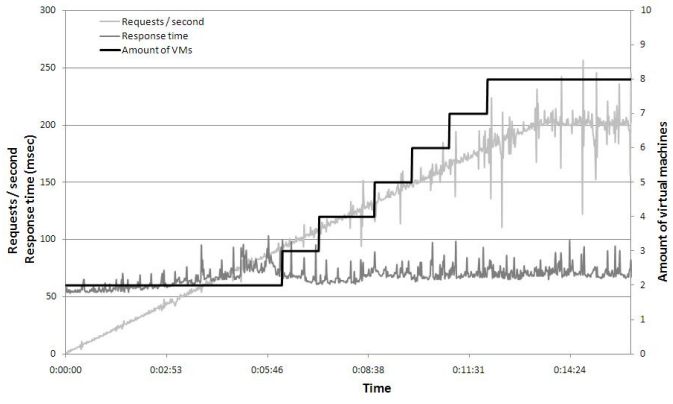

Figure 7. Average response time and amount of VMs in function of the requests per second

was increased over time. The test results are presented in Figure 7. As can be seen in this figure, the response time initially increases as the requests per second increases. When the upscaling thresholds are reached, VMs are added, stopping the response time from further increasing as the load is now balanced over an expanded infrastructure.

VI. CONCLUSIONS

This paper highlights the opportunities of hybrid cloud computing for SMEs and presents a virtual infrastructure management tool that can be used by SMEs to set-up and manage their hybrid cloud. Different reasons can drive SMEs to maintain their own data center instead of becoming "data center-less".

Providing a tool to easily set-up and manage these hybrid clouds takes into account the technical possibilities, the SMEs perspective, and the economic tradeoff between the different business models such as classic data centers, private cloud computing and public cloud computing. We are aware of the fact that these hybrid clouds are not the best solution for every SME. If the restrictions of the public cloud not apply and the SMEs only have a limited IT infrastructure and expertise, then outsourcing to public clouds can be very interesting due to economy of scale.

Results clearly illustrate that current SME data centers can be optimized to compete with the public cloud. As the number of SMEs far exceeds the number of large and very large organizations in almost every country all over the world, this approach results in interesting business benefits. By using hybrid clouds, SMEs critical or latency sensitive applications are kept on the infrastructure (collocated or not) in which they have already invested, and applications are moved toward cloud computing enabled data centers in order to handle occasional peak requests. This methodology allows SMEs to freeze capital investments and move applications toward cloud computing enabled data centers. Hybrid cloud computing may therefore become a very important competitive feature of SME data centers to leverage the economies of scale that the "public clouds" offer.

We will continue the design of more advanced load balancing algorithms by taking into account the different pricing and performance models. Future work also includes the development of additional plug-ins for our proxy so on one hand more private and public cloud providers are supported, and on the other hand also security can be managed. Especially the latter is challenging as classical security models may be insufficient as data is replicated and distributed in potentially worldwide infrastructures [8].

REFERENCES

[1] N. Leavitt, Is Cloud Computing Really Ready for Prime Time?, Computer, 42(1), pp. 15-20, 2009.

[2] D. Hilley, Cloud computing: A taxonomy of platform and infrastructure-level offerings, Georgia Institute of Technology, Technical Report GIT-CERCS-09-13, April 2009.

[3] M. Armbrust, A. Fox, R. Griffith, A.D. Joseph, R. Katz, et al., Above the Clouds: A Berkeley View of Cloud Computing, Technical Report, February 10, 2009.

[4] J.A. Welsh and J.F. White, A Small Business Is Not a Little Big Business, Harvard Business Review, 59(4), pp. 18-32, 1981.

[5] N.A. Sultan, Reaching for the cloud: How SMEs can manage, International Journal of Information Management, 31, pp. 272-278, 2011.

[6] J. Bolton, Small Firms: Report of the Committee of Inquiry on Small Firms, Cmnd 4811, HMSO, London, 1971.

[7] B. Sotomayor, R.S. Montero, I.M. Llorente, and I. Foster, Virtual Infrastructure Management in Private and Hybrid Clouds, IEEE Internet Computing, 13(5), pp. 14-22, 2009.

[8] K. Jeffery and B. Neidecker-Lutz, The Future of Cloud Computing: Opportunities for European Cloud Computing Beyond 2010, Expert Group Report, European Commission, 2010.

[9] J. Thong, C.S. Yap, and K.S. Raman, Top management support, external expertise and information systems implementation in small businesses, Information Systems Research, 7(2), pp. 248-267, 1996.

[10] C.S. Yap, C. Soh, and K.S. Raman, Information systems succes factors in small business, Omega - The International Journal of Management Science, 5(6), pp. 597-609, 1992.

[11] P.C. Heckel, Hybrid Clouds: Comparing Cloud Toolkits, Seminar Paper Business Informatics, University of Mannheim, 2010.

[12] S. Fraser, R. Biddle, S. Jordan, K. Keahey, B. Marcus, et al., Cloud computing beyond objects: seeding the cloud, Proceeding of the 24th ACM SIGPLAN Conference on Object oriented programming systems languages and applications, pp. 847-850, 2009.

[13] B.A. Aubert, M. Patry, and S. Rivard, A Framework for Information Technology Outsourcing Risk Management, Database for Advances In Information Systems, 36(4), pp. 9-28, 2005.

[14] P. Sempolinski and D. Thain, A Comparison and Critique of Eucalyptus, OpenNebula and Nimbus, Proceedings of the IEEE International Conference on Cloud Computing Technology and Science, pp. 417-426, 2010.

[15] J.G. Devos, H. Van Landeghem, and D. Deschoolmeester, IT/IS and Small and Medium-Sized Enterprises: Literature Overview, paper submitted for publication in the Journal of Small Business Management.

 ISBN: 978-1-61208-153-3