Business Models for Cloud-based High Perfomance Computing Service Provision
Insights from the Swiss Higher Education Sector

Markus Eurich and Roman Boutellier
D-MTEC, ETH Zurich, Scheuchzerstrasse 7, 8092 Zurich, Switzerland

Keywords: Business Model, Cloud Computing, High Performance Computing, Higher Education, Public Sector.

Abstract: In 2013, the Swiss Federal Institute of Technology in Zurich (ETH Zurich) and the University of Zurich jointly set up a Cloud stack in order to experiment with high performance computing (HPC) service provision in higher education institutions and its corresponding business model alternatives. This project demonstrated that Cloud-based service provisioning is possible for HPC and can be applied to big data problems as well. On this basis and against the background of new public management reforms, this study aims to foster the understanding of the business model aspects: value proposition and revenue mechanisms. Therefore, 14 interviews were conducted on the potential use of Cloud HPC services and revenue mechanisms. The results show that HPC service providers appreciate Cloud computing providing shorter time to service and more customized services; and eventually becoming more transparent and efficient, i.e. complying with new public management concepts. However, the service consumers do not see a real need to consume Cloud-based services as there is hardly any "Cloud-only" application at the moment. Finally, the three revenue mechanisms 'pay per use', subscription, and 'pay for a share' are discussed.

1 INTRODUCTION

The emergence of Cloud computing contributed a great deal to the digitization of the public sector (e.g., Chandrasekaran and Kapoor, 2011; Kundra, 2010; Lifka et al., 2013). In 2011, Norwich University's College of Graduate and Continuing Studies conducted a survey on Cloud computing among government and higher education institution professionals: the results show that almost half of the respondents indicated that their organizations are in the process of implementing Cloud computing services (Norwich University, 2011). There are already some successful cases of Cloud computing in the public sector: Australia's national science agency virtualized its business applications so that they can be managed and shared across all its locations; the regional government of Castilla in Spain is using Cloud-based services to accelerate the rollout of e-government applications for taxes and driving licenses; and the Chinese University of Hong Kong centralized its data center and network resources on a private Cloud platform (Macias and Greg, 2011). However, the public sector is still significantly lagging behind the private sector in terms of Cloud deployments (Baldwin, 2012). In a 2011 study on the future of Cloud computing in the public and private sectors, over 1,500 interviews were conducted with professionals from organizations in Europe, North America, and Asia. The interviews showed that only 23% of public sector organizations are using Cloud-based hosted data or remotely hosted apps compared to 42% of the organizations in the private sector. The study indicates that European organizations are particularly slow in adopting Cloud services and appear to be behind Asian and US organizations (Red Shift Research, 2011).

This study now aims to provide insights on, and discusses some implications of, the use and implementation of Cloud computing in the European higher education sector as part of the broader public sector. The goal is to foster the understanding of business models for Cloud-based high performance computing services in higher education.

Before the research questions are given, the terms "Cloud computing", "high performance computing", and "business model" are defined for this paper.

Eurich M. and Boutellier R..
Business Models for Cloud-based High Perfomance Computing Service Provision - Insights from the Swiss Higher Education Sector.
DOI: 10.5220/0005050801010110
In *Proceedings of the 11th International Conference on e-Business* (ICE-B-2014), pages 101-110
ISBN: 978-989-758-043-7

Cloud Computing refers to the use of computing resources, which are available in a remote location and accessible over a network. Therefore, the Cloud is an operational model, a usage model, and a business model. Cloud computing services are often divided into three layers: software as a service (SaaS), platform as a service (PaaS), and infrastructure as a service (IaaS) (Eurich, Giessmann, Mettler, and Stanoevska-Slabeva, 2011; Vaquero, Rodero-Merino, Caceres, and Lindner, 2009; Weinhardt et al., 2009).

A *Business Model* describes an organization's value creation, proposition, and capture (McGrath, 2010; Osterwalder and Pigneur, 2010; Teece, 2010). Value creation includes resources, activities, and business partners. Value proposition refers to the benefits that a customer can expect from the product or service. Value capture is comprised of revenue streams and pricing (Osterwalder and Pigneur, 2010; Timmers, 1998). In this study, value capture focuses on the revenue mechanism (cf. e.g., Hedman and Kalling, 2003; Johnson, Christensen, and Kagermann, 2008).

High Performance Computing (HPC) refers to all computations that need high processing power or memory capacity. HPC uses resources that are optimized for a massive parallel workload computing in the back-end. The HPC service consumer typically interacts with the HPC resources only via a front-end system (Calleja et al., 2010; Reuther and Tichenor, 2006).

HPC assists researchers in solving complex problems in a variety of different areas like weather forecasts, earthquake simulations, biomedicine, nanotechnology, materials science, environmental modeling, and disaster simulation (Calleja et al., 2010). This selection of HPC applications shows that HPC can be important for the public sector. However, deploying and maintaining HPC resources is expensive and knowledge-intense. As most public HPC services are consumed by members of higher education institutions and as they also typically possess the necessary knowledge and experience, many higher education institutions run their own HPC infrastructure. Today, HPC service provisioning is almost exclusively organized at the institutional level. Even though most higher education institutions in Switzerland are directly, or at least indirectly, controlled by the government and paid for with tax money, they cannot provide each other with HPC services. This can be a problem when HPC infrastructures are specialized for specific purposes: a researcher from institution A cannot use the service from institution B although institution B possesses the most adequate HPC resources for A's problem. The leading higher education institutions' resources are the most exhaustive while some smaller universities cannot afford any HPC resources at all. However, they cannot just use or buy some of the resources that belong to other higher education institutes. One reason is that the institutional and funding structures are very heterogeneous. There are two additional major obstacles: first, the value proposition of services provided via a Cloud solution compared to the traditional way of service provisioning is unclear. Second, there is no pricing mechanism to charge other institutions. In order to discuss the first obstacle, we describe the case of tests on a private Cloud infrastructure that were conducted in the course of the Swiss Academic Compute Cloud project (Kunszt et al., 2013); to discuss the latter obstacle, revenue mechanisms are discussed in this study. The need for reasonable revenue mechanisms will become increasingly important in the course of new public management reforms. Higher education institutions are typically public, non-profit organizations. However, new public management reforms require these institutions to adopt for-profit management concepts in order to have higher accountability, transparency, and efficiency (De Boer, Enders, and Leisyte, 2007; Schubert, 2009). It may be reasonable to assume that higher education institutions could provide Cloud-based HPC services to other public sector institutions in the long run. In this way, the higher education institutions could provide services to, e.g., biomedicine or disaster and earthquake simulations, public organizations like hospitals or regional governmental institutions.

We use two preconditions for the study: first, Cloud-based service provisioning is possible for HPC. For instance, SGI Cyclone is a supercomputer on demand that provides elastic, scalable, and cost transparent services and that gives a service consumer immediate access to the resources and computing capabilities (SGI, 2014). Some Clouds support virtual machines that have several hundred cores and a considerable amount of memory. Second, the Cloud model can be applied to big data problems as well. If a remote Cloud should be used, the data transfer might be cumbersome or even prohibitive, but a local Cloud can deal with large data volumes or can even be explicitly designed in order to manage big data problems through Hadoop (Kunszt et al., 2014).

The business model part "value creation" is of rather technical nature and is presented in detail in Kunszt et al. (2013; 2014). Pilot tests were conducted in a private Cloud, which is a Cloud infrastructure that is operated for a single organization. In the course of the Swiss Academic Compute Cloud project (Kunszt et al., 2013; 2014), more sophisticated tests followed in a hybrid Cloud, which is a composition of two or more (in our case private) Clouds, which remain distinct entities but are bound together (cf., Mell and Grance, 2011; Sotomayor, Montero, Llorente, and Foster, 2009). The Swiss Academic Compute Cloud project has access to OpenStack Cloud installations at ETH Zurich, the University of Zurich, SWITCH, and the Zurich University of Applied Sciences (Kunszt et al., 2013; 2014), which manage the Cloud installations as self-run IT centers. As the Cloud installations are operated by self-run IT centers, they are capital intensive and a reasonable life-cycle management must be applied to keep the IT infrastructure up-to-date (Eurich, Tahar, and Boutellier, 2011; Eurich, Calleja, and Boutellier, 2013). The current Clouds in the project are relatively small: they range from 100 to 400 Central Processing Unit (CPU) cores each, but are equipped with quite decent memory and storage. The choice of OpenStack as a reference Cloud software stack has emerged from an evaluation done in the Swiss Academic Compute Cloud project (Kunszt et al., 2013).

The relation between technological innovations and management decisions is complex, but both aspects must be aligned and they can iteratively influence each other (Hedman and Kalling, 2003): the same technology commercialized in different ways may result in different economic outcomes (Chesbrough, 2010). Against the technological background (Kunszt et al., 2014), this paper is dedicated to the managerial aspects. The research questions focus on two major business model parts: value proposition and revenue mechanism.

- *Value proposition*: What are the benefits of scientific Cloud-based HPC services?
- *Revenue mechanism*: How can Cloud-based HPC services be priced?

These questions are discussed for both the service consumer and the service provider.

To this end, this article is structured as follows. The next section clarifies the research methodology. The following two sections describe the results of our study and are structured in accordance with the research questions: value proposition (section 3) and pricing mechanisms (section 4). After these descriptions, the results are discussed in section 5. The paper concludes with a summary and an outlook on future research.

2 METHODOLOGY

The aim of this study is to analyze two major business model components: value proposition and revenue mechanism for Cloud computing services that are meant to be jointly provided by and accessible to several higher education institutions. The study was conducted as part of the Swiss Academic Compute Cloud project (Kunszt et al., 2013).

For the analyses of value proposition and revenue mechanism, the study was based on an inductive qualitative research design (Bryman and Bell, 2007; Creswell, 2013). Information was gathered by means of semi-structured interviews. Eight interviews were conducted with academic service consumers, who are currently using HPC service that are provided in a traditional manner. They are the potential buyers of Cloud-based HPC services and are in charge of making investments and taking decisions. Six interviews were conducted with service providers at ETH Zurich and at the University of Zurich. Additional information was used, which was collected from interviews with representatives from the Swiss National Supercomputing Centre, the Swiss Initiative in Systems Biology (System X), and the Friedrich Miescher Institute.

In a previous study (Eurich, Calleja, and Boutellier, 2013), revenue streams of HPC services were analyzed; and as a project-internal pre-assessment step, three revenue mechanisms were identified as being acceptable in terms of economic sustainability and convenience: ‘pay per use’, subscription, ‘pay for a share’.

The interviews with scientific service consumers at ETH Zurich and at the University of Zurich were all conducted in 2013; six were conducted face-to-face and two via phone. Information was gathered by interviewing research groups, which are currently using central computing services. The interviewed service consumers were asked what they would use the computing capacity for; how they perceive the pricing approaches from a service consumer perspective ('pay per use', subscription, 'pay for a share'); and what advantages and disadvantages they expect from these approaches.

The interviews with the service providers at ETH Zurich and at the University of Zurich were also conducted in spring 2013: four were conducted face-to-face and two by phone. These interviews also followed an interview guide. The service providers were asked in what way Cloud Computing resources could improve service provisioning, how they perceive the pricing approaches from a service provider perspective ('pay per use', subscription, 'pay for a share'), and what are advantages and disadvantages they expect from these approaches.

Obviously, some questions were the same for both the service consumers and providers. This was done on purpose in order to reveal a potential difference in perception between consumers and providers.

The data gathered from the interviews was analyzed using open, axial, and selective coding techniques (Urquhart, 2001). The extracted key statements and assertions were then grouped along the research questions and interviewee categories (service consumers, service providers, HPC experts), which resulted in a grid that allowed to identify patterns (Campbell, 1966).

3 VALUE PROPOSITION

Several studies have analyzed the value propositions of Cloud computing in the public sector (e.g., Forst & Sullivan, 2011; Kundra, 2010; Macias and Greg, 2011). The identified benefits of Cloud computing services for the public sector include amongst others: simple scalability, labor optimization, capital expenditure reduction, fast deployment, assured service levels, access to up-to-date technology, and reduced maintenance effort. However, it is noticeable that the identified value proposition for the public sector is more or less the same as for the private sector (Baldwin, 2012).

With a particular focus on scientific Cloud applications, a large-scale user survey revealed several benefits of Cloud computing services like computing elasticity, data elasticity, and rapid prototyping (Lifka et al., 2013). Like the more general studies on Cloud computing for the public sector, this survey does not differentiate between advantages for service consumers and for service providers. A reason for this lack of discrimination might be ascribed to the issue that public sector institutions are typically only perceived as service consumers of Cloud computing services.

In our case, however, the public sector organizations are not only service consumers, but also service providers. Therefore, we aimed to gain insights into service consumers' (3.1) and service providers' (3.2) perception of Cloud computing benefits.

3.1 Service Consumer Perspective

Service providers need to understand the needs of, and the number of potential service consumers. The service consumers reported that they could mainly use Cloud computing services for:

- *Testing and experimenting*: So far, academic service consumers see the major benefit of Cloud service in conducting tests and experiments on the Cloud infrastructure. In this way, they would use Cloud services only in a pre-phase of an actual research project. With the tests, service consumers aim to produce preliminary results that they can use to write a fact-based research proposal in order to get a grant to buy their own infrastructure.
- *Training for students*: In a scientific context, senior staff is often not very pleased to see juniors and students experimenting with their high-end, sometimes fragile IT infrastructure. Therefore, they would appreciate Cloud computing services that are totally separated from the operational computing resources. Students could use this test environment in the Cloud to gain some experience. For the same reason, workshops and classes for students could also be conducted on Cloud resources. Cloud resources are particularly useful and convenient when workshops take place infrequently; in this case the teacher does not need to spend much time for setting up a test and demonstrating the IT environment.
- *Some special applications*: Only a few consumers see a need for Cloud computing resources for particular applications. Cloud applications mentioned in the interviews include medical IT services, like on-the-fly services during surgery or ultrasound image recognition, or some sort of easy simulations.
- *Storage*: Scientific service consumers are particularly interested in Cloud-based storage services, which would allow them to access their data whenever and from wherever they want. However, they would very much appreciate a trustworthy and reliable European storage service. Trustworthiness and reliability are demanded because the users are worried about their sensitive data. They fear a potential data and knowledge leakage as well as being spied on. Recent laws and regulations, like the US Patriot Act, spurred further unease and uncertainty among the academic users.

3.2 Service Provider Perspective

The results of the interviews show that the picture is in fact different when you ask service providers about the usefulness of Cloud services. Regardless of the Cloud computing layer, they expect that Cloud computing could deliver the following value to the consumers:

- *Flexibility*: By making use of a Cloud-based service provision approach, the service provider can decouple the service provisioning and the actual hardware infrastructure that the services depend on. The provider can develop the underlying infrastructure with no or very little impact on the services provided, which results in a much higher service availability as well.
- *Provisioning of additional services*: Cloud resources would give service providers the opportunity to react faster to consumers' demands and also to provide services that are currently not in the provider's service portfolio. With Cloud computing, service providers could draw on ready-made solutions from the Cloud, assembling their services based on customer demand. For example, applications that run only on certain operating systems that are currently not supported can be added with ease.
- *Time to service*: Cloud resources have the potential of improving the time to service for the users. Users could get the service immediately. Currently, the users have to wait weeks or even months until the service is set up, tested, and ready to be consumed. When users require a specific service, it can take the service provider a long time to provide it, particularly when additional resources need to be procured. Cloud computing could help to bridge the time until the service is ready.
- *Self-serving aspect and increased automation*: Cloud computing solutions could be provided directly to the users. At least some experienced users could consume the services from the Cloud, which could increase the level of automation and reduce the time and money for administrative overhead.
- *Elasticity*: Service providers would be given the chance to define the capacity of services provided in accordance with the actual demand in a short amount of time. Cloud computing enables a dynamic scaling up and down; when less capacity is needed or service consumers do not require a particular service anymore, the Cloud should provide the option to give the capacity and resources back when not needed.
- *Balance workload*: Linked to the option to pay for services only on demand, the workload could be balanced better, i.e., Cloud resources can be used for topping up capacity and for boosting the capacity in times of peak demand in the short run.

4 REVENUE MECHANISMS

In a previous study (Eurich et al., 2013), revenue streams of Cloud computing services were analyzed; in a project-internal pre-assessment three revenue mechanisms have been identified as being acceptable in terms of economic sustainability and convenience:

- *'Pay per use'*: Service consumers are charged a fee according to the time and volume of a computing service that has been consumed.
- *Subscription*: The service consumer pays a fee on a regular basis for the usage of a service. Subscriptions allow services to be sold in bundles.
- *'Pay for a share'*: Service consumers buy a share in order to get a corresponding amount of service.

Both, service consumers (4.1) and service providers (4.2) were asked for an assessment of the three revenue mechanisms, 'pay per use', subscription, and 'pay for a share'.

4.1 Perception of Revenue Mechanisms by Service Consumers

4.1.1 Pay per Use

Pro

Service consumers appreciate the 'pay per use' revenue mechanism as a fair approach. Costs are transparent and you only pay for what you get. Service consumers could imagine to 'pay per use' for some kind of small jobs, like for testing or experimenting with a service. Especially when a research activity contains uncertainties, service consumers do not want to spend their money upfront. In this case, they neither want to buy a package that lasts for a month (like in the subscription model) nor do they want to spend money upfront (like in the 'pay for a share' scheme). The 'pay per use' scheme gives them the freedom to initiate research activities and to immediately quit the ones they do not want to pursue any longer. Service consumers perceive the possibility of terminating the consumption and the payments of a service advantageous to other schemes, e.g., the 'pay for a share' scheme. They do not want to subsidize other users due to being locked in a contract when they do not need any further services. Finally, the 'pay per use' revenue mechanism sounds particularly attractive to service consumers, who develop applications themselves, i.e., who operate more on the PaaS or IaaS layer. These researchers often come from computer science or physics departments. The 'pay per use' model would give them the opportunity to flexibly and quickly make use of just the services they need.

Contra

On the other hand, service consumers do not want to be bothered with billing and accounting. They do not want to spend time on administrative tasks like checking invoices and accounting activities. The 'pay per use' revenue mechanism is perceived as expensive and not paying off for a longer time of service consumption. Finally, academic service consumers are concerned about acquiring money. In the current academic budget allocation model it is almost impossible to get money to be spent on a 'pay per use' basis. Researchers cannot ask for money if they do not know on what it will be spent. Currently, they can only successfully receive grants when they submit a proposal for funding in which they ask for money that will be spent on a specific hardware or to buy a share.

4.1.2 Subscription

Pro

Academic service consumers acknowledge 'subscription' to be fair. Like the 'pay per use' scheme it is easy to understand and you pay for what you get. Some interviewees intuitively liked it because it is also the way they pay for their private mobile or smart phone services and Internet connections. They like the idea of having some predefined packages from which they can choose and would be fond of having the option to upgrade to another package if necessary. Service consumers consider this a flexible and easy (because there are predefined packages) approach.

Contra

The same arguments as for the 'pay per use' scheme hold also true for the subscription approach. There would be fewer invoices and accounting activities than in the 'pay per use' scheme, but it would still be less convenient than the 'pay for a share' solution. Problems with the current academic budget allocation model would also arise when applying the subscription revenue mechanism.

4.1.3 Pay for a Share

Pro

Among our three revenue mechanisms, the 'pay for a share' is perceived as the most convenient by academic service consumers. They only need to pay once and then are all set for the next couple of years. Researchers can focus on their actual research and do not have to care about comparing prices, squaring accounts, and other kinds of accounting tasks. An interviewee stated, “maybe in the end it does not pay off, but despite that you pay for convenience. You do not need to pay every time you need a service. If you have some capacity left you can do some extra research”. Moreover, it is well aligned with current academic budget allocation models and, therefore, it might be relatively easy to get money for buying a share for services in a Cloud stack.

Contra

Some of the arguments that favor the 'pay per use' and the subscription approach can be interpreted as disadvantages of the 'pay for a share' approach. For example, a service consumer mentioned, “if you do not make full use of your share, you subsidize others and waste money for something that you do not need”. In addition, there might be a loss in flexibility to change the volume or the kind of service consumption compared to the 'pay per use' and the subscription approach. Finally, the renewal of a share might be a problem because this means a massive investment at one time.

4.2 Perception of Revenue Mechanisms by Service Providers

4.2.1 Pay per Use

Pro

First, the providers of academic services acknowledge that the ‘pay per use’ scheme could provide an added value for their service users. They find the ‘pay per use’ scheme a feasible approach when users do not know their demand and when service consumption is only required for a few weeks. In the communication with the service consumer, this approach is probably the easiest to explain because everyone basically knows how it works. Since it is very transparent what is provided and what is used, reporting and accounting is also easily understandable.

Second, an advantage for the service providers themselves could be a better handling of spare capacity. This approach could be particularly interesting in case of external organizations' demand.

Contra

A disadvantage of the 'pay per use' scheme is that service consumption is highly unpredictable. A service provider mentioned that "if you want to give your consumers the choice to consume a service whenever they want, service delivery becomes difficult to plan and schedule". In the scientific environment, the situation is probably even more unpredictable than on the free market. Unlike the free market, at a university with a self-run IT center you may only have a limited number of users, which may have a demand at the same time, e.g., right after the examination period or at the beginning of a new term. An interviewee described that "on the free market, you may have a higher number of users from different regions and with different needs, and as a result demand might be better balanced. At the university you need to invest a lot upfront into hardware if you want to offer your consumers an immediate response to their service requests". However, due to the high unpredictability, there is no guarantee that the investment is amortized.

Like the service consumers, the providers also assume an extra effort for billing and accounting. Some algorithms need to be set up to monitor the users' consumption and to calculate the price of it. However, once done successfully, billing and monitoring could be automated and there is no extra effort for this anymore.

4.2.2 Subscription

Pro

Compared to the 'pay per use' scheme, the subscription model is characterized by a more stable and predictable source of income. This helps service providers in calculating and forecasting revenues and expenses and provides them with a long-term frame for planning, adjustments, and procurements.

Contra

The subscription scheme features some unpredictability about users' demand in terms of quantity and type of services, which is a problem for higher education institutions with capital intensive self-run IT centers. Especially if users are given the opportunity to flexibly up- or downgrade to another package, some big upfront investments are necessary to guarantee a high level of service availability and a reduced service time. Another problem is that service providers of higher education institutions typically cannot move money from one year to the next; this means that even if a massive increase in service consumption is predicted, the service provider can only invest in its IT infrastructure or service portfolio once it has got the money.

4.2.3 Pay for a Share

Pro

The 'pay for a share' scheme is in line with the current budget allocation model. In addition, service providers appreciate that this model gives shareholders a sense of ownership and community. Compared to the 'pay per use' and the subscription scheme, the 'pay for a share' approach guarantees a certain degree of income that lasts for three to four years, the typical life cycle of hardware.

Contra

On the downside, there is no transparent service fee. It is not intuitively clear what and how much a consumer gets. The service consumer may get confused what the share is worth: it can be time on the Cloud stack, performance, or another service (like storage, consulting). The 'pay for a share' approach is subject to unpredictability. This approach is particularly prone to drop outs. Each shareholder contributes a considerable amount of money at the beginning of a long-term shared ownership. This initially paid sum is higher than a monthly fee in the subscription model or a daily fee in the 'pay per use' approach, because it is meant to last for a much longer time; as a consequence, this sum for the share can have quite some impact on the service portfolio or on the Cloud infrastructure. One problem is that a major chunk of revenues comes from new professors, who may get a share as bonus for joining the new university. However, it is quite unpredictable whether they are able to pay for a renewal once the lifecycle of the hardware is over.

5 SUMMARY AND DISCUSSION

Cloud computing has emerged as new way to digitize the public sector. However, the public sector lags considerably behind the private sector in terms of Cloud deployments. To facilitate the Cloud-enabled digitization of the public sector, research is needed on both the service consumption and the provisioning side.

The perceived advantages and disadvantages of different models of Cloud-based HPC service consumption/delivery are summarized in Table 1.

Table 1: Perception of different models of Cloud-based HPC service consumption/delivery.

	pay per use	subscription	pay for a share
Service consumers	+ trans-parent cost – admini-stration	+ fair and well known – budget allocation	+ con-venience – subsidy of other users
Service providers	+ easy to sell – not predictable	+ stable income – big upfront investment	+ long-term income – renewal of hardware

On the service consumption side, the interviews show that scientific service consumers want to focus on their research tasks. They expect an easy approach and want the higher education institution to clear obstacles from their paths. The service consumers currently perceive Cloud services as a playground for testing, experimenting, and training students. Up until now, most scientific service consumers could not see any useful Cloud computing applications. Indeed, a really convincing Cloud application that could make scientific consumers move to the Cloud is missing. As there is no "Cloud-only" application at the moment, consumers do not see a real need to consume Cloud-based services instead of traditional services. Decisions are based on beliefs not on facts. “Cloud computing doesn’t pay off for us,” reported an interviewed service consumer, and several other HPC service consumers join him in the preception that Cloud computing services are too expensive for the amount of CPU they need. Some service consumers reported that they would use Cloud services if they come at a reasonable price. In the end, they care little about pricing strategies. The three revenue mechanisms, ‘pay per use’, subscription, and ‘pay for a share’, are perceived differently among the interviewees and there is no clearly preferred revenue mechanism. Many service consumers do not know much about Cloud computing and, in fact, do not even care much about what computing resources are used and how they are priced. An interviewed service consumer put it straight: "I just need something powerful to run” [my computations on].

On the service provisioning side the interviews revealed that the service providers tend to perceive Cloud-based services as an additional resource in the short run, but not as a replacement of traditional HPC service provision. A service provider assumes that “Cloud provisioning is [currently] most useful when a user needs 10,000 or more processors for only a week to a maximum of three months”. Service providers appreciate Cloud-based HPC service provisioning for its flexibility, its elasticity, its self-serving aspect accompanied with an increased automation, the chance to provide additional services, the potential to shorten the time to service, and the opportunity to balance the workload.

This study focuses on the private and hybrid Cloud services provisioning of self-run IT centers of higher education institutions. Recently, ETH Zurich carefully opened up towards the public Cloud in times of peak consumption even though there are concerns about confidentiality and security. Consuming additional computing power from the public Cloud relieves the pressure from university IT centers’ decision makers to purchase equipment and manage the infrastructure. In case of high usage unpredictability, decisions can be postponed. This development promotes the pay-per-use pricing scheme, which is particularly burdened with the disadvantages of usage unpredictability in the case of self-run private Cloud centers.

6 OUTLOOK AND CONCLUSION

Currently, Cloud-based HPC services are actually not necessary from a service provisioning perspective because powerful private infrastructures already exist. Therefore, there is a lack of motivation to establish and invest into Cloud-based HPC services. It remains to be seen to what extent Cloud-based HPC service provision can reap the benefits that service providers expect. However, the Cloud is real; it is here and it is growing. Higher education instutions might be well advised to at least gather some experiences with the Cloud because the IT infrastructure has become an essential requirement in attracting the best reseachers (Drucker, 2002). Not uncommonly, applications only emerged years after a new technology was introduced, e.g., computer simulations (Drucker, 1999). We assume that currently Cloud computing is only a means to optimize service provisioning (cf., Pring, 2010) while truely innovative applications on the basis of the Cloud may only emerge later.

In order to reap optimization benefits, government bodies must put some incentives in place or enforce public institutions to move to the Cloud by the means of new public management reforms. However, the realization of cost and service advantages of the Cloud requires a holistic approach. Training has to be provided to both the service consumers and the service providers. In addition, government bodies need to support different pricing schemes. For example, at the moment it is almost impossible to get a grant for a research project that is based on a per-pay-use or subscription approach. Finally, there needs to be some regulations, which control the exchange of computing services for money among public institutions. To conclude, Cloud-based service provisioning is most advantageous at an organizational level, but the realization and acceptance depends on the involvement and support from government bodies and the service consumers.

This study sheds light on both the service consumers' and the service provider's opinion on the value of Cloud-based service provision. However, the findings are limited to insights that contribute to facilitate Cloud-based HPC service provisioning among higher education institutions. ETH Zurich has already received inquiries for the usage of HPC services from public and private organizations. Because of technical, legal, and regulatory issues, none of these requests could have been granted. Future research could focus on incentive schemes, legal and regulatory aspects, and technological requirements to enable service exchange among organizations. In this context, some of the findings of this study could be tested and transferred to other public organizations.

We assume that the importance of services will increase at ETH Zurich, especially in the area of HPC (Eurich, Tahar, Boutellier, 2011). Therefore, there might be an overemphasis on service provisioning in the assessment of pricing schemes. Future work could discuss the possible relation between the specific use (e.g., experimentation, storage) and the suitable pricing scheme.

Finally, it should be considered that HPC service provisioning can no longer be subsidized the way it used to be. The rapid increase in the demand of computing resources has pushed higher education institutions and also other public organizations to their limits in terms of computing service provisioning and its financing. Energy costs grow exponentially: the current way of HPC service provisioning must be rethought. Decision makers need to reflect on the different types of consumers and their ties to the national infrastructure.

REFERENCES

Baldwin, H. Public Sector Cloud Computing: The Good, the Bad and the Ugly, *Computerworld*, May 9, 2012, http://www.computerworld.com/s/article/9226932/Public_sector_cloud_computing_The_good_the_bad_and_the_ugly, accessed June 2014.

Bryman, A. and Bell, E., 2007. *Business Research Methods*, Oxford University Press. Oxford, UK, 2nd edition.

Calleja, P., Gardiner, C., Gryce, C., Guest, M., Lockley, J., Parchment, O., and Stewart, I. HPC-SIG Report 2010, *UK High Performance Computing Special Interest Group*, 2010, http://www.hpc-sig.org/Publications?action=AttachFile&do=get&target=HPC-SIG_Report2010.pdf, accessed June 2014.

Campbell, D. T., 1966. Pattern Matching as an Essential in Distal Knowing. In Hammond, K. (ed.), *The Psychology of Egon Brunswik*. Holt, Rinehart and Winston. New York, pp. 81-106.

Chandrasekaran, A. and Kapoor, M. State of Cloud Computing in the Public Sector – A Strategic analysis of the business case and overview of initiatives across Asia Pacific, *Frost & Sullivan*, May 11, 2011. http://www.frost.com/prod/servlet/cio/232651119, accessed June 2014.

Chesbrough, H., 2010. Business Model Innovation: Opportunities and Barriers. *Long Range Planning*, 43(2-3), pp. 354-363.

Creswell, J., 2013. *Qualitative Inquiry and Research Design: Choosing Among Five Approaches*. Sage Publications. Thousand Oaks, CA, 3rd edition.

De Boer, H. F., Enders, J., and Leisyte, L., 2007. Public Sector Reform in Dutch Higher Education: The Organizational Transformation of the University. *Public Administration*, 85(1), pp. 27-46.

Drucker, P. F., 1999. Beyond the Information Revolution. *Atlantic Monthly*, 284, pp. 47-57.

Drucker, P. F., 2002. They're not employees, they're people. *Harvard Business Review*, 80(2), pp. 70-77.

Eurich, M., Calleja, P., and Boutellier, R., 2013. Business Models of High Performance Computing Centres in Higher Education in Europe. *Journal of Computing in Higher Education*, 25(3), pp. 166-181.

Eurich, M., Giessmann, A., Mettler, T., and Stanoevska-Slabeva, K., 2011. Revenue Streams of Cloud-based Platforms: Current State and Future Directions. Proceedings of the 17th Americas Conference on Information Systems (AMCIS 2011). Detroit, MI, August 4-7, 2011.

Eurich, M., Tahar, S., and Boutellier, R., 2011. Effizienzdruck und technologische Innovation im Hochschul-IT Management: Strukturwandel der ETH-Informatikdienste. *Hochschulmanagement*, 2/2011, pp. 36-41.

Hedman, J. and Kalling, T., 2003. The Business Model Concept: Theoretical Underpinnings And Empirical Illustrations. *European Journal of Information Systems*, 12(1), pp. 49-59.

Johnson, M. W., Christensen, C. M., and Kagermann, H., 2008. Reinventing Your Business Model. *Harvard Business Review*, 86(12), pp. 57-68.

Kundra, V. State of Public Sector Cloud Computing, *CIO Council Indonesia*, May 20, 2010, http://dosen.narotama.ac.id/wp-content/uploads/2012/01/State-of-Public-Sector-Cloud-Computing.pdf, accessed June 2014.

Kunszt, P., Maffioletti, S., Flanders, D., Eurich, M., Schiller, E., Bohnert, T. M., Haug, S., et al., 2014. Towards a Swiss National Research Infrastructure. Proceedings of the Euro-Par 2013: Parallel Processing Workshops, Aachen, Germany, August 26-30, 2013, Springer. Heidelberg, pp. 157-166.

Kunszt, P., Maffioletti, S., Messina, A., Flanders, D., Mathys, S., and Murri, R., 2013. Academic Cloud Provisioning and Usage Project, https://wiki.systemsx.ch/display/cloudresult, accessed June 2014.

Lifka, D., Foster, I., Mehringer, S., Parashar, M., Redfern, P., Stewart, C., and Tuecke, S. XSEDE Cloud Survey Report. *Extreme Science and Engineering Discovery Environment*, September 2013, https://www.cac.cornell.edu/technologies/XSEDECloudSurveyReport.pdf, accessed June 2014.

Macias, F. and Greg, T. Cloud Computing Advantages in the Public Sector. *Cisco*, 2011, https://www.cisco.com/web/strategy/docs/c11-687784_cloud_omputing_wp.pdf, accessed June 2014.

McGrath, R., 2010. Business Models: A Discovery Driven Approach. *Long Range Planning*, 43(2-3), pp. 247-261.

Mell, P. and Grance, T., 2011. The NIST Definition of Cloud Computing. *US Department of Commerce, National Institute of Standards and Technology*, NIST Special Publication 800-145, September 2011, http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf, accessed June 2014.

Norwich University. A Pulse on Virtualization & Cloud Computing, *Norwich University*, April 2011, http://www.thecre.com/fisma/wp-content/uploads/2011/05/ Norwich_Survey_Findings1.pdf, accessed June 2014.

Osterwalder, A. and Pigneur, Y., 2010. *Business Model Generation: A Handbook for Visionaries, Game Changers, and Challengers*. Wiley. Hoboken, NJ.

Pring, B. Cloud Computing: The Next Generation of Outsourcing, Gartner, November 10, 2010, http://www.gartner.com/DisplayDocument?id=1460416, accessed June 2014.

Red Shift Research. Adoption, Approaches & Attitudes: The Future of Cloud Computing in the Public and Private Sectors, *Red Shift Research*, June 2011, http://whitepaper.techweekeurope.co.uk/adoption-approaches-attitudes-the-future-of-cloud-computing-in-the-public-and-private-sectors-340.html, accessed June 2014.

Reuther, A. and Tichenor, S., 2006. Making the Business Case for High Performance Computing: A Benefit-Cost Analysis Methodology. *CTWatch Quarterly*, 2(4A), pp. 2-9.

Schubert, T., 2009. Empirical Observations on New Public Management to Increase Efficiency in Public Research - Boon or Bane?. *Research Policy*, 38(8), pp. 1225-1234.

SGI. SGI Cyclone: Results On Demand, *Silicon Graphics International Corp.,* 2010, https://www.sgi.com/pdfs/4205.pdf, accessed June 2014.

Sotomayor, B., Montero, R. S., Llorente, I. M., and Foster, I., 2009. Virtual infrastructure management in private and hybrid clouds. *Internet Computing, IEEE*, 13 (5), pp. 14-22.

Teece, D. J., 2010. Business Models, Business Strategy and Innovation. *Long Range Planning*, 43(2-3), pp. 172-194.

Timmers, P., 1998. Business Models for Electronic Markets. *Electronic Markets*, 8(2), pp. 3-8.

Urquhart, C., 2001. An Encounter with Grounded Theory: Tackling the Practical and Philosophical Issues. In Trauth, E. (ed.), *Qualitative Research in IS: Issues and Trends*. Idea Group Publishing. London, UK, pp. 104-140.

Vaquero, L. M., Rodero-Merino, L., Caceres, J., and Lindner, M., 2009. A Break in the Clouds: Towards a Cloud Definition. *ACM SIGCOMM Computer Communication Review*, 39(1), pp. 50-55.

Weinhardt, C., Anandasivam, A., Blau, B., Borissov, N., Meinl, T., Michalk, W., and Stößer, J., 2009. Cloud Computing: A Classification, Business Models, and Research Directions. *Business & Information Systems Engineering*, 1(5), pp. 391-399.