Revenue Streams and Value Propositions of Cloud-Based High Performance Computing in Higher Education

Markus Eurich(✉) and Roman Boutellier

Department of Management, Technology, and Economics, ETH Zurich, Scheuchzerstrasse 7, 8092 Zurich, Switzerland
meurich@ethz.ch, roman.boutellier@sl.ethz.ch

Abstract. Most higher education institutions can no longer subsidize academic high performance computing (HPC) services the way they used to. New business models are needed. At the same time, Cloud computing has emerged as a new way to digitize the public sector, but thus far there is only little experience in this domain. Therefore, ETH Zurich and the University of Zurich jointly set up a Cloud stack to experiment with HPC service provision and corresponding business model alternatives. On this basis of an interview series, this study aims to foster the understanding of Cloud-based HPC services and revenue streams. The results suggest that service providers appreciate Cloud computing as a means to become more transparent and efficient, i.e. to comply with new public management concepts. However, service consumers can hardly see the need to consume Cloud-based services, because convincing "Cloud-only" applications are still missing. Different revenue streams are discussed.

Keywords: Business model · Cloud computing · High performance computing · Higher education · Public sector

1 Introduction

In higher education HPC service provisioning can no longer be subsidized the way it used to be. There is a rapid increase in the demand of computing resources which has pushed higher education institutions and also other public organizations to their limits in terms of computing service provisioning and its funding. Recently, Cloud computing has emerged as a new way to digitize services in the public sector [1–3]. In 2011, Norwich University's College of Graduate and Continuing Studies conducted a survey on Cloud computing among government and higher education institution professionals: almost half of the respondents indicated that their organizations are in the process of implementing Cloud computing services [4]. There are already some successful cases of Cloud computing in the public sector: 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; the Chinese University of Hong Kong centralized its data center and network resources on a private Cloud platform; and Australia's national science agency virtualized its business applications so that they can be managed and

M. S. Obaidat and A. Holzinger (Eds.): ICETE 2014, CCIS 554, pp. 61–75, 2015.
DOI: 10.1007/978-3-319-25915-4_4

shared across all its locations [5]. However, the public sector is still significantly lagging behind the private sector in terms of Cloud deployments [6]. 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, the United States, 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 [7].

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 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:

Cloud computing refers to the use of computing resources, which are available in an outside 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) [8–10].

A *business model* describes an organization’s value creation, proposition, and capture [11–13]. 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 comprises revenue streams and pricing [13, 14]. In this study, value capture focusses on the revenue streams [15, 16].

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 [17, 18].

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 [17]. 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 which 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, and second, there is no pricing mechanism to charge other institutions. In order to discuss the first obstacle, we describe the case of our private Cloud stack and in order to discuss the latter obstacle, revenue streams are discussed in this study. The need for reasonable revenue streams 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 [19, 20]. 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 to disaster and earthquake simulations, to 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 [21]. Some Clouds support virtual machines that have several hundred cores and a lot 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 [22].

The business model part “value creation” is of rather technical nature and is presented in detail [22, 23]. 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 [22, 23] 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 bundled [24, 25]. 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 [22, 23], which manage 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 [26, 27]. 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 as an evaluation done in the Academic Cloud Provisioning and Usage project [23].

The relation between technological innovations and management decisions is complex, but both aspects must be aligned and they can iteratively influence each other [15]: the same technology commercialized in different ways may result in different economic outcomes [28]. Against the technological background [22], this paper is dedicated to the managerial aspects. The research questions focus on two major business model parts: revenue stream and value proposition:

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

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: revenue streams (Sect. 3) and value proposition (Sect. 4). After these descriptions, the results are discussed in Sect. 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 streams for Cloud computing services that are meant to be jointly provided by and accessible to several higher education institutions.

For the analysis of revenue stream and value proposition, the study is based on an inductive qualitative research design [29, 30]. 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 that was collected from interviews with representatives from the Swiss National Supercomputing Centre (CSCS), the Swiss Initiative in Systems Biology (System X), and the Friedrich Miescher Institute in Basel.

In a previous study [27], revenue streams of HPC services were analyzed; and as a project-internal pre-assessment step, three revenue streams 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 for they would use the computing capacity; 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 how 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 [31]. 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 [32].

3 Revenue Streams

In a previous study [27] revenue streams of Cloud computing services were analyzed; and in a project-internal pre-assessment three revenue streams 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", i.e. a part of infrastructure or even a part of a complete package of software, hardware and services in order to get a corresponding amount of service.

Both, service consumers (Sect. 3.1) and service providers (Sect. 3.2) were asked for an assessment of the three revenue streams, 'pay per use', subscription, and 'pay for a share'.

3.1 Perception of Revenue Streams by Service Consumers

3.1.1 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 stream 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 services they just need.

Contra. On the other side, 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 long time of service consumption. Finally, academic service consumers are concerned about acquiring money. In the current academic budget allocation model you cannot get money to be spent on a 'pay per use' basis. Researchers cannot ask for money if they do not know for 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 is about to be spent on a specific hardware or for buying a share.

3.1.2 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 could choose and would be fond of having the option to upgrade to another package if needed. Service consumers consider this a flexible, yet 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 stream.

3.1.3 Pay for a Share

Pro. Among the three revenue streams, the 'pay for a share' is perceived as the most convenient by the academic service consumers. They only need to pay once and then they 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. 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. On the downside, 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, 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 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.

3.2 Perception of Revenue Streams by Service Providers

3.2.1 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. 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 peak-demand at the same time, e.g. right after the examination period or at the beginning of a new term. 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 into hardware upfront 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 for it. However, once done successfully billing and monitoring could be automated and there is no extra effort for this anymore.

3.2.2 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 time 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 institutes with capital intensive self-run IT centers. Especially if users are given the opportunity to flexibly up- or down-grade to another package, some big upfront investment are necessary to guarantee a high level of service availability and low time to service. Another problem is that service providers can typically not move money from one year to the next: this means, 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.

3.2.3 Pay for a Share

Pro. The 'pay for a share' scheme is in line with most current budget allocation models. In addition, service providers appreciate that this model gives shareholders a sense of ownership and community. They are the owner of a part of the big computer. 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; and as a consequence this sum for the share can have quite some impact on the service portfolio or on the Cloud infrastructure. A problem is that a major chunk of 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 life-cycle of hardware is over.

4 Value Proposition

Several studies have analyzed the value propositions of Cloud computing in the public sector [2, 3, 5]. 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 are more or less the same as for the private sector [6].

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 [1]. 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' (Sect. 4.1) and service providers' (Sect. 4.2) perception of Cloud computing benefits.

4.1 Service Consumer Perspective

Service providers need to understand the needs 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 like to 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 and expensive infrastructure. Therefore, they would appreciate Cloud computing services which are separated from their operational computing resources. Students could use this test environment in the Cloud to get 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 a surgery or ultrasound image analysis; or some sort of simple simulations.
- *Storage*: Scientific service consumers are particularly interested in Cloud-based storage services, which would allow them to access their data when and 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 privacy. They fear a potential data and knowledge leakage. Recent laws and regulations, like the US patriotic act, spurred further unease and uncertainty among the academic users.

4.2 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 de-couple 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 the 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 in house 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.
- *Scaling*: 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 on the short run.

5 Summary and Discussion

The emergence of Cloud computing contributed much to the digitization of 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.

On the service consumption side the interviews show that scientific service consumers want to focus on their research tasks and want to reduce their administrative load. They expect an easy approach and want the higher education institution to clear

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	+ Fair and well known	+ Convenience
	– Administration	– Budget allocation	– Subsidy of other users
Service providers	+ Easy to sell	+ Stable income	+ Long term income
	– Not predictable	– Big upfront investment	– Renewal of hardware

obstacles from their paths. The service consumers currently perceive Cloud services as a playground for testing, experimenting, and training students. Up to until now, most scientific service consumers cannot see any useful Cloud computing applications. Indeed, a really convincing Cloud application that could make scientific consumers move to the Cloud is missing; even the coverage of peak is not perceived as a big advantage. 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. We should never forget that research is not under the same time and cost pressure as the business world. Researchers are used to wait.

"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 consumer reported that they would use Cloud service if they come at a reasonable price. In the end, they care little about pricing strategies and they typically do not consider the full cost, because education institutes subsidize. 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]. Computers are cheap compared to personnel cost and in many cases separately paid for in grants.

A special concern is confidentiality: many researchers distrust the secret services of big nations and want to keep their data under their direct control. They still believe that data stored on their own computer are much safer than data stored in the Cloud, even if their computers are connected to the Internet.

On the service provisioning side the interviews revealed that the service providers tend to perceive Cloud-based services as an additional resource on 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 peak demand for a week to a maximum of three months". Service providers sell Cloud-based HPC service provisioning with the arguments of 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 focusses on 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. Most institutions trust Moore's law: machines get smaller, cheaper and more powerful. Thus delay helps to reduce investments. 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 Conclusions

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 which extent Cloud-based HPC service provision can reap the benefits that service providers expect. A general increase in cost pressure may be needed to speed up acceptance. However, the Cloud is real, it is here and it is growing. Higher education institutions 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 researchers [33]. Not uncommonly, applications only emerge years after a new technology has been introduced, like e.g. computer simulations [34]. We assume that currently Cloud computing is only a means to optimize service provisioning [35], while truly innovative applications on the basis of the Cloud may only emerge later.

In order to reap optimization benefits, governmental bodies must put some incentives in place or enforce public institutions to move to the Cloud, e.g. with the help 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, governmental bodies need to support different pricing schemes: 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 need to be some regulations which enable the paid exchange of computing services among public institutions. There seems to be no way around better cost transparency. To conclude, Cloud-based service provisioning is most advantageous on an organizational level, but the realization and acceptance depends on the involvement and support from governmental bodies and the service consumers.

This study sheds light on both the service consumers' as well as on the service providers' opinion on the value and about revenue streams of Cloud-based services. 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 enquiries for the usage of HPC services from public and private organizations. Because of technical, legal, and regulatory issues, only few of these requests have been granted. Future research could focus on incentive schemes, legal and regulatory aspects, and technological requirements to enable service exchange among organizations. Another challenge for future research is the specialization of big computers through co-design for specific problems. Top performance for a reasonable amount of money can only be achieved through a systematic balance of hardware and software. The general purpose computer cannot provide the power needed to solve the most complex problems e.g. in chemistry and climate science. Finally, 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 [26]. Therefore, there might be an overemphasis on service provisioning in the assessment of revenue streams. Future work could discuss the possible relation between the specific use (e.g., experimentation, storage) and the suitable revenue stream.

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 to their limits in terms of computing service provisioning and its financing. Decision makers need to reflect on the different types of consumers and their ties to the national infrastructure.

References

1. Lifka, D., Foster, I., Mehringer, S., Parashar, M., Redfern, P., Stewart, C., Tuecke, S.: XSEDE Cloud Survey Report (2013)
2. Kundra, V.: State of Public Sector Cloud Computing (2010). http://dosen.narotama.ac.id/wp-content/uploads/2012/01/State-of-Public-Sector-Cloud-Computing.pdf
3. Chandrasekaran, A., 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 (2011). http://www.frost.com/prod/servlet/cio/232651119
4. Norwich University: A Pulse on Virtualization & Cloud Computing. (2011). http://www.thecre.com/fisma/wp-ontent/uploads/2011/05/Norwich_Survey_Findings1.pdf
5. Macias, F., Greg, T.: Cloud Computing Advantages in the Public Sector. Cisco (2011). https://www.cisco.com/web/strategy/docs/c11-687784_cloud_omputing_wp.pdf
6. Baldwin, H.: Public Sector Cloud Computing: The Good, the Bad and the Ugly (2012). http://www.computerworld.com/s/article/9226932/Public_sector_cloud_computing_The_good_the_bad_and_the_ugly
7. Red Shift Research: Adoption, Approaches & Attitudes: The Future of Cloud Computing in the Public and Private Sectors (2011). http://whitepaper.techweekeurope.co.uk/adoption-approaches-attitudes-the-future-of-cloud-computing-in-the-public-and-private-sectors-340.html
8. Eurich, M., Giessmann, A., Mettler, T., Stanoevska-Slabeva, K.: Revenue streams of cloud-based platforms: current state and future directions. In: Proceedings of the 17th Americas Conference on Information Systems (AMCIS 2011), Detroit, MI (2011)

9. Weinhardt, C., Anandasivam, A., Blau, B., Borissov, N., Meinl, T., Michalk, W., Stößer, J.: Cloud computing: a classification, business models, and research directions. Bus. Inf. Syst. Eng. **1**, 391–399 (2009)
10. Vaquero, L.M., Rodero-Merino, L., Caceres, J., Lindner, M.: A break in the clouds: towards a cloud definition. ACM SIGCOMM Comput. Commun. Rev. **39**, 50–55 (2009)
11. McGrath, R.G.: Business models: a discovery driven approach. Long Range Plan. **43**, 247–261 (2010)
12. Teece, D.J.: Business models, business strategy and innovation. Long Range Plan. **43**, 172–194 (2010)
13. Osterwalder, A., Pigneur, Y.: Business Model Generation: A Handbook for Visionaries, Game Changers, and Challengers. Wiley, Hoboken (2010)
14. Timmers, P.: Business models for electronic markets. Electron. Markets **8**, 3–8 (1998)
15. Hedman, J., Kalling, T.: The business model concept: theoretical underpinnings and empirical illustrations. Eur. J. Inf. Syst. **12**, 49–59 (2003)
16. Johnson, M., Christensen, C., Kagermann, H.: Reinventing your business model. In: Harvard Business Review on Business Model Innovation, pp. 47–70. Harvard Business Press, Boston (2008)
17. Calleja, P., Gardiner, C., Gryce, C., Guest, M., Lockley, J., Parchment, O., 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_Report
18. Reuther, A., Tichenor, S.: Making the business case for high performance computing: a benefit-cost analysis methodology. CTWatch Q. **2**, 2–9 (2006)
19. De Boer, H.F., Enders, J., Leisyte, L.: Public sector reform in dutch higher education: the organizational transformation of the university. Public Adm. **85**, 27–46 (2007)
20. Schubert, T.: Empirical observations on new public management to increase efficiency in public research - boon or bane? Res. Policy **38**, 1225–1234 (2009)
21. SGI: SGI Cyclone: Results On Demand. Silicon Graphics International Corp, Fremont, CA (2014). https://www.sgi.com/pdfs/4205.pdf
22. Kunszt, P., Maffioletti, S., Flanders, D., Eurich, M., Schiller, E., Bohnert, T.M., Edmonds, A., Stockinger, H., Jamakovic-Kapic, A., Haug, S.: In Proceedings of Euro-Par 2013: Parallel Processing Workshops Aachen, pp. 157–166. Germany (2014)
23. Kunszt, P., Maffioletti, S., Messina, A., Flanders, D., Mathys, S., Murri, R.: Academic Cloud Provisioning and Usage Project. Zurich, Switzerland (2013). https://wiki.systemsx.ch/display/cloudresult
24. Sotomayor, B., Montero, R.S., Llorente, I.M., Foster, I.: Virtual infrastructure management in private and hybrid clouds. IEEE Internet Comput. **13**, 14–22 (2009)
25. Mell, P., Grance, T.: The NIST Definition of Cloud Computing. NIST Special Publication 800-145. US Department of Commerce, National Institute of Standards and Technology (2011). http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf
26. Eurich, M., Tahar, S., Boutellier, R.: Effizienzdruck und technologische Innovation im Hochschul-IT Management: Strukturwandel der ETH-Informatikdienste. Hochschulmanagement 2/2011, pp. 36–41 (2011)
27. Eurich, M., Calleja, P., Boutellier, R.: Business models of high performance computing centres in higher education in Europe. J. Comput. High. Educ. **25**, 166–181 (2013)
28. Chesbrough, H.: Business model innovation: opportunities and barriers. Long Range Plan. **43**, 354–363 (2010)
29. Bryman, A., Bell, E.: Business Research Methods. Oxford University Press, Oxford (2007)
30. Creswell, J.W.: Qualitative Inquiry And Research Design: Choosing Among Five Approaches. Sage Publications, Thousand Oaks (2012)

31. Urquhart, C.: An encounter with grounded theory: tackling the practical and philosophical issues. In: Qualitative Research in IS: Issues and Trends, pp. 104–140 (2000)
32. Campbell, D.T.: Pattern matching as an essential in distal knowing. In: Hammond, K.R. (ed.) The Psychology of Egon Brunswik, pp. 81–106. Holt, Rinehart and Winston, New York (1966)
33. Drucker, P.F.: They're not employees, they're people. Harvard Bus. Rev. **80**, 70 (2002)
34. Drucker, P.F.: Beyond the information revolution. Atlantic Mon. **284**, 47–59 (1999)
35. Pring, B.: Cloud Computing: The Next Generation of Outsourcing. Gartner (2010). http://www.gartner.com/DisplayDocument?id=1460416