

**This is an electronic reprint of the original article.
This reprint *may differ* from the original in pagination and typographic detail.**

Author(s): Mazhelis, Oleksiy; Fazekas, Gabriella; Tyrväinen, Pasi

Title: Impact of Storage Acquisition Intervals on the Cost-Efficiency of the Private vs. Public Storage

Year: 2012

Version:

Please cite the original version:

Mazhelis, O., Fazekas, G., & Tyrväinen, P. (2012). Impact of Storage Acquisition Intervals on the Cost-Efficiency of the Private vs. Public Storage. In R. Chang (Ed.), Proc. of the Fifth IEEE International Conference on Cloud Computing (Cloud 2012) (pp. 646-653). IEEE Computer Society. IEEE International Conference on Cloud Computing. <https://doi.org/10.1109/CLOUD.2012.101>

All material supplied via JYX is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

Impact of Storage Acquisition Intervals on the Cost-Efficiency of the Private vs. Public Storage

Oleksiy Mazhelis, Gabriella Fazekas, and Pasi Tyrväinen
Department of Computer Science and Information Systems
University of Jyväskylä
Jyväskylä, Finland
mazhelis@jyu.fi, fazekas.gabriella@gmail.com, pasi.tyrvaainen@jyu.fi

Abstract—The volume of worldwide digital content has increased nine-fold within the last five years, and this immense growth is predicted to continue in foreseeable future reaching 8ZB already by 2015. Traditionally, in order to cope with the growing demand for storage capacity, organizations proactively built and managed their private storage facilities. Recently, with the proliferation of public cloud infrastructure offerings, many organizations, instead, welcomed the alternative of outsourcing their storage needs to the providers of public cloud storage services. The comparative cost-efficiency of these two alternatives depends on a number of factors, among which are e.g. the prices of the public and private storage, the charging and the storage acquisition intervals, and the predictability of the demand for storage. In this paper, we study how the cost-efficiency of the private vs. public storage depends on the acquisition interval at which the organization re-assesses its storage needs and acquires additional private storage. The analysis in the paper suggests that the shorter the acquisition interval, the more likely it is that the private storage solution is less expensive as compared with the public cloud infrastructure. This phenomenon is also illustrated in the paper numerically using the storage needs encountered by a university back-up and archiving service as an example. Since the acquisition interval is determined by the organization’s ability to foresee the growth of storage demand, by the provisioning schedules of storage equipment providers, and by internal practices of the organization, among other factors, the organization owning a private storage solution may want to control some of these factors in order to attain a shorter acquisition interval and thus make the private storage (more) cost-efficient.

Keywords—storage services; cloud storage; cost model; acquisition interval.

I. INTRODUCTION

According to IDC, the global volume of digital content exhibits exponential growth, and from 1.8ZB in 2011 it will grow to 2.7ZB in 2012 and will reach 35ZB by 2020 [1], [2]. As the volume of digital content grows, the global need for storage capacity rapidly increases, too.

To cope with the growing demand for storage, organizations may proactively build their private storage capacity; in this scenario, the organizations periodically estimate their future demand for storage and then acquire and manage the storage infrastructure internally. Alternatively, the organizations may opt for outsourcing their storage to the providers

of public cloud infrastructure services, such as Amazon Simple Storage Service (S3), Box.com, Apple iCloud, etc. The use of such cloud-based storage services gives the organization the flexibility of rapidly increasing its storage capacity as the demand for the storage grows, as well as the possibility to pay only for the volume of storage the organization actually uses within each charging period. These benefits have rendered the cloud storage services popular among companies: for instance, in October 2011 Amazon reported half a trillion of stored objects in its S3 service, representing an immense 187% combined annual growth rate (CAGR) between 2006 and 2011 [3].

As the cloud infrastructure capacity is usually paid for only when used, the cloud infrastructure providers include the so-called utility (pay-per-use) premium into their pricing [4]. As a result, the unit price per unit of time of a public cloud infrastructure capacity is usually more expensive as compared with the unit cost of the private capacity [4], [5]. Still, if the demand for infrastructure services exhibits periodical or random peaks, the adoption of public cloud infrastructure is likely to offer the organizations cost advantages over the private infrastructure: this is due to the fact that the high premium charged by the public cloud provider is compensated by avoiding the extensive periods of time when the private infrastructure would remain idle [4], [6].

However, as opposite to the fluctuating demand for computing resources, the demand for storage is often accumulating over time, since newly created digital content only partially supersedes the already stored files. As a result, in a longer term, the use of the public storage services may prove more expensive as compared with the private solutions [7].

The cost-efficiency of the public vs. private storage depends on a number of factors, such as the premium charged by the provider of public cloud infrastructure, the charging period (for the public storage) and the storage acquisition interval (for the private storage), the intensity of incurred data communications, the predictability of the growth of storage needs, etc. [7], [8]. In this paper, we study the effect of the private storage acquisition interval on the cost-efficiency of the private vs. public storage. This interval can be determined by the organization’s ability to foresee

the growth in storage demand, by the provisioning schedule of the storage equipment provider, by internal practices of the organization, etc. The paper analytically shows that the shorter the interval at which the organization re-assesses its storage needs and acquires additional storage, the more likely it is that the private storage solution is less expensive as compared with the public cloud infrastructure. Numerical experiments are employed in order to illustrate this dependency using the storage needs encountered by a university back-up and archiving service as an example.

The analysis of storage costs in the paper focuses on the storage needs and their growth and predictability, storage acquisition interval, as well as the costs incurred due to the transfer of data to and from the storage location. It should be noted that the storage costs may be affected also by additional factors, such as the economies and diseconomies of scale, the cost of capital, and the required level of availability and durability. Combined, these and other factors are likely to have a complex, non-linear effect on the overall costs making them difficult to analyze [6]. In order to simplify the analysis, in this paper it is assumed that these additional factors either have a minor effect or affect similarly the costs of both the private and the public storage solutions, and hence are left outside of the scope of the paper.

The remainder of the paper is organized as follows. In the next section, the related works in the domain of the cost-efficient use of cloud infrastructure are overviewed. In section III, an analytical model for comparing the cost efficiency of the private vs. public storage is introduced where the effect of the acquisition interval is taken into account. Numerical experiments illustrating the effect of the acquisition interval and its interplay with some other factors are provided in section IV. Finally, section V summarizes the obtained results and outlines the directions for further work.

II. RELATED WORKS

In recent years, extensive research efforts have been devoted to the issue of cost-efficient usage of cloud infrastructure services in general, and cloud storage services in particular. A short overview of the recent research in this domain is provided below.

A number of works have focused on the cost-efficient use of cloud infrastructure and on the factors affecting it. In particular, the cost benefit of using cloud bursting, i.e. offloading the computing load during peak times to a public cloud infrastructure have been analytically investigated in [4], [9]. The cost efficient allocation of computing load to the private and the public portions of a hybrid cloud infrastructure was also studied in [10], [11] and in [12], [6] where the communication overheads were taken into account as well. It has been shown that, assuming negligible data communication overheads, the cost-optimal time of using public cloud is the inverse of the premium charged by the public cloud provider.

The economies of scale, i.e. the decline in the cost per unit of a service with the number of units produced [13], may as well affect the cost-efficiency of the private vs. public cloud infrastructure. Such economies of scale are manifested e.g. in the volume discount offered for the cloud infrastructure capacity and, in the presence of such discounts, the cost of a hybrid cloud may exceed the cost of a private or a public cloud infrastructure [6].

The cost-optimal allocation of individual computing tasks to the private and public cloud resources was also approached as a multi-integer linear programming problem in [14]. Based on the results of a simulation study, the authors found little or no cost benefits in offloading the peaks of the workload, though the preliminary character of the study and the complex nature of the optimization model make it difficult to interpret the results.

Walker [7] compares the acquisition and leasing of storage as alternative investment decisions based on their Net Present Value (NPV). The estimation of NPV takes into account the dynamics of the demand for storage, the gradual decline of acquired and leased storage prices, the disk replacements due to possible disc failures, and the salvage value of the acquired discs at the end of their use time. Using numerical examples, the authors illustrated that leasing represents the cost-optimal alternative for small and medium sized enterprises, whereas for large enterprises, acquiring storage is likely to be less expensive in a longer term.

Mastroeni & Naldi [15] further revised the Walker's model by replacing the deterministic estimation of the pricing dynamics and disc failure dynamics in [7] with probabilistic models. Based on these models, the authors arrive at a probabilistic distribution of differential NPV values, and use its median to determine the economically justifiable alternative. Note that both in [7] and [15] the costs are accounted on a yearly basis; thus, the role of the acquisition intervals shorter than a year is not visible in these models.

Uttamchandani et al. [16] introduced BRAHMA, a tool that applies constraint-based optimization to cost-optimally supply the storage demand with a mixture of in-house and public cloud storage resources. Based on customer storage needs and projected growth thereof over a lookahead period, as well as associated service level objectives, the tool suggests an optimal placement both for the storage and for the system administrators. The tool helps in identifying the optimal sourcing in case the customer and the storage service provider have a heterogeneous set of devices and human resources having different costs. However, to the best of our knowledge, the tool assumes a perfect knowledge of the customer demand growth and fails to consider the storage acquisition intervals; as a result, the cost of the over-provisioned storage is not visible when using the tool.

Constraint-based optimization is also employed by Trummer et al. [17] to optimally allocate applications, along with their storage resources, to the clouds. The authors' approach

assumes the resource requirements to be known in advance, i.e., similarly to the BRAHMA tool, the effects of imperfect knowledge and resulting storage over-provisioning are not taken into account.

Finally, Weinman [8] considers the delay with which the required resource is provisioned, and analyses both the cost of over-provisioning (i.e. unused resources) and under-provisioning (i.e. the opportunity cost of unserved demand). The author discusses the role of provisioning time given a possibility to predict the future demand over a specific forecast visibility; however, only the cases with zero forecasting visibility were considered in the paper.

In summary, while a number of works focus on the cost-efficient use of the private and public infrastructure resources, relatively little attention has been devoted to the role of the acquisition intervals in the cost-efficient use of private vs. public storage capacity. Therefore, below, a storage cost model is introduced where the effect of the acquisition interval is taken into account.

III. STORAGE COST MODEL

In this section, the cost constituents of alternative storage approaches are considered and their total costs are compared. Different cost constituents need to be taken into account depending on whether the storage solution is owned and managed privately by the organization or is offered by a public cloud infrastructure provider.

For the private storage, the relevant cost constituents include the cost of hardware and software acquisition, integration, configuration, upgrade costs, as well as the recurring costs of renting floor space, power, bandwidth, and the cost of administration and maintenance. Thus, the cost of the private storage is a function of the demand as well as its growth pattern and its predictability, the time interval between storage acquisitions, and the pricing of the needed equipment, software, and personnel, along some other expenses.

On the other hand, for the public cloud storage, the cost constituents include the usage-dependent costs of storage capacity, data transfer, and input/output requests (based on the pricing set by Amazon S3). Depending on the charging policy of the provider, the cost of the storage may be determined by the maximum volume of storage occupied during the charging period: for instance, Amazon Web Services (AWS) offerings apply monthly charges based on the maximum storage capacity used within the month.

Besides the difference in cost constituents, there is an important difference in the way the storage is acquired, provisioned, and charged for. Namely, the private storage needs to be acquired in advance to meet the expected demand growth until the next acquisition time and it incurs volume-dependent costs irrespectively of whether the storage is used or not, whereas the public storage can be deployed virtually instantly as the demand grows and it is charged based on

the volume of the storage actually used within the charging period. The in-house storage needs to be acquired in advance and with some excess depending on how accurately the future may be predicted (not needed in public cloud storage), but on the other hand, the price of a TB in-house can be significantly lower than the pricing in the public cloud. We therefore suggest that the cost efficiency of private vs. public storage depends on both the price difference of the private and public storage, the interval at which the storage can be acquired, and the accuracy with which the future needs for the storage can be predicted.

Let s_i denote the storage needs accumulated over time period i (e.g. during month i). Let $p_p(s_i)$ denote the price of a unit of storage set by the public storage provider, and let $p_o(s_i)$ denote the total cost of owning a unit of private storage capacity over time i . Both of the prices are shown as the functions of the volume of storage capacity used or acquired s_i , to indicate the fact that the prices can be a subject to volume discounts, as is e.g. in the case of AWS storage. Note that $p_p(s_i)$ can be found by consulting price lists of public IaaS vendors, whereas $p_o(s_i)$ needs to be estimated by summing the total costs of acquiring and using the storage over the total period of planned usage T (e.g. the depreciation period) and then deriving the share of the total costs during the time i by dividing the total costs by T .

Let us consider first the case of using private storage capacity. Let us assume that the organization is acquiring private storage capacity with acquisition interval τ . Then, when deciding on how much storage to acquire, the organization needs to predict how much storage it would need within time τ , i.e. until the next acquisition time. For instance, if $\tau = 12$ month, the firm needs to predict how greatly its storage needs are likely to increase within the next year, and acquire the storage accordingly. Then, the cost of acquiring own storage capacity c_o can be estimated as

$$c_o = \hat{s}_{\max} p_o(\hat{s}_{\max}) \tau, \quad (1)$$

where \hat{s}_{\max} is the organization's estimate of the maximum storage needed within the next acquisition interval.

We assume that the firm will acquire the storage capacity sufficient to meet the maximum storage needs. Furthermore, since it is difficult to predict the future storage needs with 100% accuracy, we assume that the organization is likely to over-estimate its storage needs and over-provision its storage capacity, to avoid the situation when it would not be able to meet customer expectations, i.e.

$$\hat{s}_{\max} = k_e k_s s_{\max}, \quad (2)$$

where $k_e \geq 1$ represents an estimation error. The coefficient of redundancy $k_s \geq 1$ is introduced to account for the fact that a portion of storage capacity is used for the purposes other than storing the data - for instance, to maintain the level of redundancy sufficient for the required level of failure-resistance.

Consider now the case of using public storage capacity. Let us assume for simplicity that time interval i corresponds to the charging interval set by the public storage provider; for instance, in case of AWS, i corresponds to the monthly charging interval. Furthermore, we will assume for simplicity that the charging interval is shorter than the organization's acquisition intervals τ by factor N . Then, we can express the length of acquisition interval τ in terms of the charging intervals; for instance, monthly charging periods and yearly acquisition intervals would correspond to $\tau = 12$. Thus, the cost of public storage c_p accumulated over the acquisition interval is:

$$c_p = \sum_{i=1}^N s_i p_p(s_i). \quad (3)$$

Let us also assume that the price of a unit of public storage capacity is greater as compared with the cost of a unit of private storage. This is justified by the fact that the public storage provider charges a premium for the organization's flexibility in rapidly provisioning and de-provisioning the resources [4]; as a result, some organizations found it significantly less expensive to host own storage facilities than to use the storage capacity of Amazon, with the difference up to the factor of 26 [18]. Thus, it can be written that $p_p = u_s p_o$, where u_s is the utility premium ratio, or in short, utility premium of the public storage vendor.

Assuming for simplicity that the prices are not subject to volume discounts, equations (1) and (3) can be rewritten as:

$$c_o = k_e k_s s_{\max} p_o (\hat{s}_{\max}) \tau, \quad (4)$$

$$c_p = u_s p_o \sum_{i=1}^N s_i. \quad (5)$$

To assess whether the public or private storage is less expensive, let us compare the relation c_p/c_o with 1:

$$\frac{c_p}{c_o} = \frac{u_s p_o \sum_{i=1}^N s_i}{k_e k_s s_{\max} p_o \tau} = \frac{u_s}{k_e k_s} \times \frac{\sum_{i=1}^N s_i}{s_{\max} \tau}. \quad (6)$$

Recall that $u_s \geq 1$ and $k_e k_s \geq 1$. Then, assuming that i) a notable premium charged by public storage vendors (not all organizations have the scale and capabilities required to attain the unit storage cost 26 times cheaper than Amazon, but attaining 10-fold savings appears to be a feasible assumption), ii) the estimation error is a fraction of storage needs ($k_e < 2$), and iii) a reasonable degree of overheads is present in own storage (e.g. $k_s \leq 2$), it is likely that $u_s \geq k_e k_s$. Furthermore, since $\sum_{i=1}^N s_i \leq s_{\max} \tau$, it follows that the value of the ratio depends on how great the acquisition interval is, in relation to the charging interval:

$$\begin{aligned} \text{If } \tau \rightarrow 1 &\Rightarrow \frac{\sum_{i=1}^N s_i}{s_{\max} \tau} \rightarrow 1 \Rightarrow \frac{c_p}{c_o} \geq 1; \\ \text{If } \tau \gg 1 &\Rightarrow \frac{\sum_{i=1}^N s_i}{s_{\max} \tau} \ll 1 \Rightarrow \frac{c_p}{c_o} < 1. \end{aligned} \quad (7)$$

Thus, in case the organization's acquisition intervals are significantly longer than the charging periods of the public storage vendor, the use of the public storage is likely to be economically justifiable. On the other hand, if the organization is able to shorten the acquisition intervals to be similar to the intervals of the public storage vendor, then acquiring and maintaining own storage is likely to be less expensive.

It should be noted that besides the costs of storage capacity itself, the cost of a storage solution includes also the costs incurred due to the transfer of data to and from the storage location, namely:

- the initial transfer of new data being saved (that includes also the modified versions of the previously saved items);
- the transfer of stored data back to the user in response to the occasional reading requests (including also the rare retrievals of the backup data).

Cheng et al. [19] analyzed the usage pattern of Youtube videos and modeled the growth of the number of views with a power-law distribution. The authors define the active life span of the videos, stating that after a short period of popularity, the videos are rarely watched again. Therefore we will assume for simplicity that the data is intensively used shortly after it is initially saved, but only occasionally requested afterwards, and that the charging period concurs with the life span of the saved data. Then, the volume of data transfer within a charging period i can be approximated as $k_b \Delta s_i$, where k_b indicates how many times a byte of stored data is transferred in average within a charging interval, and $\Delta s_i = s_i - s_{i-1}$ is the increment in storage needs during charging interval i .

In case of private storage, the price of a unit of bandwidth p_{bo} is likely to depend on the maximum bandwidth required during a charging period [20]. Thus, the private storage transfer costs can be estimated as a function of the maximum storage added within a charging period:

$$c_{bo} = k_b \Delta s_{\max} p_{bo} \tau, \quad (8)$$

where $\Delta s_{\max} = \max_i \Delta s_i$.

The bandwidth costs in case of using a public storage provider, on the other hand, are based on the actual data transfer needs within each charging period. Assuming again that the volume of transferred data is proportional to the volume of data stored by the public storage provider, the cost of data transfer in case of using a public storage provider can be approximated as:

$$\begin{aligned} c_{bp} &= \sum_{i=1}^N k_b \Delta s_i p_{bp} = k_b p_{bp} \sum_{i=1}^N \Delta s_i \\ &= k_b p_{bp} (s_N - s_0) = k_b p_{bp} (s_{\max} - s_0). \end{aligned} \quad (9)$$

Then, the ratio (6) above can be rewritten as:

$$\frac{c_p}{c_o} = \frac{u_s p_o \sum_{i=1}^N s_i + k_b p_{bp} (s_{\max} - s_0)}{k_e k_s s_{\max} p_o \tau + k_b \Delta s_{\max} p_{bo} \tau}. \quad (10)$$

Observe that $s_{\max} - s_0 \leq \Delta s_{\max} \tau$. If we assume that the unit pricing of data communication is roughly equal for the private and the public storage ($p_{bp} \approx p_{bo}$), it follows that

$$k_b p_{bp} (s_{\max} - s_0) \leq k_b \Delta s_{\max} p_{bo} \tau.$$

Furthermore, it can be seen that

$$\text{If } \tau \rightarrow 1 \Rightarrow k_b p_{bp} (s_{\max} - s_0) \approx k_b \Delta s_{\max} p_{bo} \tau; \quad (11)$$

$$\text{If } \tau \gg 1 \Rightarrow k_b p_{bp} (s_{\max} - s_0) < k_b \Delta s_{\max} p_{bo} \tau.$$

Thus, the presence of the data communication costs strengthens the dependency of the ratio $\frac{c_p}{c_o}$ on the length of acquisition interval τ ; in other words, for larger τ , the economic advantage of using public storage increases with the growths of the data communications.

IV. ILLUSTRATIVE NUMERICAL EXAMPLES

In the previous section, it was shown that the interval at which the organization is re-evaluating its storage needs and acquires additional storage capacity affects the cost of the private storage, as compared with the public storage. Namely, the longer the acquisition interval, the more likely it is that the use of the public storage would be less expensive. In this section, the effect of the acquisition interval will be illustrated by using the example of demand profile of the back-up and archiving service provided by the Oxford University to its senior members, postgraduates and staff members [21].

The historical traces of the growth of backup storage provided by the Oxford University Computing Services (OUCS) are documented in the OUCS annual reports available at the OUCS website¹. The growth profile over the period 1996-2011 is shown in Fig. 1. As evidenced by the figure, the demand for the data storage at OUCS exhibits an exponential growth, increasing roughly by 50% on a yearly basis.

With the exception of the first year of observations when a three-digit growth was recorded, the yearly increase during 1998-2011 has been below 100% and in most of the years it was fluctuating between 30% and 70%. Let us therefore assume that the organization acquires the storage capacity sufficient for serving the maximum expected growth in the storage demand, with the maximum expected growth being 100% a year. Let us further assume that the volume of initially acquired capacity is 10TB, that 10% of storage is reserved for redundancy purposes ($k_s = 1.1$), and that the additional capacity is acquired in 5TB chunks.

The unit price of the public storage can be estimated for example by consulting the price list of Amazon S3²:

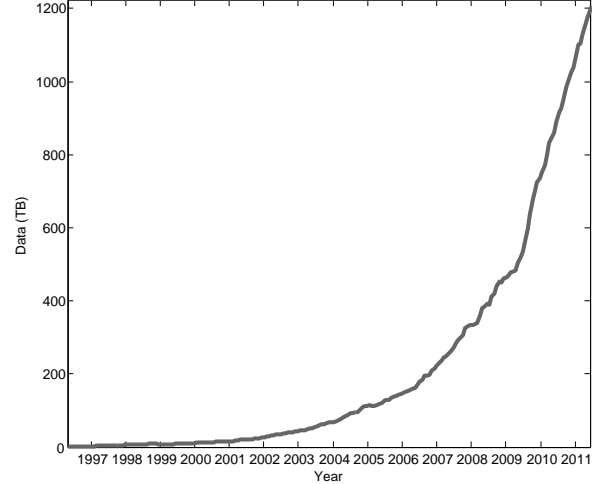


Figure 1. Growth of the OUCS back-up and archiving storage during 1996-2011 [22]

assuming the Reduced Redundancy Storage (RRS) is used, it costs \$0.093, \$0.083, and \$0.073 per GB per month to store the first TB, the next 49TB, and the next 450 TB of data, respectively. Thus, the RRS price per TB per month is \$95.23, \$84.99, and \$74.75 for the first TB, the next 49TB, and the next 450 TB of data, respectively. Note that the request pricing is not taken into account for simplicity.

The unit price of the private storage for newly designed storage solutions can be approximated using the costs incurred by Backblaze [18]: to provision a PB of storage, Backblaze reportedly spends \$94563 over three years for hardware, space, power, and bandwidth. The maintenance costs shall also be accounted for; according to Backblaze, an engineer is maintaining the company's 16PB storage facilities. We will therefore assume that employing an engineer with the yearly salary of \$44973³ will be sufficient to maintain the storage capacity necessary to meet the needs of most of the companies. This results in the total cost of \$162023 per PB over 3 years, i.e. \$6.23 per TB per month. It should be noted that besides the storage hardware, software solutions for managing the storage (such as IBM Tivoli Storage Manager) are likely to be needed, thus increasing the cost of the storage solution further; however, for the sake of simplicity we will assume that either inexpensive or open-source software is going to be used and that its costs may be neglected.

Alternatively, the unit price of the private storage can be found using the charges set by the OUCS back-up and archiving service for its research project customers. According to OUCS service level description [23], the storage cost is £842 (\$1321.5) per TB per year, thus resulting in the storage cost of £70.17 (\$110.32) per TB per month.

¹Available at <http://www.oucs.ox.ac.uk/internal/annrep/>

²Available at <http://aws.amazon.com/s3/>

³Based on <http://swz.salary.com/SalaryWizard/Installation-Maintenance-Technician-I-HRSalary-Details.aspx>

As could be seen based on these two reference examples, the utility premium u_s may vary depending on how cost-efficient the private solution is: e.g. for 100TB of storage, the premium varies from $\$79.97/\$110.32 = 0.7$ (OUCS) to $\$79.97/\$6.23 = 12.8$ (Backblaze). We will therefore explore a set of different values of $u_s = \{0.7; 1.0; 2.0; 4.0; 16.0\}$.

It should be noted that the cost estimates above take into account neither the gradual price decline nor the effect of the net present value of the assets. Both of these are important factors affecting the total cost of a storage solution; however, as their effect has been studied elsewhere [7], in this study, in order to focus on and better illuminate the effect of the acquisition interval on the total costs, we have decided to exclude these factors from the analysis.

In Fig. 2, the total costs of private and public storage solutions accumulated over the period 1996-2011 are compared for different levels of utility premium u_s . As could be seen from the figure, the private storage cost increases along with the storage acquisition interval. Given the utility premium value $u_s \leq 1$ (implying that the cost of using a unit of storage capacity per unit of time is lower in the public cloud), the total cost of the private storage always exceeds the cost of the public storage, in line with the analysis in [4]. On the other hand, given $u_s \geq 1$, the private storage is less expensive when the acquisition interval is short, but becomes more expensive than the public storage as the length of the interval grows, thus supporting the analytical reasoning presented earlier in section III (cf. (7)).

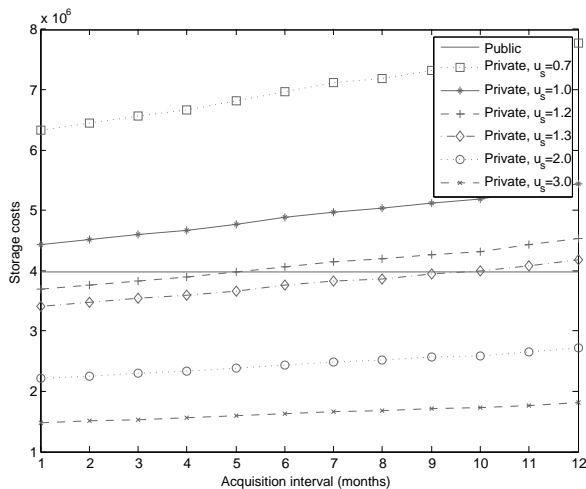


Figure 2. Storage costs vs. acquisition intervals for different values of utility premium

The figure also manifests that the length of the acquisition interval and the utility premium have a compound effect on the cost-efficiency of the private storage: the smaller the utility premium, the shorter should be the acquisition interval for the private storage to be less expensive than the public storage. For instance, while for $u_s = 1.3$, the acquisition

intervals shorter than ten months render the private storage solution cost-efficient, for $u_s = 1.2$, the acquisition interval should be reduced to less than five months for the private storage to be less expensive.

The cumulative effect of the acquisition interval and the level of redundancy on the storage cost is illustrated in Fig. 3. As could be seen, the increase in the required redundancy shortens the acquisition interval at which the private storage remains cost-efficient. Furthermore, for the redundancy above a certain threshold (1.55 in this example), the cost of the private storage always exceeds the public storage cost even for the shortest interval.

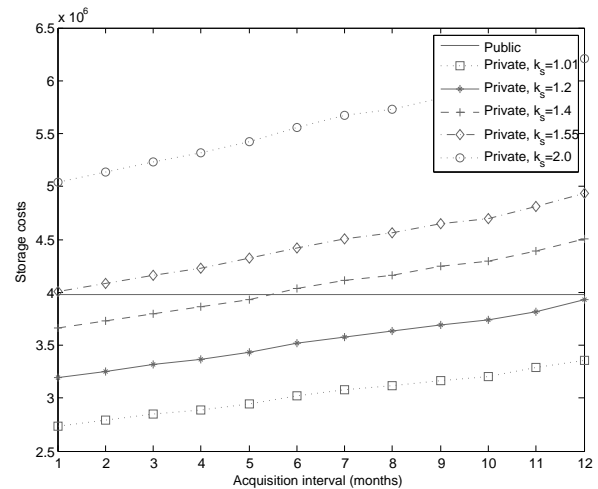


Figure 3. Storage costs vs. acquisition intervals for different levels of redundancy

The sum of the storage and the data communication costs are portrayed in Fig. 4 as a function of the acquisition interval. As could be seen from the figure, the intensity of data communications (manifested in the value of k_b) has an effect similar to the effect of the level of redundancy: namely, the greater the volume of data transfer incurred due to storing the data, the shorter the acquisition intervals that need to be maintained for the private storage to remain less expensive than the public storage. Thus, the figure supports the argumentation in Section III (cf. (7) and (11)) where it was analytically illustrated that the data communication costs strengthen the effect of the acquisition interval on the cost-efficiency of the private storage as compared with the public storage.

Finally, Fig. 5 illustrates, how the estimation errors can be compensated with shorter acquisition intervals. For example, given the estimation error $k_e = 1.3$, an acquisition interval shorter than eight months is needed to ensure that the private storage is cheaper than the public solution. However, if the estimation error is bigger, e.g. 1.4, the acquisition interval should be reduced to less than two months in order to keep the private solution cost-efficient enough. It can be seen

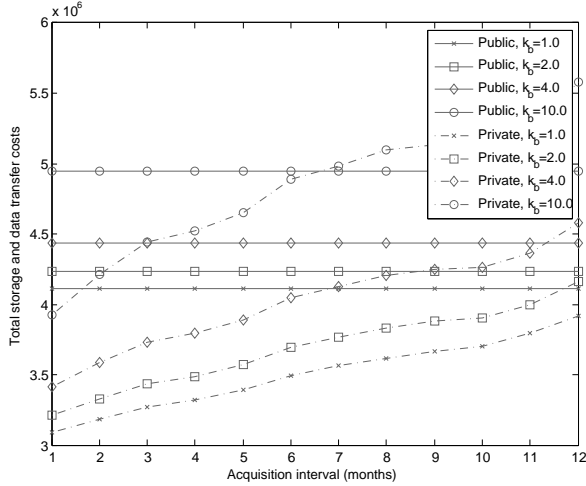


Figure 4. Storage and data transfer costs vs. acquisition intervals for different levels of data communication intensity

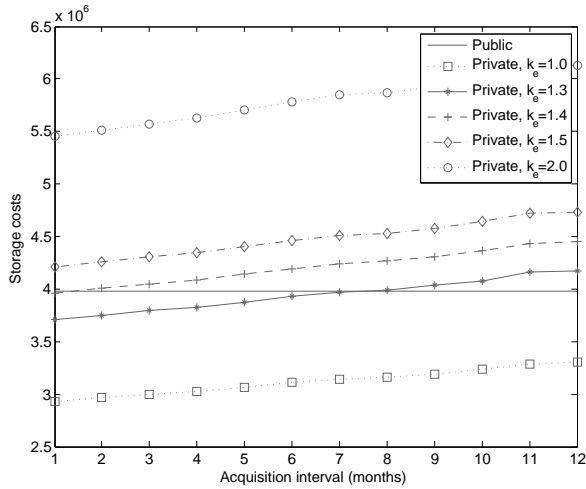


Figure 5. Storage and data transfer costs vs. acquisition intervals for different estimation error levels

also, that if the storage demand is well-known ($k_e \approx 1$), the private solution is cheaper than the public. On the other hand, if the needs are not easily estimable, the public solution is the cheaper alternative.

V. CONCLUSIONS

Contemporary organizations need to cope with rapidly growing demand for data storage. When deciding on the approach to meet the increasing storage needs, these organizations may decide to build and manage private data storage facilities or, alternatively, to utilize the on-demand storage services offered by the providers of public cloud infrastructure. The comparative cost-efficiency of these two alternatives depends on a number of factors, among which are e.g. the pricing difference between the public and private storage, the charging period (for the public storage) and

the storage acquisition interval (for the private storage), the predictability of the demand for storage, etc.

In this paper, an analytical tool to support the organizations in assessing whether the use of private or public storage solution will result in a cost-efficient solution has been introduced. It has been shown analytically that the use of a public storage is likely to be cost-efficient whenever the organization's acquisition cycles are relatively long, e.g. once per year. On the other hand, should the organization have a possibility to re-assess its storage needs and acquire additional storage often - say, every second month - the use of the private storage capacity is likely to prove less expensive. Since the acquisition interval is determined by the organization's ability to foresee the growth of storage demand, by the provisioning schedules of storage equipment providers, and by internal practices of the organization, among other factors, the organization owning a private storage solution may want to control some of these factors in order to attain a shorter acquisition interval and thus make the private storage (more) cost-efficient.

The paper has also illustrated that, along with the acquisition interval, the other factors such as the utility premium charged by the public storage provider, the level of needed storage redundancy, and the incurred data communications have a compound effect on the cost efficiency of the private vs. public storage and in particular on the threshold length of the acquisition interval below which (above which) the private (public) storage is more cost-efficient. More specifically, a decline in the utility premium, an increase in the storage redundancy, or an increase in incurred data transfer reduce the maximum length of the acquisition interval that needs to be maintained for the private storage to be less expensive as compared with the public storage.

It should be noted that the cloud providers were assumed in the paper to charge their customers based on the maximum storage usage within a charging period, in line with Amazon pricing. However, storage providers may apply also other pricing models: for example, the customers of Windows Azure Storage are charged by Microsoft based on the daily average usage during the month. It is therefore important to explore the effect of the alternative pricing models on the cost-efficiency of the private vs. public storage; this has been left for further study.

In further work, the proposed approach could be extended in several directions. First, the time dimension of the analytical tool shall be expanded to account for the declining pricing trends, and the pricing estimates themselves could be revised to include volume discounts visible e.g. in Amazon AWS offerings, as well as additional incurred costs such as the costs of input-output requests. The analytical tool shall be also elaborated to take into account the net present value of the assets. Besides deterministic storage growth profile, probabilistic profiles could be studied in future work. For a more holistic view, probabilistic communication

patterns should be taken into account as well. Finally, when estimating the data communication overheads, the specifics of possible organization's architectural solutions could be explored, since the architectural choices may have a significant influence on the data communication costs.

ACKNOWLEDGMENT

The research reported in this paper was carried out within the framework of the Cloud Software Program which is governed by TIVIT Oy nominated to organize and manage the programs of the Strategic Center for Science, Technology and Innovation in the field of ICT funded by the Finnish Funding Agency for Technology and Innovation (TEKES).

REFERENCES

- [1] J. Gantz and D. Reinsel, "Extracting value from chaos," 2011, IDC iView, available from <http://idcdocserv.com/1142>, last retrieved on 29.2.2012.
- [2] F. Gens, "IDC predictions 2012: Competing for 2020," IDC, Report, Dec 2011.
- [3] J. Barr, "The cloud scales: Amazon S3 growth," 2011, amazon Web Services Blog, available from <http://aws.typepad.com/aws/2011/10/amazon-s3-566-billion-objects-370000-requestssecond-and-hiring.html>, last retrieved on 28.2.2012.
- [4] J. Weinman, "Mathematical proof of the inevitability of cloud computing," January 8 2011, working paper, available from <http://www.joeweinman.com> (last retrieved on February 28, 2012).
- [5] A. Khajeh-Hosseini, D. Greenwood, J. W. Smith, and I. Sommerville, "The cloud adoption toolkit: supporting cloud adoption decisions in the enterprise," *Software: Practice and Experience - Special Issue on Software Architectures and Application Development Environments for Cloud Computing*, 2011. [Online]. Available: <http://dx.doi.org/10.1002/spe.1072>
- [6] O. Mazhelis and P. Tyrväinen, "Economic aspects of hybrid cloud infrastructure: User organization perspective," *Information Systems Frontiers*, pp. 1–25, 2011. [Online]. Available: <http://dx.doi.org/10.1007/s10796-011-9326-9>
- [7] E. Walker, W. Brisken, and J. Romney, "To lease or not to lease from storage clouds," *Computer*, vol. 43, pp. 44–50, April 2010. [Online]. Available: <http://dx.doi.org/10.1109/MC.2010.115>
- [8] J. Weinman, "Time is money: The value of "on-demand"," January 7 2011, working paper, available from <http://http://www.joeweinman.com> (last retrieved on February 28, 2012).
- [9] —, "Cloudonomics: A rigorous approach to cloud benefit quantification," *The Journal of Software Technology*, vol. 14, pp. 10–18, October 2011.
- [10] M. Hedwig, S. Malkowski, C. Bodenstein, and D. Neumann, "Datacenter investment support system (daisy)," in *Proceedings of the 2010 43rd Hawaii International Conference on System Sciences*, ser. HICSS '10. Washington, DC, USA: IEEE Computer Society, 2010, pp. 1–10. [Online]. Available: <http://dx.doi.org/10.1109/HICSS.2010.129>
- [11] C. Bodenstein, M. Hedwig, and D. Neumann, "Strategic decision support for smart-leasing infrastructure-as-a-service," in *ICIS*, D. F. Galletta and T.-P. Liang, Eds. Association for Information Systems, 2011. [Online]. Available: <http://dblp.uni-trier.de/db/conf/icis/icis2011.html#BodensteinHN11>
- [12] O. Mazhelis and P. Tyrväinen, "Role of data communications in hybrid cloud costs," in *Proceedings of the 37th EUROMICRO Conference on Software Engineering and Advanced Applications*, 2011.
- [13] G. Stigler, "The economies of scale," *Journal of Law and Economics*, vol. 1, pp. 54–71, 1958.
- [14] J. Strebel and A. Stage, "An economic decision model for business software application deployment on hybrid cloud environments," in *Multikonferenz Wirtschaftsinformatik 2010*, M. Schumann, L. M. Kolbe, M. H. Breitner, and A. Frerichs, Eds., Universitätsverlag Göttingen, 2010, p. 195206.
- [15] L. Mastroeni and M. Naldi, "Long-range evaluation of risk in the migration to cloud storage," in *Commerce and Enterprise Computing (CEC), 2011 IEEE 13th Conference on*, sept. 2011, pp. 260–266.
- [16] S. Uttamchandani, K. Voruganti, R. Routray, L. Yin, A. Singh, and B. Yolken, "Brahma: Planning tool for providing storage management as a service," in *IEEE International Conference on Services Computing (SCC 2007)*, july 2007, pp. 1–10.
- [17] I. Trummer, F. Leymann, R. Mietzner, and W. Binder, "Cost-optimal outsourcing of applications into the clouds," in *Cloud Computing Technology and Science (CloudCom), 2010 IEEE Second International Conference on*, 2010, pp. 135–142.
- [18] T. Nufire, "Petabytes on a budget v2.0: Revealing more secrets," July 2011, Backblaze blog, available from <http://blog.backblaze.com/2011/07/20/petabytes-on-a-budget-v2-0revealing-more-secrets/>, last retrieved on 15.2.2012.
- [19] X. Cheng, C. Dale, and J. Liu, "Statistics and social network of youtube videos," in *Quality of Service, 2008. IWQoS 2008. 16th International Workshop on*, June 2008, pp. 229–238.
- [20] B. Stiller, P. Reichl, and S. Leinen, "Pricing and cost recovery for internet services: Practical review, classification and application of relevant models," *NETNOMICS*, vol. 3, no. 1, March 2001.
- [21] C. Morris, "Data Backup and Archiving on the HFS," 2011, Oxford University Computing Services, available from <http://www.oucs.ox.ac.uk/hfs/>, last retrieved on 23.2.2012.
- [22] OUCS, "OUCS Annual Report 2004-2005," 2006, Oxford University Computing Services, available from <http://www.oucs.ox.ac.uk/internal/annrep/annrep0405/>, last retrieved on 23.2.2012.
- [23] I. Smith and M. Fraser, "Service level description: Backup and archiving service (HFS)," 2011, Oxford University Computing Services, available from <http://www.oucs.ox.ac.uk/internal/sld/hfs.xml>, last retrieved on 23.2.2012.