

Balancing Leasing and Insurance Costs to Achieve Total Risk Coverage in Cloud Storage Multi-Homing

Maurizio Naldi

Università di Roma Tor Vergata,
Dpt. of Computer Science and Civil Engineering
Via del Politecnico 1, 00133 Roma, Italy
naldi@disp.uniroma2.it
<http://www.maurizionaldi.it>

Abstract. Cloud storage multi-homing (a.k.a. as multicloud) represents a possible solution to achieve enhanced availability of one's own data and build an extremely reliable cloud environment out of relatively unreliable single cloud platforms. However, subscribing with multiple cloud providers comes at a cost. The cost of multiple subscriptions must be compared against the damage resulting from the possible temporary unavailability of the data stored on the cloud. A complementary approach is to protect against the economic damage resulting from data unavailability by subscribing an insurance policy. In this paper we investigate the complementary use of cloud multi-homing and insurance to obtain a total risk coverage against data unavailability. We provide a total cost function, which incorporates the cost of cloud subscriptions and the insurance premium, and investigate the possibility of trade-offs to achieve a total risk coverage at minimum cost.

Keywords: Cloud storage, Multicloud, Multi-homing, Insurance, Costs, Network Economics,

1 Introduction

Cloud storage is a fast growing service, whereby an individual or a company stores its data on a storage facility owned and managed by a third party (the cloud provider). Resorting to cloud storage allows users to eliminate their own storage infrastructure. Migrating from an owned infrastructure to a leased one has immediate benefits but also raises several concerns, so that several issues have to be considered in the migration decision as well as a comprehensive cost model [9] [4]. A major benefit in switching to the cloud consists in moving from a cost structure made of capital investments to a more flexible one based on operational expenses only [24] [12] [10]. However, capital investments are one-off in nature but they may lead to savings in the long run [14]. In addition, switching to the cloud may expose the cloud user to the lock-in phenomenon and price rises [13] [18].

Among the several issues, an outstanding one is that related to reliability, for which very few measurement data exist and concerns have been raised [17]. Storage on multiple clouds (a.k.a. as multicloud or cloud multi-homing) has been suggested to enhance reliability [20][15].

However large the improvement in reliability due to the use of multiple clouds, we must recognize that this approach cannot zero the probability of not being able to access its own data and comes anyway at the considerable cost of multiplying the storage costs. In [11], it has been shown that an alternative strategy to protect against the risk of an unavailable platform (in [11] it was a network, but the reasoning may be applied equally well to a cloud platform) can be subscribing an insurance policy. The two solutions (multicloud and insurance) can therefore be combined to achieve total risk coverage against cloud unavailability: using multiple clouds reduces the unavailability time and the insurance policy covers the residual risk.

In this paper, we investigate this combined approach. We consider simple models for the availability of a multi-cloud configuration and the insurance premium to be paid for a given availability level and arrive at an overall cost of the multi-cloud-plus-insurance solution. The total cost model allows us to formulate the conditions under which total risk coverage can be achieved by minimizing the total cost.

The paper is organized as follows. We review reliability analyses of cloud platforms and provide a simple model for the availability of a multicloud configuration in Section 2. The two cost components due to storage and insurance are described respectively in Section 3 and 4. The optimization of costs for a combined approach (multi-homing plus insurance) is investigated in Section 5.

2 Availability of Cloud Storage Service

The visibility a user has of its data on a cloud is very limited. Despite the cost savings that appear to be a major reason to move to the cloud, the user wishes the availability of its data to be at least equal to what an in-house solution would provide. This may not be the case. In this section, we provide a very basic analysis of the reliability of a cloud solution. After reviewing some definitions of availability, we first review the current (few) works on the actual availability of commercial clouds and then report a very simple model to predict the enhanced availability resulting from a multi-homing approach to clouds. In this paper, we deal neither with specific techniques to achieve a higher availability nor with standards concerning data management interfaces or engineering guidelines (see, e.g., the work carried out within the ENSI NFV Reliability and Availability WG). We adopt a user-centric view, where the cloud storage service is either available when the user wants to retrieve its data or it is not, denying the user the possibility to access its data.

Venkatesan and Iliadis have shown that the reliability of a data center depends very little on the assumptions adopted for the time-to-failure of individual storage nodes (i.e., the assumption of independent and exponentially-distributed

failure times) as long as the individual nodes are generally reliable (i.e., their mean time to failure is much longer than their mean time to repair) [23].

Baker et alii have investigated the behaviour of large-scale storage systems in the long run, assessing a wider set of failure root causes than that traditionally examined in reliability analyses [1]. They formulated strategies to improve long-term reliability, by acting primarily on the reduction of the mean time to repair.

Though addressing the related case of a data center network, Gill et alii have performed a thorough reliability analysis. They found that the data center network exhibits high reliability with more than four 9's of availability for about 80% of the links and for about 60% of the devices in the network [6].

Other data concerning the reliability of storage platforms tell us that the MTTR (Mean-Time-To-Repair) of individual storage nodes is in the order of tens of hours [23], while the MTTR of visible faults (those that are detected shortly after they occur) is 1.4 hrs for a mirrored disk and 4.4 hours for a mirrored archive [1].

One of the few reliability studies concerning instead the specific case of a cloud storage platform is that performed on Google's main storage infrastructure [5]. In that study, Ford et alii investigated the availability of single components of the complex cloud environment and the impact of correlated failures, arriving at an estimate of the unavailability of data stripes (data is divided into a set of stripes, each of which comprises a set of fixed size data and code blocks called chunks), with figures in the range $10^{-7} \div 10^{-3}$ depending on the time of the day.

Since quantitative assessments of the availability of clouds based on extensive measurement campaigns are typically not reported by cloud providers, an alternative approach has been taken in [17] to try a third-party assessment based on outage as reported by customers themselves. After collecting data from a number of sources about 5 major cloud providers, it has been shown that the average duration of outages was over 8 hours for 3 out of 5 providers (the actual figures ranged from 95 to 607 minutes).

A third party assessment of cloud storage availability has been provided in [7] for Amazon and Google, whose availability has resulted respectively equal to 99.565% and 99.783% (higher figures result if successive retries after the first failures are included). Some limits of that analysis have been reported in [16].

Most cloud storage providers specify the availability of their service, though they typically fail to provide guarantees or include tight performance levels in their SLA. In the survey reported in [3], 15 providers out of 17 declared at least 99.9% availability, with 12 providers declared 100% availability.

Whatever the availability expected from a single cloud, it has been suggested that we can achieve a higher reliability by employing multiple clouds at the same time. For example, in [15], erasure coding has been employed to create redundant data shares and spread the original data to different cloud providers. In that case, a combination of up to 7 storage providers was considered and shown to provide a lower read latency. Zhai et alii have also highlighted that the availability resulting from a multi-cloud strategy may not be as high as expected according to the independence hypothesis, since hidden interdependencies may appear due to

third-party infrastructure components shared by redundant deployments [26], and have suggested the use of a recommender system to take into account those hidden interdependencies.

In this paper, we do not consider a specific strategy to disseminate data among several clouds to optimize the overall reliability. Instead we consider, for sake of simplicity, a general strategy relying on the simple replication of the same volume of data among several totally independent cloud providers. For each cloud, we assume that the service follows a simple ON-OFF process, where the data are either available for the user to retrieve or they are (temporarily) not. Unavailability of data does not mean data loss: the user will be able to retrieve again its data after the outage has ended. Under such basic assumptions, we can compute the availability in a multi-homing approach with a very simple model. For our purposes, we consider the overall availability of data over a cloud. For a single cloud, we define A_1 as the probability that a data retrieval operation is successful. Let's consider a customer subscribing a contract with n cloud storage providers. For simplicity we assume that all cloud providers guarantee the same availability A_1 and that failures take place at providers independently of one another. The customer can access its data if even a single provider is operating. The probability $1 - A_n$ of not being able to retrieve the data with any of the n clouds (a classical parallel configuration), where A_n is the resulting availability of the data over n clouds, is therefore

$$1 - A_n = (1 - A_1)^n \implies A_n = 1 - (1 - A_1)^n. \quad (1)$$

If the availability of a single cloud is expressed as the number w_1 of nines (e.g., if $A_1 = 0.999$ then $w_1 = 3$), the availability of a group of n clouds exhibits the number w_n of nines. In fact, $A_1 = 1 - 10^{-w_1}$ and

$$A_n = 1 - 10^{-nw_1} = 1 - 10^{-w_n}, \quad (2)$$

with $w_n = nw_1$. Therefore the addition of a cloud increases the number of nines by the number of nines of a single cloud. If we start with a cloud with two nines, the solution with 2 clouds will exhibit four nines, and so on.

On the other hand, Equation (1) can be inverted to compute the number of clouds needed to achieve a desired value of availability

$$n = \frac{\ln(1 - A_n)}{\ln(1 - A_1)}, \quad (3)$$

though Equation (3) should be modified by employing a ceiling function so as to get integer solutions.

3 Storage Leasing Costs

Enhancing availability through multi-homing has the obvious disadvantage of subscribing to multiple clouds and paying the pertaining costs. In this section,

we review the current pricing policies of commercial clouds and provide a model to arrive at the overall cost of a multi-homing approach. We rely on the current market situation, though more advanced pricing models have been considered, especially for cloud computing [2][22]

In [19], a survey of current cloud storage pricing practices has been presented. The survey considered 8 commercial cloud storage providers, which offer pricing plans both for consumers and business customers. It has been shown that the pricing models adopted by them can be classified into just two pricing scheme, namely Block rate pricing and Bundling pricing, with just Amazon following the former and the remaining 7 adopting the latter.

In Block rate pricing, the range of consumption is subdivided into subranges, and the unit price is held constant over each subrange. More formally, in a block rate tariff the overall price p charged to the customer for a storage volume d is

$$p = \begin{cases} v_1 d & \text{if } 0 < d \leq q_1 \\ v_1 q_1 + v_2 (d - q_1) & \text{if } q_1 < d \leq q_2 \\ \dots & \\ \sum_{i=1}^{m-1} v_i q_i + v_m (d - q_{m-1}) & \text{if } q_{m-1} < d \leq q_m \end{cases} \quad (4)$$

where the v_i 's are the sequence of marginal prices, and the q_i 's bracket the subranges over which the marginal price is held constant. In Equation (4), we assume that the cloud provider does not provide more than q_m units of storage ($m \geq 2$). In turn, block rate pricing can be seen as a special form of multi-part tariff, where the fixed fee has been set equal to zero.

The overall charge is then a piecewise linear function of the amount of storage capacity (see Figure 1). Diminishing prices at the margin stimulate consumption, which in turn permits the construction of large scale capacity.

The pricing model adopted by all cloud providers but one is instead bundling pricing, where several pricing schemes are proposed for a fixed fee. Each scheme allows the customer to store data up to a maximum volume. Since the pricing schemes can be sorted by the allowed maximum volume, we can derive a price-volume relationship, which is a step function. A sample price function is shown in Fig. 2.

In [19], we have also shown that a suitable approximation of such pricing models, convenient for a mathematical treatment, is the two-part tariff. In the two-part tariff scheme, where the customer pays an initial fixed fee f for the first block of data (often justified as a subscription, access, or installation charge), plus a smaller constant price for each unit [25]. This model is equivalent to the block-rate scheme is just a single range is considered in Equation (4). It is as well a linear approximation of the staircase pricing relationship applying in bundling pricing. The two-part tariff is therefore a suitable approximation of both dominant pricing plans. The overall price charged to the customer is then

$$p_{sto} = f + v \cdot d, \quad (5)$$

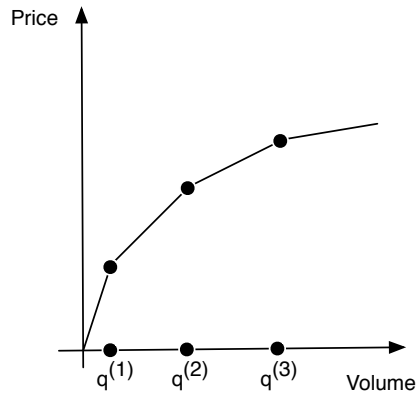


Fig. 1. Price-volume relationship in block rate pricing




Fig. 2. Sample price-volume relationship in bundling pricing

where v is the marginal price, and d is the volume of consumption, i.e., the amount of storage volume.

If we consider a multi-homing approach, the customer is actually subscribing n cloud storage contracts. If we assume, for simplicity, that all providers adopt the same two-part tariff pricing scheme, the overall expense for the customer is

$$p_{sto} = n(f + v \cdot d). \quad (6)$$

4 Insurance against Unavailability

In Section 2, we have seen that the availability of the cloud can be improved by resorting to multiple clouds, up to the point of creating a reliable storage environment out of many relatively unreliable clouds. However, in Section 3 we have also seen that this involves a multiplication of storage costs. Since the ultimate negative consequence of unreliable storage is the money loss deriving from the inability to access its own data for the duration of the storage blackout, a possible remedy is the subscription of an insurance policy. In this section, we describe such a possibility and provide the formula which relates the insurance premium to the expected unavailability.

Let's assume that the customer has subscribed to n different cloud storage providers, where n can even be just 1. This allows the customer to expect an overall availability A_n provided by Equation (1). Though the customer can make this value arbitrarily high by increasing the number of cloud providers, this approach cannot be pushed too far. In fact, the number of providers is anyway limited and the addition of a further provider increases the cost. Though adding clouds increases quite fast the overall availability of the cloud as embodied by Equation (2), full availability can never be reached.

A complementary approach is to accept that the availability of the cloud can never be 1, though increasingly close, and subscribe an insurance policy to cover the damage resulting from the remaining unavailability periods. The insurance solution comes of course at a cost, represented by the insurance premium, i.e., the price to be paid to subscribe the insurance policy.

In [11], the premium expression has been derived for the case where the failing system is a network and a simple ON-OFF continuous-time Markov process is assumed to model the occurrence of failures and the restoring of normal operations. The expected duration of the ON period is the MTBF (Mean Time Between Failures), while the expected duration of the OFF period is the MTTR (Mean Time To Restore). The pricing principle adopted follows the expected utility model, with the utility function exhibiting Constant Absolute Risk Aversion (CARA) property [8] [21]. Since no specific features of a network are exploited, we can adopt the same model here. If we consider that the insurance policy is subscribed over the same period of time for which the cloud subscription fee applies, and we normalize all time measures to that time scale, the premium p_{ins} for a volume of data d is

$$p_{ins} = kd \frac{1 - A_n}{A_n} \left(1 + \alpha k \frac{\text{MTTR}}{n} \right), \quad (7)$$

where α is the risk aversion coefficient, k is the money loss per data unit (e.g., per each GB) over the whole policy validity period, and the MTTR refers to a single cloud (expressed as a fraction of the overall policy validity period). The inclusion of the factor k allows to account for the different value we attach to the unavailability of data, which may depend on their use: delaying the listening of an MP3 song bears probably a much lower damage than delaying the access to an X-ray image needed for a remote diagnosis.

5 Trade-offs in Multi-homing

In Sections 3 and 4 we have described the two components of costs, respectively the cost of leasing storage space on several clouds and that of subscribing an insurance policy to cover against the risk of data unavailability. By combining the two operations, the user can achieve a total risk coverage strategy, where the insurance policy covers the money loss resulting from the less-than-100% availability. In fact, the user may decide to subscribe to more clouds achieving a higher availability: the cost of storage leasing will grow, but at the same time the insurance premium will go down because of the increased availability. The overall cost of the risk coverage strategy is the sum of the cost of the two components. A trade-off can be pursued where to achieve a lower overall cost by a careful choice of the number of clouds. In this section, we examine the overall multi-homing cost and investigate the possible trade-offs.

If we recall the cost of leasing storage from n cloud providers, given by Equation (6), and the cost of subscribing an insurance policy, given by Equation (7), we obtain the following overall cost in a multi-homing strategy

$$\begin{aligned} p_n &= p_{\text{sto}} + p_{\text{ins}} = n(f + v \cdot d) + kd \frac{1 - A_n}{A_n} \left(1 + \alpha k \frac{\text{MTTR}}{n}\right) \\ &= n(f + v \cdot d) + kd \frac{(1 - A_1)^n}{1 - (1 - A_1)^n} \left(1 + \alpha k \frac{\text{MTTR}}{n}\right). \end{aligned} \quad (8)$$

If we consider that the availability of even a single cloud is typically very close to 1 (we can assume at least a two nine availability), we have

$$p_n \simeq n(f + v \cdot d) + kd(1 - A_1)^n \left(1 + \alpha k \frac{\text{MTTR}}{n}\right). \quad (9)$$

We see that the two components have a different behaviour when the number of clouds grows: the cost of storage grows linearly, while the cost of insurance decreases (roughly exponentially fast). Though $\lim_{N \rightarrow \infty} p_{\text{ins}} = 0$, the multi-homing approach cannot be pushed too far, because of the linear growth of storage costs which makes the overall cost grow linearly in the end. It is to be seen if we can lower the overall cost for a moderate use of multi-homing with respect to a single cloud.

Let's consider first the double-homing case, so that we have to compare p_1 and p_2 . For this case, we can formulate the following theorem

Theorem 1. *The double homing approach is the cheapest if the unit price of storage satisfies the following inequality*

$$\frac{f}{d} + v < kd \frac{1-A_1}{2A_1} \frac{2+(3-A_1)\alpha \cdot k \cdot \text{MTTR}}{2-A_1}$$

Proof. Let's define the gain achieved by switching to double-homing as $\Delta_1 = p_1 - p_2$. We recall the general expression of the total cost from Equation (8) for the cases $n = 1, 2$ we have

$$\Delta_1 = kd \left[\frac{1-A_1}{A_1} (1 + \alpha k \text{MTTR}) - \frac{(1-A_1)^2}{1-(1-A_1)^2} \left(1 + \alpha k \frac{\text{MTTR}}{2} \right) \right] - f - vd, \quad (10)$$

which after some straightforward manipulation becomes

$$\Delta_1 = kd \frac{1-A_1}{2A_1} \frac{2+(3-A_1)\alpha k \cdot \text{MTTR}}{2-A_1} - f - vd. \quad (11)$$

The double-homing strategy is economically convenient if $\Delta_1 > 0$, i.e., if

$$kd \frac{1-A_1}{2A_1} \frac{2+(3-A_1)\alpha k \cdot \text{MTTR}}{2-A_1} > f + vd, \quad (12)$$

from which follows

$$\frac{f}{d} + v < kd \frac{1-A_1}{2A_1} \frac{2+(3-A_1)\alpha k \cdot \text{MTTR}}{2-A_1} \square \quad (13)$$

Some simple corollaries follow from Theorem 1, concerning the possibility of the double homing strategy being the cheapest one. We formulate the first one concerning the value of the data stored on the cloud

Corollary 1. *For any amount d of data stored and any couple of prices (f, v) , there exists a value of the data stored on the cloud for which the double homing strategy is the cheapest*

Proof. Since $A_1 < 1 < 2 < 3$, the right hand side of the inequality (13) has the form $ak + bk^2$, where $a, b > 0$, and is a growing function of the unit value k of the data stored, . Whatever the value of d, f, v appearing in the left hand side of that inequality, the right hand side can become so large as to satisfy the inequality. In fact, by solving the associated quadratic form, it can be seen that the inequality (13) is satisfied for any k such that

$$k > \frac{-a + \sqrt{a^2 + 4b(f/d + v)}}{2b}. \square \quad (14)$$

The second corollary concerns the amount of data stored.

Corollary 2. *If the marginal price v satisfies the inequality $v < kd \frac{1-A_1}{2A_1} \frac{2+(3-A_1)\alpha k \cdot \text{MTTR}}{2-A_1}$, there exists a number d_1 , such that the double homing strategy is the cheapest for any amount of data stored $d > d_1$*

Proof. Let ϵ be the positive quantity defined as $\epsilon = kd \frac{1-A_1}{A_1} \frac{2+(3-A_1)\alpha k \cdot \text{MTTR}}{2-A_1} - v$. Then, for $d_1 = f/\epsilon$, we have

$$\frac{f}{d_1} + v = \epsilon + v = kd \frac{1-A_1}{A_1} \frac{2+(3-A_1)\alpha k \cdot \text{MTTR}}{2-A_1}. \quad (15)$$

Hence, for any $d > d_1$ we have

$$\frac{f}{d} + v < \frac{f}{d_1} + v < kd \frac{1-A_1}{A_1} \frac{2+(3-A_1)\alpha k \cdot \text{MTTR}}{2-A_1}, \quad (16)$$

and the condition of Theorem 1 is met. \square

If we turn to multi-homing, we can define the quantity

$$\Delta_i = p_i - p_{i+1} \quad i = 1, 2, \dots, \quad (17)$$

which is the gain achieved when we add a cloud to a storage platform already consisting of i clouds. For $i = 1$, we find the quantity already defined in Theorem (1). Adding a cloud leads to an overall lower cost as long as that gain is positive. We can therefore identify the optimal number n^* of clouds as

$$n^* = \max i \quad : \quad \Delta_i > 0. \quad (18)$$

If we recall Equation (8), we can write the general expression of the gain

$$\Delta_i = -(f+vd) + kd \left[\frac{(1-A_1)^i \left(1 + \alpha k \frac{\text{MTTR}}{i}\right)}{1 - (1-A_1)^i} - \frac{(1-A_1)^{i+1} \left(1 + \alpha k \frac{\text{MTTR}}{i+1}\right)}{1 - (1-A_1)^{i+1}} \right], \quad (19)$$

which after some algebraic manipulation becomes

$$\Delta_i = -(f+vd) + kd(1-A_1)^i \frac{A_1 + \frac{\alpha k \cdot \text{MTTR}}{i(i+1)} [1 + iA_1 - (1-A_1)^{i+1}]}{[1 - (1-A_1)^i][1 - (1-A_1)^{i+1}]}. \quad (20)$$

Since Equation (20) appears a bit cumbersome to obtain the optimal number of clouds, we can obtain a useful approximation by considering that $A_1 \simeq 1$, so that $iA_1 \simeq i$ and


$$\Delta_i \simeq -(f+vd) + kd(1-A_1)^i \left[A_1 + \alpha k \frac{\text{MTTR}}{i} \right], \quad (21)$$

so that solving the inequality $\Delta_i > 0$ is roughly equivalent to solving the simpler equation

$$n^* = [i] \quad : \quad kd(1-A_1)^i \left[A_1 + \alpha k \frac{\text{MTTR}}{i} \right] = f + vd. \quad (22)$$

We can get a feeling of how adding clouds may actually help achieve a total risk coverage at a reduced cost by computing the total cost through Equation

Parameter	Value
Availability	0.999
MTTR [hrs]	5
Data unav. loss [€/GB]	35
Storage fixed fee [€]	4
Marginal price [€/GB]	0.05
Data volume [TB]	10
α	5

Table 1. Value of parameters for the sample case**Fig. 3.** Total cost

(8) for a sample configuration. In Table 1, we report a typical set of parameter values for a month's time period. In this configuration we have considered that a single cloud delivers a three 9's availability and that not being able to access data provokes a loss of 35 € for each GB. The storage prices are those reported in [19] for the two-part tariff fitting of Amazon's prices.

In Fig. 3, we report the total cost as we increase the number of clouds. We see that: 1) the minimum cost is achieved for a double homing strategy; 2) adding a further cloud makes the insurance premium negligible with respect to the storage cost, so that for $n \geq 3$ the total cost grows linearly. The cost saving in switching from a single cloud to 2 clouds is a remarkable 21%.

6 Conclusions

Storing data on multiple clouds (multicloud or cloud multi-homing) provides a higher reliability than a single cloud, but does not guarantee that data are always accessible. Complementing the use of multiple clouds with the subscription of an insurance policy against the money loss deriving from being unable to access data allows to achieve total risk coverage against cloud unavailability.

In this paper we have derived a basic cost model for this combined approach to identify the optimal configuration leading to the minimum overall cost. We have provided the conditions for double homing to be economically convenient with respect to a single cloud solution and derived the equation that leads to the optimal number of clouds. We have also shown the behaviour of the cost curve for a typical configuration, where double homing plus insurance provides total risk coverage at a cost that is 21% less than that of a single cloud.

This study can be considered as a first step that proves the economical viability of a combined approach (multicloud + insurance) to achieve total risk coverage. Some simplistic assumptions should be removed in future works, namely assuming that data replication is complete (and therefore storage costs grow linearly), neglecting the occurrence of correlated failures (which leads to a faster growth of availability), and assuming a simple insurance pricing model (which leads to a vanishing insurance premium when the availability grows).

References

1. Baker, M., Shah, M., Rosenthal, D.S.H., Roussopoulos, M., Maniatis, P., Giuli, T., Bungale, P.: A fresh look at the reliability of long-term digital storage. *SIGOPS Oper. Syst. Rev.* 40(4), 221–234 (Apr 2006)
2. Berndt, P., Maier, A.: Towards sustainable IAAS pricing. In: Vanmechelen, K., Altmann, J., Rana, O.F. (eds.) *GECON 2011, LNCS*, vol. 7150, pp. 173–184. Springer, Heidelberg (2011)
3. Casalicchio, E., Silvestri, L.: Mechanisms for SLA provisioning in cloud-based service providers. *Comput. Netw.* 57(3), 795–810 (Feb 2013)
4. Dutta, A., Hasan, R.: How much does storage really cost? Towards a full cost accounting model for data storage. In: Vanmechelen, K., Altmann, J., Rana, O.F. (eds.) *GECON 2013, LNCS*, vol. 8193, pp. 29–43. Springer, Heidelberg (2013)
5. Ford, D., Labelle, F., Popovici, F.I., Stokely, M., Truong, V.A., Barroso, L., Grimes, C., Quinlan, S.: Availability in globally distributed storage systems. In: 9th *USENIX Symposium on Operating Systems Design and Implementation, OSDI 2010*, pp. 61–74. Vancouver, BC, Canada (2010)
6. Gill, P., Jain, N., Nagappan, N.: Understanding network failures in data centers: Measurement, analysis, and implications. *SIGCOMM Comput. Commun. Rev.* 41(4), 350–361 (Aug 2011)
7. Hu, Z., Zhu, L., Ardi, C., Katz-Bassett, E., Madhyastha, H., Heidemann, J., Yu, M.: The need for end-to-end evaluation of cloud availability. In: Faloutsos, M., Kuzmanovic, A. (eds.) *15th International Conference on Passive and Active Measurement PAM 2014, LNCS*, vol. 8362, pp. 119–130. Springer, Heidelberg (2014)
8. Kaas, R., Goovaerts, M., Dhaene, J., Denuit, M.: *Modern Actuarial Risk Theory*. Springer (2004)
9. Kashef, M.M., Altmann, J.: A cost model for hybrid clouds. In: Vanmechelen, K., Altmann, J., Rana, O.F. (eds.) *GECON 2011, LNCS*, vol. 7150, pp. 46–60. Springer, Heidelberg (2011)
10. Laatikainen, G., Mazhelis, O., Tyrväinen, P.: Role of acquisition intervals in private and public cloud storage costs. *Decision Support Systems* 57, 320–330 (2014)
11. Mastroeni, L., Naldi, M.: Network protection through insurance: Premium computation for the on-off service model. In: 8th *International Workshop on the Design of Reliable Communication Networks, DRCN 2011, Krakow*, pp. 46–53. IEEE (2011)

12. Mastroeni, L., Naldi, M.: Storage Buy-or-Lease Decisions in Cloud Computing under Price Uncertainty. In: 7th EuroNF Conference on Next Generation Internet NGI 2011, Kaiserslautern. IEEE (2011)
13. Mastroeni, L., Naldi, M.: Pricing of insurance policies against cloud storage price rises. *SIGMETRICS Performance Evaluation Review* 40(2), 42–45 (2012)
14. Mastroeni, L., Naldi, M.: Long-range evaluation of risk in the migration to cloud storage. In: 13th IEEE Conference on Commerce and Enterprise Computing, CEC 2011, Luxembourg, pp. 260–266. IEEE (2011)
15. Mu, S., Chen, K., Gao, P., Ye, F., Wu, Y., Zheng, W.: μ LibCloud: Providing High Available and Uniform Accessing to Multiple Cloud Storages. In: 2012 ACM/IEEE 13th International Conference on Grid Computing. pp. 201–208. Beijing, China. ACM (2012)
16. Naldi, M.: A note on "The Need for End-to-End Evaluation of Cloud Availability". ArXiv e-print 1408.0510, available at <http://arxiv.org/abs/1408.0510> (Aug 2014)
17. Naldi, M.: The availability of cloud-based services: Is it living up to its promise? In: 9th International Conference on the Design of Reliable Communication Networks DRCN 2013, pp. 282–289. IEEE (2013)
18. Naldi, M.: Forecast Uncertainty in Procurement Decisions for Cloud Storage. In: 2014 UKSim-AMSS 16th International Conference on Computer Modelling and Simulation, Cambridge, pp. 237–242. IEEE (2014)
19. Naldi, M., Mastroeni, L.: Cloud storage pricing: A comparison of current practices. In: Proceedings of the 2013 International Workshop on Hot Topics in Cloud Services, HotTopiCS '13, Prague, pp. 27–34. ACM, New York, NY, USA (2013)
20. Papaioannou, T.G., Bonvin, N., Aberer, K.: Scalia: An adaptive scheme for efficient multi-cloud storage. In: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, SC '12. pp. 20:1–20:10. IEEE/ACM (2012)
21. Pratt, J.W.: Risk Aversion in the Small and in the Large. *Econometrica* 32(1-2), 122–136 (Jan-Apr 1964)
22. Roovers, J., Vanmechelen, K., Broeckhove, J.: A reverse auction market for cloud resources. In: Vanmechelen, K., Altmann, J., Rana, O.F. (eds.) GECON 2011, LNCS, vol. 7150, pp. 32–45. Springer (2011)
23. Venkatesan, V., Iliadis, I.: A General Reliability Model for Data Storage Systems. Research Report RZ3817, IBM Research, Zurich (March 2012)
24. Walker, E., Brisken, W., Romney, J.: To lease or not to lease from storage clouds. *IEEE Computer* 43(4), 44–50 (2010)
25. Wilson, R.: *Nonlinear Pricing*. Oxford University Press (1997)
26. Zhai, E., Chen, R., Wolinsky, D.I., Ford, B.: An untold story of redundant clouds: Making your service deployment truly reliable. In: Proceedings of the 9th Workshop on Hot Topics in Dependable Systems HotDep '13. pp. 3:1–3:6. ACM, New York, NY, USA (2013)