On the Network Performance of Amazon S3 Cloud-storage Service

Valerio Persico[1], Antonio Montieri[2], Antonio Pescapè[1,2]
[1]University of Napoli "Federico II" (Italy) and [2]NM2 s.r.l. (Italy)
{valerio.persico, pescape}@unina.it, montieri@nm-2.com

***Abstract*—The advances in networking technologies and the increase in the need for storage resources have prompted many companies to outsource their storage needs. Cloud-storage providers offer clean and simple file-system interfaces, abstracting away the complexities of direct hardware management. At the same time, however, such services eliminate the direct oversight of performance that final users with high service-level requirements traditionally expect. While several works in literature have addressed security-related issues (such as privacy, integrity, availability, etc.) few of them have targeted the network performance of this kind of services.**

In this work we propose the analysis of the performance of the network associated to the storage service offered by Amazon: S3. Thanks to a large-scale distributed campaign performed by leveraging the Bismark measurement platform, we have characterized how the performance of the network may impact the quality of service experienced by final users on the basis of their location and the configuration of services. We found how performance heavily changes (up to 1553 KiB/s) according to the location of the customers and the cloud region they rely on (up to 2117 KiB/s), also deriving a number of usage guidelines for the customers. In addition we characterize the impact of leveraging the Amazon CDN service to distribute contents, finding that while it guarantees up to a 275-percent performance improvement, cases exist for which additional costs may lead to worse performance.

I. INTRODUCTION

An increasing number of users and organizations more and more depends on the cloud as cloud-based infrastructures are today leveraged to deliver a growing set of services and applications. This conspicuous demand drives to huge investments by providers, to suitably implement the public cloud paradigm and make both remarkable technical and economical benefits available to customers. Cloud storage (also known as *storage as a service*) denotes a family of increasingly popular on-line services for archiving and backup data. Thanks to cloud storage services customers are able to deposit, access, and distribute different kinds of data, with no need of any upfront plan, provision, or investment and being subjected to a pay-as-you go payment model. In accordance to customers' reluctance in outsourcing critical information and resources, cloud providers usually give high guarantees in terms of service availability and reliability that are often well regulated by Service Level Agreements (SLAs). On the other hand, they rarely make promises about the performance expected from this kind of services [5].

From this perspective, the network is a key component for cloud services indeed, as also remarked by the recent literature [13]. As the cloud is accessible only from remote, the characteristics and the performance figures of the collection of paths between cloud datacenters and hosts on the Internet potentially impact the performance of the overall system and the quality of the service (QoS) experienced by customers. However, while all providers grant high-performance network connectivity to their customers, they provide no more than qualitative information about either its design or expected performance, due to security and commercial reasons [19]. The actual performance of the service experienced by the customers clearly depends upon their specific location with respect to the cloud resources. Indeed, the performance of the network interconnecting the customer to the set of pooled resources composing the cloud (i.e., the *cloud-to-user network*) is hard to accurately predict also for the cloud provider, because it is made up of different parts some of which are not under its direct control. This creates a strong motivation for providers for geographically distributing datacenters around the world to reduce speed-of-light delays to reach customers and improve the QoS accordingly. Thus, all the major cloud providers host the offered services at different locations all around the globe. This notwithstanding, customers have no information about the performance associated to the storage service. In addition, no guideline is provided about how performance varies with service configuration, location, and related costs.

In this paper, we aim at improving the knowledge on the performance of cloud storage services, by focusing on the performance of the cloud-to-user network experienced when relying on *Amazon Simple Storage Service (S3)*. S3 is the general purpose storage as a service provided by Amazon [2], where customer data is organized by means of *objects* stored in *buckets*. A bucket is a logical unit of storage uniquely identified and belonging to one of the locations in which the provider has deployed its storage infrastructures (hereafter *cloud regions*). Costs for the customer depend upon the storage class (standard, infrequent access, or long-term archive) and the cloud region in which the bucket is placed, according to a pay-as-you-go model. In more details, cost is calculated as the sum of three quotas depending on the size of the stored object, the number of download requests, and the volume of the traffic transferred [2]. *CloudFront (CF)* is the global Content Delivery Network (CDN) service offered by Amazon [1] and integrates with S3 in order to distribute contents to the end users with low latency and high data transfer speeds. Data is distributed to the users through the global network composed of the Amazon *edge locations* spread all over the world [3]. CF can be leveraged in combination with S3, by simply activating it with no need for further configuration. Notably, CF is associated to roughly the same storage and data transfer costs as S3, but to markedly higher cost for download requests (around 2×).

Previous works only marginally analyzed the network performance of Amazon storage service, considering few limited scenarios and without clarifying the impact of a number of choices available to the customers. To fill this gap, we provide a detailed characterization of the QoS final users deal with when leveraging Amazon S3 in different configurations and from a number of distinct locations spread worldwide. To this aim, we have performed an experimental campaign by leveraging the Bismark distributed measurement platform [22]. It is worth noting that differently from some previous works, our study is based on *non-cooperative* approaches [18] and only leverages active measurements that inject measurement traffic into the network to evaluate its characteristics and performance. Therefore, the enforced methodology does not need any privileged point of view beyond the one available to the general user, such as those of either the network operator or the cloud provider itself.

Our study drives to the following main contributions: (i) a general assessment of the performance of Amazon S3 is provided; (ii) the impact of a set of factors of interest to the general user is investigated, comprising the cloud region, the location of the user, and the size of the stored file; (iii) the benefits of leveraging the CF distribution is evaluated, also considering the enhancement with respect to the performance of the standard service and the shortcomings of the architecture implemented by the provider. The paper is organized as follows: Sec. II positions the paper against the related work; Sec. III describes the methodology we propose and adopt for this type of analysis; Sec. IV discusses the main results; finally, Sec. V ends the paper with concluding remarks.

II. RELATED WORK

Recently, cloud performance has attracted more and more interest. While a number of works has investigated cloud performance with respect to specific classes of applications [24], [10] or considering a broad set of performance indexes [11], [21], only few works have explicitly focused on the performance of the network of these complex systems. Most of these works took into consideration the characteristics and the performance of either the intra-datacenter or the inter-datacenter networks [25], [17], [16], [15]. Only few of them have focused on the network interconnecting the cloud to the user, mainly considering the latency as the only parameter of interest [20], [23]. In this work, we go beyond the state of the art by investigating the performance of the cloud-to-user network and considering Amazon S3 as an interesting use case.

As S3 has been the first Amazon web service publicly available [4], a number of works has tried to shed light on its performance [8], [14], [12], [6], [7]. Differently than most of the studies above [8], [14], [12], [11], our work aims at focusing on the performance of remote data delivery, i.e. it aims at investigating the quality of service perceived by users that perform content retrieval from vantage points (VPs) not placed into the cloud. With this goal in mind our approach only leverages active measurements. On the one hand, this characteristic frees our study from the need of any privileged point of view [6], [7], thus also guaranteeing easier repeatability. On the other side, the methodology we propose allows to better evaluate the impact of factors under the control of the customers, not limiting the validity of the study neither to the service usage patterns observed in traffic captures [6], [7] nor to a specific geographic zone [8], [14]. Indeed, in order to obtain a significant characterization of the performance of the cloud-to-user network—and of the service leveraging it—with respect to geographically distributed users, differently than previous works [12], [11] we leverage a set of geographically distributed VPs. In addition, our study provides an up-to-date view of the performance of the service under investigation, in front of both the evolution of the cloud network infrastructure over the time and the new services not considered in most of the previous works as not available at that time (e.g., a larger number of regions in which buckets can be placed and the integration of the cloud storage with CDN services).

TABLE I: Summary of factors and considered values.

Factor	Values
Cloud Region	North Virginia (US), Ireland (EU), Singapore (AP), Sao Paulo (SA)
File Size	1 B, 1 KiB, 1 MiB, 16 MiB, 100 MiB
Storage Class	Standard (S3), CloudFront (CF)
Source Region (VPs)	United States (US), Europe (EU), Asia-Pacific (AP), South Africa (ZA), Central-South America (CSA)

III. METHODOLOGY

In this section we describe the setup and the choices implemented in our experimental study. We have taken into account a number of factors that may impact the service performance experienced by users. These factors are summarized in Tab. I and will be briefly discussed in the following.

A. Factors of interest

As of today, Amazon has datacenters in 12 regions around the world. Due to experimental cost constraints, for our experimental campaigns we have identified a subset of 4 **cloud regions** among all possible ones: North Virginia (hereafter US), Ireland (EU), Singapore (AP), and Sao Paulo (SA). We have picked a region per continent, in order to ensure geographical diversity to our dataset. In each of these regions, we have created a bucket that contains files of various sizes, from 1 B to 100 MiB as showed in Tab. I. **File sizes** have been selected to assess network performance against objects of different nature, possibly related to diverse use cases. We have used the standard *HTTP GET* method to download these files from the buckets, in order to emulate the common behavior of the vast majority of cloud-storage customer applications [14].

In order to take into account different application needs and use cases, we have considered two different types of **storage classes**, i.e. Amazon S3 standard and CF. While the former is a solution suitable for a large variety of applications that do not have strict requirements on data transfer latency, the latter allows to distribute contents to the final users through CDNs, at both higher performance and cost.

Moreover, to take into consideration the heterogeneity of users of the cloud-storage services, and therefore the ability of these services to serve users spread worldwide, we have leveraged the facilities made available by the *Bismark infrastructure* to emulate the usage made by customers spread worldwide [22]. We have selected 77 geographically distributed Bismark nodes (*vantage points, VPs*) also according to their availability (nodes consist of home routers hosted on a volunteer basis). VPs are located in the United States (US, 36 VPs), Europe (EU, 16), Central-South America (CSA, 4), Asia-Pacific Region (AP, 12), and South Africa (ZA, 9). In our dataset, each VP is assigned an identifier made up of (i) the geographic set it belongs to and (ii) an incremental number. In the following we will refer to each of these geographic sets as **source region**.

It is worth noting how the experimental setup remarkably extends the analyses previously presented in the scientific literature for what concerns both the VPs leveraged and the factors considered.

B. *Experimental campaign and dataset*

All the results presented in this section refer to experimental campaigns conducted in May 2016. In order to collect the dataset we refer to, the VPs have been instructed as detailed in the following. Each VP has performed repeated *download cycles* over 7 days. Each cycle is composed of 40 sequential download requests spaced out by 10 seconds and uniquely identified by a combination of factors in Tab. I, i.e. cloud region, file size, and storage class. Downloads within cycles are randomly scheduled and repeated from each VP every 2 hours. After every download, VPs have run *TCP-traceroute* toward the IP address that has served the request in order to trace the information related to the path and estimate the RTT to the bucket (note that this information is not always available, due to the version of the firmware of the Bismark nodes and to the measurement tools available on them). The collected dataset is publicly available.[1]

Note how—according to the non-cooperative monitoring approach implemented—none of the information is obtained leveraging privileged information or the collaboration of the provider.

IV. RESULTS

In this section we discuss in details the most interesting results stemming out from our experimental campaign.

General overview of the performance. Our experimentations confirm that the size of the object to retrieve heavily impacts the measured performance of the network [11], [14], independently from the VP.

We report here the distribution of the performance in terms of the average *goodput* per download calculated over all the dataset. As shown in Fig. 1, the larger the size, the higher the measured goodput is. In more particular, 1 MiB, 1 KiB, and 1 B reported on average goodput values 1, 3, and 6 orders of magnitude lower than those registered with the download of 100 MiB-sized objects, respectively. While this trend was expected, as being related to TCP dynamics that make instantaneous throughput values grow over time, it may still heavily impact user experience. Indeed, users deal with different performance levels in terms of goodput when adopting the S3 service to retrieve objects of different sizes. For instance, from the performance point of view, it is convenient for the user to download a single 100 MiB blob containing clustered contents (e.g., an archive containing a hundred pictures) instead of a hundred 1 MiB files (e.g., single pictures), as the total time needed to retrieve contents would be less.

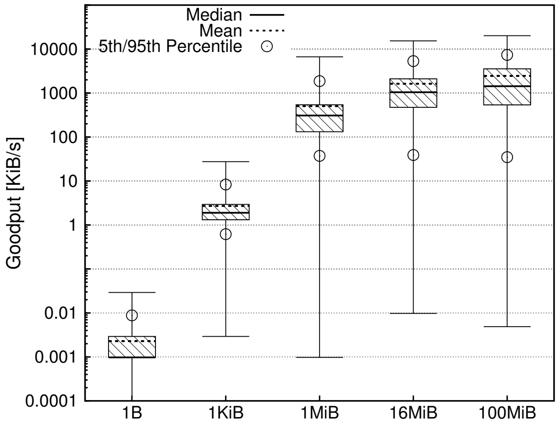

Fig. 1: General overview of S3 performance grouped by object size. Measured goodput heavily depends on size.

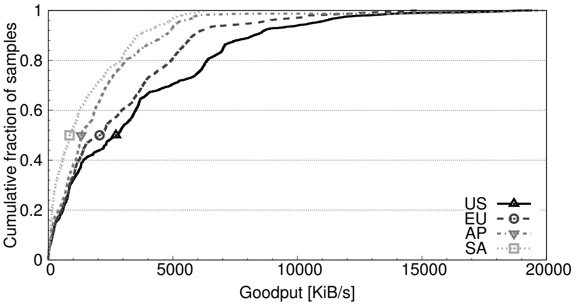

Fig. 2: S3 performance for objects of 100 MiB size, grouped by cloud region. Average goodput for US, EU, SA, and AP cloud regions was 3562, 2791, 1445, and 2018 KiB/s, respectively.

From the monitoring viewpoint, the slight variation observed on average between 16 MiB and 100 MiB sizes, suggests that 100 MiB is enough to obtain a good estimation of the maximum goodput achievable. Considering 100 MiB results as a reference, a limited error is done when considering 16 MiB objects (33.6%, on average), while a larger one when referring to 1 MiB contents (79.5%, on average). In the following analyses, if not explicitly stated otherwise, we will restrict the results presented to 100 MiB objects, as being the size reporting the best performance observed, on average.

Impact of the geographic region. Our results reported that the measured performance may be heavily impacted by the placement of both the bucket and the VP.

In order to evaluate how the performance changes when relying on cloud datacenters placed in different geographic regions, in this section we first compare the performance of the four cloud regions considered, as observed from the 77 distributed VPs. Considering the goodput average values from all the VPs (see Fig. 2), US, EU, SA, and AP cloud regions reported 3562, 2791, 1445, and 2018 KiB/s, respec-

---

[1] http://traffic.comics.unina.it/cloud

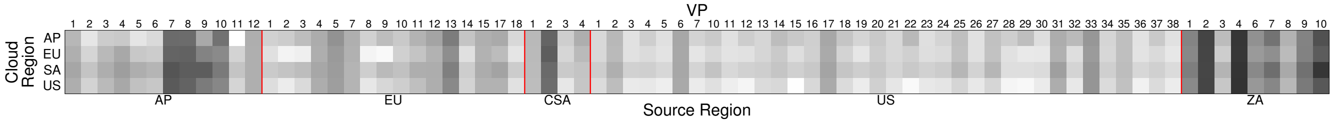

Fig. 3: Average S3 goodput from each VP when downloading 100 MiB objects from each cloud region (rows) for users placed at different locations (columns). Values vary from 10 KiB (black) to 100 000 KiB (white).

tively. Counterintuitively, AP and SA cloud regions are also associated to higher network-transfer costs with respect to EU and US. Considering the global mean as a threshold value, two performance classes can be identified: US and EU (whose values are higher than the threshold, on average) versus SA and AP (lower than the threshold, on average), where the former performs 45.5% better than the latter. On the basis of the adopted cloud deployment strategies that often see customers leveraging a single cloud region [9], [7] these results could be of interest for optimally choosing the cloud region to rely on. Indeed, if the cloud customer has no knowledge of the location of the users willing to retrieve contents, US and EU represent the best available choices, on average. Configuring the service to leverage the AP or SA cloud region leads to higher cost and lower performance when considering globally distributed users.

As expected, performance also strongly depends upon the placement of the VPs. The heatmap in Fig. 3 reports detailed information on how the goodput is subjected to changes on varying source VP and destination cloud region. While in some cases, for a fixed VP the performance is not subjected to significant variations when relying on different cloud buckets (e.g., AP12, EU6, US14), in other cases a non negligible discrepancy is measured on changing buckets (e.g., AP11, EU8, US13). It is evident in some of these cases how the globally optimal choice (EU or US region) is dramatically outperformed by local optimum choices.

Fig. 4 summarizes the general trends observed by considering the aggregated values for both VPs and bucket regions. As already observed for detailed results, lighter boxes placed on the diagonal show that often the best performance for a given VP placement region is obtained relying on a bucket placed in the same geographic zone. CSA sources (composed at 50% of nodes placed in Mexico and at 50% of nodes placed in Brazil) represent a notable exception as the best average performance associated to this source region is obtained when retrieving objects from a bucket placed in the US region. The best average performance is obtained considering VPs placed in the US region retrieving objects from a bucket in the same region. Conversely, the worst performance, on average, is associated to VPs placed in AP retrieving objects from SA, and vice versa (CSA from AP).

When considering homologous regions (i.e. VPs and cloud buckets placed in the same geographic region) the performance grows ($+44.36\%$, on average; see Fig. 5), with the AP region showing the major improvement ($+77.30\%$). Our dataset shows how this behavior is not completely generalizable, however. Notably, when considering CSA source vs. SA cloud region for instance, an improvement an order of magnitude less than the average is obtained ($+1.79\%$). Since CSA sources exhibit bad performance toward every cloud region, this phenomenon can be due to the access networks of the Bismark nodes available in CSA that are not under our control.

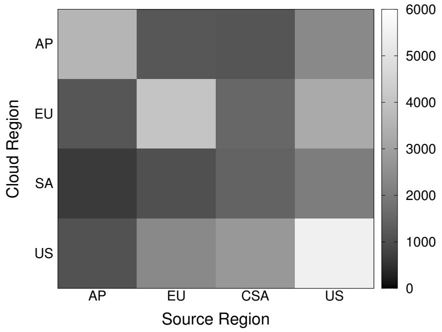

Fig. 4: S3 performance for objects of 100 MiB size. For each couple (source region, cloud region) the mean goodput is showed. ZA sources have been excluded as not having a homologous counterpart in Amazon S3 infrastructure.

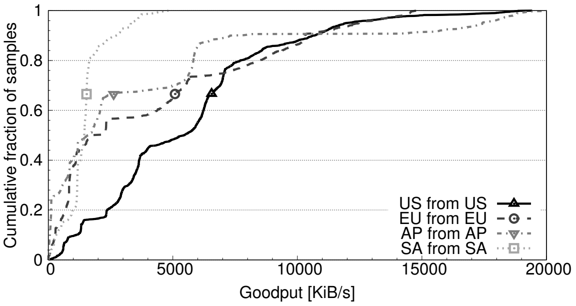

Fig. 5: S3 performance for objects of 100 MiB size for homologous source region and cloud region. Average goodput for US, EU, SA, and AP was 5453, 4046, 1464, and 3573 KiB/s, respectively. Compared to values reported in Fig. 2, average goodput values improve up to 77.30% (AP from AP).

Additional analyses are needed to investigate how much this finding is generalizable. This being said, this last analysis shows how cleverly choosing the cloud region according to the placement of the potential users allows the cloud customer to obtain best performance at an even lower cost.

Impact of the adoption of the CDN service. In our experimentations we also tested the performance variation achieved by enabling the content distribution through the CF service. Note that from the customer point of view CF does not require any additional configuration but only its explicit activation and implicates additional costs. Although CF generally leads to a performance enhancement, we found a number of cases for which this assumption does not hold.

Fig. 6 shows the distributions of the goodput obtained with both S3 standard and CF, considering two different sizes (1 MiB and 100 MiB). The beneficial impact of CF in terms of performance is evident, as it is able to deliver better performance on average ($+274.69\%$ and $+104.65\%$ when considering 1 MiB and 100 MiB content size, respectively).

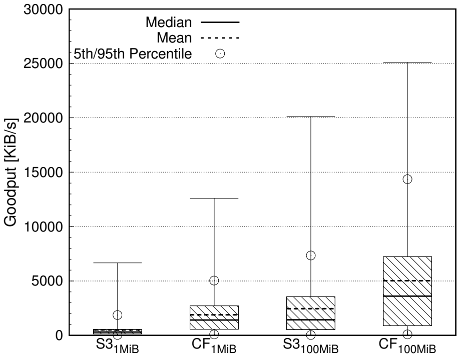

Fig. 6: Comparison of S3 and CF goodput for 1 MiB and 100 MiB file sizes. Adopting CF higher goodput values can be achieved on average.

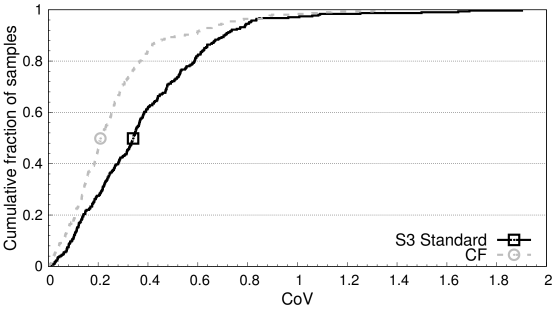

Fig. 7: Performance variability in terms of CoV for S3 standard and CF. Relying on CF guarantees lower variability in the perceived performance (-28%, on average).

Adopting CF is convenient also in terms of the variability of the performance, that appears to be reduced by one third. Fig. 7 reports the distribution of the performance variability, calculated as the CoV (Coefficient of Variation) over each pair (VP, destination).[2] CoV associated to CF is markedly lower (0.38 vs. 0.26, on average).

Each edge location observed has been associated to a geographic region leveraging DNS names. Fig. 8 graphically shows how edge locations have been assigned to CF download requests for each VP. In our experimentations we have been served by 38 out of the 54 available edge locations advertised by Amazon [1] ($18/21$ in US, $14/16$ in EU, $1/2$ in SA, and $5/15$ in AP). Indeed, in most of the cases the content is downloaded from an edge location placed in the same geographic region, although some variability in the association has been observed. When a VP is not served by an edge location placed in the same geographic zone, the content is always downloaded from a US bucket. This is in accordance with the policies advertised by the provider, explicitly saying that requests may be redirected to an edge location belonging to a cheaper geographic zone when needed. Interestingly, we found that the VPs placed in Mexico (CSA1 and CSA3) and in Brazil (CSA2) have been always served by edge locations placed in the US. Considering the performance associated to each edge location seen from each VP, we found that only for 18 out of the 48 VPs the best-performing edge location, on average, is the one seen with the highest frequency. In other words, for 62.5% of the cases, the strategy for associating an edge location to a VP leads to suboptimal results in terms of goodput [1]. We believe that the observed phenomenon is the result of the load balancing policies implemented by the provider to distribute requests across edge locations and are probably caused by edge-location overhead.

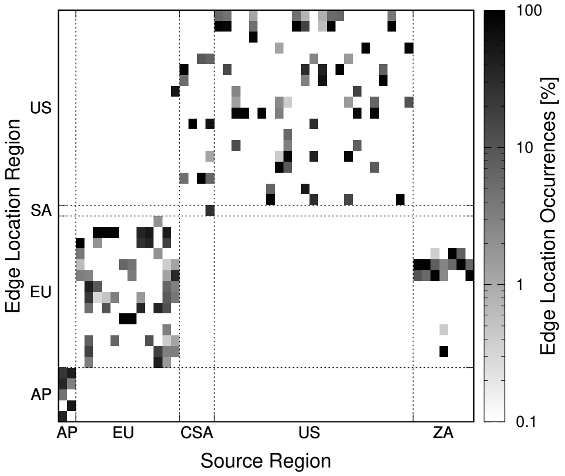

Fig. 8: Occurrences of the associations of different edge locations (on the rows) to each VP (on the column) for files of 100 MiB size. 38 out of 54 edge locations have been encountered. In most of the cases, objects are downloaded from an edge location placed in the same geographic region.

The distribution of the performance enhancement obtained enabling CF with respect to the different file sizes is reported in Fig. 9. CF leads to a 2215 KiB/s mean improvement considering 1 MiB, 16 MiB, and 100 MiB file sizes (i.e. $+144.11\%$, on average). Improvements in terms of goodput up to 20 371 KiB/s have been observed for 100 MiB objects, whereas the gain observed for 1 MiB objects is always lower than 7630 KiB/s (see Fig. 9a). However, more than 30% of the samples reported a negligible average improvement (i.e. lower than 100 KiB/s) with respect to the others. Most of them are related to VPs placed in the AP region.

Interestingly we found a number of cases for which adopting CF leads to worse performance than S3 (see the distribution of negative values in Fig. 9b). Considering all the combinations of file sizes, VPs, and cloud regions, we found 119 out of 924 cases in which S3 delivers, on average, better performance than CF, in terms of goodput and/or download elapsed time. The majority of these cases (more than 70%) refers to data retrieval from the bucket placed in the same geographic region of the VP, and involves 45 distinct VPs. While in some cases the performance degradation is negligible, in the worst case the negative variation in terms of mean goodput observed adopting CF, is equal to -43.59% (US38 retrieving a 1 MiB object from a bucket placed in US). In these cases, additional costs for enabling the CDN service don't generate any advantage for the cloud customer, even leading to performance degradation.

V. CONCLUSION

Cloud storage services offer a fast and convenient way to archive and share objects of different nature. The high-level management interface however, while guaranteeing ease

---

[2] The Coefficient of Variation, $CoV = \frac{\sigma}{\mu}$, is defined as the ratio of the standard deviation over the mean.

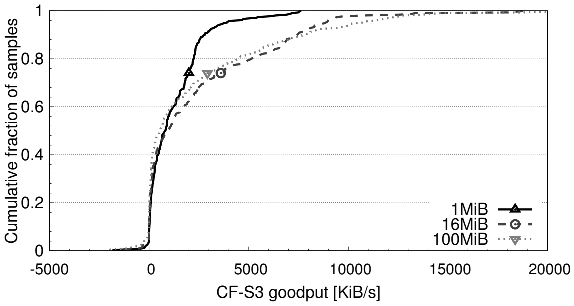

(a) Goodput gain distribution.

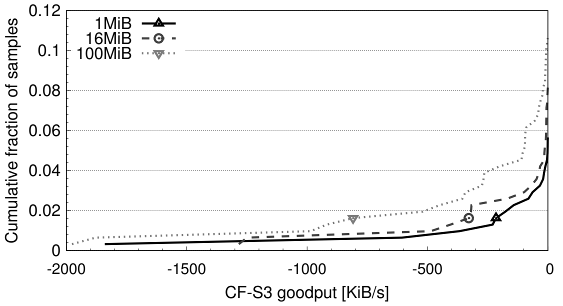

(b) Samples for which CF leads to worse performance (CF-S3<0).

Fig. 9: Performance improvement obtained when adopting CF. CF generally leads to a performance enhancement, but in some cases this does not hold.

of use, hides system implementation details and performance figures. In this work we have performed an experimental study about the performance of the cloud-to-user network for the Amazon S3 cloud-storage service, as it is perceived by a set of home users distributed all over the globe. Thanks to the dataset obtained leveraging the Bismark platform we report a general assessment of the performance of this service. We found that the US and EU cloud regions are able to offer better performance in terms of goodput ($+45.5\%$, on average) even at a lower cost, although sometimes this choice leads to suboptimal performance. Enabling CF leads to an average performance improvement ($+144.11\%$). However, a number of cases has been found for which relying on the CDN service is detrimental, generating up to a -43% performance decrement, even in presence of higher costs.

ACKNOWLEDGMENTS

This work is partially funded by art. 11 DM 593/2000 for NM2 srl (Italy). The experimental work in this paper was also supported by a grant provided by Amazon AWS in Education. We gratefully thank the Bismark Project and G. Martins for his precious help in configuring the Bismark measurement platform for our experimentations.

REFERENCES

[1] Amazon CloudFront – Content Delivery Network (CDN) website. http://aws.amazon.com/cloudfront/, Mar. 2016.

[2] Amazon Simple Storage Service website. http://aws.amazon.com/s3/, Mar. 2016.

[3] Amazon Web Services Global Infrastructure. http://aws.amazon.com/about-aws/global-infrastructure/, Feb. 2016.

[4] Ten Years in the AWS Cloud – How Time Flies! https://aws.amazon.com/blogs/aws/ten-years-in-the-aws-cloud-how-time-flies/, Mar. 2016.

[5] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, et al. A view of cloud computing. *Communications of the ACM*, 53(4):50–58, 2010.

[6] I. Bermudez, S. Traverso, M. Mellia, and M. Munafó. Large scale observation and analysis of Amazon AWS traffic. (October), 2012.

[7] I. Bermudez, S. Traverso, M. Mellia, and M. Munafo. Exploring the cloud from passive measurements: The Amazon AWS case. *2013 Proceedings IEEE INFOCOM*, pages 230–234, 2013.

[8] S. L. Garfinkel. Technical Report TR-08-07 : An Evaluation of Amazon ' s Grid Computing Services : EC2 , S3 and SQS. *Applied Sciences*, pages 1–15, 2006.

[9] K. He, A. Fisher, L. Wang, A. Gember, A. Akella, and T. Ristenpart. Next stop, the cloud: Understanding modern web service deployment in ec2 and azure. In *Proceedings of the 2013 Conference on Internet Measurement Conference*, IMC '13, pages 177–190, New York, NY, USA, 2013. ACM.

[10] Z. Hill, J. Li, M. Mao, A. Ruiz-Alvarez, and M. Humphrey. Early observations on the performance of windows azure. In *Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing*, pages 367–376. ACM, 2010.

[11] A. Li, X. Yang, S. Kandula, and M. Zhang. Cloudcmp: Comparing public cloud providers. In *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement*, IMC '10, pages 1–14, New York, NY, USA, 2010. ACM.

[12] H. Liu and S. Wee. Web server farm in the cloud: Performance evaluation and dynamic architecture. In *Cloud Computing*, pages 369–380. Springer, 2009.

[13] J. C. Mogul and L. Popa. What we talk about when we talk about cloud network performance. *SIGCOMM Comput. Commun. Rev.*, 42(5):44–48, Sept. 2012.

[14] M. Palankar, A. Iamnitchi, M. Ripeanu, and S. Garfinkel. Amazon S3 for Science Grids: a Viable Solution. *Proc. ACM Int. Workshop on Data-aware Distributed Computing*, pages 55–64, 2008.

[15] V. Persico, A. Botta, A. Montieri, and A. Pescapé. A first look at public-cloud inter-datacenter network performance. In *2016 IEEE Global Communications Conference: Communication QoS, Reliability and Modeling (Globecom2016 CQRM)*, Washington, USA, Dec. 2016.

[16] V. Persico, P. Marchetta, A. Botta, and A. Pescapè. Measuring network throughput in the cloud: The case of amazon ec2. *Computer Networks*, 93:408–422, 2015.

[17] V. Persico, P. Marchetta, A. Botta, and A. Pescapé. On network throughput variability in microsoft azure cloud. In *Global Communications Conference (GLOBECOM)*, 2015.

[18] V. Persico, A. Montieri, and A. Pescapé. CloudSurf: a platform for monitoring public-cloud networks. In *2016 IEEE 2nd International Forum on Research and Technologies for Society and Industry Leveraging a better tomorrow (RTSI) (IEEE RTSI 2016)*, Bologna, Italy, Sept. 2016.

[19] C. Raiciu, M. Ionescu, and D. Niculescu. Opening up black box networks with cloudtalk. In *Proc. 4th USENIX Conference on Hot Topics in Cloud Computing*, page 6, 2012.

[20] H. Saljooghinejad, F. Cuadrado, and S. Uhlig. Let latency guide you: Towards characterization of cloud application performance. In *2015 IEEE 7th International Conference on Cloud Computing Technology and Science (CloudCom)*, pages 99–106. IEEE, 2015.

[21] J. Schad, J. Dittrich, and J.-A. Quiané-Ruiz. Runtime measurements in the cloud: Observing, analyzing, and reducing variance. *Proc. VLDB Endow.*, 3(1-2):460–471, Sept. 2010.

[22] S. Sundaresan, W. de Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescapè. Broadband internet performance: A view from the gateway. *SIGCOMM Comput. Commun. Rev.*, 41(4):134–145, Aug. 2011.

[23] O. Tomanek and L. Kencl. Claudit: Planetary-scale cloud latency auditing platform. In *Cloud Networking (CloudNet), 2013 IEEE 2nd International Conference on*, pages 138–146. IEEE, 2013.

[24] R. Tudoran, A. Costan, G. Antoniu, and L. Bougé. A performance evaluation of azure and nimbus clouds for scientific applications. In *Proceedings of the 2nd International Workshop on Cloud Computing Platforms*, page 4. ACM, 2012.

[25] G. Wang and T. Ng. The impact of virtualization on network performance of amazon ec2 data center. In *INFOCOM, 2010 Proceedings IEEE*, pages 1–9, March 2010.