Computer Networks 55 (2011) 1100–1113

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Towards the design of optimal data redundancy schemes for heterogeneous cloud storage infrastructures ☆

Lluis Pamies-Juarez *, Pedro García-López, Marc Sánchez-Artigas, Blas Herrera

Department of Computer Engineering and Maths, Universitat Rovira i Virgili, Av. Paisos Catalans 26, 43007 Tarragona, Spain

ARTICLE INFO

Article history:
Received 12 May 2010
Received in revised form 17 September 2010
Accepted 10 November 2010
Available online 19 November 2010
Responsible Editor: K. Kant

Keywords:
Cloud storage
Distributed storage
Data availability
Erasure codes

ABSTRACT

Nowadays, data storage requirements from end-users are growing, demanding more capacity, more reliability and the capability to access information from anywhere. Cloud storage services meet this demand by providing transparent and reliable storage solutions. Most of these solutions are built on distributed infrastructures that rely on data redundancy to guarantee a 100% of data availability. Unfortunately, existing redundancy schemes very often assume that resources are homogeneous, an assumption that may increase storage costs in heterogeneous infrastructures – e.g., clouds built of voluntary resources.

In this work, we analyze how distributed redundancy schemes can be optimally deployed over heterogeneous infrastructures. Specifically, we are interested in infrastructures where nodes present different online availabilities. Considering these heterogeneities, we present a mechanism to measure data availability more precisely than existing works. Using this mechanism, we infer the optimal data placement policy that reduces the redundancy used, and then its associated overheads. In heterogeneous settings, our results show that data redundancy can be reduced up to 70%.

1. Introduction

We are witnessing today the rapid proliferation of cloud storage services as means to provision reliable storage and backup of files. Amazon S3 [1] is a representative example but also Mosso [27], Wuala [38] or Cleversafe [10]. All of these services offer users clean and simple storage interfaces, hiding the details of the actual location and management of resources. Most of these clouds (e.g., [1,27]) are built on well-provisioned and well-managed infrastructures, typically data centers, that are responsible for provisioning users with storage services. Very often, these data centers are controlled exclusively by cloud providers (e.g., Amazon, Google, Microsoft, etc.), whereas the user pays a price for the use of their resources. There is also the notion of resource-performance guarantee between the cloud provider and the user, that ensures that the user sees the performance he/she expects to see.

Current cloud storage infrastructures are focused on providing users with easy interfaces and high performance services. However, there are some classes of storage services for which the current cloud model may not fit well. For example, consider a research institute that wishes to freely share its results with others institutions as a "public service" (e.g., in form of a digital repository to make protein research more accessible to scientists), but it requires deployment resources. Since this service may not be commercial, the service deployers may not want to pay the cost of running the service.

To host such services, Chandra and Weissman in [7] proposed the idea of using voluntary resources (those donated by end-users in @home systems [2] and peer-to-peer

☆ An early version of this work was presented at ICPP'09 [26]. This work has been partially funded by the Spanish Ministry of Science and Innovation through project DELFIN, Ref. TIN2010-20140-C03-03.

* Corresponding author.
E-mail addresses: lluis.pamies@urv.cat (L. Pamies-Juarez), pedro.garcia@urv.cat (P. García-López), marc.sanchez@urv.cat (M. Sánchez-Artigas), blas.herrera@urv.cat (B. Herrera).

1389-1286/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.comnet.2010.11.004