The Journal of Systems and Software 86 (2013) 1907–1922

Contents lists available at SciVerse ScienceDirect

The Journal of Systems and Software

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

An approach for constructing private storage services as a unified fault-tolerant system

J.L. Gonzalez[a,c,d,*], Jesus Carretero Perez[a], Victor Sosa-Sosa[d], Juan F. Rodriguez Cardoso[a], Ricardo Marcelin-Jimenez[b]

[a] Carlos III University, Computer Science Department, Madrid, Spain
[b] UAM-Iztapalapa, Electrical Engineering Department, Mexico City, Mexico
[c] Cd. Valles Institute of Technology, San Luis Potosi, Mexico
[d] Information Technology Laboratory, Center of Research and Advanced Studies of the National Polytechnic Institute (CINVESTAV), Ciudad Victoria, Mexico

ARTICLE INFO

Article history:
Received 7 September 2012
Received in revised form 26 February 2013
Accepted 26 February 2013
Available online 16 March 2013

Keywords:
Storage services
Web-based storage
Fault-tolerance
Virtualization
Federation

ABSTRACT

Organizations are gradually outsourcing storage services such as online hosting files, backup, and archival to public providers. There are however concerns with this process because organizations cannot access files when the service provider is unavailable as well as they have no control and no assurance on the management procedures related to data. As a result, organizations are exploring alternatives to build their own multi-tenant storage capacities.

This paper presents the design, implementation and performance evaluation of an approach for constructing private online storage services. A hierarchical multi-tier architecture has been proposed to concentrate these services in a unified storage system, which applies fault-tolerant and availability strategies to the files by passing redundant information among the services or tiers. Our approach automates the construction of such a unified system, the data allocation procedure and the recovery process to overcome site failures. The parameters involved in the performance of the storage services are concentrated into intuitive metrics based on utilization percentage, which simplifies the administration of the storage system. We show our performance assessments and the lessons learned from a case study in which a federated storage network has been built from four trusted organizations spanning two different continents.

1. Introduction

Web-based storage solutions are a set of value-added services such as content synchronization among different devices (Dropbox; sugarsync) (e.g. tablets, phones, computers), online backup (Garg, 2009; Cleversafe) and remote file access via web dashboards (Dropbox).

This kind of services have become a common solution for the large majority of users – individuals or organizations – that need to store or retrieve their files *any time from anywhere* (ATAW).

The users can assess web-based storage by choosing either a Pay As You Go pricing model in which some quality of service can be contracted (Windows azure data and storage; Amazon) or free file hosting systems (FHS) (Dropbox; sugarsync) in which the quality, security and privacy are based on a best effort. Both storage solutions are becoming a business success case (Annette Jump, 2011; Richard et al., 2011), which has increased both the web and file sharing traffic on the Internet (Waterloo, 2011).

However, the users of this kind of services have expressed their concerns about the lack of control on the data management procedures in charge of their information (Chow et al., 2009). There is also uncertainty of legal aspects as they delegate a private content to a *third-party* (Smith, 1979; Salveggio, 2004), for instance, the possibility of performing data mining on private data (Google privacy and policies), and the information vulnerability, in case that the storage provider were compromised (Nytimes; Google, 2011).

In addition, regulations in different countries require organizations to guarantee the integrity and confidentiality of their files during large periods of time such as six years in the case of medical contents (Law, 2011) or three for financial information (U. IRS; T.R.R. Guide). Under the current growing rates, this is enough time to achieve massive volumes that will have an impact on middle- and long-term storage costs (Gantz and Reinse, 2011; Walker et al., 2010).

Price is a aspect to weigh up when users choose a storage provider (Ion et al., 2011; fujitsu. and com Masaharu Sato,

* Corresponding author at: Carlos III University, Computer Science Department, Madrid, Spain.
E-mail addresses: joseluig@ac.upc.edu, joseluig@arcos.inf.uc3m.es (J.L. Gonzalez).

0164-1212/$ – see front matter © 2013 Elsevier Inc. All rights reserved.
http://dx.doi.org/10.1016/j.jss.2013.02.056

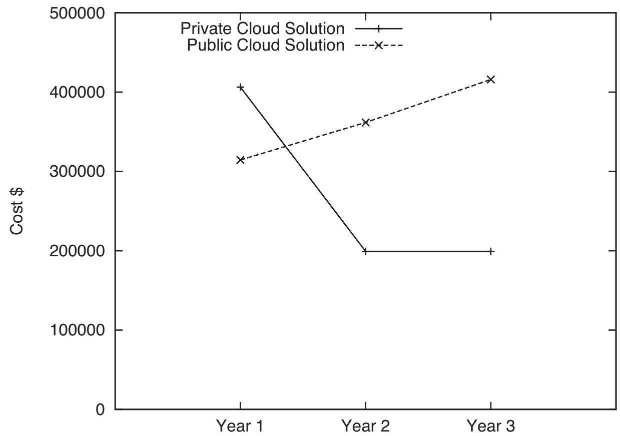

Fig. 1. Costs for private and public solutions.

2010), specifically, *long-term costs and penalizations* are considered because prices are set per unit of storage and time. The longer some content is stored in the cloud, the higher the price to be charged.

The costs associated to private and public storage solutions have been widely studied (Mazhelis et al., 2012; Chen and Sion, 2011; Walker et al., 2010; Singh and Jangwal, 2012).

In the case of private storage solutions, the most significant cost is the capital investment (site infrastructure and the space required for it), which is observed in the first year and also includes the investment required when the storage must be replaced because of its life expectancy (Mazhelis et al., 2012; Chen and Sion, 2011). The common operating expenses in private solutions are energy, personal and maintenance, which represent a constant cost that depend on the country where the site has been installed.

In the case of storage solutions offered by public providers, the operating expenses are commonly the service costs such as storage, network, databases (server hours and transactions), and data transfer. Additional charges could be applied to the organization's account when contracting functionalities such as geographical redundancy, auditing and monitoring systems.

Fig. 1 shows the results of a cost comparison over three years (Singh and Jangwal, 2012) for both public (100 EC2 Amazon instances and S3 storage) and the same solution developed in a private site.

The comparison shown in Fig. 1 of course is based on assumptions. The assumptions affecting the costs of the public solution are 24/7 availability and an increment up 15% in the resources for facing up a given increment of the demand.

The costs assumed in the private solution are the virtual server deployment costs, the yearly support cost they occur as well as the allocated maintenance. It also assumes the organization did not include the professional services to install the private cloud infrastructure because organization technical staff had the skills to do it themselves.

Fig. 1 illustrates the costs of public and private cloud solutions with a budget estimate. This means the costs can be different when the organizations establish different scenarios with different assumptions. The organizations can reproduce this type of exercises either by using a budget estimate web site Planforcloud in which the administrators can either choose cloud services from several providers or use a mathematical model (Walker et al., 2010).

The private storage scenarios are acceptable for organizations having space to allocate infrastructure, handling large volume of data for several years and that are interesting in the development of additional functionalities to avoid both the privacy issues of the outsourcing and the service unavailability produced by the vendor lock-in.

As a result, organizations have started to consider the possibility of building their own in-house storage capacities as information is becoming a very important asset, and this is a promising alternative to keep in control of this valuable resource (Peters, 2011).

This decision is fostering studies that aim to develop new fault-tolerant storage systems based on commodity components as well as guidelines for configuring such environments in a cost-effective manner (Annette Jump, 2011).

This paper presents the design, implementation and performance evaluation of an approach for constructing private online storage systems by using commodity components with minimal proprietary requirements.

This approach introduces the following contributions:

- *A multi-tier architecture of storage services*. A multi-tier hierarchical architecture has been proposed to address user/organization concerns about data confidentiality and integrity. The first tier or *Client tier* is the ambit of the user and it includes local storage service. The second or *Trusted tier* is the ambit of the organizations and includes online hosting file service. The third one or *Trusted Shared tier* is a federated ambit including backup and archival services. That architecture preserves the control on the data assets management and offers options to decide the best suited place to store/retrieve each file. The tiers of this architecture perform analogous functions to traditional storage but they are applied to the online storage service. This means that the *Client tier* can be regarded as the cache, the *Trusted tier* works on the role of the disk, and the *Trusted Shared tier* plays the role of tape system. We have defined that hierarchy according to the latency and costs to achieve data availability at each tier.
- *A virtualization mechanism*. A mechanism automates the virtualization of the multi-tier architecture of storage services as a single unified system. It also automates the placement and distribution of data as well as the recovery procedure to overcome site failures.
- *A data allocation and load balancing method.* This data placement method has been designed for passing anonymized files through tiers in a balanced manner. The goal of this method is preserving the autonomy of the storage services and keeping metadata in-house.
- *Intuitive metrics to simplify the administration of the storage system.* Our approach proposes to concentrate the parameters involved in the performance of the storage services into easily-understandable metrics called *Utilization Factors* (UF), which represent the percentage of consumed capacity of a disk, server, tier or the whole system.
- *A set of fault-tolerant strategies.* We propose reactive and proactive redundancy strategies to deliver data availability when suffering from failures such as disk/server crashing, virtual machine blackout and even a multi-site disaster.

To show the feasibility of our proposal, this paper also presents a case study describing the implementation of four private online storage systems based on four trusted organizations, which make part of a federated storage network.

The approach provides an isolated storage system for each trusted organization in which users store and retrieve their files in ATAW manner by using the unified organization domain. The users are also able to store/retrieve files by using the federated shared storage when the site of their organization is down.

In the implementation, each organization keeps in-house control on metadata management, file allocation, resource location and access methods, as well as a given data availability and reliability degree. Moreover, organizations decide the amount of resources to be assigned to the federated storage. We show our performance