Data-intensive Storage Services on Clouds: Limitations, Challenges and Enablers

Elliot K. Kolodner
IBM Haifa Research Lab

Abstract

The emergence of cloud environments has made feasible the delivery of Internet-scale services by addressing a number of challenges such as live migration, fault tolerance, and quality of service. However, current approaches do not tackle key issues related to cloud storage, which are of major importance given the amount of data produced by various sources (e.g., smart phones, cameras, social networks, etc). In this paper we discuss the challenges to providing optimized data-intensive storage services and enablers to meet them. We are building a scalable and flexible infrastructure facilitating a new data model to raise the abstraction level of storage, data mobility, computational and content-centric access to storage as well as mechanisms for cost-efficiency with provisions for QoS and security guarantees.

1. Introduction

The explosion of personal and organizational digital data has been recognized as one of the most significant characteristics of the decade. Generated data is growing faster than we can store it, while in parallel our society has become critically dependent on services requiring this data. In the emerging era of the Future Internet, the explosion of raw data and the dependence on data services is expected to be further amplified due to the strong proliferation of data-intensive services and the digital convergence of telco, media and ICT.

Emerging cloud-based infrastructures for storage have been widely accepted as the next-generation solution to address the data proliferation and the reliance on data. Such infrastructure promises essentially infinite scalability and capacity, with continuous and ubiquitous availability at substantially lower costs than traditional storage solutions. Nevertheless, today's cloud environments need to overcome specific limitations. In this paper, we start by highlighting these limitations and the corresponding challenges through a sample scenario. We then focus on enablers to overcome them.

2. Objectives

To demonstrate the limitations of today's infrastructures and to motivate the objectives of our work we present a futuristic scenario examining a day in the life of a citizen, Ann, who strongly relies on the wealth of services enabled by the Future Internet in 2013. Ann is a journalist for a large European broadcasting company, that delivers news via TV, radio and web, and depends on reliable, secure and efficient access to data from anywhere at any time, regardless of the data's physical location and representation.

1. Ann's company has recently acquired two new media-rich news channels aiming to improve news coverage for subscribers across Europe. The content of these channels must be integrated into the company's network and be supported by the consolidated ICT infrastructure including the company's website.
2. As new items are being posted by journalists and read by viewers, the channels analyze their popularity. While popular items have multiple replicas in the relevant geographic locations, non-popular items may be moved to storage in less expensive locations to reduce storage costs.
3. Ann's management decides to launch a campaign to publish the new integrated content offering. Ann is asked to prepare a short film to advertise the acquisition. The film will include a short historic summary of the new channels, using news items taken over their 20 years of operation.

3. Methodology and challenges

To support such a scenario, current cloud infrastructures need to address specific research challenges, such as:

- Raising the abstraction level of storage: The broadcasting company operates in an environment where it deploys new media services, based on content that is distributed and stored across distant locations in a multitude of files and formats. Data objects and associated rich metadata would simplify management of this content.
- Data mobility and federation: Ann needs to integrate the content repositories of the two new channels acquired by her company with the existing ones. Federation mechanisms and protocols are needed to allow transparent federation and interoperability across different providers without reformatting the new content.
- Computations: As a part of the channel integration, content in the acquired channels will be analyzed as part of the content delivery process. Separately, the news channels collect statistics to determine the popularity of local content according to content type, viewer profiles, viewing hours, etc. The ability to perform computations close to storage would facilitate the execution of efficient and cost-effective computations.
- Access to storage: Ann searches for key events of interest which are captured and stored over the years in the archives of the company. The ability to access objects, old or new, in a federated cloud, based on their content and relationships regardless of their physical location, representation and type, would enable simplified and efficient searching in the archived data.
- Other capabilities of cloud-based storage: For best viewing experience, media items are served from the replicas closest to the viewers' geographic location. The infrastructure should provide management mechanisms to ensure QoS by dynamically optimising the location and the number of object replicas.

4. Technology

We suggest meeting the above challenges through five enablers: (1) raising the abstraction level of storage, (2) data mobility and federation without boundaries, (3) computational storage, (4) content-centric access and (5) advanced capabilities for cloud-based storage.

4.1 Raising the abstraction level of storage

Recently, initial object models for the cloud have emerged, replacing traditional file systems and adapting them to cloud scale. These models have flattened the tree hierarchy of the file system, which is no longer relevant, and relaxed semantics and consistency. However, these models do not support rich metadata, and their security model is weak, depending on Access Control Lists (ACLs). They are also proprietary.

We are working on a new **data model** which fits the scale of the cloud, yet has rich metadata and a supports a strong yet flexible security model. The metadata will be an integral part of the storage, allowing the description of object content and its handling. This new metadata model supports effective access, management and manipulation of the storage and stored data and also the derivation of new metadata.

4.2 Data mobility and federation without boundaries

A key aspect missing from today's large storage systems is the capability of data mobility. Users cannot easily migrate their data across providers and thus suffer from data lock-in with a specific provider [1].

We are developing new approaches including efficient methods for large-scale data migration with extensions to lazy copying techniques and de-duplication, as well as platform-independent data migration and methods for federated access to data with emphasis on federated sets of data objects.

4.3 Computational storage

Compute and storage are, in most cases, treated as two different resources in a decoupled manner. Given that bandwidth is neither infinite nor networking costs negligible, it is clear that for many applications it is better to move the computation to the data, rather than bring the data to the computation. As the cloud model has emerged, this idea of bringing compute to storage has been applied for restricted programming paradigms, e.g., MapReduce [2] .

We are developing a construct called **storlet**, which specifies a general computation to be executed on a storage object. Such computations can be supplied by the user or an application and can be shipped to the storage for execution, or can be located in the cloud and scheduled to be executed near the storage upon invocation. They may involve user driven computational tasks as well as autonomous dynamic data derivation and transformation.

4.4 Content-centric Access

To access a traditional storage system, the user must know the location of the system and the place of the artifact or file in a storage hierarchy. Current storage systems for cloud applications overcome these limitations, but they are tailored to a specific domain, for example, Google books, Flickr and YouTube.

In contrast, we are developing a generic storage cloud system with methods to facilitate **access to storage by its content and relationships**, rather than details of the underlying storage containers. These methods will support any domain by allowing the definition of domain-specific storage optimizations. They will scale to the cloud, allowing storage to be spread out across multiple data centers.

4.5 Capabilities for Cloud-based Storage

Although cloud storage as available from providers such as Amazon offers useful features such as demand-based access to raw storage resources, it is not ready to store the critical data of individuals, businesses and governments with the required reliability and QoS as evidenced, for example by the rudimentary SLA of Amazon S3 [3,4].

We are addressing quality, cost and security assurance aspects as required for business critical and sensitive applications. Our approach will support multi-tenancy, where multiple organizations share the same storage infrastructure, guarantee secure and authorized access to the data and services, provide tools for checking compliance with standards and regulations and support SLAs for reliability and QoS.

5. Conclusions

We are building a scalable and flexible infrastructure to meet the challenges based on the enablers that we outlined: a new data model to raise the abstraction level of storage, data mobility and federation, computational and content-centric access to storage, and advanced capabilities for cloud-based storage.

6. References

[1] "Above the Clouds: A Berkeley View of Cloud Computing", February 2010, http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.pdf

[2] "MapReduce: simplified data processing on large clusters", Jeffrey Dean and Sanjay Ghemawat, http://portal.acm.org/citation.cfm?doid=1327452.1327492

[3] "Amazon Web Services Goes Down, Takes Many Startup Sites with It", Tech Crunch Feb 2009.

[4] Amazon S3 Service Level Agreement, http://aws.amazon.com/s3-sla/