Journal of Computer and System Sciences 82 (2016) 282–309

Contents lists available at ScienceDirect

Journal of Computer and System Sciences

www.elsevier.com/locate/jcss

# Exploiting user metadata for energy-aware node allocation in a cloud storage system

Cengiz Karakoyunlu [*], John A. Chandy

*Department of Electrical and Computer Engineering, University of Connecticut, Storrs, CT 06269, United States*

ARTICLE INFO

*Article history:*
Received 4 August 2015
Accepted 21 September 2015
Available online 21 October 2015

*Keywords:*
Energy efficiency
Cloud storage
Cloud systems
Load balancing
Metadata aware placement

ABSTRACT

Cloud computing has gained popularity in recent years delivering various services as cost-effective platforms. However, the increasing energy consumption needs to be addressed in order to preserve the cost-effectiveness of these systems. In this work, we target the storage infrastructure in a cloud system and introduce several energy efficient storage node allocation methods by exploiting the metadata heterogeneity of cloud users. Our proposed methods preserve load balance on demand and switch inactive nodes into low-energy modes to save energy. We provide a mathematical model to estimate the outcome of proposed methods and conduct theoretical and simulative analyses using real-world workloads.

## 1. Introduction

Cloud systems have gained popularity over the last decade by enabling users to store data, host applications and perform computations over the network. Cloud systems significantly decrease the cost on the user end as management, maintenance and administration tasks are typically handled by the cloud providers. Cloud providers also benefit from this scheme as they can utilize system resources more efficiently through techniques, such as virtualization, enabling them to achieve better performance and energy efficiency. There are numerous cloud providers offering a broad range of services [1–3].

Even though cloud systems tend to have lower energy costs compared to traditional HPC clusters due to better utilization techniques, the increasing energy consumption of cloud systems still needs to be addressed as the amount of data stored and the number of computations and applications in cloud increase steadily. According to [4], the estimated energy consumption of U.S. data centers was more than 100 billion Kilowatt hours in 2011. Between 2005 and 2010, data centers consumed between 1.7% and 2.2% of all electricity in the U.S. [5], matching the energy consumption of the aviation industry. Increasing energy consumption also means higher cooling costs and additionally, a cloud system with an underperforming cooling mechanism may have reliability issues potentially causing violations of service-level agreements (SLA) with the cloud users. It is, therefore, very important to have an energy-efficient cloud system not only for lower energy costs, but to also meet performance demands of the cloud users as well.

A typical cloud system is shown in Fig. 1. In this work, we refer to any application using the backend storage of a cloud system as the *user* of that storage system. There are several components in a cloud system contributing to the overall energy consumption – namely processing units, network components and storage systems. In this work, we specifically target the energy consumption of the cloud storage infrastructure forming the backend of the cloud computing units, since the

---

[*] Corresponding author.
*E-mail addresses:* cengiz.k@uconn.edu (C. Karakoyunlu), john.chandy@uconn.edu (J.A. Chandy).

http://dx.doi.org/10.1016/j.jcss.2015.09.003
0022-0000/

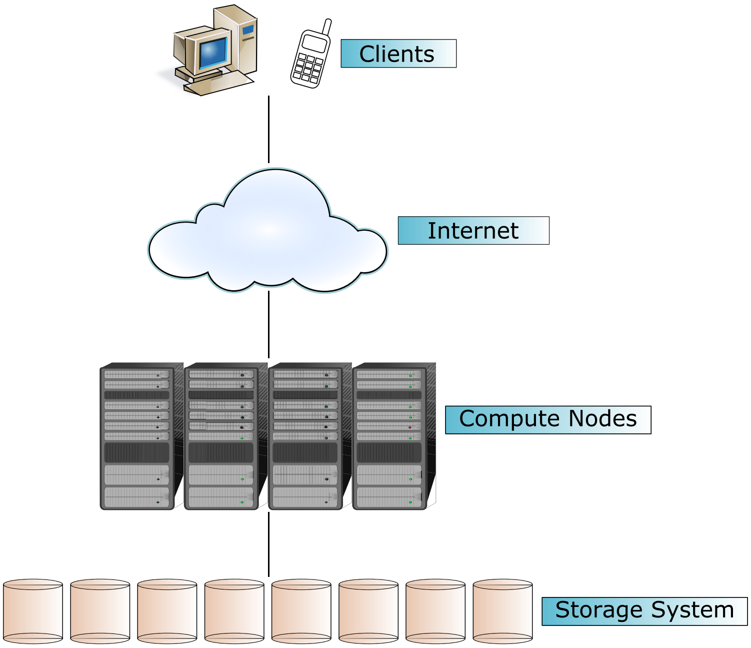

**Fig. 1.** Typical cloud system architecture.

storage system costs constitute an important fraction (between 25–35%) of overall cloud system costs [6–8]. Storage systems also have more idle periods since data stored is usually redundant and archival, written once and not touched again [9]. There have been many studies to reduce the energy consumption of other components of cloud systems. However, since idleness is not usually available in the network and processing units of clouds, these studies have been mostly on virtual machine consolidation, workload characterization, data migration or scheduling. In this work, we propose methods to have an energy-aware cloud storage system while at the same time trying to achieve uniform system utilization on demand. In particular, we take advantage of idleness existing in cloud storage systems and try to switch inactive nodes into low power modes. Our work is driven by two key assumptions: first, the cloud storage system suffers from incasting, and second, most of the data stored in the cloud (as much as 75%) is not heavily accessed [8], creating idle periods. It is important to note that our methods can also be implemented in the storage systems that form the backend of computational platforms (i.e. Hadoop clusters), where there might be idle periods [10,11].

*Incasting* is a condition that occurs because of queue limitations in most network switches [12]. As a result of the incasting behavior, there is a limit to the number of storage nodes (e.g. 4 servers in a cluster-based storage network as in [13]) across which data can be striped for parallel access. Beyond this limit, the I/O bandwidth no longer scales and in fact deteriorates. If $M$ represents the number of nodes at which the performance maxes out, from a performance point of view, there is no point in using more than $M$ nodes to increase parallelism. While for performance scaling, we would only need $M$ storage nodes in the system, however, for storage capacity reasons, we may need more than $M$ nodes. In such systems, we would need to keep these extra nodes active and thus waste energy since the network resources are maxed-out by a *subset* (of size $M$) of the storage nodes.

To save energy, one could turn off these extra storage nodes, and activate them when necessary. The difficulty comes while trying to identify which storage nodes should be turned on or off. Our approach is to distribute cloud users across the storage nodes, such that each user is allocated only $M$ nodes – i.e. the limit at which performance is maxed out. Grouping data on a subset of the storage nodes and putting the remaining nodes into low power modes has been studied in many related studies [10,4,14]. However, the majority of these studies used data classifications, redundant data or a hot-cold zone approach to group data on a subset of the storage nodes. Our approach is different from existing studies in that we try to group *cloud users* on a subset of the storage nodes without any data classification. Not classifying data enables our methods to be implemented in cloud storage systems with any kind of redundancy scheme. It is important to also note that, because of the job processing structure (batch mode or intelligent schedulers) in clouds, we can effectively turn on and off storage nodes based on the user job submissions since we know a priori which $M$ nodes are going to be used at any time. Therefore, any latency due to transitioning a storage node from inactive to active mode can be hidden.

The challenge becomes now how to allocate a subset of cloud storage nodes for each user in order to reduce the energy footprint, while at the same time trying to preserve uniform distribution of the system resources if demanded. We exploit the heterogeneity in the user metadata for energy-aware storage node allocation. As an example, if we assume a storage system consists of $N$ storage nodes and $M$ of them max out the incast bandwidth ($N > M$), then each user can be assigned a separate subset of $M$ storage nodes based on a certain metadata (i.e. user id, usage pattern) and any storage node that is