Global Journal of Information Technology

Volume 06, Issue 1, (2016) 01-10

http://sproc.org/ojs/index.php/gjit

Challenges of big data storage and management

Rajeev Agrawal *, Department of Computer Systems Technology, North Carolina A&T State University, Greensboro, USA.

Christopher Nyamful, Department of Computer Systems Technology, North Carolina A&T State University, Greensboro, USA.

Suggested Citation:

Agrawal, R. & Nyamful, C. (2016). Challenges of big data storage and management. *Global Journal of Information Technology*. *6*(1), 01-10.

Received 08 January, 2016; revised 04 February, 2016; accepted 17 March, 2015.

Selection and peer review under responsibility of Prof. Dr. Adem Karahoca, Bahcesehir University, Turkey

Abstract

The amount of data generated daily by industries, large organizations and research institute is increasing at a very fast rate. These huge volumes of data need to be kept not just for analytic purposes, but also in compliance with laws and service level agreements to protect and preserve data. Storage and management are major concern in this era of big data. The ability for storage devices to scale to meet the rate of data growth, enhance access time and data transfer rate is equally challenging. These factors, to a considerable extent, determine the overall performance of data storage and management. Big data storage requirements are complex and thus needs a holistic approach to mitigate its challenges. This paper examines the challenges of big data storage and management. In addition, we also examines existing current big data storage and management platforms and provide useful suggestions in mitigating these challenges.

Keywords: big data, storage systems, challenges, performance.

*ADDRESS FOR CORRESPONDENCE: **Rajeev Agrawal,** Department of Computer Systems Technology, North Carolina A&T State University, Greensboro, USA. *E-mail address*: ragrawal@ncat.edu

1. Introduction

Huge data expounders like Google, Flickr, Yahoo and Facebook, run a hyper-scale computing environment. YouTube alone receives over 1 billion unique users each month, and 100 hours of video are uploaded to YouTube every minute [17]. With this explosive amount of data being generated, storage capacity and scalability has become a major issue. The storage demand for big data at the organizational level is reaching petabytes (PB) and even beyond.

Storing and maintaining large sets of data over time at the rate of growth can be challenging. Factors such as capacity, performance, throughput, cost, and scalability are involved in any ideal storage solution system. In addition, storage devices play an important role in mitigating big data challenges. Reliability is of equal concern for big data storage. Reliability is basically the retrieval of data in its original form without any loss. The issue of reliability takes into account both internal and external system failures, and vulnerabilities. With the scale of data, the probability of losing some data during retrieval can be very high. Large data intensive applications such as Google map, and Facebook requires high Input-Output-Operations-Per Second (IOPS) to maintain performance in order to stay in business.

The I.T departments of most large organizations are facing strict and tight budget, which is limiting their ability to manage the huge data at their disposal effectively. With the limited funds, organizations are now required to design techniques that falls within their budget. The ability to maximize performance and capacity while minimizing cost has become a headache for organizations operating big data scale, and research focus for academia. Their existing storage systems are inadequate to meet the stringent requirements of big data storage and management.

The amount of unstructured data is rapidly increasing with daily upload of high-definition videos and pictures with our mobile devices. There is a compelling need for intensive research on big data storage and management. Given the current volume of data, large organizations are employing techniques such as data compression, deduplication, object storage, and cloud storage. This paper looks at these various techniques and their corresponding trade-offs, and provides suggestions for big data storage and management. Rest of the paper is organized as follows. Section 1 examines the storage medium issues for data storage and management. Section 1.2 covers current storage architectures. Section 2 explores challenges of big data storage, section 3 deals with the management issues, and the fourth section provides some useful suggestions to meet these challenges, and finally the conclusion.

1.1. Storage Mediums Issues

Mechanical disk drives (HDD) and Solid State Disk (SSD) drives are the major trending storage mediums, with hard disk drives forming the basis of the bulk of big data storage. Solid State Drives and Hard Disk Drives are mostly used by organizations as their storage device, with their capacity density expected to increase at a rate of 20 percent [3]. HDDs characteristics are completely different from SSDs. I/O subsystems designed for HDDs in traditional storage systems do not work for SSDs (Agrawal et al., 2008). Reliability risk such as overheating and magnetic faults, and disk access overhead has made HDDs undesirable for big data storage, though price per gigabyte is relatively low. SSDs on the other hand can service I/O processing request at a much faster rate than HDDs, because there is no mechanical part, reducing access time hence increase in I/O rate. SSDs are more resistant to physical shock, hence more reliable. The problem with SSDs is the price per gigabyte. The cost of replacing all mechanical drives with SSDs for a big data storage is unreasonably high. Table 1 below compares the various storage mediums based on primary performance metrics. The hybrid, which is a combination of HDDs and SSDs are what some experts say provides the best price per performance ratio. High-capacity HDDs and a small percentage of SSDs should provide the needed capacity and performance for big data systems, provided the technical expertise is available to do the tiering.

Table 1. Storage Mediums Characteristics

	MAGNETIC STORAGE	OPTICAL STORAGE	SOLID STATE	HYBRID STORAGE
CAPACITY	Up to 6TB	Up to 50GB	Up to 2TB	Up to 6TB
ACCESS TIME	Relatively low latency	Slow access time	Very low latency	Low latency
COST	Less expensive	Relatively cheaper	Very expensive	Less expensive
DATA TRANSFER RATE	Relatively low transfer rate	Good transfer rate	High transfer rate	High transfer rate

1.2 Current Storage Architectures for Big Data Storage

- Cloud providers such as Drop-box, Google Drive, Google Docs, and Microsoft cloud provide storage for their clients in the form of Storage-as-a-Service (StaaS). Storage-as-a-Service allows organizations and individuals to rent a storage space at a relatively minimal cost. Cloud providers achieve this, by leveraging several hundreds to even thousands of servers across multiple data centers. Cloud storage relieves enterprises the burden of hardware, software, and storage. The cloud systems allow storage for large scale data, for efficient data analysis and management [5]. Storage in cloud computing plays a vital role in the large-scale services rendered to clients. It provides virtual storage infrastructure in the form of pay-as-you-go manner, such as Amazon S3. In order to meet the challenges imposed by big data storage, cloud providers employ a huge number of commodity disk, across several data centers. Amazon S3 aims to provide scalability, high availability, and low latency at commodity costs. The system provides efficient data loading from different sources, flexible data partitioning scheme, index and parallel sequential scan [6].
- Network Attached Storage (NAS) is designed for file sharing. It enhances storing, retrieving, and accessing files for applications and clients. It uses a high level abstraction that enables cross-platform data sharing. NAS comes with a processor and software for management and backup of data. NAS is efficient and reliable, and comes with huge storage capacity [11].

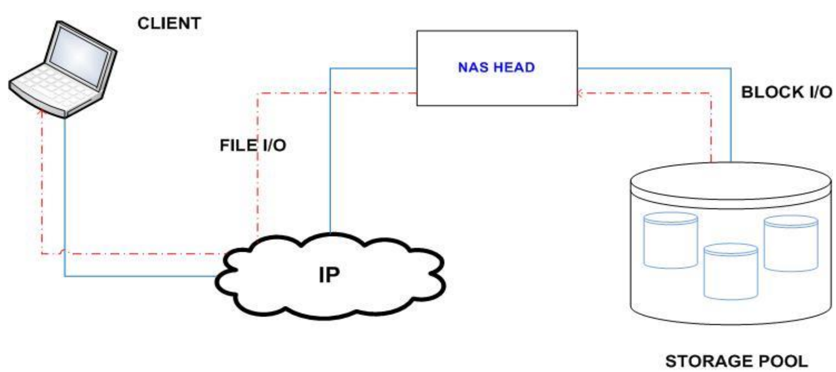

Figure 1. Architectural Overview of NAS

Figure 1 shows the NAS head connected to IP network and storage pools, converting file I/O to block I/O and vice versa. The benefits of this system is very compelling, in that it provides performance in I/O and throughput. Mostly, NAS systems are confined to corporate data centers, and accessible through Network File Systems (NFS) and Common Internet File Systems (CIFS) protocols. The massive increase of data and widespread of mobile devices means limited connectivity to files on this system, which poses concern. Another issue with these system is that, it's not able to scale beyond the capacity of the system, pushing customers to purchase additional separate boxes.

- Storage Area Network over IP (IP-SAN) provides the platform where thousands of computers connect to share a large amount of storage devices that range from simple disks to large, high-performance, high-functioning storage system. The widespread of IP networks makes IP-SAN attractive to many organizations. IP networks facilitate ease deployment of IP-SAN technology. It is less expensive, since existing IP-based network infrastructure can be leveraged to implement the technology. . IP-SAN can span over a wide geographical area, making organizations to consolidate their widespread storage infrastructure, and manage it centrally. Remote backup, replication, archiving, and recovery procedures are enhanced under the performance of IP-SAN. Three standard protocols – internet Small Computer Systems Interface (iSCSI), Fibre Channel over IP (FCIP), and internet Fibre Channel Protocol (iFCP) [1] enables the deployment of SANs using IP as a transport mechanism. They use encapsulation method to send and receive data payload. The major challenge of this system in the midst of huge data sets is high latency, because it uses commodity Ethernet.
- Object-based Storage: According to [11], storage object is said to be collection of bytes on a storage device, with methods for accessing data, and security policies to prevent unauthorized access. . Because of the variable-length nature of object, it makes it ideal to store the different types of data. Objects can be seen as the union of both file and block technologies.

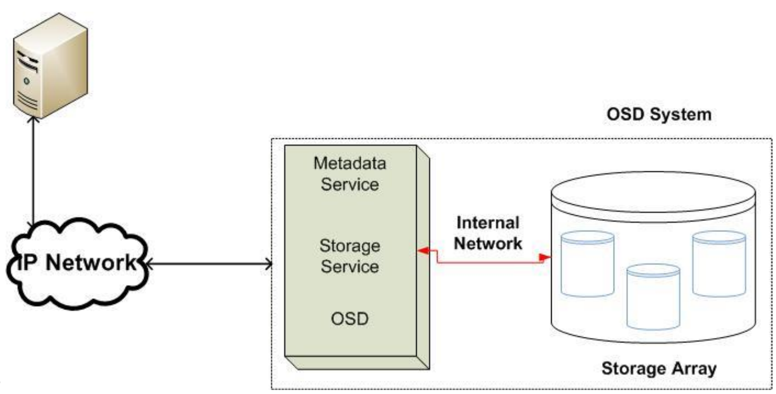

Figure 2. Object-Storage Architecture

Figure 2 shows each object has both data – a sequence of bytes, and metadata – a set of attributes describing the object. Object- Storage device gives scalable, self-management, and shared storage efficiency. . It can handle large amounts of data without impacting performance.

2. Challenges of big data storage and management

With the rate of data explosion, storage systems of organizations and enterprises are facing major challenges from huge quantities of data, and the ever increasing of generated data. Data, irrespective of its size, plays a vital role in the industry. Value can be created from large data set. For example, [8] Facebook increases its ad revenue by mining its users personal preferences and creating profiles, showcasing advertisers which products they are most interested in. Google also uses data from google search, google hangouts, YouTube, and Gmail accounts to profile users' behaviour.

In spite of the numerous benefits that can be gained in large data set, big data demand for storage and processing poses a major challenge. The total size of data that will be generated by the end of 2015 is estimated at 7.9 zettabytes (ZB), and by 2020, is expected to reach 35 ZB. It is clear that big data has outgrown its current infrastructure, and pushes the limit on storage capacity and storage network. Existing traditional techniques cannot support and perform effective analysis, due to the large scale of data.

2.1 Characteristics of Big Data

Big data translate into important variable that characterizes its features: 1) Volume - Large data set from business, scientific research, government, and online social networking are generated daily. The issue of storage capacity and scale, are important variables of consideration. The demand for big data storage, are reaching petabytes and Exabyte. These amounts of data are definitely difficult to handle by the traditional system. 2) Variety - All data generated are totally different. Big data is composed of various forms data such as raw, structured, semi-structured and unstructured data. The complex relationship between this data types cannot be efficiently processed by the traditional. 3) Velocity - Big data is characterized with velocity, which deals with the speed at which data is received and transmitted from various devices. The transmission of data in high real time environments is very critical to processes. The high velocity capture and flow of data cannot be processed by the existing traditional system. 4) Complexity - Data is complex and diverse in types and representation. There has to be a connection and correlation between hierarchies of data. Any change of data can impact lot of data.

2.2 Data Backup and Archiving

Data backup plays an indispensable role in the computing system. Backup is one way to ensure data protection. By keeping copies of production data, backup protects data from potential loss such as hardware and software loss, human errors, and natural disasters. The huge amount of data needing backup and archiving has reached several petabytes [9] and may soon reach tens, or even hundreds of petabytes. The massive amount of data in today's IT environment may consume much storage as well as making data backup inefficient. The deployment of backup systems can affect performance and availability of a system. Backup processes can compete with system processes for resources, such as disk input/output. There are a number of techniques for data backup such as online or offline, full or partial backup, frequencies, and the choice of storage media. In online backup, production process may be going on during backup operations, ensuring consistency in data. In offline backup, production system is shutdown while backing up data. It is the safest way, since it prevents backing up data in the process of updating. Another common approach employed by organizations is a combination of full and partial backup [16] In full backup, a complete system image is taken at once or more between specific time interval and kept on secondary storage device. Partial backup can be categorized as differential and incremental. Incremental backup copies new generated data or data that has changed since the last full or incremental backup. Differential or cumulative backup saves changes in data since last full backup. Incremental backup uses relatively less time in its backup process, but slower restoration process. Full backup process is slower, but provides a faster restoration. For large datasets, the major challenge confronting backups are capacity, availability, and performance. Many suitable backup techniques have been proposed to minimize these challenges while ensuring data protection. De-duplication technique is deployed to identify and remove redundant data content. In [2] deduplication is proposed, where duplicate chucks of data, from stream of data, is identified and eliminated. Since the volume of data can affect backup time, de-duplication help reduce the overall backup time. According to [4], a backup scheduling approach, based on integer programming, is proposed, to minimize the overall backup time. Data backups can be performed in distributed system, where data can be backed up onto multiple storage media in a distributed system. In [7], a distributed backup strategy is proposed to improve backup efficiency. They do so by optimizing network bandwidth between storage systems. Archiving data ensures proper data retention, saves space on storage systems and eases the backup burden.

2.3 Data Replication

A replica set is a group of instances that host the same data set. Replica sets provide redundancy and high availability, and are the basis for all production deployments. Unlike backups, Replicas are readily accessible by production systems. Replica consistency ensures that, buffered data in the host is captured on the disk when the replica is created. Consistency enhances replica usage, and is very pivotal in all replication technologies.

There are many replica consistency model being used by organizations and researchers. In [10] a decentralized copy update consistency model, based on replica identification of timestamp is proposed – which effectively substitutes master copy nodes to slave copy nodes to start update process. Two consistency protocols, aggressive copy coherence, and lazy copy coherence is proposed in [18]. This is based on the master-slave replica model, where the root node represents the master copy. Data is sent to all nodes, once data in the root node updates. In big data environment, because of the huge volumes of data, the aggressive copy coherence won't be able to handle the load pressure. In the lazy copy coherence, by comparing timestamps, the copy nodes can be accessed and new data updated. Though this can reduce network load, in big data environment, there will be high access latency.

2.4 Data Deduplication

Deduplication is being adopted almost everywhere to reduce cost and free storage space in data centres. This technology identifies data blocks with identical content and eliminates redundancy, and thus dramatically reduces the need to store or transfer data in the overall capacity [14].

Deduplication is used in backups and archiving of data, and also employed in primary workloads. In big data environment, a major portion of a large dataset workload is redundant. Applying deduplication to big data storage systems seems desirable, since 1) it will reduce a huge amount of redundant data; 2) also reduces disk IO traffic to enhance performance. In [19], three causes of redundancy in big data workloads are identified. They include, deploying of more nodes, the use of replication mechanism, and expansion of dataset. There are three methods of deduplication from different level – byte, file, and block level. File-level deduplication, also known as Single Instance Storage (SIS), identifies two files with same content but different names or from different directories, and thus deleting the duplicate.

Block-level deduplication reduces stream data into blocks, check the data block with already stored data blocks to determine whether its unique or same data block exist, then removes the redundant copy. This process is made possible through the use of specific hash value, which is generated using a hash algorithm such as MD5, for each data block.

In attempt to achieve a high accuracy and high -precision matching, byte-level deduplication is introduced. It compares stream of data with stored data byte by byte. The five phases commonly referred to in data de-duplication are; 1) Data blocks is divided into blocks of fixed length or variable length. 2) a block identifier is computed for each data block. 3) block identifiers are compared with existing ones to check for unique value. 4) all duplicate data are removed before data integrity check is performed. 5) New data block is stored and previously occupied disk space is released for other usage.

2.5 Business Continuity and Disaster Recovery

Information availability ensures smooth running of business. Organizations are facing an increasing amount of data that is even out-scaling its storage capacity. Since data is fundamental to business processes and operations, it's critical for businesses and organizations to outline measures to protect and ensure data availability in case of any downtime eventualities. Business Continuity (BC) ensures the continuation of business processes, regardless of the circumstance. BC Planning entails preparing for, responding, and recovering from the unexpected such as planned and unplanned outages. The lack of proper BC planning can put people, and organizations out of business. A report from the US labour department, suggest that 40% of the companies facing such disasters go out of business, and 25% of the remaining companies close within two years [13]. BC planning objectives are to protect human lives, minimize financial and reputational losses, continue serving the customers, and remain in compliance with the regulatory laws and service level agreements. Planning and implementing BC factors into account backup and replication. In the face of ever-growing amount of data, the traditional system methodologies of periodic backup remain an issue. Issues such as backup window, downtime production systems during backups, and security of open files, have rendered the traditional backup methods not to meet data recovery point objectives (RPO), and recovery time objectives (RTO). An optimal business continuity solution should have the ability to 1) to maintain different states of data, 2) retrieve an application consistent image data, 3) and to make use of remote replication.

3. Big Data Storage Management

Due to the massive increase, and the heterogeneous nature of application data, one main challenge of big data is effectively, manage the petabyte (PB) of data being generated daily. Storage management encompasses technologies and process organization to improve data storage performance. Big data requires efficient technologies in processing large quantities of data within an acceptable time frame. A wide range of techniques and technologies have been developed and adapted to manipulate, analyze, and visualize big data [12]. Technologies such as massive parallel processing (MPP) database, data mining grids, distributed file system, cloud computing platforms, and scalable storage systems are highly desirable. The deployment of Map-Reduce, with Yahoo's Pig, alongside Facebook's Cassandra applications, has gotten the attention of the industry.

Google file system, GFS, is designed to meet the increasing demands of big data, such as scalability, reliability, and availability. GFS is composed of clusters, which is made up of hundreds of storage servers that support several terabytes of disk space. This meets the scalability issue of big data. Hadoop is a free version of Map-reduce implementation by Apache Foundation [15]. Hadoop distributed file system (HDFS), is a distributed file system designed to run on commodity hardware. HDFS can store data across thousands of servers. All data in HDFS is reduced into block-size chunks, and distributed across different nodes and are managed by the Hadoop cluster.

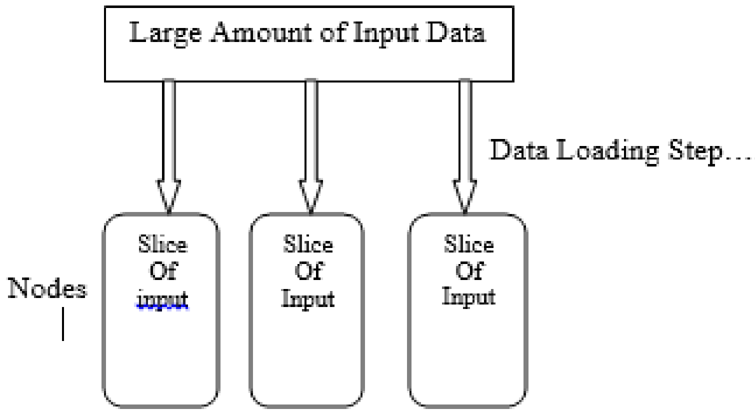

Figure 3. Distributed data across nodes at load time

Figure 3 shows the distribution of data across different data nodes to enhance performance of the entire Hadoop system. Storage vendors such as NetApp, EMC, Hitachi Data Systems, and many more are offering storage management solutions to big data inclined companies. EMC VPLEX enables manageability of storage area network, through a virtual storage infrastructure that consolidate heterogeneous storage devices.

4. Suggestions For A Big Data Storage And Management System

Big data storage and management has been a challenge in the face of the increasing volume of organizational data. Addressing the issue of big data is not trivial. In order to adapt to the massive, multi-source, and complex data, this paper suggest approaches that we think in our opinion efficiently mitigates challenges of big data storage and management.

4.1. Big Data Storage Mediums

All storage management solutions have at their end, storage devices. The qualities of these storage devices can have a significant impact on the entire storage system. The ability of the storage device to scale, its access time, data transfer rate, and cost-effectiveness, can be critical in the big data environment. We suggest the use of a hybrid storage device, which is an aggregate of hard disk drive (HDDs), and solid state drives (SSDs). HDDs provide huge storage capacity, at a relatively cheap price. This characteristic of the HDDs allows storage systems to scale to meet the rate of growth of data. The disadvantage here is that, HDDs has a slow data transfer rate – becoming a bottleneck for performance. On the other hand, SSDs provides avenue for high performance, and reliability. They have low latency, thus providing a much faster, random access. SSDs are very expensive for the storage capacity they provide. The combination of these storage devices into a logical unit in an array, should solve the storage demands posed by large datasets.

4.2. Backup Strategies

Recovery is the main objective for backup. The ability of production systems to recover, and in a timely manner is very crucial in the era of big data. From this perspective, this paper recommend full back up as a favourable choice over the others. Full backup ensures speedy recovery, though it takes a considerable amount of time to backup a large dataset. Applying full backup to large datasets may increase the rate of data block repetition. Data deduplication technology, significantly reduces the volume of stored data blocks for every single full backup, and allow users to backup, and recover data within a relatively short period of time. In an efficient storage system, backups are usually done from a replication system, rather than directly from the production system. Replication keeps copy of production data in real time.

New generated data between separate full backups, can be recompensed from the replication system, in case of downtime

4.3. Business Continuity and Disaster Recovery

An optimal business continuity solution, takes into account, two parameters, to a negligible level – Recovery Point Objective, (RPO), which is the point in time that a production system, and data must be recovered after a disaster. Recovery Time Objective (RTO), is the time frame within which production system, and data must be recovered after a disaster. In a large dataset, the complexity of business continuity increases, with the influx of a variety of data, which must be maintained in their formats. Business continuity planning requires, saving a copy, or multiple copies of production data, through backups, local or remote replication. The use of enterprise software such as EMC Powerpath can be beneficial. EMC Powerpath provides features such as cluster support, dynamic load balancing, configuration and management, automatic path failover.

5. Conclusion

Big data era has brought about an explosive growth of data. The increase of mobile applications, social media, and big data analytic initiatives has caused big data storage challenges to become even greater. Choosing the right storage devices, management tool, and efficient techniques is relevant and determines the rate of growth. The approach to big data storage and management can significantly, affect an entire organization. Organizations and businesses are now more concerned about how to efficiently keep and retain all their data. This paper examines and summarizes existing current storage technologies for big data applications. Variables such as capacity, scalability, data transfer rate, access time, and cost of storage devices, are also highlighted. Finally some suggestions are made to curb the problems posed by big data storage.

References

[1] Bouganim, L., Jónsson, B., & Bonnet, P. (2009). uFLIP: Understanding flash IO patterns. *arXiv preprint arXiv:0909.1780*.

[2] Chervenak, A., Vellanki, V., & Kurmas, Z. (1998). *Protecting file systems: A survey of backup techniques.* Paper presented at the Joint NASA and IEEE Mass Storage Conference.

[3] Fontana, R. E., Hetzler, S. R., & Decad, G. (2012). Technology Roadmap Comparisons for TAPE, HDD, and NAND Flash: Implications for Data Storage Applications. *Magnetics, IEEE Transactions on, 48*(5), 1692-1696. doi: 10.1109/TMAG.2011.2171675

[4] Geer, D. (2008). Reducing the storage burden via data deduplication. *Computer, 41*(12), 15-17.

[5] He, X., & Zhao, L. (2013). *A Data Management and Analysis System in Healthcare Cloud.* Paper presented at the Service Sciences (ICSS), 2013 International Conference on.

[6] Ji, C., Li, Y., Qiu, W., Awada, U., & Li, K. (2012). *Big data processing in cloud computing environments.* Paper presented at the Proceedings of the International Symposium on Parallel Architectures, Algorithms and Networks, I-SPAN.

[7] Nakamura, S., Nakayama, K., & Nakagawa, T. (2009). *Optimal backup interval of database by incremental backup method.* Paper presented at the Industrial Engineering and Engineering Management, 2009. IEEM 2009. IEEE International Conference on.

[8] Policroniades, C., & Pratt, I. (2004). *Alternatives for Detecting Redundancy in Storage Systems Data.* Paper presented at the USENIX Annual Technical Conference, General Track.

[9] Reed, D. A., Gannon, D. B., & Larus, J. R. (2011). Imagining the future: Thoughts on computing. *Computer*(1), 25-30.

[10] Renuga, K., Tan, S., Zhu, Y., Low, T., & Wang, Y. (2009). *Balanced and efficient data placement and replication strategy for distributed backup storage systems.* Paper presented at the Computational Science and Engineering, 2009. CSE'09. International Conference on.

[11] Rui-Xia, Y., & Bo, Y. (2012). *Study of NAS Secure System Base on IBE.* Paper presented at the Industrial Control and Electronics Engineering (ICICEE), 2012 International Conference on.

[12] Singhal, R., Bokare, S., & Pawar, P. (2010). *Enterprise storage architecture for optimal business continuity.* Paper presented at the Data Storage and Data Engineering (DSDE), 2010 International Conference on.

[13] Sun, G.-Z., Dong, Y., Chen, D.-W., & Wei, J. (2010). *Data backup and recovery based on data de-duplication.* Paper presented at the Proceedings of the 2010 International Conference on Artificial Intelligence and Computational Intelligence-Volume 02.

[14] Sun, Y., & Xu, Z. (2004). *Grid replication coherence protocol.* Paper presented at the Parallel and Distributed Processing Symposium, 2004. Proceedings. 18th International.

[15] White, T. (2012). *Hadoop: The definitive guide*: " O'Reilly Media, Inc.".

[16] You, L. L., Pollack, K. T., & Long, D. D. (2005). *Deep Store: An archival storage system architecture.* Paper presented at the Data Engineering, 2005. ICDE 2005. Proceedings. 21st International Conference on.

[17] YouTube Data Statistics. (2015). Retrieved 01-15-2015, 2015, from http://www.youtube.com/yt/press/statistics.html

[18] Zheng, S., Li, M.-C., & Sun, W.-F. (2011). DRCSM: a Novel Decentralized Replica Consistency Service Model. *Journal of Chinese Computer Systems, 32*(8), 1622-1627.

[19] Zhou, R., Liu, M., & Li, T. (2013). *Characterizing the efficiency of data deduplication for big data storage management.* Paper presented at the Workload Characterization (IISWC), 2013 IEEE International Symposium on.