Efficient Resource Mapping by means of Storage Tiering using RAID in Cloud

P.Sarala
Associate Professor, Department of Computer Science And Engineering,
Nandha Engineering College, Erode, Tamil Nadu, India.

Dr N.Shanthi
Professor and Dean, Department of Computer Science And Engineering,
Nandha Engineering College, Erode, Tamil Nadu, India.

Abstract - Recent Technologies that reach end users are those which provide services for all their needs. Cloud computing is a technology where service is provided for all end user's needs and requests. As we move deeper into an era in which data usage and the storage capacity are increased double year by year, some of the major companies are mainly concentrated on demand storage option like cloud storage. Cloud Storage enables users to remotely store their data and enjoy the on-demand high quality cloud applications without the burden of local hardware and software management. Cloud Storage Systems commonly use replication of stored data sets to ensure high reliability and availability. However, the high storage overhead of replication becomes increasingly unacceptable with the explosive growth of data stored in cloud. The objective of this work is to afford Imaginative cloud storage with well-designed Cloud RAID mechanism which should achieve the right tradeoffs between storage efficiency, performance, and reliability. As there exists no widely-accepted methods for Cloud RAID, we present a new breed of storage solutions, which are featured – rich RAID arrays that provide highly optimized I/O processing capabilities. These storage systems are configured with a large amount of memory (Cache) and multiple I/O paths and use sophisticated algorithms to meet the performance sensitive applications.

Index Terms – Cloud Computing, Storage Tiering, Thin Provisioning, RAID, LUN Masking, Sub LUN Tiering, Storage Performance, Hybrid Clustering Algorithm.

I. INTRODUCTION

Data created by individuals or businesses must be stored so that it is easily accessible for further processing. Devices designed for storing data are termed as Storage Devices or simply Storage. The type of storage used varies based on the type of data and the rate at which it is created and used. A storage device uses magnetic, optic or solid state media.

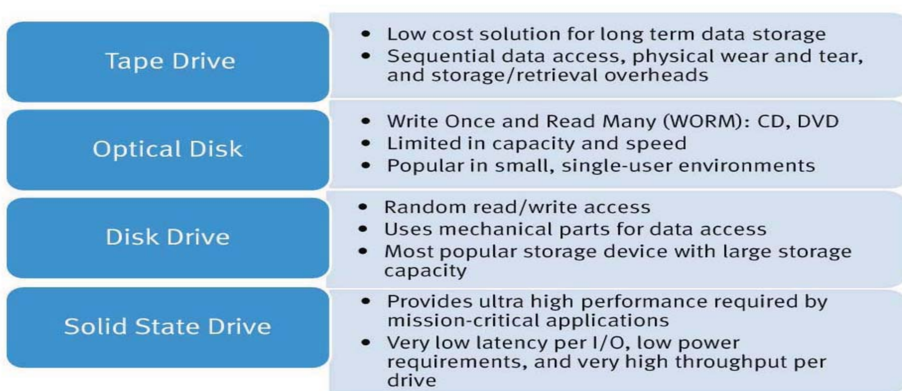

Fig. 1 Storage Device Options

Business-critical applications require high levels of performance, availability, security, and scalability. A disk drive is a core element of storage that governs the performance of any storage system. Some of the older disk array technologies could not overcome performance constraints due to the limitations of a disk drive and its mechanical components. RAID technology made an important contribution to enhance storage performance and reliability, but disk drives, even with a RAID implementation could not meet performance requirements of today's applications.

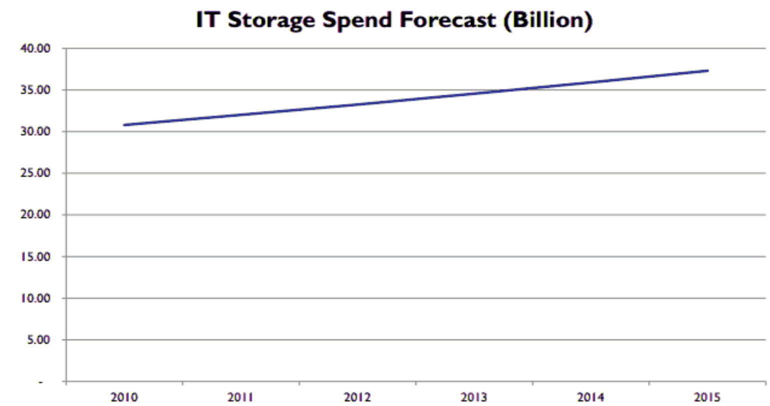

Fig. 2 Storage cost sensex

With advancements in technology, a new breed of storage solutions, known as an Intelligent Storage System, has evolved. These intelligent storage systems are feature-rich RAID arrays that provide highly optimized I/O processing capabilities. These storage systems are configured with a large amount of memory (called cache) and multiple I/O paths and use sophisticated algorithms to meet the requirements of performance-sensitive applications. These arrays have an operating environment that intelligently and optimally handles the management, allocation and utilization of storage resources.

II.PROPOSED CLOUD ARCHITECTURE

A. RAID Technology and Intelligent Storage System

Disk drives are susceptible to failures due to mechanical wear and tear and other environmental factors. A disk drive failure could result in data loss. A disk drive has a projected life expectancy before it fails. Today, data centers house thousands of disk drives in their storage infrastructures. Greater the number of disk drives in a storage array, greater the probability of a disk failure in the array. RAID technology was developed to mitigate this problem. RAID is an enabling technology that leverages multiple disks as part of a set. The technology also provides data protection against drive failures.

B.RAID Techniques and their Levels

In general, RAID implementations also improve storage system performance by serving I/Os from multiple disks simultaneously. Storage systems with flash drives are also benefited in terms of protection and performance by implementing RAID. Striping, Mirroring, and Parity are RAID techniques that form the basis for defining various RAID levels. These techniques determine the data availability and performance of a RAID set. RAID controller helps implementing these RAID techniques.

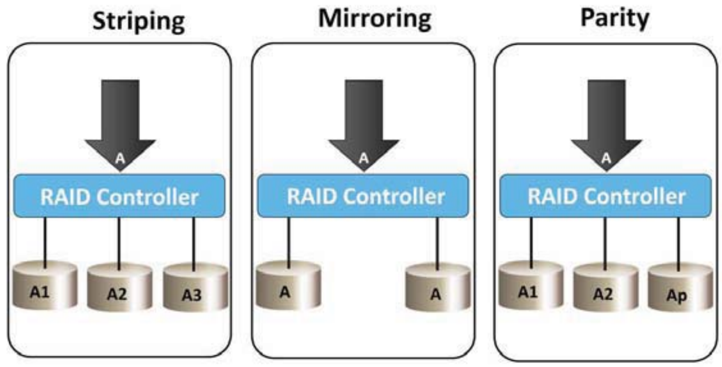

Fig. 3 RAID Techniques

Application performance and data availability requirements determine the RAID level selection. These RAID levels are defined on the basis of the RAID technique(s) implemented. Some RAID levels use one technique, whereas others use a combination of techniques.

Table 1 RAID Levels and their Functionalities

RAID Levels	Definition
RAID 0	Striping with no fault tolerance
RAID 1	Disk mirroring
Nested	Combinations of RAID levels; Example: RAID 1 + RAID 0
RAID 3	Parity RAID with dedicated parity disk
RAID 5	Parity RAID with distributed parity across all the disks in the set
RAID 6	Distributed parity RAID with dual parity

C. Components of Intelligent Storage System

In the Proposed System, we have introduced an intelligent storage system which consists of four key components: Front-end, Cache, Back-end, and Physical Disks. The figure on the slide illustrates these components and their interconnections. An I/O request received from the compute system at the front-end port is processed through cache and the back end to enable storage and retrieval of data from the physical disk. A read request can be serviced directly from cache if the requested data is found in cache.

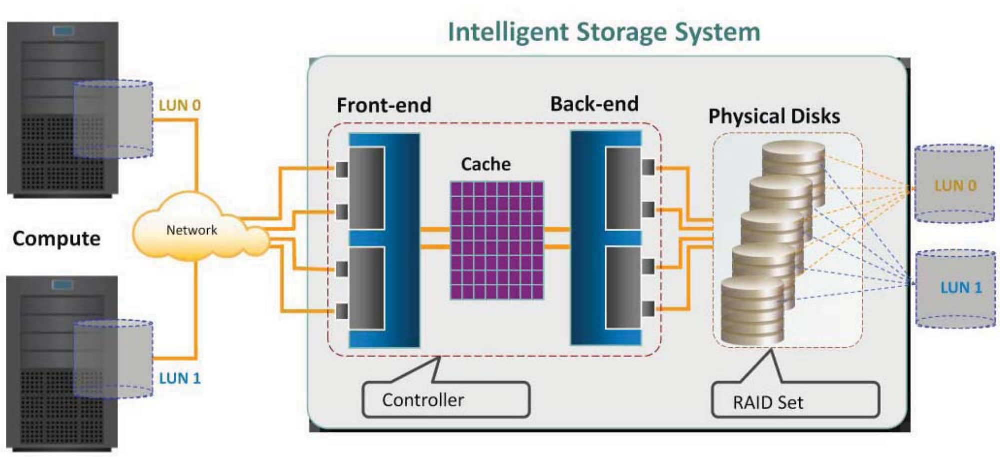

Fig. 4 Cloud Storage System Architecture

In an intelligent storage system, physical disks are logically grouped together to form a set, called RAID set, on which a required RAID level is applied. The number of drives in the RAID set and the RAID level determine the availability, capacity, and performance of the RAID set. It is highly recommended that the RAID set be created from the same type, speed, and capacity drives to ensure the maximum usable capacity, reliability, and consistent performance. For example, if drives of different capacities are mixed in a RAID set, then the capacity of the smallest drive will be used from each disk in the set to make up for the RAID set's overall capacity. The remaining capacity of the larger drives will remain unused. Likewise, mixing higher Revolutions per minute (RPM) drives with lower RPM drives lowers the overall RAID set's performance.

RAID sets usually have large capacities because they combine the total capacity of individual drives in the set. Logical Units are created from the RAID sets by partitioning (seen as slices of RAID set) the available capacity into smaller units. These units are then assigned to the compute system for their storage requirements.

Logical units are spread across all the physical disks, which belong to that set. Each logical unit created from the RAID set is assigned a unique ID called Logical Unit Number (LUN). LUNs hide the organization and composition of RAID set from the compute systems. The diagram on the slide shows a RAID set consisting of five disks that have been sliced or partitioned into two LUNs: LUN 0 and LUN 1. These LUNs are then assigned to compute system 1 and compute system 2 for their storage requirements.

It is also possible to control access of LUNs by a compute system. This is done with the help of "LUN masking". LUN masking is a process that provides data access control by defining which LUNs a compute system can access. LUN masking function is implemented on the storage processor/controller. This ensures that the volume access by servers is controlled appropriately, preventing unauthorized or accidental use in a shared environment.

III.SYSTEM MODEL AND DESIGN GOALS

A. Virtual Provisioning

One of the biggest challenges for storage administrators is balancing the storage space required by various applications in their data centers. Administrators typically allocate storage space based on anticipated storage growth. They do this to reduce the management overhead and application downtime required to add new storage later on. This generally results in the over-provisioning of storage capacity, which leads to higher costs, increased power, cooling, and floor space requirements, and lower capacity utilization. In our work, these challenges are addressed by Virtual Provisioning.

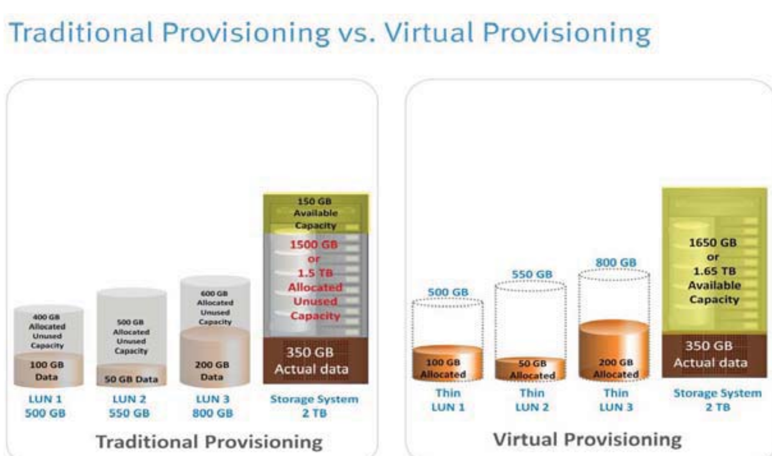

Fig. 5 Virtual Provisioning compared with Thin Provisioning

Virtual Provisioning is the ability to present a logical unit (Thin LUN) to a compute system, with more capacity than what is physically allocated to the LUN on the storage array. Physical storage is allocated to the application "on-demand" from a shared pool of physical capacity. This provides more efficient utilization of storage by reducing the amount of allocated, but unused physical storage.

B. Thin LUN

Thin LUNs are logical devices. Physical storage need not be completely allocated to them at the time of creation. Physical storage is allocated to the Thin LUNs from the Thin pool. 'Thin LUN extent' is the minimum amount of physical storage that is consumed at a time by a 'Thin LUN' from a 'Thin pool'. From the operating system's perspective, Thin LUNs appear as traditional LUNs. Thin LUNs are best suited for situations where space efficiency is paramount. They are used for applications when the storage space consumption is difficult to predict.

C. Thin Pool

A Thin pool comprises physical drives that provide the actual physical storage used by Thin LUNs. A Thin pool is created by specifying a set of drives and a RAID type for that pool. Thin LUNs are then created out of that pool (similar to traditional LUN created on a RAID group). All the Thin LUNs created from a pool share the storage resources of that pool. Multiple pools can be created within a storage array. Adding drives to a Thin pool increases the available shared capacity for all the Thin LUNs in the pool. Drives can be added to a Thin pool while the pool is used in production. The allocated capacity is reclaimed by the pool when Thin LUNs are destroyed.

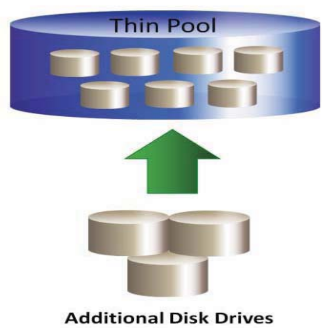

Fig. 6 Block Diagram of Thin Pool

D. Automated Storage Tiering

Organizations are experiencing tremendous data growth, which increases their storage requirements. They are also required to meet regulatory requirements. The cost of storing data and meeting SLA is a serious concern. Buying more high-end storage is not a cost-efficient solution for the growing data storage needs. Organizations require solutions that enable storing the right data, at the right cost, with the right access.

In our paper , Storage tiering has emerged as a means to address these challenges. It is an approach to establish a hierarchy of storage types. It helps identify active or inactive data to relocate them to an appropriate storage type. This enables in meeting service level requirements at an optimal cost. Each tier has different levels of protection, performance, data access frequency, and other considerations. For example, high performance FC drives may be configured as tier 1 storage to keep frequently accessed data to improve performance and low cost SATA drives as tier 2 storage to keep the less frequently accessed data. Moving the active data (frequently used data) to Solid-state drive (SSD) or FC improves the application performance. Moving the inactive data (less frequently used) to SATA can free up storage capacity in high performance drives and reduce the cost of storage. This movement of data happens based on defined tiering policies. The tiering policy may be based on parameters such as file type, frequency of access, performance, etc. For example, if a policy states "move the files which are not accessed for last 30 days to lower tier", then the files matching this condition are moved to the lower tier.

There are two types of Storage tiering: manual storage tiering and automated storage tiering. The manual storage tiering is the traditional method where the storage administrator has to monitor the storage workloads periodically and move the data between the tiers. A traditional storage tiering process is manual, repetitive, and takes few hours to few days to complete. Automated storage tiering automates the storage tiering process. Data movement between tiers is performed non-disruptively, without affecting business continuity. Automated storage tiering eliminates manual tiering when the application workload characteristic changes over time. It improves application performance at the same cost or provides the same application performance at a lower cost. Data movements between tiers can happen within or between storage arrays.

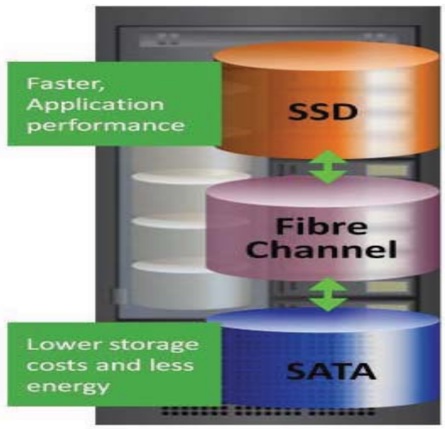

Fig. 7 Automated Storage Tiering – Intra Array

E. Automated Storage Tiering Building Blocks

There are three major building blocks of automated storage tiering: Storage Type is a combination of a drive technology (SSD, FC, or SATA) and a RAID protection type. This slide displays a simple example with three storage types – one for SSD, one for FC, and one for SATA. Storage Groups are a logical collection of LUNs that are to be managed together. In this example, three storage groups represent three different business applications. Typically, each application will have its own service level requirements.

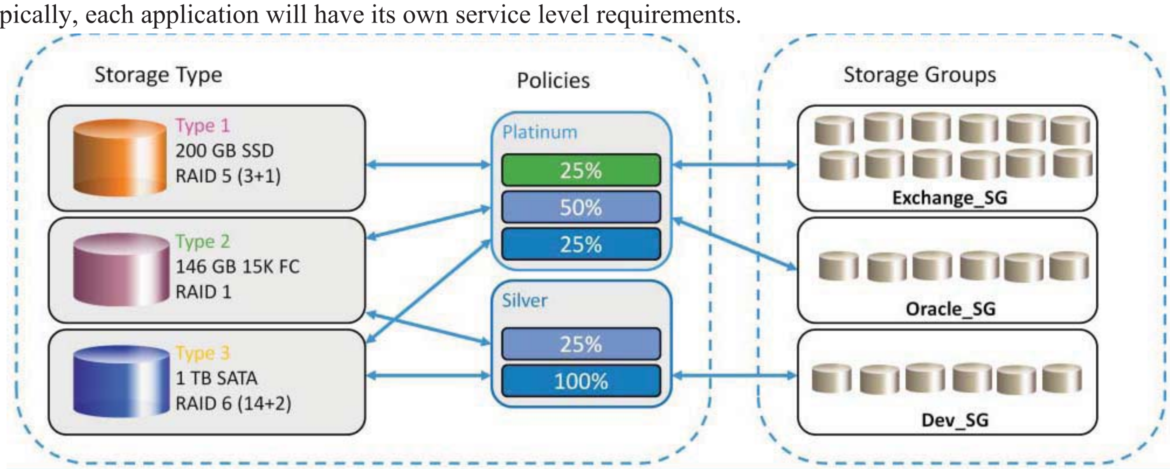

Fig. 8 Storage Types and their Policies

Policies manage data movement across storage types. The data movement is done by specifying a set of tier usage rules. A policy may be applied to one or more previously defined storage groups. In our example, the Platinum policy implies that up to 25% of the Storage Group's capacity is allowed to reside on SSD, up to 50% on Fibre Channel, and up to 25% on SATA. Note that the percentages for tier usage within a given policy need not add to 100%. For example, the "Silver" policy implies that all of the Storage Group's capacity (i.e. up to 100% of its

capacity) may reside on SATA. Only the Fibre Channel capacity is restricted to 25% of the total capacity allocated to this Storage Group. This gives the flexibility to move the volumes included in these Storage Groups entirely to SATA, if that can produce optimal
results.

F. Sub-LUN tiering

Traditional storage tiering moves an entire LUN from one tier of storage to another. This movement includes both active and inactive data in that LUN. This method does not give effective cost/performance benefits. In sub-LUN tiering, a LUN is broken down into smaller segments and tiered at that level. Movement of data with much finer granularity (ex., 8MB) greatly enhances the value proposition of automated storage tiering. Tiering at the sub-LUN level moves more active data to faster drives and less active data to slower drives effectively.

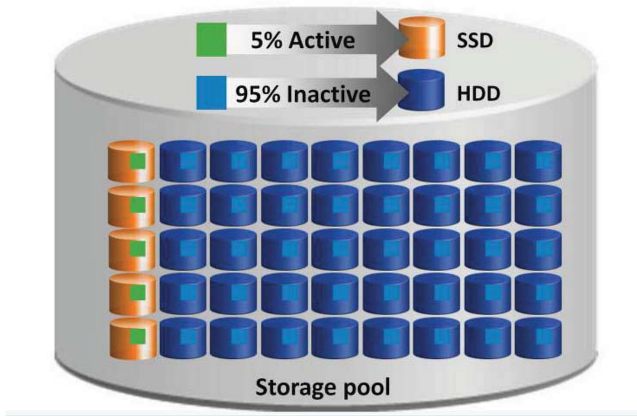

Fig. 9 Tiering at the Sub LUN levels.

IV.OPTIMAL SOLUTION AND CLUSTERING ALGORITHM FOR STORAGE TIERING

The problem of clustering has been studied widely in the database and statistics literature in the context of a wide variety of data mining tasks [29, 30]. The clustering problem is defined to be that of finding groups of similar objects in the data. The similarity between the objects is measured with the use of a similarity function. The problem of clustering can be very useful in the text domain, where the objects to be clusters can be of different granularities such as documents, paragraphs, sentences or terms. Clustering is especially useful for organizing documents to improve retrieval and support browsing [31, 32]. Traditional methods for clustering have generally focused on the case of quantitative data [23, 24, 25, 30, 26], in which the attributes of the data are numeric. The problem has also been studied for the case of categorical data [27, 28, 29], in which the attributes may take on nominal values. A broad overview of clustering (as it relates to generic numerical and categorical data) may be found in [29, 30].

Hierarchical clustering algorithms have been studied extensively in the clustering literature [29, 30] for records of different kinds including multidimensional numerical data, categorical data and text data.
In particular, the effectiveness of this method in improving the search efficiency over a sequential scan has been shown in [21, 22]. While hierarchical clustering methods tend to be more robust because of their tendency to compare all pairs of documents, they are generally not very efficient, because of their tendency to require at least O (n2) time. On the other hand, k-means algorithms which start off with a set of k seeds from the original corpus and assign documents to these seeds on the basis of closest similarity. Normally k-means algorithms are more efficient than hierarchical algorithms, but may sometimes not be very effective because of their tendency to rely on a small number of seeds.

Hence this work used both hierarchical and k-means clustering algorithms to good effect. Specifically, it uses a hierarchical clustering algorithm on a sample of the corpus in order to find a robust initial set of seeds. This robust set of seeds is used in conjunction with a standard k-means clustering algorithm in order to determine good clusters. The size of the sample in the initial phase is carefully tailored so as to provide the best possible effectiveness without this phase becoming a bottleneck in algorithm execution.

V.PERFORMANCE EVALUATION

A.Service Time taken by various Storage Devices

Table 3-Time taken by SATA,FC and SSD

Drive Type	Protocol	Seek	Latency	Total
SATA	1?	9	4.1	14
10k RPM Fibre	0.3?	4.7	3	8
15k RPM Fibre	0.2?	3.6	2	5.8
SSD	0.2?	0	0	0.2

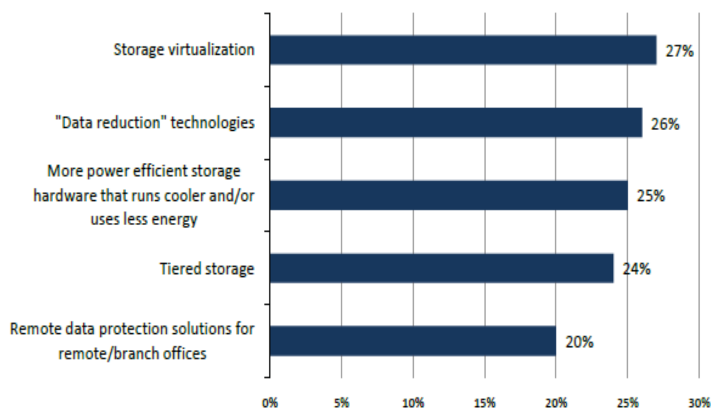

Fig. 10 Measurement of performance on implementing Storage Tiering in Cloud Storage System.

B.Thin Provisioning Capacity

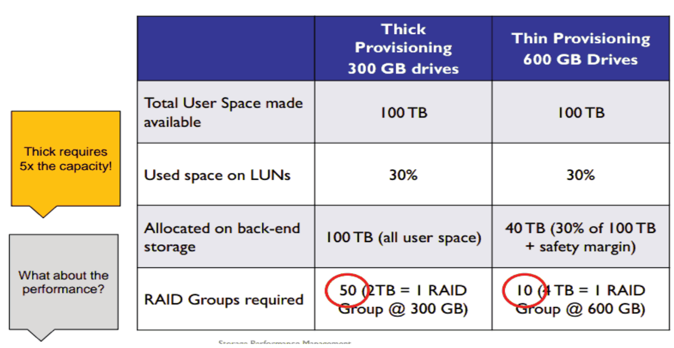

	Thick Provisioning 300 GB drives	Thin Provisioning 600 GB Drives
Total User Space made available	100 TB	100 TB
Used space on LUNs	30%	30%
Allocated on back-end storage	100 TB (all user space)	40 TB (30% of 100 TB + safety margin)
RAID Groups required	50 (2TB = 1 RAID Group @ 300 GB)	10 (4 TB = 1 RAID Group @ 600 GB)

VI.CONCLUSION

Businesses are always looking to gain an edge by ensuring the most demanding applications have the highest level of performance. Traditional Storage System has not been able to keep pace with performance demands, but with the introduction of Intelligent Storage System, organizations can now achieve significantly higher levels. And this is not just for high-performance computing (HPC) environments—organizations would like to apply the benefits to many different applications.

In this paper we have proposed an effective Storage tiering strategy which positively impact the bottom line by only utilizing power hungry, high performing disks where they are required and energy efficient SATA for the majority of data, reducing power, cooling, and floor space requirements. Enterprise Flash Drives provide the lowest $/IOPS while SATA drives provide the lowest $/GB and are both very space and energy efficient usage rate from dropping down should be considered. We apply proposed method to the actual environment by resolving these challenges.

REFERENCES

[1] Angel Cuevas , Manuel Uruen , Gustavo de Veciana , Rube´n Cuevas , Noe¨l Crespi, "Dynamic Data-Centric Storage for long-term storage
[2] in Wireless Sensor and Actor Networks," Wireless Network, DOI 10.1007/s11276-013-0598-5 © Springer Science+Business Media New
[3] York 2013.
[4] Dong Yuan,Yun Yang,IEEE,Xiao Liu,IEEE,Wenhao Li,Lizhen Cui,Meng Xu,and jinjun Chen,IEEE, "A Highly Practical Approach toward
[5] Achieving Minimum Data Sets Storage Cost in the Cloud", IEEE Transaction On Parallel and Distributed Systems,Vol 24,No 6,June
[6] 2013.
[7] Xuefeng Liu, Yuqing Zhang, Member, IEEE, Boyang Wang, and Jingbo Yan," Mona: Secure Multi-Owner Data Sharing for Dynamic
[8] Groups in the Cloud," in IEEE Transactions on Parallel And Distributed Systems, Vol. 24, No. 6, June 2013.
[9] Hyuck Han,Young Choon Lee ,IEEE,Woong Shin ,Hyungsoo Jung, Heon Y,Yeom ,IEEE,and Albert Y,Zomaya ,Fellow,IEEE,"Cashing in
[10] on the Cache in the cloud", IEEE Transaction On Parallel and Distributed Systems,Vol 23,No 8,August 2012.
[11] Alshawabkeh M., Riska A. , Sahin A. , Awwad M,"Automated Storage Tiering Using Markov Chain Correlation Based Clustering ",
[12] Machine Learning and Applications (ICMLA), 2012 11th International Conference on Volume: 1 , 2012.
[13] Rajkumar Buyyaa,b, Chee Shin Yeoa, Srikumar Venugopala, James Broberg a, Ivona Brandic c, "Cloud computing and emerging IT
[14] platforms: Vision, hype, and reality for delivering computing as the 5th utility" , Future Generation Computer Systems, Elsevier 2009.
[15] *Bob* Laliberte*, "*Automate and Optimize a Tiered Storage Environment – FAST!", December**,** 2009, Enterprise Strategy Group,EMC2.
[16] Vidhya Suryanarayana V. ,Mylar Balasubramanya K. , Pendse R.,"Cache isolation and thin provisioning of hypervisor caches", Local
[17] Computer Networks (LCN) IEEE 37th Conference , 2012
[18] Jianzong Wang , Weijiao Gong , Varman P. ,Changsheng Xie , "Reducing Storage Overhead with Small Write Bottleneck Avoiding in
[19] Cloud RAID System ", Grid Computing (GRID), ACM/IEEE 13th International Conference , 2012
[20] Emaru, H. , Takai, N.K.,"Performance Management for the Dynamic Storage Tiering by Virtual Volume Clustering ", Utility and Cloud
[21] Computing (UCC), Fourth IEEE International Conference 2011
[22] Jiguang Wan ,Jibin Wang , Jianzong Wang , Zhihu Tan , Maliang Liu, "RSA: RAID system with self- healing and active data migration
[23] ", Intelligent Computing and Intelligent Systems (ICIS),International Conference,Volume:3,2010
[24] Jun Liu , Baojiang Cui ,"Performance Prediction of Distributed RAID Storage System ", Information Science and Engineering (ICISE),
[25] 1st International Conference,2009
[26] Cong Wang, "Toward Secure and Dependable Storage Services in Cloud Computing", IEEE trans on Services Computing Vol. 5 No. 2,
[27] Jun 2012.
[28] Qian Wang and Cong Wang "Enabling Public Auditability and Data Dynamics for Storage Security in Cloud Computing", IEEE Tran Vol.22 No.5, May 2011.
[29] Shuai Han ,"Ensuring data storage security through a novel third party auditor scheme in cloud computing" , IEEE conference on Cloud
[30] Computing and Intelligence Systems pp. 264 - 268, Sep 2011.
[31] Cloud Bridge ,White Paper ," Four Key Elements of a Secure, High-Performance Hybrid Cloud Networking Strategy," in citrix.com
[32] Jiwu Shu,Bigang Li, and Weimin Zheng ,"Design and Implementation of an SAN System Based on the Fiber Channel Protocol",IEEE
[33] Transaction on Computers, Vol 54,No 4,April 2005
[34] Eric A. Hibbard, CISSP, CISA, "STORAGE SECURITY - THE ISO/IEC STANDARD", Storage Networking Industry Association, 2012.
[35] Brett Allison, IntelliMagic," Storage Performance Management Overview ", Storage Networking Industry Association, 2012.
[36] Anil Vasudeva, President & Chief Analyst, IMEX Research, "NextGen Infrastructure for Big Data ", Storage Networking Industry
[37] Association, 2012.
[38] Philippe Nicholas, "Massively Scalable File Storage ", Storage Networking Industry Association, 2013.
[39] Roger Cummings, Antesighanus, "Combining SNIA Cloud, Tape and Container Format Technologies for the Long Term Retention of Big
[40] Data", Storage Networking Industry Association, 2013.
[41] Chad Thibodeau, Cleversafe, Sebastian Zangaro, "Cloud Archive & Long Term Preservation Challenges and Best Practices ", Storage

[42] Networking Industry Association, 2012.
[43] Mark S Fleming, IBM, "Understanding High Availability in the SAN ", Storage Networking Industry Association, 2012.
[44] Mike Fishman, "Trends in Data Protection and Restoration Technologies ", Storage Networking Industry Association, 2012.
[45] Thomas Rivera, "Advanced Data Reduction Concepts ", Storage Networking Industry Association, 2013.
[46] A. Jain, R. C. Dubes," Algorithms for Clustering Data", Prentice Hall, Englewood Cliffs, NJ, 1998.
[47] L. Kaufman, P. J. Rousseeuw," Finding Groups in Data: An Introduction to Cluster Analysis", Wiley Interscience, 1990.
[48] [31] P. Anick, S. Vaithyanathan,"Exploiting Clustering and Phrases for Context-Based Information Retrieval",
[49] *ACM SIGIR Conference*, 1997.
[50] [32] D. Cutting, D. Karger, J. Pederson,"Constant Interaction-time Scatter/ Gather Browsing of Large Document Collections", ACM
[51] SIGIR Conference, 1993.