Advancements in Tape outpace Disk Technology

Isabel Schwerdtfeger
System Services, Global Technology Services
IBM Germany
Hamburg, Germany
email: schwerdtfeger@de.ibm.com

***Abstract*—Tape technologies are able to effectively address many new data intensive market opportunities. Tape has traditionally been a primary backup device, but it no longer does that alone. Recent tape advances have made greater progress than disk technologies over the past 10 years. Today's modern tape technology is able to efficiently store huge amounts of data in a cost and efficient way. This paper examines the hypothesis that tape shall be viewed as an active or "near-line" storage component, essential in any petabyte (PB) or exabyte (EB) storage environment.**

Keywords-Archive; Capacity; Costs; Disk; High Performance Computing (HPC); High Performance Storage System (HPSS); Linear Tape Open (LTO); Performance; Tape; Total Cost of Ownership (TCO).

I. INTRODUCTION

By 2019, 75 percent of organizations will treat archived data as an active and "near line" data source, and not simply as a separate repository to be viewed or searched periodically, up from less than 10 percent today [1]. Therefore, storage systems continue to play a central role in strategic planning in organizations despite effective optimization tools for storage software, tiered storage architectures, and effective data management policies in place. Back in 2006, the worldwide storage market including the total disk and tape storage hardware segment increased to $28.2 billion [2]. At that time, IBM was the #1 seller in storage technology, exactly 50 years after the first IBM System RAMAC 350 storage product was launched in 1956 with a maximal capacity of 5 megabytes (MB) (for 5 million characters) [3].

Despite the enormous data growth in the last ten years, the storage market has come down to a fragmented market view, where today the measurements are based solely on Big Data, Virtualization, Enterprise Information Archives, or just by Tape Drive level. The Big Data market is expected to have a compound annual growth rate of 28.5 percent that will reach $5.89 billion in 2018 in disk storage systems [4]. That is a pure fraction of what has been shipped and installed into disk systems since 2006 and it shows how sharp the costs for disk storage systems have come down, despite the huge data growth needs spurred on by digital-natives, smartphones, and supercomputers up to this day.

Already in 2006, and long before it was widely known, tape technology was already "green", to conserve energy, power, and cooling costs. Tape is the cheapest storage for storing large amount of data. Tape cartridges only consume power when read or written [3]. In a recent memo by the Tape Storage Council (TSC) the current trends, usages, and technology innovations occurring within the tape storage industry were described as cartridge capacity increase, much longer media life, and improved bit error rates [5]. The publication of this memo by leading vendors suggests that today's modern tape technology is nothing like the tape of the past.

For tape technologies two main types exists: Linear Tape Open (LTO) and Enterprise. LTO is a joint specification for ½ inch magnetic tape and its corresponding tape drives. IBM, HP and Seagate jointly worked on this development since the late 1990s. Due to the sale of Seagate to Quantum, Quantum is now part of the LTO consortium. The LTO consortium issues in regular intervals the LTO Ultrium Roadmap. On September 10th, 2014, the newest "Generation 10" Roadmap was released. The newest LTO Generation 10 will be able to store up to 120 Terabyte (TB) on a single tape cartridge and a read/write performance up to 2750 megabyte per second (MB/s) [6]. As for Enterprise tape, that is IBM 3592, a series of tape drives and corresponding magnetic tape data storage media formats developed by IBM.

Significant technology innovations in tape technologies address the constant demand for improved reliability, higher capacity, power efficiency, ease of use and the lowest cost per gigabytes (GB) of any storage solution [5]. A management of PB or multi-PB disk infrastructures is very difficult to build and maintain as things break much more frequently [7]. Furthermore, high demands for upcoming EB storage systems, and the need for efficient data management solutions were already addressed by the U. S. Department of Energies (DOEs) in 2009 and summarized in a report to the U.S. Government. Alone one PB of storage is roughly equivalent of 210,000 DVDs [8]. Fig. 1 shows the exponential data growth by the DOEs of archived data projected until 2022. It shows that the DOEs will have single archive systems that retain between 2 and 50 EB of data.

Numerous sources, including analysts, consultants, IBM and Oracle, predict in a 10-year projection for technology costs per terabytes a steady decline in price for flash, disks and tape from 2014 until 2023. The compound growth rate

 ISBN: 978-1-61208-416-9

(CGR) for flash is predicted to reduce by -30%, for disk by -15% and for tape -23% [9].

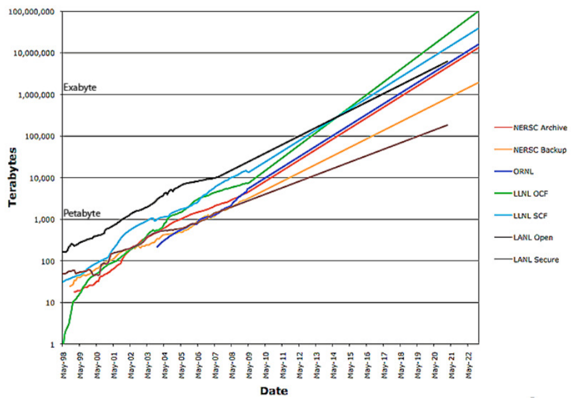

Figure 1. Archived Data stored by the DOE Lab through 2018-2022 [13]

While daily data grows exponentially, there will be an increasing gap, if this data should be managed on disks systems alone. In three years two and a half more disks are necessary to cope with storage needs. Today a disk backup with 600TB with about 300 disks at 7,5 kilowatt (kw) will be in three years 1,8 PB with ca. 750 disks at 18,7kw [10]. Together including the recent tape technology advancements, this gap could be better addressed with tape instead.

The introduction in Section I illustrates the dynamics of huge data growth and describes market and trends of storage systems. In addition, basic tape technologies are being described. Section II lists the recent developments of tape advancements, while Section III compares tape benefits to disk technologies including the aspect of Total Cost of Ownership (TCO). Section IV describes an example of an active archive by its benefit of using tape technologies with use of optimized software. Big Data software, such as the High Performance Storage System (HPSS) used by scientists and researchers today, supports the growing data needs as illustrated in Section I. This paper closes in Section V with a short summary of the examination and conclusion.

II. LIST OF RECENT DELVEOPMENTS OF TAPE TECHNOLOGIES

In 2013 and 2014, at least eight new developments and products for the tape business were announced and available for the market. The memo of the TSC lists the recent developments in chronological order [5]:

- Sept. 16, 2013 Oracle announced a new enterprise tape drive StorageTek T10000D with 8,5TB native capacity and a data rate of 252 MB/s native.
- Jan. 16, 2014 Fujifilm Recording Media U.S.A. reported over 100 million manufactured LTO Ultrium data cartridges, equivalent to over 53 EB in data capacity.
- Apr. 30, 2014 Sony Corporation independently developed a soft magnetic under layer to successfully demonstrate the world's highest areal recording density for tape storage media of 148 GB/inch².
- May 19, 2014 Fujifilm and IBM successfully demonstrated a record areal data density of 85.9 GB/inch² on linear magnetic particulate tape using the NANOCUBIC and Barium Ferrite (BaFe) particle technologies. This breakthrough equates to a standard LTO cartridge capable of storing up to 154 TB of uncompressed data, making it 62 times greater than today's available LTO-6 cartridge capacity.
- Sept. 9, 2014, IBM announced Linear Tape File System (LTFS) Enterprise Edition (EE) Version 2.1.4.4 extending LTFS to the tape library support.
- Sept. 10, 2014, the LTO Consortium announced the extended roadmap with LTO-9 and LTO-10.
- Oct. 6, 2014, IBM announced the TS1150 enterprise drive with a data rate up to 360 MB/sec and a native cartridge capacity of 10 TB.
- Nov. 6, 2014 HP announced its new release of StoreOpen Automation that delivers a solution for LTFS in automation environments with Windows OS, available as free download.

This list includes an impressive number of vendors and suppliers that are committed to investing time and resources to tape technologies. The diversity of vendors and suppliers also suggests that the industry is not solely for a niche group of players but is developing into a storage market as a whole.

It should be widely known, that the following aspects of tape vs. disk apply today:

- Tape is cheaper to acquire,
- Tape is less costly to own and operate,
- Tape is more reliable,
- Tape now has media partitions for faster "disk-like" access,
- The capacity of a tape cartridge is higher than a disk drive's capacity, and the media life for tape is 30 years or more for all new data [11].

Fig. 2 shows the areal density developments of disks and tape since the 1990s. While for disks the improvement was a 35% improvement from 2003 to 2009, the future outlook of that development is now slowing down, while the tape roadmap stays stable growing with a 33,15% per year until 2022.

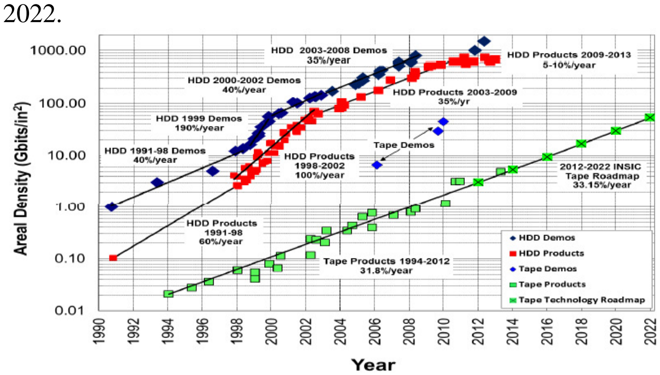

Figure 2. Areal Density of Hard Disk and Tape –Laboratory Demonstrations and Products [10] [11]

III. Analysis of Tape vs. Disk

As section II listed an overview of recent tape technology advancements, mainly covering the terms of capacity increase and of cost decline, this section brings the aspects of costs per terabyte, capacity and performance together. As already demonstrated, costs per terabyte in disk and tape will decrease over the next 10 years. While data volumes continue to grow, at least at the same or even higher speed, a combination of these two developments show, that in four years the break even for tape in terms of cost effectiveness per terabyte will be reached.

Fig. 3 provides the clear picture of this development, reached by 2019: the red trend line is the decrease in cost per terabyte for disks, the blue trend line is the decrease in cost per terabyte for tape, while the yellow trend line shows, as an example, the archived data growth from Oak Ridge National Laboratory.

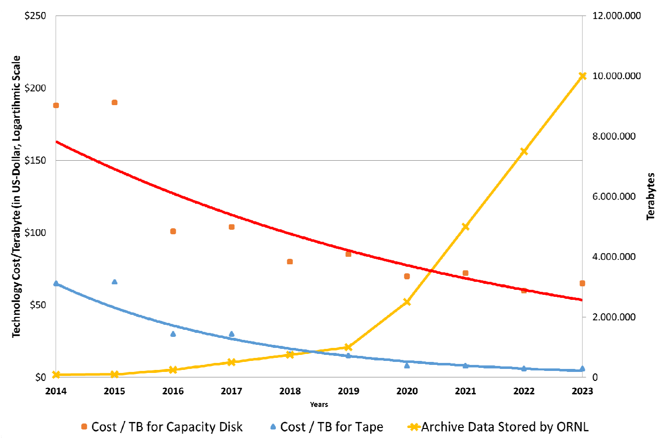

Figure 3. Development of Technology Cost per Terabyte vs. Terabyte of data growth [9] [13]

Not just for high performance research organizations, such as Jülich Supercomputing Centre (JSC), the third aspect of performance is vital for IT-managers to define their storage strategy and layout. The Fig. 3 combines only cost per terabyte and data growth rate. However, the criteria of performance must be viewed as another third aspect of a tape vs. disk comparison.

A tape data rate using a 10 TB native capacity on the IBM TS1150 is up to 360 MB/s [5]. A standard commercially available hard disk drive of 4 TB SAS has a performance up to 6 gigabytes per second. Comparing the aspect of performance alone for tape vs. disk, the disk will always win this criteria. As a result, tape will need a "little helper", realized by tape storage software, to perform much better than disk. As the access and interchange capabilities of tape have to be improved, Linear Tape File System (LTFS), a long awaited file system specification for LTO, is available since mid of 2010 [11].

LTFS provides a dual partitioning functionality, allowing the tape to be self-describing. Metadata operations, such as browse directory tree structures and file-name search, are performed more quickly and do not require physical tape movement [11].

Using the Linear Tape File System, files can be created on tape and accessed similar to the process of creating and accessing files on an external hard drive or a USB flash drive. Applications, such as file browsers, image viewers and media players can directly browse and access files on tape. LTFS enables easy and simple use of tapes in desktop computers and embedded systems. One drive can read and write data at an impressive 1.2 Gigabytes per second, (approximately two times the speed of a hard drive), when using LTFS [12]. This is still not quite that fast than a single HDD, but tape comes close. The following section IV describes, how performance can be highly achieved by using tape and respective tape software technologies.

IV. Big Data Storage Software

In Spring 2010, a couple of leading storage vendors and suppliers formed a collaborative industry alliance called the Active Archive Alliance in order to educate end user organizations on the evolving new technologies that enable reliable, online and efficient access to their archived data [14]. "Active Archive is a combined solution of open systems applications, disk, and tape hardware that allows users to access all of their data, and gives you an effortless solution that stores and manages all of your data." [15] The National Center for Supercomputing Applications (NCSA) successfully deployed an active archive for the world's largest active file repository in 2014. This active archive has a 99,99% availability, consists of 224 TS1140 tape drives at 240 MB/s each with a 52,7GB/s aggregate I/O throughput and an ability to grow up to 380 PB of data storage over the course of five years [16]. Data-intensive needs of scientists and engineers are driving this massive scale of a data archive in conjunction with the use of its Blue Waters high performance supercomputer at NCSA. The solution consists of a File-based active archive to enterprise tape, including the following hardware and software components:

- IBM HPSS software solution with DataDirect Networks disk cache
- Spectra Logic T-Finity enterprise tape library
- IBM TS1140 3592 JC enterprise data tape with 4.0 TB native capacity

"As the size of the archives increases, the tiering in these systems could be a big driver for a variety of digital storage technologies, from flash memory to hard disk drive, to magnetic tape or optical discs and even cloud storage (which can be any combination of these technologies)". [18] A leading concept has been long-time established in the storage industry for this: it is called hierarchical storage management (HSM).

Due to lack of performing software components (or very custom designed software in combination with an proprietary file system at a high price or at high installation complexity levels), difficulties in ensuring data end-to-end integrity, high implementation efforts needed for a high availability deployment, and access to data with an adequate speed in

 ISBN: 978-1-61208-416-9

Gigabytes per second hindered IT Departments to introduce in their storage concepts also tape technologies as active part in their storage infrastructure.

Simpler HSM deployments in using a combination of higher valued fibre channel disks with lower commodity type of SATA disks and only using tape as the backup of this storage infrastructure is a very common deployment in most of organizations today. “Easy-Tiering” software technologies are in place to move within the disk pools older data to lesser expensive disk hardware or can be found in file systems supporting policies for hierarchical storage management (i.e., in IBM General Parallel File System (GPFS)). In periodically terms that data gets then stored onto tapes in order to be archived. When those data arrived at the tape pool, that data was very sporadic or only on special events accessed (i.e., huge data loss due to catastrophic damages, or for special audit purposes, etc.). A high velocity of that data access was not needed frequently or the data was not a very business critical item. This will change.

Even with the development of Flash technology arrived, this type of technology has not solved all performance problems for small block files [15]. For increased metadata performance in i.e., for database storage systems, Flash and solid-state-disks (SSDs) will play in future a vital role and will be commercially priced in an attractive way. But for massive data scale up to multiple PB and EB with the need to access that data in a faster way, the current standard HSM concepts will fail to provide a professional solution and/or surge higher costs due to higher software licenses fees while storage capacity increase. It is time that IT Department managers look at tape in a different way. For archival storage systems, Flash could improve metadata performance. It also has potential use as a low latency cache for user data to be used in the hierarchy of storage that most archival storage systems offer today with disk and tape alone [13].

IBM is the leader for deploying a proven highly scalable archive and Hierarchical Storage Management system with its HPSS offering. HPSS has been used successfully for very large digital image libraries, scientific data repositories, university mass storage systems, and weather forecasting systems, as well as defense and national security applications. Major compute-intensive and data-intensive sites such as Lawrence Livermore National Laboratory (LLNL), Los Alamos National Laboratory (LANL), Brookhaven National Laboratory (BNL), Oakridge National Laboratory (ORNL), Argonne National Laboratory (ANL), the European Center for Medium-range Weather Forecasts (ECMWF), the German Climatic Research Center (DKRZ), the German Meteorological Service (DWD), and the National Oceanic and Atmospheric Administration (NOAA) have all chosen HPSS for their mission-critical data assets.

HPSS is a network-centered, cluster-based software offering that provides for stewardship and access of many petabytes of data. HPSS is capable of concurrently accessing hundreds of disk arrays and tape drives for extremely high aggregate data transfer rates, thus enabling HPSS to easily meet otherwise unachievable demands of total storage capacity, file sizes, data rates, and number of objects stored. In recent installations HPSS demonstrates an aggregated throughput for read/write from clients from disk to the tape archive with up to 18Gigabytes per second [8]. HPSS is designed to scale horizontally and consists of a HPSS Core Server, that manages the metadata, and HPSS Data Movers that move the data between the disk cache and the tape pool. Fig. 4 shows the architecture of HPSS.

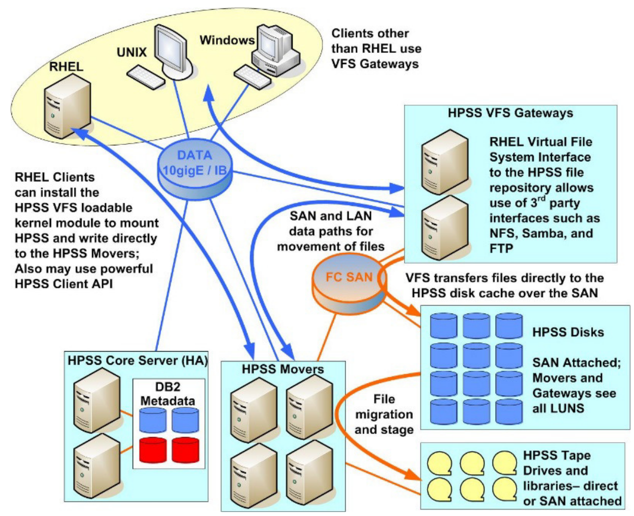

Figure 4. HPSS Architecture, IBM Corp.

In order to come to a start configuration for an HPSS installation, it is necessary to provide a “rule-of-thumb” on how to measure the demands for a HSM concept for an individual organization. The DOE Lab report states that there is a correlation of capacity with the amount of main memory of the server systems [13].

However, today business organizations that have increasing data growth needs, with smaller but high capable server environments, need to consider on what journey they want to continue with their storage environment if they reach certain storage capacity levels. Starting with an active data storage capacity of 3 PB and an expected exponentially data growth rate for the next 3 to 5 years, a professional HSM concept should be evaluated.

Fast and secure access to data is a key requirement in business environments today, as more data is being created, reused or reactivated for a different purpose. Those storage environments should be enabled with a flexible HSM concept in order to drive efficiencies. HPSS is being installed on commodity x86_64 node servers and with a standard storage disk system i.e., an IBM DS3860 or with a NetApp’s E-Series 5500 disk storage system or similar, with the operating system Red Hat Enterprise Linux.

Table I shows for HPSS performance values for a very small designed HSM configuration. It consists of one core server with optional manual cold-standby functionality, and two standard x86 Intel server as Data Movers. Three PB of data would be efficiently managed by this HSM configuration. The HPSS configuration would achieve already with this rather simple server hardware configuration the “standard” disk performance of today’s 4 TB HDDs.

 ISBN: 978-1-61208-416-9

TABLE I. AGGREGATE THROUGHPUT RATES OF HPSS, P. SCHAEFER, HPSS ENGINEER, IBM CORP.

Configuration	Disk*	Fibre*	10GigE network*	Tape system*
Small LTO/6	2458	1638	3414	896
Small TS1150	2458	1638	3414	1152
5Gbps LTO/6	6145	4095	9104	3072
5Gbps TS1150	6145	4095	9104	3168
* all values in Megabytes/second (Mb/s)				

As a start for a small HPSS, i.e., as a start basic configuration having one core server, one data mover and a disk cache with dual controllers of 844 TB of usable space (8+2p LUNs and 10% sparing), and for a HPSS system with a moderate 5Gigabyte per second, using one core server, two data movers and a disk cache with 985 TB of usable space (8+2p LUNs and 10% sparing), both disk cache each including the performance turbo option, achieve an reasonable aggregated throughput performance at lower costs.

V. CONCLUSION

There is no doubt, that data, especially unstructured data, will continuously grow, as individuals, organizations and researchers want more capacity for storage, access information almost everywhere, and conduct more valuable data experiments at any time with any interface. Tape is still being seen as a laggard in its development in the past five years. But in the past 18 months, tape technology vendors and developers have reached new levels of improvement, so that tape can be viewed for future storage system requirements differently.

From 2019 tape will play a greater role, as leading analysts, consultants and vendor predict. Fig. 3 clearly indicates the breakeven point, when to consider tape as an active storage component and not just as a backup device. Section III described the development in regard of cost per terabyte and terabyte of data growth.

For tape and performance, intelligent Big Data storage software is being used to improve the performance and speed of such tape systems. In addition to that, new data center service models are evolving, i.e., Cloud-based data centers. For that, the next stage is to enhance those cloud data centers operated by service providers with a “Big Data Cloud Storage” offering. This would be consisting of large tape environments and the use open source software such as OpenStack, to provide an easy to use enterprise-class cloud based storage service. This would be very suitable for very large archives (petabytes to exabytes) with occasional or significant retrieval needs, and it would offer very low cost storage together with affordable and predictable retrieval costs.

REFERENCES

[1] A. Dayley, G. Landers, A. Kros, and J. Zhang, “Magic Quadrant for Enterprise Information Archiving 2014”, Gartner, Inc., Nov. 2014.

[2] IBM Press Release, 2007, Report: IBM Passes All Major Vendors to Take #1 Position in Disk and Tape Storage Hardware Market Share Worldwide for 2006, URL: http://www-03.ibm.com/press/us/en/pressrelease/21643.wss, [accessed 2015-02-12].

[3] T. Pearson, 2007, “IBM passes all major vendors in storage”, URL: https://www.ibm.com/developerworks/community/blogs/InsideSystemStorage/entry/ibm_passes_all_major_vendors?lang=en, [accessed 2015-02-12].

[4] L. Dubois and A. Nadkarni, “Worldwide Storage in Big Data 2014 – 2018 Forecast”, IDC Opinion, Aug. 2014.

[5] Tape Storage Council, 2014, “2014 TSC Memo: Data Growth and Technology Innovation Fuels Tape’s Future”, URL: http://tapestorage.org/wp-content/uploads/2014/12/2014-Tape-Storage-Council-Memo.pdf, [accessed 2015-02-11].

[6] LTO® Ultrium, (2014), LTO® Program Further Extends Product Roadmap Through Generation 10, URL: http://www.lto.org/2014/09/lto-program-further-extends-product-roadmap-through-generation-10/, [accessed 2015-02-13].

[7] T. Olavsrud, 2012, “How to Implement Next-Generation Storage Infrastructure for Big Data”, URL: http://www.computerworld.com/article/2503174/enterprise-applications/how-to-implement-next-generation-storage-infrastructure-for-big-data.html, [accessed-2015-05-10].

[8] M. Boettinger, J. Cadmus, J. Meyer, and J. Nachtsheim, 2015, “German Researchers Use IBM Big Data Solution To Manage World's Largest Trove Of Climate Data”, URL: http://www.prnewswire.com/news-releases/german-researchers-use-ibm-big-data-solution-to-manage-worlds-largest-trove-of-climate-data-300080089.html, [accessed 2015-05-08].

[9] D. David Floyer, 2014, “The Emergence of a New Architecture for Long-term Data Retention”, URL: http://wikibon.org/wiki/v/The_Emergence_of_a_New_Architecture_for_Long-term_Data_Retention, [accessed 2015-05-14].

[10] S. Weingand, 2014, “Tape ist lebendiger denn je”, URL: http://de.slideshare.net/JosefWeingand/ibm-dpr-tape-ist-lebendiger-denn-je-0814, [accessed 2015-05-13].

[11] F. Moore, 2014, “Tape. New Game. New Rules”, Horison Information Strategies, URL: http://horison.com/wp-content/uploads/2014/10/Fujifilm-Tape-New-Game-New-Rules-21.pdf, [accessed 2015-02-13].

[12] StorageDNA, 2012, Facts about LTO and LTFS, URL: http://www.storagedna.com/media/documents/StorageDNA_whitepaper_Facts-about-LTO-LTFS_web.pdf, [accessed 2015-05-14]

[13] National Energy Research Scientific Computing (NERSC) Facility, Extreme Scale Workshop, 2009, HPSS in the Extreme Scale Era, URL: http://www.nersc.gov/assets/HPC-Requirements-for-science/HPSSExtremeScaleFINALpublic.pdf, [accessed 2015-02-13].

[14] Active Archive Alliance, 2015, Introduction, URL: http://www.activearchive.com/, [accessed 2015-02-13].

[15] Spectra Logic, 2015, What is Active Archive?, URL: http://www.spectralogic.com/index.cfm?fuseaction=solutions.showContentAndChildren&CatID=1690, [accessed 2015-02-13].

[16] Active Archive Alliance Case Study, 2014, (NCSA) National Centers for Supercomputing Applications, NCSA Employs Active Archive for World’s Largest Active File Repository, URL: http://www.activearchive.com/common/pdf/AA-Case-Study-NCSA.pdf, [accessed 2015-02-13].

[17] T. Coughlin, 2014, Archiving Storage Tiers Part 2, URL: http://www.forbes.com/sites/tomcoughlin/2014/12/30/archiving-storage-tiers-part-2/, [accessed 2015-02-13].

[18] T. Couglin, 2015, Flash Memory in Archives, URL: http://www.forbes.com/sites/tomcoughlin/2015/01/31/flash-memory-in-archives/, [accessed 2015-02-13].

 ISBN: 978-1-61208-416-9