

WHITE PAPER

New Possibilities for Tape Storage in the Age of Cloud and Big Data

Sponsored by: Fujifilm

Masaaki Moriyama

February 2014

IN THIS WHITE PAPER

Many IT managers still have the "old view" of tape storage. Tape, however, has made great progress in terms of reliability, long-term data retention, and ease of use. Moreover, it has evolved into a medium that is well suited to the low-cost, long-term retention of the immense quantities of data generated today. As we enter the age of cloud and big data, when the use of large quantities of data is key to both business and social activities, rethinking our "old view" of tape storage can lead to more appropriate and effective investments in IT. This IDC White Paper examines new opportunities for tape storage, focusing specifically on storage conditions in the age of cloud and big data, as well as the innovations that have been made in tape storage technology. It also describes the efforts made by Fujifilm to innovate in this field and develop new demand, and identifies both the opportunities and challenges for its tape storage business.

SITUATION OVERVIEW

Revisiting the Value and Potential of Tape Storage

Tape storage was the main backup storage system for many years, but its usage has declined recently due to the penetration of disk-based backup and the emergence of purpose-built backup appliances with deduplication functions. The drawbacks of tape storage are well known: long restoration times and complex media management compared with those of disk storage. Most companies continue to use tape storage in combination with disk-based backup, for example, disk-to-tape or disk-to-disk-to-tape, but a growing number are no longer considering tape storage for new backup installations. Some IT managers consider tape storage as "old technology." However, new demand is starting to come from corporate and service provider datacenters, driven by cloud, big data, innovations in tape storage technology, and other changes in the IT market. This paper looks at user case studies on tape storage to analyze the changes that are taking place in the IT market, the storage that these changes demand, and the technical innovations that are occurring in tape storage. It also identifies new values and potential for tape storage.

User Case Study: Recognizing the Value of Tape Storage

The case studies in this section involve a user that recognizes the value of tape storage and continues to invest in it, and a user that is investing in tape storage for the first time.

METEO-FRANCE

One example of an end user with rapidly increasing data capacity is Météo-France, the French national weather forecast entity. They project their creation of data per day will increase from 40TB today to 400TB by 2018 and their total capacity will increase from 15PB today to 150PB by 2018.

Météo-France has a hierarchical storage management (HSM) system consisting of disk and tape storage. According to Météo-France, in order to manage their fast growing data capacity at low cost, 80% of their data is on tape storage today and they project it will increase to 90% in the future.

Météo-France pays particular attention to high levels of tape security and sustainability. They therefore appreciate the new Barium Ferrite (BaFe) technology currently being used in IBM TS1140, Oracle T10000 C/D and LTOG6 that assures a roadmap for rapid capacity increases, higher recording stability, and longer archival life. Météo-France claims that energy consumption and floor space are becoming increasingly critical when considering their storage options. They see that when the data is projected to grow 10 times over in 5 years, increasing the storage split in favor of tape is the only viable option one has based on lower cost/TB, lower energy consumption, and less floor space. Météo-France plans to continue to invest in tape storage and is looking forward to seeing further development on tape technology to accelerate the increase in capacity per tape cartridge.

Magic TV

End users that require tape storage are not limited to very large enterprises which store more than tens or hundreds of petabytes.

Magic TV, a French production house that creates one of the most famous TV programs in France called *Le Grand Cabaret*, has also realized benefits from adding tape storage into their system.

According to Mr. Oliver Durand, post-production manager of Magic TV, they were storing their data onto 2 HDD servers, 83TB on one server and 37TB on the other. In 2011, Magic TV experienced a failure in one of their HDD servers that had 37TB of data (approximately 400 programs) due to overheating of the server caused by an unexpected air conditioning issue within their server room. All 400 programs were lost and were only recovered after spending a very long time at a considerable cost.

After Magic TV experienced this disaster, one of their resellers offered them an LTO tape storage based system to better preserve their data and to support future increasing capacities at a lower cost. Magic TV found that the cost to install tape

storage was approximately equal to the cost of recovering the data lost from their HDD system and decided to invest in tape storage to achieve better data security.

Presently, they archive their data on tape while still keeping their HDD server system in place, recognizing that tape and HDD can play different roles in supporting Magic TV's activities. Mr. Oliver Durand is fully confident that they have made the right step forward.

Magic TV plans to create 18 prime time shows annually which equals approximately 3TB or more data each year and to continue investing in tape storage.

Storage Requirements in the Age of Cloud and Big Data

IDC views the changes taking place in today's IT market as a shift from the 2nd Platform (PCs, client/server systems) to the 3rd Platform (cloud, mobility, big data, and social technology). In the future, it will be the 3rd Platform that will drive growth in IT spending and will create many business opportunities and business values. The major generational changes in platforms, from the 1st Platform (mainframes), to the 2nd, and on to the 3rd can be observed in magnitude-order changes in the numbers of accessing users and devices, numbers of applications used, and types and volumes of data created. With the emergence of the 3rd Platform, companies continue to accumulate enormous numbers of applications and volumes of data in their datacenters and clouds (i.e., service provider datacenters). The data that is being accumulated is also changing from human-generated to machine-generated (data for measurement, sensor, and control), and this further accelerates diversification and volume growth of data.

Big data is a topic that has gained much attention, and refers not only to the large volume and variety of data accumulated, but also attempts to analyze and use this data to create new values in business and social activities. Not only will there be greater use of cloud and big data, but corporate and social activities themselves will become more data-oriented, and this leads to forecasts of continued growth in storage investments in datacenters.

The challenge for storage in the age of cloud and big data is to contribute to the "analysis" and "retention" of this immense, diverse data. On the "analysis" side, there is growing demand for significantly better storage I/O performance because of the need for fast processing. On the "retention" side, there is an incomparably greater quantity of data to be stored than in the past, and it is important that it is archived at low cost for long periods of time. "Low cost" in this context refers not only to initial investments, but also to operations management and energy costs.

In the past, companies archived data for compliance and security purposes. In a growing number of cases, disk storage was preferred for these archival activities because of the ease in searching for files. With the advent of cloud and big data, however, the old criteria may no longer apply to archiving investment decisions because today's data is more diverse, there is more data to store, capacity is increasing with greater speed, and retention periods are becoming longer.

Service providers, research institutes, universities, broadcasters, video services, and other companies handling large volumes of data have been among the first to face the challenges of storing enormous amounts of data for long periods at low cost. One of the solutions that such companies have found to address archiving in the age of cloud and big data is tape storage. Many companies are taking another look at the benefits that tape storage offers to archiving in terms of high capacity, low cost (price per GB), low power consumption, and long life.

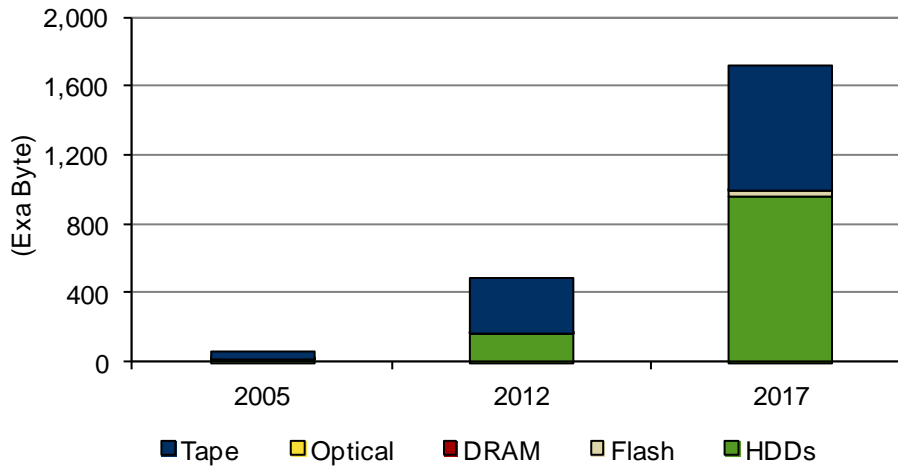
With the emergence of Linear Tape File Systems (LTFS), which give tape storage the same ability to use and manipulate file data as disk storage, tape storage is no longer limited just to backups, disaster recovery, and archiving, but is also an easy choice for active archives. "Active archives" refers to tiered archive solutions made up of multiple storage media with different properties (cost, capacity, I/O performance, etc.). These systems enable data to be referred to and searched online, while also restraining storage costs by moving data to lower-cost tiers. For example, lower-cost tape storage is used in the lower tiers of file storage and object storage services that keep large volumes of data in the cloud, allowing archive capacity to expand while also restraining retention and electric power costs.

Installed Raw Tape Storage Capacity Continues to Grow

IDC has provided forecasts for installed raw storage capacity up to 2017 in four markets (enterprises, PCs, mobile devices, and entertainment/others) broken down by medium type (HDD, optical, tape, flash, dynamic random-access memory [DRAM]). The enterprise market includes storage systems, servers, and libraries. Figure 1 shows 2005 and 2012 actual data and 2017 forecasts of installed raw storage capacity in the global enterprise market segmented by medium type. In 2005, the total installed raw storage capacity stored on all media in the enterprise market was 56 EB (exabytes), which grew 8.8-fold to 489 EB in 2012, and is forecast to grow another 3.5-fold to 1,718 EB in 2017. As cloud and big data gets used more and more, the volume of data stored on servers and storage systems in user companies and service provider datacenters will only continue to grow.

FIGURE 1

Worldwide Installed Raw Storage Capacity in Enterprise Market by Medium Type 2005, 2012, and 2017



Note: Based on the IDC report *Where in the World Is Storage: A Look at Byte Density Across the Globe* (IDC #243338, October 2013).

Source: IDC, 2013

The breakdown by medium type indicates that in 2012, tapes had doubled the installed raw capacity of HDDs in the worldwide enterprise market. By 2017, HDDs will hold the largest volume of data due to increased use both as primary storage and secondary storage. Flash is forecast to have the highest growth rate of any media. This is because the enterprise market demands better storage I/O performance to improve the performance of applications and facilitate big data analysis. Nonetheless, flash will still account for low single digits of the total installed raw storage capacity in the worldwide enterprise market in 2017.

Installed raw capacity of tape storage will grow from 327 EB in 2012 to 726 EB in 2017 (2.2 fold). The ratio of tape storage to total installed raw storage capacity in the enterprise market will decline, but will still be above 40%, and it will continue to play a major role in the retention of data. The benefits that it offers in terms of low cost, large capacity, and long life are expected to create new demand for tape in the long-term archiving of large-volume data.

One observation about storage investments that can be seen from the forecasts of installed raw storage capacity in the enterprise market by medium type is that enterprises are increasingly making use of different kinds of storage media depending on suitability for specific purposes, and among the criteria used in the selection of media are I/O performance, price per GB, data retention periods, access frequency, and energy costs.

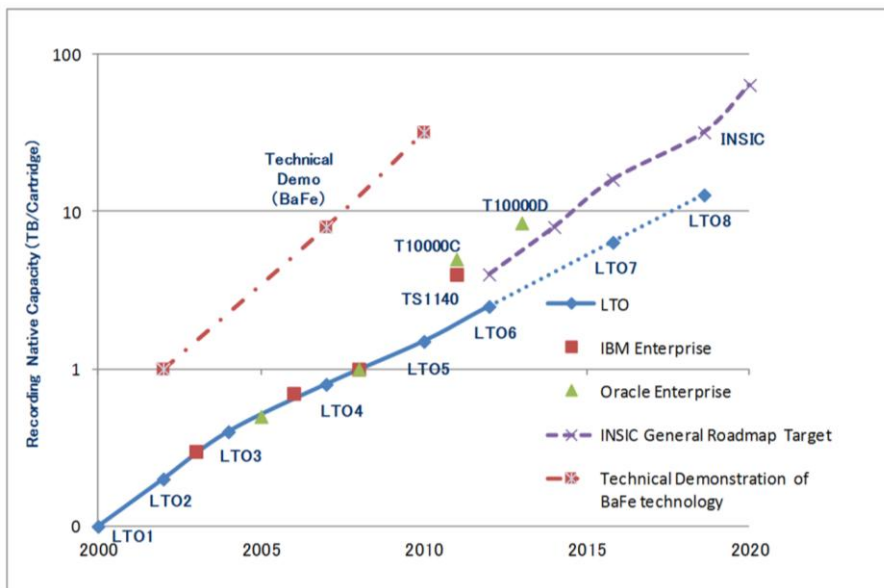
Innovations in Tape Storage

Tape Storage Systems Make Advances in Capacity and Speed

Tape storage offers large capacity and high speed. In the LTO storage market, the most popular tape technology, the sixth-generation "LTO-6" was launched in November 2012, with non-compressed capacity of 2.5 TB (terabytes) per cartridge, and a data transfer rate of 160 MB/second. In the past, there were a number of competing standards in the tape storage market, but LTO carved out a position as a general-use medium, and the technology has been developed and brought to market on the basis of a published standardization roadmap. Currently, the roadmap up to the eighth generation has been published. LTO-7 will have non-compressed capacity of 6.4 TB/cartridge and data transfer rate of 315 MB/second; LTO-8, non-compressed capacity of 12.8 TB/cartridge and data transfer rate of 472 MB/second (Figure 2).

FIGURE 2

Capacity Gains for Tape Cartridges, Future Roadmap, and BaFe Technology R&D Trends



Notes:

- The recording capacities of individual cartridges are created on the basis of published information from vendors and the LTO Consortium. Recording capacities have been published for LTO-7/8, but the date of launch has not been announced
- The dotted "INSIC" line in the graph illustrates the General Roadmap Targets for Recording Media Technology published by the Information Storage Industry Consortium.
- The "Technical Demo (BaFe)" dotted line in the graph illustrates the recording capacity R&D results using the Barium Ferrite (BaFe) technology jointly developed and announced by IBM and Fujifilm in 2010.

Source: IDC Japan, 2014

Vendors are also using proprietary standards to bring large capacity and high speed to high-end tape storage. IBM began shipping the "TS1140," which has non-compressed capacity of 4 TB/cartridge and data transfer rate of 250 MB/second, in June 2011. Oracle launched the "T10000D" with non-compressed capacity of 8.5 TB/cartridge and data transfer rate of 252 MB/second in September 2013 (Figure 2).

Further advances in technology will be required to increase the capacity and speed of tape storage. Crucial to this will be the development of new tape media in addition to hardware (drives and libraries) and software. Higher recording densities and high durability of tape media will translate directly into higher capacity and speed. Among the recent innovations in tape media recording density is the use of a new generation of magnetic particles.

For many years, tape media used metal particles, but the latest generations use barium ferrite particles. Both the IBM TS1140 and the Oracle T10000D use tapes with barium ferrite particles. In the LTO space, metal particle tapes were offered in addition to barium ferrite particle tapes from LTO-6. Fujifilm, the developer and vendor of barium ferrite tapes, is working jointly with drive vendors to further extend capacity.

WORM and Encryption Support

Innovations have also been seen in the way tape storage operates. For example, Write Once Read Many (WORM) functions were added to LTO from LTO-3, and encryption functions from LTO-4. WORM is a function that prevents falsification of data after it has been written to tape media; encryption uses AES (Advanced Encryption Standard, 256 bit) to encrypt data when it is written to tape media. Both functions were developed to address security and compliance concerns for the data stored on tape media. LTO-5 also added partitioning functions that create logical partitions on tape media.

New Possibilities for Tape Storage Opened Up by LTFS

The Linear Tape File System (LTFS) is a technology that was developed by IBM and then later opened to the public. LTFS gives tape storage the same ease of operation and manipulation as disk storage by allowing direct access from the host to file data on the tape media. In LTO, the most commonly used tape storage format, LTFS has been supported since LTO-5, which also brought support for partitioning functions.

LTFS is able to store data on tape media in file format because it uses partitioning functions to make partitions on the tape and write the metadata of the file to the index partition and file data to the data partition according to the LTFS format. Users have random access to the file data they require by accessing the metadata in the index partition of the tape media, the same way they access files in disk storage. Some of the benefits of LTFS are as follows:

- ☒ **Easier operation.** Prior to LTFS, tape storage wrote data to tape media in proprietary formats determined by backup software and the like, and to read the data, it was necessary to use the same software to restore the data to the host

disk. By contrast, data written in LTFS is able to read the data in file format as long as the system has an access interface that supports the LTFS format. This makes it possible to exchange tape media among different systems just like USB memory sticks. Tape media containing large volumes of file data can be transported to other locations for editing in a manner similar to how optical discs and USB memory are used today.

- ☒ **Ease of management.** With LTFS, users can take advantage of the file management techniques that they have grown familiar with from disks, for example, drag-and-drop copying. This makes it easier to manage tape storage and lessens the load on managers.
- ☒ **Expanded use.** LTFS allows hosts to access files directly on the tape, making it easier to use tape storage for both near line and active archive storage. File searching takes longer than with disk storage, but depending on the type of data, users can take benefit from directly referencing and searching for large files stored on inexpensive tapes from the host.

LTFS opens new possibilities for tape by providing the ease of use and operation similar to disk storage but with traditional strengths of large capacity, low cost, long media life, encryption and WORM functions. IDC believes that new demand can be generated for LTFS-based tape storage both at service providers that provide cloud services, and also in the areas where there are growing needs for online archives of unstructured data like healthcare, broadcasting/video, surveillance cameras, universities, and research institutes.

Fujifilm's Contributions to Innovation in Tape Storage

We examined Fujifilm's contributions to innovation in tape storage in two aspects: in terms of the development of high-density tape media, and evangelism of tape storage itself.

Development of Barium Ferrite Tape Media

Barium ferrite particles have begun to be used in today's tape media as magnetic particles. Fujifilm developed tape media using barium ferrite particles to improve the reliability of tape while also providing higher density and capacity.

The surface of the tape media is coated with magnetic particles, and the density with which tape media is able to record information is largely dependent on how fine the magnetic particles are and to what degree of precision a thin coating of these particles can be applied. Barium ferrite particles are finer than conventional metal particles and are well suited to high-density recording. They are also oxides, which avoid the problem of oxidation and degradation of magnetic recording characteristics experienced with conventional metal particles. Even at very fine sizes, they continue to hold their magnetic force for long periods, so they are suitable for long-term data retention. These magnetic particles also achieve low noise and high output, resulting in a large signal and regenerative recording and noise ratio (SN ratio) from the head of the tape drive. This in turn improves the stability of data writing and reading.

Fujifilm expects to increase the recording capacity of single cartridges to 100 TB or more in the future by innovation in high-rigidity base film and precision servo technology besides further developing barium ferrite particles.

Evangelism of Tape Storage

Tape Storage Education

Fujifilm holds a regular "Global IT Executive Summit" in the United States for major users of tape storage. The objectives of this program include introducing users to the latest tape storage technologies, discussing case studies, recognizing the value of tape storage, and sharing best practices. The company is also considering a similar program in Europe. In Japan, it publicizes the latest advances in tape technology and archiving, and encourages users to take a new look at the value of tape storage.

"d:ternity" Archiving Service

Fujifilm offers the "d:ternity" tape-based archiving service at three locations: Japan, Europe, and the United States. There are slight differences in the services offered by each of the locations, but the basic concept is to give users an inexpensive, long-term solution to archiving challenges.

The service converts user data into a format that is not dependent on the proprietary formats of backup and archive software vendors, transfers it to a storage system that has advanced a security and data protection functions, and then archives it in a tape library. The "d:ternity" service means that users do not need to invest in their own hardware, backup software, or staff for long-term archiving, nor do they need to spend manpower or time on data migration as tape generations advance, which has been one of the biggest challenges in long-term tape archiving. The result is substantial cost savings compared with conventional in-house archiving.

When archive data is required, it is easy for users to search data files in their archive over the web and issue return requests. Data can be returned either in physical media (tapes, HDDs, SSDs, optical discs, etc.), or over the Internet. The "d:ternity" service leverages the value of tape storage as an archiving service and also solves many of the challenges in tape storage.

Fujifilm's Business Opportunities and Challenges

The advent of cloud and big data has emphasized the importance of being able to generate, retain, and analyze large volumes of data and use them to identify values for business and society. This has caused people to take another look at the value of tape storage which, using state-of-the-art technology, is able to archive large volumes of data at low cost for periods of more than 30 years. In light of these circumstances, IDC believes that it will be important for Fujifilm to address the following challenges in expanding the business opportunities for tape storage.

- ☒ **Encourage potential users to reexamine the value of tape storage.** The most important challenge is to encourage user companies to reexamine the value of tape storage. This will require collaborating with vendors of tape drives,

automation systems, and software to collect and introduce many user success stories (particularly in archiving). Having a number of different examples and different sizes of company, different sectors and different amounts of tape media usage will enable the company to provide a broader range of corporate users with the materials they need to make decisions and encourage them to consider the values that tape storage offers. If at all possible, it will be important to provide numerical data that demonstrates the superiority of tape storage for long-term archiving in terms of total cost of ownership (TCO).

- ☒ **Leverage expertise in the tape business.** One of the challenges facing the company is how to leverage the expertise it has accumulated in the tape media business to encourage users to reexamine the value of tape storage and to solve the challenges that confront them. The "d:ternity" service that the company offers in Japan, Europe, and the United States is a good example of leveraging expertise to solve users' tape storage challenges. Either on its own, or in collaboration with vendors of tape drives, automation systems and software, the company should encourage the use of tape by providing services that address the challenges facing users in archiving and data protection.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2014 IDC. Reproduction without written permission is completely forbidden.