

**Draft Draft Draft Draft v2.1 9/1/98**

## **Technology Assessment of High Capacity Data Storage Systems: Can We Avoid A Data Survivability Crisis?**

**M. Halem, F. Shaffer, N. Palm, E. Salmon, S. Raghavan\*, and L. Kempster\***  
**Earth and Space Data Computing Division, Code 930**  
**NASA Goddard Space Flight Center**  
**Greenbelt, MD 20770**

### **EXECUTIVE SUMMARY**

This technology assessment of long-term high capacity data storage systems identifies an emerging crisis of severe proportions related to preserving important historical data in science, healthcare, manufacturing, finance and other fields. For the last 50 years, the information revolution, which has engulfed **all** major institutions of modern society, centered itself on data - their collection, storage, retrieval, transmission, analysis and presentation (Drucker, Forbes Magazine, 8/24/98). The transformation of long term historical data records into information concepts, according to Drucker, is the next stage in this revolution towards building the new information based scientific and business foundations. For this to occur, data survivability, reliability and evolvability of long term storage media and systems pose formidable technological challenges. Unlike the Y2K problem, where the clock is ticking and a crisis is set to go off at a specific time, large capacity data storage repositories face a crisis similar to the social security system in that the seriousness of the problem emerges after a decade or two. The essence of the storage **crisis** is as follows: *since it could take a decade to migrate a peta-byte of data to a new media for preservation, and the life expectancy of the storage media itself is only a decade, then it may not be possible to complete the transfer before an irrecoverable data loss occurs.*

Over the last two decades, a number of anecdotal crises have occurred where vital scientific and business data were lost or would have been lost if not for major expenditures of resources and funds to save this data, much like what is happening today to solve the Y2K problem. A prime example was the joint NASA/NSF/NOAA effort to rescue eight years worth of TOVS/AVHRR data from an obsolete system, which otherwise would have not resulted in the valuable 20-year long satellite record of global warming. Current storage systems solutions to long-term data survivability rest on scalable architectures having parallel paths for data migration. This adds significantly to the complexity of storage management systems and their long-term evolvability in three ways:

---

\* Raytheon STX Corporation

\* Raytheon STX Corporation

- (i) Hardware components become obsolete much sooner than the storage media, leaving behind massive volumes of data in legacy systems.
- (ii) Software systems need to evolve maintaining backward compatibility in terms of device readability, portable file management databases, network connections, and other operability issues.
- (iii) The reliability of the system can not be compromised either by inadvertent data corruption or deliberate penetration.

In addition, these systems need to be insulated from blackout and brown outs similar to the telephone and internet system overloads. Technological growth across all these areas has not been impressive in relation to the data growth, since the return on investment by the industry in new technology is often based on incremental demand-supply relationships.

The data growth experienced in the recent past has been of staggering proportions. With the Moore's law of doubling clock speeds every 18 months, even a conservative estimate suggests that a corresponding doubling of storage requirements arise during the same period. However, the data transfer speeds increase at a rate of about 1.3 times every 18 months, and thus fall behind data growth rates at least by a factor of 3. According to a recent Computer Technology Review article, the total storage at a typical Fortune 1000 site is projected to escalate from just 10 TB last year to 1 PB by the year 2000. A typical large data base system for US Government agencies is expected in the next 5 years to accept 5 TB per day, be able to maintain 300 TB on-line (within 15 sec to 1 min. access time), and be able to archive from 15 to 100 PB. In addition to these numbers, data intensive programs such as NASA's Earth Observation System (EOS) and the intelligence data archival systems at the Rome Air Development Center, and scientific laboratories such as Thomas Jefferson National Accelerator Facility will have enormously large scientific databases with very large storage requirements.

Meanwhile, the technology growth in mass storage systems, though impressive in some areas, has barely kept up with the growth requirements in others. Today a single cartridge can hold 10 GB on 128 tracks compared with the 1960 vintage tapes of less than 1 MB capacity, a 10,000 fold increase in 30 years. Magnetic tapes are now available for 1GB/\$1 range with an expected life of 10 years. Optical tape holds forth the promise of very large capacity in a single media unit fairly great information volume density, and very low cost per GB. One optical tape/cassette is equivalent to 250 to 500 magnetic tapes such as 3480 cartridges. However, only now, optical tape technology is beginning to show commercial potential. While revolutionary media such as optical tape storage systems, holographic storage and biological storage such as Halobacterium have not proved commercially viable yet, the prospects of tremendous capacity increases with retrieval and transfer speeds orders of magnitude faster than today appear to be within reach.

Considering the impact of mass storage technology in promoting information growth in global economies of the world, U.S. needs to maintain its leadership and technological edge in this field.

This technology assessment leads to two sets of recommendations to achieve this goal. From a short-term perspective, the Government should:

1. Develop a long-term strategic plan to maintain the U.S technological edge in high capacity storage systems.
2. Require agency CIOs to collect annual reports internally on data growth and related statistics (mounts, access, media, etc.) from all its internal organizations managing more than a TB of data. This information should be used to identify storage media at risk and estimate future requirements within the agencies but also across the government.
3. Encourage agency investment in hardware and software storage system migration to safeguard large volume data through effective, timely, lossless transfer of data from aging systems to the new dense media with minimum impact on the production systems. Encourage national exchanges of technical innovations and sharing of software performance by large repositories of data holders and support national storage standards.

From the long-term perspective, the Government should:

1. Establish a national goal to develop the next generation mass storage technologies towards building an affordable modular exabyte ( $10^{18}$  byte capacity) system. This system should be scalable, distributed, evolvable, and reliable.
2. Forge a strategic alliance of government agencies, industry and academia to maintain a vibrant research program in support of these critical storage technology goals.
3. The Government should serve as a catalyst by promoting storage technology application demonstrations and test bed facilities that help vendors integrate software and hardware to stress end-to-end performance of their systems. The Government needs to develop opportunities for the industry to foster growth through vendor-clientele partnerships.

## 1. INTRODUCTION

This report is in response to a request by the NASA Associate Administrator/OES to assess the long term high capacity data storage media and system related issues of survivability, reliability and evolvability. A preliminary review is presented of the potential impacts to agencies and the business sector, who like NASA, have responsibilities for the management and preservation of long term digital data in mass storage systems. Unlike the Y2K problem, where the clock is ticking and a crisis is set to go off at a specific time, large capacity storage systems face a crisis similar to the social security system in that the seriousness of the problem emerges after a decade or so. The longer one waits to take steps to address the problem, the more draconian become the required measures to save the data.

The information revolution has placed a great demand on the storage industry to develop media that can keep up with exponential growth at constant costs. Recent technology developments have given us the ability to store a peta-byte ( $8 \times 10^{15}$  bits) of data into silos for under a million dollars, about \$1/GB for the media alone. This local capacity is large enough today to handle most of the demand for all known civilian systems except for a few large projects such as the Earth Observation System. However, if one had to face the need to migrate a few petabytes of data to a new media for preservation or economy, it could take more than 10 years to transfer this volume of data in the current environment. This then becomes the crux of the **crisis**: *if it takes a decade to copy data to new denser media and the life expectancy of the new storage media itself is also a decade then one may not have time to save all data on the new media?* Unless major breakthroughs in the speed of I/O transfers occur or new scalable transferring processes are developed, it may not be possible to permanently store all data collected in the future. Considering the explosive growth of our observational and computational systems, even if the process of migrating key information starts as soon as new storage media becomes available, the old storage media needs to have a life expectancy of at least twice that needed to transfer data.

This past decade has seen microprocessor speeds increase according to Moore's law from 10MHz to 1000 MHz, more than a 100-fold increase. Storage densities have also increased from 200 MB to 20 GB on to more compact 4 x4 in. magnetic cartridges. Unfortunately, data transfer speeds, required to read these new media to newer technologies have only increased from 3MB/s to 12MB/s in this same time frame, only a 4-fold increase. With the advent of high-speed robotic silos and software to manage these mass storage repositories, one would have expected the industry to have provided solutions to the data transfer rate problem as well. Unfortunately, data transfer rates are only one of the problems faced by the mass storage user community; other

problems involving performance issues, information security and portable database software methodologies offer other daunting challenges to pushing the technology envelope.

In what follows, we present a historical perspective of attempts to preserve digital data through migration to new storage systems in section 2. While most of the experiences cited were drawn from NASA/GSFC computer experiences in evolving legacy related storage systems, they are not dissimilar from other agency or commercial anecdotal experiences related. Storage growth requirements are estimated in section 3 and a technology assessment is described in section 4. Section 6 presents some suggested recommendations to deal with future data storage crisis. This initial study will be expanded in a later version to include the potential impact of the deterioration or loss of large volumes of data to other agencies and the business community as well as to national security, as the nation moves to a digital society.

## **2.    ASTORAGE CRISIS PERSPECTIVE**

Over the last two decades, a number of situations have occurred where vital scientific and business data were lost or would have been lost if not for major expenditures of resources and funds to save this data, much like what is happening today to solve the Y2K problem. Data loss can be prevented with better data management and investment strategies by computing center managers. Unfortunately, organizations prefer to dealing with future storage crisis by deciding in favor of investing on upgrades to increase computing power. A number of anecdotal examples for crisis situations can be cited to demonstrate the challenges faced by the mass storage community in preventing data loss.

One such crisis was the experience in 1978, prior to launching TIROS N. This observing system carried two major data producing instruments, the Advanced Very High Resolution Radiometer (AVHRR) and TIROS Operational Vertical Sounder (TOVS). Nearly 7 years of continuous data would have been lost but for a dedicated NASA/NOAA/NSF reclamation project at considerable inter-agency expense. The inter-agency project involved migration of the data from the terabit mass storage system(TBM) developed by Ampex Corp., an industry leader in seventies. Ampex entered the large scale high performance systems market in the early 1970's with a high speed tape system designed to hold 1 terabit of data and rapidly locate and transfer the data to supercomputing systems. Five systems were sold to the community principally in the atmospheric modeling area. The largest system was installed at NCAR at a cost of between one and two million dollars over a period of several years and another system was installed at NOAA for archival of TIROS N and future NOAA weather satellite data. Ampex did not have a market for the product and eventually dropped support. The agencies which purchased this equipment were left with their only copy of satellite data on this media with no operational backup. NCAR had the biggest system and was able to keep the tape drive hardware operational for the longest period. Later, NOAA shipped their system to NCAR for spare parts to maintain the archive. This provided NCAR the capability to recover critical satellite data when their systems failed. As a result, after 19 months, 88% of the AVHRR data and 49% of the TOVS data were rescued

at a cost of more than \$500K. Nevertheless, with these seven years, mankind now has a 20 year continuous long data record for the study of global warming. Notable science results have already appeared that would never have been possible but for this effort. In particular, the microwave sounder global warming studies by Spencer at MSFC, unique TOVS decade long sounding products by Susskind at GSFC, changing desert boundaries studies by Tucker at GSFC, and many other valuable studies.

Another good case to cite is the experience at CERN. According to a January 1998 article written by Mark Ferelli, one of the largest data collection efforts in the world today is going on at CERN, the world's largest high-energy particle physics laboratory. It occupies a 30,000 square meter facility in Geneva Switzerland. One of the main problems faced by CERN is a COTS storage management software solution. The solution must be scalable since their upcoming Large Hadron Collider will be generating petabytes of data along several paths at an average of 100 MBps. CERN generates one "event" per second, which creates either 2 or 0.5 MBps at each of the four large experiment sites. Total data collection from these four experiments is over 20 TB per year. The total data generation rate on all experiments in the facility may soon reach 200 TB per year, with a potential to attain a peak rate of 1 petabyte a year. All this data is captured on disk and moved off to tape cartridges. The scientists there are working to capture data they have on old tape formats. For 6 months, they have been copying 200 MB 3480 tapes onto the D3 (Redwood) media. They expect that at the end of 4-6 months, they will have migrated over about 150,000 tape images. At peak performance, they can copy over 1000 tapes per day. When this is done, they plan to begin copying the 3490 tapes with over 20,000 tapes of 800MB capacity in the archive. With such an unending task in mind, CERN has bought a variety of systems with no viable long term solution in sight.

Finally, from experience gained at the NASA Center for Computational Sciences (NCCS) we can relate migration challenges faced at GSFC. Recently, acquisition and introduction of an IBM 3494 robotic mass storage system into its production mass data storage and delivery system was completed in November 1996. It had already reached the saturation point with its StorageTek subsystem, comprised of six robotic 4400 silos holding nearly 33,000 800 MB tapes. A major repacking effort of copying 200 MB tapes onto 36-track 400 MB capacity tapes, and later again onto 800 MB tapes had already been completed. With the indefinite delay in StorageTek delivery of its new "eagle" technology (20 GB tapes with a transfer rate of 12 MB/second), the NCCS had no choice but to purchase yet another robotic device, this time the IBM technology solution.

The NCCS began writing all new data to the IBM Magstar robotic subsystem in November 1996. For a cost of nearly \$250,000 the NCCS upgraded the remaining five StorageTek silos to faster robotics (2.5x), in preparation for mounting more than 33,000 800 MB tapes for repacking onto the more dense, 10 GB tapes. In August 1997, requiring more tape mounts to support an average monthly rate of 1.5 TB of new data into the system, the NCCS installed IBM Magstar tape drives in its StorageTek silos, creating the ability to store 12 times the amount of data in a silo.

In March of 1998, with the creation of a specialized software function for making a duplicate copy of existing data on a remote backup silo while simultaneously repacking that same data on the 10 GB Magstar tapes, the NCCS began in earnest to migrate off the 800 MB tapes. This migration/repacking and duplication effort is taking place simultaneously on the same production system that is still acquiring data at the rate of over 1.5 terabytes of new data a month and transferring over 200 GB of data into/out of the system daily. After four months of concentrated effort, the NCCS has been able to free up 3,500 800 MB tapes with nearly 3 TB's worth of data. To complete this migration effort of 35,000 800 MB tapes will require over 3.5 year's worth of processing. With the life expectancy of the IBM Magstar 10 GB tapes to be 10 years, the NCCS would have finished migrating 35,000 800 MB tapes by the year 2002, only to begin again the process of migrating off the 10 GB media at that point.

### **3. STORAGE GROWTH**

When we think of large storage requirements, we tend to consider the large repositories at the NASA, IRS, CIA (Central Intelligence Agency) or NSA (National Security Agency). Other major users are associated with agriculture, forestry, banking, medicine, weather, surveillance, military and entertainment. Mass storage requirements for many of these applications have grown dramatically high recently due to two major reasons. First, with doubling of clock speeds in chip technology and computer upgrades, there is an ever-growing demand for mass storage. Second, at least for earth science applications, there is a continuous demand placed on the mass storage systems in that doubling the spatial and spectral resolutions increase the storage requirements in a cubic fashion. For example, at the NCCS facility, currently we have been experiencing a growth of 1.5 to 2 TB every month, with an acceleration factor of 2 for every instrument mission increasing in data resolution. Such stories on increasing demand are becoming common across a variety of applications.

#### ***Satellite Remote Sensing***

Geospatial Information Systems (GIS) are the new horizon in Earth exploration (Reported at AFCEA TechNet 96 panel entitled, "Space Segment Data Providers -- U.S. Private Sector Sources") with 10,000 commercial remote sensing systems in 95 countries (ibid. reported by L. E. Jordan III, president of ERDAS Inc.). Timeliness is critical because old information is bad information. Products set to launch in the next three years will downlink intelligent information of two-, three- or four-dimensional images that a person can fly through, navigate through, query via the television, telephone or computer in real-time to stations world wide (ibid). Data storage requirements for this kind of applications will be on the order of near-Terabytes per day.

#### ***Banking and Finance Industry***

The government will be forced by an unnoticed provision of the recent budget compromise to abandon paper checks almost entirely before 1999. All 1 billion government checks, including Social Security, Medicare and Medicaid, will have to be made through electronic transfers. When compared to the check processing volume load at Chase-Manhattan, which processes 150,000 checks daily, we can see that this will place an enormous burden on the imaging system for storing the checks digitally. A recent study predicts that the number of consumers who bank on-line will more than double to nearly 2 million by the end of this year, from 754,000 in 1995. By the year 2000, estimates are that there will be some 13 million on-line banking customers (A study by Jupiter Communications LLC in New York reported in Investor's Business Daily 10 Sep 96.). The recent announcement that 15 banks will team with IBM to offer electronic services is the first signal that banks are ready to change the way they do business. All of this should result from an aggressive move by the banking industry to reduce operating costs.

### ***Civilian Government Applications***

The National Archives and Records Administration (NARA) is in an unenviable plight. Just the cost of space to store the government's paper records required nearly half the agency's 1997 budget. NARA suffers from being caught in the middle of the national transition from paper-based to electronic-based information. To further complicate this issue, in the fall of 1997, NARA has been directed by Congress to implement their recommendations for the storage of electronically created documents. The policy had been to have agencies create a hard copy of the documents to store them for archival purposes until now. The new directive is to develop a plan to quickly migrate the archives to a digital media.

### ***Military Applications***

The Pentagon established the Advanced Concept Technology Demonstration (ACTD) candidates for fiscal 1997, with information technologies for the battlefield leading the list. The ACTD program is designed to field advanced technologies quickly. For FY 1997, Rapid Battlefield Visualization with a 3-D view of the battlefield is the most important of the 18 projects. Commensurably, 3-D graphics performance will be increased 10 times over the next three years. Other military applications will develop as the data begins to gather from the effort in place to care for the nuclear test stockpile. For example, the Air Force Rome Labs has 350 TB on line and 20,000 TB (20 petabytes), of intelligence and imagery data off line. Since nuclear weapons can no longer be tested above or below ground, the explosions will be simulated on a Teraflop machine being built by Sandia, Los Alamos and Lawrence Livermore. The data will be a national resource. One of the agencies involved is the Department of Energy. In concert with the support of this project, DOE has been a contributor to the development of the 1 TB cartridge optical tape system coming from LOTS Technology.

### ***Medical Informatics***



The emergence of medical informatics as a new discipline is due in large part to advances in computing and communications technology, and to an increasing awareness that the knowledge base of medicine is essentially unmanageable by traditional paper-based methods. Tele-radiology is a clear example of strong growth in this direction. For instance, to store a 2048 X 2048 mammogram data with 16-bit resolution, a typical hospital incurs a data rate of 10 Gbytes per day for mammograms alone. The FDA stipulates that the data can only be compressed using loss less compression techniques, which hardly provide a compression level of 5-6 for mammograms. A quick analysis shows that a typical hospital attempting to store MRI, PET scan, mammograms and X-Rays can potentially experience a data growth rate of 1 TB a month. Unlike any other field, for clinical experimentation, this data needs to be at least near-line and should be preserved forever.

### ***World Wide Web: E- Commerce***

The emergence of the World Wide Web as the corporate and consumer information access medium of choice will combine with enhanced database software and new server and browser technologies to enable powerful data management services. Millions of people are traveling the Internet and using web browsers to retrieve information. The Internet soon will be the de facto medium through which people retrieve information. Over 60 million PCs are already Internet-ready and that number will grow to 265 million by 2000. Three years from now, over 82 percent of corporations expect to support fully-connected internal networks.

How does this translate into data storage? First, consider that data files of the future will contain images at the least and probably audio or video clips. This will translate into files that range from 1 to 1,000 MB in size. Second, assume that by 2000 we have 10 Mbps service to the home. Third, assume that a user wishes to retain the content of some information. Internet traffic through the five largest network providers and metropolitan-area exchanges exceeded 250 TB per month in 1997. Estimates put the Web at 50 million pages on 475,000 sites and the entire Internet at anywhere from 1 to 10 TB of data with typical Web sites averaging 10 to 20 MB of data. Will the choice be to store the data on-site or to leave the data at the web site? The question is how much storage is available locally verses how long does it take to download the data. Couple this with the most common problems associated with the Internet: speed (80.9 %), organizing retrieved information (33.6 %), and finding information (32.4 %). Solving this equation will be an individual matter (user-pull) contingent on affordable systems (technology-pull). The eventual impact the Internet will have on data storage will depend upon the size of individual data files being transferred, the speed with which those files are transferred, and the cost of data storage.

### ***The Scientific Community***

Several scientific applications have experienced tremendous data growth. For instance, at the Thomas Jefferson National Accelerator Facility, on-line data increased in 1997 by a factor of 5 from 100 to 500 GB -- this amount should double in 1998 to 1 TB and double again in 1999 to 2

TB. Over the same time period, near-line data increased by a factor of 30 from 5 to 150 TB -- it should double in 1998 to 300 TB and quadruple in 1999 to 1.2 PB. At the same time, on-line data should increase by a factor of 20 from 1996 to 1999, near-line data should increase by a factor of 240. This represents just one of the hundreds of national and international laboratories and research institutions, which will be storing their data electronically.

In addition to total overall data storage requirements increasing, there is an upward shift in the hierarchy of currently stored data. More data is being stored, and more data is moving from off-line to near-line; and from near-line to on-line. Some data bases will grow by gigabytes while others will grow 1 to 3 terabytes daily. Some data bases will grow to petabytes over time and data bases with 30 year retention requirements could grow to 100's of petabytes. All of these data bases will demand a user-friendly software to accommodate potentially hundreds, thousands or tens of thousands of user transactions per day without heavy operator interaction

### ***Overall Market Growth***

The worldwide data storage market in 1996 is approximately \$100 billion including media and devices. This includes data, audio and video applications, 40% of which are in the US. The computer data storage portion in 1996 is about \$65B. The digital magnetic tape storage system market was about \$7 billion in 1996. The magnetic hard disk market has been growing at the 25% level for several years despite large increases in bytes per drive and modest increases in unit volume (price per byte is declining rapidly).

There were 6,454 large tape libraries in 1994. With the rapid introduction of many smaller libraries the total number of libraries is expected to increase to about 90,050 by the year 2000 for a compounded growth rate of 55 percent (1995 Freeman Associates Report). According to Fred Moore of STK in the 1995 yearly market report, over 600 PetaBytes will be stored on-line in computer-readable formats by the year 2000. While this sounds impressive, it accounts for only 5% of the expected world's data by that time. According to LargeStorage Configurations, by the year 2000, each of the Fortune 500 will have over a PetaByte under management; and from Network Computing magazine (Fred Richardson, "Windows NT Storage Management," Computer Technology Review, Vol. XVI, Number 11), the total storage at a typical Fortune 1000 site is projected to escalate from just 10 TB last year to 1 PB by the year 2000.

Only a decade or so ago, megabytes of storage were considered more than adequate for most large data centers. Then for several years gigabyte sized systems offered apparently unlimited storage capacity for almost any conceivable project. More recently, Terabyte systems have become available for large data requirements. The growth of data storage needs has outstripped the impressive advancements in computer system performance and appears to require even larger and faster data storage systems. Recent articles in Defense Electronics and the Journal of Electronic Defense describe the growing need for petabyte storage capacities. A typical large data base system for US Government agencies is expected in the next 5 years to accept 5 terabytes per day, be able to maintain 300 Terabytes on-line (within 15 sec to 1 min. access time), and be able

to archive from 15 to 100 petabytes. These large government systems are expected to presage an analogous requirement in the commercial sectors.

## **4. ASSESSMENT OF TECHNOLOGY**

Undoubtedly the mass storage systems have been experiencing impressive growth in technology in many areas including storage density new forms of storage media, though relatively a smaller level of progress has been noticed in other areas such as access speeds. To provide a comprehensive perspective, we begin with the storage technology of the past.

### **4.1 Storage Media in the Past:**

Goddard opened its first science computing center in the basement of building 1 in July of 1960 with a IBM 7090 using 200/556 bits per inch 7 track tape drives these tape drives were upgraded to 200/556/800 BPI capability. In 1965 the IBM 360 systems were introduced to the Goddard Community. Tape compatibility was maintained with both the 7 track tapes and 9 track 800/1600 bits (or bytes) per inch tape being used. Toward the end of 1960 it was noticed that tapes which had not been used for several years had deteriorated significantly to the point where in the worst case the media was dissolving back into the original components. This process damaged the tape drives as well as the loss of science data that had not been backed up in many cases. In the 1970's the tape subsystems were upgraded to STK 1600/6250 BPI drives with much higher reliability. The computing center in conjunction with NSSDC led an effort to have the science community voluntary move the data off of the older media to the new tape system with mixed response. Much was saved in this process other data were lost.

Magnetic tapes have been around for about 50 years and still are the preferred digital storage media for long term archiving. In the last decade, removable disk storage has become another contender for long-term storage as their costs have come down with increasing capacities. The three probable governing factors for the choice of magnetic tapes are (i) cost in terms of bytes per dollar, (ii) survivability and (iii) technological evolvability to increased storage density per unit volume. Its major deficiencies result from imposing sequential access thereby limiting scaleable I/O transfer rate increasing in performance as density increases. Magnetic disks of varying size, optical disk, CD- ROMs, optical tapes and other media have been tried by various communities but have not yet gained wide acceptance over magnetic tapes. For the last 25 years of electronic computing, seven track magnetic tapes of 800 bits per inch and 2400 feet of length were the standard format for the data storage industry. Standards have been put into place by most major manufacturers who built tape controllers. The tape media has evolved over time to 9 tracks with 1600 BPI and then to cartridges with more tracks and denser storage. Today a single cartridge can hold 10 GB on 128 tracks compared with the 1960 vintage tapes of less than 1 MB capacity, a 10,000 fold increase in 30 years.

## 4.2 Current Storage Media

The NCCS is currently using 3480/3490, Magstar and Redwood tape storage media in climate controlled autonomous access silos. This system eliminated human dependencies ( i.e: mishandling; tape label errors;)from the media managing process and it was believed that it would eliminate the largest source of error from our mass storage system. Operationally, active long term data storage systems have a number of challenges that can introduce some potentially catastrophic errors. The initial 3480/3490 technologies were designed to be used on IBM mainframes using the IBM flagship operating system, MVS. This system has capabilities not currently available on the more cost effective Unix systems, such as tracking temporary failures and notifying operations when a tape needed to be replaced before it failed permanently. We now experience as many as two hard failures a month with a total activity of 30,000 tape mounts per month on average. Depending on the need of the community some tape media may be much more active than others, this also is harder to track on the present operating system. Tape drive maintenance is critical, a failing or out of adjustment drive can introduce permanent failure to new data or destroy existing data media. Good environmental control and protect from physical damage such as leaking roofs or failing sprinkler pipes is critical. Coordinating multi-vendor shops to assure that IEEE standards are met and interpreted in the same way is necessary to keep the data integrity. Instances of differing interpretation while not common can be problematic. The process of unobtrusive (to the science Community) upgrades introduces a new level of complexity necessary to meet the growing, near exponential demand for storage. The NCCS presently has three types of media 3480/3490, MAGSTAR and Redwood, each with several densities or size of tape available. The 3480/3490 has the capacity of 200 MB, 400 MB, or 800 MB on a tape, presently most are 800 MB. The Magstar is 10 GB. The Redwood is 25GB or 50 GB per tape. The age of the media and activity is:

Type	Age	% of Mounts
3489/3490	4-5 years	45%
Magstar	1-2 years	45%
Redwood	1 year	10%

The NCCS computing facility is backing up the active data on the Redwood equipment located in the EOS Building 32, 1/2 mile from the main computer center in building 28.

The NCCS has the problem of replacing the HP/Convex 3830, a Y2K nor compliant system now 5 years old to assure that the user community does not experience any impact to their work. The NCCS is evaluating 4 other vendors' data servers and will be evaluating several alternative operating systems to replace the Unitree system that currently manages the data. Computing center management has a good track record for benchmarking upgrades. Because other computing facilities are in the process of doing similar changes the possibility of problems is reduced greatly. For active growing data storage facilities the evolution of media, supporting hardware and

software has to continue in order to support the community dependent on it. The facility must maintain its inventory of people skills to enable this process.

### **4.3 Future Storage Media**

Data storage technology is now allowing higher density storage with faster speeds and lower costs than ever before. At the same time, more data storage technologies options exist now than ever before. Types of media technology can now be chosen to best suit the applications to which they are intended. The choices are exciting but for our purpose we will focus on magnetic tapes and optical media. As pointed out earlier, magnetic discs are a viable option for long-term storage if the costs were to come down a factor of 10 or more. It is highly conceivable that this will occur in the next few years, since this technology benefits from leveraging the mass market for PC's.

#### **4.3.1 Short-term Technology Outlook**

The infusion of new technology into linear tape systems provides new life and a continued migration path to higher area density and capacity. Track-follow servos and MR (thin film magneto-resistive) read head technology are easily applicable to all forms of linear tape technology. The 3480 was the first to introduce MR head technology to any magnetic data storage system, nearly a decade ago.

Choosing a magnetic disk drive (MDD) has never been easy. There is a constantly changing product line brought on by rapid technological advantages. And while the dizzying pace of MDD technology makes products "obsolete" sooner, it also provides a myriad of choices. You will always wish you had more storage space. We expect these technology improvements to continue for the next commercial market.

Optical tape holds forth the promise of very large capacity in a single media unit (cassette or reel), fairly great information volume density, and very low cost per MB. The generally long access times have been improved and are really not a problem in some sequential types of applications, such as entertainment presentation (audio and video) and data archiving.

Compared to magnetic tape, optical tape offers orders of magnitude greater capacity per media unit. In fact, capacity comparisons with virtually all other storage media are impressive. One DOT cassette is equivalent to 250 to 500 3480 cartridges. One reel of tape for one system holds the data of 5000 conventional computer tapes, 100 12 inch instrumentation tapes, or 2000 5.25-inch optical disk. Optical proponents also feel that their media will last much longer in archives. Tests appear to confirm advantages of optical over magnetic tape with regard to aging and sensitivity to environmental deterioration. Use of the popular 3480/3490 format for the LaserTape system will certainly be advantageous for both users and manufacturers trying to establish a marketplace. Since the manufacturing cost for optical tape should approximate that for magnetic, whereas the optical storage density is far greater, the media cost per capacity unit

will be far lower. In fact, optical tape media cost appears to be lower than that for any storage competitor at this time.

Optical tape technology is not yet demonstrated to be commercially competitive for our purpose. It is worth sharing some previous experiences in waiting for new technology to arrive. In 1984, large scale government systems managers waiting for new technology but under pressure to provide more storage space immediately for use by their high speed processor customers received a gift in terms of an announcement several days before Christmas. A major manufacturer of storage devices withdrew from the optical disk market after promising for years the delivery of these optical disc storage systems. These systems had been under beta testing at super computer centers such as NCAR for several months when the vendor withdrew the technology from the market place because of the difficulty in manufacturing the large 14 in. media platters. The failure of this promised upgrade path delayed the introduction of optical disk systems for large scale use by 5 or more years. However, optical disk drives find their biggest success in the international industry. For comparisons not all the numbers provided by the manufacturers are believable and commercial success at this point is at the PC level and not suitable for our purpose.

Generally speaking, the 12-inch optical disks should hold 60 GB in a dual-head format while the single-head 5.25-inch are scheduled to reach 10.2 GB. On the small end, the 2.5-inch optical has caught up with the 3.5-inch optical and each hold 640 MB. Pit depth modulation technology will triple the DVD technology after it tops out at the turn of the century. Tapes will reach 1 TB in a 3480-type format using laser optical media. Magnetic tapes can reach 660 GB in a single large-format cassette, while the small 4 mm may find a comfortable market niche at the 12 GB level. The 8 mm from Sony and Exabyte will top out at 100 GB and 300 GB respectively. Other improvements include partitioning and the addition of the MIC chip.

#### **4.3.2 Long-term Technology Leaps**

Over the next 15 years, scientists will narrow the width of lines etched into semiconductors to less than one-tenth of a micron, meaning that electrical signals running through those circuits will contain so few electrons that adding or subtracting a single one could make a difference in the computer's functions. To control the movements of very small groups of electrons, researchers are developing quantum dots that can corral rambunctious electrons, allowing them to escape only when zapped by a precisely sized boost of energy from outside. Such "quantum confinement" could lead to tiny, very high-powered lasers that could make it possible to store 15,000 times more data on a computer chip the same size as those produced today.

Current nano-technology storage capacities available are 125 MB/in<sup>2</sup>. The National Storage Industry Consortium (NISC) is hoping to have production units at a level of 1.25 GB/in<sup>2</sup> by 2000. This should be reachable because there has been a 60 percent growth per year for the last 5 years and if the trend continues, by 2005, the capacities will reach 12.5 GB/in<sup>3</sup> and 125 GB/in<sup>3</sup> by 2010.

Holographic storage could reach 125 GB/cm<sup>3</sup>. In principle, this technology is fairly established. Unfortunately, there are no good recording mediums, the holograms are hard to stabilize, and the laser beam is hard to steer. The file maintenance system is difficult and as more holograms are added to the system, the older ones fade. Finally, holographic materials offer a low defraction efficiency. Emerging 3-D two-photon photochromatic storage technologies are being funded by Rome Laboratories and developed by the University of Southern California, San Diego that will read pages of data stored in a cube format. 3-D storage cube technology cannot come into reality until some of the components reach maturity. The impediment to implementation has been that no one has found a suitable medium that is sensitive to light so that it does not require a high powered laser to etch the hologram yet stable enough that it does not deteriorate. In addition, the medium should also have a high signal to noise ratio. In short, the medium material for holographic storage is a physics/chemistry problem yet to be solved.

Presently, the Earth and Space Data Computing Division (ESDCD) at NASA/GSFC is sponsoring research in this area at the University of Maryland, Baltimore County (UMBC). Specifically, UMBC Professor M. Hayden has recently been successful in storing a rather permanent

hologram in a polymer matrix which he has designed. This arrangement yields a permanent storage because the 675 nm recording light beam actually breaks some molecular bonds in the dye. The medium is said to be photochromic. Subsequent illumination with the same beam for many hours does not alter the hologram, and, in this sense, the image is termed permanent. Hayden has also been successful in temporarily storing a hologram in the same polymer matrix which is sensitive to a longer wavelength of 800 to 900 nm. This medium is said to be photorefractive and the recording light will erase an image in a few tens of minutes. Thus, by virtue of this polymer research, he is now able to both permanently store images at the shorter wavelength and also store

erasable images at the longer wavelength in the same medium. Another holographic storage problem being sponsored by the ESDCD is that of multiplexing data in the medium; this is a systems engineering and data encoding problem; we know we can do it, but it is a search for the best to optimize the process and achieve an acceptable bit error rate. In this area we are using the NASA Small Business Innovative Research (SBIR) program to fund Optitek, Inc. who are exploring an innovative, proprietary parallel optical architecture to angularly multiplex holograms into the same photorefractive crystal volume which would be a design breakthrough if successful. Optitek will be delivering a holographic storage device to the ESDCD for evaluation in the near future. Additionally with SBIR funding, both Optitek and Arizona State University are studying how to optimize multiplexing schemes with regard to output signal-to-noise by a global encoding of the stored data. Another important area of research of critical components necessary to holographic storage is that of advancing the technologies associated with input/output devices such as spatial light modulator and charge coupled devices (CCDs). This technology development is already well funded by industry which will use these devices for other purposes. For example, spatial light modulators are the basis for flat-screened TVs and CCDs are used in CamCorders and all sorts of scientific instruments.

Electron etching on HD-ROM (High Density - Read Only Memory) storage media was demonstrated by Los Alamos National Laboratory (LANL). The technology incorporates the ability of liquid-metal focused electron-beam milling in an ultrahigh vacuum to achieve feature and hole aspect ratios on the sub-micron (nanometer) scale. The method developed by LANL uses a Focus Ion Beam (FIB) micromilling process for storing data in one of three different formats: 1) binary (image) at densities of 2.9 GB/in<sup>2</sup>, 2) alphanumeric (ASCII) at optical or non-optical densities, and 3) graphical (video) at optical and non-optical densities. This media may carry all three formats and thus increase the utility of the recorded media. Using an electron beam that is 150 billionths of an inch wide to record on a 10 micron-thick steel tape could result in a storage capacity of 50 TB/in<sup>3</sup>. If a 3  $\mu$ m tape were used, the capacity goes to 190 TB/in<sup>3</sup>. The projected migration path is 100 GB/in<sup>2</sup> by 1999 and 400 GB/in<sup>2</sup> by 2003. The practical limit is reported to be 1.4 TB/in<sup>2</sup>. Read and write speeds are 2 GB/sec and 4 MB/sec respectively with 200 millisecond access time. This technology is now being promoted by Norsam Technologies ([www.norsam.com](http://www.norsam.com)), and the first demonstrated product was shown at the AIIM show in May 1998. A 2-inch metal disk had been etched with 90,000 analog images. Each page was 100 x 200 microns and each letter was 1.5 x 2 microns. The image could be viewed through a powerful microscope developed by IBM, and viewed on a computer screen, printed or transmitted. This shrunken version of microfilm represent 9 four-drawer file cabinets or 4.5 GB. The next product should be ready by 2000 and offer 165 GB on a single 5.25-inch platter.

Biological solutions are on the horizon as well. Researchers are considering storage mediums based on the phosphorescent properties of Halobacterium found in the lagoons off of San Francisco Bay or in the iridescent covering of jelly fish. Much like promises of organic computing, it remains to be seen how far these promising technologies are commercially viable.

## **5. RECOMMENDATIONS**

The mass storage arena is facing a crisis of severe proportions that merits a national level attention. Considering the impact of mass storage technology in promoting information growth in globalized economies of the world, U.S. needs to maintain its leadership and technological edge in this field. A compromise in its investment on mass storage technologies could spell disaster not only for the scientific community but also could result in loss of critical information to the commercial sector. Industry invests in new technology based on incremental demand-supply relationships and promise of high returns in relation to the investment. The Government has a clear near term need that could lead to a long-term demand and thus should play a critical role similar to that played in high performance computing, i.e., of a catalyst and a ready buyer to stimulate growth.

### **5.1 SHORT TERM RECOMMENDATIONS**



Considering the impact of mass storage technology in promoting information growth in global economies of the world, U.S. needs to maintain its leadership and technological edge in this field. This technology assessment leads to two sets of recommendations to achieve this goal. From a short-term perspective, the Government should:

1. Develop a long-term strategic plan to maintain the U.S technological edge in high capacity storage systems.
2. Require agency CIOs to collect annual reports internally on data growth and related statistics (mounts, access, media, etc.) from all its internal organizations managing more than a TB of data. This information should be used to identify storage media at risk and estimate future requirements within the agencies but also across the government.
3. Encourage agency investment in hardware and software storage system migration to safeguard large volume data through effective, timely, lossless transfer of data from aging systems to the new dense media with minimum impact on the production systems. Encourage national exchanges of technical innovations and sharing of software performance by large repositories of data holders and support national storage standards.

## **5.2 LONG TERM RECOMMENDATIONS**

From the long-term perspective, the Government should:

1. Establish a national goal to develop the next generation mass storage technologies towards building an affordable modular exabyte ( $10^{18}$  byte capacity) system. This system should be scalable, distributed, evolvable, and reliable.
2. Forge a strategic alliance of government agencies, industry and academia to maintain a vibrant research program in support of these critical storage technology goals.
3. The Government should serve as a catalyst by promoting storage technology application demonstrations and test bed facilities that help vendors integrate software and hardware to stress end-to-end performance of their systems. The Government needs to develop opportunities for the industry to foster growth through vendor-clientele partnerships.

In summary, the untamed growth in information storage and exchange is a harbinger of a looming crisis similar to the Y2K problem. This problem deserves an early attention to preserve all the valuable historic data sets in science, medicine, commerce, and industry. This report will be continually revised to include additional studies and results from other agencies as well as the industry.