WORM (Write-Once-Read-Many) technology is used in many applications because of the integrity of the data and the accepted legal admissibility of files stored using the technology. Fundamentally there are two basic kinds of writeable media today, sector based and non-sector based media. An example of sector-based media is Ultra Density Optical. UDO technologies are reasonably fast, store 30 GB per disk and are available in both write-once and rewritable formats. Examples of non-sector based media are CD-R and DVD-R. DVD-R is also reasonably fast however the capacity is only about a third of that of UDO.

DVD-RAM is a sector-based optical medium with the same capacity of DVD-R, however it is not known as a write-once technology. So does it match up to other archive media?

Pegasus Disk Technologies
2333 San Ramon Valley Road
San Ramon, California 94583
Phone (925) 314-1800
Fax (925) 314-1801
www.pegasus-afs.com

Lets first discuss price. Sector-based media are more expensive than DVD-R. However, now that UDO is available, the difference in cost is not great when you compare cost per MB. That said, cost still seems to be the big attraction to DVD-R. When considering the importance of archived data one would think this is a small consideration. Why pay more for sector-based media solutions? Ask yourself, how valuable is your data archive? The data archive is likely to be one of a company's biggest corporate assets; therefore can you afford to loose data for any reason if it's avoidable? Now, ask yourself what is the implication of lost archive data? If your in a regulated industry, the penalty could be just fines, if your in a public company, however, the implication could be jail time for executives depending on the data.

Optical File System

One current standard file system for authoring optical is UDF (Universal Disk Format). ANSI, ECMA, and ISO set out to develop an international universal format for information interchange. Initially known as Non-Sequential Recording or NSR, it was developed as an information interchange standard to succeed the CD standard 9660. Thus the International Standard ISO/IEC 13346 was created. The standard is different from ISO 9660 in that it was created to write as well as read information among different operating systems. The initial plan was designed to support non-sequential media and would enable information to be exchanged over a variety of operating systems. It was expected to gain acceptance based on the popular CD-ROM based technology. OSTA set out to develop a specification as a subset of the NSR standard that could be brought to market quicker with the ability to address data interchange between systems. Thus, the Universal Disk Format was created.

A consortium of leading companies in the optical storage industry developed UDF. Most of these companies are the driving force behind the various DVD standards today. UDF is based on the ISO/IEC 13346 standard. It is in fact a subset of that standard and is fully compliant with ISO/IEC 13346. UDF was designed as a follow-on standard to support DVD based on CD-ROM, CD-R and CD-RW. It was conceived that through the use of UDF, DVD based content can reside on the same disk and be accessed by a wide variety of computer systems as well as the consumer DVD player in the home.

No Sector vs. Sector Based Media

DVD-R, along with all the CD technologies, is a sequential-based medium that spirals out from the center of the disk. Starting from the center of the disk, each bit is written next to the last bit written in a sequential order, regardless if the information is related or not, not unlike tape. This is done, because there are no sectors on the media. Unlike tape the data can be read in a random sequence once written, however, the write process is sequential.

With CD technology, disk at once or session writes were not an issue because the disks were only 650 MB, a manageable amount of data for most applications. Now, however optical technology can store up to 15 GB per side. This is, by most standards, quite a lot of data to have to cache or to have to fill prior to closing the media out.

DVD-R uses a session-based write format that requires the disk stay in "open session" so long as the user would like to continue to write to it. This means that if the application creates files smaller than about 4 GB, which is the majority of data out there today, and the user would like to fill the disk before closing the session, the user has 2 choices.

The user can cache information on a hard drive until a capacity threshold is met and then write all the data at once. A problem with caching prior to writing to the disk is, in many applications hard drive caching does not meet various U.S. and International regulations governing applications such as financial, securities and medical records storage. The other option is to write to the disk, opening new sessions as necessary. Once the disk is full, the disk must be "closed". Once the disk has been closed, it can no longer be written to and is now read only.

Both situations can be a problem in an archive environment. In the first case, data is not truly archived until the data is on the archive media. Caching prior to writing to the disk in many cases does not meet the regulations governing write-once applications. Data can, in some cases, stay in cache up to months before enough data is collected in order to fill a disk. The data is still vulnerable while on the hard drive and even during the writing process. In the second situation, all the data on the disk is quite vulnerable while the disk is "open". As in the caching situation, it may take several weeks or months to fill a disk, and as a result, the disk will have to be kept "open". While the disk is in an "open" session mode, it cannot be universally read outside the system that created it. Therefore the information cannot be exported or shared outside the network. Additionally any time during the process a bad write can occur and corrupt the entire disk. If this happens, the disk becomes unreadable. A bad write can occur on any session of an open disk for a number of reasons--dust, dirt, fingerprints, scratches, or even a defect during the manufacturing process. Any one of these possibilities are issues to be considered. It must also be noted at this point that all DVD-R media are unprotected and exposed to the open environment – no cartridges.

In addition to the "open session" issues, there are performance issues to be considered as well. Recordable media has directory structures describing where data is stored on the media along with other important information. In the case of DVD-R media, the directory structures have to be updated with each new session, however, because the media must be written in a sequential manner, with each new session, a new updated directory is written. This has the effect that the drive must search several directory locations in order to find the data on the disk. With each new directory entry, the read performance goes down as the drive must skip around to find the relevant directory information containing the information necessary to retrieve the data the user is looking for.

In the case of sector based UDO or DVD-RAM, an area of the disk can be set aside for the directory information. This allows all appended information to be in the same area on the disk ensuring the read speeds to maintain maximum performance throughout the life of the media regardless how many writes are necessary to fill the disk.

There also is a problem with appending files. If a file is appended on DVD-R media, the new file must be written in the next open block in the sequence. This means that related files would not be next to each other; in fact, they may not even be on the same piece of media. This can be a real problem in a records management environment such as the medical industry, where patient's records are scattered over a number of volumes in a storage library, thus causing the library to jump from one piece of media to the next in order to collect all the data, should a doctor need to review the patient history. With sector based media, blocks can be allocated for future appending files, thus allowing related files to be all in the same area and on the same disk. When considering the media swap time of most libraries, this can be a very big time advantage, and can

enhance the reliability of the system, as the library does not have to do as many swaps on average.

Another area where sector based media is far more robust for archiving of critical information is media defect management. CD-R and DVD-R media have historically been unforgiving when it comes to recovering from a write failure. Users have reported frustration with having to queue up enough information to write a complete disk, only to face write errors during the disk imaging process resulting in corrupt media. This problem has fueled an attitude among users that the data stored on –R media is not necessarily safe and reliable. On the other hand, sector based media and the drives that write to them have much more extensive defect management software built into the on-board firmware to retry, map-around and manage any possible bad areas on the media to ensure no user data or directory information is written in those areas and the media remains accessible over long periods of time. Sector based drives have defect management as part of device interface. Vendors writing software for sector based media can store duplicate directory entries, map bad sectors out of existence, and retry errors when reported. However, probably most important of all and the biggest differentiator is, if a sector does go bad, the software controlling the drive and media continues to give the user access to the remaining massive amounts of information stored on that disk.

	Sector Based Media	Non-Sector Based Media
Sequential Writes	Should be written sequentially with WORM, however, blocks can be set aside for future updates	Must be written sequentially with no open block allocation for future updates.
Appending Files	Appended files can be put on allocated areas of the disk	Appended files must be written to the next open block on the disk.
Directory Structures	Open blocks can be set aside for directory updates	Directory updates must be written to the next open block on the disk.
Volume Spanning	New disks can be added to a volume set	Once the session is closed, no further surfaces can be added to the volume set.
Performance	No Performance issues based on the number of writes to the disk	Performance goes down as the number of sessions on the disk goes up.

When is WORM WORM?

Write-once is just that. WORM has been historically tied to the storage technologies such as optical disk because the first optical disks required a physical change on the recording surface that was impossible to undo. Today, most WORM hardware technologies are fundamentally rewriteable with an electronic write protect process through firmware in the storage device. The result is an archive system that will only allow a bit sector to be written once and then never changed again. The system includes hardware, media, software and procedures on how to handle data and media. In an archive system a bit sector on the media within can be accessed an infinite number of times for reads, but the data cannot be changed. If the file system used in the archive system is a rewriteable file system such as NTFS or UDF, data is treated in a different way than with a WORM file system. An archive file system provides no method of overwriting data or directory entries. While a rewriteable file system has to be told by the device that it cannot change data already written.

Magnetic disk and tape will always be fundamentally erasable because they are magnetic media. If the disk is subject to strong enough magnetic field, the media will become unreadable and need to be reformatted. This has been one attraction to optical storage. Regardless of the amount of magnetic field subjected to an optical disk, the data remains unchanged.

Additionally, file systems designed for magnetic disks do not work well with WORM optical media. Magnetic file systems can update information stored in the directory or data areas. WORM media does not allow this functionality by design. There have been software vendors that have modified versions of NTFS or NFS that work with WORM media, yet the fundamental problem continues to be file system overhead. These file systems must rewrite entire directories or files to new locations on WORM media, because the original directories cannot be modified.

Another crucial issue here is the migration to other operating system versions or platforms over the life of the data. While this is rarely considered when a new system is purchased, it becomes a significant concern for the user as operating system vendors evolve product over time or produce new systems, which often times change file formats on magnetic media. Even the slightest changes in the file system can cause difficulties for a user that has massive amounts of archive information stored on long-term removable media.

In most cases, WORM or other archive-based applications are transactional-based applications. Data or directory information should be able to be written at *any time* without increasing disk usage or overhead beyond the additional file space required for the data file itself and the small space required for the given entry to the directory describing that data. Data and directory information should be safe from intentional or accidental file changes or deletion. Sector based optical media, both UDO and DVD-RAM, provide this capability when combined with a file system designed specifically for WORM and archive functionality applications.

A Better Way to Archive

Pegasus is known for its file system and volume management software for optical disk technologies. InveStore® is storage management software that recognizes and supports many different third party and industry standard file formats in use with various forms of optical media. One such format written by Pegasus specifically to manage data that is write once in nature is the Archive File System (AFS), or Pegasus-AFS™. In use in Financial, Medical, Insurance, Transportation, Government, Law Enforcement and many other archive applications for the past 15 years, the Pegasus-AFS was designed specifically for the unique characteristics of optical media. Matched with UDO WORM or DVD-RAM sector based media, this creates a solid reliable high performance archive solution with top-notch security and data tracking mechanisms. To offer users of DVD and CD a more reliable and robust solution for DVD, Pegasus provides DVD-Archive™ for use with DVD-RAM based hardware. The same reliable file system that has been running with erasable-based optical media without failure for the past 15 years is now offered for DVD-RAM media and libraries, providing the user with considerable savings on the library hardware. Media are priced relatively the same as DVD-R for the same capacity of double-sided disks. And, the Pegasus DVD-Archive™ software supports mixing of single-sided CD-ROM and DVD-ROM media in the same library as the double-

sided writeable DVD-RAM media – all accessible under a single user interface and drive letter access! Each media or spanned volume set appears as sub-directories to the application or user.

The Advantages to Using an Archive File System in This Way Are Manifold.

The Pegasus- Archive File System™ (AFS) provides some special features that uniquely identify it as a true archive file system. Unlike more typical file systems, the Pegasus-AFS™ provides no method of overwriting data or directory entries. The AFS, by design, provides no ability to reclaim space. It is a fully historical file system that provides a complete transaction based audit trail of any and all changes to either the data or the meta-data. Most other so-called WORM file systems only provide snapshots of the system. The user must decide when they want to take a snapshot and then instruct the file system to do so. ALL changes to the disk between snapshots are lost with no audit trail. The Pegasus-AFS™ incorporates a highly specialized meta-data structure that allows the user to easily and quickly see, not only what has changed, but when it was changed as well. The user can determine if the attributes of a given file have changed, whether a file has been renamed or not (the user can quickly see both the old path\name and the new path/name) and what sections of the file data have been changed. The Pegasus-AFS™ provides all the features required to implement a complete change/document control system.

- Legacy data is often subject to technology innovation since, as technology evolves, older product is no longer supported. Therefore, older hardware needs to be upgraded or replaced. If the old product has write-once applications but the new product does not, the customer may not be able to access legacy data. By relying on a WORM file system rather than hardware, the data will be accessible no matter what the supported technology as the file system can manage data on all supported devices. InveStore® has the added benefit that it can manage them virtually from a single access point, whether it be a drive letter or IP address. This means that a DVD library can run under the same drive letter as the CD library and even the UDO library, all as a single volume or repository. All devices are accessible and manageable as if they were a single device, and individual disks or volume sets are accessible as sub-directories under the assigned drive letter. All data can be exchanged between devices by means of drag and drop. This ability to access legacy data allows for migration of data over technologies over time: from CD to DVD to UDO or whatever comes next. This translates to a lower cost of management over time.

- A write-once file system tracks EVERY change made to the media. The Pegasus-AFS does not allow overwrites of any data or directory structures on media. A complete audit trail is kept for all writes and appends made to the media. Utilizing the Pegasus-AFS, the system prevents any file(s) from being physically deleted from UDO or DVD-RAM media.

- Once data is written to a UDO or DVD-RAM disk utilizing the Pegasus-AFS, data will only be readable by networks with access to a system running the Pegasus software. If media are removed and taken to a system operating some other optical disk management software, the disk is protected either by the physical characteristics of the media (WORM), or by the setting of a hidden system bit, as on DVD-RAM media, which prevents any changes to the media outside of the Pegasus-AFS system control, such as formatting or unwanted writes or deletes.

- The Pegasus-AFS software is and will remain operating system independent and transportable. This means that disks written to under the first MS DOS and Novell-based systems are still readable today by our latest software. Applications of optical media boil down to long-term data retention. "Long-term" data often outlives the operating system where the optical storage hardware is installed. Also, optical disk media is removable, creating transportability possibilities from operating system to operating system that require an OS independent and consistent data access approach.

There is significantly less likelihood of changing the contents of a record once it has been stored on an archive file system. The smaller the chance for alteration, the fewer possible challenges to record integrity; the fewer challenges to record trustworthiness in the litigation context, the faster and cheaper disputes get resolved. A write-once archive file system is simply a more prudent decision for the safety and retention of corporate assets. This AFS provides the security that hardware vendors currently cannot support and responsible corporate consumers demand.

Pegasus Disk Technologies
2333 San Ramon Valley Road
San Ramon, California 94583
Phone (925) 314-1800
Fax (925) 314-1801
www.pegasus-afs.com

References:

Further information on optical standards can be found by contacting the following organizations:

ANSI and ISO Global Engineering Documents
15 Inverness Way East
Englewood, CO 80112
Phone: (303) 792-2181
Fax: (303) 792-2192

DVD Forum Office of the Secretary
Shiba Shimizu Bldg. 5F
2-3-11 Shibadaimon Minato-ku
Tokyo 105-0012, Japan
Phone +03-5777-2881
Fax +03-5777-28882
http://www.dvdforum.com/forum.shtml

European Computer Manufactures Assoc.
114 Rue du Rhone
CH 1204 Geneva
Switzerland
Phone: +44 22 735 36 34
Fax: +44 22 786 52 31
http://www.ecma.ch

Optical Storage Technology Association
19925 Stevens Creek Blvd.
Cupertino, CA 95014
Phone: (408) 253-3695
Fax: (408) 253-9938
Email: info@osta.org
http://www.osta.org