

An Emerging Network Storage Management Standard
Media Error Monitoring and Reporting Information (MEMRI) -
To Determine Optical Tape Data Integrity

Fernando Podio

Information Technology Laboratory
National Institute of Standards and Technology
Bldg. 225, Room B255
Gaithersburg, MD 20899
fernando.podio@nist.gov
Tel: +1 301-975-2947
Fax: +1 301-869-7429

William Vollrath

LOTS Technology
1274 Geneva Drive
Sunnyvale, CA 94089-1122
WilliamVollrath@msn.com
Tel: +1 408-747-1357
Fax: +1 408-747-0245

Joel Williams

Systems Engineering and Security
7474 Greenway Center Drive. Suite 700
Greenbelt, MD 20770
joel.williams@ses-inc.com
Tel: +1 301-441-3694
Fax: +1 301-441-3697

Ben Kobler

National Aeronautics and Space Administration
Goddard Space Flight Center, Code 505
Greenbelt, MD 20771
ben.kobler@gsfc.nasa.gov
Tel: +1 301-614-5231
Fax: +1 301-614-5267

Don Crouse

LSC Inc.
4211 Lexington Ave No.
Arden Hills, MN 55126
Don@lsci.com
Tel: +1 612-482-2348
Fax: +1 612-482-4595

Abstract: Sophisticated network storage management applications are rapidly evolving to satisfy a market demand for highly reliable data storage systems with large data storage capacities and performance requirements. To preserve a high degree of data integrity, these applications must rely on intelligent data storage devices that can provide reliable indicators of data degradation. Error correction activity generally occurs within storage devices without notification to the host. Early indicators of degradation and media error monitoring

and reporting (MEMR) techniques implemented in data storage devices allow network storage management applications to notify system administrators of these events and to take appropriate corrective actions before catastrophic errors occur. Although MEMR techniques have been implemented in data storage devices for many years, until 1996 no MEMR standards existed. In 1996 the American National Standards Institute (ANSI) approved the only known (world-wide) industry standard specifying MEMR techniques to verify stored data on optical disks. This industry standard was developed under the auspices of the Association for Information and Image Management (AIIM). A recently formed AIIM Optical Tape Subcommittee initiated the development of another data integrity standard specifying a set of media error monitoring tools and media error monitoring information (MEMRI) to verify stored data on optical tape media. This paper discusses the need for intelligent storage devices that can provide data integrity metadata, the content of the existing data integrity standard for optical disks, and the content of the MEMRI standard being developed by the AIIM Optical Tape Subcommittee.

1. Introduction

Network storage management applications provide for high data accessibility by means of backup and archiving, remote vaulting, file mirroring, hierarchical storage management, and storage device and server mirroring (Peterson [1]). The amount of stored digital data is increasing for certain applications to Terabytes (10^{12} bytes) and even Petabytes (10^{15} bytes). The amount of stored data per media unit is also increasing. It is currently possible to utilize magnetic tapes that store tens to hundreds of Gigabytes per tape cartridge and optical disks with capacities of tens of Gigabytes. It is anticipated that when optical tape products become available they will reach Terabyte capacities per media unit.

The demand for network storage management applications that provide high data integrity is also increasing. To preserve a high degree of data integrity, these network storage management systems can make use of intelligent data storage devices that incorporate sophisticated MEMR techniques. These techniques can verify the integrity of stored data and are able to provide the user early warnings of data degradation through the network storage application. MEMR techniques have been implemented in data storage devices for many years. MEMRI information is defined to be data integrity metadata derived from MEMR techniques.

In 1995, organizations involved in the developing of optical tape technology expressed interest in addressing data integrity issues related to this emerging data storage technology. In August 1995 AIIM formed the Optical Tape Study Group (OTSG) to promote discussions among users and industry on media characteristics, testing methods, data integrity and future standards for optical tape technology. The Study Group attracted broad industry and user participation (*e.g.*, Eastman Kodak Company, the Library of Congress, LOTS Technology, the National Archives and Records Administration, National Media Lab/Imation, National Storage Industry Consortium, NIST, Polaroid Corp., StorageTek, Systems Engineering & Security, Inc., Terabank Systems, and the Univ. of Arizona, Optical Sciences Center). This user and industry group expressed interest in developing a standard specifying MEMR techniques and the associated MEMRI metadata for optical tape. Data integrity issues, MEMR techniques and proposed MEMRI metadata have been discussed and analyzed in a number of papers submitted to AIIM OTSG (Podio [2]), (Silberstein [3]), (Podio [4]), (Manavi [5]), (Thibodeau [6]), (Vollrath [7] [8]).

The AIIM OTSG completed its work during 1997 by submitting to AIIM two project proposals for the development of ANSI standards: (a) a proposal for the development of an ANSI media interchange standard for digital data interchange in write-once read many times (WORM) optical tape cartridges (AIIM OTSG [9]) and (b) a project proposal for the

development of an ANSI standard specifying media error monitoring and reporting information (MEMRI) to verify the integrity of stored data on optical tape media (AIIM OTSG [10]).

In August 1997, AIIM formed the Optical Tape Subcommittee C21.3 working under the Storage Devices and Applications Committee C21. This Subcommittee started work in November 1997 on the two proposed standards mentioned above. MEMRI's project scope is for optical tape drives. However it is the current thinking of C21.3 to extend the scope of the proposed standard to *any type* of sequential storage media if the magnetic tape drive industry wishes to join in the development of this standard.

The paper is organized as follows. Section 2 focus on the need for media error monitoring and reporting techniques and discusses the power of error correcting codes currently implemented in data storage devices. Section 3 describes existing network storage management data integrity standards. Section 4 describes an emerging standard that specifies metadata to determine optical tape data integrity (MEMRI) and the means of transporting the media error information in a technology-and-interface-independent manner.

2. Media Error Monitoring and Reporting and Error Correcting Codes

2.1 Need for Media Error Monitoring

The need for data integrity metadata is apparent. Error Correcting Code (ECC) systems implemented in data storage devices are very powerful. They are designed to correct burst errors and can also easily correct random errors. The user is undisturbed by (but also usually unaware of) the level of corrections taking place. Depending upon the number, length and location of burst errors; however, some of these errors may be uncorrectable. When errors become uncorrectable data loss will occur.

In the absence of early warning indicators of data degradation and other media error monitoring capabilities, data storage devices cannot provide network storage management applications with the information required to maintain data storage subsystems with a high degree of data integrity.

2.2 Error Correcting Codes for Data Storage Devices

One ECC commonly used on rewritable magneto-optical disks is an interleaved ECC (ISO [11]). In this type of optical disk, each sector contains 512 bytes which are 5-way interleaved. (If the sector contains 1024 bytes the format is similar but a 10-way interleave is used.) This example of a 5-way interleaved data field format where each of the five columns represents a codeword has been described in detail in (Podio, Vollrath, Kobler [12]).

Figure 1 depicts how bytes are interleaved within a 512 byte sector. As shown, other bytes, in addition to the 512 user-data bytes, are stored on the media. The bytes are recorded onto the disk from left to right and from top to bottom ($SB_1, SB_2, SB_3, D_1, D_2, \dots, E_{4,16}, E_{5,16}$). Each of the five columns shown in Figure 1 represents a codeword. In this implementation, a Reed-Solomon code is used to calculate the ECC bytes, $E_{1,1}, E_{2,1}, \dots, E_{5,16}$ for the first codeword and so on. Interleaving the bytes in different codewords reduces the probability of a burst error on the media (n consecutive bytes in error) to produce uncorrectable data.

The sector bytes in this interleaved configuration represent the following:

- D_1 to D_{512} are the user data bytes.
- $P_{1,1}, P_{1,2}, \dots, P_{3,4}$ are the data management pointers (DMP) bytes. These bytes are information for sector reallocation (the link between defective and replacement sectors).
- FF bytes. Two filler bytes.
- C_1 to C_4 bytes. These are cyclic redundancy check (CRC) bytes. CRC bytes are computed from the user, DMP and FF bytes.
- SB_1 to SB_3 bytes. These are bytes that synchronize the data signal and the drive clock.
- RS_1 to RS_{40} bytes. These bytes preserve synchronization within the sector.
- $E_{1,1}, E_{2,1}, \dots, E_{5,16}$ The ECC bytes.

CODEWORD NO.			1	2	3	4	5	ROW NO.
SB_1	SB_2	SB_3	D_1	D_2	D_3	D_4	D_5	105
			D_6	D_7	D_8	D_9	D_{10}	104
			D_{11}	D_{12}	D_{13}	D_{14}	D_{15}	103
		RS_1	D_{16}	D_{17}	D_{18}	D_{19}	D_{20}	102
			D_{21}	D_{22}	D_{23}	D_{24}	D_{25}	101
			D_{26}	D_{27}	D_{28}	D_{29}	D_{30}	100
		RS_2	D_{31}	D_{32}	D_{33}	D_{34}	D_{35}	99
		
		RS_{33}	D_{496}	D_{497}	D_{498}	D_{499}	D_{500}	6
			D_{501}	D_{502}	D_{503}	D_{504}	D_{505}	5
			D_{506}	D_{507}	D_{508}	D_{509}	D_{510}	4
		RS_{34}	D_{511}	D_{512}	$P_{1,1}$	$P_{1,2}$	$P_{1,3}$	3
			$P_{1,4}$	$P_{2,1}$	$P_{2,2}$	$P_{2,3}$	$P_{2,4}$	2
			$P_{3,1}$	$P_{3,2}$	$P_{3,3}$	$P_{3,4}$	(FF)	1
		RS_{35}	(FF)	C_1	C_2	C_3	C_4	0
			$E_{1,1}$	$E_{2,1}$	$E_{3,1}$	$E_{4,1}$	$E_{5,1}$	-1
			$E_{1,2}$	$E_{2,2}$	$E_{3,2}$	$E_{4,2}$	$E_{5,2}$	-2
		RS_{36}	$E_{1,3}$	$E_{2,3}$	$E_{3,3}$	$E_{4,3}$	$E_{5,3}$	-3
			$E_{1,4}$	$E_{2,4}$	$E_{3,4}$	$E_{4,4}$	$E_{5,4}$	-4
		
			$E_{1,14}$	$E_{2,14}$	$E_{3,14}$	$E_{4,14}$	$E_{5,14}$	-14
		RS_{40}	$E_{1,15}$	$E_{2,15}$	$E_{3,15}$	$E_{4,15}$	$E_{5,15}$	-15
			$E_{1,16}$	$E_{2,16}$	$E_{3,16}$	$E_{4,16}$	$E_{5,16}$	-16

Figure 1. An example of a 5-way interleaved data field format where each of the five columns represent a codeword

The power of this interleave scheme and ECC is shown in the following example. If a burst error of length 16 occurs on the media (for example, from bytes D_{16} to D_{31}), each codeword (column) has no more than four bytes in error (caused by this burst error). The ECC can easily correct this burst error. This ECC implementation has the capability of directly correcting 8 bytes in error per codeword.

Different data storage devices utilize different ECC implementations. One standard (ANSI [13]) describes the use of two Reed-Solomon codes which operates between the drive data buffer and the read/write heads. The outer ECC appends 8 check bytes to each 128 byte data block and is defined as an RS (136, 128) code. The inner ECC encoder appends 8 check bytes to each 85 byte data block with two ID bytes. This is known as an RS (95, 87) code. Examples of other ECC implementations can be found in (ISO [14] [15]).

3. Network Storage Management Data Integrity Standards

3.1 MEMR – A Standard to Determine Optical Disks Data Integrity

In 1996 ANSI approved the only known (world-wide) voluntary industry standard specifying media error monitoring and reporting techniques to verify data stored on data storage media (ANSI [16]). This standard, developed under AIIM C21 Committee, “Storage Devices and Applications” and known as “ANSI/AIIM MS59” specifies a set of MEMR tools and the associated media error metadata to verify stored data on optical digital data disks. An accompanying Technical Report, known as “ANSI/AIIM TR39” (ANSI [17]), provides guidelines for the use of these MEMR tools and the associated metadata. An equivalent international standard based on ANSI/AIIM MS59 is also being developed. Users of optical digital data storage drives that implement the techniques specified in ANSI/AIIM MS59 can gather statistical information on media errors. This metadata allows them to highlight data integrity trends on particular selected disks or across their entire data sets. These techniques provide data recovery and media error monitoring tools with different levels of sophistication. ANSI/AIIM MS59 specifies two MEMR levels:

High level:

- A set of functional commands that are operating system (e.g., Unix, Windows™) and interface (e.g. SCSI-2, IPI) independent. In addition, this high level interface is optical disk media type (e.g, Write-Once Read Many Times or rewritable media) and size independent.

Interface level:

- An implementation of SCSI-2 (Small Computer Systems Interface – version 2) commands. These commands allow network storage management applications to include data verification tools at the SCSI level which are drive type and size independent, through the use of MEMR techniques.

Table 1 lists the types of data integrity metadata provided by MEMR techniques specified in the standard (ANSI/AIIM MS59). Table 2 describes the content of the technical report (ANSI/AIIM TR39). Standard metadata related to media errors allows network storage management applications to retrieve the same information even if storage subsystems consist of different optical disk drives. Decisions on the frequency of use of these tools are not specified in the standard.

Table 1 – Metadata Specified in ANSI/AIIM MS59 Standard

- Early warnings of an increasing number of media errors
- List of reallocated sectors
- Corrections above a defined media error level
- Total number of bytes in error, number of bytes in error per sector and maximum number of bytes in error in any sector codeword
- The uncorrected or corrected sector content
- Errors encountered reading header information such as the sector address, sector marks and synchronization signals
- Maximum length of contiguous defective bytes

Table 2 – Scope of ANSI/AIIM TR39

- Interpretation of the metadata provided by ANSI/AIIM MS59
- Statistical sampling methods
- Automation of media testing
- Graphical methods for media error representation
- Use of error distributions and statistical models to evaluate data integrity
- A summary of the ANSI/AIIM MS59 command set

3.2 Associated Industry Efforts

Two additional industry efforts are identified: (a) a proposed standard for tape drive management developed by Hewlett-Packard (Hewlett-Packard TapeAlert™ [18]); and (b) Imation's Media Performance Manager (Imation [19]).

TapeAlert™ is a tape drive status monitoring and messaging utility that makes it easy to detect problems, which could have an impact on backup quality. To reach agreement on an enhanced specification (Gold [20]), Hewlett-Packard has formed an industry Working Group, incorporating representatives from major tape drive and backup software companies. Imation's Media Manager is a storage management software product that

allows system administrators to perform data storage management from the user's desktop in real-time. Media Manager was co-developed by Imation and Sterling Software.

4. MEMRI – An Emerging Standard to Determine Optical Tape Data Integrity

4.1 Scope and Objectives of the Standard

This standard is currently under development in AIIM Optical Tape Subcommittee C21.3. The current thinking of C21.3 is to establish an ANSI/AIIM standard specifying a set of MEMRI metadata to verify stored data on optical tape media and the means of transporting the media error information in a technology-and-interface-independent manner. The high-level interface approach proposed in this new standard is independent of the host operating system and the interface between the storage device and the host. It is also media type and size independent and it can be applied to systems that use any optical tape media size and type.

The draft standard (AIIM Optical Tape Subcommittee [21]) being developed specifies:

- A high-level interface that will allow a drive to indicate to a network storage management application whether a given media should continue to be used or should be replaced by fresh media. This information is an early warning of data degradation. The supplier of MEMRI compliance (*server*) can choose to offer this recommendation either with or without defining the underlying chain of reasoning that generated the conclusion.
- A standardized means of communicating the information that led to the drive to recommend media replacement. If a supplier chooses to define the underlying chain of reasoning that generated the early warning of data degradation, the MEMRI Standard supplies a standardized means of communicating the analyses that were used to generate the recommendation.
- The standard means of describing and accessing through a high-level functional interface, standardized or vendor-specific drive accumulators and registers that store media error information.
- A high-level standard set of metadata elements that represent the content of drive accumulators that store statistics about media errors.

Currently, the draft standard specifies transport functions (query, response, attention and clear attention) and the query layer (data types and volume-related queries). The transport functions and the query layer are briefly described below.

4.2 Transport Functions

The transport functions are the means by which the *client* communicates with the *server*. Standard communications needs are transacted across a given interface. The three proposed transport functions are:

- **Query.** Provide for communication of *queries* from the *client* to the *server*. These queries will consist of a string of characters taken from the *printable ASCII characters*. The client may send a query at any time.

- **Response.** Provide for communication of the *responses* to the *queries* from the *server* to the *client*. These responses will consist of strings of characters taken from the *printable ASCII characters*. The server must generate exactly one single response to each query sent to it by the client, and cannot otherwise generate responses.
- **Attention.** Provide for the *server* to indicate “attention” to the *client*. This attention will consist of a boolean indicator raised by the server. The server may assert attention at any time.
- **Clear Attention.** Provide for the client to recognize and clear the attention condition.

4.3 Query Layer

As currently specified, any query can have the string “.DATA_TYPE” appended to it. If the query is valid the server will respond with “INTEGER”, “STRING”, or “CONTAINER” to indicate the data type returned by that query. Table 3 below shows some examples using fictitious queries.

- **Integer Data Type**
 - Values of integers (a 32-bit signed integer) are set by appending ‘=’ and a value to the integer’s query. Values can be set either in decimal or in hex, using the “C” notation. When an integer value is queried, the server normally responds with its decimal value. If the query has .AS_HEX appended to it, the value will be returned in hex.
- **String Data Type**
 - It is a string of printable ASCII characters. Values can be set by appending ‘=’ and a value to the string’s query. Strings can be set to any sequence of byte values using the *printable characters of the ASCII character set and escape sequences*.

Table 3 – Fictitious Query Examples

Client	/MEMRI/X.DATA_TYPE	The client requests the data type of X.
Server	INTEGER	The server responds with the data type.
Client	/MEMRI/VOLUMES.DATA_TYPE	The client requests the data type of VOLUMES
Server	CONTAINER	The server responds with the data type.

- **Container Data Type**

- An object capable of holding and organizing several sub-objects. The sub-objects are typically related to each other. Container values cannot be set in a single step. Container values can be polled.

4.4 Volume-Related Queries (/MEMRI/VOLUMES)

The /MEMRI/VOLUMES object is a *container*. The members of this container will be a list of all volumes for which this *server* has MEMRI information available. Table 4 shows examples of volume-related queries. The volume-related queries include Volume Parameters (/MEMRI/VOLUMES/nn/PARAMETERS) and Volume Media Parameters (/MEMRI/VOLUMES/nn/PARAMETERS/MEDIA):

- Volume Parameters (/MEMRI/VOLUMES/nn/PARAMETERS):
- This container object contains all *parameters* for the specified volume. Compliant devices must provide this container object and it must contain the CASETYPE, MANUFACTURER and MEDIA sub-objects. Additional sub-objects can be present.
- Volume Media Parameters (/MEMRI/VOLUMES/nn/PARAMETERS/MEDIA):
- This container object contains all media-related *parameters* for the specified volume. Compliant devices must provide this container object and it must contain the TYPE and MEMRI_STORAGE_LOCATION sub-objects. Additional sub-objects can be present.

Table 4 - Examples of Volume-Related Queries

<i>Client</i>	/MEMRI/VOLUMES	<i>The client requests a list of all volumes</i>
<i>Server</i>	{AA23 THX-1138 "REFERENCE TAPE"}	<i>The server returns the list of volumes</i>
<i>Client</i>	/MEMRI/VOLUMES/THX-1138	<i>The client accesses information about volume "THX-1138"</i>
<i>Server</i>	{PARAMETERS EVENTS READINGS ACCUMULATORS RECOMMENDATIONS}	<i>The server responds by listing all of the areas for which information is available concerning the volume</i>

5. Conclusions

To provide for high accessibility and high data integrity of stored data, network storage management applications require reliable indicators of data integrity from data storage devices. To satisfy this need, data storage devices are becoming more intelligent. Reliable indicators of data degradation include early warnings, media error monitoring and reporting techniques and data integrity metadata. Error correction activity generally occurs without notification to the host. With data integrity metadata available at the network storage management level, these applications can monitor the degree of data degradation and error correction activity taking place in data storage devices. Highly reliable storage systems can use these network storage management applications to significantly decrease the possibility of data loss. Network storage management industry standards that specify MEMR techniques and the associated data integrity metadata are emerging. An industry standard to verify the data integrity of stored data on optical digital data disks exists and others standards are being developed.

References

- [1] M. Peterson. *Opportunities in Network Storage Management*, A document produced for Avail Systems, Nov. 1995, <http://www.sresearch.com/search/105275.html>.
- [2] F. Podio, *Optical Storage Media Data Integrity Issues*, Association for Information and Image Management International, Optical Tape Study Group, AIIM OTSG¹/95-007, Aug. 1995.
- [3] S. Silberstein. *Analysis of ANSI/AIIM MS59, Media Error Monitoring and Reporting Techniques to Verify Stored Data on Optical Digital Data Disks*, AIIM OTSG/95-008, Aug. 1995.
- [4] F. Podio. *Monitoring and Reporting Techniques for Error Rate and Error Distribution on Optical Disks Systems*, AIIM OTSG/96-013, Feb. 1996
- [5] M. Manavi. *Tape Defect Size Characterization*, AIIM OTSG/96-015, Feb. 1996
- [6] K. Thibodeau, *Position Letter in Support of the Development of Media Error Monitoring and Reporting Standards for Optical Tape*, AIIM OTSG/96-029, Apr. 1996.
- [7] W. Vollrath. *MEMRI (TM), Media Error Monitoring, Reporting and Information*, OTSG/96-030, Apr. 1996.
- [8] W. Vollrath. *MEMRI (TM) Information Structure Proposal*, OTSG/96-050, Oct. 1996.
- [9] AIIM Optical Tape Study Group. *Project Proposal for the development of an American National Standard, Digital Data Interchange in 12.65 mm (0.5 in) Write-Once Read Many Times (WORM) Optical Tape Cartridges*, April 15, 1997
- [10] AIIM Optical Tape Study Group. *Project Proposal for the Development of an American National Standard, Media Error Monitoring and Reporting Information*

¹ Association for Information and Image Management International, Optical Tape Study Group documents can be obtained from Fernando Podio, NIST.

(MEMRI) to Verify the Integrity of Stored Data on Optical Tape Media, AIIM OTSG/97-022 Rev. 1 (Revised August 7, 1997).

- [11] International Standards Organization (ISO). *Information Technology - 130 mm Rewritable Optical Disk Cartridge for Information Interchange*, ISO/IEC 10089-1991.
- [12] F. Podio, W. Vollrath, B. Kobler, *Media Error Monitoring and Reporting Information (MEMRI) - Metadata for Intelligent Digital Data Storage Devices*, Proceedings of the Second IEEE Metadata Conference, September 1997, <http://computer.org/conferen/proceed/meta97/papers/fpodio/fpodio.html>.
- [13] American National Standard for Information Systems. *Information Technology - Helical-Scan Digital Computer Tape Cartridge, 12.65 mm (0.498 in), for Information Interchange*, ANSI X3.267-1996
- [14] International Standards Organization (ISO). *Information technology -- Data interchange on 90 mm optical disk cartridges -- Capacity: 640 Mbytes per cartridge*, ISO/IEC 15041:1997.
- [15] International Standards Organization (ISO). *Information technology -- 3,81 mm wide magnetic tape cartridge for information interchange -- Helical scan recording -- DDS format using 60 m and 90 m length tapes*, ISO/IEC 12247:1993.
- [16] American National Standard for Information Systems. *Media Error Monitoring and Reporting Techniques for Verification of Stored Data on Optical Digital Data Disks*, ANSI/AIIM MS59, March 1996.
- [17] American National Standard for Information Systems. *Guidelines for the Use of Media Error Monitoring and Reporting Techniques*, ANSI/AIIM TR39, January 1996.
- [18] Hewlett Packard Co. *TapeAlert™ - A New Approach to Tape Drive Management*, <http://www.hp.com/tape/tapealert/ta.html>.
- [19] Imation Corp. *The Imation Media Performance Manager*, Imation's Press Release, April 1997.
- [20] S. Gold. *TapeAlert Specification Version 1.2*, September 1997. <http://www.hp.com/tape/tapealert/taspec12.html>
- [21] AIIM Optical Tape Subcommittee. *Preliminary MEMI Standard*, AIIM C21.3/97-013.