# Techniques and Metadata for Intelligent Data Storage Devices

**Fernando Podio**
**Information Technology Laboratory, NIST**
**Bldg. 225 Room A248**
**Gaithersburg, MD 20899**
**Phone +1-301-975-2947; FAX: +1-301-869-7429**
**e-mail: fernando.podio@nist.gov**

**Presented at the THIC Meeting at the Amberley Suite Hotel Albuquerque NM**

**April 22, 1998**

NIST

# Outline

- Techniques and metadata for intelligent data storage devices
- Data storage subsystem reliability and data integrity
- Existing and proposed standards
- The “TapeAlert” effort
- The “MEMRI” effort
- Expansion of the “MEMRI” effort
- Conclusions

NIST

# Techniques and Metadata for Intelligent Data Storage Devices?

There are many techniques that can make a storage device "intelligent" (e.g., network-attached, compression, data integrity monitoring and reporting).

***This presentation will focus on:***

- ***techniques and metadata to verify the stored data on sequential storage media (magnetic or optical).***
- ***emerging storage management standards that specify these techniques and the associated metadata.***

NIST

# Data Storage Subsystem Reliability and Data Integrity

Some of the questions that need to be answered continually:

- Is stored data safe and available?
- What is the level of correction that is taking place in the storage device?
- Should the given media continue to be used, or should the data be copied to fresh media?

NIST

# Data Storage Subsystem Reliability and Data Integrity

To preserve data integrity and availability, storage devices and advanced storage management software need incorporate standardized techniques and metadata for verification of the stored data. They need to provide:

- Early warnings of data degradation
- Advanced media error monitoring and reporting
- Robust disaster-recovery solutions

NIST

# Standard Techniques and Metadata for Verification of Stored Data

- *For **optical disk** subsystems, there is an ANSI standard and an ANSI Technical Report (both approved in 1996):*
  - ANSI/AIIM MS59:
    - "Media Error Monitoring and Reporting Techniques for Verification of Stored Data on Optical Digital Data Disks"
  - ANSI/AIIM TR39:
    - Accompanying guidelines

- *An equivalent ISO standard is being developed based on ANSI/AIIM MS59 (DIS 12142).*

NIST

# Standard Techniques and Metadata for Verification of Stored Data (Continued)

*For **sequential** storage media subsystems, there exist no formal ANSI or ISO standards describing techniques and metadata that would allow the host to identify potential errors and problems that may compromise storage reliability or integrity.*

# Proposed Industry Storage Management Standards for Sequential Storage Media Subsystems

- Two industrial groups are developing "standardized" techniques to monitor the reliability of sequential storage devices and the data integrity of stored data:

❑ **TapeAlert™**

- A specification being developed under the TapeAlert Working Group.
- Developed by Hewlett-Packard as a proposed new industry standard for magnetic tape drive management.

❑ **"Media Error Monitoring and Reporting Information" (MEMRI)**

- A standard being developed under the Association for Information and Image Management (AIIM) International, Optical Tape Subcommittee.
- Current scope is techniques and metadata to verify stored data on optical tapes.

NIST

# TapeAlert™

- For magnetic tapes, drives and libraries:
  - Performs constant self-diagnostics.
  - Specifies event driven monitoring.
  - Interprets this information into “standard” high-level error flags.
  - Accesses the error flags via SCSI Log and Mode Pages.
  - Accumulation of trend metadata is not explicitly specified.
  - Specification versions through 2.0 are fully written.
  - Compliance to the specification is certified by the TapeAlert Logo Program Licensed by HP.

# MEMRI

- Defines a means of transporting media error monitors, reports and information in a technology and interface independent manner.
- Defines information-level and report-level content.
- Defines a standard means for describing and transporting standardized and vendor-specific metadata contained in accumulators and registers.

# MEMRI (Continued)

- Accumulates extensive statistics:
  - Across media and drive lifetime
  - Across many media and drives
  - The metadata can be collected into a single database containing data integrity information.

- Data integrity trends can be observed from the metadata:
  - Across long time periods
  - Across many media use sessions

# Current Status of the “MEMRI” Effort

- Currently is being developed under Subcommittee C21.3, “Optical Tape”.
- Currently proposed for optical tape devices.
- A proposal to expand the “MEMRI” effort to any type of sequential storage media devices will soon be submitted to AIIM.

# Strategy to Expand the "MEMRI" effort

- Interested organizations will soon submit to AIIM a project proposal for the development of a ANSI/AIIM standard specifying techniques and metadata to verify stored data on sequential (magnetic and optical) storage media.
  - They will propose a new C21 Subcommittee as the home for this activity.

- AIIM Committee C21, "Data Storage Subsystems"
  - C21.1, "FSMS"
  - C21.2, "COLD Technologies"
  - C21.3, "Optical Tape"
  - **C21.4, "Intelligent Data Storage Devices" (Proposed)**

# Perceived Qualities of an Expanded ANSI Standard

- Existing quality figures and methods (e.g., MTBF, BER) are not sufficient to satisfy user's requirements (e.g., high-end applications).
- Will provide media error monitoring and reporting and early warnings of data degradation for any type of sequential storage media.
- Will specify techniques to verify the state of the stored data on the media.
- Will specify how to transport the related metadata in a technology and interface-independent manner.
  - Auxiliary standards (or an extension of the standard) perhaps required to specify transport across popular interfaces.
- May provide different levels of conformance to support different levels of system sophistication and requirements.
- Data integrity trends can be observed from the collected metadata.

NIST

# Current Thinking

- The standard intends to specify:
  - A simple high-level summary interface

    Should the host continue using the media, or should it be copied to fresh media?

  - Reports backing up the high-level summary
    - Address the questions:
      - Why is this media still usable?
      - Why should this media be copied to fresh media?

  - Access to accumulators and registers containing data integrity metadata
    - The content of a set of accumulators/registers will be defined.
    - Support for vendor-specific accumulators/registers will be provided.

NIST

# Current Thinking (Continued)

- Alternatives open to discussion are where the metadata is stored:
  - on the media
  - in a cartridge-intrinsic memory device
  - in the drive
  - in the host system
  - a combination of the above?

# Alternative 1: Storage On The Media

- Data logs would be written to a dedicated area of the media
  - Rewritten on read/write drives.
  - Appended or written to a special area on write-once drives written (for example, immediately before eject).
- The data log would be inseparable from the media surface it represents.

# Alternative 1: Storage On The Media (Continued)

- No host system management would be required.
  - The host would only poll the data log when it needs to.
  - Some hosts might extensively analyze the logs while others completely ignore them.
- Requires the capability to write new logs after read operations.
  - Not possible with read-only drives
  - May not be acceptable behavior to write on a write-protected tape, even if to a special log area.

NIST

# Alternative 2: Storage On A Cartridge-intrinsic Memory Device

- Data logs would be written to memory device intrinsic to the tape cartridge.
- The data log would be inseparable from the media surface it represents.
- The data logs to describe the health of many gigabytes of media may be larger than the amount of memory feasible to embed in a cartridge.
- Most existing cartridges and cassettes do not include this feature.

# Alternative 2: Storage On A Cartridge-intrinsic Memory Device (Continued)

- No host system management would be required.
  - The host would only poll the data log when it needs to.
- Memory device limitations
  - Device must be low- or zero-retention power.
  - Likely devices have a limited number of write cycles.
  - Likely devices have limited data retention times.

# Alternative 3: Storage Inside The Drive

- Data logs would be maintained in non-volatile storage within the drive. Possibilities:
  - Non-volatile SRAM
  - Internal hard drive
  - Other?
- Would allow read-only drives and write-protected media to be monitored.
- No host system management would be required, at least at the time of the record/playback operation.

# Alternative 3: Storage Inside The Drive (Continued)

- Data logs could possibly become separated from their media or split across multiple drives.
- Media history fragmentation
  - The full history of a particular media would not be available as a single entity.
  - Separate histories available piecemeal from each drive that ever mounted the media.
  - Possibly unclear which drives mounted which media.

NIST

# Alternative 4: Storage In The Host System

- Data logs would be maintained on the host.
- Data logs could get separated from their media, but only if physically moved to another host system.
- Would require extensive management of data logs on the host system.
- Read-only and write-once media could be monitored.

# Alternative 4: Storage In The Host System (Continued)

- The data logs for a particular media could be kept unified unless the media were in continuous use by multiple hosts.
- Full history of media is continuously available.
- If the media is moved, its logs could easily be moved with it.

# Conclusions

- *To preserve data integrity and availability, storage devices and advanced storage management software need to provide for standardized techniques and metadata that allows for verification of the stored data.*
- *Storage subsystems that provide standardized techniques and metadata increase user's confidence in the use of these subsystems.*
- *No such standards exist for sequential storage media subsystems.*
- *An extended standard based on the MEMRI effort will provide the required techniques and data integrity metadata. Possible data degradation trends can be observed from the collected metadata.*

NIST