




US 20150178006A1

(19) **United States**(12) **Patent Application Publication**  
**Stripling**(10) **Pub. No.: US 2015/0178006 A1**(43) **Pub. Date: Jun. 25, 2015**(54) **DETERMINING, DISPLAYING AND USING  
TAPE DRIVE SESSION INFORMATION**(71) Applicant: **KIP CR P1 LP**, New York, NY (US)(72) Inventor: **Jeffrey R. Stripling**, Austin, TX (US)(21) Appl. No.: **14/640,072**(22) Filed: **Mar. 6, 2015****Related U.S. Application Data**(63) Continuation of application No. 12/025,300, filed on  
Feb. 4, 2008, now Pat. No. 9,015,005.**Publication Classification**(51) **Int. Cl.**  
**G06F 3/06** (2006.01)(52) **U.S. Cl.**CPC ..... **G06F 3/0619** (2013.01); **G06F 3/0653**  
(2013.01); **G06F 3/0686** (2013.01); **G06F**  
**2003/0698** (2013.01)

(57)

**ABSTRACT**

Methods and systems for monitoring libraries and library drives. In one embodiment, a method is provided which includes polling the library to determine whether a storage media has been loaded in to, or unloaded from, a drive. When the media has been loaded, the method includes beginning a media monitoring session. The method may include determining which tape has been loaded in to which drive of the library and associating the session performance with that tape and that drive. Some embodiments include using a Read Element Status command to perform the polling. In some examples, the library is a magnetic tape library and the polling is performed out-of-band. Information regarding the media sessions can be output in order of session performance. Some embodiments include determining whether the cause of a slow session is external to the library and correcting the external cause.






FIG. 1




FIG. 3




FIG. 2

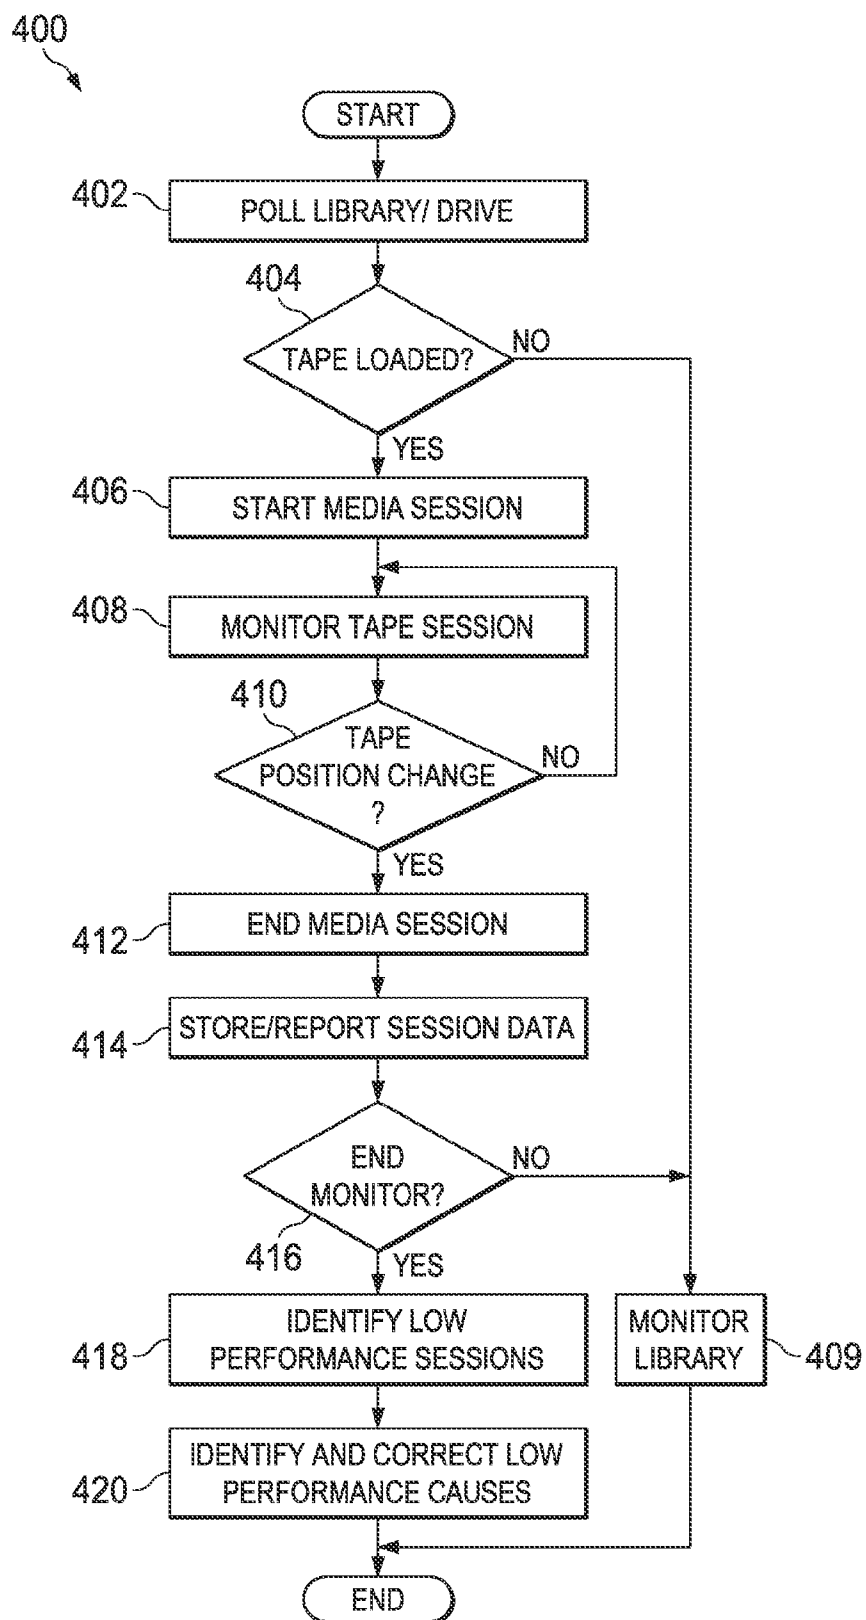


FIG. 4




FIG. 5




FIG. 8

FIG. 6

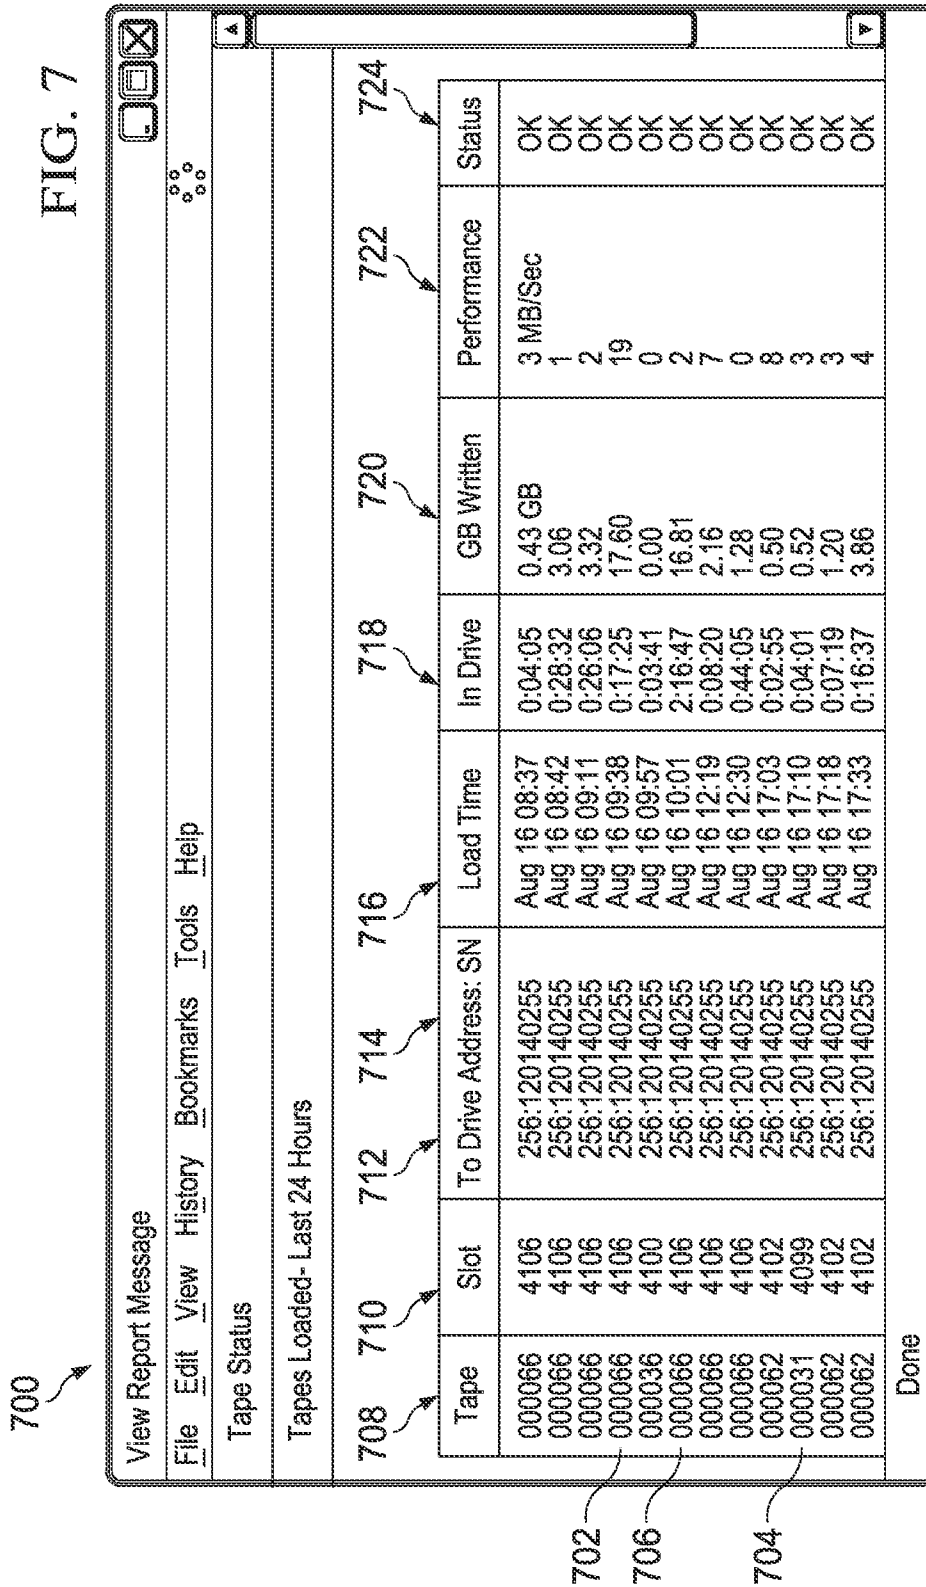
600

View Report Message									
File Edit View History Bookmarks Tools Help									
Tape Status									
Tapes Loaded- Last 24 Hours									
608	610	612	614	616	618	620	622	624	
Tape	Slot	To Drive Address: SN	Load Time	In Drive	GB Written	Performance	Status		
000036	4100	257:1210139351	Aug 16 01:26	0:39:36	74.50 MB/Sec	33 MB/Sec	OK		
000036	4100	257:1210139351	Aug 16 02:06	1:01:35	120.64	33	OK		
000058	4097	257:1210139351	Aug 16 03:08	0:02:04	0.00	0	OK		
000036	4100	257:1210139351	Aug 16 03:11	0:07:29	0.38	1	OK		
000058	4097	257:1210139351	Aug 16 03:20	0:17:31	13.25	13	OK		
000058	4097	257:1210139351	Aug 16 03:38	0:46:14	128.02	47	OK		
000058	4097	257:1210139351	Aug 16 04:25	0:08:26	2.16	7	OK		
000036	4100	257:1210139351	Aug 16 04:34	0:08:59	10.06	30	OK		
000058	4097	257:1210139351	Aug 16 04:44	0:07:15	1.26	4	OK		
000058	4097	257:1210139351	Aug 16 04:53	0:04:03	0.49	8	OK		
000058	4097	257:1210139351	Aug 16 04:57	0:07:54	1.21	3	OK		
000031	4099	257:1210139351	Aug 16 05:06	0:04:02	0.51	3	OK		
000058	4097	257:1210139351	Aug 16 05:11	0:08:00	1.43	5	OK		
000058	4097	257:1210139351	Aug 16 05:20	0:08:44	1.20	3	OK		
000058	4097	257:1210139351	Aug 16 05:29	0:05:57	0.42	1	OK		
000058	4097	257:1210139351	Aug 16 05:36	0:18:55	3.83	4	OK		
000058	4097	257:1210139351	Aug 16 05:55	0:06:14	0.53	2	OK		
000058	4097	257:1210139351	Aug 16 06:02	0:08:01	0.48	1	OK		
000058	4097	257:1210139351	Aug 16 06:11	0:05:02	0.45	2	OK		
000058	4097	257:1210139351	Aug 16 06:17	1:27:12	2.96	0	OK		
000058	4097	257:1210139351	Aug 16 06:45	0:07:05	1.15	3	OK		
Done									

602

606

604



## DETERMINING, DISPLAYING AND USING TAPE DRIVE SESSION INFORMATION

### CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application is a continuation of, and claims a benefit of priority under 35 U.S.C. 120 of the filing date of U.S. patent application Ser. No. 12/025,300, filed on Feb. 4, 2008, "Determining, Displaying, and Using Tape Drive Session Information", the entire contents of which are hereby expressly incorporated by reference for all purposes.

### TECHNICAL FIELD OF THE INVENTION

[0002] The present disclosure relates to methods and systems for monitoring storage components. More particularly, an embodiment relates to methods and systems for monitoring library performance by identifying library tape sessions in a non-intrusive manner using an out-of-band appliance.

### BACKGROUND

[0003] Data represents a significant asset for many entities. Consequently, data loss, whether accidental or caused by malicious activity, can be costly in terms of wasted manpower, loss of goodwill from customers, loss of time and potential legal liability. To ensure proper protection of data for business and legal purposes, many entities back up data to a physical storage media such as magnetic tapes or optical disks. Traditionally, backup would occur at each machine controlled by an entity. As the sophistication of network technology increased, many entities turned to enterprise level backup in which data from multiple machines on a network is backed up to a remote library. Typically, a library includes a variety of components which include a plurality of media for data storage, e.g., multiple magnetic tapes. Centralized data backup has the advantage of increased volume, efficiency, and redundancy.

[0004] In many systems, the data to be backed up and backup commands are sent over a network from multiple machines on the network to a library. In many instances, the data to be backed up and the backup commands are routed to the library through a switch.

[0005] One example of a library commonly used in enterprise backup systems is a magnetic tape library. A magnetic tape library can comprise components such as tape cartridges (containing magnetic tape), robots, tape slots and tape drives. A typical magnetic tape library contains multiple cartridge slots in which tape cartridges can be stored. Tape cartridges, commonly referred to as tapes, are physically moved between cartridge slots and tape drives by a robot. The robot is controlled by commands received from the host devices on the network. When specific data is required, a host device determines which cartridge slot contains the tape cartridge that holds the desired data. The host device then transmits a move-element command to the robot and the robot moves the tape cartridge to a tape drive which reads the desired data from the tape cartridge.

[0006] In a SCSI tape library, devices that are part of the library are typically addressed by target number. Thus, each drive and robot of a tape library typically has a target number. Cartridge slots, on the other hand, are addressed by element numbers that are used by the robot to locate the slots. Because the robot also places tape cartridges in the drives, each drive is also associated with an element number.

[0007] Components of a library are subject to wear and other forms of degradation which may degrade individual library component(s) to a point where the degraded component(s) can no longer be reliably used. Degradation of library components may render a library unreliable and prone to data loss, diminishing the library's usefulness as a backup device. The failing of even one library component can cause the loss of valuable data. Furthermore, the process of determining which library component is failing and rendering the library unreliable can be time consuming, uncertain and expensive. Networks and other equipment external to the library may also affect the performance of the library. For instance, if data is being backed up from a remote location, transport delays and interruptions along the transmission path may cause the data to arrive relatively slowly (or not at all) at the library. While the library may store the data just as fast as it arrives, the performance of the library suffers because the storage rate in these circumstances is nonetheless slow.

### SUMMARY

[0008] Embodiments disclosed provide methods and systems of monitoring library performance that eliminate, or at least substantially reduce, the shortcomings of prior art library systems and methods. Some embodiments include polling a library to determine whether a media (e.g., tape, optical disk, or other media) has been loaded or unloaded between polls. When it is discovered that a media has been loaded a media session is begun. Performance data (e.g., data transfer speed, error rates, etc) associated with the media session is monitored. When it is determined that the tape has been unloaded, the media session ends. Information regarding the media session for instance, performance related information, can be output and stored for subsequent use.

[0009] Back up system topologies vary widely. For instance, in some embodiments, data may flow to the back up system, its library(s), and drive(s) from multiple points of a network. In many back up systems, each tape may be placed in any drive capable of reading that tape and any drive may be used to read any compatible tape in the library. During a back up job, and over the course of time, an individual tape may be loaded in to many different drives and an individual drive may have many different tapes loaded in to it. These factors may contribute to, mask, or have little effect on the cause(s) of differing library performance between tape sessions. For instance, jobs (e.g., a data write to the tape or a data read from the tape) associated with a particular tape might suffer comparatively low library performance in which case the tape becomes suspect. In other circumstances, jobs associated with a particular drive might suffer comparatively low library performance in which case the drive becomes suspect. In other cases, the low performance jobs may exhibit no apparent pattern with regard to either the tapes or the drives in which case it may be useful to associate the low performance with the tape session rather than the tape or the drive involved. Users and owners of storage area networks, back up systems, libraries, drives, tapes, and the data thereon are interested in knowing the performance of these systems as a whole as well as the individual components of these systems. This knowledge can aid in troubleshooting failures, troubleshooting slow performing sessions, and in optimizing system performance and utilization.

[0010] The methods of some embodiments monitor the performance of back up systems, libraries, and operations and associate information with identified tape sessions. A tape



session can begin when a tape is loaded in to a drive and can continue until the tape is unloaded. Information can be collected regarding the tape session to allow a user to optimize the performance and utilization of the systems and components involved. The methods of some embodiments gather information such as the amount of data transferred per time period, the types and numbers of errors occurring during a time period, and the identities of the tapes and drives used during a time period. The gathered information can allow interested parties to understand and improve the performance and utilization of the back up system.

**[0011]** In some embodiments that use a robot (or other automated tape handling mechanism) to load and unload tapes in a library, the library can be periodically polled to determine the position of the tapes in the library. Information from consecutive polls can be examined and used to generate a list of tape movements in the library. For instance, if a tape in a first poll was not in a tape drive and the tape upon a second poll is in a tape drive, the beginning of a tape session may be declared. If a tape was in a tape drive in a first poll and the tape is not in the tape drive upon a second poll, the end of a tape session can be declared. In some embodiments, the system can use the SCSI Read Element Status command to implement the polls.

**[0012]** In some embodiments without automatic tape handling mechanisms (e.g., embodiments which include a free-standing tape drive) the tape drive can be periodically polled to determine if a tape has been loaded in to any of the drives. By examining differences between consecutive polls, a list of tape movements in to and out of the tape drives can be generated. If a tape drive in a first poll reported that no tape was in the drive and the drive reports that a tape is in the tape drive upon a second poll, the beginning of a tape session can be declared. If a tape drive in a first poll reported that a tape was in the drive and the drive reports that no tape is in the tape drive upon a second poll, the end of the tape session can be declared. In some embodiments, the SCSI Read Media Serial Number command can be used to implement the polls.

**[0013]** During the tape sessions, in some embodiments, information can be gathered regarding the session. Examples of such information include, but are not limited to, the amount of data written and read, the number of errors encountered, and the start and end times of the session. The gathered information can be stored in a database and additional information can be gathered or calculated. Examples of calculated information include (but are not limited to) the duration and the average data transfer rate of the tape sessions. Additional measurements of system performance can also be gathered or derived.

**[0014]** The tape drive and the tape used for a given tape session can be identified in some embodiments. Information can be gathered during the tape sessions and can be associated with the tape drive, the tape, the tape session, or any combination thereof. In some embodiments, a list of tape sessions grouped by tape drives, a collection of tape drives, tapes, time periods, or other considerations can be displayed. For example, errors can be grouped by tape session and can be used to pinpoint the source of problems associated with the tape drives and tapes in use for the tape sessions. Some embodiments provide informational listings of tape sessions. Users can examine the listings to determine whether particular tape sessions were the source of most performance issues or whether the performance issues correspond to some other

phenomenon. Information can be gathered that may be useful in distinguishing between low performance tapes, tape drives, and tape sessions.

**[0015]** Some embodiments provide outputs such as user selected tables of tape session information. Notifications based on user selected tape session characteristics can also be provided. Systems of some embodiments can sort the sessions by performance and list the worst performing sessions for the user. Users can analyze this tape session information to identify low performing sessions. For example, the user (e.g., the library administrator) can determine the source of the data that was being backed up during the low performing session. The user can improve the data connection speed between the data source and the tape drive. The next time the same data source writes to a tape, the user can examine the speed of the new tape session to verify whether system performance improved. By eliminating the low performing sessions, the user can forestall purchase of new library equipment by getting more performance out of existing equipment.

**[0016]** Embodiments of the disclosure provide many advantages. For instance, some embodiments allow the user to determine the performance of a set of tapes and tape drives. For another example, some embodiments allow the user to identify low performing tape sessions and eliminate the causes thereof from the back up environment. The identification of low performing sessions also allows users to identify inefficiencies in backup jobs. Users may also improve the performance of components in the backup system. Other advantages include allowing users to identify data sources that can be backed up by a drive closer to the data source. Still further advantages include allowing users to improve the time it takes to complete a set of back ups and to avoid damaging drives by running the drives below their recommended streaming rates. Overall back up performance can also be improved by identifying the worst sessions and eliminating the causes thereof.

#### BRIEF DESCRIPTION OF THE FIGURES

**[0017]** A more complete understanding of the disclosure and the advantages thereof may be acquired by referring to the following description, taken in conjunction with the accompanying drawings in which like reference numbers generally indicate like features and wherein:

**[0018]** FIG. 1 is a diagrammatic representation of one embodiment of a library.

**[0019]** FIG. 2 is a diagrammatic representation of one embodiment of a backup system using a library and having a monitoring appliance.

**[0020]** FIG. 3 is a diagrammatic representation of one embodiment of a controller which can be used in a monitoring appliance.

**[0021]** FIG. 4 is a flowchart illustrating an embodiment of a method for monitoring libraries.

**[0022]** FIG. 5 is a diagrammatic representation of a tape session.

**[0023]** FIG. 6 illustrates a report of a monitoring appliance.

**[0024]** FIG. 7 illustrates another report of a monitoring appliance.

**[0025]** FIG. 8 is a flowchart illustrating another embodiment of a method of for monitoring libraries.

## DETAILED DESCRIPTION

[0026] Preferred embodiments of the disclosure are illustrated in the FIGURES, like numerals being used to refer to like and corresponding parts of the various drawings.

[0027] Various embodiments provide a method and system for classifying information from monitoring media drives based on when a media is in a media drive. By monitoring tape sessions, the degradation and reliability of system components can be monitored, allowing unreliable or unacceptably degraded components to be identified, thus allowing for the unreliable or degraded components to be bypassed or replaced, enhancing the reliability of the system and proactively preventing data loss.

[0028] Embodiments can be utilized in conjunction with monitoring systems for various media including magnetic tapes and optical disks. While various embodiments are disclosed with reference to tape systems, embodiments can be used with other media systems. FIG. 1 is a diagrammatic representation of one embodiment of a tape library, as would be understood by one of ordinary skill in the art. Library 100 can comprise drives 140a-140e, media changer 125 and associated robot 130, import/export element(s) 145 and slots 135a-135j. Drives 140a-140e can read/write data from/to magnetic tape (contained within cartridges), eject tape cartridges, and perform other operations. Slots 135a-135j store the magnetic tape cartridges when they are not in a drive and robot 130 moves the magnetic tape cartridges between drives 140a-140e and slots 135a-135j. For example, robot 130 may move a tape cartridge stored at slot 135a to drive 140b so that data can be written to the tape cartridge. It should be noted that some libraries may employ a single robot or multiple robots in an expandable or modular configuration.

[0029] A Read Element Status (RES) command is a command which is used to poll the state of a library. That is, a RES command is sent to a library and in response, the library returns its current state, including the locations of individual tape cartridges. Thus, a RES command provides a snapshot of a library at any one time. Examples of a RES command can be found in section 6.5 of NCITS T10/999D Revision 10a Specification, section 6.10 of INCITS T10/1383D Revision 7 Specification and section 6.10 of INCITS T10/1730-D Revision 5 Specification; all of the above-mentioned Specifications (the "T10" specifications) are hereby incorporated by reference.

[0030] A Read Media Serial Number (RMSN) command is a command which is used to determine the serial number of a tape in a drive. The RMSN command is sent to a drive and in response, the drive returns the vendor assigned serial number of the tape, if a tape is in the drive. If the tape has no serial number, the response indicates that no serial number is present even though a tape is in the drive. If the drive contains no tape, the response includes an indication that no media is present in the drive (a.k.a. as a "Check Condition" response). Examples of a RMSN command can also be found in the T10 specifications.

[0031] A Log Sense (LS) command is a command which is used to obtain data associated with a particular drive. A LS command is sent to a particular drive of a library and in response, the drive returns data associated with the drive and/or the media contained in the drive depending on the actual LS command. For example, such data might include: read errors, write errors, utilization and performance data, data regarding the data written and read to a media, media loaded, detail tables or other data. In one embodiment, the

amount of data written over a period of time may be derived from data returned in response to one or more LS commands. More specifically, data returned in response to an LS command may include Log Sense page 0xC, 0x32 or 0x39 which may include data from which the amount of data written by a particular drive over a period of time may be derived (e.g., calculated). Examples of a LS command can be found in "SCSI Primary Commands-3 (SPC-3)", (Project T10/1416-D), Revision 22a, Mar. 25, 2005, propagated by the T10 Technical Committee of the InterNational Committee on Information Technology Standards (INCITS), which is hereby incorporated by reference.

[0032] An Inquiry command is a command that is used to query relatively static information (e.g., serial number, manufacturer or other relatively static information) from components of a library such as a drive or a media changer. According to one embodiment, Inquiry commands are used to query individual library components. That is, an individual Inquiry command may query a particular library component. Examples of an Inquiry command can be found in "SCSI Primary Commands-3 (SPC-3)", (Project T10/1416-D), Revision 22a, Mar. 25, 2005, propagated by the T10 Technical Committee of the InterNational Committee on Information Technology Standards (INCITS), referenced above. Embodiments of the invention can utilize a monitoring appliance which can be a Read Verify Appliance (RVA). The monitoring appliance polls the state of a library over time by sending RES commands to the library at intervals. Returned library states are compared. Based on the comparison between the returned library states, the movement of tape cartridges over time can be tracked and the beginning and ending of tape sessions can be identified. One example of tracking movement of media is found in U.S. patent application Ser. No. 11/801,809, filed on May 11, 2007, entitled "Method and System for Non-Intrusive Monitoring of Library Components", by Foster et al., which is incorporated herein as if set forth in full.

[0033] Various embodiments identify tape (or other media) sessions which represent the period of time that a tape is in a tape drive or other selected location. Tracking tape sessions allows sources of errors to be identified. For example, an individual tape cartridge may suffer from an increasing number of data errors in any drive: this implies that the tape cartridge is degrading and may require replacement. In another example, a tape cartridge may encounter errors when used with a tape drive A but may encounter no errors when used with other tape drives. If other tape cartridges encounter errors when used with tape drive A, but encounter no errors when used with other tape drives, this implies that tape drive A is unreliable. Accordingly, tape drive A may be bypassed or replaced. Some errors, though, do not appear to correlate to either a particular tape or to a particular drive. In many instances, these types of errors may correlate with influences external to the tapes and tape drives. By tracking tape sessions, errors can be correlated to these other influences.

[0034] FIG. 2 is a diagrammatic representation of a system 200 in which a plurality of hosts 202a-202d have access to library 100 over network 205. Network 205 can comprise the Internet, a SAN, a LAN, a WAN, a wireless network, or any other communications network known in the art. Hosts 202a-202d are coupled to library 100 via network 205 and switch 207. Similarly, library 100 can communicate with switch 207 over any suitable communications link or network known in

the art and can use additional protocols such as iSCSI, TCP/IP, or other protocol known in the art. Switch 207 is connected to monitoring appliance 220.

[0035] Switch 207 is connected to library 100. Thus, switch 207 can forward RES commands from monitoring appliance 220 to library 100. Switch 207 receives the library state generated in response to the RES command from library 100 and forwards the library state to monitoring appliance 220 which stores the library state and/or compares it with one or more previous library states. Thus monitoring appliance 220 can continually poll the state of library 100, track tape movements and monitor the components of library 100. It should be noted that the intervals between polling can be fixed periods of time or of variable duration or a combination of the two. In one embodiment, the monitoring appliance polls a library every 30 seconds. In another embodiment, the time duration between polling is less than the backup time or the tape movement time. In embodiments of the invention, the intervals can be determined by computer algorithm or user input. Because monitoring appliance 220 tracks tape movements using RES commands, monitoring appliance 220 does not need to intercept commands or responses from network 205 to track tape movements. Thus, in one embodiment, monitoring appliance 220 can be an out-of-band appliance. This allows monitoring appliance 220 to be a non-intrusive device which does not monitor or interfere with commands from and responses to hosts 202a-202d. Consequently, monitoring appliance 220 can be a compartmentalized device which can be coupled to a switch and which does not have to be integrated into network 205. An advantage of this out-of-band methodology is that a monitoring appliance can be used to monitor library components without cumbersome configuring.

[0036] While shown as a physical media library in FIG. 2, library 100 can be a virtual media library that is a virtual representation of one or more physical media libraries as presented by switch 207, a library controller, or other component. Examples of library virtualization are described in U.S. patent application Ser. No. 10/704,265, entitled "System and Method for Controlling Access to Multiple Physical Media Libraries", and U.S. patent application Ser. No. 10/703,965, entitled "System and Method for Controlling Access to Media Libraries", both of which are hereby incorporated by reference herein.

[0037] FIG. 3 is a diagrammatic representation of a monitoring appliance Controller 300 ("controller 300"). Controller can include a processor 302, such as an Intel Pentium 4 based processor (Intel and Pentium are trademarks of Intel Corporation of Santa Clara, Calif.), a primary memory 303 (e.g., RAM, ROM, Flash Memory, EEPROM or other computer readable medium known in the art) and a secondary memory 304 (e.g., a hard drive, disk drive, optical drive or other computer readable medium known in the art). A memory controller 307 can control access to secondary memory 304. Controller 300 can comprise a communications interface 306 (e.g., fibre channel interface, Ethernet port or other communications interface known in the art) to connect controller 300 to switch 207. An I/O controller 312 can control interactions with switch 207. Similarly, an I/O controller 314 can control interactions over I/O interfaces 308 and 310. Controller 300 can include a variety of input devices. Various components of controller 300 can be connected by a bus 326.

[0038] Secondary memory 304 can store a variety of computer instructions that include, for example, an operating

system such as a Windows operating system (Windows is a trademark of Redmond, Wash. based Microsoft Corporation) and applications that run on the operating system, along with a variety of data. More particularly, secondary memory 304 can store a software program 330 that monitors tape sessions. During execution by processor 302, portions of program 330 can be stored in secondary memory 304 and/or primary memory 303.

[0039] FIG. 4 is a flow chart illustrating method 400 for monitoring tape sessions associated with library 100 or individual tape drives 130. Method 400 can begin with monitoring appliance 220 of FIG. 2 polling library 100 via a RES command (for example) at step 402. In response to the poll of step 402, library 100 returns a snapshot of the status of the components of library 100 including information regarding slots 135, drives 140, and the tapes therein. Successive polls of library 100 can be compared at step 404. If the current poll, when compared with the previous poll, indicates that a tape has been loaded into a drive 140 a new tape session may be initiated at step 406. Monitoring application 220 may receive the tape number, the tape drive 140 number, and the tape drive 140 serial number involved in the information returned with the RES command response. If no tape has been loaded, then monitoring appliance 220 can continue polling library 100 as shown at step 402 until a change of status occurs with respect to the loading and unloading of tapes in to and out of drives 140. If a tape was already in drive 140, or if a tape session began at step 406, monitoring appliance 220 can monitor the tape session at step 408. In step 409 monitoring of library 100 as a whole, drives 140, or the components of library 100 can continue at all times during method 400.

[0040] FIG. 8 is a flow chart illustrating one embodiment of a method for collecting data and organizing the collected data so that it can be accessed by or displayed to a user. According to one embodiment, the method of FIG. 8 can be implemented as a set of computer executable instructions stored on a computer readable medium at, for example, monitoring appliance 820. The set of computer executable instructions can, when executed, collect data associated with library components in a non-intrusive manner. At collect data step 810, data is collected by a monitoring appliance by querying library components of a library for data by sending commands corresponding to the library components to the library components. In response, the library components return data associated with the library components. More specifically, in one embodiment, RES, Log Sense (LS), and Inquiry commands are sent to library components and, in response, the library components return corresponding data. Data may be collected over time by collecting data from library components at intervals. For example, various commands can be sent every 10 seconds, every 20 seconds or with any other frequency. Thus, over time, library components may be queried a plurality of times, and as a result, the repository may contain data regarding library components at different points in time.

[0041] The monitoring appliance, at compile data step 820, can compile the returned data. For example, the monitoring appliance can compile the data in defined structures (e.g., XML structures or other structures). A structure may contain data associated with a library component returned in response to one or more commands (e.g., RES, LS, or Inquiry commands). For example, a XML structure can include data from RES commands and serial numbers of library components determined from Inquiry commands issued to corresponding library components. At store data step 830, the monitoring

appliance stores the returned data in a repository. Storing data in a repository may comprise storing structures in the repository or may comprise unpacking returned data compiled at compile step 820 and storing the data in the repository. In one embodiment, the data may be stored in records in a database at the repository.

[0042] At organize data step 840, the data in the repository is organized. As part of step 840, data in the repository can be arranged in chronological order and/or superfluous or redundant data can be purged. Corresponding data can be grouped together. For example, data structures corresponding to a particular drive may be arranged in chronological order within the repository.

[0043] At process data step 850, data in the repository may be processed to display desired data to a user. For example, a user may only want to see data corresponding to a particular library component. At process data step 850, the data in the repository is processed such that the desired data is selected to be displayed to a user. Similarly, data may be processed into graphs or charts, or in accordance with any number of purposes or user desires.

[0044] In one embodiment, processing data in the repository can comprise comparing returned library states or data regarding the locations of individual tape cartridges stored in the repository to track the movement of one or more tape cartridges. For example, data in the repository corresponding to different times can be compared and the movement of tape cartridges in a library tracked by differencing the locations of tape cartridges in a library at different times. In embodiments in which data is stored in records in a database, records associated with different times may be compared to track the movement of tape cartridges. Processing data can further comprise correlating errors with a particular library component based on the movement of one or more tape cartridges within the library. Additional information regarding monitoring libraries is disclosed in U.S. patent application Ser. No. 11/801,809, filed on May 11, 2007, entitled "Method and System for Non-Intrusive Monitoring of Library Components," by Foster et al. and of U.S. patent application Ser. No. 12/024,755, filed on Feb. 1, 1998, entitled "Media Library Monitoring System and Method," by Sims, now U.S. Pat. No. 7,908,366, which is incorporated herein as if set forth in full.

[0045] Performance monitored in accordance with the embodiment of FIG. 8, or other methods of monitoring performance data, can be correlated with tape sessions in some embodiments. Turning briefly to FIG. 5, FIG. 5 illustrates two tape sessions 502 and 504 and performance 505 and 508 associated therewith. As FIG. 5 shows, tapes A1234 and B5678 were involved in tape sessions 502 and 504. Tape session 502 shows a higher and steady performance 506 when compared to the lower and erratic performance 508 of tape session 504. Performance 508 exhibits a peak 510 after a gradual build up and shifting low-level performance thereafter which leads to a low performance plateau 512 with which tape session 504 ends. Tape session 504 may adversely affect the performance of system 200 of FIG. 2 whether the cause of the low performance of tape session 504 is the tape, the drive 140, or some other factor either internal or external to library 100. For instance, any of a number of components of system 200 could slow the data delivery rate to library 100.

[0046] Returning to FIG. 4, monitoring appliance 220 can continue monitoring performance 506 and 508 of library 100 until, at step 410, when it detects that the tape involved in tape session 502 and 504 is unloaded from the drive 140. Moni-

toring appliance 220 may then end (at step 412) tape session 502 or 504 depending on which tape A1234 or B5678 moved. The performance information gathered during tape session 502 or 504 may be output, stored, or further processed at step 414. If the user desires no further tape session monitoring, step 416 shows that monitoring application 220 of FIG. 2 may stop monitoring library 100. If the user wishes to continue monitoring tape sessions, step 416 shows that such monitoring may continue by returning to step 402 and cycling through selected portions of method 400 to detect tape loads and unloads and to monitor performance during tape sessions.

[0047] Users can review performance 506 and 508 related information gathered during tape sessions 502 and 504 to identify low performance tape sessions such as tape session 504 at step 418. Users can also identify the causes for the low performance session 504 at step 420 from information gathered during tape sessions 502 and 504 as well as other information available to the users. For instance, if tape session 504 was associated with host 202b of FIG. 2 and shows low performance 508, host 202b and system 200 components between host 202b and library 100 may be examined to determine if they might be the cause of low performance 508. Underperforming portions of system 200 may be replaced to improve performance 508.

[0048] FIGS. 6 and 7 show reports 600 and 700 which monitoring appliance 220 may display for users of system 200. FIG. 6 shows a tape status report 600 for the last 24 hours for library 100 and FIG. 7 shows tape status report 700 for the last 24 hours for a single drive 140. Report 600 can be sorted by time of tape insertion into a drive and shows performance data associated with a collection of drives. Each line of report 600 can represent a tape session such as tape sessions 602, 604, and 606. Report 600 can also display, for each tape session 602, 604, and 606, the identity 608 of a tape in a drive, the slot 610 in which the tape is stored when not in use, the drive number 612 of the drive involved, the serial number 614 of the drive, the time 616 at which the tape was loaded, the total time 618 during which the tape was loaded in the drive, the amount of data 620 read to or written from the tape, performance 622 (here data transfer rate in GBytes), and the overall status 624 of tape session 602, 604, or 606.

[0049] FIG. 6 illustrates a query from a database in to which information obtained from polling library 100 is stored. For instance, tape identities 608 may originate from bar code tags on the tapes. For another example, slot numbers 610 can correspond to slots 135 of FIG. 2. In-drive time 618 can be computed by subtracting the clock time of the poll which detected the beginning of session 602, 604, or 606 from the clock time of the poll which detected the ending of tape session 602, 604, or 606. Another way in which in-drive time 620 can be calculated is to subtract the times reported by library 100 at which a tape was loaded and unloaded for tape sessions 602, 604, or 606. Many other types of performance information can be reported in report 600. For instance, instead of reporting amount of data transferred 620, the number of read or write errors could be reported. Another example of information which may be reported is performance 622 which is shown as a rate determined by dividing data amount 620 by in-drive time 618. For instance, tape session 602 is shown as having a performance 622 of 33 MB/sec. Tape session 604 shows performance 622 of 7 MB/sec. Tape session 606 illustrates a case in which performance 622 was 0 MB/Sec.

[0050] Other performance related information may be reported by report 600 such as the room temperature, room humidity, or the personnel on call during various tape session 602, 604, and 606. Tape session status 624 could correspond to the status of tape drive 140, the tape, library 100, etc. Tape session status 624 could also report information such as the number of errors which occurred during tape sessions 602, 604, or 606. Other features of report 600 could serve to call attention to information reported therein. For instance, if the number of errors during a session rose beyond a user selected limit, the row of report 600 corresponding to relevant tape session 602, 604, or 606 could be color coded, could be configured to flash, or could be configured to draw attention to itself in some other manner.

[0051] Users can peruse information 608, 610, 612, 614, 616, 618, 620, 622, and 624 (for example) in report 600 and identify low performing sessions 606. Users can also correct or optimize system 200 performance related to low performance tape session 606. Such activity can obviate the desire to purchase additional libraries 100 when library 100 may not be the reason for low performance tape sessions 606. For instance, library 100 to which FIG. 6 pertains has a specified data transfer rate of 80 MB/sec. An inspection of FIG. 6 reveals that only a few tape sessions achieved even about half of the specified rate. Reports such as report 600 can reveal many other aspects of system 200 performance. For instance, while some back up operations can span several tapes, FIG. 6 reveals that library 100 loaded and unloaded tape number 00058 numerous times. It is possible that determining why such activity occurred might reveal an aspect of system 200 which could be corrected or optimized.

[0052] FIG. 7 illustrates that for a single tape drive (of library 200 or a stand-alone drive) report 700 can be a list of tapes that were loaded into the drive. In FIG. 7, report 700 is sorted by the time the tapes were loaded in to drives 140. Monitoring appliance 200 could also generate reports such as reports 600 or 700 for each individual tape in system 200 showing when the tapes were loaded into each of a variety of drives 140. Appliance monitor 220 could also display additional information regarding the tapes, drives 130, and tape sessions 502 and 504. For example, performance curves 506 and 508 of FIG. 5 could be displayed with reports 600 and 700. For another example, monitoring appliance 220 could display which tapes were in which drives 140 during the time that the data for performance curves 506 and 508 were gathered. In some embodiments, reports 600 and 700 and other performance related information can be delivered by a world-wide-web based interface, e-mail, SCSI or SNMP reporting mechanisms, etc.

[0053] As used herein, the terms “comprises,” “comprising,” “includes,” “including,” “has,” “having” or any other variation thereof, are intended to cover a non-exclusive inclusion. For example, a process, product, article, or apparatus that comprises a list of elements is not necessarily limited only those elements but may include other elements not expressly listed or inherent to such process, product, article, or apparatus. Further, unless expressly stated to the contrary, “or” refers to an inclusive or and not to an exclusive or. For example, a condition A or B is satisfied by any one of the following: A is true (or present) and B is false (or not present), A is false (or not present) and B is true (or present), and both A and B are true (or present).

[0054] Additionally, any examples or illustrations given herein are not to be regarded in any way as restrictions on,

limits to, or express definitions of, any term or terms with which they are utilized. Instead, these examples or illustrations are to be regarded as being described with respect to one particular embodiment and as illustrative only. Those of ordinary skill in the art will appreciate that any term or terms with which these examples or illustrations are utilized will encompass other embodiments which may or may not be given therewith or elsewhere in the specification and all such embodiments are intended to be included within the scope of that term or terms. Language designating such nonlimiting examples and illustrations includes, but is not limited to: “for example”, “for instance”, “e.g.”, “in one embodiment”.

[0055] Although embodiments have been described in detail herein, it should be understood that the description is by way of example only and is not to be construed in a limiting sense. It is to be further understood, therefore, that numerous changes in the details of the embodiments and additional embodiments will be apparent, and may be made by, persons of ordinary skill in the art having reference to this description. It is contemplated that all such changes and additional embodiments are within scope of the claims below.

What is claimed is:

1. A monitor appliance for a media library, the monitor appliance comprising:

a processor; and

a computer memory electronically connected to the computer processor, the computer memory storing a set of computer instructions executable by the computer processor for:

polling the media library to determine the position of a storage media in the media library;

establishing a media session associated with the storage media corresponding to a time period in which the storage media is loaded in a drive of the media library;

polling the media drive during the media session for performance parameters representing the performance of the media drive; and

associating performance data based on the performance parameters with the media session.

2. The monitor appliance of claim 1, wherein the instructions are further executable for presenting the performance data to a user in a manner that allows the use to identify the media session, associated storage media and associated performance data.

3. The monitor appliance of claim 1, wherein the polling the media library further comprises sending a Read Element Status command to the media library.

4. The monitor appliance of claim 1, wherein the media library is a magnetic tape library and the storage media is a magnetic tape.

5. The monitor appliance of claim 1, wherein the polling is out-of-band.

6. The monitor appliance of claim 1, wherein the instructions are further executable for:

determining in to which of a plurality of media drives the storage media has been loaded;

determining which of a plurality of storage medium has been loaded in to the media drive; and

associating the performance with the storage media which has been loaded in to the media drive and with the media drive in to which the storage media was loaded.

7. The monitor appliance of claim 1, wherein the instructions are further executable for creating a list of media sessions sorted by performance.

**8.** The monitor appliance of claim **1** wherein the instructions are further executable for, when the performance is slow, determining whether the cause of the slow performance is external to the media library.

**9.** The monitor appliance of claim **8** wherein the instructions are further executable for correcting the external cause of the slow performance.

**10.** A monitor system for a library, the monitor comprising:  
a processor; and

a computer memory electronically connected to the computer processor, the computer memory storing a set of computer instructions executable by the computer processor to:

poll the library at intervals to determine a position for a storage media in the library,

determine from the polled positions of the media whether the media has been loaded in to, or unloaded from, a drive of the library,

when the media has been loaded in to the drive, begin a media session;

monitor the performance of the library while the media is loaded in the drive and associate performance data determined from the monitoring with the media session;

when the media has been unloaded from the drive, end the media session; and

output information associated with the media session.

**11.** The monitor system of claim **10**, wherein the computer instructions are further executable to determine in to which of a plurality of drives the media has been loaded.

**12.** The monitor system of claim **10**, wherein the computer instructions are further executable to determine which of a plurality of medium has been loaded in to the drive.

**13.** The monitor system of claim **10**, wherein polling the library comprises sending a Read Element Status command to the library.

**14.** The monitor system of claim **10**, wherein the library is a magnetic tape library and the media is a magnetic tape.

**15.** The monitor system of claim **10**, wherein the monitor system is at least a portion of an out-of-band appliance.

**16.** The monitor system of claim **10**, wherein the polling is out-of-band.

**17.** The monitor system of claim **10**, wherein the computer instructions are further executable to:

determine in to which of a plurality of drives the media has been loaded;

determine which of a plurality of medium has been loaded in to the drive; and

associate the performance with the media which has been loaded in to the drive and the drive in to which the media was loaded.

**18.** The monitor system of claim **10**, wherein the computer instructions are further executable to create a list of media sessions sorted by performance.

\* \* \* \* \*