(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
11 July 2002 (11.07.2002)

PCT

(10) International Publication Number
WO 02/054221 A2

(51) International Patent Classification[7]: G06F 3/06

(21) International Application Number: PCT/IL02/00009

(22) International Filing Date: 3 January 2002 (03.01.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/259,158 3 January 2001 (03.01.2001) US

(71) Applicant *(for all designated States except US)*: **NICE SYSTEMS LIMITED** [IL/IL]; 8 Hapnina Street, P O B 690, 43107 Ra'anana (IL).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **KASMIRSKY, Yehoshaphat** [IL/IL]; 25/8 Aharoni Street, ., 76281 Rehovot (IL). **KOR, Ilan** [IL/IL]; 7 Hanafa Street, Tel Mond 40600 (IL).

(74) Agent: **BODNER, Marc**; Plinner, Bodner & Co., Beit Agish-Ravad, 13 Noah Mozes, 67442 Tel Aviv (IL).

(81) Designated States *(national)*: AE, AG, AL, AM, AT, AT (utility model), AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DE (utility model), DK, DK (utility model), DM, DZ, EC, EE, EE (utility model), ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SK (utility model), SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZM, ZW.

[Continued on next page]

(54) Title: CONTENT-BASED STORAGE MANAGEMENT

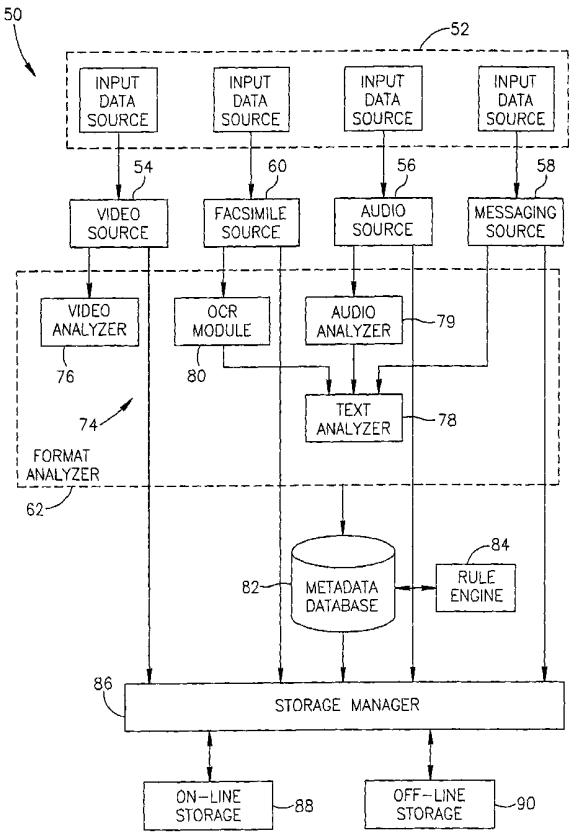

(57) Abstract: A system and method for data management according to the content of the data. The present invention enables data to be stored in one of a plurality of different storage options according to at least one characteristic of the data, in which the at least one characteristic is related to the content of the data. The present invention comprises a rule-based storage management mechanism for the processes of archiving and/or retrieving data. It should be noted that at least one storage option according to the present invention is optionally deletion and/or destruction of the data, such that the data may optionally be removed from storage media or may optionally not be stored initially on the storage media. Optionally and more preferably, the data is stored for a time interval according to the at least one characteristic of the data. Most preferably, the data is moved to a different type of storage option after an event occurs, for example the time interval has elapsed.

WO 02/054221 A2

(84) Designated States *(regional)*: ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published:

— *without international search report and to be republished upon receipt of that report*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

CONTENT-BASED STORAGE MANAGEMENT

FIELD OF THE INVENTION

The present invention relates to a system and a method for content-based storage management, and in particular, for such a system and method in which decisions concerning the location and/or timing of storage of data are based upon the content of the data.

BACKGROUND OF THE INVENTION

Storage facilities for digital information are a critical resource. The demand for storage space for both conventional data, such as text documents and other human readable files, and multimedia streams, such as audio and/or video data, has increased significantly. Such an increase results from a number of different factors, such as legal requirements to store and maintain certain types of information; an increase in the different types of data which are being stored; and even an increase in the size of individual units of data, such as word processing document files, video data files and so forth. This increased demand has in turn resulted in a higher demand for storage space, and in particular for storage space which is accessible "on-line.

As the demand for on-line storage space increases, a number of options are possible to fulfill that demand. For example, additional hardware, such as magnetic media devices ("hard disk drives"), may be purchased to increase the available amount of electronically accessible storage space. However, as the quantity of such hardware devices increases, the management problem for electronic management of these devices also increases. Furthermore, merely increasing the storage space may be both wasteful and unnecessary, since not all of the data may be required, or at least not required for immediate access.

The problem may be partially alleviated through the use of a mixture of different types of storage facilities. For example, on-line storage refers to direct-access, permanently mounted storage areas, such as magnetic (or other types of media) disk drives and disk arrays. The time required for access to such storage areas to be made is typically measured in fractions of a second. Since not all data may need to be stored in on-line storage, which is fast but also expensive, near-line storage is also available. Near-line storage is based upon an automatically (machine) operated storage area, such as optical disks residing on a disk "jukebox" or tapes in an automatic tape library. Such automatically operated storage devices are able to store and

automatically access a relatively large amount of data with fewer physical reading devices, or drives, for reading the data. This type of storage is less expensive, but also somewhat slower for accessing the data, such that access times are measured in seconds to minutes, or even longer, depending upon the availability of physical drives for reading the storage media. On-line or near-line storage may also feature a system with a plurality of physical drives, connected together, for example in a LAN (local area network) or WAN (wide area network).

Off-line storage is the least expensive type of storage, but is also the slowest for access, as it does not permit automatic electronic access. Instead, manual operation of the storage devices and physical drives is required by a human operator. The number of physical drives is greatly reduced compared to the number of storage devices (or at least the amount of available storage space). However, the access time for data from such devices is measured from minutes to hours, depending upon the availability of the human operator and the location of the storage devices, as well as the availability of the physical drives.

Other types of storage devices and functions may also be used, in addition to, or in replacement for, the above-described devices and functions. In any case, the difficulty with a mixed system, or a system in which different types of storage areas (topology) are used, with different types of storage devices and different accessibility (particularly with regard to access time), is the management of the data. Certain types of data may be more important, or at least more time-critical for access, such that the access time may be very important for some types of data, and much less important for other types of data. Cost is also an important factor. Also, decisions must be made concerning the number and type of storage devices to be purchased, along with any required supporting devices and/or system support, such as human operators for example. Currently, these systems are designed and constructed manually, and decisions are made on the basis of some type of policy. However, the operation of the actual system and even the design itself may not be optimal for a particular organization.

SUMMARY OF THE INVENTION

The background art does not teach or suggest a solution to the problem of efficiently managing data storage. The background art also does not teach or suggest a solution to the problem of managing data storage for both cost efficiency and for suitable access times. The background art also does not provide a solution for storing data according to the content of the data, such that important data can be stored in a more accessible location/type of file storage. In addition, the background art does not teach or suggest a system and method for managing data,

such that data is correctly stored, migrated and/or deleted, according to the content thereof.

The present invention overcomes these problems of the background art by providing a system and method for data management according to the content of the data. The present invention enables data to be stored in one of a plurality of different storage options according to at least one characteristic of the data, in which the at least one characteristic is related to the content of the data. It should be noted that at least one storage option according to the present invention is optionally deletion and/or destruction of the data, such that the data may optionally be removed from storage media or may optionally not be stored initially on the storage media. Therefore, a "storage option" according to the present invention includes any type of storage media, device, system or combination thereof, or deletion (removal) of the data.

Optionally and more preferably, the data is stored for a time interval according to the at least one characteristic of the data. Most preferably, the data is moved to a different type of storage option after an event has occurred, for example after the time interval has elapsed. It should be noted that movement or migration of the data may also include deletion or removal of the data.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention is herein described, by way of example only, with reference to the accompanying drawings, wherein:

FIG. 1 is a schematic block diagram of an exemplary system and flow of operations according to the present invention;

FIG. 2 is a schematic block diagram of another exemplary system and flow of operations according to the present invention; and

FIG. 3 is a schematic block diagram of a detailed exemplary system according to the present invention.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

The present invention is of a system and method for data management according to the content of the data. The present invention enables data to be stored in one of a plurality of different storage options according to at least one characteristic of the data, in which the at least one characteristic is related to the content of the data. It should be noted that at least one storage option according to the present invention is optionally deletion and/or destruction of the data, such that the data may optionally be removed from storage media or may optionally not be stored

initially on the storage media. Optionally and more preferably, the data is stored for a time interval according to the at least one characteristic of the data. Most preferably, the data is moved to a different type of storage option after an event has occurred, for example after the time interval has elapsed. It should be noted that movement or migration of the data may also include deletion or removal of the data.

According to preferred embodiments of the present invention, the at least one characteristic of the data, according to which the storage option is selected, is examined by a rule engine. Preferably, the rule engine compares the at least one characteristic of the data to at least one rule, and then selects the storage option (or options) according to that rule. The rule engine therefore more preferably operates as a filter, for determining which storage option(s) is most appropriate for the examined data. The storage decision is then preferably implemented by a storage manager.

The rules according to which the rule engine operates are optionally manually entered by a human operator, or alternatively may optionally be generated automatically according to a predefined business rule or according to an automatically generated business rule, or a combination thereof.

Preferably, the present invention is operative with a system featuring a plurality of different storage options. More preferably, these different storage options include at least two different storage options having different types of accessibility. Examples of storage options having different types of accessibility include but are not limited to on-line storage, near-line storage and off-line storage. The type of storage media which is used for any particular storage option is not limited according to the present invention, as the present invention is operable with any suitable type of storage media, including but not limited to DAT (tape-based storage), AIT (also tape-based storage), magnetic storage media, optical disks, CD-ROM or a mass storage device of any type, or any type of storage system, or any combination thereof.

The at least one characteristic of the data, which is related to the content of the data, may optionally and preferably be obtained in a number of different ways. For example, the data may optionally and preferably have associated metadata, which is related to the content of the data. The metadata is more preferably added through annotation of the data itself. Such annotation is optionally performed manually, through human intervention, but is preferably performed automatically. More preferably, automatic annotation is performed after the data is automatically analyzed. The associated metadata is then preferably used to determine which storage option should be used for the data, and more preferably also the time interval during which the data

should be placed in that storage option.

As previously described, the data is more preferably filtered by a rules engine, according to at least one characteristic of the data. For the implementation of the present invention with metadata, the filtering process is more preferably performed according to the associated metadata.

Automatic analysis of the data is more preferably performed according to the type of data being analyzed. Examples of different types of automatic analysis processes which may optionally be performed include but are not limited to, speech-to-text conversion for voice communication data, a video analyzer for video data, OCR (optical character recognition) for printed matter which has been electronically scanned, image analysis for image data, text analysis for textual data, and analyzers for user interface data. These different types of data analysis processes are preferably performed according to the source of data, which may optionally be any suitable data source. Examples of different types of data sources include but are not limited to, video data, audio data (including also voice communication data such as voice over IP (VoIP) data, streaming audio data and any other type of audio-related data), coded data, e-mail messages and/or attachments, chat and other types of messaging system messages, documents transmitted by facsimile and user interface data.

In addition, the present invention is useful for the collection of data about substantially any type of user interface function. Examples of such user interface functions include but are not limited to any type of GUI window activity; activity with GUI gadgets such as buttons, sliders or any function provided through a GUI window; the display of any image and/or text, including but not limited to Web pages and/or any component thereof; information provided through an audible interface such as a synthesized voice; information provided through the display of video data; and any type of information which is provided through, or otherwise detectable by, the operating system of the user computational device.

According to preferred embodiments of the present invention, the data is preferably "migrated" or moved from a first storage option to a second, different storage option after a time interval has expired. The time interval is preferably determined according to the metadata.

The principles and operation of the method according to the present invention may be better understood with reference to the drawings and the accompanying description.

Referring now to the drawings, Figure 1 shows a first exemplary system **10** for managing data according to the content of the data, with regard to the flow of operations through system **10**. As shown, system **10** features at least one input data source **12**. Examples of different types

of input data sources **12** include but are not limited to, video data, audio data (including also voice communication data such as voice over IP (VoIP) data, streaming audio data and any other type of audio-related data), coded data, e-mail messages and/or attachments, chat and other types of messaging system messages, documents transmitted by facsimile and user interface data. With regard to user interface data, in which the action of a user upon operating a computational device and/or a peripheral device and/or input device thereof causes the data to be generated, optionally and preferably the data is in the form of an event. Each action of the user preferably causes an event to be generated. The event may then optionally form the data to be captured.

The captured data is optionally and preferably passed to a format analyzer **14** for rendering the captured data into a common format for analysis. Format analyzer **14** preferably features a plurality of format modules **16**, each of which is suitable for data of a different type of format. For example, if the input data is voice communication data, then preferably a format module **16** converts the voice communication to textual data, for speech-to-text conversion. Different format modules **16** preferably handle other types of input data, as explained in greater detail with regard to Figure 3 below.

According to a preferred embodiment of the present invention, the common data format is optionally and preferably textual data. For this preferred embodiment, textual data is optionally not further preprocessed by a format module **16**, or alternatively is only minimally processed. Other types of data in different data formats are then preferably converted to textual data by format module **16**, as described with regard to Figure 3 below.

Next an analysis module **18** preferably analyzes the data, once the data is in the common format. It should be noted that optionally analysis module **18** is able to handle a plurality of different data formats, through a single module or alternatively from a set of such modules (not shown). Preferably, analysis module **18** operates on a plurality of different types and/or sources of data simultaneously, for example as a multi-thread application. For the non-limiting, illustrative example, as described above, textual data is analyzed, as is well known in the art. One non-limiting example of a text analyzer software is the IntelligentMiner™ product (IBM Corp, see http://www-4.ibm.com/software/data/iminer/fortext/tatools.html as of December 31, 2001 for further details). This product is useful for analyzing text for a number of functions, such as locating information related to a topic, categorization of information and classification of information. Text analyzer tools are generally known in the art for extracting content and/or information related to the subject matter of text, for example according to one or more keywords,

concepts or any other type of organization and/or analysis scheme.

According to another preferred embodiment of the present invention, the uniform format features a uniform data structure, with a plurality of different types and/or categories of information, for example data from screen events and voice data combined in a single file. This uniform data structure preferably is able to contain the different types or categories of characteristics which are of interest for being associated with the data, in order to determine the content of the captured data. Non-limiting examples of a uniform data structure which may optionally be implemented according to the present invention include structures which use XML (extensible mark-up language) or ASF (Advanced Streaming Format, from Microsoft Corp., USA).

Analysis module **18** preferably extracts and/or creates, or otherwise determines, at least one characteristic of the captured data, preferably obtained in the uniform format. Optionally and more preferably, analysis module **18** obtains the at least one characteristic from the captured data in the form of metadata. This metadata is then optionally stored in a metadata database (not shown, see Figure 3).

According to an optional but preferred embodiment of the present invention, analysis module **18** also gives feedback for improving the performance of format analyzer **14** and/or format module **16**, in order to improve the operation of these components.

The captured data, with the at least one characteristic related to the content of data, which is optionally and more preferably in the form of metadata, is then passed to a rule engine **20**. Alternatively, rule engine **20** only receives the at least one characteristic, more preferably in the form of metadata. The captured data is then passed directly to a storage manager **22**. Rule engine **20** more preferably compares the metadata to at least one rule, which is most preferably a business rule specified by a manual human user, or alternatively may optionally be generated automatically according to a predefined business rule or according to an automatically generated business rule, or a combination thereof. Optionally, one or more rules may be fed to rule engine **20** through an interface **24**, such as a GUI (graphical user interface) for example. As a non-limiting example, interface **24** may optionally be a simple Web browser-based interface. Rule engine **20** then preferably determines the type of storage option (or options) according to one or more rules, as selected through the comparison of the metadata (or characteristic) of the captured data to the rule(s). Additionally or alternatively, the output of rule engine **20** is optionally and preferably fed back to format analyzer **14**, and/or format module **16**, and/or analysis module **18**.

Rule engine **20** optionally and more preferably determines both the type of storage option (or options), which should be selected for the particular captured data, and also the term of storage. Most preferably, the captured data is initially stored with a first storage option, and then is migrated (moved) to at least one additional storage option after a period of time has elapsed. This period of time is also most preferably determined according to at least one rule by rule engine **20**.

Rule engine **20** may optionally perform an action according to a rule and/or event, in which the event may optionally and preferably trigger automatic application of the rule. Examples for actions based on rules are given below. One example of actions which are based on rules include but are not limited to manipulations of stored data. For example, the compression of the data may optionally be altered after an event has occurred, such as a period of time has elapsed. A non limiting example of a reason for altering such compression is to enable more rapid playback of the data.

Previously stored data may optionally and preferably be updated with business data. Non limiting examples of such data include the addition of social security or identification number, customer identifier, preferred customer status information and so forth.

Current transaction data may also optionally and preferably be linked to the previously stored data file. Previously stored data may also optionally be updated. A non limiting example of such linking may optionally be performed by linking transactions performed by a certain high-status or preferred customer to past transaction by that customer.

Another event/action example may optionally be performed with multiple mirroring of data, for example by distributing identical data to several destinations and/or pre-defined locations and/or storage options. Such mirroring may optionally be performed for redundancy purposes, for example for security of the stored data, by duplicating to multiple storage locations/options, and/or for general availability reasons.

Storage manager **22** preferably then receives the output of rule engine **20** and the captured data. The output of rule engine **20** preferably includes at least one storage option for the captured data. As previously noted, this storage option could be a type of storage media or deletion and/or removal and/or destruction of the captured data. More preferably, the storage option includes a particular storage device **26** into which the captured data should be placed. Alternatively, storage manager **22** could determine the identity of the particular storage device **26** for storing the data.

Figure 2 shows a different flow arrangement of the system of Figure 1, as another example of a system according to the present invention. In this arrangement, system **10** again features a rule engine **40**. However, in this implementation, rule engine **40** is the initiating process or component for subsequent actions which are performed by system **10**. As show, rule engine **40** takes the input source according to metadata obtained from an analysis module **38**. Rule engine **40** then preferably sends the captured data, or alternatively only selected captured data, to a storage manager **42**. Storage manager **42** sends the captured data to the correct storage option, shown as preferably being a selected storage device **44**, according to a request for action by rule engine **40**. Optionally and more preferably, and most preferably as necessary, rule engine **40** feeds back the captured data, and/or information about the captured data, into the input sources. Improved metadata may optionally be obtained from analysis module **38**.

According to an optional implementation of the present invention, the user defines a task in rule engine **40** for archiving certain types of information, such as information about specific telephone calls. Rule engine **40** then preferably uses analysis module **38** to select specific data. Analysis module **38** may optionally be implemented as a call management server, for example. The selected specific data may optionally be any one or more of voice data, data captured from user interface actions, video data, an e-mail transaction, facsimile data, VoIP, Web-co browsing data (obtained from two or more users viewing the same Web page(s) through different Web browser processes), or any coded data or any combination of any type of input sources. The data is obtained from input sources **36**, which may optionally be implemented as an input sources logger. The captured data is then transferred into a storage manager **42**.

This data can optionally be retrieved as required from storage devices **44**, more preferably directly by using rule engine **40** and/or storage manager **42**. Such retrieved data may then optionally and more preferably be fed into input sources **36** in order for the retrieved data to be widely accessible (available), and/or for further manipulation.

Figure 3 shows an exemplary, detailed implementation of a system according to the present invention, which is related to the flow of operations shown in Figure 1. As shown, a system **50** preferably features a plurality of input data sources **52**. For the purposes of explanation only and without any intention of being limiting, input data sources **52** are shown as optionally and preferably including a video source **54** for video data, an audio source **56** for audio data, a messaging source **58** for e-mail messages (including attachments), instant messaging and/or chat data, and a facsimile source **60** for data which is transmitted by facsimile.

The input data from input data sources **52** is preferably then fed to at least one format analyzer **62** for rendering the captured data into a common format for analysis. Format analyzer **62** preferably features a plurality of format modules **74**, each of which is suitable for data of a different type of format. For the purposes of illustration only and without any intention of being limiting, format modules **74** are shown as optionally and preferably including a video analyzer **76**, a text analyzer **78**, an audio analyzer **79** and an OCR (optical character recognition) module **80**.

An example of text analyzer **78** was previously described with regard to Figure 1. OCR module **80** may optionally be implemented as is well known in the art, for example through the use of OCR software having an algorithm which could easily be selected by one of ordinary skill in the art.

Video analysis may optionally be performed by video analyzer **76** as follows. Video data is obtained, for example from a camera as a non-limiting example of video source **54**. A frame-grabber is then preferably used to obtain at least one frame from the video data. The frame is preferably analyzed. More preferably, only a portion of the frame is stored as captured data. For example, if a video camera is used to monitor the entrance to a secure area, then optionally only those frames, or alternatively those portions of each frame, which feature a human subject near the actual entrance are of interest. Additionally or alternatively, changes in the background of each frame may optionally be detected and tracked, as being of interest.

One example of a type of analysis which may be performed with the video data is a motion detection algorithm, which is well known in the art. Another example is face recognition algorithms, which are also well known in the art. Non-limiting examples of video analysis algorithms are described at http://www.cs.rochester.edu/u/nelson/research/motion/motion.html as of December 31, 2001, for motion detection algorithms and at http://www-white.media.mit.edu/vismod/demos/facerec as of December 31, 2001, for face recognition algorithms. More preferably, such analyses are performed with firmware, such as a DSP (digital signal processor) for example. The results may then optionally be stored as the captured data.

The output of format analyzer **62** preferably features at least one characteristic of the captured data which more preferably is metadata, as previously described. The metadata is more preferably stored in a metadata database **82**. A rule engine **84** then preferably analyzes the metadata from metadata database **82** (or alternatively obtained directly from format analyzer **62**), in order to apply one or more rules to the captured data. Rule engine **84** may optionally be

implemented with the BlazeSoftware Advisor product of Blaze Software (see http://www.blazesoft.com/products as of December 31, 2001 for details). Rule engine **84** may also optionally be implemented as a rule/task engine from Nice Systems Ltd. (Ra'anana, Israel), for example based on business data. Regardless of the specific implementation, rule engine **84** preferably operates according to at least one business rule.

Rule engine **84** preferably compares the at least one characteristic of the data to at least one rule, and then selects the storage option (or options) according to that rule. Rule engine **84** therefore more preferably operates as a filter, for determining which storage option(s) is most appropriate for the examined data. The storage decision is then preferably implemented by a storage manager **86**.

Preferably, storage manager **86** is able to select from a plurality of different storage options. More preferably, these different storage options include at least two different storage options having different types of accessibility. Examples of storage options having different types of accessibility include but are not limited to on-line storage, near-line storage and off-line storage. The type of storage media which is used for any particular storage option is not limited according to the present invention. For the purposes of illustration only and without any intention of being limiting, storage manager **86** is shown as being able to select from a plurality of storage devices, shown as an on-line storage device **88** and an off-line storage device **90**. Of course, other types of storage devices and/or systems could be used in place of, or in addition to, these examples of storage devices.

Rule engine **84** is optionally and preferably able to feed back information to format analyzer **62**, for improving the performance of format analyzer **62**.

A non-limiting example of the operation of system **50** may be performed as follows. System **50** could optionally be implemented at a service center which processes services requests from customers remotely, such that the customer is not physically present at the service center. The customer therefore contacts service center personnel, for example through voice communication (such as a telephone call for example), e-mail messages, facsimiles and so forth. A plurality of business rules has been defined and implemented by rule engine **84**, which could optionally include the following rules: a record is kept for every customer contact that refers to a financial transaction for at least three months even if no transaction occurred, and is kept for each contact in which a financial transaction occurred for at least seven years. In addition, the record for each contact resulting in an actual financial transaction is first stored in on-line storage

for one month, and then in an off-line storage for the remainder of the term to seven years.

Once the customer has contacted a service center operator, for example through the telephone, data is provided through audio source **56** as an example of input source **52**. This captured audio data is analyzed, for example in order to determine if the financial transaction occurred during the contact. If such a transaction occurred, then metadata associated with the captured audio data indicates such an occurrence. Format analyzer **62**, and particularly audio analyzer **79**, preferably analyzes the captured audio data to obtain such metadata.

The data itself from the call is preferably handled according to one or more business rules, which may optionally be defined manually and/or generated automatically, through the operation of rule engine **84**. Preferably, rule engine **84** then generates an action to be performed by storage manager **86**. For example, storage manager **86** may optionally store the data from the call, migrate the data to a new type of storage, or delete the data, or any other action or any combination of actions which should be performed according to one or more events. Thus, rule engine **84** is able to generate one or more instructions for execution by storage manager **86**.

While the invention has been described with respect to a limited number of embodiments, it will be appreciated that many variations, modifications and other applications of the invention may be made.

WHAT IS CLAIMED IS:

1. A method for managing data storage according to content of the data, comprising:

determining at least one characteristic of the data according to the content;

selecting one of a plurality of storage options according to said at least one characteristic of the data; and

placing the data into said selected storage option.

2. The method of claim 1, wherein said selected storage option causes deletion of the data.

3. The method of claims 1 or 2, wherein said plurality of storage options include storage options having at least two different types of devices.

4. The method of claim 3, wherein said different types of devices have different access times.

5. The method of any of claims 1-4, wherein at least one storage option includes an on-line storage device.

6. The method of any of claims 1-5, wherein at least one storage option includes an off-line storage device.

7. The method of any of claims 1-6, wherein at least one storage option includes a near-line storage device.

8. The method of any of claims 1-7, wherein said at least one characteristic of the data includes metadata associated with the data.

9. The method of claim 8, wherein said metadata is obtained by analyzing the data.

10. The method of claim 9, wherein the data is analyzed manually.

11. The method of claims 9 or 10, wherein the data is analyzed automatically.

12. The method of claim 11, wherein the data is analyzed automatically according to a type of the data.

13. The method of claims 9-12, wherein the data includes a plurality of different types of data, and said plurality of different types of data is analyzed concurrently.

14. The method of claims 11-13, wherein the data is rendered into a common format before being analyzed automatically.

15. The method of claims 11-13, wherein the data is rendered into a common format after being analyzed automatically.

16. The method of any of claims 1-15, wherein said at least one characteristic of the data is compared to at least one rule for determining said storage option to be selected.

17. The method of claims 15 or 16, wherein said at least one rule includes a time interval for holding the data in said selected storage option.

18. The method of claim 17, wherein the data is migrated from a first selected storage option to a second selected storage option after said time interval has elapsed.

19. The method of any of claims 16-18, wherein said at least one rule is entered manually.

20. The method of any of claims 16-18, wherein said at least one rule is generated automatically.

21. The method of claim 20, wherein said at least one rule is generated automatically according to business data.

22. The method of any of claims 16-21, wherein said at least one rule includes an

action to be performed on the data according to an event, wherein said event is related to said at least one characteristic of the data.

23. The method of any of claims 1-22, further comprising receiving data from an input source before determining said at least one characteristic of the data, wherein said input source includes at least one of video data, audio data, coded data, e-mail messages, e-mail attachments, chat messages, other types of messaging system messages, documents transmitted by facsimile and user interface data, and a combination thereof.

24. The method of any of claims 1-7, wherein the data is analyzed manually to determine said at least one characteristic of the data.

25. The method of any of claims 1-7, wherein the data is analyzed automatically to determine said at least one characteristic of the data.

26. The method of claims 24 or 25, wherein the data is obtained from a plurality of input sources, each input source providing a different type of data.

27. The method of claim 26, wherein the data is analyzed automatically according to a type of the data.

28. The method of claims 26 or 27, wherein the data is rendered into a common format before being analyzed automatically.

29. The method of claims 26 or 27, wherein the data is rendered into a common format after being analyzed automatically.

30. The method of any of claims 9-29, wherein feedback from an analysis of the data is used for determining said at least one characteristic of the data.

31. A system for data management according to content of the data, comprising:

(a) an analyzer for analyzing the data to determine at least one characteristic of the

data;

(b) a rule engine for comparing said at least one characteristic of the data to at least one rule; and

(c) a storage manager for receiving a decision from said rule engine for at least storing the data.

32. The system of claim 31, further comprising a client, wherein said rule engine determines if data is to be retrieved to said client.

33. The system of claims 31 or 32, further comprising an input source for providing data to said analyzer, wherein said rule engine determines if the data is to be used as feedback to said input source.

34. The system of any of claims 31-33, wherein an operation of said rule engine is manually triggered.

35. The system of any of claims 31-33, wherein an operation of said rule engine is automatically triggered.

36. The system of claim 35, wherein said rule engine is an initiator of a process for at least storing the data.

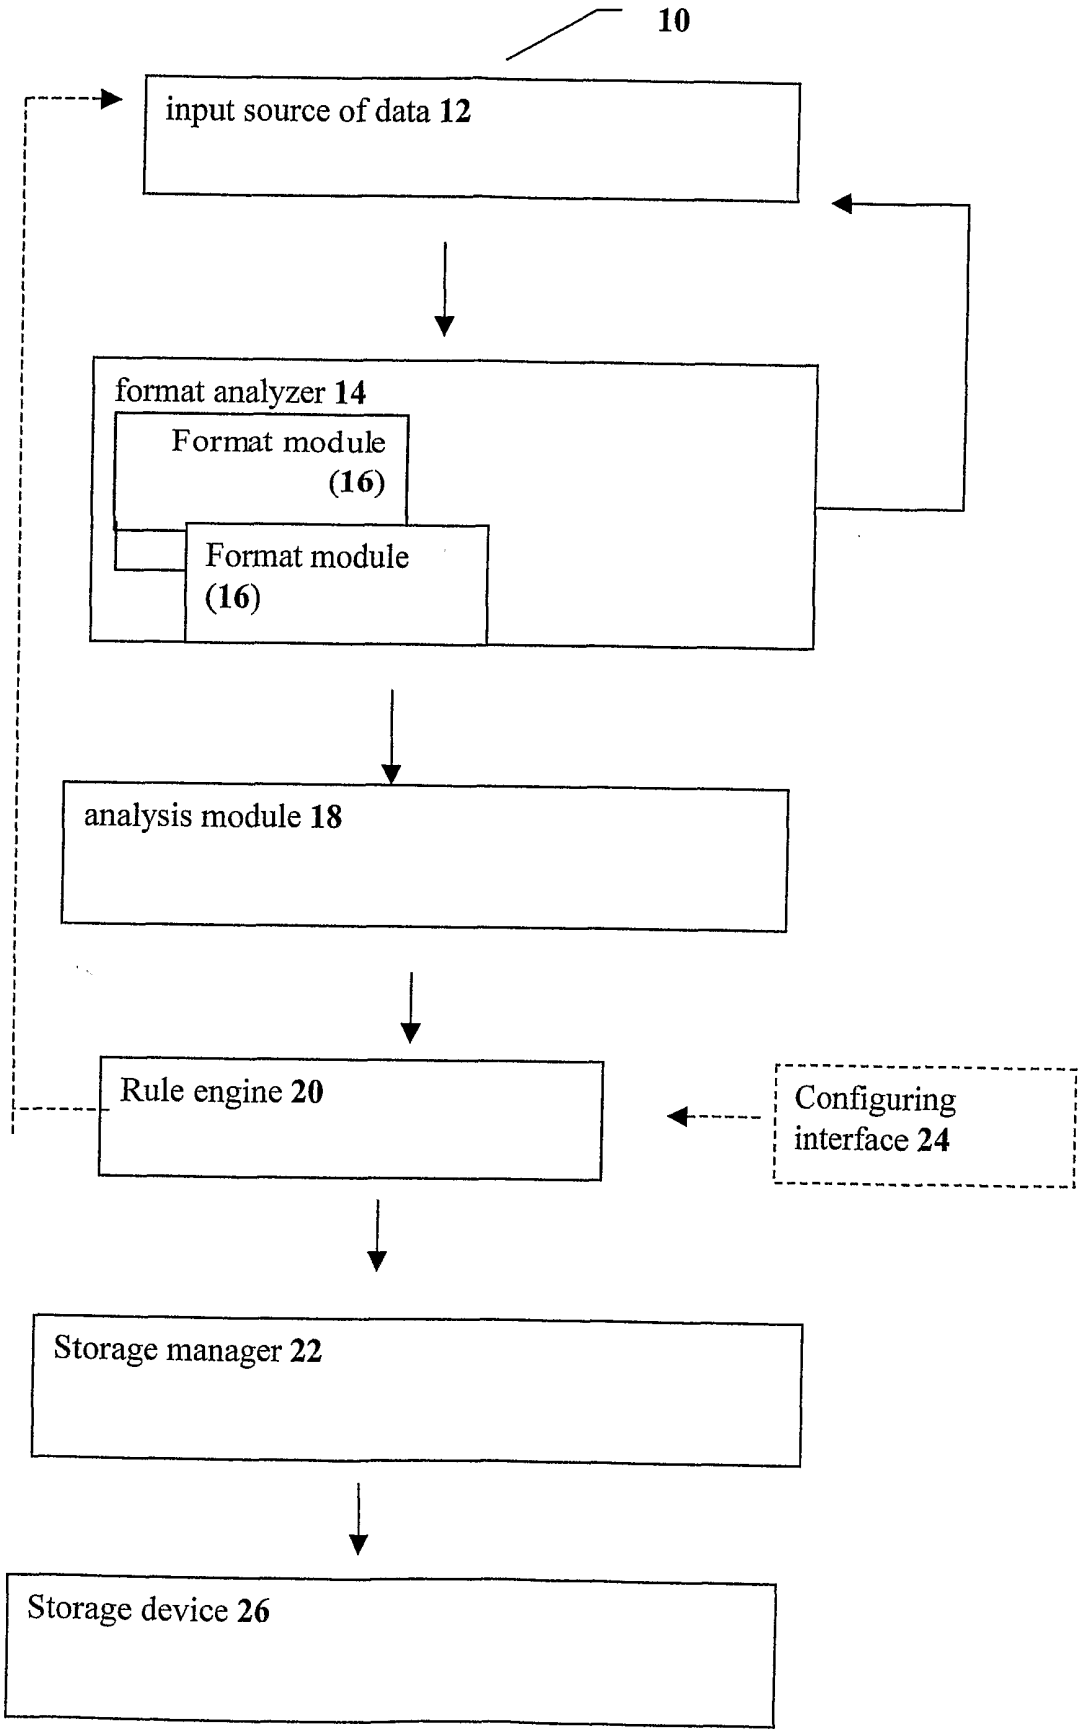
10
input source of data 12
format analyzer 14
Format module (16)
Format module (16)
analysis module 18
Rule engine 20
Configuring interface 24
Storage manager 22
Storage device 26

Figure 1

Figure 2

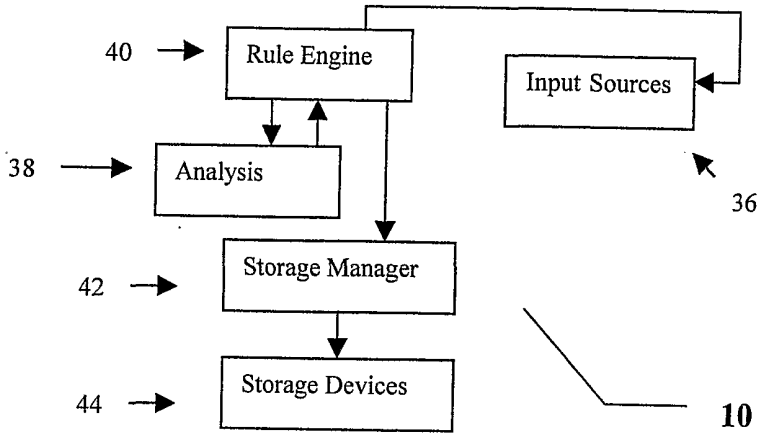
40
Rule Engine
Input Sources
38
Analysis
36
Storage Manager
42
Storage Devices
44
10

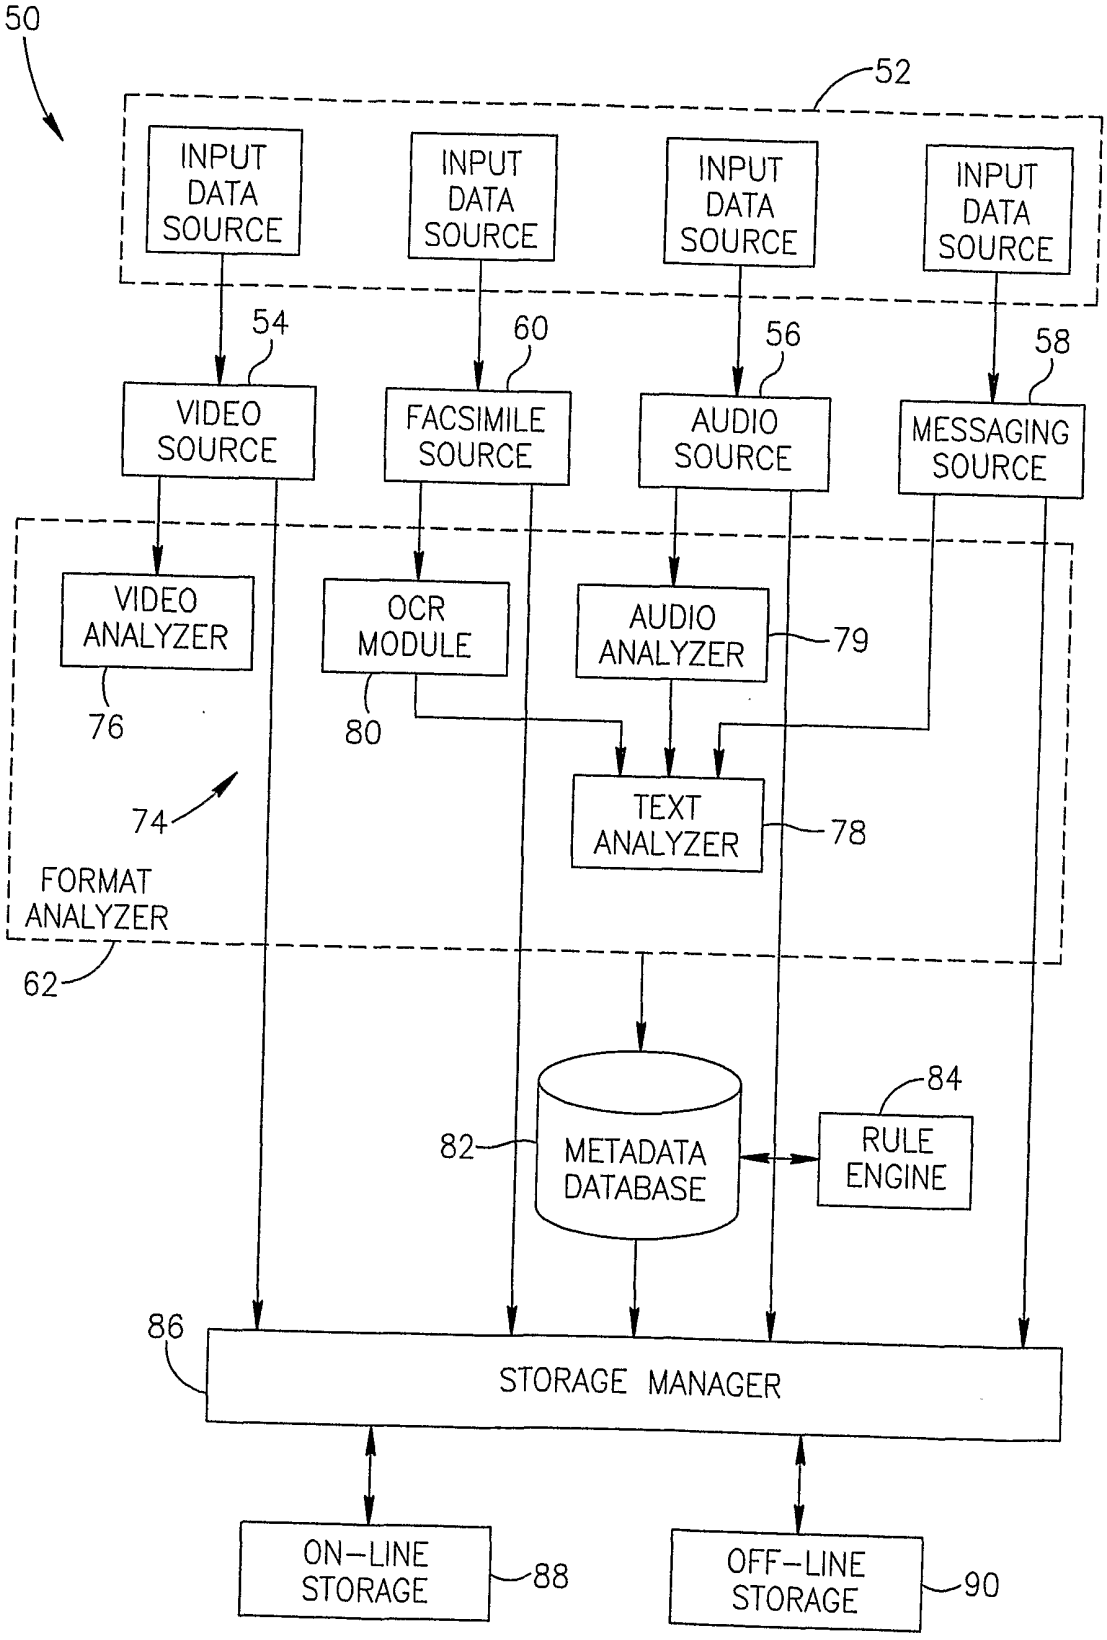
50
52
INPUT DATA SOURCE
INPUT DATA SOURCE
INPUT DATA SOURCE
INPUT DATA SOURCE
54
60
56
58
VIDEO SOURCE
FACSIMILE SOURCE
AUDIO SOURCE
MESSAGING SOURCE
VIDEO ANALYZER
OCR MODULE
AUDIO ANALYZER
79
76
80
74
TEXT ANALYZER
78
FORMAT ANALYZER
62
84
82
METADATA DATABASE
RULE ENGINE
86
STORAGE MANAGER
ON-LINE STORAGE
88
OFF-LINE STORAGE
90

FIG.3