



- (51) **International Patent Classification:**
G06F 13/14 (2006.01)
- (21) **International Application Number:**
PCT/US20 14/035 134
- (22) **International Filing Date:**
23 April 2014 (23.04.2014)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (71) **Applicant:** HEWLETT-PACKARD DEVELOPMENT COMPANY, L.P. [US/US]; 11445 Compaq Center Drive W., Houston, Texas 77070 (US).
- (72) **Inventors:** BALLARD, Curtis, C ; 3404 E Harmony Road, Fort Collins, Colorado 80528 (US). PICKETT, Seth; 3404 E Harmony Road, Fort Collins, Colorado 80528 (US).
- (74) **Agents:** ORTEGA, Arthur et al; Hewlett-Packard Development Company, L.P., Intellectual Property Administration, Mail Stop 35 3404 E. Harmony Road, Fort Collins, Colorado 80528 (US).
- (81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM,

DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

- (84) **Designated States** (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- as to the identity of the inventor (Rule 4.1 7(I))
- as to applicant's entitlement to apply for and be granted a patent (Rule 4.1 7(H))

Published:

- with international search report (Art. 21(3))

(54) **Title:** DETERMINING CONNECTIVITY INFORMATION FOR A STORAGE DEVICE

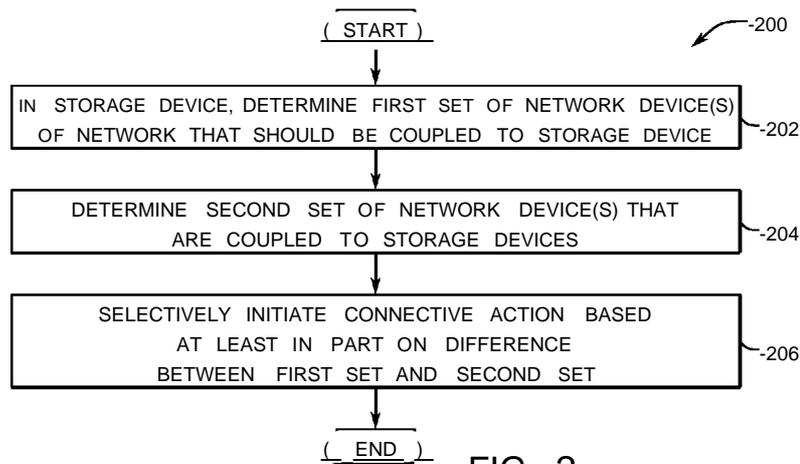


FIG. 2

(57) **Abstract:** A method includes, in a storage device coupled to a storage network, determining connectivity information for the storage device. The determining the connectivity information includes determining a first set of at least one network device coupled to the storage device and determining a second set of at least one network device that is coupled to the storage device. The technique includes selectively initiating corrective action based at least in part on a comparison of the first set with the second set.



- 1 -

DETERMINING CONNECTIVITY INFORMATION FOR A STORAGE DEVICE

Background

[0001] A computer may access a storage area network (SAN) for purposes of storing and retrieving large amounts of data. The typical SAN includes a consolidated pool of mass storage devices (magnetic tape drives, hard drives, optical drives, and so forth); and the SAN typically provides relatively high speed block level storage, which may be advantageous for backup applications, archival applications, database applications and other such purposes.

- 2 -

Brief Description of the Drawings

[0002] Fig. 1 is a schematic diagram of a computer system according to an example implementation.

[0003] Figs. 2 and 3 are flow diagrams depicting techniques to determine connectivity information for a storage device of a storage area network (SAN) according to example implementations.

[0004] Fig. 4 is a flow diagram depicting a technique to detect and report a failed connection path according to an example implementation.

[0005] Fig. 5 is a schematic diagram of a storage device of Fig. 1 according to an example implementation.

Detailed Description

[0008] Referring to Fig. 1, in accordance with example implementations, a computer system 100 includes host resources 102 and storage resources 150 that may be configured in numerous ways to provide datacenter or cloud services for clients (not shown). As examples, the computer system 100 may provide such services as Software as a Service (SaaS), Infrastructure as a Service (IaaS) or Platform as a Service (PaaS). For the example of Fig. 1, the storage resources 150 include P physical storage devices 154 which are identified by reference numerals 154-1, 154-2, . . . 154-P in Fig. 1); and the host resources 102 include N hosts 110 (hosts 110-1, 110-2, . . . 110-N-1 and 110-N, being depicted as examples in Fig. 1). In general, the computer system 100 may be locally disposed at a given site or may be geographically distributed at multiple locations, depending on the particular implementation.

[0009] Clients (not shown in Fig. 1) of the system 100, such as thin clients, tablets, portable computers, smartphones, desktop computers, and so forth may access (via a network fabric not shown in Fig. 1) the computer system 100 for purpose of using the hosted services. Depending on the particular service for a given client, one or multiple storage devices 154 may be pooled together. Moreover, a given host 110 may be shared by clients, a given storage device 154 may be shared by multiple clients, and so forth.

[0010] As depicted in Fig. 1, the storage resources 150 are part of a storage area network (SAN) 120. The SAN 120 has several points where network connections are made; and in general, a new connection is formed whenever a given network device (a host 110 or storage device 154, as examples) is attached, or coupled, to the SAN 120. The computer system 100 may have hundreds of storage devices 154 and may also have a relatively large number of hosts 110. Due to the resulting complexity of the computer system 100, a given host 110 may miss a network notification that a new storage device 154 has been attached and as a result, the host 110 may fail to connect to the storage device 154. Moreover, a given host 110 may miss a momentary disconnection with a storage device 154 and fail to

- 4 -

reconnect to the device 154 when the device 154 becomes available. It is also possible that the same result may occur due to other causes (hardware or software defects, for example). It is noted that these are merely examples of points of the SAN 120 where network connections are made.

[001 1] One solution for a host 110 failing to connect to a storage device or for the failure of a connection in the SAN 120 preventing such a connection, is for a human storage administrator to perform relatively complex troubleshooting for purposes of identifying the underlying cause of the connectivity problem. In this troubleshooting, if a new storage device 154 is attached and the connections do not work correctly, then an assumption may be incorrectly made that the newly-attached storage device 154 is faulty, even though another problem of the SAN 120 may be preventing the connection.

[001 2] In accordance with example systems and techniques that are disclosed herein, the storage device 154 is constructed to initiate and perform a connectivity analysis so that this connectivity analysis may be used (by a human storage administrator or by an automated component, as examples) to resolve the connectivity problem relatively quickly and accurately. As a result, replacement of a non-defective storage device 154 may be avoided.

[001 3] More specifically, referring to Fig. 1 in connection with Fig. 2, in accordance with example implementations, a given storage device 154 is constructed to perform a technique 200 for purposes of initiating and performing a connectivity analysis for the device 154. Pursuant to the technique 200, the storage device 154 undertakes measures to determine (block 202) one or multiple network devices that should be coupled to the storage device 154. The storage device 154 further undertakes measures to determine (block 204) a second set of one or multiple network devices that are coupled to the storage device 154. Based at least in part on the difference(s) between the first and second lists, the storage device 154 selectively initiates corrective action, pursuant to block 206. As further described herein, this corrective action may include, as examples, generating a report identifying the network device(s) that should be but are not coupled to the storage

- 5 -

device 154; communicating an alert message to a storage administrator; alerting a connectivity analysis or repair engine; and so forth.

[0014] Referring back to Fig. 1, as a more specific example, in accordance with example implementations, the computer system 100 may contain switch fabric 124, which, in general, represents the network cables, switches, gateways, bridges, routers and so forth that couple the storage devices 154 to the hosts 110. The switch fabric 124 may contain one or multiple types of network fabric, such as a local area network (LAN) fabric, wide area network (WAN) fabric, Internet-based fabric, Fibre Channel (FC) fabric, Small Computer System Interface (SCSI) fabric, Fibre Channel over Ethernet (FCOE), a combination of one or more these fabrics, and so forth.

[0015] For the following example implementation, at least one of the storage devices 154 contains a connectivity analysis and reporting engine 160, or "engine 160." In accordance with example implementations, the engine 160 is constructed to use a management interface (a management application programming interface (API), for example) of the SAN 120 for purposes of acquiring connectivity information data for the storage device 154 and use the acquired connectivity information data to at least make a preliminary analysis of any connectivity issues associated with the storage device 154.

[0016] More specifically, although a management interface may be used by network components (a switch, for example) of the SAN 120 other than the storage devices 154 to initiate logins with the storage devices 154, in accordance with example implementations, the engine 160 makes use of the management interface to initiate a connection with the SAN 120 for purposes of accessing the SAN's network management functions. In this manner, in accordance with example implementations, the engine 160 may initiate a connection between the storage device 154 and the SAN 120 using, for example, a port login request for a FC SAN.

[0017] As a more specific example, in accordance with example implementations, the switch fabric 124 may include a switch 130 that provides a network management service, such as a "Nameserver" service 132, which may be accessed by the engine 160 for purposes of acquiring connectivity information. In this manner, the switch

- 6 -

130 contains a memory 134 that stores connectivity data 136 that is acquired by the Nameserver service 132 and which identifies the network devices that should be coupled to the storage device 154, among other connectivity data. The engine 160 accesses the Nameserver service 132 to retrieve the connectivity information data 136.

[001 8] In accordance with example implementations, the storage device 154 contains a hardware interface 155 that is constructed to receive connections from host computers. The interface 155 is further constructed to initiate connections, such as initiating a connection to a switch, and in accordance with further example implementations, the interface 155 is constructed to recover a lost connection to a host by reinitiating the connection. As an example, the interface may be constructed from logic that forms state machines to receive and initiate connections, according to the protocol that is used by the switch fabric 124.

[001 9] In accordance with example implementations, after retrieving the connectivity information data 136, the engine 160 parses the data 136 for purposes of building a table of network devices that should be (according to the data 136) coupled to the storage device 154. Using this constructed table, the engine 160 compares the network devices that should be coupled to the storage device 154 with a list of network devices that actually are coupled to the storage device 154. In this manner, in accordance with example implementations, the list of network devices that are actually coupled to the storage device 154 may be stored in an internal memory of the storage device 154.

[0020] By comparing the list of network devices that are coupled to the storage device 154 with the list of network devices that should be coupled to the storage device 154, the engine 160 identifies any differences and flags these differences as potential connectivity issues. Moreover, in accordance with example implementations, the engine 160 may generate data (data representing a connectivity report, for example) that highlights any connectivity issues that are identified by the engine's analysis.

[0021] The engine 160 may initiate one or multiple corrective actions based on the detected differences between which network devices should be coupled to the

- 7 -

storage device 154 and which network devices are coupled to the storage device 154. For example, in accordance with some implementations, in response to detecting discrepancies, the engine 160 generates an alert message that is communicated via the switch fabric 124 to a monitoring station 114 for the SAN 120. In accordance with example implementations, in response to receiving the alert message, the monitoring station 114 generates an alert message 115 (an SMS message, an electronic message (email), multiple different type messages and so forth) to a storage administrator so that the administrator may take the appropriate action(s). In further implementations, the monitoring station 114 may communicate an alert message to an automated component for purposes of addressing the connectivity issue. Moreover, in accordance with example implementations, the engine 160 may store or communicate data representing a report that identifies any potential connectivity issues. This report data may accompany the alert message; or in accordance with further example implementations, the engine 160 may store data representing the report in a memory of the storage device 154 for subsequent retrieval by a storage administrator, for example.

[0022] Thus, referring to Fig. 3 in conjunction with Fig. 1, in accordance with example implementations, the engine 160 may perform a technique 300. Pursuant to the technique 300, the engine 160 uses (block 302) a storage device of a storage network to initiate a connection to the storage network (using the interface 155, for example) for purposes of accessing (block 304) management functions to retrieve connectivity information data. The engine 160 then parses (block 306) the retrieved connectivity information data to build a table of one or multiple network devices that should be coupled to the storage device 154. The engine 160 further retrieves (block 308) connectivity data from an internal memory of the storage device 154 to compare the network device(s) that are coupled to the storage device 154 with the network device(s) that should be coupled to the storage device. The engine 160 selectively communicates (block 310) an alert to a management station based at least in part on this comparison.

[0023] In accordance with further example implementations, the engine 160 performs a corrective action that includes identifying one or possibly multiple

- 8 -

connection paths that caused the failure of the storage device 154 to be coupled to a given network device. For example, in accordance with example implementations, the engine 160 may attempt to connect to a given network device and collect a record detailing the step(s) in which the connection process that succeeded and the step(s) in which the connection process failed.

[0024] As a more specific example, in accordance with example implementations, the engine 160 is constructed to detect "hops" in which a connection request may be communicated through multiple connection paths, or points, to a given network device. In this manner, by analyzing the hops, the engine 160 may generate a report identifying connection path(s) in which the request failed. If a request is successful for a given connection path, then the problem may have been automatically resolved or at least the physical connections that are part of the successful connection path may be eliminated as the source of the problem.

[0025] Thus, referring to Fig. 4 in conjunction with Fig. 1, in accordance with example implementations, the engine 160 performs a technique 400 that includes attempting (block 402) to communicate with a network device that should be coupled to the storage device 154 but is not coupled to the storage device 154. The engine 160 determines (block 404) which connection path(s) failed based at least in part on hops data. From this information, the engine 160 may generate (block 406) a report identifying the failed connection path(s).

[0026] Referring to Fig. 5, in accordance with example implementations, the storage device 154 is a sequential access medium device, such as a magnetic tape drive, and as such, the device 154 includes a bay to receive removable media, such as a physical magnetic tape cartridge 550. It is noted that the storage device 154 may be a storage device other than a magnetic tape drive or a sequential access medium device, in accordance with further example implementations.

[0027] In general, the storage device 154 is a physical machine that includes actual hardware and actual machine executable instructions, or "software." For example, the hardware may include a controller 520 that, in general, controls the overall operations of the device 154. The controller 520 may include one or multiple processors 522 (one or multiple central processing units (CPUs), microcontrollers,

- 9 -

processing cores, and so forth), as well as a memory 524 (a non-transitory memory, such as semiconductor storage, optical storage, and so forth) that may store data, program instructions and so forth for processing by the processor(s) 522. Through the execution of the machine executable instructions, the controller 520 forms an instance of the connectivity analysis and reporting engine 160, in accordance with example implementations.

[0028] For the example implementation of Fig. 5, the storage device 154 includes a drive interface 540 for purposes of writing data to and reading data from the physical cartridge 550. In this regard, the drive interface 540 may include such features as motors coupled to reels of the physical tape cartridge 550, read elements, write elements, servo elements and various other components, such as sense amplifiers, positioners, pulse detectors, error correction code (ECC) engines, and so forth, as can be appreciated by the skilled artisan.

[0029] Among its other features, the storage device 154 may include a read data path 530, a write data path 532, a drive motor interface 534, and one or multiple interfaces 155 (for redundancy purposes) that, as described above, may be coupled to the switch fabric 124 to receive as well as initiate connections for the storage device 154..

[0030] Among the advantages of the systems and techniques that are disclosed herein, storage devices may be used to automatically initiate and perform connectivity analyses for purposes of detecting and reporting connectivity issues more rapidly, thereby preventing lengthy and costly service events and customer escalations. Other and different advantages may be achieved using the systems and techniques that are disclosed herein in accordance with further example implementations.

[0031] While the present invention has been described with respect to a limited number of embodiments, those skilled in the art, having the benefit of this disclosure, will appreciate numerous modifications and variations therefrom. It is intended that the appended claims cover all such modifications and variations as fall within the true spirit and scope of this present invention.

- 10 -

What is claimed is:

1 1. A method comprising:
2 in a storage device coupled to a storage network, determining connectivity
3 information for the storage device, wherein determining the connectivity information
4 comprises determining a first set of at least one network device that should be
5 coupled to the storage device and determining a second set of at least one network
6 device that is coupled to the storage device; and
7 selectively initiating corrective action based at least in part on a comparison of
8 the first set with the second set.

1 2. The method of claim 1, wherein determining which network device
2 should be coupled to the storage device comprises using the storage device to
3 initiate a connection with the storage network.

1 3. The method of claim 1, wherein determining which network device
2 should be coupled to the storage device comprises retrieving connectivity
3 information data from the storage network.

1 4. The method of claim 3, wherein determining which network device
2 should be coupled to the storage device further comprises parsing the connectivity
3 information data and building a table to identify the at least one device of the first list.

1 5. The method of claim 1, wherein determining the second list comprises
2 retrieving connectivity information from an internal memory of the storage device.

1 6. The method of claim 1, wherein selectively initiating the corrective
2 action comprises communicating an alert message to a system monitoring station.

- 11 -

1 7. The method of claim 1, further comprising:
2 attempting to communicate over multiple connection paths to at least one
3 device that is not coupled to the storage device but should be coupled to the storage
4 device and generating a report identifying at least one of the multiple connection
5 paths that failed.

1 8. The method of claim 7, wherein generating the report comprises
2 analyzing hops associated with the multiple connection paths.

1 9. An article comprising a non-transitory computer readable storage
2 medium to store instructions that when executed by a computer cause the computer
3 to:
4 initiate a connection from a storage device to a storage network;
5 access at least one management function of the storage network to retrieve
6 first connectivity information data for the storage device, the first connectivity
7 information indicating at least one device that should be coupled to the storage
8 device;
9 retrieve second connectivity information data from an internal memory of the
10 storage device, the second connectivity information indicating at least one device
11 that is coupled to the storage device;
12 identify at least one device that should be coupled to the storage device but is
13 not coupled to the storage device based on the first connectivity information and the
14 second connectivity information; and
15 selectively communicate an alert to a management station for the storage
16 network based at least in part on the identification.

1 10. The article of claim 9, the storage medium storing instructions that
2 when executed by the computer cause the computer to communicate with a service
3 provided by a switch of the storage network.

- 12 -

1 11. The article of claim 9, the storage medium storing instructions that
2 when executed by the computer cause the computer to attempt to communicate over
3 multiple connection paths to at least one device that is not coupled to the storage
4 device but should be coupled to the storage device and generating a report
5 identifying at least one of the multiple connection paths that failed.

1 12. A storage device comprising:
2 storage media interface;
3 a network interface; and
4 a controller to:
5 initiate a connection with a storage network using the network interface
6 to determine connectivity information for the storage device, wherein determining the
7 connectivity information comprises determining a first set of at least one network
8 device that should be coupled to the storage device and determine a second set of
9 at least one network device that is coupled to the storage device; and selectively
10 initiate corrective action based at least in part on a comparison of the first set with
11 the second set.

1 13. The apparatus of claim 12, wherein the network interface:
2 accepts network connections for the storage device; and
3 initiates a network connection for the storage device to determine the first set
4 of at least one network device that should be coupled to the storage device.

1 14. The storage device of claim 12, wherein the storage network comprises
2 a storage area network (SAN).

1 15. The storage device claim 12, further comprising a memory, wherein the
2 controller determines which devices should be coupled to the storage device based
3 at least in part on data retrieved from the memory.

100

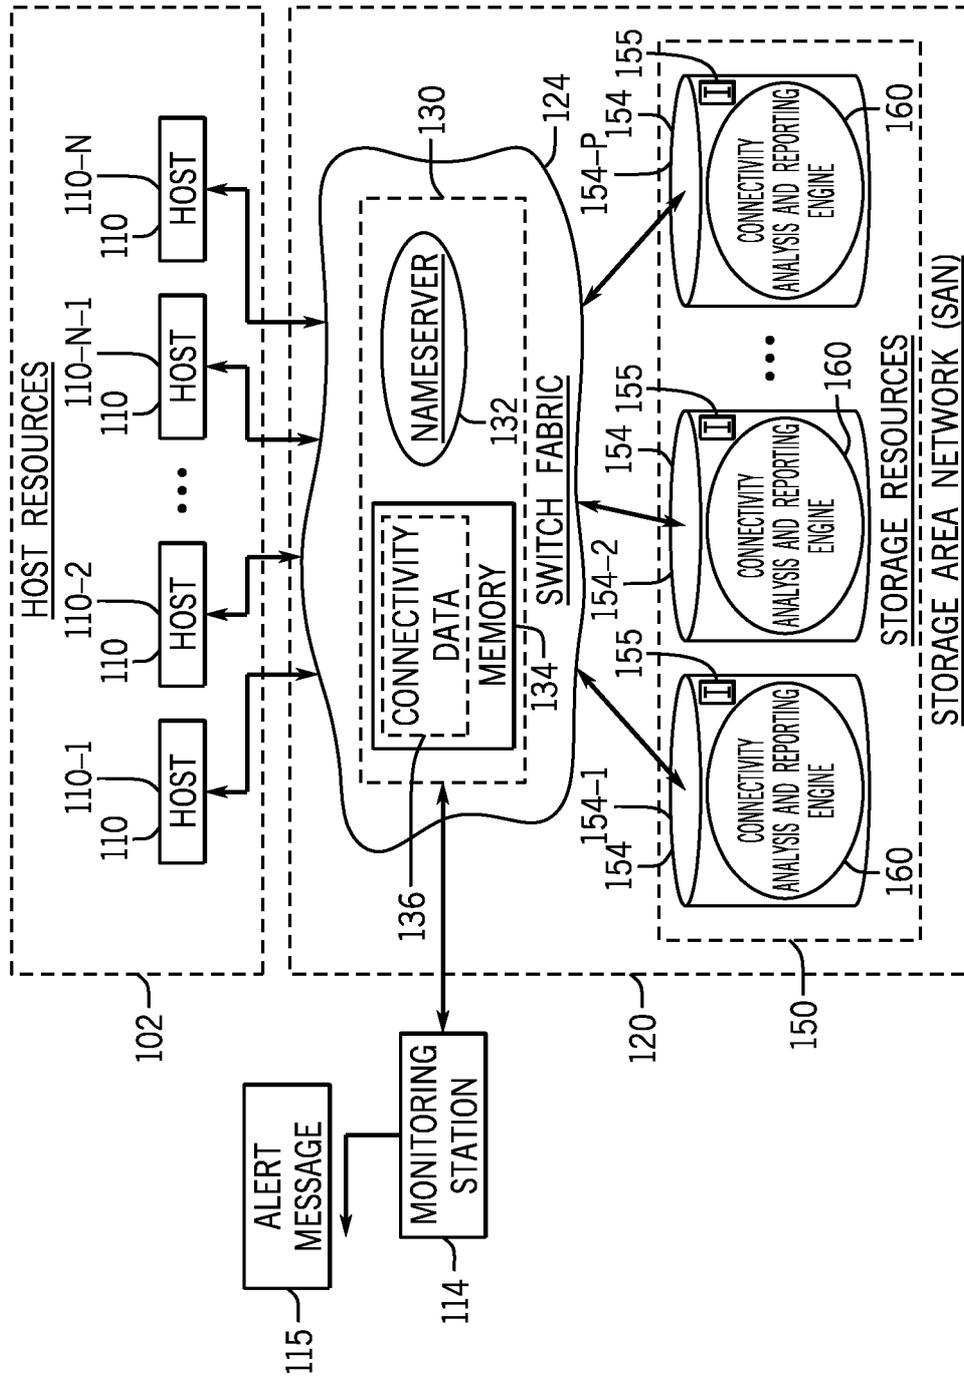


FIG. 1

2 / 4

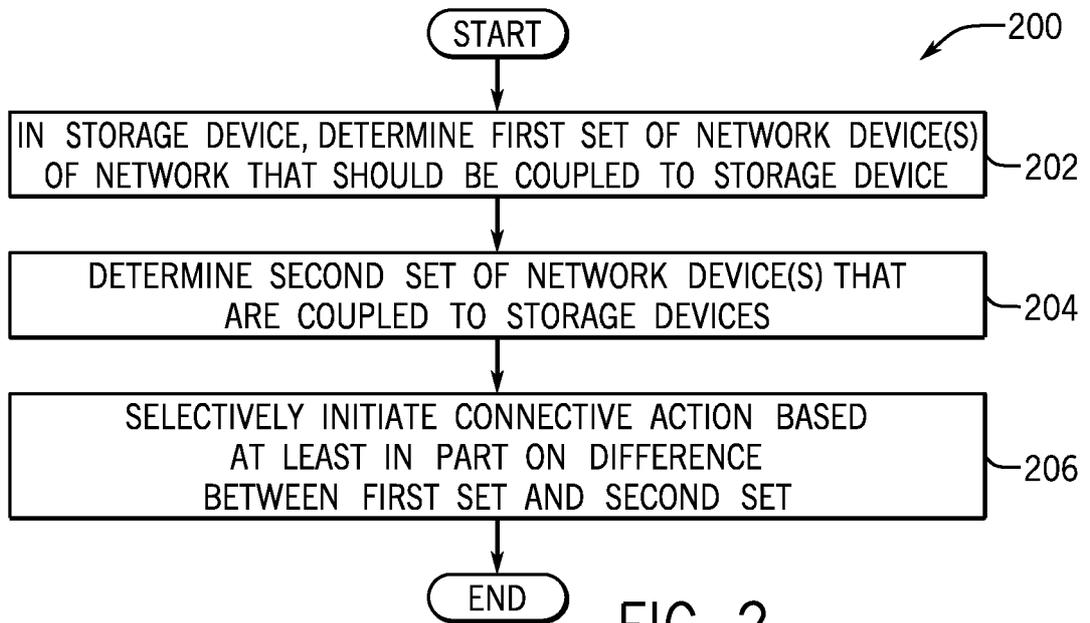


FIG. 2

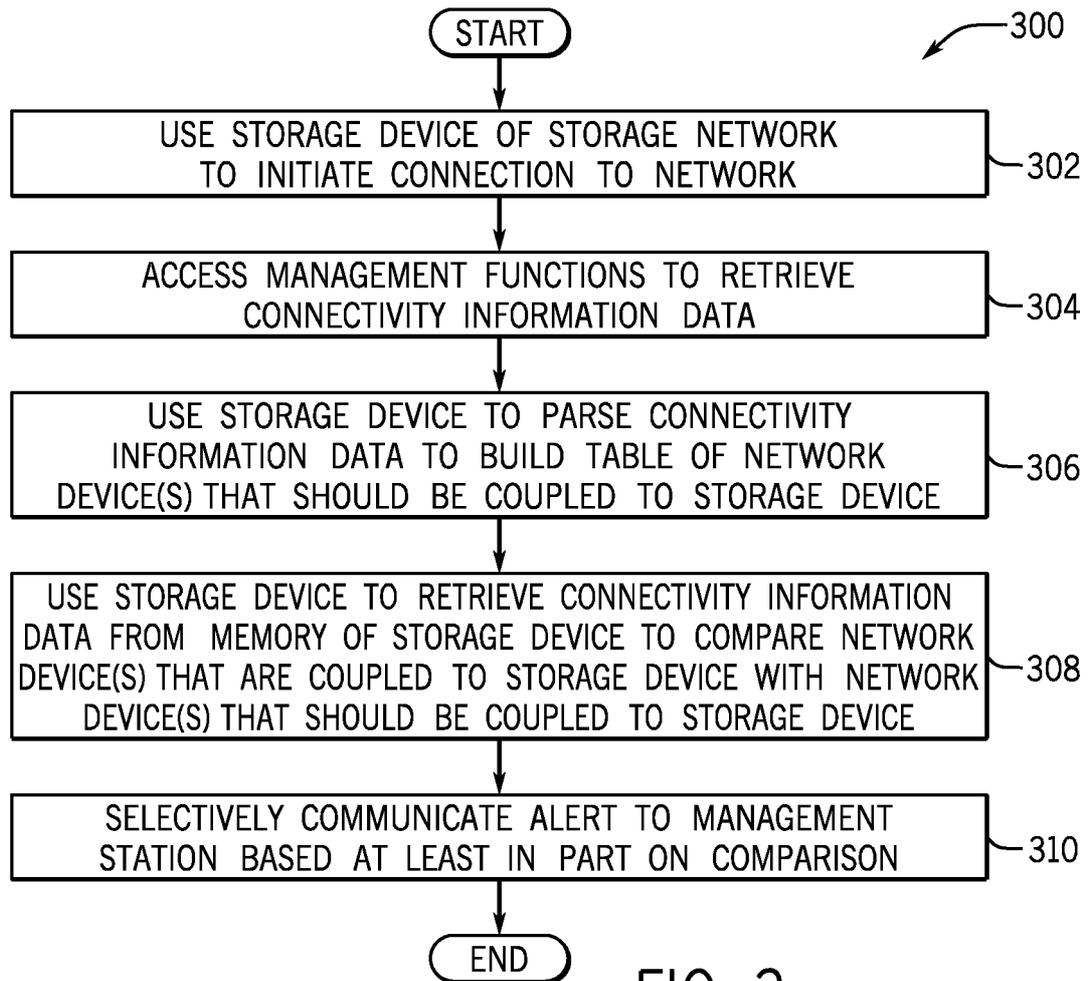


FIG. 3

3 / 4

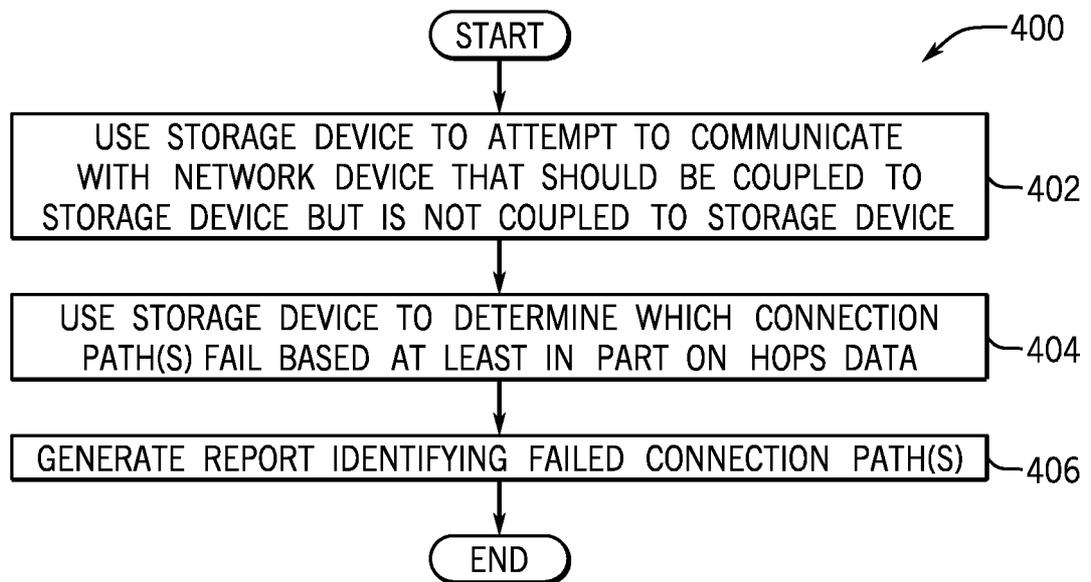


FIG. 4

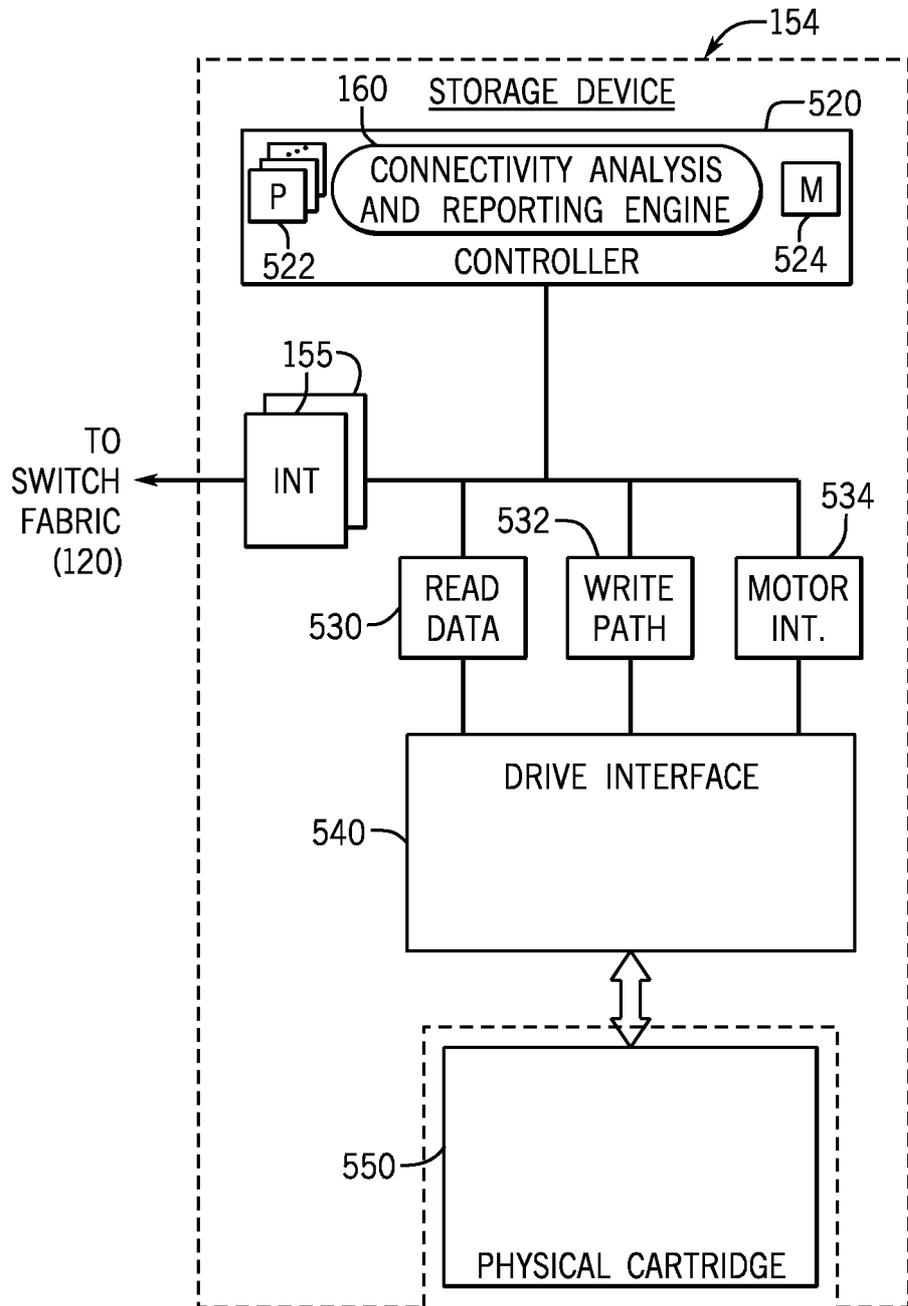


FIG. 5

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2014/035134**A. CLASSIFICATION OF SUBJECT MATTER****G06F 13/14(2006.01)i**

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHEDMinimum documentation searched (classification system followed by classification symbols)
G06F 13/14; G06F 15/173; G06F 17/30; G06F 12/00; G06F 15/167; G06F 15/177; G06F 13/00Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean utility models and applications for utility models
Japanese utility models and applications for utility modelsElectronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & Keywords: storage area network, storage device, network device, connectivity analysis and reporting engine, presence list, and similar terms.**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2008-0109584 AI (GEORGE ALEXANDER KALWITZ) 08 May 2008 See paragraphs [0019], [0023], [0027H0031], [0035], and [0040]; claims 1 and 11-12; and figures 1-2.	1,3-4,6
A		2,5,7-15
A	US 2013-0151646 AI (SRIRAM CHIDAMBARAM et al.) 13 June 2013 See paragraphs [0035], [0240], and [0385].	1-15
A	US 2010-0293316 AI (VIVEK MEHROTRA et al.) 18 November 2010 See paragraphs [0004], [0023]-[0025], and [0035]-[0036].	1-15
A	US 2006-0047850 AI (HARINDER PAL SINGH BHASIN et al.) 02 March 2006 See paragraphs [0017H0019], [0192], [0202], [0209], and [0219].	1-15
A	US 2006-0080416 AI (SHREYAS P. GANDHI) 13 April 2006 See paragraphs [0011], [0014], [0071], and [0213].	1-15

 Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

30 December 2014 (30.12.2014)

Date of mailing of the international search report

30 December 2014 (30.12.2014)

Name and mailing address of the ISA/KR


International Application Division
Korean Intellectual Property Office
 189 Cheongsa-ro, Seo-gu, Daejeon Metropolitan City, 302-701,
 Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

NHO, Ji Myong

Telephone No. +82-42-481-8528



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2014/035134

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2008-0109584 AI	08/05/2008	US 7793145 B2	07/09/2010
US 2013-0151646 AI	13/06/2013	US 7843906 B1	30/11/2010
		US 8218538 B1	10/07/2012
		US 8743872 B2	03/06/2014
US 2010-0293316 AI	18/11/2010	US 8054763 B2	08/11/2011
US 2006-0047850 AI	02/03/2006	wo 2006-026708 A2	09/03/2006
		wo 2006-026708 A3	17/01/2008
		wo 2006-026708 A8	13/03/2008
US 2006-0080416 AI	13/04/2006	US 7516214 B2	07/04/2009
		wo 2006-026677 A2	09/03/2006
		wo 2006-026677 A3	26/02/2009