




- (51) **International Patent Classification:**
G06F 11/14 (2006.01) *H04L 29/06* (2006.01)
G06F 15/16 (2006.01)
- (21) **International Application Number:**
PCT/US2015/028540
- (22) **International Filing Date:**
30 April 2015 (30.04.2015)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (71) **Applicant:** HEWLETT PACKARD ENTERPRISE DEVELOPMENT LP [US/US]; 11445 Compaq Center Drive West, Houston, TX 77070 (US).
- (72) **Inventors:** NADARAJAH, Navaruparajah; 995 Patrick Circle, Folsom, California 95630 (US). PUTTAGUNTA, Krishna; 8000 Foothills Blvd., Roseville, California 95747 (US). MOHAN, Rupin T.; 153 Taylor Street, Littleton, Massachusetts 01460-1407 (US). VOIGT, Douglas L.; 11311 Chinden Blvd., Boise, Idaho 83714-0021 (US).
- (74) **Agents:** ORTEGA, Arthur S. et al.; Hewlett Packard Enterprise, 3404 E. Harmony Road, Mail Stop 79, Fort Collins, CO 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, 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, ST, 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).

[Continued on next page]

(54) **Title:** DATA REPLICATION100
FIG. 1

(57) **Abstract:** In one example, disclosed is a storage system with multicast-based data replication. The storage system includes a server executing a host application and a primary storage array communicatively coupled to the server. The system also includes a plurality of secondary storage arrays configured as replication targets for the primary storage array and coupled to the primary storage array through a network. Upon receipt of a data storage request from the host, the primary storage array sends a replication data packet to the secondary storage arrays via a multicast operation.

**Declarations under Rule 4.17:**

- *as to the identity of the inventor (Rule 4.17(i))*
- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

Published:

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

DATA REPLICATION

BACKGROUND

[0001] Data replication is an important feature in modern data centers and increases the reliability, fault tolerance, and availability of stored data. Today's enterprise networks rely on data replication for a variety of purposes, including data backup, disaster recovery, data mining, archive functions, and others. Data movement between virtual hosts has led to increased demand for replication solutions within and between data centers.

DESCRIPTION OF THE DRAWINGS

[0002] Certain exemplary embodiments are described in the following detailed description and in reference to the drawings, in which:

[0003] Fig. 1 is a block diagram of an example of a storage network;

[0004] Fig. 2 is a block diagram of an example of a replication data packet;

[0005] Fig. 3 is a block diagram of an example multicast based replication configuration; and

[0006] Fig. 4 is a process flow diagram summarizing a method of replicating storage operations;

[0007] Fig. 5 is a block diagram showing a non-transitory, computer-readable medium that stores code configured to replicate storage operations.

DETAILED DESCRIPTION

[0008] The present disclosure describes techniques for replicating storage data. In a typical data center, host systems are able to store data to one or more networked storage arrays. In a data replication system, the data received by the primary array is stored to the primary storage array and also replicated to one or more secondary storage arrays.

[0009] The primary storage array connects to secondary storage arrays through dedicated array ports. For every secondary storage array in the replication setup, the primary storage array sends the same Input/Output (I/O) and manages all the replication protocol used to manage both Synchronous and

Asynchronous requirements. Therefore, the array based replication solutions consume a great deal of network resources and are not scalable.

[0010] The replication techniques disclosed herein leverage the multicast ability of the Ethernet network to duplicate a single input frame on a port to one or more output ports using the network switch's ability to duplicate frames. The storage array offloads the duplication of data to the network thereby leaving the storage array resources available to better serve the host applications executing on the servers. Multicast in the network is less costly than the storage array itself duplicating the I/O and sending multiple copies of the I/O. This results in a replication technique that is efficient and scalable. Additionally, the data replication is handled by the primary storage array and is transparent to the host applications.

[0011] Fig. 1 is a block diagram of an example of a storage network. The storage network is generally referred to by the reference number 100. As shown in Fig. 1, the storage network 100 may include data centers 102, which may be geographically dispersed. Each data center 102 may include a number of servers 104 operatively coupled by a communications network 106, for example, a wide area network (WAN), local area network (LAN), virtual private network (VPN), the Internet, and the like. The communications network 106 may be a TCP/IP protocol network or any other appropriate protocol. Any number of clients 108 may access the servers 104 through the communications network 106. Each server 104 may host one or more virtual machines or client applications. The applications running on the server 104 are referred to herein as host applications, and can include operating systems, file systems, database applications, and others.

[0012] Each data center 102 may also include one or more storage arrays 110, each of which can include an array of storage devices, such as an array of physical storage disks, an array of solid state drives, and others. The servers 104 may access the storage arrays 110 through a front-end network, which includes a number of switches 112 coupled by data links. Communication between the servers 104 and the storage arrays 110 can be accomplished using any suitable communication protocol including, Ethernet, Fibre Channel

(FC), Fibre Channel over Ethernet (FCoE), Internet Small computer System Interface (iSCSI), and others. The front-end network is the primary storage array access mechanism for host servers 104. The front-end network may be a block based (Storage Area Network) SAN or a file system based Network Attached Storage (NAS).

[0013] In some examples, the storage arrays 110 are also coupled together through a backend network, which can be used for data replication. The backend network can be an Ethernet network and includes the backend switches 114. Data stored to the data storage systems 110 may be replicated between the storage arrays 110 through the backend network using synchronous or asynchronous replication. The two data centers 102 may be communicatively coupled through a network 116, which uses Internet Protocol (IP) over Ethernet.

[0014] Each server 104 can be configured to store data to one or more storage arrays 110 within the corresponding data center 102. The storage array targeted by the server 104 for storing data in the first instance is referred to as the primary storage array for that server 104. Different servers 104 may be configured to target a different storage 110 array as its primary storage array. Each storage array 110 can also serve as a replication target for the other storage arrays 110. Replication operations are handled by a replication agent 118. The replication agent 118 can include a master component and a listener component as described further in relation to Fig. 3. The replication agent 118 may be implemented as hardware or as a combination as hardware and software. For example, the replication agent 118 may be implemented as a general purpose processor executing computer code, an Application Specific Integrated Circuit (ASIC), and others.

[0015] When a primary storage array receives a data storage request from a server 104, the I/O associated with the data storage request is replicated to one or more secondary storage arrays in accordance with the replication configuration that has been set up by a system administrator. Regardless of the number of secondary storage arrays, the primary storage array only sends the I/O once as a network-assisted multicast operation. In addition, a TCP/IP based

control path is also used for ensuring reliable delivery of data. The replication may be file-level replication or block-level replication depending on the design of a particular implementation.

[0016] Data replication in the storage system 100 is handled by the storage arrays 110 and is transparent to the applications executing on the servers 104. One or more of the storage arrays may be configured to replicate I/O received from a server 104 to one or more associated storage arrays in accordance with a configuration set up by a system administrator. Although only three storage arrays 110 are shown, the system 100 can include several additional storage arrays. For one or more of the storage arrays 110, the system administrator can create a replication group, which identifies specific storage arrays 110 as secondary storage arrays that are to receive replicated I/O.

[0017] In some examples, the Internet Group Management Protocol (IGMP) is used to configure the replication groups. The switches 114 are equipped with a mechanism to process IGMP joins and register the endpoint as being a member of the replication. IGMP snooping controls the flow of multicast updates only to ports that have joined the group. This ensures that other ports that are not in the replication path will not have packets forwarded to them by the switch 114. Configuring replication groups via IGMP limits the multicast of I/O replication data to the members of the replication group, thereby reducing the possibility of packet flooding. Furthermore, any number of secondary servers can be included within the replication group, which makes the replication techniques described herein fully scalable.

[0018] In some examples, the replication groups are configured using virtual local area network (VLAN) segregation. VLAN segregation can be used to create isolated broadcast domains within the backend network. The replication group can be configured as a segregated VLAN, such that multicast operations would forward I/O packets to ports within the configured VLAN. In some examples, IGMP replication groups may be configured within a particular VLAN.

[0019] The configuration of the storage network 100 is one example of a network that can be implemented in accordance with the present techniques. Those of ordinary skill in the art would readily be able to modify the described

storage network 100 based on design considerations for a particular system. For example, a storage network 100 in accordance with the present techniques may include any suitable number of data centers 102 and each data center 102 may include any suitable number of physical servers 104 and any suitable number of data storage systems 110.

[0020] Fig. 2 is a block diagram of an example of a replication data packet. The replication data packet 200 is created by the primary storage array for replicating data and is transmitted to the secondary storage arrays within a replication group. The replication data packet 200 includes data 202, which is the data to be replicated. The data 202 may correspond with a particular file or storage block stored to the storage array. The replication data packet 200 also includes a location identifier 204 that identifies the storage location of the data. The location identifier may be file for file-level replication or a Logical Unit Number (LUN) for block-level replication.

[0021] The replication data packet 200 also includes a Primary Array identifier 206 that identifies the source of the replication data packet 200. In some examples, the replication data packet 200 also includes Cyclical Redundancy Check (CRC) value 208 that can be used by the secondary storage array for integrity verification of the packet 200.

[0022] The replication data packet 200 also includes a tag 210. The tag 210 is used by the secondary storage array to enable the secondary storage array to determine if a packet has been lost. The tag 210 is incremented for each successive replication data packet 200 sent by the primary storage array. The secondary storage array can use the tag 210 to re-order packets and to determine if one or more transmitted packets failed to be received. If a packet fails to be received after a time threshold, the secondary storage array can request that the packet be resent using the control path established previously. The request would identify the lost packet by its tag 210 and the primary storage array would resend the packet to the requested secondary storage array.

[0023] Fig. 3 is a block diagram of an example multicast based replication configuration. The example of Fig. 3 shows a server 104 coupled to a primary storage array 300 through one of the switches 112 of the front-end network.

The primary storage array 300 can be any of the storage arrays 110 shown in Fig. 1. The primary storage array 300 is associated with a replication group that includes three secondary storage arrays. The secondary storage arrays 302 may be located in the same data center or geographically dispersed. Each of the secondary storage arrays 302 may be one of the storage arrays 110 shown in Fig. 1.

[0024] The primary storage array 300 is coupled to the secondary storage arrays 302 through the switches 304 and 306. In some examples, the switches 304 and 306 are a part of a back-end network corresponding with the switches 114 of Fig. 1. In some examples, the switches 304 and 306 are a part of the front-end network corresponding with the switches 112 of Fig. 1. Accordingly, switch 112 and switches 304 and 306 may belong to the same switch fabric or separate switch fabrics.

[0025] The server 104 may send data to the primary storage array 300 to be stored. The data may be data produced by an application or virtual machine executing on the server 104. For example, the data may be user data files that a person might use in the course of business, performing a job function, or for personal use, such as business data and reports, Web pages, user files, image files, video files, audio files, software applications, or any other similar type of data that that a user may wish to save to storage.

[0026] Data sent for storage to the primary storage array 300 by the server 104 is automatically replicated by the primary storage array 300. The replication process is managed by the primary storage array 300 without involvement of the host applications executing on the server 104. The data replication process is transparent to the server 104. The replication processes is managed in the primary storage array 300 by a master replication agent 308, which communicates with corresponding listener replication agents 310 executing on each of the secondary storage arrays 302. When data is received at the primary storage array 300, the master replication agent sends the IO-data with a tag added via send Multicast socket (multicast replication group). Switches 304 and 306 multicast this data to all the listening secondary storage arrays in that multicast replication group.

[0027] During the configuration of the replication group, all end nodes in the replication network, including primary storage array port and secondary storage arrays ports register themselves into a multicast group by IGMP registration. The registration establishes a data path between the master replication agent 308 of the primary storage array 300 and the listener replication agents 310 of the secondary storage arrays 302. The configuration of the replication group can be controlled by an administrator through a management program. For example, the management program can be executing on a client device coupled to the network 106 (Fig. 1). The management program can connect all listener replication agents 310 to the master replication agent 308 by TCP/IP connection and establish a control path in accordance with the replication group specified by the administrator. Through the management program, the user can also select a replication mode (Synchronous or Asynchronous) between the primary storage array 300 and each secondary storage array 302. Each secondary storage array 302 can use a different replication mode.

[0028] In the example shown in Fig. 3, there are three secondary storage arrays 302. To replicate data to the three storage arrays 302, four members of the switching network are registered in the multicast group. The four members are labeled A through D. Member A is the source array on 304 and is the only multicast traffic generator. Members B, C, and D are the three receiving secondary arrays. The configuration of the replication group as shown in Fig. 3 is one example of a configuration that may be implemented in a storage network. For example, additional members of the switch 304 can be included in the multicast group, and fewer or additional members of the switch 306 can be included in the multicast group. Additionally, the coupling between the primary storage array 300 and the secondary storage arrays 302 can be achieved through a single switch or multiple switches. Furthermore, any suitable number of secondary storage arrays 302 may be included in the replication group, including one, two, three, or more.

[0029] When data is received by the primary storage array 300 from the server 104, the storage array's master replication agent 308 commits the data to local cache and/or stores the data to the disk or other storage media of the

primary storage array 300. The master replication agent 308 then creates a replication data packet such as the replication data packet 200 of Fig. 2. The replication data packet includes the data to be replicated and the tag 210. The master replication agent 308 sends this replication data packet over the multicast data path. Copies of the replication data packet are automatically created in the switches 304 and 305 as appropriate. For example, switch 306 would create three copies of the replication data packet, one for each output port included in the multicast group. Thus, the primary storage array 300 can issue the replication data packet once as a multicast operation and allow the network components to create the relevant copies of the replication data packet.

[0030] Once the data is sent by multicast, all listening replication agents 310 receive the replication data packet. Upon receipt, the listening replication agent 310 commits the data to local cache and/or stores the data to a non-volatile storage media, such as a hard disk. After storing the data to the storage media, the listening replication agent 310 send a good status via the control path to the master replication agent 308.

[0031] The master replication agent manages the replication process differently depending on the replication mode. In synchronous replication, the master replication agent 308 waits for a good status to be received for all of the synchronous replication data packets, after which the master replication agent 308 sends a good status to the host running on the server 104. The good status sent to the host confirms successful processing of the corresponding storage I/O received from the host. In asynchronous replication, the master replication agent 308 sends the good status for the host I/O without waiting for the results of the replication operations.

[0032] If a listener replication agent doesn't receive one or more replication data packets, the listener replication agent sends the relevant details to the master replication agent 308 through the control path and asks the master replication agent 308 to resend them. The master replication agent 308 will resend the requested replication data packets to that particular listener replication agent 310 by a unicast transmission, which will go to only that particular listener replication agent 310.

[0033] Fig. 4 is a process flow diagram summarizing a method of replicating storage operations. The method may be referred to by the reference number 400, and is performed by a controller residing on a storage device, such as one of the storage arrays 110 shown in Fig. 1. For example, the storage device may include an ASIC or other processor configured to perform the method 400. The data replication process described herein is transparent to the host applications that generate the data.

[0034] At block 402, a storage instruction is received from a host application executing on a server. The storage instruction includes data to be stored to the storage media of the storage device. For example, the storage media may be any type of a non-volatile memory, including a hard disk or a solid state memory such as flash memory. The data may be a file or a block of data to be stored to one or more specific LUNs.

[0035] At block 404, the data is stored to the storage media. Although not shown in block 404, the data may be first stored in a volatile cache memory and stored to the storage media at a later time.

[0036] At block 406, a replication data packet is generated. In some examples, the replication data packet includes the components shown in Fig. 2. Each time a new replication data packet is generated, the tag included in the packet is incremented so that no two replication data packets include the same tag.

[0037] At block 408, the replication data packet is transmitted to a plurality of secondary storage arrays via network assisted multicast. The replication data packet is transmitted only once by the storage array. The network switches generate copies of the replication data packet and route the copies to ports that are registered in an Internet Group Management Protocol (IGMP) group.

[0038] At block 410, the storage array receives a request to resend a dropped or corrupt packet. Block 410 will be performed, for example, if the packet's cyclical redundancy check fails or if the packet does not arrive at one of the intended replication targets, as indicated by the tags received at a particular secondary server.

[0039] At block 412, replication data packet is resent to the secondary storage array that requested the resending. The replication data packet is send by a unicast operation that targets the specific secondary storage array that issued the request.

[0040] It is to be understood that the process flow diagram of Fig. 4 is not intended to indicate that the method is to include all of the blocks shown in Fig. 4 in every case. Further, any number of additional blocks can be included within the method, depending on the details of the specific implementation. In addition, it is to be understood that the process flow diagram of Fig. 4 is not intended to indicate that the method is only to proceed in the order indicated by the blocks shown in Fig. 4 in every case.

[0041] Fig. 5 is a block diagram showing a non-transitory, computer-readable medium that stores code configured to replicate storage operations. The computer-readable medium is referred to by the reference number 500. The computer-readable medium 500 can include RAM, a hard disk drive, an array of hard disk drives, an optical drive, an array of optical drives, a non-volatile memory, a universal serial bus (USB) drive, a digital versatile disk (DVD), a compact disk (CD), and the like. In some embodiments, the computer-readable medium 500 stores code to be executed by a processor residing on a storage array. The computer-readable medium 500 may be accessed by a processor 502 over a communication path 504.

[0042] As shown in Fig. 5, the various components discussed herein can be stored on the computer-readable medium 500. A region 506 can include an I/O processing engine that processes I/O requests received from a host applications running on a server. For example, processing I/O requests can include storing data to a storage drive in accordance with a storage instruction or retrieving data from a storage drive and sending it to a host application that requested it. A region 508 can include a master replication agent that manages the replication of storage instructions as described above. A region 510 can include a listener replication agent that receives replication data packets and process them as described above. The master replication agent and the listener

replication agent may be components of the replication agent 118 shown in Fig. 1.

[0043] While the present techniques may be susceptible to various modifications and alternative forms, the techniques discussed above have been shown only by way of example. It is to be understood that the technique is not intended to be limited to the particular examples disclosed herein. Indeed, the present techniques include all alternatives, modifications, and equivalents falling within the scope of the following claims.

CLAIMS

What is claimed is:

1. A system comprising:
a server executing a host application;
a primary storage array communicatively coupled to the server; and
a plurality of secondary storage arrays to be configured as replication targets for the primary storage array and coupled to the primary storage array through a network;
wherein after receipt of a data storage request from the host, the primary storage array is to send a replication data packet to the secondary storage arrays via a multicast operation.
2. The system of claim 1, the secondary storage arrays to be configured as a replication group by being registered as an Internet Group Management Protocol (IGMP) group.
3. The system of claim 1, the secondary storage arrays to be configured as a replication group by being configured as a Virtual Storage Area Network (VSAN).
4. The system of claim 1, comprising a front-end network to couple the primary storage array to the server, and a back-end network to couple the primary storage array to the secondary storage arrays.
5. The system of claim 1, comprising a network to couple the primary storage array to the server and to couple the primary storage array to the secondary storage arrays.
6. The system of claim 1, the secondary storage arrays configured to perform a combination of synchronous replication and asynchronous replication.

7. The system of claim 1, wherein the primary storage array is to increment a tag and including the tag in the replication data packet, the tag used by the secondary storage arrays for packet ordering and identification of dropped packets.

8. The system of claim 1 wherein, after receipt of a request from one of the plurality of secondary storage arrays to resend the replication data packet, the primary storage array is to resend the replication data packet to the one of the plurality of secondary storage arrays via a unicast operation.

9. A method of replicating storage instructions, the method comprising:
receiving a storage instruction from a host application executing on a server, the storage instruction comprising data to be stored;
storing the data to a storage media; and
replicating the data to a plurality of secondary storage arrays by including the data in a replication data packet and transmitting the replication data packet to the plurality of secondary storage arrays via network assisted multicast;
wherein replicating the data is transparent to the host application.

10. The method of claim 9, wherein transmitting the replication data packet comprises transmitting the replication data packet once, and wherein the switches in the network generate copies of the replication data packet and route the copies to ports that are registered in an Internet Group Management Protocol (IGMP) group.

11. The method of claim 9, comprising incrementing a tag and including the tag in the replication data packet.

12. The method of claim 9, comprising receiving a request from one of the plurality of secondary storage arrays to resend the replication data packet

and resending the replication data packet to the one of the plurality of secondary storage arrays via a unicast operation.

13. A non-transitory, computer-readable medium, comprising code configured to direct a processor to:


- receive a storage instruction from a host application executing on a server, the storage instruction comprising data to be stored;
- store the data to a storage media; and
- replicate the data to a plurality of secondary storage arrays by including the data in a replication data packet and transmitting the replication data packet to the plurality of secondary storage arrays via network assisted multicast;

wherein the data replication is transparent to the host application.

14. The non-transitory, computer-readable medium of claim 13, comprising code configured to direct the processor to increment a tag and include the tag in the replication data packet.

15. The non-transitory, computer-readable medium of claim 13, comprising code configured to direct the processor to receive a request from one of the plurality of secondary storage arrays to resend the replication data packet and resend the replication data packet to the one of the plurality of secondary storage arrays via a unicast operation.

1/5




100
FIG. 1




200
FIG. 2

3/5




300
FIG. 3

4/5



400
FIG. 4

5/5



500
FIG. 5

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2015/028540**A. CLASSIFICATION OF SUBJECT MATTER****G06F 11/14(2006.01)i, G06F 15/16(2006.01)i, H04L 29/06(2006.01)i**

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F 11/14; G06F 15/16; G06F 15/173; G06F 13/28; G06F 15/167; H04L 12/18; G06F 11/00; H04L 12/56; H04L 29/06

Documentation 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 models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

eKOMPASS(KIPO internal) & Keywords: data, replication, multicast, unicast, primary, second, command, storage array, network, switch, 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.
Y	US 2005-0256972 A1 (ROBERT COCHRAN et al.) 17 November 2005 See paragraphs [0018]-[0021], [0030], [0041], and [0054]-[0085]; claims 10, 32, and 36; and figures 1 and 5-11.	1-15
Y	US 7,631,052 B2 (GENE FEIN et al.) 08 December 2009 See column 4, line 1 - column 5, line 7; claim 1; and figure 5.	1-15
Y	US 2014-0204940 A1 (NEXENTA SYSTEMS, INC.) 24 July 2014 See paragraphs [0159], [0293], and [0332]-[0336]; and figure 8.	2,10
A	US 7,376,859 B2 (KENNETH WAYNE BOYD et al.) 20 May 2008 See column 2, line 59 - column 3, line 62; and figure 1.	1-15
A	US 2013-0329605 A1 (JUNIPER NETWORKS, INC.) 12 December 2013 See paragraphs [0038]-[0060] and figure 2.	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

25 February 2016 (25.02.2016)

Date of mailing of the international search report

25 February 2016 (25.02.2016)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

COMMISSIONER

Telephone No. +82-42-481-3578



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2015/028540

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2005-0256972 A1	17/11/2005	JP 2005-327283 A US 8185663 B2	24/11/2005 22/05/2012
US 7631052 B2	08/12/2009	AU 2009-225859 A1 AU 2009-225859 B2 CA 2718913 A1 CA 2718913 C CN 101978676 A EP 2255519 A2 EP 2255519 B1 JP 2011-519442 A JP 5211235 B2 KR 10-1160504 B1 KR 10-2010-0125431 A US 2009-0238167 A1 US 2009-0240782 A1 US 2013-262632 A1 US 8458285 B2 US 8909738 B2 WO 2009-117259 A2 WO 2009-117259 A3	24/09/2009 13/06/2013 24/09/2009 17/12/2013 16/02/2011 01/12/2010 16/07/2014 07/07/2011 12/06/2013 13/07/2012 30/11/2010 24/09/2009 24/09/2009 03/10/2013 04/06/2013 09/12/2014 24/09/2009 12/11/2009
US 2014-0204940 A1	24/07/2014	US 2014-0204941 A1 US 2014-0207849 A1 US 2014-0207850 A1 WO 2014-116875 A2 WO 2014-116875 A3	24/07/2014 24/07/2014 24/07/2014 31/07/2014 16/10/2014
US 7376859 B2	20/05/2008	US 2005-097391 A1	05/05/2005
US 2013-0329605 A1	12/12/2013	CN 103514245 A CN 104521196 A EP 2672668 A1 EP 2859694 A1 EP 2882150 A1 EP 2882151 A1 US 2013-0329548 A1 US 2013-0329584 A1 US 2013-0329725 A1 US 2013-0332399 A1 US 2013-0332577 A1 US 2013-0332601 A1 US 2013-0332602 A1 US 2014-0129700 A1 US 2015-0244617 A1 US 8750288 B2 US 8755377 B2 US 8953441 B2 US 8959185 B2	15/01/2014 15/04/2015 11/12/2013 15/04/2015 10/06/2015 10/06/2015 12/12/2013 12/12/2013 12/12/2013 12/12/2013 12/12/2013 12/12/2013 12/12/2013 08/05/2014 27/08/2015 10/06/2014 17/06/2014 10/02/2015 17/02/2015

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2015/028540Patent document
cited in search reportPublication
datePatent family
member(s)Publication
date

US 9064216 B2	23/06/2015
US 9094308 B2	28/07/2015
US 9100289 B2	04/08/2015
WO 2013-184846 A1	12/12/2013