



- (51) **International Patent Classification:**
H04L 29/08 (2006.01) *G06F 9/455* (2006.01)
- (21) **International Application Number:**
PCT/US2015/013813
- (22) **International Filing Date:**
30 January 2015 (30.01.2015)
- (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:** LUMAN, Ronald John; 2604 Roe Drive, Lewis Center, Ohio 43035 (US). BONDURANT, Matthew David; 2580 55th Street, Boulder, Colorado 80301 (US). TALGERY, Hrishikesh; 4245 Technology Drive, Fremont, California 94538 (US). HADA, Craig M.; 4245 Technology Drive, Fremont, California 94538 (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, LI, 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 4.17(i))

[Continued on next page]

- (54) **Title:** VIRTUAL MACHINES AND FILE SERVICES

(57) **Abstract:** In one example, a storage area network (SAN) with file services is disclosed. The method includes invoking, via a processor, a command to activate file services on at least one node of the SAN, where the SAN comprises a plurality of nodes. The method includes installing, via the processor, a virtual machine infrastructure on each of the plurality of nodes, where the virtual machine infrastructure includes a plurality of virtual machines. The method includes enabling, via the processor, the file services on each of the virtual machines of the plurality of nodes.




FIG. 1A



— *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))*

VIRTUAL MACHINES AND FILE SERVICES

BACKGROUND

[0001] When a physical computer and a host operating system are configured appropriately, a virtual machine can access resources located on a storage area network (SAN). A SAN is a set of interconnected devices and servers that are connected to a common communication and data transfer infrastructure such as Fibre Channel. The purpose of the SAN is to allow multiple servers access to a pool of block storage in which any server can potentially access any storage unit with the block storage.

BRIEF DESCRIPTION OF THE DRAWINGS

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

[0003] Fig. 1A is an illustration of an example of a computer system including a storage area network (SAN) with virtual machines located in each node of the SAN;

[0004] Fig. 1B is an illustration of an example of a computer system including a storage area network (SAN) with a pair of virtual machines located in a pair of nodes of the SAN;

[0005] Fig. 2A is a process flow diagram of an example of a method of enabling file services within a SAN;

[0006] Fig. 2B is a process flow diagram of an example of a method of generating and clustering virtual machines (VMs) on a SAN; and

[0007] Fig. 3 is a block diagram of a tangible, non-transitory, computer-readable medium that holds code to direct a processor to enable file services on a virtual machine of a storage area network (SAN).

DETAILED DESCRIPTION OF SPECIFIC EXAMPLES

[0008] Examples disclosed herein provide techniques for enabling file services using a virtual management system located in a storage area network (SAN). As the information technology (IT) field is faced with challenges associated with data

growth, the ability to merge block-level storage and file-level storage to provide a stream-lined management system becomes increasingly important.

[0009] Generally, a SAN provides block-level storage, e.g., blocks of data, that can be accessed by applications running on network servers connected to the SAN. However, a SAN does not provide file-level storage, e.g., storage of files/folders using a file directory protocol. Instead, a file system may be built as a separate entity, on top of a SAN, to perform file services. The combination of both the file system and the SAN, which includes the block storage, may be considered as a SAN file system or shared disk file system.

[0010] However, the use of separate, dedicated hardware for both block-level storage and the file services often involves duplicate server hardware and may slow system performance for larger applications. Other techniques to provide both block-level storage and file services may include converging the two storage techniques to create an integrated service system under a single server and operating system. However, such integrated solutions may lead to a reduction in security, system performance, and fault isolation.

[0011] As disclosed herein, an infrastructure that includes file services isolated from block-level storage services, where both services are included in an existing SAN appliance, may be used as a storage management solution. The SAN may be embedded with dormant file services software that may be enabled on nodes within the SAN. Once the file services are invoked for use, virtual machines may be generated and located within the nodes of the SAN using a template approach. Each virtual machine, as a file services server, may have the ability to substantially maintain the activities of the file services so as to isolate the file services from the block services of the SAN.

[0012] Fig. 1A is an illustration of an example of a computer system 100 including a storage area network (SAN) with virtual machines located in each node of the SAN. The computer system 100 may be used to provide hosted services, for example, cloud services for clients. In particular, the computer system 100 may be used by one or more host cloud services, such as Software as a Service (SaaS), Infrastructure as a Service (IaaS), and Platform as a Service (PaaS). The computer system 100 may include one or more servers, for example, application host servers 102, that communication with one or more clients 104 and with storage, for example,

a SAN 106. The SAN 106 is a network of multiple storages devices that connects the storage devices with the application host servers 102 using an interconnection technology, such as SCSI switches, or in the case of the present examples, both a SAN fabric and a file services network fabric 108. In operation, applications running on the application host servers 102 may access the storage devices of the SAN 106 to retrieve data. The network of storage device within the SAN 106 can include hard drives, tape libraries, and, more commonly, disk-based devices, for example, RAID hardware.

[0013] In the present examples, the SAN 106 provides both block-level storage and file services for the computer system 100. Specifically, the SAN 106 includes physical nodes 110 used for the storage of data. The number of nodes 110 may vary depending on the particular implementation. In the example depicted in Fig. 1, the SAN 106 includes six nodes 110-1 through 110-6 that are intercoupled by a network communications channel 112 to provide secure and efficient management of data among the nodes 110. In some examples, the network communication channel 112 may include point-to-point connections in a mesh network, among others. The nodes 110 are physical machines that include computing resources and storage resources. For example, the computing resources may provide management and end-to-end control of input/output from the application host servers 102 to storage. The storage resources of the nodes 110 may include a physical mass storage device, such as a magnetic tape drive, a hard drive, an optical drive, among others. In an example, the storage resources may be coupled to the nodes 110. For example, one or more mass storage devices may be coupled to a given node 110 using a Serial Attached Small Computer System Interface (SCSI), i.e., a "SAS."

[0014] The nodes 110 of the present examples may be configured to provide dual services. In some examples, the nodes 110 are configured to perform block-level storage services and at least one node 110 is configured to provide both file and block services within the SAN 106. The node 110 providing file services may be redundantly provisioned, e.g., with two or more nodes 110 providing identical or overlapping file services within the SAN 106. Thus, the host application servers 102 may communicate with the nodes 110 using the file services network fabric 108 to retrieve either block-level storage or file-level storage.

[0015] As previously discussed, the network communications channel 112 of the SAN 106 provides storage to be shared amongst the nodes 110. This shared storage, along with the writing of data to persistent storage in a consistent state, allows the nodes 110 to be grouped into failover arrays as part of a redundancy approach, where a given failover array includes two or more nodes 110, depending on the implementation. For example, the failover arrays may include nodes 110 grouped in pairs. Accordingly, if one of the nodes 110 within a failover pair array enters a fail state, another storage node 110 with the failover pair array may take over. The fail state of the example may include a system shutdown or unreliable system performance due to a power failure, hardware failure, or software failure, among others. A failed storage node may be serviced or replaced for purposes of providing reliable storage services.

[0016] As previously stated, in addition to block-level storage, the SAN 106 of the present example also provides file-level storage, e.g., the file services. In general, file services enable servers of the SAN 106 to index and share data files, e.g., word processing documents, spreadsheets. In this manner, the data files may be stored on file services servers within the SAN 106 and accessed by the host application servers 102 of the computer system 100. The file services servers of the SAN 106 may run basic services, such as data storage, folder sharing, and share permissions. In particular, the file services servers of the present example may permit the host application servers 102 to use a file protocol, such as a Network File System (NFS) or Service Message Block (SMB) protocol, to create, share, and to retrieve data files that are stored on the SAN 106.

[0017] To create the file services servers, the present examples may create virtual machines (VMs) 114, e.g., 114-1 to 114-6, within the nodes 110 of the SAN 106 using a virtualization infrastructure. As discussed herein, the VMs 114 may be automatically created using a configuration template. The newly created VMs 114 may claim file services resources that have been reserved by the SAN 106 and attach the resources to the VMs 114. Accordingly, each of the VMs 114 may be considered as a file services VM 114 that is executed on a respective node 110 within the SAN 106 to provide file-level storage. In this manner, the SAN 106 provides a physical platform for the files services VMs 114 to be created and to enable and provide file services for the storage of files and folders that can be

directly accessed and managed. Such file services may include file-sharing capabilities such as server message block (SMB) file sharing, or network file system (NFS) file sharing, among other type of file sharing techniques. The file sharing systems may run as a virtualization guest. In the present techniques, the file services may be substantially isolated from the block-level storage services since the files services VMs 114 are located in an isolated environment within the SAN 106.

[0018] As previously described, the network communications channel 112 may be used to provide communication between the nodes 110 using hardware of the SAN 106, such as interconnect hardware. In this case, communications between each of the file services VMs 114 may also be established through the network communications channel 112. In examples, the channel 112 may act as standard network interface for the file services VMs 114. The network communications channel 112 may be under VM guest control, with minimal diagnostic information provided by the SAN 106.

[0019] The SAN 106 may include a service processor 116, which may be an actual physical machine or a virtual machine. As shown in Fig. 1, the service processor 116 is in communication with the nodes 110 via a network fabric 118. The service processor 116 may perform such functions as reporting and supporting diagnostic and maintenance activities. Further, each file services VM 114 may be coupled to a management processor 120 via the network fabric 118 to perform file services management including generating commands such as creating, showing, setting, and removing file shares.

[0020] Fig. 1B is an illustration of an example of a computer system 122 including a storage area network (SAN) with a pair of virtual machines located in a pair of nodes of the SAN. Like numbers are as described with respect to Fig. 1A. As shown in Fig. 1B, the number of files services VMs 114 within the nodes 106 may be limited. In examples, the VMs 114 may be provided in pairs to automatically enable file services within the nodes 106. The pairing of files services VMs 114 provides a redundancy approach in the case of power failure or disruption. For example, a failed VM may be replaced by the second VM in the pair to continue the file services.

[0021] The examples described provide for the management of reserved resources and the creation and management of virtual machines that are capable of enabling and providing file services capabilities on a SAN that also includes block

storage. The virtual machines of the SAN may provide an isolated and separate environment in which file services are enabled for data storage and file sharing.

[0022] Although the computer systems 100 and 122 of Fig. 1A and Fig. 1B are depicted as self-contained entities within a box, the components of the systems 100 and 122 may be locally disposed at a given site or may be distributed at multiple locations, depending on the particular implementation. Further, it is to be understood that the illustrations of Fig. 1A and Fig. 1B are not intended to indicate that the computer systems 100 and 122 are to include all of the components shown in the figures in every example. Further, any number of additional components can be included within the computer systems 100 and 122 depending on the details of the specific implementations, including SAN discovery collectors, switches, servers, among others.

[0023] Fig. 2A is a process flow diagram of an example of a method of enabling file services within a SAN. The SAN may be configured to include server-computing capabilities on each of its nodes. For example, the SAN may include hardware for block-level storage services, hardware for providing a file sharing data path, and a physical infrastructure to implement interconnections between the nodes. Further, the software of the SAN may be capable of acting as a virtualization host and providing file services that are capable of running as a virtualization guest.

[0024] When a user of a computer system identifies a need to use file services, at block 202, a command may be invoked to activate dormant file services. Based on the command, at block 204, the computer system may automatically create and install a virtual machine (VM) infrastructure where a VM may be located on each of the nodes of the SAN. In some examples, the VMs may be created on less than all of the nodes. Instead, the VMs may be created on at least two of the nodes of the SAN so as to be created in pairs for a redundancy approach. In some examples, the VMs may take control of the file services and thus, may be considered as file services VMs. At block 206, an enable command starts each file services VM and enables a files services system on each of the file services VMs for use. In the present examples, the file services system may be substantially maintained within the file services VMs and isolated from block services. In some examples, after using file services, a user may decide to disable file services, at which point the file services VMs are destroyed.

[0025] In some cases, file services software of the SAN, may be dormant when delivered to a user. For example, dormant resources for file services may be identified in advance and reserved for later use. The use of such pre-allocated file resources may insure that file services may be enabled within the SAN on an as-needed basis. This may conserve energy when file services are not in use, as opposed to running file services on a continuous basis. Further, the pre-allocation of resources for file services ahead of their use may allow for large contiguous blocks of memory to be reserved. In some cases where the setting aside of resources is delayed until deployment, the resources may be fragmented so that smaller blocks of memory may exist.

[0026] The reserved file resources may include central processing unit (CPU) cores, memory, and network interfaces, among other resources. In some cases, the reservation for the resources may be removed and the file services resources may be claimed by the files services VMs for attachment prior to activating the files services VMs.

[0027] Based on the present examples, the SAN can provide both block-level storage and file-level storage using a virtualized infrastructure created directly within the nodes of the SAN. Accordingly, file access and sharing of data can be handled directly within the SAN since file services software may be enabled on virtual machines located within the SAN. Accordingly, the SAN of the present examples provide mass file storage, as well as the ability to connect to storage via a file level protocol.

[0028] It is to be understood that the process flow diagram of Fig. 2A is not intended to indicate that the method is to include all of the blocks shown in Fig. 2A, 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. 2A is not intended to indicate that the method is only to proceed in the order indicated by the blocks shown in Fig. 2A in every case.

[0029] Fig. 2B is a process flow diagram of an example of a method of generating and clustering virtual machines (VMs) on a SAN. At block 208, a number of VMs may be generated in the SAN using a configuration template. The configuration template is the configuration of a single virtual machine that may be

shared among other users to create new virtual machines (VMs) and to compile common configuration details and software tools to be delivered across each new VMs. In some examples, the new VMs may inherit the same contents and configuration from the configuration template. For example, the configuration template may contain basic configuration, such as the number of virtual CPUs, the size of memory, virtual disks, virtual network interfaces, and so on. Various techniques may use a configuration template including a Clone to Template method or Convert to Template method. The Clone to Template method duplicates a virtual machine and changes the machine to match a template format. The Convert to Template method does not change the duplicated virtual machine into a template format but leaves the original virtual machine intact for further use. In the present examples, each generated VM has the capacity to attach the file services resources reserved by the SAN, and thus, create a file services VM in order to store and share file-level data.

[0030] A network communications channel connects the nodes of the SAN in order to provide a pathway for the nodes to communicate with one another. The network communication channel may also establish communications between the file services VMs. In addition to the network communication channel, a clustering method may be used to establish communications and maintain common configurations across the file services VMs. In examples, the clustering method may cluster the groups of VMs together into groups. Accordingly, at block 210, the file services VMs may be grouped into cluster formations. In operation, after the file services VMs are installed within the nodes of the SAN, the file services VMs may identify a network interface using a media access control (MAC) address and may assign an internet protocol (IP) address to each file services VMs. The IP addresses may be used by the file services VM to locate one another and to form groups of file services VMs into cluster formations.

[0031] It is to be understood that the process flow diagram of Fig. 2B is not intended to indicate that the method is to include all of the blocks shown in Fig. 2B 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. 2B is not intended to indicate

that the method is only to proceed in the order indicated by the blocks shown in Fig. 2B in every case.

[0032] Fig. 3 is a block diagram of a tangible, non-transitory, computer-readable medium that holds code to direct a processor to enable file services on a virtual machine of a storage area network (SAN). The computer-readable medium 300 can be accessed by a processor 302 over a system bus 304. In some examples, the code may direct the processor 302 to perform the steps of the current method as described with respect to FIGS. 2A and 2B. In examples, the SAN may implement a virtualized management system to provide block-level storage services and file-level storage services, where the file-level storage services operate in an isolated environment separate from the block-level services.

[0033] The computer-readable medium 300 can include an invoke module 306 to activate file services on at least one node of the SAN. The computer-readable medium can include an install module 308 to install a virtual machine infrastructure on each of the nodes of the SAN. In some examples, the virtual machine infrastructure includes a plurality of virtual machines. Accordingly, the computer-readable medium 300 may include a generate virtual machines module 310. For example, after a user has invoked a command to use file services and to install a virtual machine infrastructure on the nodes of the SAN, the SAN may generate multiple virtual machines using a template approach using the virtual machine infrastructure. In some cases, a virtual machine may be installed on each of the nodes of the SAN.

[0034] The computer-readable medium 300 may include an enable module 312 within the SAN to enable file services, on an as needed basis, on each of the virtual machines. In some examples, the enabling of the files services on the virtual machines may create file services virtual machines (VMs), which may provide an isolated environment to process a file services request while providing separation from block services. A virtualized SAN may provide file services that run independently in terms of management, availability, and migration of resources between file services virtual machines. Overall, by implementing a storage array with both block storage and files services, a computer system may stream-line its storage allocation and improve data protection and storage management.

[0035] The computer-readable medium 300 may include an activate module 314 where the files services may be activated to run on each of the virtual machines. In some cases, file services resources may be reserved in advance on a node of the SAN. For example, the SAN may reserve resources in anticipation of performing a service, such as file-level storage services, file sharing services, and the like. In the present examples, the operating system (OS) software of the SAN may reserve file services resources to several nodes of the SAN in anticipation of enabling a file services system within the SAN. In some cases, the OS software may be modified to specifically support the reservation. The reserved file services resources may be claimed by the virtual machines via a claim resources module 316 at the time of activating the file services.

[0036] The computer-readable medium 300 may include a cluster virtual machines module 318 to cluster the virtual machines, for example, into pairs or groups to provide communication among the virtual machines and to provide a redundancy approach. As a result, a two-node cluster, a three-node cluster, a four-node cluster, or a five-node cluster and so on, may be established where the number of virtual machines is based on the specific implementation.

[0037] The block diagram of FIG. 3 is not intended to indicate that the computer-readable medium 300 is to include all of the components or modules shown in FIG. 3. Further, any number of additional components may be included within the computer-readable medium 300, depending on the details of the specific implementations as described herein.

[0038] While the present techniques may be susceptible to various modifications and alternative forms, the examples discussed above have been shown only by way of example. However, it should again be understood that the techniques are 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 appended claims.

CLAIMS

What is claimed is:

1. A method, comprising
invoking, via a processor, a command to activate file services on at least one node of a storage area network (SAN), wherein the SAN comprises a plurality of nodes;
installing, via the processor, a virtual machine infrastructure on each of the plurality of nodes, wherein the virtual machine infrastructure comprises a plurality of virtual machines; and
enabling, via the processor, the file services on each of the virtual machines.
2. A method of claim 1, comprising,
activating, via the processor, the file services to run on each of the virtual machines, wherein file services resources are reserved in advance on at least one node of the SAN; and
claiming, via the processor, the file services resources for attachment to the plurality of virtual machines at the time of activating the file services.
3. The method of claim 1, comprising,
generating, via the processor, the plurality of virtual machines via a template, wherein at least two or more virtual machines include file service capabilities; and
clustering, via the processor, the plurality of virtual machines to enable communications between the plurality of virtual machines.
4. The method of claim 1, comprising,
establishing, via the processor, a network communications channel between each node of the SAN, wherein the network communications channel is exposed to the plurality of virtual machines as a network interface.

5. The method of claim 1, comprising,
running, via the processor, block services on the SAN, wherein the block services are isolated from the file services via a virtual machine infrastructure.
6. The method of claim 1, comprising,
substantially maintaining, via the processor, the file services within each of the virtual machines.
7. The method of claim 1, comprising grouping, via the processor, the plurality of virtual machines in pairs to provide a redundancy approach, wherein the file services are attached to a pair of virtual machines.
8. A method of claim 4, where the network communications channel is utilized by the processor to detect and notify each node of the SAN in the event of failure.
9. The method of claim 3, wherein the clustering comprises automatically identifying and assigning an Internet Protocol (IP) address to each of the virtual machines.
10. A system, comprising
a storage array comprising reserved file service resources;
a plurality of nodes located within the storage array; and
a plurality of virtual machines located within the storage array, wherein
at least two of the virtual machines comprise file service capabilities.
11. The system of claim 10, wherein the file service capabilities are isolated from block services of the storage array.
12. The system of claim 10, wherein the plurality of virtual machines are created based on a configuration template.

13. The system of claim 10, comprising an interconnect channel to provide network communications between the plurality of virtual machines.

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

invoke a command to activate file services on at least one node of a storage area network (SAN), wherein the SAN comprises a plurality of nodes;

install a virtual machine infrastructure on each of the plurality of nodes, wherein the virtual machine infrastructure comprises a plurality of virtual machines; and

enable the file services on each of the virtual machines.

15. The tangible, non-transitory, computer-readable medium of claim 14 comprising code to direct a processor to:

activate the file services to run on each of the virtual machines, wherein file services resources are reserved in advance on each node of the SAN; and

claim the file services resources for attachment to the plurality of virtual machines at the time of activating the file services.




FIG. 1A





FIG. 1B



3/5




200-A
FIG. 2A





4/5



200-B
FIG. 2B



5/5




FIG. 3

A. CLASSIFICATION OF SUBJECT MATTER**H04L 29/08(2006.01)i, G06F 9/455(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)

H04L 29/08; G06F 15/16; G06F 3/06; G06F 12/00; G06F 9/455; H04L 9/00

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:SAN, block-level, file service, VM(virtual machine), separation, storage

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 8910156 B1 (DEEPAK KENCHAMMANA-HOSEKOTE et al.) 09 December 2014 See column 4, lines 41-47; column 5, lines 17-59; column 6, lines 17-20; column 8, lines 41-45; column 9, lines 38-40; and figures 1-2.	1-15
A	US 2010-0154054 A1 (KENNETH S. BECK) 17 June 2010 See paragraphs [0007], [0011]; and figures 1, 7.	1-15
A	US 2012-0278450 A1 (STEVEN R. SOLTIS) 01 November 2012 See paragraphs [0074]-[0080], [0082]-[0084]; and figures 4-5.	1-15
A	WO 2009-145764 A1 (HEWLETT-PACKARD DEVELOPMENT COMPANY, L.P.) 03 December 2009 See paragraphs [0010]-[0015], [0025]-[0026]; and figures 1-2.	1-15
A	US 2014-0165062 A1 (BROCADE COMMUNICATIONS SYSTEMS, INC.) 12 June 2014 See paragraphs [0025]-[0029]; and figures 1, 2A.	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 September 2015 (25.09.2015)

Date of mailing of the international search report

29 October 2015 (29.10.2015)

Name and mailing address of the ISA/KR


International Application Division
Korean Intellectual Property Office
189 Cheongsu-ro, Seo-gu, Daejeon Metropolitan City, 35208,
Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

KIM, Seong Woo

Telephone No. +82-42-481-3348



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2015/013813

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 8910156 B1	09/12/2014	US 2015-046926 A1	12/02/2015
US 2010-0154054 A1	17/06/2010	US 2003-0028514 A1	06/02/2003
		US 2003-0078946 A1	24/04/2003
		US 2004-0139125 A1	15/07/2004
		US 2004-0143607 A1	22/07/2004
		US 2004-0210673 A1	21/10/2004
		US 2004-0249904 A1	09/12/2004
		US 2004-0250113 A1	09/12/2004
		US 2008-0005196 A1	03/01/2008
		US 2008-0010284 A1	10/01/2008
		US 2010-0146045 A1	10/06/2010
		US 2012-0192270 A1	26/07/2012
		US 2014-0032958 A1	30/01/2014
		US 6950833 B2	27/09/2005
		US 7593968 B2	22/09/2009
		US 7617292 B2	10/11/2009
		US 7640582 B2	29/12/2009
		US 7765329 B2	27/07/2010
		US 7814058 B2	12/10/2010
		US 8001222 B2	16/08/2011
		US 8151347 B2	03/04/2012
		US 8396908 B2	12/03/2013
		US 8578478 B2	05/11/2013
US 2012-0278450 A1	01/11/2012	US 2002-0083120 A1	27/06/2002
		US 2007-0094354 A1	26/04/2007
		US 2009-0240784 A1	24/09/2009
		US 7165096 B2	16/01/2007
		US 7552197 B2	23/06/2009
		US 8219639 B2	10/07/2012
		US 8583760 B2	12/11/2013
WO 2009-145764 A1	03/12/2009	CN 102047237 A	04/05/2011
		CN 102047237 B	10/12/2014
		US 2011-0078682 A1	31/03/2011
		US 8966476 B2	24/02/2015
US 2014-0165062 A1	12/06/2014	US 2011-0173608 A1	14/07/2011
		US 2011-0173609 A1	14/07/2011
		US 8719069 B2	06/05/2014