



(51) International Patent Classification:

*H04L 12/933 (2013.01)*

(21) International Application Number:

PCT/US2013/035228

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, 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.

(22) International Filing Date:

4 April 2013 (04.04.2013)

(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, ML, MR, NE, SN, TD, TG).

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

13/442,727 9 April 2012 (09.04.2012) US

(71) Applicant: DELL PRODUCTS L.P. [US/US]; One Dell Way, Round Rock, Texas 78682-2244 (US).

(72) Inventors: BRUNDIDGE, Michael, A.; 112 Shady Elm, Georgetown, Texas 78628 (US). CHANDRASEKHAR, Babu; 3233 Ranch Park Trail, Round Rock, Texas 78681 (US).

(74) Agent: OAKS, Brian, W.; Baker Botts L.L.P., 98 San Jacinto Blvd., Suite 1500, Austin, Texas 78701 (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, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD,

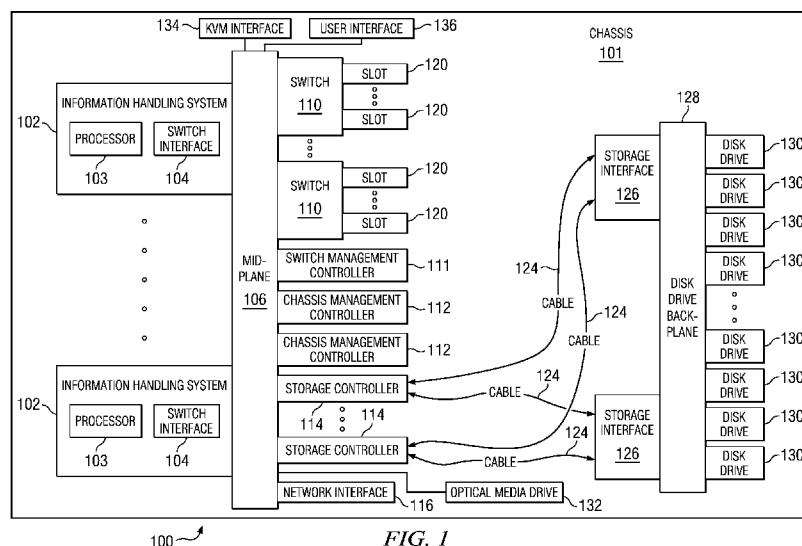
**Declarations under Rule 4.17:**

- 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))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

**Published:**

- with international search report (Art. 21(3))
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

(54) Title: METHODS AND SYSTEMS FOR VIRTUALIZATION OF STORAGE SERVICES IN AN INTEGRATED CHASSIS



100

FIG. 1

(57) **Abstract:** In accordance with embodiments of the present disclosure, a system may include a chassis, one or more chassis management controllers housed in the chassis, and a switch management controller. The chassis may be configured to receive a plurality of modular information handling systems. The one or more chassis management controllers may be configured to receive a storage management command, encapsulate the storage management command in a first datagram, and communicate the first datagram to a switch management controller housed in the chassis. The switch management controller may be configured to extract the storage management command from the first datagram, identify a storage controller associated with the storage management command, and communicate an input/output control request to the storage controller based on the storage management command.

**METHODS AND SYSTEMS FOR VIRTUALIZATION OF STORAGE  
SERVICES IN AN INTEGRATED CHASSIS**

**TECHNICAL FIELD**

The present disclosure relates in general to  
5 information handling systems, and more particularly to  
virtualization of storage services in an integrated  
chassis.

**BACKGROUND**

As the value and use of information continues to  
10 increase, individuals and businesses seek additional ways  
to process and store information. One option available  
to users is information handling systems. An information  
handling system generally processes, compiles, stores,  
and/or communicates information or data for business,  
15 personal, or other purposes thereby allowing users to  
take advantage of the value of the information. Because  
technology and information handling needs and  
requirements vary between different users or  
applications, information handling systems may also vary  
20 regarding what information is handled, how the  
information is handled, how much information is  
processed, stored, or communicated, and how quickly and  
efficiently the information may be processed, stored, or  
communicated. The variations in information handling  
25 systems allow for information handling systems to be  
general or configured for a specific user or specific use  
such as financial transaction processing, airline  
reservations, enterprise data storage, or global  
communications. In addition, information handling  
30 systems may include a variety of hardware and software  
components that may be configured to process, store, and  
communicate information and may include one or more

computer systems, data storage systems, and networking systems.

Existing server architectures either provide a single monolithic server capable of running one operating system and input/output (I/O) resources at a time, or 5 bulky blade server chassis providing multiple servers and I/O control modules in a single chassis. A system chassis with multiple information handling systems with various peripheral and input/output capabilities common 10 to the chassis as a whole may provide advantages, as it allows a blade server chassis in a small form factor, thereby providing a blade server chassis with a size comparable to the size of a monolithic server.

Implementation of a system chassis with multiple 15 information handling systems with various peripheral and input/output capabilities common to the chassis as a whole presents numerous challenges.

#### SUMMARY

In accordance with the teachings of the present 20 disclosure, the disadvantages and problems associated with virtualizing storage services in an integrated chassis have been reduced or eliminated.

In accordance with embodiments of the present disclosure, a system may include a chassis, one or more 25 chassis management controllers housed in the chassis, and a switch management controller. The chassis may be configured to receive a plurality of modular information handling systems. The one or more chassis management controllers may be configured to receive a storage management command, encapsulate the storage management command in a first datagram, and communicate the first 30 datagram to a switch management controller housed in the

chassis. The switch management controller may be configured to extract the storage management command from the first datagram, identify a storage controller associated with the storage management command, and 5 communicate an input/output control request to the storage controller based on the storage management command.

In accordance with these and other embodiments of the present disclosure, a method may include receiving a 10 storage management command at a chassis management controller housed in a chassis configured to receive a plurality of modular information handling systems, encapsulating the storage management command in a first datagram, communicating the first datagram from the 15 chassis management controller to a switch management controller housed in the chassis, extracting the storage management command from the first datagram, identifying a storage controller associated with the storage management command, and communicating an input/output control 20 request from the switch management controller to the storage controller based on the storage management command.

Technical advantages of the present disclosure will be apparent to those of ordinary skill in the art in view 25 of the following specification, claims, and drawings.

#### BRIEF DESCRIPTION OF THE DRAWINGS

A more complete understanding of the present embodiments and advantages thereof may be acquired by referring to the following description taken in 30 conjunction with the accompanying drawings, in which like reference numbers indicate like features, and wherein:

FIGURE 1 illustrates a block diagram of an example system chassis with multiple information handling systems and with various peripheral and input/output capabilities common to the chassis as a whole, in accordance with 5 certain embodiments of the present disclosure; and

FIGURE 2 illustrates a flow chart of an example method for virtualization of storage services in an integrated chassis, in accordance with certain embodiments of the present disclosure.

10 DETAILED DESCRIPTION

Preferred embodiments and their advantages are best understood by reference to FIGURES 1 and 2, wherein like numbers are used to indicate like and corresponding parts.

15 For the purposes of this disclosure, an information handling system may include any instrumentality or aggregate of instrumentalities operable to compute, classify, process, transmit, receive, retrieve, originate, switch, store, display, manifest, detect, record, reproduce, handle, or utilize any form of information, intelligence, or data for business, scientific, control, entertainment, or other purposes. For example, an information handling system may be a personal computer, a PDA, a consumer electronic device, a 20 network storage device, or any other suitable device and may vary in size, shape, performance, functionality, and price. The information handling system may include memory, one or more processing resources such as a central processing unit (CPU) or hardware or software 25 control logic. Additional components of the information handling system may include one or more storage devices, one or more communications ports for communicating with 30

external devices as well as various input and output (I/O) devices, such as a keyboard, a mouse, and a video display. The information handling system may also include one or more buses operable to transmit 5 communication between the various hardware components.

For the purposes of this disclosure, information handling resources may broadly refer to any component system, device or apparatus of an information handling system, including without limitation processors, busses, 10 memories, input-output devices and/or interfaces, storage resources, network interfaces, motherboards, electro-mechanical devices (e.g., fans), displays, and power supplies.

For the purposes of this disclosure, computer-readable media may include any instrumentality or aggregation of instrumentalities that may retain data and/or instructions for a period of time. Computer-readable media may include, without limitation, storage media such as a direct access storage device (e.g., a 20 hard disk drive or floppy disk), a sequential access storage device (e.g., a tape disk drive), compact disk, CD-ROM, DVD, random access memory (RAM), read-only memory (ROM), electrically erasable programmable read-only memory (EEPROM), and/or flash memory; as well as 25 communications media such wires, optical fibers, microwaves, radio waves, and other electromagnetic and/or optical carriers; and/or any combination of the foregoing.

Information handling systems often use an array of 30 physical storage resources (e.g., disk drives), such as a Redundant Array of Independent Disks (RAID), for example, for storing information. Arrays of physical storage

resources typically utilize multiple disks to perform input and output operations and can be structured to provide redundancy which may increase fault tolerance. Other advantages of arrays of physical storage resources 5 may be increased data integrity, throughput and/or capacity. In operation, one or more physical storage resources disposed in an array of physical storage resources may appear to an operating system as a single logical storage unit or "logical unit." Implementations 10 of physical storage resource arrays can range from a few physical storage resources disposed in a chassis, to hundreds of physical storage resources disposed in one or more separate storage enclosures.

FIGURE 1 illustrates a block diagram of an example 15 system 100 having a chassis 101 with multiple information handling systems 102 and with various peripheral and input/output capabilities common to chassis 101 as a whole, in accordance with certain embodiments of the present disclosure. As depicted in FIGURE 1, system 100 20 may comprise a chassis 101 including a plurality of information handling systems 102, a mid-plane 106, one or more switches 110, a switch management controller 111, one or more chassis management controller 112, a network 25 interface 116, one or more slots 120, one or more cables 124, one or more storage interfaces 126, a disk drive backplane 128, a plurality of disk drives 130, an optical media drive 132, a keyboard-video-mouse (KVM) interface 134, and a user interface 136.

An information handling system 102 may generally be 30 operable to receive data from and/or communicate data to one or more disk drives 130 and/or other information handling resources of chassis 101 via mid-plane 106. In

certain embodiments, an information handling system 102 may be a server. In such embodiments, an information handling system may comprise a blade server having modular physical design. In these and other embodiments, 5 an information handling system 102 may comprise an M class server. As depicted in FIGURE 1, an information handling system 102 may include a processor 103 and one or more switch interfaces 104 communicatively coupled to the processor 103.

10 A processor 103 may include any system, device, or apparatus configured to interpret and/or execute program instructions and/or process data, and may include, without limitation a microprocessor, microcontroller, digital signal processor (DSP), application specific 15 integrated circuit (ASIC), or any other digital or analog circuitry configured to interpret and/or execute program instructions and/or process data. In some embodiments, processor 103 may interpret and/or execute program instructions and/or process data stored in a memory, a 20 hard drive 130, and/or another component of system 100.

A switch interface 104 may comprise any system, device, or apparatus configured to provide an interface between its associated information handling system 102 and switches 110. In some embodiments, switches 110 may 25 comprise Peripheral Component Interconnect Express (PCIe) switches, in which case a switch interface 104 may comprise a mezzanine card configured to create a PCIe-compliant interface between its associated information handling system 102 and switches 110. In other 30 embodiments, a switch interface 104 may comprise an interposer. Use of switch interfaces 104 in information handling systems 102 may allow for minimal changes to be

made to traditional servers (e.g., M class servers) while supporting the overall system architecture disclosed herein. Although FIGURE 1 depicts an implementation including a single switch interface 104 per information 5 handling system 102, in some embodiments each information handling system 102 may include a plurality of switch interfaces 102 for redundancy, high availability, and/or other reasons.

Mid-plane 106 may comprise any system, device, or 10 apparatus configured to interconnect modular information handling systems 102 with information handling resources of chassis 101. Accordingly, mid-plane 106 may include slots and/or connectors configured to receive information handling systems 102, switches 110, switch management 15 controller 111, chassis management controllers 112, storage controllers 114, network interface 116, optical media drive 132, KVM interface 134, user interface 136, and/or other information handling resources. In one embodiment, mid-plane 106 may include a single board 20 configured to interconnect modular information handling systems 102 with information handling resources. In another embodiment, mid-plane 106 may include multiple boards configured to interconnect modular information handling systems 102 with information handling resources. 25 In yet another embodiment, mid-plane 106 may include cabling configured to interconnect modular information handling systems 102 with information handling resources.

A switch 110 may comprise any system, device, or 30 apparatus configured to couple information handling systems 102 to storage controllers 114 (e.g., via mid-plane 106) and slots 120 and perform switching between information handling systems 102 and various information

handling resources of system 100, including storage controllers 114 and slots 120. In certain embodiments, a switch 110 may comprise a PCIe switch. In other embodiments, a switch may comprise a generalized PC bus switch, an Infiniband switch, or other suitable switch. 5 As shown in FIGURE 1, chassis 101 may include a plurality of switches 110. In such embodiments, switches 110 may operate in a redundant mode for shared devices (e.g., storage controllers 114 and/or devices coupled to slots 120) and in non-redundant mode for non-shared/zoned devices. As used herein, shared devices may refer to those which may be visible to more than one information handling system 102, while non-shared devices may refer to those which are visible to only a single information 10 handling system 102.

15

A chassis management controller 112 may be any system, device, or apparatus configured to facilitate management and/or control of system 100, its information handling systems 102, and/or one or more of its component 20 its component information handling resources. A chassis management controller 102 may be configured to issue commands and/or other signals to manage and/or control information handling system 102 and/or information handling resources of system 100. A chassis management controller 112 may comprise a microprocessor, 25 microcontroller, digital signal processor (DSP), application specific integrated circuit (ASIC), field programmable gate array (FPGA), erasable programmable read-only memory (EPROM), or any combination thereof. As shown in FIGURE 1, a chassis management controller 112 may be coupled to mid-plane 106. Also as shown in FIGURE 30 1, system 100 may include a plurality of chassis

management controllers 112, and in such embodiments, chassis management controllers 112 may be configured as redundant. In some embodiments, a chassis management controller 112 may provide a user interface and high 5 level controls for management of switches 110, including configuring assignments of individual information handling systems 102 to non-shared information handling resources of system 100. In these and other embodiments, a chassis management controller may define configurations 10 of the storage subsystem (e.g., storage controllers 114, storage interfaces 126, disk drives 130, etc.) of system 100. For example, a chassis management controller may provide physical function configuration and status information that would normally occur at the driver level 15 in traditional server implementations. Examples of physical functions include disk drive discovery and status, RAID configuration and logical volume mapping.

In addition or alternatively, a chassis management controller 112 may also provide a management console for 20 user/administrator access to these functions. For example, a chassis management controller 112 may implement Intelligent Platform Management Interface (IPMI) or another suitable management protocol permitting a user to remotely access a chassis management controller 25 112 to configure system 100 and its various information handling resources. In such embodiments, a chassis management controller 112 may interface with a network interface separate from network interface 116, thus allowing for "out-of-band" control of 100, such that 30 communications to and from chassis management controller 112 are communicated via a management channel physically isolated from an "in band" communication channel with

network interface 116. Thus, for example, if a failure occurs in system 100 that prevents an administrator from interfacing with system 100 via network interface 116 and/or user interface 136 (e.g., operating system 5 failure, power failure, etc.), the administrator may still be able to monitor and/or manage system 100 (e.g., to diagnose problems that may have caused failure) via a chassis management controller 112. In the same or alternative embodiments, chassis management controller 10 112 may allow an administrator to remotely manage one or parameters associated with operation of system 100 and its various information handling resources (e.g., power usage, processor allocation, memory allocation, security privileges, etc.). Although FIGURE 1 depicts chassis as 15 having two chassis management controllers 112, chassis 101 may include any suitable number chassis management controllers 112.

A storage controller 114 may and include any system, apparatus, or device operable to manage the communication 20 of data between one or more of information handling systems 102 and one or more of disk drives 130. In certain embodiments, a storage controller 114 may provide functionality including, without limitation, disk aggregation and redundancy (e.g., RAID), input/output 25 (I/O) routing, and error detection and recovery. As shown in FIGURE 1, a storage controller 114 may be coupled to a connector on mid-plane 106. Also as shown in FIGURE 1, system 100 may include a plurality of storage controllers 114, and in such embodiments, storage controllers 114 may be configured as redundant. In addition or in the alternative, storage controllers 114 30 may in some embodiments be shared among two or more

information handling systems 102. As also shown in FIGURE 1, each storage controller 114 may be coupled to one or more storage interfaces 126 via cables 124. For example, in some embodiments, each storage controller 114 5 may be coupled to a single associated storage interface 126 via a cable 124. In other embodiments, each storage controller 114 may be coupled to two or more storage interfaces 126 via a plurality of cables 124, thus permitting redundancy as shown in FIGURE 1. Storage 10 controllers 114 may also have features supporting shared storage and high availability. For example, in PCIe implementations, a unique PCIe identifier may be used to indicate shared storage capability and compatibility in system 100.

15 In embodiments in which switches 110 comprise PCIe switches or switches configured in accordance with another communication standard, chassis management controllers 112 may not be able to communicate directly with switches 110. Available chassis management 20 controllers 112 may not be configured in accordance with PCIe or other communication standards available in switches 110, and it may not be desirable to create a proprietary chassis management controller 112 configured in accordance with PCIe or such other communication 25 standards for use in system 100 due to cost concerns. Accordingly, challenges may exist in enabling management functionality of storage devices (e.g., disk drives 130) via chassis management controllers 112 as direct PCIe communications (or communications via another 30 communication protocol) may not be available between

Accordingly, a switch management controller 111 may be communicatively interfaced between chassis

management controllers 112 and switches 110, and configured to virtualize management communications between chassis management controllers 112 and switches 110 related to management of storage components (e.g., 5 storage controllers 114, disk drives 130) and/or other components of system 100. Switch management controller 111 may, in some embodiments, interface with switches 110 via a private network (e.g., an Ethernet network) internal to chassis 101. In such embodiments, each 10 switch 110 and switch management controller 111 may establish a Transmission Control Protocol/Internet Protocol (TCP/IP) socket for communication. In these and other embodiments, switch management controller 111 may comprise a power PC management processor or processor 15 similar in structure and/or function.

In operation, a chassis management controller 112 may receive a storage management command (e.g., a storage application programming interface (API) call) from a management console or other interface. Chassis 20 management controller 112 may encapsulate such command in a network datagram (e.g., an Ethernet packet, frame, or other datagram) and communicate such datagram via the private network to switch management controller 111. Switch management controller 111 may receive such 25 datagram and extract the command from the datagram. Based on analysis of the command, storage management controller 111 may identify a storage controller 114 associated with the command and execute an input/output control request to such storage controller 114 via an 30 appropriate switch 110. The storage controller 114 may communicate a reply (e.g., a return code) via an appropriate switch to storage management controller 111,

which reply storage management controller 111 may encapsulate into a network datagram (e.g., an Ethernet packet, frame, or other datagram) and communicate such datagram via the private network to a chassis management controller 112. The chassis management controller 112 may extract the reply from the datagram and interpret such reply and/or forward the reply to a management console interfaced to the chassis management controller.

As depicted in FIGURE 1, a switch 110 may have coupled thereto one or more slots 120. A slot 120 may include any system, device, or apparatus configured to allow addition of one or more expansion cards to chassis 101 in order to electrically couple such expansion cards to a switch 110. Such slots 120 may comprise any suitable combination of full-height risers, full-height slots, and low-profile slots. A full-height riser may include any system, device, or apparatus configured to allow addition of one or more expansion cards (e.g., a full-height slot) having a physical profile or form factor with dimensions that practically prevent such expansion cards to be coupled in a particular manner (e.g., perpendicularly) to mid-plane 106 and/or switch 110 (e.g., the proximity of information handling resources in chassis 101 prevents physical placement of an expansion card in such manner). Accordingly, a full-height riser may itself physically couple with a low-profile to mid-plane 106, a switch 110, or another components, and full-height cards may then be coupled to full-height slots of full-height riser. On the other hand, low-profile slots may be configured to couple low-profile expansion cards to switches 110 without the need for a full-height riser.

Slots 120 may also include electrically conductive elements (e.g., edge connectors, traces, etc.) allowing for expansion cards inserted into slots 120 to be electrically coupled to switches 110. In operation, 5 switches 110 may manage switching of communications between individual information handling systems 102 and expansion cards coupled to slots 120. In some embodiments, slots 120 may be nonshared (e.g., each slot 120 is associated with a single information handling 10 system 102). In other embodiments, one or more of slots 120 may be shared among two or more information handling systems 102. In these and other embodiments, one or more slots 120 may be configured to be compatible with PCIe, generalized PC bus switch, Infiniband, or other suitable 15 communication specification, standard, or protocol.

Network interface 116 may include any suitable system, apparatus, or device operable to serve as an interface between chassis 101 and an external network (e.g., a local area network or other network). Network 20 interface 116 may enable information handling systems 102 to communicate with the external network using any suitable transmission protocol (e.g., TCP/IP) and/or standard (e.g., IEEE 802.11, Wi-Fi). In certain embodiments, network interface 116 may include a network 25 interface card (NIC). In the same or alternative embodiments, network interface 116 may be configured to communicate via wireless transmissions. In the same or alternative embodiments, network interface 116 may provide physical access to a networking medium and/or 30 provide a low-level addressing system (e.g., through the use of Media Access Control addresses). In some embodiments, network interface 116 may be implemented as

a local area network (LAN) on motherboard (LOM) interface.

In some embodiments, various components of chassis 101 may be coupled to a planar. For example, a planar 5 may interconnect switches 110, chassis management controller 112, storage controllers 114, network interface 116, optical media drive 132, KVM interface 134, user interface 136, and/or other modular information handling resources of chassis 101 to mid-plane 106 of 10 system 100. Accordingly, such planar may include slots and/or connectors configured to interconnect with such information handling resources.

Storage interfaces 126 may include any system, device, or apparatus configured to facilitate 15 communication between storage controllers 114 and disk drives 130. For example, a storage interface may serve to permit a relatively small number of communication links (e.g., two) between storage controllers 114 and storage interfaces 126 to communicate with a greater 20 number of disk drives 130. Thus, a storage interface 126 may provide a switching mechanism and/or disk drive addressing mechanism that allows an information handling system 102 to communicate with numerous disk drives 130 via a limited number of communication links and/or 25 channels. Accordingly, a storage interface 126 may operate like an Ethernet hub or network switch that allows multiple systems to be coupled using a single switch port (or relatively few switch ports). A storage interface 126 may be implemented as an expander (e.g., a 30 Serial Attached SCSI (SAS) expander), an Ethernet switch, a FibreChannel switch, Internet Small Computer System Interface (iSCSI) switch, or any other suitable switch.

In order to support high availability storage, system 100 may implement a plurality of redundant storage interfaces 126, as shown in FIGURE 1.

Disk drive backplane 128 may comprise any system, 5 device, or apparatus configured to interconnect modular storage interfaces 126 with modular disk drives 130. Accordingly, disk drive backplane 128 may include slots and/or connectors configured to receive storage 10 interfaces 126 and/or disk drives 130. In some embodiments, system 100 may include two or more backplanes, in order to support differently-sized disk 15 drive form factors. To support redundancy and high availability, a backplane 128 may be configured to receive a plurality (e.g., 2) of storage interfaces 126 which couple two storage controllers 114 to each disk drive 130.

Each disk drive 130 may include computer-readable media (e.g., magnetic storage media, optical storage media, opto-magnetic storage media, and/or other type of 20 rotating storage media, flash memory, and/or other type of solid state storage media) and may be generally operable to store data and/or programs (e.g., one or more operating systems and/or one or more application programs). Although disk drives 130 are depicted as 25 being internal to chassis 101 in FIGURE 1, in some embodiments, one or more disk drives may be located external to chassis 101 (e.g., in one or more enclosures external to chassis 101).

Optical media drive 132 may be coupled to mid-plane 30 106 and may include any suitable system, apparatus, or device configured to read data from and/or write data to an optical storage medium (e.g., a compact disc (CD),

digital versatile disc (DVD), blue laser medium, and/or other optical medium). In certain embodiments, optical media drive 132 may use laser light or other electromagnetic energy to read and/or write data to an 5 optical storage medium. In some embodiments, optical media drive 132 may be nonshared and may be user-configurable such that optical media drive 132 is associated with a single information handling system 102.

KVM interface 134 may be coupled to mid-plane 106 10 and may include any suitable system, apparatus, or device configured to couple to one or more of a keyboard, video display, and mouse and act as switch between multiple information handling systems 102 and the keyboard, video display, and/or mouse, thus allowing a user to interface 15 with a plurality of information handling systems 102 via a single keyboard, video display, and/or mouse.

User interface 136 may include any system, apparatus, or device via which a user may interact with system 100 and its various information handling resources 20 by facilitating input from a user allowing the user to manipulate system 100 and output to a user allowing system 100 to indicate effects of the user's manipulation. For example, user interface 136 may include a display suitable for creating graphic images 25 and/or alphanumeric characters recognizable to a user, and may include, for example, a liquid crystal display (LCD), cathode ray tube (CRT), a plasma screen, and/or a digital light processor (DLP) projection monitor. In certain embodiments, such a display may be an integral 30 part of chassis 101 and receive power from power supplies (not explicitly shown) of chassis 101, rather than being coupled to chassis 101 via a cable. In some embodiments,

such display may comprise a touch screen device capable of receiving user input, wherein a touch sensor may be mechanically coupled or overlaid upon the display and may comprise any system, apparatus, or device suitable for 5 detecting the presence and/or location of a tactile touch, including, for example, a resistive sensor, capacitive sensor, surface acoustic wave sensor, projected capacitance sensor, infrared sensor, strain gauge sensor, optical imaging sensor, dispersive signal 10 technology sensor, and/or acoustic pulse recognition sensor. In these and other embodiments, user interface 136 may include other user interface elements (e.g., a keypad, buttons, and/or switches placed in proximity to a display) allowing a user to provide input to system 100. 15 User interface 136 may be coupled to chassis management controllers 112 and/or other components of system 100, and thus may allow a user to configure various information handling resources of system 100 (e.g., assign individual information handling systems 102 to 20 particular information handling resources).

FIGURE 2 illustrates a flow chart of an example method for virtualization of storage services in an integrated chassis, in accordance with certain embodiments of the present disclosure. According to 25 certain embodiments, method 200 may begin at step 202. As noted above, teachings of the present disclosure may be implemented in a variety of configurations of system 100. As such, the preferred initialization point for method 200 and the order of the steps 202-220 comprising 30 method 200 may depend on the implementation chosen.

At step 202, a chassis management controller (e.g., a chassis management controller 112) may receive a

storage management command (e.g., a storage application programming interface (API) call) from a management console or other interface communicatively coupled to the chassis management controller and encapsulate the storage 5 management command in a datagram (e.g., an Ethernet packet, frame, or other datagram).

At step 204, the chassis management controller may communicate the datagram to a switch management controller (e.g., switch management controller 111). 10 Such communication may be made over a private network of a chassis housing both of the chassis management controller and the switch management controller. In some embodiments, such communication may be made using TCP/IP.

At step 206, the switch management controller may extract (e.g., decapsulate) the storage management command from the datagram. At step 208, based on 15 analysis of the command, the storage management controller may identify a storage controller (e.g., a storage controller 114) associated with the command. At step 210, the switch management controller may communicate via an appropriate switch (e.g., a switch 110) an input/output control request based on the storage management command to the identified storage controller.

At step 212, the switch management controller may 25 receive from the identified storage controller via an appropriate switch a reply to the input/output control request. At step 214, the switch management controller may encapsulate the storage management command in a datagram (e.g., an Ethernet packet, frame, or other 30 datagram). At step 216, the switch management controller may communicate the datagram to the chassis management controller. Such communication may be made over a

private network of a chassis housing both of the chassis management controller and the switch management controller. In some embodiments, such communication may be made using TCP/IP.

5 At step 218, the chassis management controller may extract (e.g., decapsulate) the reply from the datagram. At step 220, the chassis management controller may communicate the reply to the sender (e.g., a management console) of the storage management controller. After 10 completion of step 220, method 200 may end.

The storage controller 114 may communicate a reply (e.g., a return code) to storage management controller 111, which reply storage management controller 111 may encapsulate into a network datagram (e.g., an Ethernet 15 packet, frame, or other datagram) and communicate such datagram via the private network to a chassis management controller 112. The chassis management controller 112 may extract the reply from the datagram and interpret such reply and/or forward the reply to a management 20 console interfaced to the chassis management controller.

Although FIGURE 2 discloses a particular number of steps to be taken with respect to method 200, method 200 may be executed with greater or lesser steps than those depicted in FIGURE 2. In addition, although FIGURE 2 25 discloses a certain order of steps to be taken with respect to method 200, the steps comprising method 200 may be completed in any suitable order.

Method 200 may be implemented using system 100, components thereof or any other system operable to 30 implement method 200. In certain embodiments, method 200 may be implemented partially or fully in software and/or firmware embodied in computer-readable media.

Although the present disclosure has been described in detail, it should be understood that various changes, substitutions, and alterations can be made hereto without departing from the spirit and the scope of the disclosure 5 as defined by the appended claims.

WHAT IS CLAIMED IS:

1. A system comprising:
  - a chassis configured to receive a plurality of modular information handling systems;
  - one or more chassis management controllers housed in the chassis and configured to:
    - receive a storage management command;
    - encapsulate the storage management command in a first datagram; and
    - communicate the first datagram to a switch management controller housed in the chassis; and
  - the switch management controller configured to:
    - extract the storage management command from the first datagram;
    - identify a storage controller associated with the storage management command; and
    - communicate an input/output control request to the storage controller based on the storage management command.
2. A system according to Claim 1, the command originating from a management console communicatively coupled to the one or more chassis management controllers.
3. A system according to Claim 1, the datagram comprising an Ethernet frame.

4. A system according to Claim 1, the one or more chassis management controllers configured to communicate the first datagram to the switch management controller via Transmission Control Protocol/Internet Protocol.

5. A system according to Claim 1, the one or more chassis management controllers configured to communicate the first datagram to the switch management controller via a private network internal to the chassis.

6. A system according to Claim 1, the switch management controller configured to communicate the input/output control request to the storage controller via a switch.

7. A system according to Claim 6, the switch compliant with Peripheral Component Interconnect Express.

8. A system according to Claim 1, the switch management controller further configured to:

- receive a reply to the input/output control request from the storage controller;
- encapsulate the reply in a second datagram; and
- communicate the second datagram to the one or more chassis management controllers.

9. A system according to Claim 8, the chassis management controller further configured to:

- extract the reply from the second datagram; and
- communicate the reply to an originator of the storage management command.

10. A system according to Claim 1, the chassis management controller further configured to:

receive a second datagram having encapsulated therein a reply to the input/output control request;

extract the reply from the second datagram; and communicate the reply to an originator of the storage management command.

11. A method comprising:

receiving a storage management command at a chassis management controller housed in a chassis configured to receive a plurality of modular information handling systems;

encapsulating the storage management command in a first datagram;

communicating the first datagram from the chassis management controller to a switch management controller housed in the chassis;

extracting the storage management command from the first datagram;

identifying a storage controller associated with the storage management command; and

communicating an input/output control request from the switch management controller to the storage controller based on the storage management command.

12. A method according to Claim 11, the command originating from a management console communicatively coupled to the chassis management controller.

13. A method according to Claim 11, the datagram comprising an Ethernet frame.

14. A method according to Claim 11, wherein communicating the first datagram comprises communicating the first datagram via Transmission Control Protocol/Internet Protocol.

15. A method according to Claim 11, wherein communicating the first datagram comprises communicating the first datagram via a private network internal to the chassis.

16. A method according to Claim 11, wherein communicating the input/output control request comprises communicating the input/output control request to the storage controller via a switch.

17. A method according to Claim 16, the switch compliant with Peripheral Component Interconnect Express.

18. A method according to Claim 11, further comprising:

receiving a reply to the input/output control request from the storage controller;

encapsulating the reply in a second datagram; and communicating the second datagram from the switch management controller to the chassis management controller.

19. A method according to Claim 18, further comprising:

extracting the reply from the second datagram; and communicating the reply to an originator of the storage management command.

20. A method according to Claim 11, further comprising:

receiving at the chassis management controller a second datagram having encapsulated therein a reply to the input/output control request;

extracting the reply from the second datagram; and communicating the reply to an originator of the storage management command.

1/2

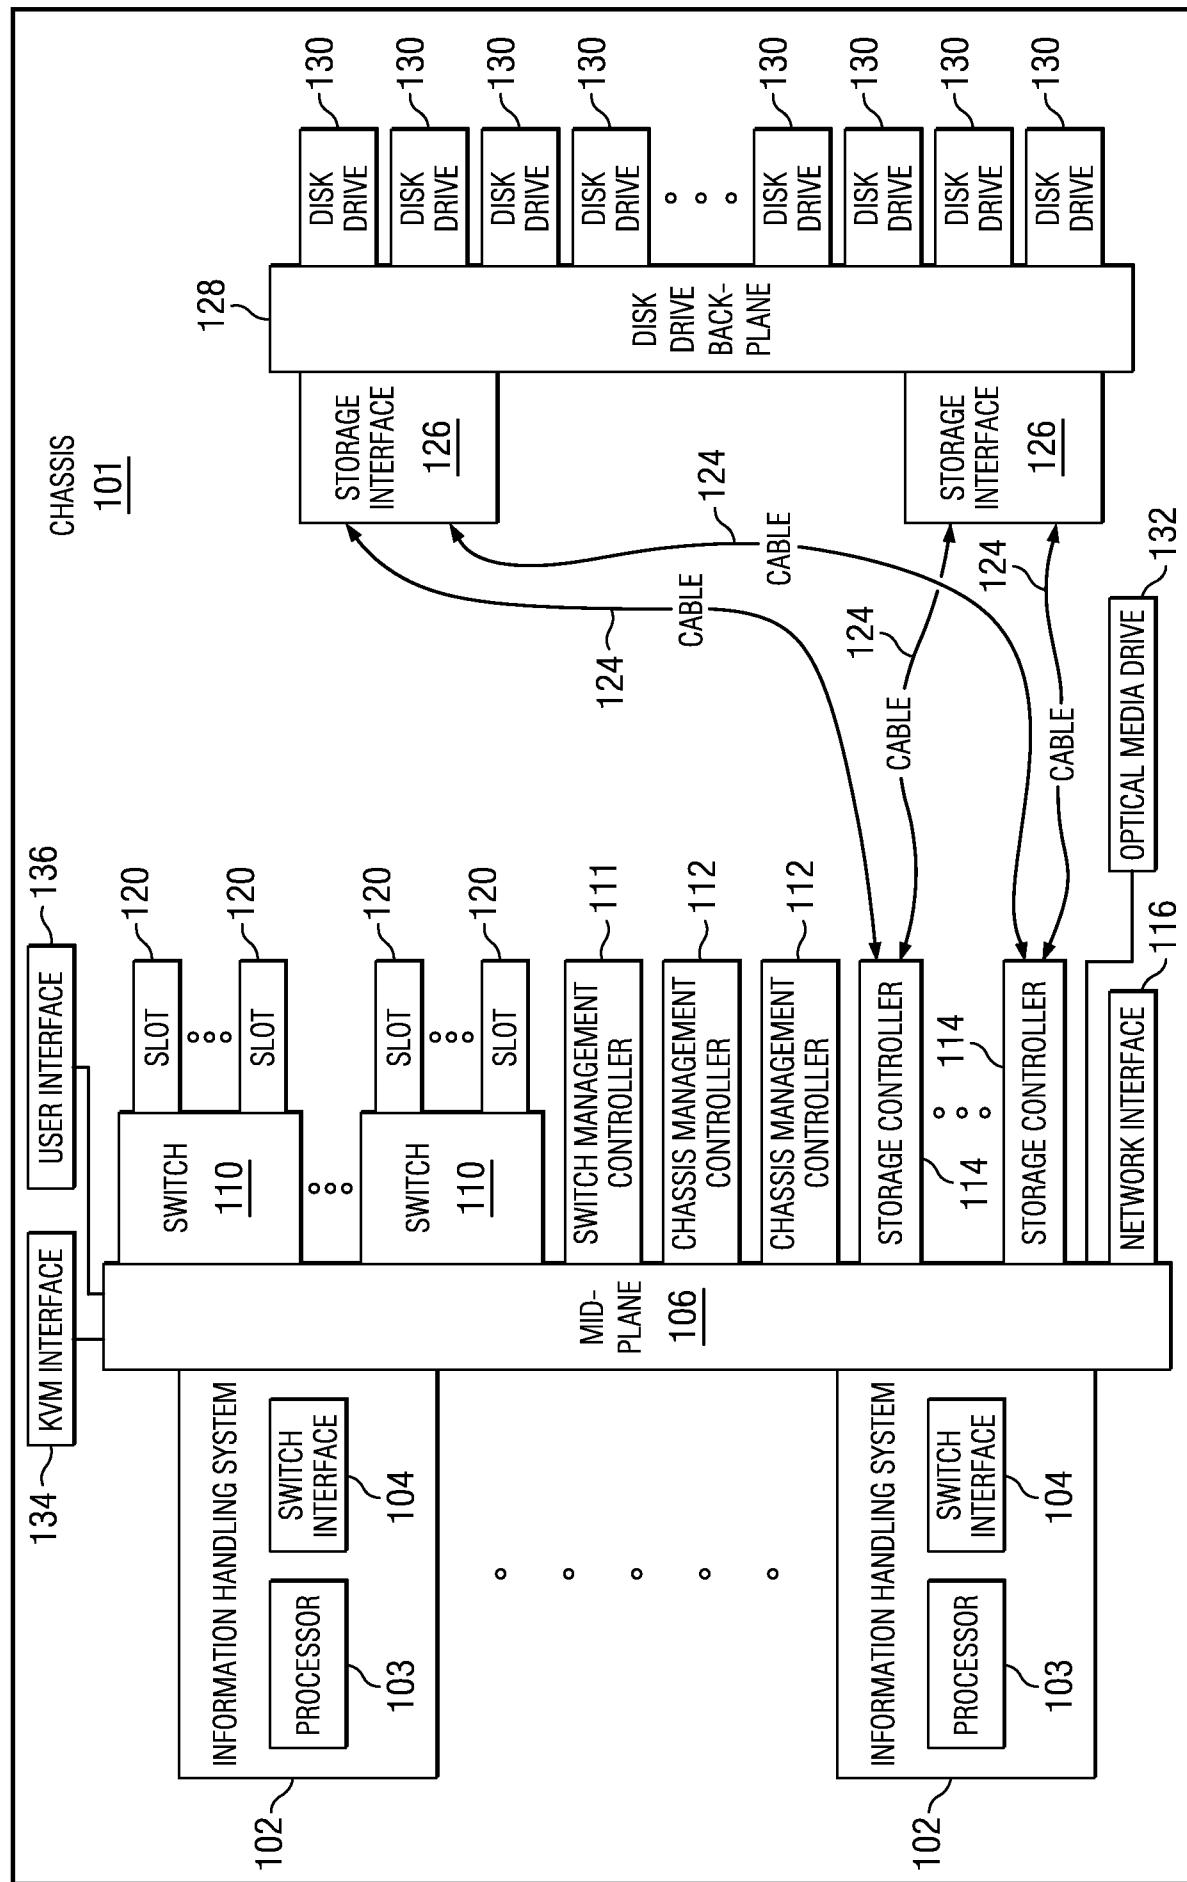


FIG. 1

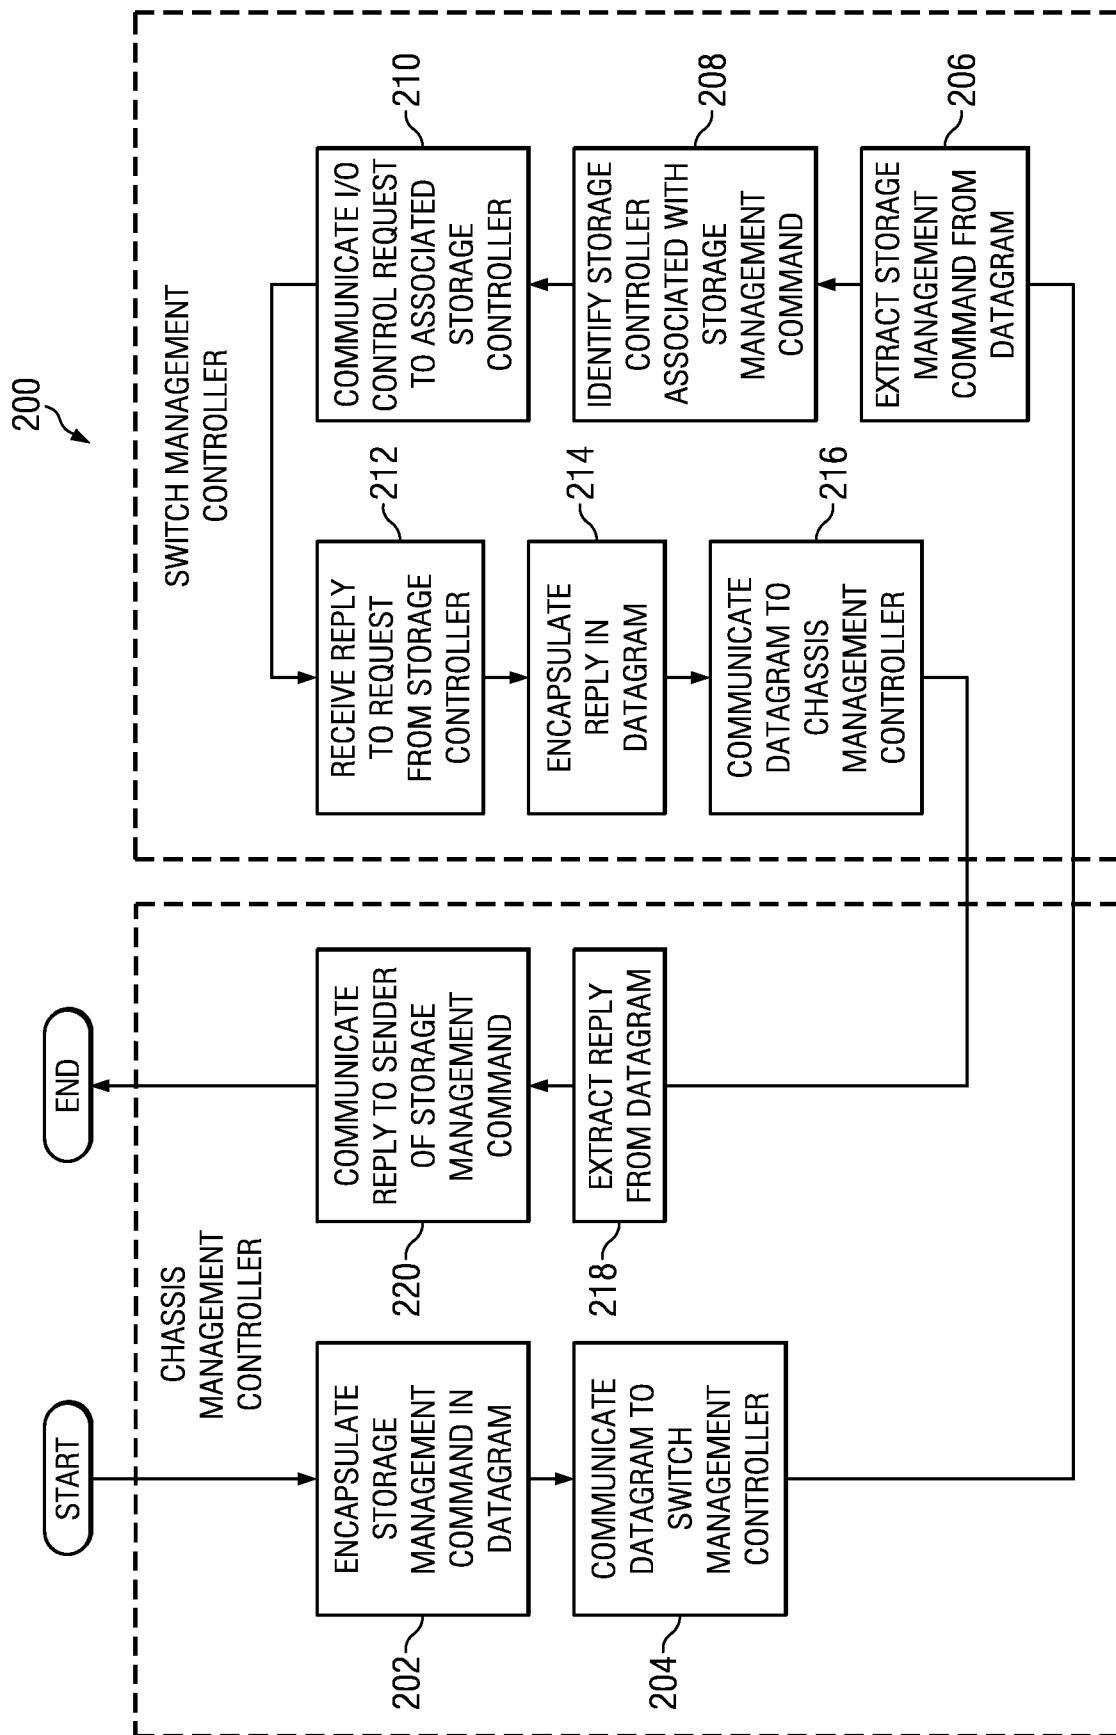


FIG. 2

# INTERNATIONAL SEARCH REPORT

International application No  
PCT/US2013/035228

**A. CLASSIFICATION OF SUBJECT MATTER**  
INV. H04L12/933  
ADD.

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

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

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

EPO-Internal

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2012/063304 A1 (GNANASEKARAN SATHISH KUMAR [US] ET AL) 15 March 2012 (2012-03-15) page 1, paragraph 2 page 2, paragraph 25 page 3, paragraph 41 - paragraph 42 page 4, paragraph 59 page 5, paragraph 68 page 6, paragraph 90 page 8, paragraph 123 page 9, paragraph 127 ----- -/--	1-20



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	Date of mailing of the international search report
5 August 2013	09/08/2013
Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer  Mircescu, Alexander

## INTERNATIONAL SEARCH REPORT

International application No
PCT/US2013/035228

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 01/16763 A1 (IBM [US]) 8 March 2001 (2001-03-08) page 1, line 5 - line 12 page 7, line 26 - line 31 page 8, line 1 - line 20 page 9, line 6 - line 22 page 23, line 16 - line 24 page 24, line 11 - line 20 ----- US 6 980 543 B1 (KASTENHOLZ FRANK [US] ET AL) 27 December 2005 (2005-12-27) column 1, line 34 - line 54 column 3, line 24 - line 37 column 8, line 33 - line 44 column 15, line 4 - line 20 -----	1-20
A		1-20

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No PCT/US2013/035228
---

Patent document cited in search report	Publication date	Patent family member(s)			Publication date
US 2012063304	A1	15-03-2012	NONE		
-----					
WO 0116763	A1	08-03-2001	AT 333678 T CN 1399738 A DE 60029467 T2 EP 1226501 A1 ES 2265971 T3 JP 4066382 B2 JP 2003508954 A JP 2006020371 A US 6985431 B1 US 2005232204 A1 US 2005232205 A1 US 2005232270 A1 US 2005243850 A1 WO 0116763 A1	15-08-2006 26-02-2003 19-07-2007 31-07-2002 01-03-2007 26-03-2008 04-03-2003 19-01-2006 10-01-2006 20-10-2005 20-10-2005 20-10-2005 03-11-2005 08-03-2001	
-----					
US 6980543	B1	27-12-2005	US 6980543 B1 US 2006007946 A1 US 2010067523 A1	27-12-2005 12-01-2006 18-03-2010	
-----					