



- (51) **International Patent Classification:**  
*G06F 13/00* (2006.01) *H04N 21/433* (2011.01)
- (21) **International Application Number:**  
PCT/US2014/061633
- (22) **International Filing Date:**  
21 October 2014 (21.10.2014)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**  
14/135,213 19 December 2013 (19.12.2013) US
- (71) **Applicant:** SONY COMPUTER ENTERTAINMENT AMERICA LLC [US/US]; 2207 Bridgepointe Parkway, San Mateo, California 94404 (US).
- (72) **Inventor:** COLENBRANDER, Roelof Roderick; c/o Sony Computer Entertainment America LLC, 2207 Bridgepointe Parkway, San Mateo, California 94404 (US).
- (74) **Agent:** ISENBERG, Joshua D.; c/o JDI Patent, 809 Corporate Way, Fremont, California 94539 (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).

**Published:**

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

(54) **Title:** MASS STORAGE VIRTUALIZATION FOR CLOUD COMPUTING

(57) **Abstract:** Aspects of the present disclosure relate to systems and methods for virtualizing mass storage using an interface card, or other special hardware unit, operatively coupled to a host system. In various implementations, the interface device and host system may collectively emulate a mass storage device for another terminal computing system in a manner that is transparent to the terminal system. Moreover, in various implementations, a mass storage device may be emulated in a manner that is transparent to the terminal system, without requiring modifications to the terminal platforms hardware or software architecture.

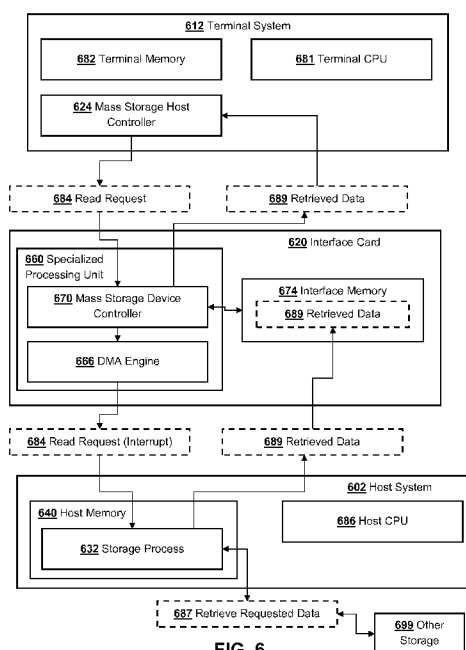


FIG. 6

## MASS STORAGE VIRTUALIZATION FOR CLOUD COMPUTING

### CLAIM OF PRIORITY

This application claims the benefit of prior to commonly-assigned, co-pending US Non-Provisional application number 14/135,213, to Roelof Roderick Colenbrander, entitled “MASS STORAGE VIRTUALIZATION FOR CLOUD COMPUTING”, (Attorney Docket No.

5 SCEA13031US00), filed December 19, 2013, the entire contents of which are herein incorporated by reference.

### CROSS REFERENCE TO RELATED APPLICATIONS

This application is related to commonly-assigned, co-pending US application number  
10 14/135,374, to Roelof Roderick Colenbrander, entitled “VIDEO LATENCY REDUCTION”, (Attorney Docket No. SCEA13037US00), filed December 19, 2013, the entire contents of which are herein incorporated by reference.

### FIELD

Aspects present disclosure relates to storage virtualization. Further aspects of the present  
15 disclosure relate to storage virtualization for cloud computing applications designed for embedded platforms.

### BACKGROUND

With bandwidths and data transfer rates of telecommunication systems continuing to advance, many of today’s computing needs are shifting towards network-based distributed computing  
20 systems, commonly referred to as “cloud computing.” Generally speaking, cloud computing involves offloading certain computing tasks to one or more remote computing resources and utilizing the remote computing resources to provide services to a local device. Cloud computing is often implemented as a client-server model, with the remote server providing desired  
25 computing capabilities to a local client device, avoiding the need for the local client device to perform certain computing tasks on its own. In many instances, the computing resources can be entirely or almost entirely offloaded to the server, with the client providing little or no more than a user interface for the cloud computing task, allowing potentially any internet connected device to serve as a client for a particular application.

Cloud computing has been adopted in a wide variety of computing services, including remote data storage solutions, music and video streaming services, and a wide variety of other consumer and business solutions. In general, cloud computing provides several well-known benefits, including increased efficiencies through the use of shared resources, increased flexibility in  
5 accessing content without restriction to a particular device, lower up-front costs when hardware upgrades become available, and other advantages.

One arena in which cloud computing has yet to achieve widespread adoption is in the ever popular field of video gaming. Video games, including personal computer (PC) games, home console games, handheld console games, and the like, remain a popular source of entertainment  
10 among consumers. With a large and ever growing collection of available video game titles, cloud gaming provides a particularly attractive solution to many video game demands. Ideally, within a cloud gaming framework a user would be able to instantly access any video game title from among a vast collection of remotely stored titles, without needing a physical copy of the game, without having to wait for the title to download, and without locally possessing the hardware  
15 system or resources needed to run the game. Unfortunately, implementing this type of service within a traditional cloud computing server design presents several technical challenges that have thus far prevented widespread adoption of cloud gaming for many video game platforms.

One challenge is that many of today's video games are very fast paced and action packed, with advanced graphics that require a significant amount of computing resources to render.

20 Particularly in a video streaming based server design, in which all of the game computations and graphics rendering computations are performed remotely, the game needs be able to perform all these tasks and deliver a compressed audio/video stream to the client device with minimal to no perceptible latency in response to control inputs. Otherwise, this may disrupt the timing for the player and fail to replicate the responsiveness of a traditional locally executed gaming session.

25 Another challenge stems from the fact that there has traditionally been a distinction between PC games on the one hand, which are designed to run on personal computers, laptops, and other computing systems of a general purpose nature, and console games on the other hand, which are designed to run on specialized systems known as video game consoles. Video game consoles are typically designed as embedded platforms having unique computing architectures specially

adapted for the video gaming environment. Using dedicated console hardware as a video gaming platform provides several benefits, particularly in the form of optimizing the hardware for the high computing resource demands of graphics rendering and gaming computations required to produce many of today's video games. As a result, while PC games have achieved some popularity, console games have traditionally dominated the video game market and may continue to do so in the future.

However, adapting video game consoles and other embedded platforms to a cloud streaming service presents several challenges. Extensive modifications of the hardware or software architecture of the specialized platform may present compatibility issues with applications designed for the specialized platform. As a result, traditional storage server and storage virtualization techniques are not a viable solution for cloud based streaming of applications designed for many of these platforms.

It is within this context that the present disclosure arises.

#### SUMMARY

An implementation of the present disclosure may include an interface device having a first communication interface, a second communication interface, and a specialized processing unit operatively coupled to said communication interfaces. The first communication interface is a mass storage interface. The second communication interface is an expansion interface. The specialized processing unit includes a mass storage device controller connected to the first communication interface and a direct memory access (DMA) engine connected to the second communication interface. The mass storage device controller is configured to receive a mass storage data request through the first communication interface. The DMA engine is configured to notify a host system of the mass storage data request through the second communication interface.

Another implementation of the present disclosure may include a system having a central processing unit; a first memory unit operatively coupled to the central processing unit; a storage process embodied in the first memory unit and executable by the central processing unit (CPU); a first communication interface; and a specialized processing unit operatively coupled to the first communication interface. The first communication interface is a mass storage interface.

The specialized processing unit has a mass storage device controller connected to the first communication interface and a memory access interface operatively coupled to the CPU. The mass storage device controller is configured to receive a mass storage data request from through the first communication interface. The memory access interface is configured to notify the storage process of the mass storage data request. The storage process is configured to process the data request by accessing a storage system.

Another implementation may include a method comprising: a) receiving, from a first computing device, a mass storage data request at a mass storage device controller; b) notifying, with a memory access interface connected to the mass storage device controller, a storage process embodied in a second computing device of the data request; c) accessing, with the storage process of the second computing device, a storage system having data corresponding to the data request; d) storing, with the storage process of the second computing device, the data corresponding to the data request on a first memory unit; e) programming, with the storage process of the second computing device, the mass storage device controller to access the data from the first memory unit in order to complete the data request.

#### BRIEF DESCRIPTION OF THE DRAWINGS

The teachings of the present disclosure can be readily understood by considering the following detailed description in conjunction with the accompanying drawings, in which:

FIG. 1 is a schematic diagram of an example streaming technique over a network.

FIG. 2 is a schematic diagram of another example streaming technique over a network.

FIG. 3 is a schematic diagram of an example host system having an interface card to provide storage virtualization for a terminal system.

FIG. 4A-4B is a schematic diagram of an example computer architecture for a host system coupled to a terminal system. FIG. 4A is a schematic diagram of an example terminal system architecture. FIG. 4B is an example host system and interface card architecture.

FIG. 5 is a schematic diagram of an example interface card design having a specialized processing unit.

FIG. 6 is a schematic diagram of an example method of emulating a mass storage device.

### DETAILED DESCRIPTION

Although the following detailed description contains many specific details for the purposes of illustration, anyone of ordinary skill in the art will appreciate that many variations and alterations to the following details are within the scope of the invention. Accordingly, the exemplary  
5 embodiments of the invention described below are set forth without any loss of generality to, and without imposing limitations upon, the claimed invention.

#### **Introduction**

Aspects of the present disclosure relate to systems and methods for virtualizing mass storage using an interface card, or other special hardware unit, operatively coupled to a host system. In various implementations, the interface device and host system may collectively emulate a mass storage device for another "terminal" computing system in a manner that is transparent to the terminal system. Moreover, in various implementations, a mass storage device may be emulated in a manner that is transparent to the terminal system, without requiring modifications to the  
15 terminal platform's hardware or software architecture.

Those skilled in the art may be familiar with 'virtual machine' software, such as VMWare, VirtualBox or Parallels, which allow, e.g., a 'virtual' session of an older operating system (e.g., Windows XP) to run within a session of a newer operating system (e.g., Windows 7). Virtual machine software is often used to simulate physical hardware components, such as disk drives,  
20 network cards, video cards in software for compatibility reasons. According to aspects of the present disclosure specialized hardware and software can virtualize a mass storage device in situations where it is impractical to do so with software alone.

#### **Implementation Details**

To better illustrate various aspects of the present disclosure, a first illustrative example of a distributed computing system **100** is depicted in FIG. 1. The distributed system **100** depicted in  
25 FIG. 1 involves a plurality of computing devices **102,104** connected over a network **106**, such as the internet.

In the example distributed system, a streaming server **102** executes a program **108** in order to provide a real-time stream of the program output to the remotely connected client device **104** through the network **106**. In the illustrated schematic diagram of FIG. 1, the program **104** may be a video game program which processes control inputs received from the client device **104** and sends a compressed audio and video stream to the client device **104** in accordance with the execution of the program **108**. A user located at the client side may interact with the program through a user interface of the client device **104**, e.g. by providing control inputs and receiving the audio and video game output at a display and/or speakers connected to the client device **104**. In the example of FIG. 1, the computational load for processing the inputs, rendering the video game graphics, and the like, may be entirely offloaded to the streaming server **102** for remote execution.

In order to implement the cloud gaming using the system **100** depicted in FIG. 1, the streaming server **102** may be configured to not only execute the game program **108**, but the streaming server **102** may also implement a “streaming server process” **110** embodied in the server **102**, which may perform various functions to coordinate the program inputs/outputs for communication between the client device **104** and the streaming server **102** over the network **106**. By way of example, the streaming server process **110** may be embodied in one or more applications which are run on the streaming server **102** in coordination with the game software **108**. Remote execution of the game **108** together with the streaming server process **110** may be performed with minimal or no latency between the inputs and outputs at the client device **104** so that the user at the client side perceives that the program **108** is being executed locally.

Moreover, a wide collection of remotely stored programs may be accessible for execution. For example, in a cloud gaming application, the user would like to be able to select any game designed for the relevant video game platform, and the remote server would be able to load the game, execute it, and streams its output to the client device in accordance with inputs received from the user and with no perceptible latency between the received inputs and audio/video outputs.

In the example depicted in FIG. 1, this may be accomplished by executing both the streaming server process **110** and the game code **108** at the same computing device, e.g., the streaming server **102**. The streaming server may access a selected program from potentially any storage

device, such as a locally connected mass storage device, a storage server, a network, and the like, through a suitable connection. For cloud gaming, and other cloud computing implementations, this may be readily implemented if the software application **108** that is executed at the server **102** for the client **104** is designed to run personal computer (PC) architecture or other computer  
5 architecture that is similar to the streamliner server **102**. By way of example, and not by way of limitation, for a PC cloud gaming implementation, a readily available x86 based server may be used to not only implement the streaming server process **110**, but also execute the PC game program **108** once it is accessed by the server **102**.

However, the example technique **100** depicted in FIG. 1 may not be suitable for certain  
10 applications in which the program **108** is designed to be executed on embedded platforms and other specialized architectures. For example, if the desired program is a video game designed for a dedicated console platform, the game software may not be able to run on the server architecture without some type of emulator. Unfortunately, in a cloud computing environment in which  
15 minimizing latency is a chief concern, a software emulator executed on a streaming server may be an undesirable solution, particularly in graphics intensive applications where interpretation of the code by the emulator at run-time would introduce a perceptible delay. Moreover, in many cases it would be impractical to build a dedicated server for each unique architecture for which a cloud service is desired. In other situations, an emulator could be a desirable solution, but a suitable emulator might not be available.

20 In order to address these challenges, it may be desirable to utilize an additional computing device at the server side for execution of the program, whereby the additional computing device executes the desired program while the streaming server coupled thereto performs its usual functions for interfacing with a client device. An illustrative example of a distributed system **200**  
25 which utilizes such a technique is depicted in FIG. 2, in which a “terminal system” **212** is connected to a streaming server **202**, and the terminal system executes a desired application **208**. The streaming server can provide resources for network storage and video compression if the terminal system does not otherwise have sufficient free resources for these functions.

In certain implementations, the terminal system **212** may be an embedded system or other computing system having a specialized architecture, such as a video game console, while the

streaming server **202** may be based on any computing architecture, such as an x86 instruction set architecture, which is adapted to provide streaming services and does not necessarily have to have a similar architecture as the terminal system **212**. The streaming server **202** may be configured to execute a streaming server process **202** and provide the usual cloud services for the terminal system program **208**, and the computing devices at the server side, e.g. the streaming server **202** and the terminal system **212**, may collectively operate to provide video game streaming services for a client device **204**.

The configuration in the simplified schematic depicted in FIG. **2** may provide a practical and flexible server solution that can be readily adapted to a variety of different software platforms designed for a variety of different computing architectures. For example, potentially any streaming server **202** could be adapted to stream software applications designed for a variety of different computing architectures by locally connecting a terminal system which runs the software.

However, executing the desired software on a terminal system **212** at the server side may present certain challenges within a cloud computing framework. For example, the terminal system **212** may be able to execute any selected program for a current streaming session from among a variety of different programs, similar to the process described above with respect to the example setup depicted in FIG. **1**. Likewise, any other data reads/writes to a storage device required by the program **208** running on the terminal system **212** could be freed from constraints of local storage by the terminal system. Storage data requests are typically accomplished through a local mass storage connection in the terminal system **212**, such as a serial Advanced Technology Attachment (SATA) connection or other mass storage interface. It would be desirable to provide a virtualized storage solution for the terminal system without having to make modifications to its hardware architecture or consuming additional computing resources with the storage virtualization process.

To better appreciate the challenge in the context of a cloud gaming environment, the following example illustrates how the server set up in FIG. **2** might be utilized in a particular cloud gaming implementation. Specifically, a user at the client device **204** may be assigned to a “cloud terminal system” **212** for a gaming session. At the time the user starts playing, the game **208** the

user wants to play should be available on the cloud terminal system **212**. Similarly, once the gaming session is over and the user is done, the terminal system **212** should be restored to a clean state for a next session by a different user, which may involve a different game.

Unfortunately, virtualizing game storage in this manner using storage virtualization software within the terminal system **212** would consume valuable computational resources (e.g. CPU and memory) and may produce compatibility issues with the game. Moreover, in an application involving a heavy computational load, such as many modern video games, which incorporate advanced graphics and other resources intensive tasks, consuming additional resources from the terminal system for storage virtualization may not be a viable solution.

An illustrative example of a storage virtualization technique according to various implementations of the present disclosure is depicted in in FIG. **3**. In the example virtualization technique of FIG. **3**, a mass storage device is emulated using a “host system” having a specialized piece of hardware, so that a mass storage device may be emulated for a terminal system connected thereto. The example technique of FIG. **3** may be implemented in a cloud streaming application similar to that depicted in FIG. **2**, in which the terminal system may be a specialized system running the desired application, and the host system for the mass storage device emulation is a streaming server. However, it is understood that storage virtualization in accordance with principles herein may be implemented in other contexts, including cloud computing applications beyond game streaming, and local implementations where storage virtualization for a computing device is desired, particularly where it is desired without making modifications to the computing device.

Turning to the example of FIG. **3** in more detail, virtualized storage is provided for the terminal system **312** using a specialized interface device **320** coupled to a host system **302** for the interface device, which is in this case the streaming server. The interface device **320** together with a storage process **322** embodied in the host system **302** may perform various mass storage device data requests received from the terminal system **312**.

In the example depicted in FIG. **3**, the interface device **320** may be a piece of hardware which looks like an actual mass storage device to the terminal system, but under the hood its data may come from a storage server (not pictured). A normal mass storage device typically uses a SATA

connection or other standard mass storage interface. In order emulate the mass storage device and receive the data requests, the interface device **320** of FIG. **3** may include a similar SATA connection or other mass storage communication interface compatible with the mass storage communication interface of terminal system, along with additional components such as a  
5 specialized processing unit, a memory, and other components.

In various implementations, in order to facilitate easy integration of the storage virtualization functionality into computing system, such as a streaming server or other “host” system, the interface device **320** may be implemented as an expansion card, such as a PCI-Express add-on card, which may be connected to the host system **302** through an expansion interface of the host  
10 system. During operation, the interface card may receive disk read and/or write requests from the terminal system through its mass storage connection (e.g. through a SATA port). These data access requests may be ultimately processed by a storage application **322** on the host system **302**, and the storage application may access data from a file, storage server, or something else.

Although not necessary for the storage virtualization functionality, the interface device **320** may  
15 also optionally include additional components to provide other functionality that may be especially useful in a cloud streaming context. For example, in order to support streaming functionality in the host system, the interface device **320** may optionally include one or more additional interfaces for audio and/or video data transfer, as well as corresponding functionality to support audio/video (A/V) streaming via a streaming process of the host system **302** in order  
20 to support cloud gaming or other cloud computing implementations. The optional A/V interface of the interface device **320** may be compatible with the terminal system **312** and may configured according to any of a variety of A/V standards, such as HDMI, composite video, component video, S-video, or another A/V communication interface.

FIGS. **4A** and **4B** provide a more detailed overview of an illustrative hardware / software  
25 architecture of an interface device and certain computing devices according to various implementations of the present disclosure. FIG. **4A** illustrates a computing architecture for an example terminal system according to various aspects of the present disclosure. FIG. **4B** illustrates a computing architecture for an example host system having an example interface device according to various implementations of the present disclosure. In some implementations,

the terminal system may be an existing embedded system, video game console, or other computing device having a specialized architecture. In some implementations, the host system may be a streaming server configured to stream an application running on the terminal system to a client device, e.g. as depicted in FIG. 2.

- 5 The example terminal system **412** may include a mass storage host controller **424** (also sometimes known as a host adapter or host bus adapter), which may be, e.g., a SATA host controller (AHCI), USB host controller, or a host controller for another mass storage communication interface standard. The terminal system **412** may also include a corresponding mass storage device driver **426** that is compatible with the mass storage device host controller
- 10 **424** to facilitate communication via the storage device interface. In order to connect to a mass storage device, normally the host controller **424** would be connected to the mass storage device using a corresponding connector (such as a SATA cable, USB cable, or other connector depending on the mass storage device interface of the terminal system). Specifically, the host controller **424** would be connected to a mass storage device controller (also sometimes known as
- 15 a “disk controller”, although the mass storage device may be a solid state drive or other storage device that does not necessarily use a disk) of the particular mass storage device to which it is connected. In the example system depicted in FIGS. **4A-4B**, the host controller **424** is instead connected to the interface device **420**, which together with its host system **402** emulates a mass storage device for the terminal system **412**.
- 20 The interface device **420** may include a communication interface for mass storage device connections that is compatible with the storage communication interface used by the terminal system **412**. The communication interface of the interface device may include one or more mass storage ports **414**, such as one or more SATA ports, USB ports, or other compatible mass storage interface hardware. The interface device **402** may include a specialized processing unit
- 25 operatively coupled to the mass storage hardware interface, with the specialized processing unit having logic **428** that is dedicated to performing functions associated with mass storage device emulation, and optionally other functions, for signals received through the storage connector from the terminal system **402**. The logic **428** may also support communication with the host system **402** through an additional communication interface, which may communicate with a
- 30 peripheral bus of the host system **402** in order to interface with a storage process embodied in the

host system. By way of example, and not by way of limitation, the interface device **420** may be an add-on card which communicates with the host system **402** process bus through an expansion interface, such as peripheral component interconnect (PCI), PCI-eXtended (PCI-X), PCI-Express (PCIe), or another interface which facilitates communication with the host system **402** e.g., via a peripheral bus. The host system may include an interface device driver **426** to support the exchange of signals via the interface device **420**.

In certain implementations, the specialized processing unit may be a field programmable gate array (FPGA), an application specific integrated circuit (ASIC), or another specialized processing unit having dedicated units of logic configured in accordance with principles described herein. The logic units **428** of the specialized processing unit may also include dedicated logic to support various functions for cloud streaming in addition to storage virtualization, such as audio/video capture of the output from an application **408** running on the terminal system **402**, which may facilitate streaming of the A/V output over the network by the host system **402** in implementations involving cloud gaming. The interface card **402** may include an A/V communication interface, such as one or more HDMI ports or other A/V ports, as well as a compatible A/V receiver **430** to receive the A/V output of the application running on the terminal system through a corresponding A/V connector.

In the example depicted in FIGS. **4A-4B**, a mass storage device controller embodied in the logic **428** of the specialized processing unit is controlled by the interface device driver **426** in the host system **402**, and a storage process **432** embodied in the host system **402**, e.g. a software application running on a central processing unit **404**, may provide data transport services for the interface device driver **426**. For example if the terminal system **412** requests to read data from its disk drive or other mass storage device connected to via the mass storage connector, this request may make it through the mass storage host controller **424**, the mass storage device controller contained in the logic **428**, the interface device driver **426**, to the storage process **432** embodied in the host system. The storage process **432** may then retrieve the data, which it could get from a variety of sources. For example, the storage process **432** may retrieve the data from a disk drive or other mass storage device locally connected to the host machine, or it may retrieve it from a memory drive or a storage server over a network. Whichever method the storage process uses is transparent to the terminal system **402**.

In certain implementations, the host system **402** may optionally be a streaming server adapted to transmit over a network a stream of output from the application **408** running on the terminal system **412**. For example, the host system **402** may include an A/V process **434**, which may be an application running on the host system **402**. The A/V capture process may coordinate with  
5 dedicated logic **428** in the specialized processing unit for capturing an A/V stream sent through an A/V connector connected to the interface card **420**. The streaming server **402** may include an Ethernet adapter or other network adapter **436**, and a corresponding Ethernet driver or other network driver **438** for the operating system of the host **402**, with a compatible network library **439** providing protocol support for the network communication. The host system may also  
10 include system memory **440**, controlled by a corresponding memory driver **442** (e.g. tmpfs) and supported by a file system library **444**. A streaming server process **446** may be run on the host system **402** to perform functions associated with provide a real time stream to a client device connected over a network (not pictured in FIGS. **4A-4B**).

In some situations, data storage may be implemented using a combination of a remote storage  
15 server **448** (shown in phantom in FIG. 4A), the memory **440** and the memory driver **442**. By way of example, and not by way of limitation, in some implementations data may only be read from the storage server **448** and written to the memory **440** by the memory driver **442**.

The terminal system **412** may include various other components to support the application **408**, which may be, e.g., video game software designed for an existing embedded platform. The  
20 terminal system **412** may include a file system layer **427** to access storage, as well various components to support graphics rendering for the application **408**. In addition to a central processing unit (CPU), the terminal system **412** may include a graphics processing unit (GPU) **450**, which may be operated by a graphics driver **452** and supported by a graphics library **454**. An A/V transmitter **456**, such HDMI transmitter, may be included to transmit the application  
25 audio and/or video output to the host interface device in the host system through a compatible A/V connector.

Turning now to FIG. 5, a schematic diagram of an example interface device **520** that may be implemented on the interface card **420**, some components on it, and internals of an example specialized processing unit **560** is depicted in accordance with various implementations of the

present disclosure. By way of example, and not by way of limitation, the interface device **520** may be configured as an add-on card having components attached to a printed circuit board (PCB), and the interface card **520** may interface with a peripheral bus of a host system through a host hardware interface **562**, such as a peripheral expansion port or other expansion communication interface which allows communication the peripheral bus of a host system when connected.

The example specialized processing unit **560** may include various blocks of logic dedicated to specialized functionality in accordance with various aspects of the present disclosure. The specialized processing unit may be implemented, e.g., as an FPGA, ASIC, or similar specialized processing unit. The specialized processing unit **560** may include a host interface block **564** which implements part of a protocol stack for the communication interface between the interface card **520** and a peripheral bus of a host system (not pictured in FIG. 5) for the interface device **560**.

Communication busses like PCI-Express can be thought of as a protocol stack having several layers. Different communication protocols have different layers. Typically there is an ‘application layer’ at the top, then some transport related layers in the middle and some physical layer at the bottom. The host interface block **564** need not implement all layers of such a protocol stack. Instead, the host interface block may take care of the physical layer, which is responsible for putting digital information on a communication link, e.g., through electrical or optical signals. The host interface block may also be responsible for portions or possibly all of the ‘transport layers’ of the protocol stack, but need not be responsible for the application layer.

By way of example, and not by way of limitation, the host interface block **564** may be a hard PCIe block for communication through a PCI-Express connection, and which embeds the protocol stack for a PCIe interface or other interface for accessing a local bus of the host system. The host interface block **564** may be integrated into a memory access interface unit **566** which, together with other logic units of the specialized processing unit **560**, may directly access system memory of a host system through the host hardware interface **562**, e.g., using an interrupt of the request to the host system.

In some implementations, the memory access interface **566** may include components that provide memory access and interrupt functionality. In particular, the host interface block **564** may be configured to provide an connection between an on-chip-interconnect **572** and the host hardware interface **562** in a way that makes any on-chip device accessible from the host system using memory mapped Input/Output (I/O). This functionality would allow the host system to program any device connected to the on-chip-interconnect **572**, such as the mass storage controller **570**, memory controller **576**, or GPIO **582**.

The memory access interface **566** may also include an interrupt connection **565** that allows any connected device, e.g., the mass storage controller **570**, to generate an interrupt upon an event (e.g., a mass storage read request) arrival). It is desirable for the memory access interface to provide this functionality if there can be only one device interfacing with the host hardware interface hardware **562**.

The memory access interface **566** may also (optionally) include a direct memory access (DMA) engine **567**. As used herein, and as is generally understood by those skilled in the art, the term direct memory access (DMA) refers to a feature that allows certain hardware subsystems within a computer to access system memory independently of the computer's central processing unit (CPU). The DMA engine **567** may implement data move operations between the host interface block **564** and the host hardware interface **562**. In some implementations, the memory access interface unit **566** may implement portions of a protocol stack (e.g., PCI Express) not provided by the host interface block **564**, such as connecting the host interface block **564** to the on-chip-interconnect **572**.

For purposes of functionality for mass storage device virtualization, the interface device **520** may include one or more mass storage device communication interfaces **568**, which may be implemented in the form of one or more SATA ports **569** and/or connectors, or other mass storage device communication interfaces, and which may be attached to a circuit board of the interface device **520**. By way of example, and not by way of limitation, the interface card **520** may contain two SATA ports to facilitate connection to two terminal systems, although it is noted that the interface device may alternatively contain a different number of mass storage device connectors so that a single interface device **520** may service a different number of

terminal systems. For each of the mass storage device connectors **568**, there may be a corresponding mass storage device controller **570** embodied in the specialized processing unit **560** that is compatible with the particular mass storage communication interface (e.g. SATA).

5 The one or more mass storage device controllers **570** may be connected to other logic units of the specialized processing unit **560** through the on-chip interconnect **572**, which may provide each of the mass storage device controllers **572** access to host system interface components (e.g. PCI-Express). The on-chip interconnect may be configured to a standard on-chip bus architecture configured to connect functional blocks on a specialized processing unit (e.g. an FPGA or ASIC). For example, if the specialized processing unit **560** is an FPGA, the components of the  
10 specialized processing unit may be interconnected using master-slave architecture, e.g., an Advanced Microcontroller Bus Architecture (AMBA), such as AXI4 or AXI4-Lite, or another suitable on-chip bus architecture. AXI4 may be used for large data transport and AXI-Lite may be used for low performance connections or for configuration purposes. The on-chip interconnections of the specialized processing unit logic blocks may be configured according to a  
15 master-slave type configuration as shown in FIG. 5. In the illustrated schematic, “M” and the corresponding bold lines represent represents a master connection, “S” and the corresponding dotted lines represent a slave connection, and “Ctrl” represents control.

The interface device **520** may include one or more memory units **574** which may be controlled by a memory controller **576** provided in the logic of the specialized processing unit **560**. The  
20 memory unit may support data transport between a terminal system connected through the mass storage interface **568** and a host system connected through the host hardware interface **562**, in accordance with data requests issued by the terminal system. For example, the memory unit **574** may be a temporary RAM unit, such as DDR3 RAM, or another volatile memory unit configured to temporarily store data requested by read requests issued by the terminal system, in accordance  
25 with principles described herein. The memory controller **576** may be connected to the on chip bus architecture **572** to perform memory read/write operations according to signals received from other logical units of the specialized processing unit **560**.

During operation, a mass storage host controller of a computing device (not pictured in FIG. 5) connected through the storage interface **560** may request the mass storage device controller **570**

to perform read, write, and other operations. Upon receiving such a request, the mass storage device controller **570** may notify an interface device driver and storage process on the host system (not shown in FIG. **5**) using an interrupt of the request, through the host hardware interface **562**, e.g. PCI-Express. The software components on the host system may then make  
5 decisions on handling the request.

For example, in the case of a read request from the terminal system, the storage application in the host system (not pictured) may retrieve this data from some storage source and store it in the interface device memory unit **574** across the host hardware interface **562**. When done, the storage application on the host system may program the mass storage device controller **570** on  
10 the specialized processing unit **560** to complete the read request by reading from the interface device memory **574**. By way of example, and not by way of limitation, suppose the mass storage host controller **424** may be a SATA controller. For a SATA transaction (read/write and there are others), the mass storage host controller **424** on the streaming server **402** raises an interrupt. The mass storage host driver **426** on the streaming server **402** checks the state of the mass storage  
15 host controller **424**. If the request is determined to be a 'write request', then the mass storage host controller **424** is programmed to write its data to the memory **574** (e.g., DDR3 memory) on the interface device **520**. The mass storage host controller **424** may have its own built-in DMA engine **567** for this purpose. Another interrupt is raised when the data has been written to the memory **574**. The mass storage host driver **426** then issues a DMA request using the DMA  
20 Engine **567** to copy from memory **574** to host system memory **440**.

A buffer in the host system memory **440** may be directly shared with the storage process. When the DMA transaction to host system memory **440** is done, the storage process is notified of the write request and it will finish the write transaction. It may write data to a disk, to a network storage device or to wherever it pleases to store the data.

25 It is noted that a memory **574** on the interface device **520** might not be needed for the read or write operations. It is possible for the mass host storage controller **424** to directly write to host system memory **440**. For example, the memory access interface unit **566** could map PCI-Express to AXI4 and also AXI4 to PCI-Express. This feature would allow any AXI4 Master device to write to host system memory.

The interface device may optionally include various other components which provide additional functionality for streaming applications run on a terminal system, such as cloud gaming streaming. For example, the specialized processing unit **560** may also include one or more audio and/or video capture blocks **578** for capturing an A/V stream from one or more terminal systems.

5 Each of the one or more A/V capture logic units **578** may be operatively coupled to a corresponding A/V receiver **530**, each of which may in turn be connected to a suitable A/V hardware interface **580**, such as an HDMI port or other A/V connection port as shown in FIG. **5**. A/V output from the terminal system may be connected to the A/V receiver **530** through the A/V interface **580** using a compatible A/V connector. The A/V capture unit **578** may communicate  
10 with the interface device driver and A/V process on the host system through the host hardware interface **562**, which may be connected to a host system bus (e.g., a peripheral bus), and the host system may then deliver the A/V stream to a client device over a network.

The interface device **582** may also include one or more general purpose input/output (GPIO) blocks to support additional functionality. By way of example, and not by way of limitation, each  
15 of the GPIO blocks may be connected to a corresponding one of the terminal system to provide additional functionality, such as power control of the terminal systems and other functionality.

As noted above, the specialized processing unit **560** may be implemented, e.g., as an FPGA, ASIC, or other integrated circuit having blocks dedicated to certain functionality, such as a mass storage device controller, memory controller, DMA engine, and the like, in accordance with  
20 various aspects of the present disclosure. In certain implementations of the present disclosure, one or more of these units may be provided as reusable units of logic or other chip design commonly referred to in the art as IP blocks or IP cores.

FIG. **6** depicts an example flow chart for emulating a mass storage device for a terminal system **612** in accordance with various aspects of the present disclosure. In the example depicted in FIG.  
25 **6**, an interface device **620** and host system **602** operatively coupled to each other are configured to emulate a mass storage device for the terminal system **612**. While the illustrative example of FIG. **6** depicts emulation for a read request issued by the terminal system for purposes of explanation, it is noted that implementations of the present disclosure may be configured to

complete write requests and other mass storage data requests issued by the terminal system using the illustrated components.

As shown in FIG. 6, a mass storage device may be emulated for a terminal system **612** or any other computing system, which may have a central processing unit (CPU) **681** and a memory unit **682** configured to run various applications (not pictured). The terminal system **612** may make various read, write, and other data requests through a mass storage device connector using a mass storage device controller **624**. Ordinarily, these requests would be completed by a local mass storage device having a corresponding device controller. In the example of FIG. 6, these data requests may be met by an interface device **620** in communication with a host computer system **602**, which may access some other storage source **699** to handle the various data requests.

In the illustration of FIG. 6, a read request **684** is issued by the mass storage host controller **624** of the terminal system **612**, which is received at a mass storage device controller **670** of the specialized processing unit **660**. The mass storage device controller **670** notify the storage process **632** embodied in the host system **602** using an interrupt of the request **684**, e.g. via a DMA engine **666** as shown in FIG. 6. The storage process may be an application embodied in the system memory **640** of the host system **602**, which may be executed by a central processing unit **686** of the host system **602** to make decisions on handling the request **684**.

As shown in FIG. 6, the storage process may retrieve the requested data **687** from some other storage source **699**, such as a storage server. The storage process may store the retrieved data **689** in the interface device memory unit **674**, e.g., through an expansion interface (not pictured), such as PCI-e, or other communication interface with a peripheral bus of the host system **602**. It is noted that, instead of being implemented as an add-on card or other peripheral component which communicates with the peripheral bus through an expansion interface, in some implementations the interface device **620** may be implemented as with the specialized processing unit **660** connected directly to the host system peripheral bus and provided as an integral part of the host system **602**.

In alternative implementations, the specialized processing unit **660**, e.g., FPGA, ASIC, or equivalent circuitry, could be placed directly on the motherboard of the streaming server **402** and

connected to a peripheral bus on the streaming server. The specialized processing unit could be connected over PCI-Express.

Once the data retrieved by the storage process is stored in the interface memory **674**, the storage process may program the mass storage device controller **670** to complete the request by  
5 accessing the stored data **689** in the interface memory and sending it to the terminal system **612**.

During a read operation the mass storage device controller **670** may raise an interrupt (which could be transferred using an interrupt functionality built into the DMA engine **666**) to a storage driver in the operating system of the host system **602**. If the request is determined to be a read request **684**, the storage driver asks the storage process **632** to 'read the data' which is stored in  
10 the host memory **640** (which may be shared with the storage driver). Once all data has been read, the DMA engine **666** may be programmed to transfer the data **689** from the host memory **640** to the interface memory **674**.

### **Conclusion**

It is noted that aspects of the present disclosure have been described with reference to cloud gaming as a particular example in which a need for mass storage virtualization in accordance  
15 with various implementations of the present disclosure may be desirable, but implementations of the present disclosure are not so limited. Implementations of the present disclosure may provide mass storage device for emulation for any computing device, which may not involve cloud computing or cloud gaming.

It is noted that implementations of the present disclosure may utilize any of a variety of mass storage communication protocols, including SATA, USB, IDE, RS232, Firewire, Thunderbolt,  
20 and other mass storage communication interface standards.

It is noted that implementations of the present disclosure may utilize any of a variety of audio and/or video standards, including HDMI, S-Video, component video, composite video, Display  
25 Port, VGA, DVI, and other A/V communication interface standards.

While the above is a complete description of the preferred embodiment of the present invention, it is possible to use various alternatives, modifications and equivalents. Therefore, the scope of

the present invention should be determined not with reference to the above description but should, instead, be determined with reference to the appended claims, along with their full scope of equivalents. Any feature described herein, whether preferred or not, may be combined with any other feature described herein, whether preferred or not. In the claims that follow, the

5 indefinite article “a”, or “an” refers to a quantity of one or more of the item following the article, except where expressly stated otherwise. The appended claims are not to be interpreted as including means-plus-function limitations, unless such a limitation is explicitly recited in a given claim using the phrase “means for.”

What is claimed is:

1. An interface device comprising:

a first communication interface, the first communication interface being a mass storage interface;

a second communication interface, the second communication interface being an expansion interface; and

a specialized processing unit operatively coupled to said communication interfaces, the specialized processing unit having a mass storage device controller connected to the first communication interface and a memory access interface connected to the second communication interface,

wherein the mass storage device controller is configured to receive a mass storage data request through the first communication interface,

wherein the memory access interface is configured to notify a host system of the mass storage data request through the second communication interface.

2. The interface device of claim 1,

further comprising a memory unit,

wherein the specialized processing unit includes a memory controller connected to the memory unit.

3. The interface device of claim 1,

wherein the memory access interface is configured to notify the host system of the data request using an interrupt of the data request.

4. The interface device of claim 1, wherein the memory access interface includes a DMA engine configured to implement data move operations.

5. The interface device of claim 1,

wherein the specialized processing unit is a field programmable gate array (FPGA).

- 1       6. The interface device of claim 1,  
2               wherein the specialized processing unit is an application specific integrated circuit  
3       (ASIC).
- 1       7. The interface device of claim 1, further comprising:  
2               a printed circuit board (PCB),  
3               wherein the first communication interface includes one or more mass storage  
4       ports fixed to the PCB.
- 1       8. The interface device of claim 1, further comprising:  
2               a printed circuit board (PCB),  
3               wherein the second communication interface includes one or more expansion  
4       connectors fixed to the PCB.
- 1       9. The interface device of claim 1,  
2               wherein the first communication interface includes one or more mass storage  
3       ports,  
4               wherein the specialized processing unit has a mass storage device controller for  
5       each said mass storage port.
- 1       10. The interface device of claim 1,  
2               wherein the first communication interface includes a plurality of mass storage  
3       ports,  
4               wherein the specialized processing unit has plurality of mass storage device  
5       controllers,  
6               wherein each said mass storage device controller is connected to a corresponding  
7       one of the mass storage ports.
- 8       11. The interface device of claim 1,  
9               wherein the first communication interface includes one or more serial Advanced  
10       Technology Attachment (SATA) ports,  
11               wherein the specialized processing unit has a SATA device controller for each  
12       said SATA port.

1 12. A system comprising:

2 a central processing unit;

3 a first memory unit operatively coupled to the central processing unit;

4 a storage process embodied in the first memory unit and executable by the central  
5 processing unit;

6 a first communication interface, the first communication interface being a mass  
7 storage interface;

8 a specialized processing unit operatively coupled to the first communication  
9 interface, the specialized processing unit having a mass storage device controller  
10 connected to the first communication interface and a memory access interface;

11 wherein the mass storage device controller is configured to receive a mass storage  
12 data request from through the first communication interface,

13 wherein the memory access interface is configured to notify the storage process of  
14 the mass storage data request,

15 wherein the storage process is configured to process the data request by accessing  
16 a storage system.

1 13. The system of claim 12, further comprising:

2 a second memory unit,

3 wherein the specialized processing unit has a memory controller connected to the  
4 second memory unit.

1 14. The system of claim 12, further comprising:

2 a second memory unit,

3 wherein the specialized processing unit has a memory controller connected to the  
4 second memory unit,

5 wherein the data request is a read request,

6 wherein the storage process is configured to retrieve data corresponding to the  
7 read request from the storage system,

8 wherein the storage process is configured to store the retrieved data on the second  
9 memory unit,

10 wherein the storage process is configured to program the mass storage device

11 controller to complete the read request by programming the mass storage device  
12 controller to read from the second memory unit.

1 15. The system of claim 12, further comprising:

2 a second communication interface, the second communication interface being an  
3 expansion interface,

4 wherein the memory access interface is connected to the second communication  
5 interface,

6 wherein the second communication interface is connected to the processor bus.

1 16. The system of claim 12, further comprising:

2 a printed circuit board (PCB), and

3 a second communication interface, the second communication interface being an  
4 expansion interface,

5 wherein the DMA engine is connected to the second communication interface,

6 wherein the second communication interface is connected to the processor bus,

7 wherein the first communication interface, the second communication interface,

8 and the specialized processing unit are fixed to the PCB.

1 17. The system of claim 12,

2 wherein the DMA engine is directly connected to the processor bus.

1 18. The system of claim 12,

2 wherein the memory access interface is configured to notify the storage process of  
3 the data request using an interrupt of the data request.

1 19. The system of claim 12, wherein the memory access interface includes a DMA engine  
2 configured to implement data move operations.

1 20. The system of claim 12,

2 wherein the first communication interface includes one or more mass storage  
3 ports,

4 wherein the specialized processing unit has a mass storage device controller for  
5 each said mass storage port.

1 21. The system of claim 12,

2 wherein the specialized processing unit is a field programmable gate array  
3 (FPGA).

1 22. The system of claim 12,

2 wherein the specialized processing unit is an application specific integrated circuit  
3 (ASIC).

1 23. The system of claim 12, further comprising:

2 a terminal system operatively coupled to the first communication interface  
3 through a mass storage connector, the terminal system having a mass storage host  
4 controller configured to issue the mass storage data request.

1 24. The system of claim 12, further comprising:

2 a plurality of terminal systems,  
3 wherein the specialized processing unit has a plurality of said mass storage device  
4 controllers,  
5 wherein the first communication interface has a plurality of mass storage ports,  
6 wherein each said terminal system is connected to a corresponding one of the  
7 mass storage ports with a mass storage connector,  
8 wherein each said mass storage port is connected to a corresponding one of the  
9 mass storage device controllers.

1 25. A method comprising:

2 a) receiving, from a first computing device, a mass storage data request at a mass storage  
3 device controller;  
4 b) notifying, with a memory access interface connected to the mass storage device  
5 controller, a storage process embodied in a second computing device of the data request;  
6 c) accessing, with the storage process of the second computing device, a storage system  
7 having data corresponding to the data request;  
8 d) storing, with the storage process of the second computing device, the data  
9 corresponding to the data request on a first memory unit;  
10 e) programming, with the storage process of the second computing device, the mass

- 11 storage device controller to access the data from the first memory unit in order to  
12 complete the data request.

1/5

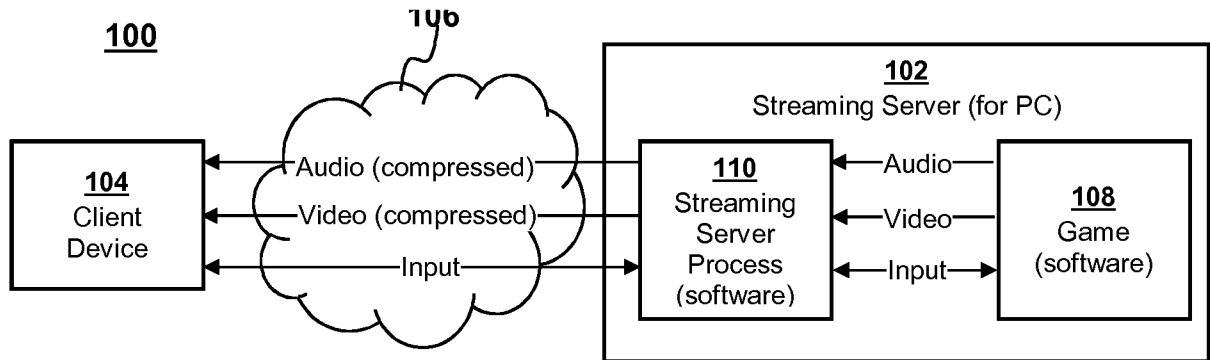


FIG. 1

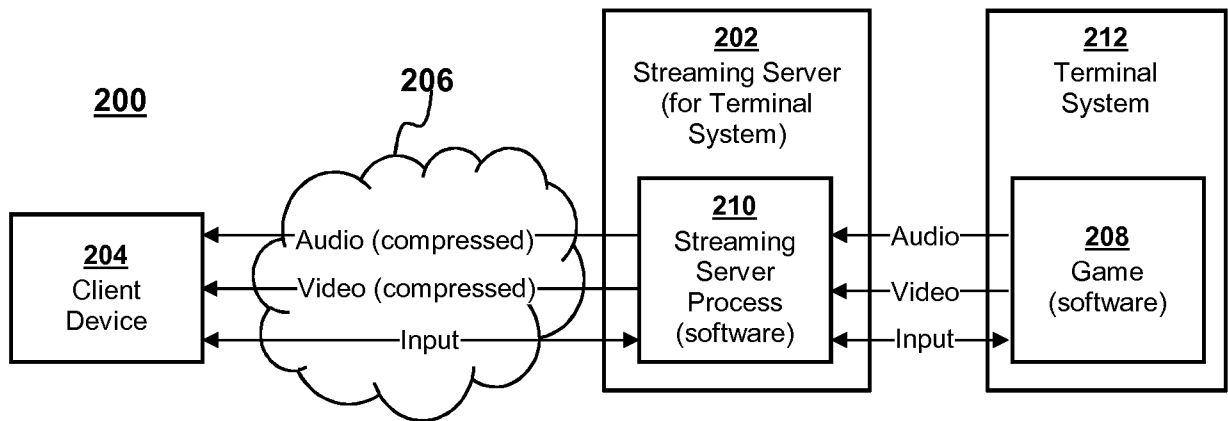


FIG. 2

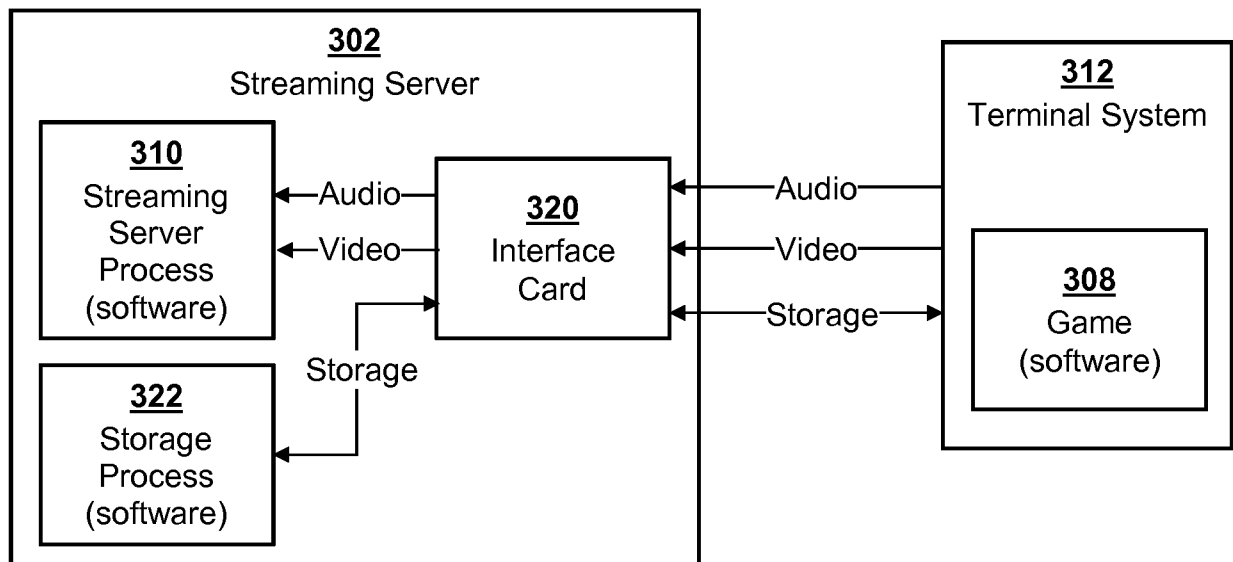


FIG. 3

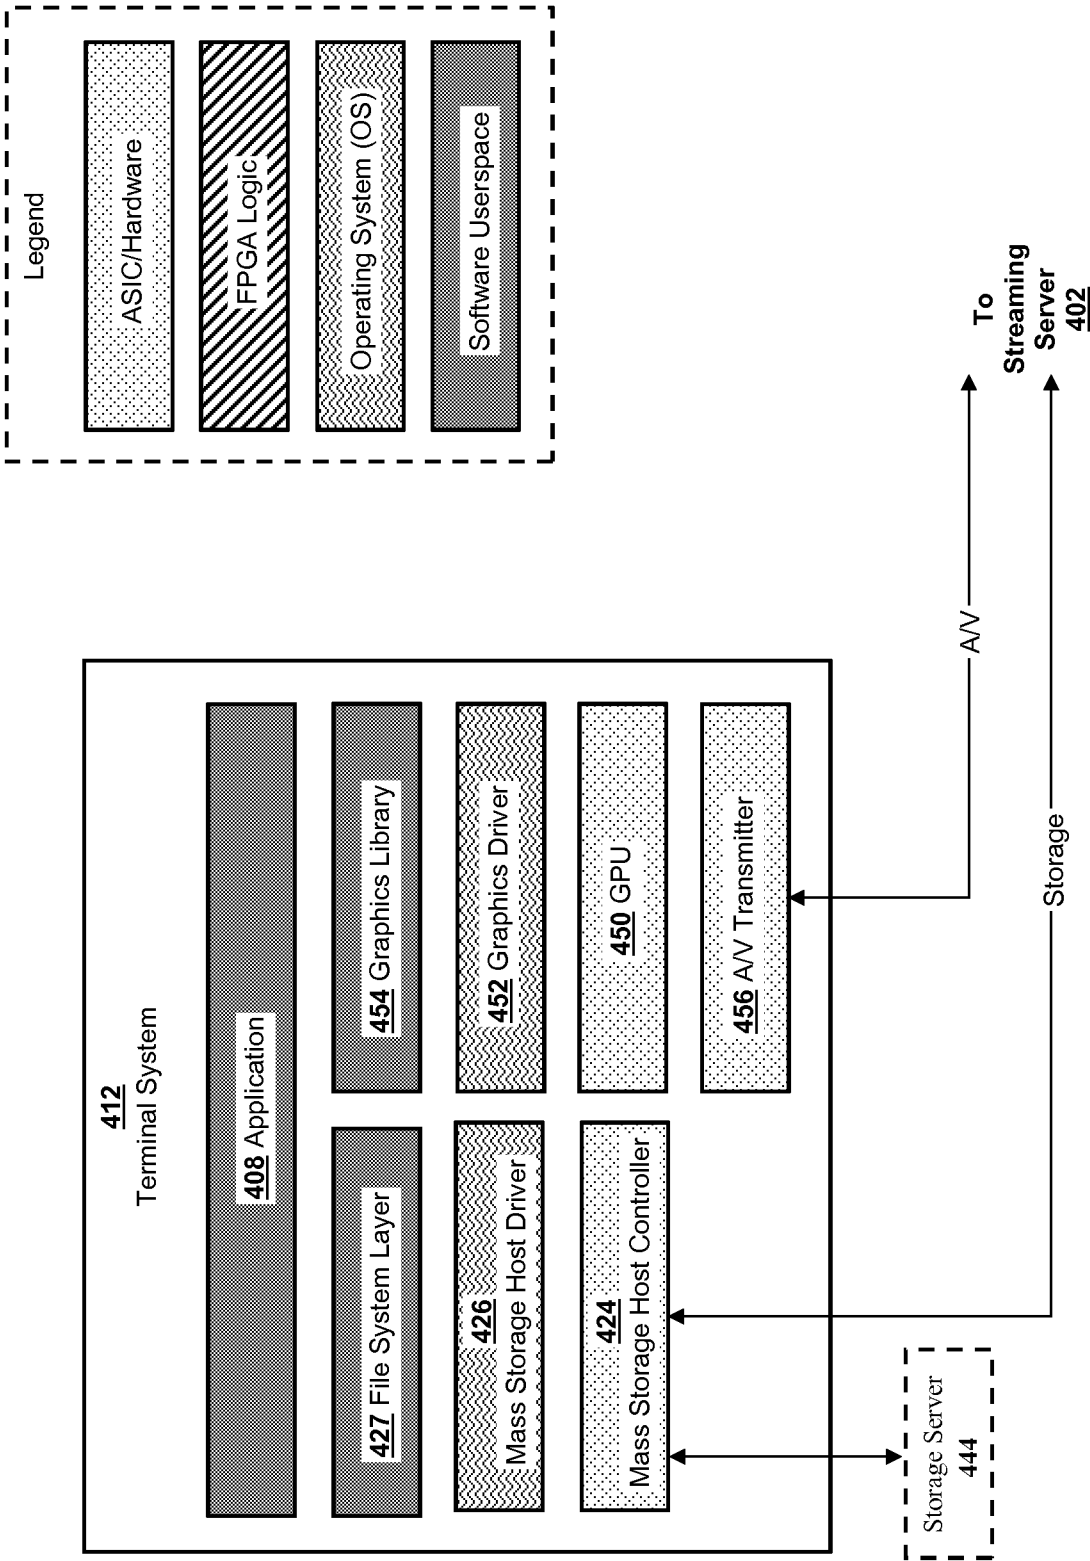


FIG. 4A

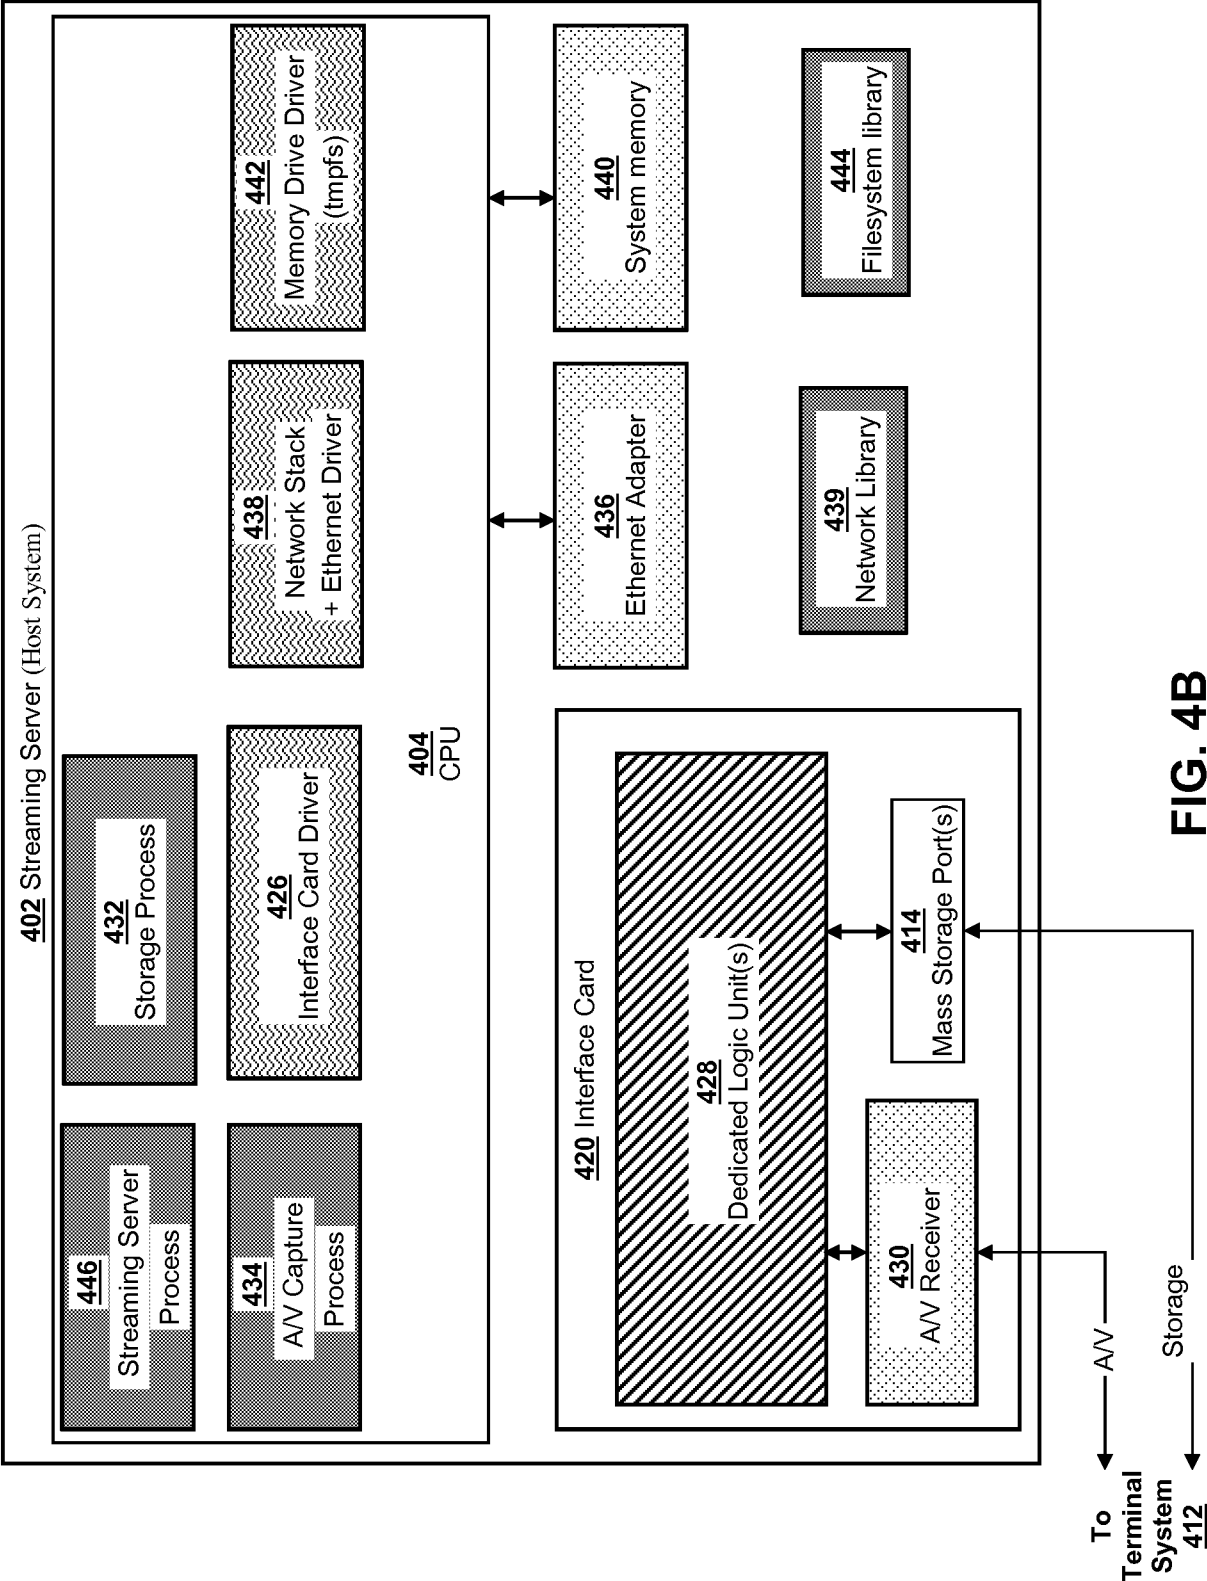


FIG. 4B

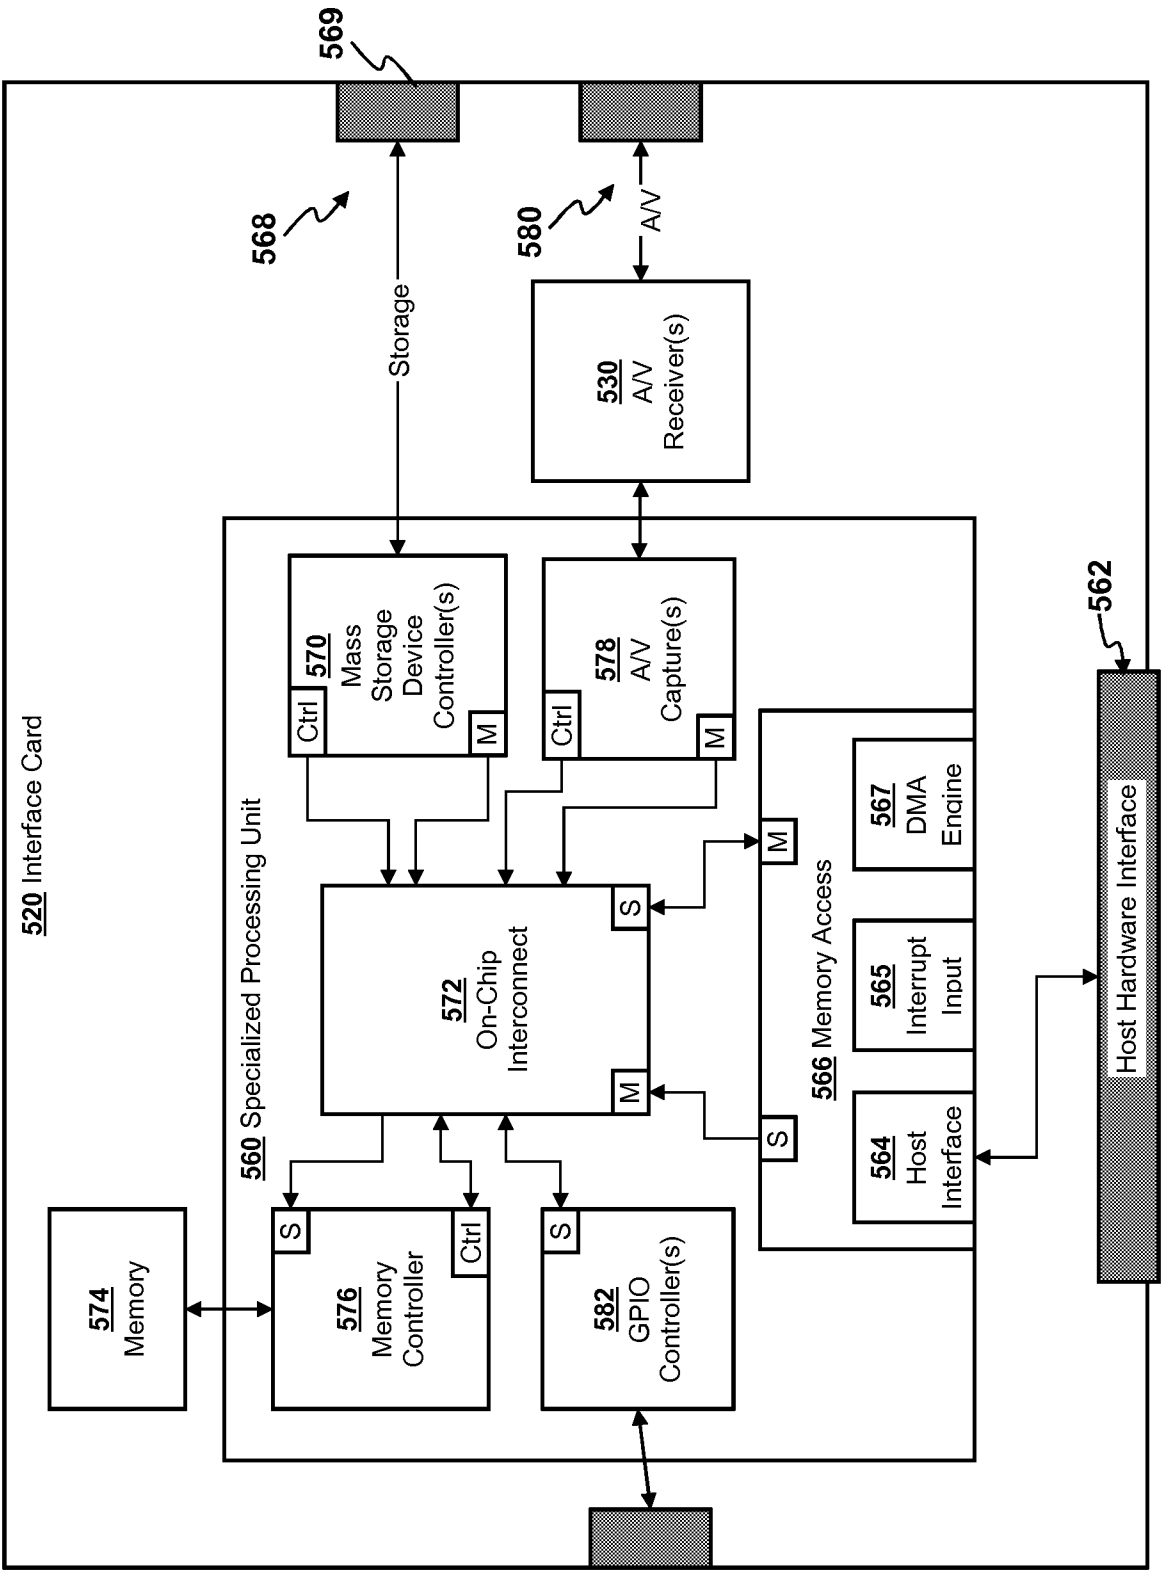


FIG. 5

5/5

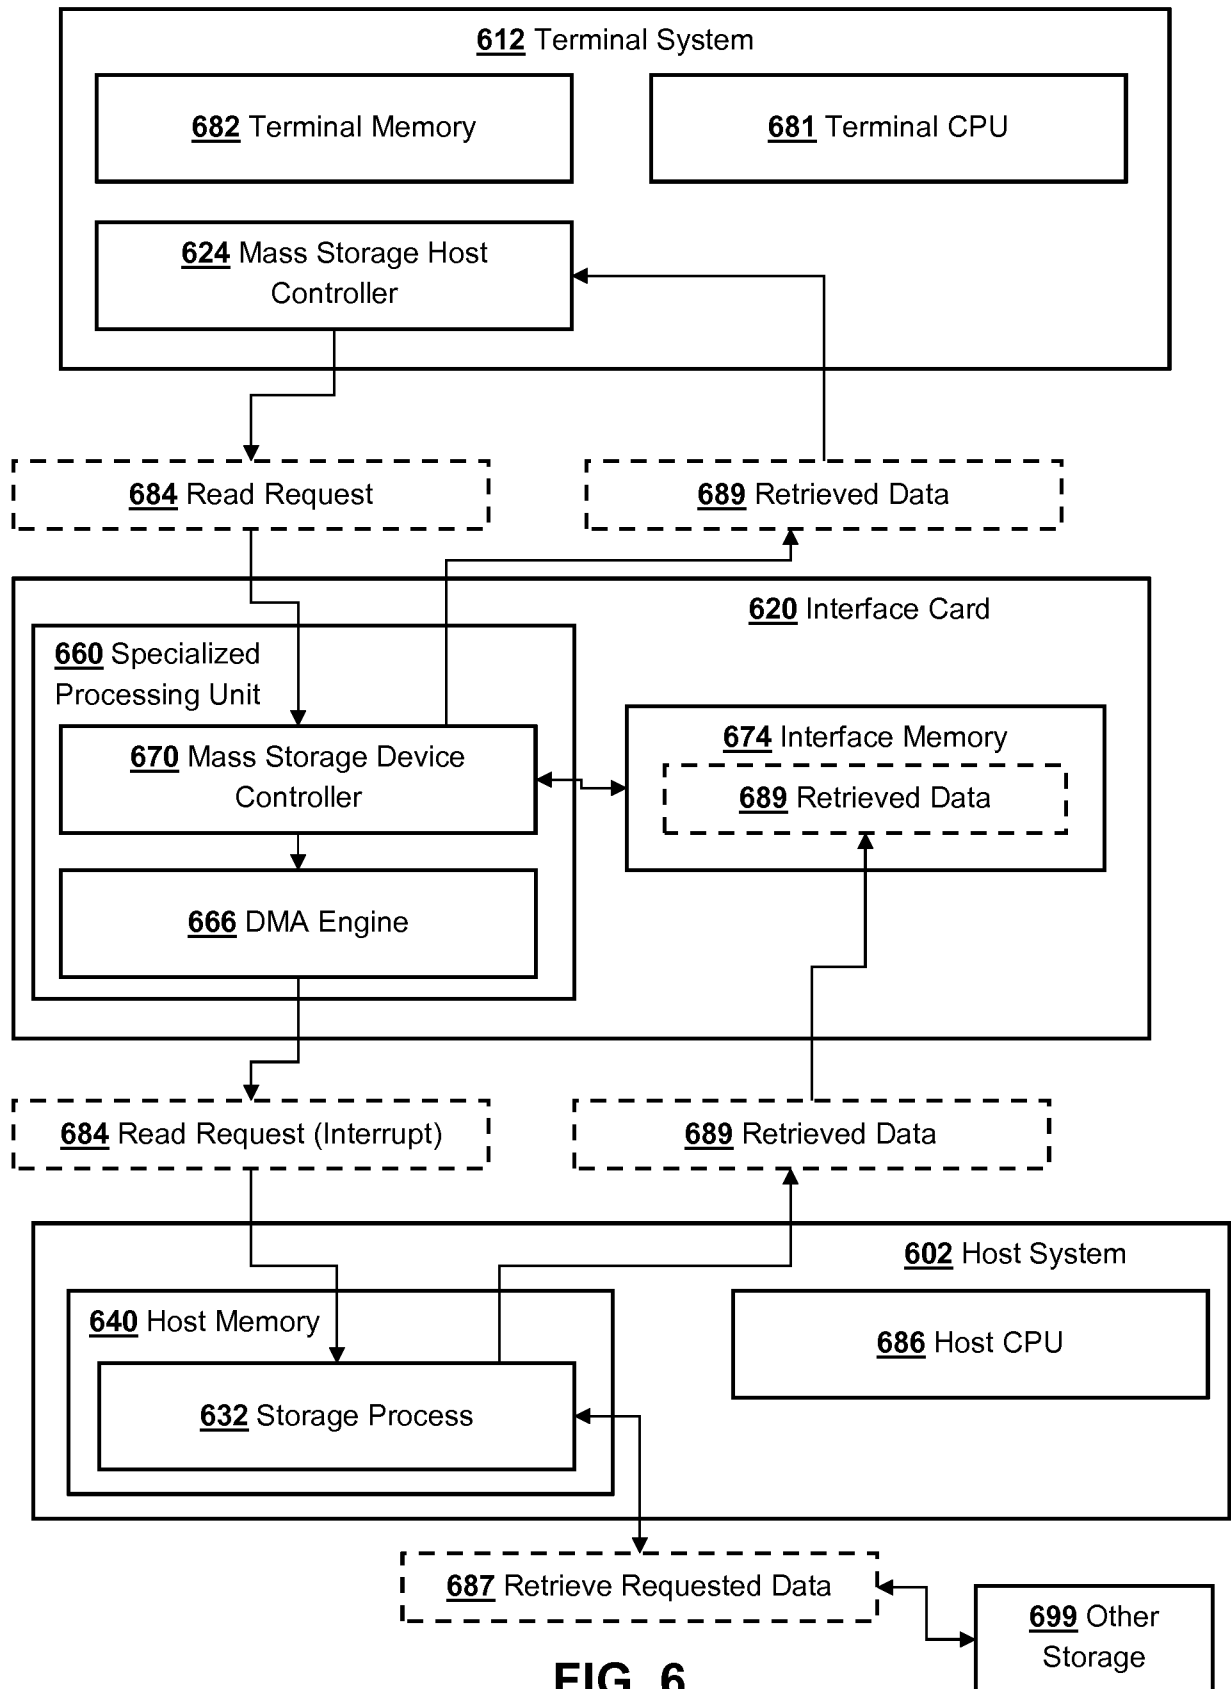


FIG. 6

## INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 2014/061633

A. CLASSIFICATION OF SUBJECT MATTER														
<p style="text-align: center;"><b>G06F 13/00 (2006.01)</b>  <b>H04N 21/433 (2011.01)</b></p> <p>According to International Patent Classification (IPC) or to both national classification and IPC</p>														
B. FIELDS SEARCHED														
Minimum documentation searched (classification system followed by classification symbols)														
G06F 13/00,13/14,13/20,13/28,13/32, 12/00,12/16, H04N 21/40,21/43,21/433,21/47,21/472														
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)														
PatSearch (RUPTO internal), USPTO, PAJ, K-PION, Esp@cenet, Information Retrieval System of FIPS														
C. DOCUMENTS CONSIDERED TO BE RELEVANT														
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.												
Y	US 2011/0271067 A1 (PIXEL8 NETWORKS, INC.) 03.11.2011, paragraphs [0005], [0020], [0021], [0030]-[0033], [0041], [0043], [0045]-[0051], [0053], [0054], claim 9	1-25												
Y	US 2010/0280817 A1 (SPANSION LLC) 04.11.2010, paragraphs [0047]-[0051]	1-25												
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input type="checkbox"/> See patent family annex.														
<table border="0"> <tr> <td>* Special categories of cited documents:</td> <td>“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</td> </tr> <tr> <td>“A” document defining the general state of the art which is not considered to be of particular relevance</td> <td>“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</td> </tr> <tr> <td>“E” earlier document but published on or after the international filing date</td> <td>“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</td> </tr> <tr> <td>“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)</td> <td>“&amp;” document member of the same patent family</td> </tr> <tr> <td>“O” document referring to an oral disclosure, use, exhibition or other means</td> <td></td> </tr> <tr> <td>“P” document published prior to the international filing date but later than the priority date claimed</td> <td></td> </tr> </table>			* Special categories of cited documents:	“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	“A” document defining the general state of the art which is not considered to be of particular relevance	“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	“E” earlier document but published on or after the international filing date	“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	“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)	“&” document member of the same patent family	“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	
* Special categories of cited documents:	“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													
“A” document defining the general state of the art which is not considered to be of particular relevance	“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													
“E” earlier document but published on or after the international filing date	“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													
“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)	“&” document member of the same patent family													
“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														
Date of the actual completion of the international search		Date of mailing of the international search report												
16 March 2015 (16.03.2015)		26 March 2015 (26.03.2015)												
Name and mailing address of the ISA/RU: Federal Institute of Industrial Property, Berezhkovskaya nab., 30-1, Moscow, G-59, GSP-3, Russia, 125993 Facsimile No: (8-495) 531-63-18, (8-499) 243-33-37		Authorized officer  I. Kryazhev  Telephone No. (499) 240-25-91												