Search report

Reference number: E.g Reference

Version 1

E.g My Idea

A solid-state drive (SSD) is a solid-state storage device that uses integrated circuit assemblies as memory to store data persistently, typically using flash memory. It is also sometimes called a solid-state device or a solid-state disk,[1] although SSDs lack the physical spinning disks and movable read-write heads used by the conventional electromechanical storage such as hard drives ("HDD") or floppy disks.[2] Compared with the electromechanical drives, SSDs are typically more resistant to physical shock, run silently, and have quicker access time and lower latency.[3] SSDs store data in semiconductor cells. As of 2019, cells can contain between 1 and 4 bits of data. SSD storage devices vary in their properties according to the number of bits stored in each cell, with single bit cells ("SLC") being generally the most reliable, durable, fast, and expensive type, compared with 2 and 3 bit cells ("MLC" and "TLC"), and finally quad bit cells ("QLC") being used for consumer devices that do not require such extreme properties and are the cheapest of the four. In addition, 3D XPoint memory (sold by Intel under the Optane brand), stores data by changing the electrical resistance of cells instead of storing electrical charges in cells, and SSDs made from RAM can be used for high speed, when data persistence after power loss is not required, or may use battery power to retain data when its usual power source is unavailable.[4] Hybrid drives or solid-state hybrid drives (SSHDs), such as Apple's Fusion Drive, combine features of SSDs and HDDs in the same unit using both flash memory and a HDD in order to improve the performance of frequently-accessed data.[5][6][7] While the price of SSDs has continued to decline over time, SSDs are (as of 2018) still more expensive per unit of storage than HDDs and are expected to remain so into the next decade. SSDs based on NAND Flash will slowly leak charge over time if left for long periods without power. This causes worn-out drives (that have exceeded their endurance rating) to start losing data typically after one year (if stored at 30 °C) to two years (at 25 °C) in storage; for new drives it takes longer.[8] Therefore, SSDs are not suitable for archival storage. 3D XPoint is a possible exception to this rule, however it is a relatively new technology with unknown data-retention characteristics. SSDs can use traditional hard disk drive (HDD) interfaces and form factors, or newer interfaces and form factors that exploit specific advantages of the flash memory in SSDs. Traditional interfaces (e.g., SATA and SAS) and standard HDD form factors allow such SSDs to be used as drop-in replacements for HDDs in computers and other devices. Newer form factors such as mSATA, M.2, U.2, and EDSFF (formerly known as Ruler SSD[9])[10] and higher speed interfaces such as NVMe over PCI Express can increase performance over HDD performance.[4]

Version 2

test

A solid-state drive (SSD) is a solid-state storage device that uses integrated circuit assemblies as memory to store data persistently, typically using flash memory. It is also sometimes called a solid-state device or a solid-state disk,[1] although SSDs lack the physical spinning disks and movable read-write heads used by the conventional electromechanical storage such as hard drives ("HDD") or floppy disks.[2] Compared with the electromechanical drives, SSDs are typically more resistant to physical shock, run silently, and have quicker access time and lower latency.[3] SSDs store data in semiconductor cells. As of 2019, cells can contain between 1 and 4 bits of data. SSD storage devices vary in their properties according to the number of bits stored in each cell, with single bit cells ("SLC") being generally the most reliable, durable, fast, and expensive type, compared with 2 and 3 bit cells ("MLC" and "TLC"), and finally quad bit cells ("QLC") being used for consumer devices that do not require such extreme properties and are the cheapest of the four. In addition, 3D XPoint memory (sold by Intel under the Optane brand), stores data by changing the electrical resistance of cells instead of storing electrical charges in cells, and SSDs made from RAM can be used for high speed, when data persistence after power loss is not required, or may use battery power to retain data when its usual power source is unavailable.[4] Hybrid drives or solid-state hybrid drives (SSHDs), such as Apple's Fusion Drive, combine features of SSDs and HDDs in the same unit using both flash memory and a HDD in order to improve the performance of frequently-accessed data.[5][6][7] While the price of SSDs has continued to decline over time, SSDs are (as of 2018) still more expensive per unit of storage than HDDs and are expected to remain so into the next decade. SSDs based on NAND Flash will slowly leak charge over time if left for long periods without power. This causes worn-out drives (that have exceeded their endurance rating) to start losing data typically after one year (if stored at 30 °C) to two years (at 25 °C) in storage; for new drives it takes longer.[8] Therefore, SSDs are not suitable for archival storage. 3D XPoint is a possible exception to this rule, however it is a relatively new technology with unknown data-retention characteristics. SSDs can use traditional hard disk drive (HDD) interfaces and form factors, or newer interfaces and form factors that exploit specific advantages of the flash memory in SSDs. Traditional interfaces (e.g., SATA and SAS) and standard HDD form factors allow such SSDs to be used as drop-in replacements for HDDs in computers and other devices. Newer form factors such as mSATA, M.2, U.2, and EDSFF (formerly known as Ruler SSD[9])[10] and higher speed interfaces such as NVMe over PCI Express can increase performance over HDD performance.[4]

Summary

Patent number	Priority date	Ranking
US2013135816 A1	2011-11-17	Similar
US2016259568 A1	2013-11-26	Similar
US2018165698 A1	2016-12-09	Related
US9720860 B2	2014-06-06	Related
US9898208 B2	2016-04-27	Related
US2018349041 A1	2017-05-31	Related
US9927975 B2	2016-08-03	Background
US9619164 B2	2014-05-06	Background
US2018260135 A1	2017-03-10	Background

US2016259568 A1

Similar

METHOD AND APPARATUS FOR STORING DATA

An SSD controller operates as an interface device conversant in a host protocol and a storage protocol supporting respective host and storage interfaces for providing a host with a view of an entire storage system. The host has visibility of the storage protocol that presents the storage system as a logical device, and accesses the storage device through the host protocol which is adapted for accessing high speed devices such as solid state drives (SSDs). The storage protocol supports a variety of possible dissimilar devices, allowing the host effective access to a combination of SSD and traditional storage as defined by the storage system. In this manner, a host protocol such as NVMe (Non-Volatile

"...or an initiator, particularly with PCIe® (Peripheral Component Interconnect Express, or PCI Express®). The NDAS system allows flexibility in abstracting various and possibly dissimilar storage devices which can include SATA (serial Advanced Technology Attachment, current specifications available at sata-io.org) HDDs (hard disk drives), SATA SSDs and PCIe/NVMe SSDs with NAND or other types of non-volatile memory. The storage devices within the NDAS system could then be used to implement various storage optimizations, such as aggregation, caching and tiering. By way of background, NVMe is a scalable host controller interface designed to address the needs of enterprise, data center and

1. An interface device, comprising: a host interface responsive to requests issued by a host, the host interface presenting a storage device for access by the host; a storage interface coupled to a plurality of dissimilar storage elements, the plurality of storage elements conversant in a storage protocol common to each of the plurality of storage elements; and a mapper connected between the host interface to the storage interface and configured to map requests received on the host interface to a specific storage element connected to the storage interface, the mapped request indicative of the specific storage element based on the storage protocol, the specific storage element independent of the

Publication date: 2015-06-04

Priority date: 2013-11-26

Inventor: GRIMSRUD Knut S. KHAN Jawad B.

Applicant: GRIMSRUD Knut S. KHAN Jawad B.

CPC class: G06F3/0613 G06F3/0604 G06F3/0607 G06F3/0661 G06F3/0683 G06F3/0685 G06F13/385 G06F13/4282 G06F2213/0026

IPC class: G06F3/06 G06F13/42

Link: Espacenet

US2013135816 A1

Similar

Method and Apparatus for Scalable Low Latency Solid State Drive Interface

An embodiment solid state drive (SSD) apparatus includes a plurality of computer processing unit (CPU) blades, a channel-interleaved interface operably coupled to the CPU blades, and an input/output (I/O) blade operably coupled to the channel-interleaved interface. In an embodiment, the CPU blades include a processor running a plurality of virtual machines that are locally switched using an Ethernet controller on a chip.

"...switch in the interface. As shown in FIG. 1, in an embodiment the SSD apparatus 10 includes several SSDs 12, a channel-interleaved interface 14, and a Peripheral Component Interconnect Express (PCIe) bridge 16. As used herein, the PCIe bridge 16 may represent or be referred to as PCIe, a PCIe bridge controller, and so on. The SSDs 12 in FIG. 1, which may also be referred to as a solid-state disk or electronic disk, are data storage devices that use integrated circuit assemblies as memory to store data persistently. The SSDs 12 do not employ any moving mechanical components, which distinguishes them from traditional magnetic disks such as hard disk drives (HDDs) or floppy disk, which are electromechanical

1. A solid state drive (SSD) apparatus, comprising: a plurality of computer processing unit (CPU) blades; a channel-interleaved interface operably coupled to the CPU blades; and an input/output (I/O) blade operably coupled to the channel-interleaved interface. 2. The SSD apparatus of claim 1, wherein the channel-interleaved interface utilizes a data frame format including a frame header, frame data, and a frame cyclic redundancy check (CRC). 3. The SSD apparatus of claim 1, wherein the channel-interleaved interface interleaves a read command between portions of write commands. 4. The SSD apparatus of claim 1, wherein the channel-interleaved interface issues write commands in multiple bursts. 5.

Publication date: 2013-05-30

Priority date: 2011-11-17

Inventor: HUANG YIREN

Applicant: FUTUREWEI TECHNOLOGIES, INC.

CPC class: G06F1/16 G06F3/0611 G06F3/0659 G06F3/0661 G06F3/0688

IPC class: G06F1/16

Link: Espacenet

US2018165698 A1

Related

METHODS AND SYSTEMS TO DETERMINE VIRTUAL STORAGE COSTS OF A VIRTUAL DATACENTER

Methods and systems that allocate the total cost of virtual storage created from hard disk drives ("HDDs") and solid state drives ("SSDs") of server computers and mass-storage devices of a cloud-computing facility are described. The virtual storage is used to form virtual disks ("VDs") of virtual machines ("VMs") comprising a virtual datacenter ("VDC"). Methods calculate a total virtual storage cost of the virtual storage from hardware costs and other costs such as labor, maintenance, facilities and licensing costs, which is used to calculate an HDD cost rate and an SSD cost rate. A cost of each VD is calculate based on virtual storage

1. A method to determine cost of virtual storage of a virtual datacenter created in a cloud-computing facility, the method comprising: calculating a total virtual storage cost based on depreciation cost of hard disk drives ("HDDs") and solid state drives ("SSDs"), licensing costs, labor costs, maintenance costs, and facility costs; calculating an HDD cost rate based on the total virtual storage cost and on depreciation costs of the HDDs and the SSDs; calculating an SSD cost rate based on based on the total virtual storage cost and on depreciation costs of the HDDs and the SSDs; calculating a cost of each virtual disk ("VD") of the virtual storage based on virtual storage policy parameters,

Publication date: 2018-06-14

Priority date: 2016-12-09

Inventor: CHANDRASHEKAR, SHRISHA SAHA, MRITYUNJOY POTLURI, VIJAY PALAVALLI,

Applicant: VMWARE, INC.

CPC class: G06Q30/0206 G06F3/0605 G06F3/0664 G06F3/067 G06F3/0685 G06F9/45558 G06F2009/4557 G06F2009/45579 G06F2009/45583

IPC class: G06Q30/02 G06F3/06 G06F9/455

Link: Espacenet

US2018349041 A1

Related

METHOD AND SYSTEM FOR IMPLEMENTING BYTE-ALTERABLE WRITE CACHE

One embodiment described herein provides a data storage system. The storage system includes a communication interface for receiving a write command from a host of the data storage system and a first write cache coupled to the host via the communication interface. The write command includes to-be-written data and address information associated with the to-be-written data. The first write cache is configured to temporarily store the to-be-written data before sending the to-be-written data to a high-capacity storage device coupled to the host, and the first write cache is configured to update its content by performing byte-level in-place writing based on the received write command.

"...NAND SSDs. First, NAND SSDs write at page level, and each page can be 16 KB or larger. Hence, changing a single byte in a page will require the entire page to be rewritten and, thus, can lead to high write amplification and significantly faster wearing of the drive. Moreover, NAND SSDs cannot be updated in-place; instead, pages must be erased before they can be written. The erase operations are typically performed at the block level, with a block being as large as several megabytes. Changing a single byte in a page can cause the whole page to be invalidated and redirected to a new page. Consequently, high-frequency small writes can cause an SSD to run out of clean page quickly, triggering

1. A data storage system, the system comprising: a communication interface configured to receive a write command from a host of the data storage system, wherein the write command comprises to-be-written data and associated address information; and a first write cache coupled to the host via the communication interface, wherein the first write cache is configured to temporarily store the to-be-written data before sending the to-be-written data to a high-capacity storage device coupled to the host, and wherein the first write cache is configured to update its content by performing byte-level in-place writing based on the received write command. 2. The data storage system of claim 1, wherein

Publication date: 2018-12-06

Priority date: 2017-05-31

Inventor: Zhou, Ping Li, Shu

Applicant: Alibaba Group Holding Limited

CPC class: G06F3/0619 G06F3/064 G06F3/0656 G06F3/0659 G06F3/0685 G06F12/0804 G06F12/0893

IPC class: G06F3/06 G06F12/0804 G06F12/0893

Link: Espacenet

US9720860 B2

Related

System and method for efficient processing of queued read commands in a memory system

A solid state drive (SSD) storage system includes a memory controller, host interface, memory channels and solid state memories as storage elements. The completion status of sub-commands of individual read commands is monitored and used to determine an optimal selection for returning data for individual read commands. The completion of a read command may be dependent on the completion of multiple individual memory accesses at various times. The queueing of multiple read commands which may proceed in parallel or out of order causes interleaving of multiple memory accesses from different commands to individual memories. A system and method is disclosed which enables the selection, firstly of completed

"...shows that sector 150 is first transferred, then sector 170, and then sector 160 before finally sector 180. The path taken (and the time taken to traverse the path) in this case is much shorter than that shown in FIG. 1, and so the data for the sectors may be transferred at a greater overall data rate as the time spent moving from sector to sector is shortened. In recent years, the use of non-volatile flash memory has become more and more prevalent. One use has been to build emulations of hard disk drives, known as Solid State Drives (SSDs), or as Solid State Disks. To allow SSDs to be used as a direct replacement for HDDs, they must offer the same interface and adhere to the same standards

1. A solid state drive memory storage system including at least one of Native Command Queueing (NCQ) and Non Volatile Memory Express (NVMe) queueing, comprising: a memory controller having an interface to receive a plurality of read commands from a host system, each read command of the plurality of read commands associated with a plurality of sub-commands for accessing a corresponding plurality of memory blocks; and memory channels to access a plurality of solid state memory devices, wherein the memory controller identifies a next read command from the plurality of read commands that will complete with the least delay or interruption by monitoring completion status of the plurality of sub-commands

Publication date: 2015-12-10

Priority date: 2014-06-06

Inventor: Rose Philip David Stephens Matthew

Applicant: OCZ Storage Solutions, Inc. Toshiba Corporation

CPC class: G06F13/1642 G06F12/0246 G06F13/1673 G06F13/4221 G06F2212/1024 G06F2212/7203 G06F2212/7208

IPC class: G06F7/00 G06F9/44 G06F12/02 G06F13/16 G06F13/42

Link: Espacenet

US9898208 B2

Related

Storage system with hybrid logical volumes utilizing in-band hinting

Systems and methods that combine a silicon storage volume with a hard disk drive (HDD) volume in a storage system that uses in band hinting to select the volume for storing actual data and meta data based on the demands of high performance computing are described. A storage system with an application processor, a storage processor, a silicon storage volume including a plurality of SSDs and an HDD with a much larger number of HDDs efficiently handles write requests from a high performance computer. A high performance computing device prepares internal meta data and actual data write requests by specifying in the internal write requests whether the data is actual data or meta data using an existing

"...than hard disk drives and have historically come in smaller capacities than hard disk drives. Additionally SSDs have a longer erase cycle than hard disk drives. Therefore overwriting previous data on SSDs is typically slower than on hard disk drives. An advantage of SSDs is that reading or writing data from random locations can be many times faster than with a hard disk drive. Hybrid drives and storage systems are a combination of flash memory, the technology SSDs are based on, and rotating media or HDDs. The flash memory on a hybrid drive is used to expand the available drive cache and also to act as a burst buffer for incoming writes. For a write, once the data is committed to flash memory

1. A data storage system comprising: an application processor, a storage processor, and random access memory (RAM) coupled to a bus; a plurality of solid state drives (SSDs) comprising a silicon storage volume, the silicon storage volume coupled to the bus; a plurality of hard disk drive (HDD) devices comprising an HDD volume, HDD volume coupled to the bus; the application processor and the storage processor each having instructions which when executed cause the application processor and storage processor to perform actions including: the application processor receiving a write request; the application processor preparing internal write requests based on the write request; the storage processor

Publication date: 2017-11-02

Priority date: 2016-04-27

Inventor: Parkhurst Michael

Applicant: DataDirect Networks, Inc.

CPC class: G06F3/0611 G06F3/061 G06F3/064 G06F3/0659 G06F3/068 G06F3/0685 G06F3/0688 G06F3/0689 G06F2212/217 G11C16/10

IPC class: G06F3/06 G11C16/10

Link: Espacenet

US9927975 B2

Background

Hybrid memory drives, computer system, and related method for operating a multi-mode hybrid drive

A multi-mode hybrid memory drive comprises a bulk memory device and a removable cache memory device. A controller of the bulk memory device may be configured to operate the bulk memory device in either a stand-alone mode or a hybrid mode responsive to detecting the removable cache memory device being coupled with a cache port of the bulk memory device. A method of operating a multi-mode hybrid drive may also comprise monitoring a cache port of a bulk memory device to determine a presence of a removable cache memory device, operating the bulk memory device as a stand-alone drive responsive to determining the removable cache memory device is not present, and operating the bulk memory device as

"...TECHNICAL FIELD The disclosure, in various embodiments, relates generally to the field of computer systems and mass storage devices. More specifically, the disclosure relates to mass storage devices configured to operate as either a hybrid drive or a stand-alone drive depending on its mode of operation. BACKGROUND Non-volatile memory is commonly used for mass storage of data, such as within consumer electronic devices. Various types of mass storage devices are currently in use, such as solid state devices (SSDs), hard disk drives (HDDs), and hybrid drives. SSDs use solid state memory devices (e.g., Flash memory), which can have advantages over the traditional electro-mechanical magnetic HDDs

1. A hybrid memory drive, comprising: a bulk memory device including: bulk storage media including non-volatile memory; a cache port; and a first controller operably coupled with the cache port and the bulk storage media; and a removable cache memory device including: cache storage media including non-volatile memory; and a second controller operably coupled with the cache storage media, wherein the first controller is configured to operate the bulk memory device in either a stand-alone mode or a hybrid mode responsive to detecting the removable cache memory device being at least one of coupled with the cache port of the bulk memory device or valid for use as a cache for the bulk memory device

Publication date: 2018-02-08

Priority date: 2016-08-03

Inventor: Pratt Thomas L.

Applicant: Micron Technology, Inc.

CPC class: G06F3/0604 G06F3/0605 G06F3/0634 G06F3/068 G06F12/0246 G06F12/0866 G06F12/0873 G06F2212/217 G06F2212/222 G06F2212/261 G06F2212/281 G06F2212/305 G06F2212/313 G06F2212/466 Y02D10/13

IPC class: G06F12/00 G06F3/06 G06F12/02 G06F12/0873

Link: Espacenet

US2018260135 A1

Background

INTEGRATED HETEROGENEOUS SOLID STATE STORAGE DRIVE

In one embodiment, a heterogeneous integrated solid state drive includes a plurality of solid state memory devices including at least one solid state memory device of a first type and at least one solid state memory device of a second type, a controller coupled to each of the plurality of solid state memory devices and an interface coupled to the controller. The controller is configured to receive at least one user-defined memory parameter and to create at least one namespace satisfying the at least one user-defined memory parameter in at least one of the plurality of solid state memory devices. In one embodiment, the at least one user-defined memory parameter is one of a group consisting of

"...(DRAM) or static random access memory (SRAM), or non-volatile memory, such as NAND flash memory. The standard NAND flash memory can be Single Level Cell (SLC) or Multi Level Cell (MLC), including enterprise MLC (eMLC), Triple Level Cell (TLC) and Quadratic Level Cell (QLC). While the higher data density of MLC memory reduces the cost per unit of storage, SLC memory has faster write speeds, lower power consumption and higher cell endurance. While early SSDs retain legacy HDD form factors and connectivity, modern SSDs break away from these limitations. Without the constraints of platter mechanics inherent in HDDs, modern SSDs have superior mechanical ruggedness and can come in much smaller

1. A heterogeneous integrated solid state drive comprising: a plurality of solid state memory devices including at least one solid state memory device of a first type and at least one solid state memory device of a second type different from the first type; a controller coupled to each of the plurality of solid state memory devices; and an interface coupled to the controller, the interface configured to communicate with a device external to the heterogeneous integrated solid state drive, the controller configured to receive at least one user-defined memory parameter and to create at least one namespace satisfying the at least one user-defined memory parameter in at least one of the plurality

Publication date: 2018-09-13

Priority date: 2017-03-10

Inventor: Hayashida, Mark Klein, Yaron

Applicant: Toshiba Memory Corporation

CPC class: G06F3/0611 G06F3/0659 G06F3/0685 G06F13/1673 G06F13/4282

IPC class: G06F3/06 G06F13/16 G06F13/42

Link: Espacenet

US9619164 B2

Background

Cluster solid state drives

Described herein are techniques for arranging a plurality of M.2 solid state drive (SSD) modules and flash storage elements into a compact form factor. On a first side of an SSD sled, a plurality of M.2 SSD modules may be communicatively coupled to a port expander. On a second side of the SSD sled, a plurality of flash storage elements (not packaged into M.2 SSD modules) may be present. A plurality of SSD sleds (with the above-described characteristics) may be sized so as to collectively fit into a single hard disk drive (HDD) compatible compartment of a chassis.

"...RELATED APPLICATIONS This application is a non-provisional patent application of and claims priority to U.S. Provisional Application No. 61/989,452, filed 6 May 2014, which is assigned to the assignee of the present invention and is incorporated by reference herein. FIELD OF THE INVENTION The present invention relates to methods for arranging a plurality of solid state drives (SSDs) into a compact form factor, and storage systems in which a plurality of SSDs are arranged into a compact form factor. BACKGROUND Most commercially available storage systems generally include those with disk drives (e.g., hard disk drives (HDDs)), those with solid state drives (SSDs) (e.g., flash drives), and

1. A storage system, comprising: a plurality of solid state drive (SSD) sleds, wherein each of the SSD sleds includes: (i) a plurality of M.2-compliant SSDs, each of the M.2-compliant SSDs comprising a serial advanced technology attachment (SATA) SSD controller and a plurality of NAND flash chips, (ii) a port expander of the SSD sled, each of the M.2-compliant SSDs being communicatively coupled to the port expander of the SSD sled, and (iii) a plurality of flash storage elements, each of the flash storage elements comprising a NAND flash chip, wherein the plurality of SSD sleds is sized so as to collectively fit into a single hard disk drive (HDD) compatible compartment of a chassis, wherein

Publication date: 2015-11-12

Priority date: 2014-05-06

Inventor: Mehta Varun McKnight Tom

Applicant: Nimble Storage, Inc.

CPC class: G06F3/0626 G06F3/061 G06F3/0655 G06F3/0658 G06F3/0688

IPC class: G06F12/00 G06F3/06

Link: Espacenet