



(51) International Patent Classification:

G06F 12/0868 (2016.01) G06F 12/0875 (2016.01)
G06F 3/06 (2006.01) G06F 11/14 (2006.01)

(21) International Application Number:

PCT/US2017/064368

(22) International Filing Date:

02 December 2017 (02.12.2017)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

15/470,787 27 March 2017 (27.03.2017) US

(71) Applicant: WESTERN DIGITAL TECHNOLOGIES, INC. [US/US]; 5601 Great Oak Parkway, San Jose, CA 95119 (US).

(72) Inventors: BHANSALI, Akhil; c/o Western Digital Technologies, Inc., 5601 Great Oak Parkway, San Jose, CA 95119 (US). MEHTA, Kirankumar; C/o Western Digital Technologies, Inc., 5601 Great Oak Parkway, San Jose, CA 95119 (US). KEYMOLEN, Bruno; C/o Western Digital Technologies, Inc., 5601 Great Oak Parkway, San Jose, CA 95119 (US). BHAIIRY, Rajendra; C/o Western Digital

Technologies, Inc., 5601 Great Oak Parkway, San Jose, CA 95119 (US).

(74) Agent: BOHN, Michel et al.; Patent Law Works LLP, 201 S. Main Street, Suite 250, Salt Lake City, UT 84111 (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, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

(54) Title: MANAGING I/O OPERATIONS FOR DATA OBJECTS IN A STORAGE SYSTEM

(57) Abstract: Various aspects for managing input/output (I/O) operations for data objects (e.g., large data objects (LOBs)), in a storage system are provided. For instance, a method may include receiving an I/O request for executing an I/O operation on a data object, determining the size of the data object, and determining the data object to be a non-conforming data object relative to a storage threshold of a cache memory device. The method may further include dividing the non-conforming data object into multiple data chunks, each of the data chunks having a size that is less than the storage threshold, moving one or more data chunks stored in one or more memory spaces of the cache memory device to a long-term memory device to free up the one or more memory spaces, and storing one or more of the data chunks of the non-conforming data object in the freed-up memory space(s).




FIG. 6

Published:

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

**MANAGING I/O OPERATIONS FOR DATA OBJECTS IN A STORAGE
SYSTEM**

Inventors:

Akhil Bhansali;

5

Kirankumar Mehta;

Bruno Keymolen; and

Rajendra Kumar Bhairy

BACKGROUND

10 **Field of the Technology**

[0001] The present disclosure relates to computing systems and, more particularly, to managing data objects in a storage system.

Description of Related Art

15 **[0002]** Various computing networks and systems are capable of storing data. A subset of these computing networks and systems are configured to store data as data objects. Some of these data objects are capable of containing large amounts of data (e.g., at least two terabytes (2TB) of data).

[0003] Typically, when a computing system receives an input/output (I/O) request to
20 write a data object to long-term storage, the data object is initially written, in its entirety, to a storage buffer of a cache. The data object in the storage buffer is subsequently written to long-term storage and may be accessed (e.g., via read requests and/or read-write requests) in the future so that the data object, in its entirety, can be read or modified (e.g., read and subsequently written with new data).

25 **[0004]** For access to a data object in long-term storage via read requests, the entire data object is read to the storage buffer. That is, all read requests to read data in the data object result in the whole data object being read to the storage buffer, even if only a subset or a small portion of the data in the data object needs reading.

[0005] For access to a data object in long-term storage via read-write requests (e.g., to
30 modify data in the data object), the entire data object is read to the storage buffer, the data desiring modification is changed, and the entire modified data object is eventually written

(re-written) to long-term storage. That is, all read-write requests to modify data in the data object result in the whole data object being read to the storage buffer (even if the data needing modification is a subset of the entire data object or only a small portion of the data in the data object), data in the data object is modified, and the entire modified data object is eventually written (re-written) to long-term storage, even though only a subset or a small portion of the data in the data object needed modification.

[0006] While some previous computing networks and systems are able to store data objects, these computing networks and systems are not as reliable and/or efficient as they could be. In some cases, practical considerations, such as cost, size limitations, storage limitations, etc., limit the size of a cache memory device of a storage system. As such, some storage systems may fail (e.g., crash, etc.) and/or underperform (e.g., have latency issues, etc.) in situations in which one or more I/O requests (i.e., read, write, or read-write requests) that correspond to or are associated with a data object that includes a size larger than the size of a cache memory device (e.g., buffer size) in the storage system. Other issues encountered by some storage systems can be rooted in the cache memory device being exhausted (i.e., full and/or storing an amount of data above a predetermined threshold amount of data) and/or not storing one or more desired data objects and/or data chunks. If a cache memory device is insufficient to satisfy read requests and/or has reached its high-water mark for write requests, performance of the storage system can drop off considerably.

[0007] For example, previous computing networks and systems can struggle to write data objects when the data objects have a size that is larger than the size of a storage buffer in the computing network and/or system. That is, at least some networks and/or systems may malfunction (e.g., crash) and/or experience latency issues when the size of a data object included in a write request is larger than the size of its storage buffer because the storage buffer is unable to write data objects larger than its storage capacity.

[0008] Some previous computing networks and/or systems may experience undue latency issues because they perform read operations by reading an entire data object to its storage buffer from long-term storage when only a subset of the data in a data object needs reading. That is, these previous computing networks and systems may take more time than necessary to perform read operations because the entire data object is read to the storage buffer when only a portion of the data in the data object (i.e., data desiring to be read or desired data) needs to be read, which increases the amount of time it takes to read the desired data to the storage buffer because superfluous data is read to the storage buffer.

[0009] Similarly, undue latency issues may be experienced by some previous computing networks and/or systems when modifying data in long-term storage because they perform read-write operations on the entire data object when only a subset of the data in a data object needs modification. That is, some pervious computing networks and systems exhibit lag when performing read-write operations because the entire data object is read to the buffer memory when only a portion of the data in the data object (i.e., data desiring to be modified or desired data) needs to be modified. This increases the amount of time it takes to modify the desired data because superfluous data is read to the storage buffer and subsequently written (re-written) to long-term storage.

10

SUMMARY

[0010] Various aspects for managing input/output (I/O) operations for data objects, particularly, but not limited to, large data objects (LOBs), in a storage system are described. In an innovative aspect, a system comprises a first non-volatile cache memory device and a storage controller coupled to the cache memory device. The storage controller includes one or more processing devices for executing instructions in one or more applications configured to cause the storage controller to perform operations comprising receiving an I/O request for a data object, determining a size of the data object, and determining the data object to be a first non-conforming data object relative to the storage threshold of the first non-volatile cache memory device. In some embodiments, the storage controller further performs operations comprising dividing the first non-conforming data object into a plurality of data chunks, each of the data chunks having a size that is less than the storage threshold, moving one or more data chunks stored in one or more memory spaces of the first non-volatile cache memory device to a long-term memory device to free up the one or more memory spaces, and storing one or more of the data chunks of the first non-conforming data object in the freed-up one or more memory spaces of the first non-volatile cache memory device.

15

20

25

[0011] In various embodiments, the I/O request is a write request and the storage controller, when performing the operations for storing the one or more data chunks of the first non-conforming data object, further performs operations comprising writing each of the one or more data chunks of the first non-conforming data object to the freed-up one or more memory spaces of the first non-volatile cache memory device. In some embodiments, the I/O request is a read request and the storage controller, when performing the operations for

30

storing the one or more data chunks of the first non-conforming data object, further performs operations comprising reading one or more data chunks stored in the long-term memory device to the first non-volatile cache memory device. In further embodiments, the I/O request is a read-write request and the storage controller, when performing the

5 operations for storing the one or more of the data chunks of the first non-conforming data object, further performs operations comprising reading one or more data chunks stored in the long-term memory device to the first non-volatile cache memory device, modifying data in the one or more read data chunks to generate one or more modified data chunks, and writing the one or more modified data chunks to the long-term memory device.

10 **[0012]** Some systems comprise the long-term memory device being coupled to the cache memory device, where determining the data object to be a first non-conforming object is further based on the size of the data object exceeding an available storage capacity of the first non-volatile cache memory device. Further, storing the one or more of the data chunks of the first non-conforming data object in the freed-up one or more memory spaces

15 includes, for each data chunk of the first non-conforming data object, determining if there is sufficient space for storing each data chunk in the first non-volatile cache memory device and, when there is not sufficient space, writing one or more existing data chunks that are associated with one or more data objects other than the first non-conforming data object and currently stored in the first non-volatile cache memory device to the long-term memory

20 device, freeing up the one or more memory spaces of the one or more existing data chunks to generate the freed-up one or more memory spaces, and writing each data chunk to the freed-up one or more memory spaces in the first non-volatile cache memory device.

[0013] In various embodiments, the system further comprises a second non-volatile cache memory device coupled to the first non-volatile cache memory device and the storage

25 controller further performs operations comprising replicating each data chunk stored in the first non-volatile cache memory device to the second non-volatile cache memory device. The first non-volatile cache memory device, in some embodiments, includes a buffer size and the storage threshold reflects the buffer size. In further embodiments, the first non-volatile cache memory device comprises a flash memory device.

30 **[0014]** The storage controller, in various embodiments, is configured to be coupled to a network that is in communication with one or more client devices. In some embodiments, the storage controller is located between the one or more client devices and the first non-volatile cache memory device when coupled to the network. In further embodiments, the

long-term memory device is a portion of a local storage network and/or a portion of an enterprise storage service.

[0015] In another innovative aspect, a system comprises a non-volatile cache memory device including a buffer size, means for receiving an I/O request for a data object, means
5 for determining a size of the data object, and means for determining the data object to be a non-conforming data object relative to the buffer size of the non-volatile cache memory device. The system, in various embodiments, further comprises means for dividing the non-conforming data object into a plurality of data chunks, each of the data chunks having a size that is less than the buffer size, means for moving one or more data chunks stored in one or
10 more memory spaces of the non-volatile cache memory device to a long-term memory device to free up the one or more memory spaces, and means for storing one or more of the data chunks of the non-conforming data object in the freed-up one or more memory spaces of the non-volatile cache memory device.

[0016] In various embodiments, the I/O request is a write request and the means for
15 storing one or more of the data chunks of the non-conforming data object further comprises means for writing each of the one or more data chunks of the non-conforming data object to the freed-up one or more memory spaces of the non-volatile cache memory device. In some embodiments, the I/O request is a read request and the means for storing one or more of the data chunks of the non-conforming data object further comprises means for reading one or
20 more data chunks stored in the long-term memory device to the cache memory device. In further embodiments, the I/O request is a read-write request and the means for storing one or more of the data chunks of the non-conforming data object further comprises means for reading one or more data chunks stored in the long-term memory device to the cache memory device, means for modifying data in the one or more read data chunks to generate
25 one or more modified data chunks, and means for writing the one or more modified data chunks to the long-term memory device.

[0017] In some embodiments, the system further comprises the long-term memory device coupled to the cache memory device and the means for determining the data object to be a non-conforming object comprises means for basing the determination on the size of
30 the data object exceeding the buffer size. In some systems, the means for storing the one or more of the data chunks of the non-conforming data object in the one or more memory spaces includes, for each data chunk of the non-conforming data object, means for determining if there is sufficient space for storing each data chunk in the non-volatile cache

memory device and, when there is not sufficient space, means for writing one or more existing data chunks that are associated with one or more data objects other than the non-conforming data object and currently stored in the non-volatile cache memory device to the long-term memory device, means for freeing up the one or more memory spaces of the one or more existing data chunks to generate the freed-up one or more memory spaces, and means for writing each data chunk to the freed-up one or more memory spaces in the non-volatile cache memory device.

[0018] In yet another innovative aspect, a method of operation in a non-volatile memory system including a set of non-volatile memory devices comprises receiving, at a storage controller coupled to a non-volatile cache memory device including a storage threshold, an I/O command for executing an I/O operation on a data object, determining a size of the data object, determining a size of the data object, and determining the data object to be a non-conforming data object relative to the storage threshold of the cache memory device. The method, in various embodiments, further comprises dividing the non-conforming data object into a plurality of data chunks, each of the data chunks having a size that is less than the storage threshold, moving one or more data chunks stored in one or more memory spaces of the non-volatile cache memory device to a long-term memory device to free up the one or more memory spaces, and storing one or more of the data chunks of the non-conforming data object in the freed-up one or more memory spaces of the non-volatile cache memory device.

[0019] In various embodiments, the I/O request is a write request and storing the one or more of the data chunks of the non-conforming data object further comprises writing each of the one or more data chunks of the non-conforming data object to the freed-up one or more memory spaces of the non-volatile cache memory device. In some embodiments, the I/O request is a read request and storing the one or more of the data chunks of the non-conforming data object further comprises reading each of the one or more data chunks of the non-conforming data object to the freed-up one or more memory spaces of the non-volatile cache memory device. In further embodiments, the I/O request is a read-write request and storing the one or more of the data chunks of the non-conforming data object further comprises reading each of the one or more data chunks of the non-conforming data object to the freed-up one or more memory spaces of the non-volatile cache memory device, modifying data in each of the one or more read data chunks to generate one or more

modified data chunks, and writing each of the one or more modified data chunks to the long-term memory device.

5 [0020] In some embodiments, determining the data object to be a non-conforming object is further based on the size of the data object exceeding an available storage capacity of the cache memory device. In some methods, storing the one or more of the data chunks of the non-conforming data object in the one or more memory spaces includes, for each data chunk of the non-conforming data object, determining if there is sufficient space for storing each data chunk in the non-volatile cache memory device and, when there is not sufficient space, writing one or more existing data chunks that are associated with one or more data objects other than the non-conforming data object and currently stored in the non-volatile cache memory device to the long-term memory device, freeing up the one or more memory spaces of the one or more existing data chunks to generate the freed-up one or more memory spaces, and writing each data chunk to the freed-up one or more memory spaces in the non-volatile cache memory device.

10 [0021] The various embodiments disclosed herein include operations to overcome or at least reduce the issues in the previous storage networks and/or systems discussed above and, accordingly, are more reliable and/or more efficient than other computing networks and/or networks. That is, the various embodiments disclosed herein include hardware and/or software with functionality to avoid or at least decrease the issues related to reliability and/or latency associated with writing large data objects (e.g., data objects including a size greater than the size or storage capacity of a cache memory device). Furthermore, the various embodiments disclosed herein include hardware and/or software to avoid or at least decrease the latency issues associated with requests to read data (e.g., read requests) and/or requests to modify data (e.g., read-write requests). Accordingly, the embodiments disclosed therein provide various improvements to storage networks and/or storage systems.

25 [0022] It should be understood that language used in the present disclosure has been principally selected for readability and instructional purposes, and not to limit the scope of the subject matter disclosed herein.

30 BRIEF DESCRIPTION OF THE DRAWINGS

[0023] The present disclosure is illustrated by way of example, and not by way of limitation in the figures of the accompanying drawings in which like reference numerals are used to refer to similar elements.

- [0024] Figure 1A is a diagram of an example computing network for storing data;
- [0025] Figure 1B is a diagram of an example storage system included in the computing network of Figure 1A;
- [0026] Figure 2 is a diagram of an example computing network for storing data;
- 5 [0027] Figures 3A through 3C are flow diagrams illustrating various examples of input/output (I/O) operations performed for large data objects;
- [0028] Figure 4 is a diagram illustrating operations for various example hardware and/or software to write a data object to long-term storage;
- [0029] Figures 5A through 5G are diagrams illustrating various examples of a least
10 recently used (LRU) technique utilized by one embodiment of a storage controller for managing a cache memory device;
- [0030] Figure 6 is a block diagram of an example method for writing a data object to a buffer memory;
- [0031] Figure 7 is a block diagram of an example method for reading a data object to a
15 buffer memory; and
- [0032] Figure 8 is a block diagram of an example method for modifying a data object in long-term storage.

DETAILED DESCRIPTION

- 20 [0033] The innovative technology disclosed herein includes various aspects, such as systems, methods, apparatuses, computer-readable media, computer program products, etc., for managing input/output (I/O) operations for data objects, particularly large data objects, in a storage system. The technology is advantageously more efficient and/or more reliable than previous computing networks and/or networks utilized to store data and, particularly,
25 utilized to store data objects, especially when the data objects include large amounts of data (are large data objects (LOBs)). A LOB includes a quantity of data or an amount of data greater than a temporary non-transitory storage medium of a storage device, as discussed in further detail below.
- [0034] The various embodiments disclosed below include hardware and/or software to
30 advantageously increase network and system reliability and/or decrease latency issues associated with writing large data objects (e.g., data objects including a size greater than the size or storage capacity of the buffer memory) to long-term storage. Furthermore, the various embodiments disclosed herein include hardware and/or software to avoid or at least

decrease latency issues associated with requests (e.g., read requests) to read data (e.g., data objects) from long-term storage and/or associated with requests (e.g., read-write requests) to modify data (e.g., data objects) in long-term storage.

[0035] Turning now to the drawings, Figure 1A is a block diagram of an example storage network 10. As shown, storage network 10 can include multiple client devices 25 capable of being coupled to and in communication with a storage system 100 via a wired and/or wireless network 75 (e.g., public and/or private computer networks in any number and/or configuration (e.g., the Internet, an intranet, a cloud network, etc.)), among other examples that may include one client device 25 or two or more client devices 25 (e.g., is not limited to three client devices 25).

[0036] A client device 25 can be any computing hardware and/or software (e.g., a thick client, a thin client, or hybrid thereof) capable of accessing storage system 100 utilizing network 75. Each client device 25, as part of its respective operation, relies on sending input/output (I/O) requests to storage system 100 to write data, read data, and/or modify data. Specifically, each client device 25 can transmit I/O requests to read, write, store, communicate, propagate, and/or transport instructions, data, computer programs, software, code, routines, etc., to storage system 100. Client device(s) 25 and storage system 100 may comprise at least a portion of a client-server model. In general, storage system 100 can be accessed by client device(s) 25 and/or communication with storage system 100 can be initiated by client device(s) 25 through a network socket (not shown) utilizing one or more inter-process networking techniques.

[0037] With reference to Figure 1B, a block diagram of an example storage system 100 is illustrated. As shown, storage system 100 may comprise, among other components, multiple storage devices 102a, 102b, ...102n and a storage controller 104 coupled to and in communication with storage devices 102a, 102b, ...102n via a wired and/or wireless network 110.

[0038] By utilizing network 110, storage system 100, in various embodiments, can comprise a cloud network (IAN), a SAN (e.g., a storage area network, a small area network, a server area network, and/or a system area network), a wide area network (WAN), a local area network (LAN), a wireless local area network (WLAN), a metropolitan area network (MAN), an enterprise private network (EPN), a virtual private network (VPN), and/or a personal area network (PAN), among other examples of computing networks and/or or sets of computing devices connected together for the purpose of sharing resources.

[0039] Storage devices 102a, 102b, ...102n (also simply referred individually, in various groups, or collectively as storage device(s) 102) may each be any type of storage device that is known or developed in the future and is capable of storing data and/or data objects. In various embodiments, a storage device 102 may include one or more non-transitory
5 computer-usable (e.g., readable, writable, etc.) mediums, which may include any non-transitory and/or persistent apparatus or device that can contain, store, communicate, propagate, and/or transport instructions, data, computer programs, software, code, routines, etc., for processing by or in connection with a computer processing device.

[0040] As shown in Figure 1B, storage devices 102a, 102b, ...102n can be implemented
10 as flash memory (e.g., a solid-state devices (SSD) or other non-volatile storage device that stores persistent data or persistent storage device). In further embodiments, storage devices 102a, 102b, ...102n may include non-transitory memory such as, for example, a dynamic random access memory (DRAM) device, a static random access memory (SRAM) device, a hard disk drive (HDD), and/or storage tape (e.g., magnetic and/or virtual), among other
15 types (e.g., non-volatile and/or persistent) of memory devices.

[0041] In various embodiments, one or more of storage devices 102a, 102b, ...102n can include non-volatile/persistent hardware and/or software configured for performing long-term data storage operations (e.g., data archiving, data back-up, data mirroring, replicating data, etc.). For instance, each of storage devices 102a, 102b, ...102n may include non-
20 volatile and/or persistent hardware and/or software configured for performing long-term data storage operations. Further, a storage device 102 may include hardware and/or software that functions to store data and/or one or more data objects as a single entry (e.g., when the size of the data object is less than or equal to the buffer size of a cache memory device 106 (e.g., the size of a buffer memory 108)).

[0042] Alternatively or additionally, a storage device 102 may include hardware and/or software to store one or more data objects divided/partitioned into multiple data chunks (e.g., a plurality of data chunks or two or more data chunks). Example data objects (and data chunks) include, but are not limited to, a group/grouping of data, a sequence of data, a data table, an index, a procedure, a routine, a view, a snapshot, and/or the like
25 contents/classes of data.

[0043] In some embodiments, a storage device 102 may comprise at least a portion of an object storage system. In various embodiments, a storage device 102 can include hardware and/or software to store data objects (e.g., data files) in containers that are

organizationally flat (e.g., non-hierarchical), and can use unique identifiers (IDs) to retrieve the data objects. In various embodiments, a storage device 102 may use metadata to store and access data objects and store the metadata with its corresponding data object.

Accordingly, in various embodiments, one or more additional storage devices 102 may be added, as desired, to storage system 100 without materially increasing processing overhead, which can provide scalability to storage system 100.

[0044] A data object may have any size capable of being stored in a storage device 102. Data objects stored in a storage device 102 may include the same predetermined size, different sizes, or a combination of the foregoing. In one example, a data object can have a size of about two or more terabytes (2TB) of data, although other sizes greater than or less than 2TB of data are also applicable. A data object may have a size (include a quantity of data) that exceeds the buffer size, which may reflect the amount of available memory/storage space in cache memory device 106 and/or buffer memory 108 (e.g., data object size > buffer size), in which case it is considered a LOB. A LOB may be stored in an inode.

[0045] An inode has a data structure that stores a data object and one or more data chunks corresponding to the data object. A data chunk is a constituent portion or partition of the data object. An inode that stores a LOB may be referred to herein as a non-conforming inode. When storing or retrieving the LOB in or from a storage device 102, the LOB in each non-conforming inode may be divided (or partitioned) into two or more data chunks. For example, one or more data chunks of a particular non-conforming inode may be stored in storage device 102a and one or more data chunks of the same non-conforming inode may be stored in storage device 102b, among other examples of data chunk distribution among storage devices 102a, 102b, ...102n.

[0046] A data object may be divided into any desired number of data chunks. In some embodiments, a data object may include the same number (or quantity) of data chunks, while in other embodiments two or more data objects can include different quantities of data chunks. In an example, a data object may include up to about ten thousand (10,000) data chunks (e.g., a number/quantity of data chunks in the range of 1 data chunk to 10,000 data chunks), among other example ranges and/or quantities of data chunks.

[0047] A data chunk may have any size suitable for long-term or non-volatile storage in a storage device 102 and storage in cache memory device 106, which is discussed further below. In some embodiments, two or more data chunks have the same or substantially

same size. For instance, two or more data chunks may include the same quantity of bytes or substantially the same quantity of bytes (e.g., within 90% or more of one another). In further embodiments, two or more data chunks may have different sizes (e.g., include different quantities of data or differing amounts of data). In some embodiments, a first data chunk and a last data chunk in a data object may have different sizes than the other data chunks in the data object, which may further include the same size or different sizes with respect to one another.

[0048] In a further example, each data chunk may include about two hundred megabytes (200MB) of data (e.g., a size of about $200\text{MB} \pm 0\text{-}10\%$), among other data chunk examples that may include greater than or less than about 200MB of data; and each data object may include about 2TB of data (e.g., is about a $2\text{TB data object} \pm 0\text{-}10\%$) that is divided into 10,000 data chunks, each of which includes about 200MB of data. Further, the first and last data chunks (e.g., data chunks 00001 and 10000) may include different sizes than the other data chunks (e.g., the other 9,998 data chunks may include about 200MB of data, while data chunks 00001 and 10000 may include less than or greater than about 200MB).

[0049] An inode for a data object may be assigned a unique object identifier, which may include any quantity and/or combination of alpha-numeric characters and/or digits. Each unique object identifier may be randomly assigned or may be assigned in a predetermined manner or order. In some embodiments, a data object may include a unique object identifier (e.g., DataObject ID) having twenty-five (25) alpha-numeric digits, among other example quantities of alpha-numeric digits that may be greater than or less than twenty-five (25) alpha-numeric digits. One particular example of a DataObject ID including twenty-five (25) alpha-numeric digits may be: 1234567890123456789012345 (i.e., twenty-five (25) numeric digits in this example), or other example combinations of alphabetic characters and/or numeric characters (or numerals).

[0050] Particular data objects that include a quantity of data or an amount of data (e.g., a data object size) that is less than or equal to the buffer size or amount of memory/storage space in cache memory device 106 and/or buffer memory 108 (e.g., $\text{data object size} \leq \text{buffer size}$) are referred to as conforming data objects. A conforming data object may be provided in an inode referred to as a conforming inode. A conforming inode may be assigned a unique object identifier and a representation of the data (e.g., metadata) in the data object. An example unique object identifier (e.g., DataObject ID) for a conforming inode may be:

2345678901234567890123456-<metadata>, wherein the inclusion of <metadata> in this example is a symbolic representation of the actual metadata for the data of an actual conforming inode.

[0051] A data chunk may include a unique data chunk identifier, which may include any quantity and/or combination of alpha-numeric characters/digits. Each unique data chunk identifier may be randomly assigned or may be assigned in a predetermined manner or order. In some embodiments, a unique data chunk identifier (e.g., DataChunk ID) may include a portion that identifies a corresponding data object (e.g., 1234567890123456789012345 in the above example), a portion (e.g., a set of alpha-numeric characters/digits) that identifies a particular data chunk (e.g., 00001 through 10000), and a portion of metadata representing/describing the data in the data chunk (e.g., represented by <metadata>), which can be generalized as: <DataObject ID>-<DataChunk ID>-<metadata>. Accordingly, each data chunk for a particular data object will have the same prefix (i.e., the DataObject ID) and a different postfix (i.e., DataChunk ID and metadata). An example of three consecutive data chunks may have the following DataChunk IDs:

1234567890123456789012345-00001-<metadata>;
1234567890123456789012345-00002-<metadata>; and
1234567890123456789012345-00003-<metadata>,

with the metadata of the data in each data chunk represented by <metadata> subsequent to the DataChunk ID in these examples. As such, storage devices 102a, 102b, ...102n may store conforming data objects as a whole and/or LOBs divided into multiple data chunks.

[0052] As shown, storage controller 104 may comprise, among other components, one or more cache memory devices 106 that may include, among other components, a buffer memory 108. Each cache memory device 106 can be referred to individually, in various groups, or collectively as cache memory device(s) 106. A cache memory device 106 may be any type of hardware and/or software memory device that is known or developed in the future and is capable of storing one or more data objects and/or data chunks for retrieval when an I/O request (e.g., write request, read request, and/or read-write request, etc.) is received at storage system 100.

[0053] A cache memory device 106 may respond with a write-through acknowledgement or a write-back acknowledgement for a write request. For write-through acknowledgement configurations, a cache memory device 106 may write a data object or

data chunk to buffer memory 108, write the data object or data chunk to one or more of storage devices 102a, storage device 102b, ... storage device 102n, and acknowledge the write request after the data object or data chunk resides on storage device(s) 102.

[0054] For write-back acknowledgement configurations, a cache memory device 106 may acknowledge a write request as the data object or data chunk is stored (e.g., buffered, held, etc.) in buffer memory 108. A write-back acknowledgement configuration may allow a cache memory device 106 to buffer data objects and/or data chunks rapidly and then write the data objects and/or data chunks to relatively slower long-term storage device(s) (e.g., one or more of storage devices 102a, 102b, ... 102n) when the long-term storage device(s) has/have bandwidth to perform I/O operations (e.g., read operations, write operations, and/or read-write operations, etc.).

[0055] A cache memory device 106 may perform read requests by providing one or more data objects and/or one or more data chunks to its buffer memory 108 for storage. For instance, cache memory device 106 may be configured to store (hold, buffer, accommodate, etc.) data objects and/or data chunks that have been previously written and/or read. This may be performed, in some cases, regardless of any pre-fetching activities storage controller 104 may have performed (e.g., read-ahead operations, as discussed elsewhere herein).

[0056] In various embodiments, a cache memory device 106 may store data objects and/or data chunks in its buffer memory 108 in accordance with a least recently used (LRU) technique. An example LRU technique removes/deletes older data objects and/or data chunks from buffer memory 108 to free up storage/buffer space in buffer memory 108 and populates the freed-up storage space with newly accessed data object(s) and/or data chunk(s). In other words, storage controller 104 may buffer each data object or each data chunk of a data object (e.g., a sequence of data chunks) as the data object(s) and/or data chunk(s) are accessed to fill available storage space in a buffer memory 108 of a cache memory device 106. Other techniques and variations are also possible and contemplated.

[0057] A cache memory device 106 may anticipate one or more future read requests (e.g., from one or more of clients 50) and store one or more data objects and/or data chunks in buffer memory 108 that may correspond to the read request(s) prior to receipt of the anticipated read request(s), which can be referred to as read-ahead functionality. In some embodiments, a cache memory device 106 may attempt to predict a particular set of data objects or a particular set of data chunk based on one or more previously received I/O requests. For example, a cache memory device 106 may predict that the next I/O request

will be a read request to read a data chunk including the DataChunk ID 1234567890123456789012345-00004-◇ because the three previous I/O requests were read requests to read data chunks including DataChunk IDs 1234567890123456789012345-00001-◇, 1234567890123456789012345-00002-◇, and 1234567890123456789012345-00003-◇, in succession.

[0058] The efficiency and extent to which a storage system may perform may be a function of the amount of memory in its cache memory device (e.g., buffer size). That is, larger cache memory devices may include a large buffer memory that can buffer more data objects and/or data chunks, which often results in better performance (e.g., as measured in I/O operations per second (IOPS), throughput, etc.). The storage technology disclosed herein advantageously addresses situations in which I/O requests correspond to LOBs, situations in which the cache memory device is full or substantially full (e.g., stores an amount of data above the predetermined threshold amount of data), and/or situations in which the cache memory device does not store desired data (e.g., one or more data objects and/or one or more data chunks).

[0059] In various embodiments, storage controller 104 may be coupled between client device(s) 25 and cache memory device(s) 106. Further, storage controller 104 may include hardware or a combination of hardware/software that can determine the size (e.g., buffer size) of an associated cache memory device (e.g., cache memory device(s) 106 and/or one or more buffer memories 108) and receive I/O requests (e.g., write requests, read requests, and read-write requests) from one or more client devices (e.g., client devices 25) via a network (e.g., network 75). Storage controller 104, in some embodiments, can determine the size of a data object (e.g., the data object size) upon receipt of an I/O request to write (e.g., a write request) the data object to long-term storage (e.g., storage devices 102a, 102b, ...102n) and compare the data object size to the buffer size of a buffer memory 108 in a cache memory device 106.

[0060] Storage controller 104 may assign conforming data objects and non-conforming data objects (LOBs) with data object identifiers. Storage controller 104 may write conforming data objects, as a whole, to a cache memory device 106 and may write LOBs by dividing the LOBs into data chunks or partitions, as discussed elsewhere herein. For example, storage controller 104 may assign a conforming data object 2345678901234567890123456-◇ as its DataObject ID, may assign no DataChunk IDs, and may include metadata in the postfix.

[0061] In various embodiments, storage controller 104 may organize data objects in a directory structure based on the unique object identifiers assigned to them. A directory structure may include any number of levels or structure suitable for locating the data object(s). In an example, data objects may be organized in a three-level directory based on the unique object identifier, among other examples implementing a fewer number of levels or a greater number of levels in a directory structure. For instance, the first set of three numbers (e.g., [123]), the second set of three numbers (e.g., [456]), and the third set of three numbers (e.g., [789]) may define the first level, second level, and third level, respectively, of the directory with the remaining numbers (e.g., sixteen (16) numbers (e.g., 0123456789012345)) providing identification of a specific data object in the directory. Accordingly, the DataObject ID for this particular instance may be: [123] [456] [789] 0123456789012345. As such, each data chunk created for this particular data object may be located in a directory addressed as: /mnt/ssd/123/456/789/.

[0062] A data chunk in a given divided/partitioned data object may include a unique object identifier, such as a DataChunk ID, which may include any quantity and/or combination of alpha-numeric and/or other type of symbols and further includes metadata, as discussed above. For example, a fifth data chunk in particular data object may be assigned 00005-< as its DataChunk ID, with 00005 as its DataChunk ID (e.g., 00001-00004 already being in use or assigned) and metadata in the postfix, which is represented by < in the DataObject ID of this example. Accordingly, the DataChunk ID for this particular data chunk would be 1234567890123456789012345-00005-<metadata> for a data object assigned the DataObject ID 1234567890123456789012345, with the actual metadata replacing the symbolic metadata in this example.

[0063] In various embodiments, storage controller 104 may write each data chunk of a data object received from a client device (e.g., client device 25) to a buffer memory 108 of a cache memory device 106. Storage controller 104, in some embodiments, may continue to write data chunks to a buffer memory 108 until a storage threshold of a cache memory device 106 has been met. For instance, the storage threshold may reflect that an aspect of a buffer memory 108 is full, which may mean that data chunks are written to the buffer memory 108 until the storage capacity of the buffer memory 108 has been met, a predetermined threshold of memory space in the buffer memory 108 is being utilized (e.g., eighty percent (80%)), among other example threshold values that may be greater than or

less than 80%), a pre-allocated storage capacity has been met for the storage of that type of data object, etc., among other examples.

[0064] A LRU technique may be included in various embodiments of storage controller 104 to reduce storage constraints in a buffer memory 108. For instance, storage controller 104 may begin writing data chunks stored in a cache memory device 106 (e.g., its buffer memory 108) for the longest amount of time (e.g., LRU or accessed) to one of storage devices 102 to free memory space (e.g., make available) in a buffer memory 108 so that the remaining data chunks in the data object can be written to a cache memory device 106.

[0065] Storage controller 104, in various embodiments, may write or continue writing data chunks to cache memory device(s) 106 as memory space is made available in their respective buffer memory 108 via the LRU technique. Writing data chunks (e.g., LRU data chunks) to one or more of the storage devices 102 to free memory space as memory space becomes available, in various embodiments, may be referred to as a push for each data chunk or pushing the data chunks (or the non-conforming inode) through cache memory device 106 to long-term storage (e.g., one or more of storage device 102a, 102b, ...102n). For instance, storage controller 104 may push a data chunk in the non-conforming inode through cache memory device 106 to long-term storage until the entire data object (i.e., each and every data chunk) is written to (e.g., buffered/stored in) storage system 100 (e.g., cache memory device 106 (e.g., buffer memory 108) and/or one or more of the storage devices 102). The pushing operations may be periodically (e.g., at predetermined intervals), substantially continuously, and/or continuously performed by storage controller 104 until there are no remaining data chunks in the non-conforming inode that are to be written to storage system 100.

[0066] In this manner, storage controller 104 may, in some cases, write data objects (e.g., non-conforming inodes) to storage system 100 by pushing smaller amounts of a data object (e.g., in the form of two or more data chunks) through a cache memory device 106 (e.g., buffer memory 108) to long-term storage (e.g., one or more of the storage devices 102) instead of writing the entire data object to storage system 100 at a single instance.

[0067] Writing (or pushing) data chunks in a data object through cache memory device(s) 106 to one or more of the storage devices 102 is more efficient than other approaches, such as those that write the entire data object to a storage buffer at a single point in time. For instance, pushing data chunks through to long-term memory is advantageously faster, uses less overhead, needs less bandwidth, etc. than writing an entire

LOB in a single operation, especially for non-conforming inodes whose size exceeds available buffer space, because writing smaller data chunks in a series of writes can allow a buffer memory 108 to have sufficient storage space over time to accommodate each individual data chunk. In other words, pushing through data chunks of a LOB may be considered an added time variable of some write operations of cache memory device 106 that can allow cache memory device 106 to more efficiently and/or reliably write and/or buffer LOBs.

[0068] In various embodiments, storage controller 104 may perform or execute I/O requests to read data objects from one or more storage devices 102 to cache memory device 106 (e.g., buffer memory 108). In some embodiments, an entire data object (e.g., a conforming inode) may be read to cache memory device 106 (e.g., buffer memory 108) when data within the non-conforming inode is read requested. In some embodiments, a subset (e.g., one or more data chunks, one or more portions, one or more partitions, etc.) of a data object (e.g., a conforming inode and/or a non-conforming inode) including the desired data may be read to cache memory device 106 from one or more of storage devices 102 when data of a conforming inode or a non-conforming inode is read requested. In various embodiments, storage controller 104 may identify the data chunk(s) stored in one or more of storage devices 102 including the desired data in a read request and reads the identified data chunk(s) from one or more storage device(s) 102 to cache memory device 106.

[0069] Similar to various embodiments discussed above, storage controller 104 may implement the LRU technique described herein, as needed, to free memory space in cache memory device 106 when buffer memory 108 is full when servicing a read request, which may mean that the storage capacity of buffer memory 108 has been met, a predetermined threshold of memory space in buffer memory 108 is being utilized (e.g., eighty percent (80%), among other example threshold values that may be greater than or less than 80%), a pre-allocated storage capacity has been met for the storage of a type of data object, etc. Storage controller 104 may recognize that buffer memory 108 is full and, in response to such recognition, begin removing/deleting one or more data objects and/or one or more data chunks from buffer memory 108 in accordance with the LRU technique.

[0070] In some embodiments, storage controller 104 may determine whether data object(s) and/or data chunk(s) that is/are subject to removal from buffer memory 108 of cache memory device 106 in accordance with the LRU technique are written to/stored in

long-term storage (e.g., one or more of storage devices 102a, 102b, ...102n). The storage controller 104 may write data objects and/or data chunks subject to removal from buffer memory 108 to storage devices 102 prior to the data object(s) and/or data chunk(s) being removed/deleted from buffer memory 108, which may further include the data object(s) and/or data chunk(s) being mirrored/replicated in another one or more of storage devices 102 consistent with the discussion below.

[0071] In some embodiments, storage controller 104 may utilize smaller amounts of data in the form of one or more data chunks to read or pull a subset of data (e.g., desired data) in a data object that is stored in one or more of storage devices 102 to buffer memory 108 instead of reading the entire data object from storage device(s) 102 to buffer memory 108. Reading (or pulling) a subset of a data object (e.g., one or more data chunks) from one or more of storage devices 102a, 102b, ...102n to buffer memory 108 may advantageously be more efficient than reading the entire data object from storage device(s) 102a, 102b, ...102n to buffer memory 108, especially when the desired data is located in a non-conforming inode, because reading a subset of a data object (e.g., a smaller amount/quantity of data in the form of one or more data chunks) to buffer memory 108 is faster, uses less buffer memory overhead, and/or needs less bandwidth to communicate/transfer data, among other advantages, than other storage solutions or services that attempt to read all of the data in a data object to a storage buffer or cache, as discussed in the Background.

[0072] In various embodiments, storage controller 104 includes functionality to perform or execute I/O requests to modify data objects (e.g., conforming inodes and/or non-conforming inodes) from one or more of storage devices 102 via cache memory device 106. In some embodiments, an entire data object may be read to cache memory device 106 and modified when data within the non-conforming inode is read-write requested.

[0073] The storage controller 104 may further re-write modified/re-written data object in cache memory devices 106 to one or more of storage devices 102, which may replace the original data object in storage device(s) 102 with the modified data object. The modified/re-written data object may then be mirrored/replicated in another one or more of storage devices 102 consistent with the discussion below.

[0074] In other embodiments, a subset of a data object including data desiring modification may be read to cache memory device(s) 106 from one or more of storage devices 102 and modified/re-written when data of a conforming inode or a non-conforming inode is read-write requested. In various embodiments, storage controller 104 may identify

the data chunk(s) stored in storage devices 102, including the data for which modification is desired in a modification request (e.g., read-write request). The identified data chunk(s) may be read (e.g., from storage device(s) 102 to cache memory device(s) 106 and subsequently modified/re-written.

5 [0075] Similar to various embodiments discussed above, storage controller 104 may implement the LRU technique, as needed, to free memory space in cache memory device 106 when buffer memory 108 is full. For instance, storage controller 104 may recognize that buffer memory 108 is full may begin removing one or more data objects and/or one or more data chunks from buffer memory 108 in accordance with the LRU technique, as
10 discussed elsewhere herein.

[0076] In some embodiments, storage controller 104 may determine whether the data object(s) and/or data chunk(s) that is/are subject to removal from cache memory devices 106 in accordance with the LRU technique are written to/stored in storage devices 102. For instance, storage controller 104 may write the data object(s) and/or data chunk(s) that is/are
15 subject to removal from a buffer memory 108 to storage device(s) 102 prior to the data object(s) and/or data chunk(s) being removed from the buffer memory 108, which may further include the data object(s) and/or data chunk(s) being mirrored/replicated in storage devices 102, as discussed elsewhere herein.

[0077] Storage controller 104, in some embodiments, may further write (or re-write)
20 modified/re-written data objects and/or data chunks to storage devices 102 subsequent to being modified/re-written in cache memory device(s) 106. In various embodiments, storage controller 104 may synchronously or asynchronously write/re-write the modified/re-written data objects and/or data chunks to storage devices 102, as discussed elsewhere herein.

[0078] In this manner, storage controller 104 may utilize one or more data chunks to
25 read-write (pull-push) a subset of data desiring modification in a data object that is stored in storage devices 102 instead of the technique in other storage solutions or services that read the entire data object and then write/re-write a subset or small portion of the data in the data object. As such, storage controller 104 may advantageously be faster, use less buffer memory overhead, needs less bandwidth, etc., among other advantages, than other storage
30 solutions or services, as discussed in the Background.

[0079] Storage controller 104 may include interface(s) for receiving requests to perform a process or method to manage storage operations in a storage system (e.g., storage system 100). In various embodiments, the storage operations are managed utilizing an application

program interface (API) for one or more of cache memory devices 106, which may be referred to herein as a cache API. The functionality of storage controller 104, in some embodiments, may provide multiple threads (e.g., a plurality of threads, two or more threads, etc.) or a thread pool configured to perform the I/O operations (e.g., write requests, read requests, and/or read-write requests) discussed above. As such, storage controller 104 may be referred to as a multi-threaded storage controller in some embodiments. In some embodiments, storage controller 104 may define the configuration (e.g., task and/or operation) of each thread and multiple threads may be started at the instantiation of cache memory device 106.

10 **[0080]** Storage controller 104 also includes functionality to create a cache entry object (e.g., a CacheEntry object) for each data object received at cache memory device 106 (e.g., received from a client (e.g., client devices 50)). The statement, CacheEntry* Create(...) is example code that may be utilized by storage controller 104 to create a CacheEntry object. The “...” in the parenthetical of this example may be replaced by a representation of a data object (e.g., a DataObject ID) in an actual cache entry object. For example, a cache entry object for a data object with a DataObject ID 1234567890123456789012345 may be created with the statement, CacheEntry* Create (DataObject ID 1234567890123456789012345) or similar coding logic that assigns a cache entry object to a data object implemented in a different computing language. As such, each data object is assigned at least a CacheEntry ID and a DataObject ID (e.g., is assigned at least two IDs).

20 **[0081]** Storage controller 104 may assign a particular thread to a particular cache entry object using a round-robin technique, among other example assigning techniques or methods. Further, storage controller 104 may assign a different thread to manage each respective cache entry object such that there is a 1:1 ratio of threads to cache entry objects.

25 **[0082]** A thread functions to manage I/O operations for its assigned cache entry object. That is, a thread may receive I/O requests and, in response thereto, perform the I/O operations in the I/O requests for its assigned cache entry object. A thread may receive write requests, read requests, and read-write requests at its assigned cache entry object and, in response thereto, execute write operations, read operations, and read-write operations, respectively, on the data object coinciding with its cache entry object that are consistent with the respective write requests, read requests, and read-write requests. Examples of a write request, a read request, and a read-write request (e.g., I/O requests) can be, respectively, as follows:

Write(CacheEntry*, (...));
 Read(CacheEntry*, (...)); and
 Read-write(CacheEntry*, (...)),

in which the “...” in these example I/O operations represent the ID of an actual cache entry object (data object).

[0083] A thread may transmit an indication or acknowledgement that the I/O operation has been performed or is completed. Non-limiting examples of an indication or acknowledgement that an I/O operation has been performed or is complete are: a WriteDone acknowledgement, a ReadDone acknowledgement, and a Read-writeDone acknowledgement for a write operation, a read operation, and a read-write operation, respectively.

[0084] A thread can manage the data chunks corresponding to or associated with its assigned cache entry object (or data object), if any. That is, the same thread that manages a particular cache entry object (or data object) also manages any data chunks that are associated with or correspond to the particular data object. For example, the thread assigned to manage a cache entry object for a data object divided into 10,000 data chunks (e.g., data chunks 00001-10000) and includes the DataObject ID 1234567890123456789012345 also manages the data chunks corresponding to the data object (e.g., the data chunks assigned DataChunk IDs 1234567890123456789012345-00001-**<metadata>** through 1234567890123456789012345-10000-**<metadata>**). In this manner, storage controller 104 and/or the multiple threads are able to avoid or at least reduce one or more concurrency penalties that may occur when the threads are performing their various respective operations or tasks.

[0085] In some embodiments, storage controller 104 may utilize pointers to the cache entry objects to post I/O operations for the threads assigned to the cache entry objects. A pointer may be used to post each write request, read request, and read-write request to the thread assigned to a particular cache entry object so that the thread can, in response thereto, perform each write operation, read operation, or read-write operation consistent with the various discussions above.

[0086] In various embodiments, storage controller 104 may be configured for redundancy to provide that data objects and/or data chunks may be stored in two or more locations in cache memory devices 106. The data objects and/or data chunks may be synchronously or asynchronously copied across one or more of cache memory devices 106

using any suitable redundancy and/or replication technique/process that is known or developed in the future. Non-limiting examples of such replication techniques may include or be referred to as: data duplication, copying data, data replication, backing-up data, and data mirroring, among other example replication techniques. The data objects and/or data
5 chunks may be replicated in the same cache memory device 106 and/or across two or more cache memory devices 106.

[0087] Storage controller 104 may be further configured for replicating the data objects and/or data chunks across one or more of storage devices 102a, 102b, ...102n. The data objects and/or data chunks may be synchronously or asynchronously copied to one or more
10 of storage devices 102 using any suitable technique that is known or developed in the future, as discussed above. The data objects and/or data chunks may be replicated in the same storage device 102 and/or across two or more storage devices 102.

[0088] Figure 2 is a block diagram of an example storage network 20 including, among other components, one or more client devices 25 and an enterprise storage service 50
15 capable of being coupled to and in communication with a storage system 200 via a wired and/or wireless network 75. Each client device 25 can be any computing hardware and/or software (e.g., a thick client, a thin client, or hybrid thereof) capable of accessing storage system 200 utilizing network 75. A client device 25, as part of its respective operation, may send I/O requests to storage system 200. For instance, a client device 25 can transmit I/O
20 requests to read, write store, communicate, propagate, and/or transport instructions, data, computer programs, software, code, routines, etc., to storage system 200. Client device(s) 25 and storage system 200 may comprise at least a portion of a client-server model. In general, storage system 200 may be accessed by client device(s) 25 and/or communication with storage system 200 may be initiated by client device(s) 25 through a network socket
25 utilizing one or more inter-process networking techniques.

[0089] Enterprise storage service 50 may include any suitable storage service and/or storage network that is known or developed in the future. An example enterprise storage service 50 is the Simple Storage Service (S3) provided by Amazon® Web Services, among other examples of an enterprise storage services. In various embodiments, enterprise
30 storage service 50 may provide a storage system that stores data objects, which can be referred to an object storage system. Enterprise storage service 50 may, in some embodiments, store data objects in a flat organization (e.g., non-hierarchical) of containers or buckets and may use unique IDs or keys to retrieve the data objects, although other

variations are also possible and contemplated. Enterprise storage service 50, in various embodiments, may include any number of storage devices. The storage devices may, in some embodiments, use metadata to store and access data objects and store the metadata with its corresponding data object similar to storage devices 102a, 102b, ...102n discussed
5 above. Accordingly, enterprise storage service 50 may function and/or operate in a manner similar to storage devices 102 discussed above and may be applicable to network 20 and/or storage system 200.

[0090] As shown, storage system 200 may comprise, among other components, a storage controller 204 including one or more cache memory devices 206 similar to storage
10 controller 104 and cache memory device(s) 106, respectively, in storage system 100, although other suitable configurations are also applicable and contemplated. Storage controller 204 and cache memory device(s) 206 are capable of being coupled to and in communication with client device(s) 25 and enterprise storage service 50 via network 75. In various embodiments, storage controller 204 may be coupled between client device(s) 25
15 and cache memory device(s) 206.

[0091] Network 75, in various embodiments, may comprise any type of computing network(s) that are known or developed in the future that may allow client devices 25, enterprise storage service 50, and storage system 200 to communicate and/or share resources with one another. Non-limiting examples of a network include, but are not
20 limited to, a cloud network (IAN), a SAN (e.g., a storage area network, a small area network, a server area network, and/or a system area network), a wide area network (WAN), a local area network (LAN), a wireless local area network (WLAN), a metropolitan area network (MAN), an enterprise private network (EPN), a virtual private network (VPN), and/or a personal area network (PAN), among other examples of computing networks
25 and/or or sets of computing devices connected together for the purpose of sharing resources.

[0092] The various embodiments of network 20 may include hardware and/or software that allow network 20 to operate in a manner similar to network 10 discussed above with reference to Figures 1A and 1B. Additionally, in network 20 a third-party storage service (e.g., enterprise storage service 50 may provide long-term storage operations and data
30 replication of data objects and/or data chunks. Further implementations may include a combination of network 20 and network 10, and/or aspects thereof, such as a network including one or more third-party storage services and one or more local or remote storage

devices, each providing long-term storage operations and data replication of data objects and/or data chunks.

5 [0093] In some embodiments, in network 10 and/or network 20, one or more of client devices 25 may be unaware of whether it/they are a part of network 10 or network 20. That is, the operations of storage system 100 and/or 200 may be unknown to or not seen by (transparent to) one or more of client devices 25, although in further implementations, some or all of client devices 25 may be aware of their membership in network 10 and/or network 20.

10 [0094] Figures 3A through 3C are a flow diagram of example operations for a storage controller (e.g., storage controller 104 and/or 204) to manage storage in accordance with various embodiments. At least in the illustrated embodiment, the operations 301 (e.g., write operations) may begin at time 1 (T1) with the creation of a data object ID and/or corresponding data chunk IDs (e.g., by storage controller 104 or 204 (e.g., via a cache API)) for a data object received (and divided/partitioned) at a storage system (e.g., storage system 15 100 or 200), which data object ID and/or data chunk IDs may be assigned in a manner consistent with the various discussions above. A new cache entry object (e.g., CacheEntry) corresponding to the data object and/or data chunks may be created (e.g., by storage controller 104 or 204 (e.g., via the cache API)) at time 2 (T2) in accordance with the various discussions above.

20 [0095] At time 3 (T3), a thread may be assigned (e.g., by storage controller 104 or 204 (e.g., via a cache API)) to the cache entry object. The assignment of the thread to the cache entry object may be the result of an assigning technique (e.g., a round-robin assigning technique, among other example assigning techniques). The thread functions to manage I/O operations for the cache entry object. That is, the thread can receive I/O requests (e.g., write requests, read requests, and read-write requests) and, in response to the I/O requests, may perform the I/O operations (e.g., write operations, read operations, and read-write operations) in the I/O requests.

30 [0096] At time 4 (T4), an I/O request (e.g., write request in this example) may be received (e.g., by storage controller 104 or 204 (e.g., via a cache API)) for a particular cache entry object (data object). The I/O request (e.g., the write request in this example) may be directed (e.g., via a pointer) to the thread assigned to the particular cache entry object (data object) and the assigned thread may perform the I/O operations (e.g., write operations in this example) on the cache entry object at time 5 (T5).

[0097] Upon completion of the I/O operations (e.g., write operations in this example), the thread may acknowledge (e.g., a WriteDone acknowledgement in this example) completion of the I/O request (e.g., write request in this example) to the storage controller (e.g., storage controller 104 or 204 (e.g., via a cache API)) at time 6 (T6).

5 [0098] In addition, the operations at time 1 through time 6 may be repeated for an I/O request received at the storage controller (e.g., storage controller 104 or 204) for each different (e.g., new) cache entry object (e.g., a different data object and/or different corresponding data chunks). The operations at time 4 through time 6 may be repeated for each I/O request received at the storage controller (e.g., storage controller 104 or 204) for the different cache entry object (e.g., different data object and/or different corresponding data chunks).

10 [0099] At time 4 (T4), an I/O request (e.g., write request in this example) may be received (e.g., by storage controller 104 or 204 (e.g., via a cache API)) for a particular cache entry object (data object). The I/O request (e.g., the write request in this example) may be directed (e.g., via a pointer) to the thread assigned to the particular cache entry object (data object) and the assigned thread may perform I/O operations (e.g., write operations in this example) on the cache entry object at time 5 (T5). Upon completion of the I/O operations (e.g., write operations in this example), the thread may acknowledge (e.g., a WriteDone acknowledgement in this example) completion of the I/O request (e.g., write request in this example) to the storage controller (e.g., storage controller 104 or 204 (e.g., via a cache API)) at time 6 (T6).

15 [00100] In addition, the operations at times 1 through 6 may be repeated for an I/O request received at the storage controller (e.g., storage controller 104 or 204) for each different (e.g., new) cache entry object (e.g., a different data object and/or different corresponding data chunks). Moreover, the operations at time 4 through time 6 may be repeated for each I/O request received at the storage controller (e.g., storage controller 104 or 204) for the different cache entry object (e.g., different data object and/or different corresponding data chunks).

20 [00101] In Figure 3B, the operations 302 (e.g., read operations) may begin at time 1 (T1) with receipt (e.g., by storage controller 104 or 204 (e.g., via a cache API)) of an I/O request (e.g., read request in this example) for a particular cache entry object (data object). The I/O request (e.g., the read request in this example) may be directed (e.g., via a pointer) to the thread assigned to the particular cache entry object (data object) and the assigned thread

may perform the I/O operations (e.g., read operations in this example) on the cache entry object at time 2 (T2).

[00102] Upon completion of the I/O operations (e.g., read operations in this example), the thread may acknowledge (e.g., a ReadDone acknowledgement in this example)

5 completion of the I/O request (e.g., read request in this example) to the storage controller (e.g., storage controller 104 or 204 (e.g., via a cache API)) at time 3 (T3). The operations at times 1 through 3 may be repeated for each read request received at the storage controller (e.g., storage controller 104 or 204) subsequent to time 3.

[00103] In Figure 3C, the operations 303 (e.g., read-write operations) may begin at time 10 1 (T1) with receipt (e.g., by storage controller 104 or 204 (e.g., via a cache API)) of an I/O request (e.g., read-write request in this example) to modify data for a particular cache entry object (data object). The I/O request (e.g., the read-write request in this example) may be directed (e.g., via a pointer) to the thread assigned to the particular cache entry object (data object) and the assigned thread may perform the I/O operations (e.g., read-write operations 15 in this example) on the cache entry object at time 2 (T2).

[00104] Upon completion of the I/O operations (e.g., read-write operations in this example), the thread may acknowledge (e.g., a Read-writeDone acknowledgement in this example) completion of the I/O request (e.g., read-write request in this example) to the storage controller (e.g., storage controller 104 or 204 (e.g., via a cache API)) at time 3 (T3).

20 The operations at times 1 through 3 may be repeated for each read-write request received at the storage controller (e.g., storage controller 104 or 204) subsequent to time 3.

[00105] In the examples of Figures 3A through 3C, upon completion of the I/O operations (e.g., write operations, read operations, read-write operations) discussed in each respective example, subsequent I/O requests may be performed by the threads (e.g., threads 25 1, 2, through n) consistent with Figures 3A through 3C. That is, subsequent to a thread acknowledgement (e.g., a WriteDone acknowledgement, a ReadDone acknowledgement, or a Read-writeDone acknowledgement) a next I/O request, which may be the same type of I/O request or a different type of I/O request, may be received by the storage controller and serviced by one of threads 1, 2, through n.

30 **[00106]** With reference to Figure 4, a diagram 400 representing example storage operations of a storage system (e.g., storage system 100 or 200) is illustrated. As shown, a newly received data object may be provided with an inode (e.g., a non-conforming inode or a conforming inode) that may be assigned (consistent with the various embodiments

discussed above) with a DataObject ID at time 1 (T1). Time 1 may further include dividing each data object into multiple data chunks (e.g., a quantity of data chunks in the range of 1 data chunk to 10,000 data chunks, among other example ranges and/or quantities), each being assigned (consistent with the various embodiments discussed above) a respective data
5 chunk ID. The inodes may be organized into a filetree including a hierarchical order (e.g., a multi-tiered or multi-level system (e.g., three levels) by the storage controller 104/204, consistent with the various embodiments discussed above).

[00107] At time 2 (T2), a cache entry object may be created (consistent with the various embodiments discussed above) for each inode in a cache memory device 106/206 such that
10 each inode is associated with a corresponding cache entry object. In other words, each cache entry object may correspond to a different data object and its corresponding data chunks, if any.

[00108] Each data chunk may be written, at time 3 (T3), to a buffer memory 108/208 in the cache memory device 106/206 and subsequently written (e.g., stored), at time 4 (T4), in
15 one or more storage devices and/or an enterprise storage service. In some embodiments, the write data chunks may be continuously transmitted to the buffer memory 108/208 until the buffer memory 108/208 is full or substantially full. After the buffer memory 108/208 is full, a LRU technique may be utilized to free space or make space in the buffer memory 108/208 so that the data chunks can continue being written to the buffer memory 108/208.
20 That is, data chunks stored in the buffer memory 108/208 for the longest amount of time may be written to long-term storage and/or enterprise storage service 75, deleted from the buffer memory 108/208, and replaced in the buffer memory 108/208 by the remaining data chunk(s). The LRU technique may, in some cases, refrain from writing (or re-writing) data chunks to long-term storage prior to deleting/removing the data chunk(s) from the buffer
25 memory 108/208 if the data chunk(s) are already written/re-written to (e.g., stored in) long-term storage (e.g., storage device(s) 102, enterprise storage service 50, etc.). In this manner, the data chunks may be pushed through the storage system (e.g., cache memory device) for long-term storage.

[00109] Data chunks may be read from long-term storage to the buffer memory 108/208
30 at time 5 (T5). In some embodiments, the write data chunks may be continuously read to the buffer memory until the buffer memory 108/208 is full. After the buffer memory 108/208 is full, the LRU technique may be utilized to free space or make space in the buffer

memory 108/208 so that the data chunks can continue being read to the buffer memory 108/208. In this manner, the data chunks may be pulled from long-term storage.

[00110] At time 6 (T6), one or more data chunks that are read from long-term storage at time 5 may be modified and subsequently re-written to the cache memory device 106/206 (e.g., buffer memory 108/208) consistent with a read-write operation. The re-written data chunk(s) may be written (re-written) to long-term storage at time 7 (T7). In a situation in which a re-written data chunk in the buffer memory becomes a LRU data chunk and is not yet written (re-written) to long-term storage, the re-written data chunk may be written (re-written) to long-term storage prior to being deleted/removed from the buffer memory 108/208.

[00111] At various times (represented as time ! (e.g., T!)), the data objects and/or data chunks in long-term storage may be replicated or mirrored in a different long-term storage location. The data object(s) and/or data chunk(s) may be synchronously or asynchronously replicated/mirrored as the data object(s) and/or data chunk(s) are written (re-written) to long-term storage. Accordingly, T! may represent any time equal to or subsequent to time 4, including a time subsequent to time 7.

[00112] Referring now to Figures 5A through 5G, Figure 5A is a block diagram illustrating an Example 500 representing multiple data chunks (e.g., data chunks 510a through 510d) shown in initial positions (e.g., positions 512a through 512d) of a LRU technique. This embodiment includes four data chunks (e.g., data chunks 510a through 510d, which may be collectively referred to as data chunks 510), although further embodiments may include other quantities of data chunks 510 (e.g., two or more data chunks). In Example 500, a left-most position 512a (e.g., the position initially occupied by data chunk 510a) is the most recently used (MRU) position and a right-most position 512d (e.g., the position initially occupied by data chunk 510d) is the LRU position, with positions 512b and 512c being intermediate positions (e.g., LRU+2 and LRU+1, respectively) and initially occupied by data chunk 510b and data chunk 510c, respectively).

[00113] In various embodiments, each time one of data chunks 510 is accessed (e.g., subject to an I/O operation (e.g., a write operation, a read operation, or a read-write operation), the accessed data chunk may move to position 512a (e.g., the MRU position) and the remaining data chunks may move, as needed, one position to the right toward the LRU position (e.g., position 512d). The following examples in Figures 5B through 5G are made with continued reference to Figure 5A and may provide greater insight and

understanding to the various LRU techniques discussed herein; however, the various LRU techniques discussed herein are not limited to the examples in Figures 5A through 5G.

[00114] Example 501 in Figure 5B: Data chunk 510b is accessed, which results in data chunk 510b being moved from position 512b to position 512a, data chunk 510a being
 5 moved from position 512a to 512b, and data chunks 510c and 510d remaining in positions 512c and 512d, respectively.

[00115] Example 502 in Figure 5C: Data chunk 510c is accessed, which results in data chunk 510c being moved from position 512c to position 512a, data chunk 510a being
 10 moved from position 512a to position 512b, data chunk 510b being moved from position 512b to position 512c, and data chunk 510d remaining in position 512d.

[00116] Example 503 in Figure 5D: Data chunk 510d is accessed, which results in data chunk 510d being moved from position 512d to position 512a, data chunk 510a being
 moved from position 512a to 512b, data chunk 510b being moved from position 512b to position 512c, and data chunk 510c being moved from position 512c to position 512d.

[00117] Example 504 in Figure 5E: Data chunk 510a is accessed, which results in data chunks 510a, 510b, 510c, and 510d remaining in positions 510, 512b, 512c, and 512d,
 15 respectively.

[00118] Example 505 in Figure 5F: Data chunk 510c is accessed, which results in data chunk 510c being moved from position 512c to position 512a, data chunk 510a being
 20 moved from position 512a to 512b, data chunk 510b being moved from position 512b to position 512c, and data chunk 510d remaining in position 512d. Subsequently, data chunk 510d is accessed, which results in data chunk 510d being moved from position 512d to position 512a, data chunk 510c being moved from position 512a to position 512b, data chunk 510a being moved from position 512b to position 512c, and data chunk 510b being
 25 moved from position 512c to position 512d.

[00119] Example 506 in Figure 5E: When a new data chunk (e.g., data chunk 510e) is accessed, data chunk 510e replaces the data chunk in the MRU position (data chunk 510a in position 512a in the example illustrated by diagram 500 of Figure 5) and the data chunk in the LRU position (e.g., data chunk 510d in position 512d in Figure 5) being deleted or
 30 removed from the buffer memory (e.g., buffer memory 108). Furthermore, the other data chunks (e.g., data chunks 510a, 510b, and 510c in Figure 5) are each moved one position to the right or one position closer to the LRU position (position 512d in Figure 5). Specifically, data chunk 510a is moved from position 512a to 512b, data chunk 510b is

moved from position 512b to position 512c, and data chunk 510c is moved from position 512c to position 512d.

[00120] Example 507 in Figure 5F: A subsequent newly accessed data chunk (e.g., data chunk 510f) would result in data chunk 510f being moved to the MRU position (e.g., position 512a), data chunk 510c being deleted/removed, and data chunks 510e, 510a, and 510b being moved to positions 512b, 512c, and 512d, respectively. Figure 5G is a continuing illustrative example of Figure 5F in which data chunk 510b is accessed subsequent to data chunk 510f, data chunk 510b is moved from the LRU position (e.g., position 512d) to the MRU position (e.g., position 512a), data chunk 510f is moved to position 512b, data chunk 510e is moved to position 512c, and data chunk 510a is moved to the LRU position (e.g., position 512d).

[00121] With reference now to Figure 6, an example method 600 of writing a data object to a buffer memory 108/208 of a cache memory device 106/206 is illustrated. As shown, method 600 may begin by receiving an I/O request (e.g., a write request), by a storage controller 104/204, to store a data object (block 602) and may determine the size of the data object (block 604).

[00122] The data object size may be compared to a storage threshold and/or the size of a buffer memory 108/208 to determine if the data object size is larger than the storage threshold and/or buffer size (block 606). If/when the data object size is smaller than the storage threshold and/or buffer size (e.g., a NO in block 606), the data object may be determined to be a conforming data object and may be written to buffer memory 108/208, in its entirety (block 608). The conforming data object may be subsequently written to one or more of storage devices 102 and/or enterprise storage service 50 (block 620).

[00123] If/when the data object size is larger than the buffer size (e.g., a YES in block 606), the data object may be determined to be a non-conforming data object and may be divided (partitioned) into a plurality of data chunks (block 610). The data object may be divided into any desired number of data chunks and each data chunk may include any size suitable for storage in buffer memory 108/208, one or more of storage devices 102, and/or enterprise storage service 50, as discussed elsewhere herein.

[00124] Method 600 may further include determining if buffer memory 108/208 is available to store (buffer) a data chunk (block 612). Buffer memory 108/208 may be considered available if/when the buffer memory is not full; otherwise, buffer memory 108/208 may be considered unavailable.

[00125] If/when buffer memory 108/208 is unavailable (e.g., a NO in block 612), one or more data objects and/or one or more data chunks buffered in buffer memory 108/208 may be written to one or more of storage devices 102 and/or enterprise storage service 50 to free space in buffer memory 108/208 (block 614) and the availability of buffer memory 108/208 may again be determined (block 612). In various embodiments, the space in buffer memory 108/208 may be freed-up by pushing one or more data chunks through buffer memory 108/208 utilizing a LRU technique, as discussed elsewhere herein. If/when buffer memory 108/208 is available (e.g., a YES in block 612), one or more data chunks may be written to buffer memory 108/208 (block 616).

10 [00126] Method 600 may further include determining if/when there are one or more remaining data chunks in the data object to be written to buffer memory 108/208 (block 618). If/when there is at least one remaining data chunk to be written to buffer memory 108/208 (e.g., a YES in block 618), the availability of buffer memory 108/208 may be determined (block 612) and blocks 612 through 618 may be repeated until there are no more remaining data chunks to be written to buffer memory 108/208.

15 [00127] If/when there are no more remaining data chunks to be written to buffer memory 108/208 (e.g., a NO in block 618), a data chunk that is not stored in storage devices 102 or enterprise storage service 50 may be written to storage devices 102 and/or enterprise storage service 50 (block 620). In some embodiments, any remaining data chunk(s) may be written to storage devices 102 or enterprise storage service 50 immediately after or within a short period of time of being modified, while in other embodiments any remaining data chunk(s) may be written to storage devices 102 or enterprise storage service 50 at or subsequent to a predetermined time of being modified. Some remaining data chunk(s) may be written to storage devices 102 and/or enterprise storage service 50 in accordance with a LRU technique.

25 [00128] Referring now to Figure 7, an example method 700 of reading a data object in storage devices 102 or enterprise storage service 50 to a buffer memory 108/208 of a cache memory device 106/206 is illustrated. As shown, method 700 may begin by receiving an I/O request (e.g., a read request), by a storage controller 104/ 204, to read data in a data object (block 702) and determining one or more data chunks in storage devices 102a, 102b, ...102n and/or enterprise data service 50 including the data (block 704). In some embodiments, the data chunks including the data desiring to be read may be stored or located in the same or different storage devices 102, as discussed elsewhere herein.

[00129] Method 700 may further include determining if buffer memory 108/208 into which the data chunk(s) are to be read is available to store (buffer) the data chunk(s) (block 706). Buffer memory 108/208 may be considered available if/when the buffer memory is not full; otherwise, buffer memory 108/208 may be considered unavailable.

5 [00130] If/when buffer memory 108/208 is unavailable (e.g., a NO in block 706), one or more data objects and/or one or more data chunks buffered in buffer memory 108/208 may be deleted from buffer memory 108/208 to free space in buffer memory 108/208 (block 708) and the availability of buffer memory 108/208 may again be determined (block 706). In various embodiments, the space in buffer memory 108/208 may be freed-up by deleting
10 one or more data objects and/or data chunks from buffer memory 108/208 utilizing a LRU technique, as discussed elsewhere herein. Deleting one or more data objects and/or data chunks from buffer memory 108/208 may further include writing any data object(s) and/or data chunk(s) to storage device(s) 102 and/or enterprise storage service 50 that are not stored in or have not previously been written to storage devices 102 and/or enterprise
15 storage service 50 prior to removal.

[00131] If/when buffer memory 108/208 is available (e.g., a YES in block 612), one or more data chunks including at least a portion of the desired data may be read to buffer memory 108/208 (block 710). Method 700 may further include determining if/when there are one or more remaining data chunks including at least a portion of the desired data to be
20 read to buffer memory 108/208 (block 712).

[00132] If/when there is at least one remaining data chunk to be read to buffer memory 108/208 (e.g., a YES in block 712), the availability of buffer memory 108/208 may be determined (block 706) and blocks 706 through 712 may be repeated until there are no more remaining data chunks including desired data to be read to buffer memory 108/208 (e.g., a
25 NO in block 712). Method 700 may further include waiting for a subsequent I/O request (block 714) and may repeat block 702 through block 714, as needed and/or desired.

[00133] With reference now to Figure 8, an example method 800 of modifying data of a data object in storage device(s) 102 or enterprise storage service 50 is illustrated. As shown, method 800 may begin by receiving an I/O request (e.g., a read-write request), by a
30 storage controller 104/204, to modify data in a data object (block 802) and determining one or more data chunks in storage device(s) 102 and/or enterprise data service 50 including the desired data (block 804). In some embodiments, the data chunks including the data desiring

modification may be stored or located in the same or different storage devices 102, as discussed elsewhere herein.

5 **[00134]** Method 800 may further include determining if a buffer memory 108/208 of a cache memory device 106/206 into which the data chunk(s) including the desired data are to be read is available to store (buffer) the data chunk(s) (block 806). Buffer memory 108/208 may be considered available if/when the buffer memory is not full; otherwise, buffer memory 108/208 may be considered unavailable.

10 **[00135]** If/when buffer memory 108/208 is unavailable (e.g., a NO in block 806), one or more data objects and/or one or more data chunks buffered in buffer memory 108/208 may be removed from buffer memory 108/208 to free space in buffer memory 108/208 (block 808) and the availability of buffer memory 108/208 may again be determined (block 806). In various embodiments, the space in buffer memory 108/208 may be freed-up by deleting one or more data objects and/or data chunks from buffer memory 108/208 utilizing a LRU technique, as discussed elsewhere herein. Deleting one or more data objects and/or data
15 chunks from buffer memory 108/208 may further include writing any data object(s) and/or data chunk(s) to storage device(s) 102 and/or enterprise data service 50 that are not stored in or have not previously been written to storage devices 102 and/or enterprise data service 50 prior to removal.

20 **[00136]** If/when buffer memory 108/208 is available (e.g., a YES in block 612), one or more data chunks including at least a portion of the data desiring modification may be read to buffer memory 108/208 (block 810). The data chunk(s) read to buffer memory 108/208 may be subsequently modified (e.g., data may be deleted from and/or new data may be written to the data chunk(s)) (block 812).

25 **[00137]** Method 800 may further include determining if there are one or more remaining data chunks including at least a portion of the data desiring modification (block 814). If/when there is at least one remaining data chunk to be modified (e.g., a YES in block 814), the availability of buffer memory 108/208 may again be determined (block 806) and blocks 806 through 814 may be repeated until there are no more remaining data chunks including data desiring modification (e.g., a NO in block 814).

30 **[00138]** The modified one or more data chunks may be eventually written (re-written) to storage devices 102 and/or enterprise data service 50 (block 816). In some embodiments, the modified data chunk(s) may be written to storage device(s) 102 and/or enterprise data service 50 immediately after or within a short period of time of being modified, while in

other embodiments the modified data chunk(s) may be written to storage devices 102 and/or enterprise data service 50 at or subsequent to a predetermined time of being modified.

Some modified data chunk(s) may be written to storage device(s) 102 and/or enterprise data service 50 in accordance with a LRU technique, as discussed elsewhere herein.

5 **[00139]** To review, the various networks, systems, and methods discussed above may include hardware and/or software to increase reliability and/or decrease latency issues associated with writing data objects, and particularly LOBs, by dividing a data object into smaller data chunks and writing each data chunks to buffer memory 108/208 in succession. That is, the various embodiments discussed herein may be characterized as performing a
10 plurality of write operations on a LOB over a period of time.

[00140] In addition, the various networks, systems, and methods discussed above may include hardware and/or software to decrease latency issues associated with reading data in data objects, and particularly LOBs, by reading one or more smaller data chunks from storage devices 102 and/or enterprise data service 50 to buffer memory 108/208 in
15 succession. That is, the various embodiments discussed herein may be characterized as performing a plurality of reads operations on a LOB over a period of time.

[00141] Furthermore, the various networks, systems, and methods discussed above may include hardware and/or software to decrease latency issues associated with modifying data in data objects, and particularly LOBs, by reading one or more smaller data chunks from
20 storage devices 102 and/or enterprise data service 50 to buffer memory 108/208 in succession, modifying the data chunk(s), and writing/re-writing the modified data chunk(s) to storage device(s) 102 and/or enterprise data service 50. That is, the various embodiments discussed herein may be characterized as performing multiple read-write operations on a LOB over a period of time.

25 **[00142]** While at least one exemplary embodiment has been presented in the foregoing detailed description of the technology, it should be appreciated that a vast number of variations may exist. It should also be appreciated that an exemplary embodiment or exemplary embodiments are examples, and are not intended to limit the scope, applicability, or configuration of the technology in any way. Rather, the foregoing detailed description
30 will provide those skilled in the art with a convenient road map for implementing an exemplary embodiment of the technology, it being understood that various modifications may be made in a function and/or arrangement of elements described in an exemplary

embodiment without departing from the scope of the technology, as set forth in the appended claims and their legal equivalents.

[00143] As will be appreciated by one of ordinary skill in the art, various aspects of the present technology may be embodied as a system, method, or computer program product.

5 Accordingly, some aspects of the present technology may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.), or a combination of hardware and software aspects that may all generally be referred to herein as a circuit, module, system, and/or network. Furthermore, various aspects of the present technology may take the form of a computer program product
10 embodied in one or more computer-readable mediums including computer-readable program code embodied thereon.

[00144] Any combination of one or more computer-readable mediums may be utilized. A computer-readable medium may be a computer-readable signal medium or a physical computer-readable storage medium. A physical computer readable storage medium may be,
15 for example, but not limited to, an electronic, magnetic, optical, crystal, polymer, electromagnetic, infrared, or semiconductor system, apparatus, or device, etc., or any suitable combination of the foregoing. Non-limiting examples of a physical computer-readable storage medium may include, but are not limited to, an electrical connection including one or more wires, a portable computer diskette, a hard disk, RAM, ROM, an
20 EPROM, an EEPROM, a Flash memory, an optical fiber, a CD-ROM, an optical processor, a magnetic processor, etc., or any suitable combination of the foregoing. In the context of this document, a computer-readable storage medium may be any tangible medium that can contain or store a program or data for use by or in connection with an instruction execution system, apparatus, and/or device.

25 **[00145]** Computer code embodied on a computer-readable medium may be transmitted using any appropriate medium, including but not limited to, wireless, wired, optical fiber cable, radio frequency (RF), etc., or any suitable combination of the foregoing. Computer code for carrying out operations for aspects of the present technology may be written in any static language, such as the C programming language or other similar programming
30 language. The computer code may execute entirely on a user's computing device, partly on a user's computing device, as a stand-alone software package, partly on a user's computing device and partly on a remote computing device, or entirely on the remote computing device or a server. In the latter scenario, a remote computing device may be connected to a

user's computing device through any type of network, or communication system, including, but not limited to, a local area network (LAN) or a wide area network (WAN), Converged Network, or the connection may be made to an external computer (e.g., through the Internet using an Internet Service Provider).

5 [00146] Various aspects of the present technology may be described above with reference to flowchart illustrations and/or block diagrams of methods, apparatus, systems, and computer program products. It will be understood that each block of a flowchart illustration and/or a block diagram, and combinations of blocks in a flowchart illustration and/or block diagram, can be implemented by computer program instructions. These
10 computer program instructions may be provided to a processing device (processor) of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which can execute via the processing device or other programmable data processing apparatus, create means for implementing the operations/acts specified in a flowchart and/or block(s) of a block
15 diagram.

[00147] Some computer program instructions may also be stored in a computer-readable medium that can direct a computer, other programmable data processing apparatus, or other device(s) to operate in a particular manner, such that the instructions stored in a computer-readable medium to produce an article of manufacture including instructions that implement
20 the operation/act specified in a flowchart and/or block(s) of a block diagram. Some computer program instructions may also be loaded onto a computing device, other programmable data processing apparatus, or other device(s) to cause a series of operational steps to be performed on the computing device, other programmable apparatus or other device(s) to produce a computer-implemented process such that the instructions executed by
25 the computer or other programmable apparatus provide one or more processes for implementing the operation(s)/act(s) specified in a flowchart and/or block(s) of a block diagram.

[00148] A flowchart and/or block diagram in the above figures may illustrate an architecture, functionality, and/or operation of possible implementations of apparatus,
30 systems, methods, and/or computer program products according to various aspects of the present technology. In this regard, a block in a flowchart or block diagram may represent a module, segment, or portion of code, which may comprise one or more executable instructions for implementing one or more specified logical functions. It should also be

noted that, in some alternative aspects, some functions noted in a block may occur out of an order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or blocks may at times be executed in a reverse order, depending upon the operations involved. It will also be noted that a block of a block

5 diagram and/or flowchart illustration or a combination of blocks in a block diagram and/or flowchart illustration, can be implemented by special purpose hardware-based systems that may perform one or more specified operations or acts, or combinations of special purpose hardware and computer instructions.

[00149] While one or more aspects of the present technology have been illustrated and
10 discussed in detail, one of ordinary skill in the art will appreciate that modifications and/or adaptations to the various aspects may be made without departing from the scope of the present technology, as set forth in the following claims.

CLAIMS

1. A non-volatile memory system, comprising:
 - a first non-volatile cache memory device including a storage threshold; and
 - a storage controller coupled to the first non-volatile cache memory device, the storage controller including one or more processing devices for executing instructions in one or more applications configured to cause the storage controller to perform operations comprising:
 - receiving an input/output (I/O) request for a data object,
 - determining a size of the data object,
 - determining the data object to be a non-conforming data object relative to the storage threshold of the first non-volatile cache memory device,
 - dividing the non-conforming data object into a plurality of data chunks, each of the data chunks having a size that is less than the storage threshold,
 - moving one or more data chunks stored in one or more memory spaces of the first non-volatile cache memory device to a long-term memory device to free up the one or more memory spaces, and
 - storing one or more data chunks of the non-conforming data object in the freed-up one or more memory spaces of the first non-volatile cache memory device.
2. The system of claim 1, wherein:
 - the I/O request is a write request; and
 - the storage controller, when performing the operations for storing the one or more data chunks of the non-conforming data object, further performs operations comprising writing each of the one or more data chunks of the non-conforming data object to the freed-up one or more memory spaces of the first non-volatile cache memory device.
3. The system of claim 1, wherein:
 - the I/O request is a read request; and
 - the storage controller, when performing the operations for storing the one or more data chunks of the non-conforming data object, further performs operations comprising reading one or more data chunks stored in the long-term memory device to the first non-volatile cache memory device.

4. The system of claim 1, wherein:
 - the I/O request is a read-write request; and
 - the storage controller, when performing the operations for storing the one or more data chunks of the non-conforming data object, further performs operations comprising:
 - reading one or more data chunks stored in the long-term memory device to the first non-volatile cache memory device,
 - modifying data in the one or more read data chunks to generate one or more modified data chunks, and
 - writing the one or more modified data chunks to the long-term memory device.
5. The system of claim 1, further comprising the long-term memory device coupled to the first non-volatile cache memory device, wherein:
 - determining the data object to be a non-conforming data object is further based on the size of the data object exceeding an available storage capacity of the first non-volatile cache memory device; and
 - storing the one or more data chunks of the non-conforming data object in the freed-up one or more memory spaces includes, for each data chunk of the non-conforming data object:
 - determining whether there is sufficient space for storing each data chunk in the first non-volatile cache memory device, and
 - when there is not sufficient space:
 - writing one or more existing data chunks that are associated with one or more data objects other than the non-conforming data object and currently stored in the first non-volatile cache memory device to the long-term memory device,
 - freeing up the one or more memory spaces of the one or more existing data chunks to generate the freed-up one or more memory spaces, and
 - writing each data chunk to the freed-up one or more memory spaces in the first non-volatile cache memory device.
6. The system of claim 1, wherein:
 - the first non-volatile cache memory device includes a buffer size; and
 - the storage threshold reflects the buffer size.

7. The system of claim 1, further comprising:
a second non-volatile cache memory device coupled to the first non-volatile cache memory device, wherein:
the storage controller further performs operations comprising replicating each data chunk stored in the first non-volatile cache memory device to the second non-volatile cache memory device.
8. The system of claim 1, wherein the long-term memory device is one of:
a portion of a local storage network; and
a portion of an enterprise storage service.
9. The system of claim 1, wherein the first non-volatile cache memory device comprises a flash memory device.
10. The system of claim 1, wherein:
the storage controller is configured to be coupled to a network that is in communication with one or more client devices; and
the storage controller is located between the one or more client devices and the first non-volatile cache memory device when coupled to the network.
11. A system, comprising:
a non-volatile cache memory device including a buffer size;
means for receiving an input/output (I/O) request for a data object;
means for determining a size of the data object;
means for determining the data object to be a non-conforming data object relative to the buffer size of the non-volatile cache memory device;
means for dividing the non-conforming data object into a plurality of data chunks, each of the data chunks having a size that is less than the buffer size;
means for moving one or more data chunks stored in one or more memory spaces of the non-volatile cache memory device to a long-term memory device to free up the one or more memory spaces; and
means for storing one or more data chunks of the non-conforming data object in the freed-up one or more memory spaces of the non-volatile cache memory device.

12. The system of claim 11, wherein:
the I/O request is a write request; and
the means for storing the one or more data chunks of the non-conforming data object further comprises means for writing each of the one or more data chunks of the non-conforming data object to the freed-up one or more memory spaces of the non-volatile cache memory device.
13. The system of claim 11, wherein:
the I/O request is a read request; and
the means for storing the one or more data chunks of the non-conforming data object further comprises means for reading one or more data chunks stored in the long-term memory device to the non-volatile cache memory device.
14. The system of claim 11, wherein:
the I/O request is a read-write request; and
the means for storing the one or more data chunks of the non-conforming data object further comprises:
means for reading one or more data chunks stored in the long-term memory device to the non-volatile cache memory device,
means for modifying data in the one or more read data chunks to generate one or more modified data chunks, and
means for writing the one or more modified data chunks to the long-term memory device.
15. The system of claim 11, further comprising the long-term memory device coupled to the non-volatile cache memory device, wherein:
the means for determining the data object to be a non-conforming data object comprises means for basing the determination on the size of the data object exceeding the buffer size, and
the means for storing the one or more data chunks of the non-conforming data object in the freed-up one or more memory spaces includes, for each data chunk of the non-conforming data object:

means for determining whether there is sufficient space for storing each data chunk in the non-volatile cache memory device, and

when there is not sufficient space:

means for writing one or more existing data chunks that are associated with one or more data objects other than the non-conforming data object and currently stored in the non-volatile cache memory device to the long-term memory device,

means for freeing up the one or more memory spaces of the one or more existing data chunks to generate the freed-up one or more memory spaces, and

means for writing each data chunk to the freed-up one or more memory spaces in the non-volatile cache memory device.

16. A method of operation in a non-volatile memory system including a set of non-volatile memory devices, the method comprising:

receiving, at a storage controller coupled to a non-volatile cache memory device including a storage threshold, an input/output (I/O) request for executing an I/O operation on a data object;

determining a size of the data object;

determining the data object to be a non-conforming data object relative to the storage threshold of the non-volatile cache memory device;

dividing the non-conforming data object into a plurality of data chunks, each of the data chunks having a size that is less than the storage threshold;

moving one or more data chunks stored in one or more memory spaces of the non-volatile cache memory device to a long-term memory device to free up the one or more memory spaces; and

storing one or more data chunks of the non-conforming data object in the freed-up one or more memory spaces of the non-volatile cache memory device.

17. The method of claim 16, wherein:

the I/O request is a write request; and

storing the one or more data chunks of the non-conforming data object further comprises writing each of the one or more data chunks of the non-conforming data object to the freed-up one or more memory spaces of the non-volatile cache memory device.

18. The method of claim 16, wherein:
the I/O request is a read request; and
storing the one or more data chunks of the non-conforming data object further comprises reading each of the one or more data chunks of the non-conforming data object to the freed-up one or more memory spaces of the non-volatile cache memory device.

19. The method of claim 16, wherein:
the I/O request is a read-write request; and
storing the one or more data chunks of the non-conforming data object further comprises:
reading each of the one or more data chunks of the non-conforming data object to the freed-up one or more memory spaces of the non-volatile cache memory device,
modifying data in each of the one or more read data chunks to generate one or more modified data chunks, and
writing each of the one or more modified data chunks to the long-term memory device.

20. The method of claim 16, wherein:
determining the data object to be a non-conforming data object is further based on the size of the data object exceeding an available storage capacity of the non-volatile cache memory device, and
storing the one or more data chunks of the non-conforming data object in the freed-up one or more memory spaces includes, for each data chunk of the non-conforming data object:
determining whether there is sufficient space for storing each data chunk in the non-volatile cache memory device, and
when there is not sufficient space:
writing one or more existing data chunks that are associated with one or more data objects other than the non-conforming data object and currently stored in the non-volatile cache memory device to the long-term memory device,
freeing up the one or more memory spaces of the one or more existing data chunks to generate the freed-up one or more memory spaces, and
writing each data chunk to the freed-up one or more memory spaces in the non-volatile cache memory device.




FIG. 1A




FIG. 1B




FIG. 2




FIG. 3A




FIG. 3B




FIG. 3C




FIG. 4

500




FIG. 5A




FIG. 5B

502 ↗




FIG. 5C




FIG. 5D




FIG. 5E




FIG. 5F




FIG. 5G




FIG. 6




FIG. 7




FIG. 8

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2017/064368

A. CLASSIFICATION OF SUBJECT MATTER

INV. G06F12/0868 G06F3/06 G06F12/0875 G06F11/14
ADD.

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F

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

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

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2015/019502 A1 (ARONOVICH LIOR [CA] ET AL) 15 January 2015 (2015-01-15)	1-3,5, 11-13, 15-18,20
A	paragraphs [0006] - [0008], [0062] - [0077]; figures 2-4 -----	4,6-10, 14,19
Y	US 9 274 956 B1 (SALYERS DAVID CARL [US] ET AL) 1 March 2016 (2016-03-01)	1-3,5, 11-13, 15-18,20
A	column 25, line 40 - column 26, line 54; figure 11 column 13, line 44 - column 13, line 65; figure 2 column 7, line 33 - column 8, line 15 column 19, line 39 - column 20, line 34 column 21, line 49 - column 23, line 5 ----- -/--	4,6-10, 14,19



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

26 February 2018

Date of mailing of the international search report

06/03/2018

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Jardon, Stéphan

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2017/064368

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2006/064564 A1 (ACHANTA PHANI GOPAL V [US] ET AL) 23 March 2006 (2006-03-23) paragraphs [0046] - [0054]; claim 1; figure 5 -----	1-20
A	US 2015/199244 A1 (VENKATACHALAM RAMKUMAR [IN] ET AL) 16 July 2015 (2015-07-16) paragraphs [0004], [0014] - [0024]; claims 1-18; figures 1-3 -----	1-20
A	US 9 405 483 B1 (WEI DANNY [US] ET AL) 2 August 2016 (2016-08-02) column 7, line 57 - column 9, line 17; figures 3, 5A-5B, 7-9 column 10, line 34 - column 19, line 48; claims 1-20 -----	1-20
A	US 2013/318300 A1 (AGRAWAL DAKSHI [US] ET AL) 28 November 2013 (2013-11-28) paragraphs [0006], [0007], [0014], [0029] - [0057]; claims 1-20; figures 3-7 -----	1-20
A	US 2015/156174 A1 (FAHEY SEAN ANTHONY [US] ET AL) 4 June 2015 (2015-06-04) the whole document -----	1-20

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2017/064368

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2015019502	A1	15-01-2015	NONE
US 9274956	B1	01-03-2016	NONE
US 2006064564	A1	23-03-2006	US 2006064564 A1 23-03-2006 US 2008133867 A1 05-06-2008
US 2015199244	A1	16-07-2015	NONE
US 9405483	B1	02-08-2016	US 9405483 B1 02-08-2016 US 2016328168 A1 10-11-2016
US 2013318300	A1	28-11-2013	NONE
US 2015156174	A1	04-06-2015	CA 2930876 A1 11-06-2015 CN 105794173 A 20-07-2016 EP 3078179 A1 12-10-2016 JP 2017504856 A 09-02-2017 US 2015156174 A1 04-06-2015 WO 2015084890 A1 11-06-2015