

(12) UK Patent

(19) GB

(11) 2586382

(13) B

(45) Date of B Publication

09.02.2022

(54) Title of the Invention: **Fast recovery from failures in a chronologically ordered log-structured key-value storage system**

(51) INT CL: **G06F 11/14** (2006.01)

(21) Application No: **2014957.1**

(22) Date of Filing: **19.02.2019**

Date Lodged: **22.09.2020**

(30) Priority Data:
(31) **15904185** (32) **23.02.2018** (33) **US**

(86) International Application Data:
PCT/IB2019/051324 En 19.02.2019

(87) International Publication Data:
WO2019/162829 En 29.08.2019

(43) Date of Reproduction by UK Office **17.02.2021**

(72) Inventor(s):
Aayush Gupta
Sangeetha Seshadri

(73) Proprietor(s):
International Business Machines Corporation
(Incorporated in USA - New York)
New Orchard Road, Armonk, New York 10504,
United States of America

(74) Agent and/or Address for Service:
IBM United Kingdom Limited
Intellectual Property Department, Hursley Park,
Winchester, Hampshire, SO21 2JN, United Kingdom

(56) Documents Cited:
US 7200623 B2 **US 20170068599 A1**
US 20100332240 A1

(58) Field of Search:
As for published application 2586382 A viz:
INT CL **G06F**
Other: **DWPI, CPRSABS, USTXT, SIPOABS, CNABS,**
CNKI, IEEE
updated as appropriate

Additional Fields
Other: **WPI, EPODOC**

GB 2586382 B




FIG. 1




FIG. 2




FIG. 3




FIG. 4

500

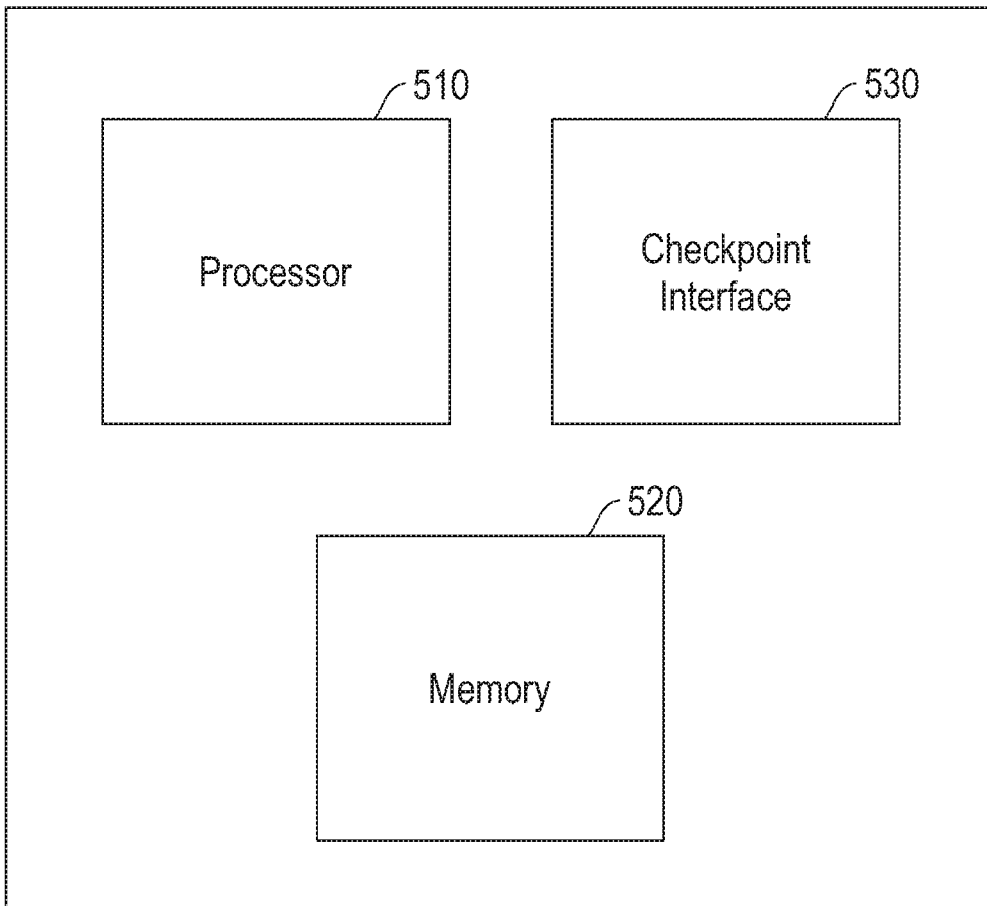


FIG. 5

600



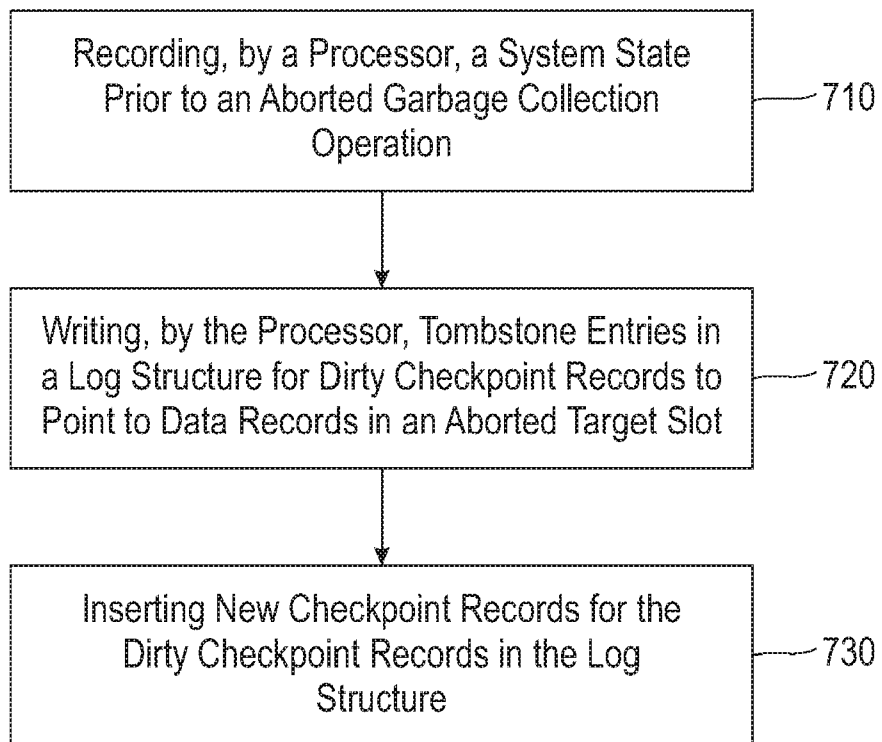


FIG. 6

700**FIG. 7**

FAST RECOVERY FROM FAILURES IN A CHRONOLOGICALLY ORDERED LOG-STRUCTURED KEY-VALUE STORAGE SYSTEM

Technical Field

[0001] The invention is generally directed to managing log-structured storage systems. In particular it provides for recovery after failure using a checkpoint in a chronological log-structured key-value store in a system.

BACKGROUND

[0002] Typical log-structured storage systems store record data in temporal order in a "log." These typical systems allow basic primitive operations, such as insert, update, delete, read. Each update of data results in a new record being inserted at the tail of the "log." Each delete results in a tombstone object being inserted at the tail of the log. Additionally, background garbage collection process compacts the data reclaiming space that does not contain valid data. However, recovery from checkpoints are problematic as checkpoints may occur concurrently with garbage collection

[0003] Therefore, there is a need in the art to address the aforementioned problem.

SUMMARY

[0004] According to the present invention there are provided a method, a computer program product, a computer program, and an apparatus according to the independent claims.

[0007] Viewed from a further aspect, the present invention provides a computer program stored on a computer readable medium and loadable into the internal memory of a digital computer, comprising software code portions, when said program is run on a computer, for performing the steps of the invention.

[0008] Embodiments relate to fast recovery from failures in a chronologically ordered log-structured key-value store. One embodiment provides a method for recovery after failure using a checkpoint in a chronological log-structured key-value store in a system including recording, by a processor, a system state prior to an aborted garbage collection operation. The processor writes tombstone entries in a log structure for dirty checkpoint records to point to data records in an aborted target slot. New checkpoint records are inserted in the log structure for the dirty checkpoint records.

[0009] These and other features, aspects and advantages of the present invention will become understood with reference to the following description, appended claims and accompanying figures.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] Embodiments of the present invention will now be described, by way of example only, with reference to the accompanying drawings in which:

FIG. 1 depicts a cloud computing environment, according to an embodiment;

FIG. 2 depicts a set of abstraction model layers, according to an embodiment;

FIG. 3 is a network architecture for retrospective snapshots in log-structured storage systems, according to an embodiment;

FIG. 4 shows a representative hardware environment that may be associated with the servers and/or clients of FIG. 1, according to an embodiment;

FIG. 5 is a block diagram illustrating a processor for fast recovery from failures in a chronologically ordered log-structured key-value store in log-structured storage systems, according to an embodiment;

FIG. 6 illustrates an example checkpoint record pointing to a garbage collection slot, according to an embodiment; and

FIG. 7 illustrates a block diagram for a process for recovery from failures in a chronologically ordered log-structured key-value store, according to one embodiment.

DETAILED DESCRIPTION

[0011] The descriptions of the various embodiments have been presented for purposes of illustration, but are not intended to be exhaustive or limited to the embodiments disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope of the described embodiments. The terminology used herein was chosen to best explain the principles of the embodiments, the practical application or

technical improvement over technologies found in the marketplace, or to enable others of ordinary skill in the art to understand the embodiments disclosed herein.

[0012] It is understood in advance that although this disclosure includes a detailed description of cloud computing, implementation of the teachings recited herein are not limited to a cloud computing environment. Rather, embodiments of the present invention are capable of being implemented in conjunction with any other type of computing environment now known or later developed.

[0013] One or more embodiments relate to fast recovery from failures in a chronologically ordered log-structured key-value store. One embodiment provides a method for recovery after failure using a checkpoint in a chronological log-structured key-value store in a system including recording, by a processor, a system state prior to an aborted garbage collection operation. The processor writes tombstone entries in a log structure for dirty checkpoint records to point to data records in an aborted target slot. New checkpoint records are inserted in the log structure for the dirty checkpoint records.

[0014] Cloud computing is a model of service delivery for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, network bandwidth, servers, processing, memory, storage, applications, virtual machines (VMs), and services) that can be rapidly provisioned and released with minimal management effort or interaction with a provider of the service. This cloud model may include at least five characteristics, at least three service models, and at least four deployment models.

[0015] Characteristics are as follows:

[0016] On-demand self-service: a cloud consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed and automatically, without requiring human interaction with the service's provider.

[0017] Broad network access: capabilities are available over a network and accessed through standard mechanisms that promote use by heterogeneous, thin or thick client platforms (e.g., mobile phones, laptops, and PDAs).

[0018] Resource pooling: the provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to demand. There is a sense of location independence in that the consumer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or data center).

[0019] Rapid elasticity: capabilities can be rapidly and elastically provisioned and, in some cases, automatically, to quickly scale out and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.

[0020] Measured service: cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active consumer accounts). Resource usage can be monitored, controlled, and reported, thereby providing transparency for both the provider and consumer of the utilized service.

[0021] Service Models are as follows:

[0022] Software as a Service (SaaS): the capability provided to the consumer is the ability to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface, such as a web browser (e.g., web-based email). The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited consumer-specific application configuration settings.

[0023] Platform as a Service (PaaS): the capability provided to the consumer is the ability to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including networks, servers, operating systems, or storage, but has control over the deployed applications and possibly application-hosting environment configurations.

[0024] Infrastructure as a Service (IaaS): the capability provided to the consumer is the ability to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls).

[0025] Deployment Models are as follows:

[0026] Private cloud: the cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on-premises or off-premises.

[0027] Community cloud: the cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be managed by the organizations or a third party and may exist on-premises or off-premises.

[0028] Public cloud: the cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.

[0029] Hybrid cloud: the cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds).

[0030] A cloud computing environment is a service oriented with a focus on statelessness, low coupling, modularity, and semantic interoperability. At the heart of cloud computing is an infrastructure comprising a network of interconnected nodes.

[0031] Referring now to FIG. 1, an illustrative cloud computing environment 50 is depicted. As shown, cloud computing environment 50 comprises one or more cloud computing nodes 10 with which local computing devices used by cloud consumers, such as, for example, personal digital assistant (PDA) or cellular telephone 54A, desktop computer 54B, laptop computer 54C, and/or automobile computer system 54N may communicate. Nodes 10 may communicate with one another. They may be grouped (not shown) physically or virtually, in one or more networks, such as private, community, public, or hybrid clouds as described hereinabove, or a combination thereof. This allows the cloud computing environment 50 to offer infrastructure, platforms, and/or software as services for which a cloud consumer does not need to maintain resources on a local computing device. It is understood that the types of computing devices 54A-N shown in FIG. 2 are intended to be illustrative only and that computing nodes 10 and cloud computing environment 50 can communicate with any type of computerized device over any type of network and/or network addressable connection (e.g., using a web browser).

[0032] Referring now to FIG. 2, a set of functional abstraction layers provided by the cloud computing environment 50 (FIG. 1) is shown. It should be understood in advance that the components, layers, and functions shown in FIG. 2 are intended to be illustrative only and embodiments of the invention are not limited thereto. As depicted, the following layers and corresponding functions are provided:

[0033] Hardware and software layer 60 include hardware and software components. Examples of hardware components include: mainframes 61; RISC (Reduced Instruction Set Computer) architecture based servers 62; servers 63; blade servers 64; storage devices 65; and networks and networking components 66. In some embodiments, software components include network application server software 67 and database software 68.

[0034] Virtualization layer 70 provides an abstraction layer from which the following examples of virtual entities may be provided: virtual servers 71; virtual storage 72; virtual networks 73, including virtual private networks; virtual applications and operating systems 74; and virtual clients 75.

[0035] In one example, a management layer 80 may provide the functions described below. Resource provisioning 81 provides dynamic procurement of computing resources and other resources that are utilized to perform tasks within the cloud computing environment. Metering and pricing 82 provide cost tracking as resources are utilized within the cloud computing environment and billing or invoicing for consumption of these resources. In one example, these resources may comprise application software licenses. Security provides identity verification for cloud consumers and tasks as well as protection for data and other resources. User portal 83 provides access to the cloud computing environment for consumers and system administrators. Service level management 84 provides cloud computing resource allocation and management such that required service levels are met. Service Level Agreement (SLA) planning and fulfillment 85 provide pre-arrangement for, and procurement of, cloud computing resources for which a future requirement is anticipated in accordance with an SLA.

[0036] Workloads layer 90 provides examples of functionality for which the cloud computing environment may be utilized. Examples of workloads and functions which may be provided from this layer include: mapping and navigation 91; software development and lifecycle management 92; virtual classroom education delivery 93; data analytics processing 94; transaction processing 95; and fast recovery from failures in a chronologically ordered log-structured key-value store processing 96. As mentioned above, all of the foregoing examples described with respect to FIG. 2 are illustrative only, and the invention is not limited to these examples.

[0037] It is understood all functions of one or more embodiments as described herein may be typically performed by the processing system 300 (FIG. 3) or the autonomous cloud environment 410 (FIG. 4), which can be tangibly embodied as hardware processors and with modules of program code. However, this need not be the case for non-real-time processing. Rather, for non-real-time processing the functionality recited herein could be carried out/implemented and/or enabled by any of the layers 60, 70, 80 and 90 shown in FIG. 2.

[0038] It is reiterated that although this disclosure includes a detailed description on cloud computing, implementation of the teachings recited herein are not limited to a cloud computing environment. Rather, the embodiments of the present invention may be implemented with any type of clustered computing environment now known or later developed.

[0039] FIG. 3 illustrates a network architecture 300, in accordance with one embodiment. As shown in FIG. 3, a plurality of remote networks 302 are provided, including a first remote network 304 and a second remote network 306. A gateway 301 may be coupled between the remote networks 302 and a proximate network 308. In the context of the present network architecture 300, the networks 304, 306 may each take any form including, but not limited to, a LAN, a WAN, such as the Internet, public switched telephone network (PSTN), internal telephone network, etc.

[0040] In use, the gateway 301 serves as an entrance point from the remote networks 302 to the proximate network 308. As such, the gateway 301 may function as a router, which is capable of directing a given packet of data that arrives at the gateway 301, and a switch, which furnishes the actual path in and out of the gateway 301 for a given packet.

[0041] Further included is at least one data server 314 coupled to the proximate network 308, which is accessible from the remote networks 302 via the gateway 301. It should be noted that the data server(s) 314 may include any type of computing device/groupware. Coupled to each data server 314 is a plurality of user devices 316. Such user devices 316 may include a desktop computer, laptop computer, handheld computer, printer, and/or any other type of logic-containing device. It should be noted that a user device 311 may also be directly coupled to any of the networks in some embodiments.

[0042] A peripheral 320 or series of peripherals 320, e.g., facsimile machines, printers, scanners, hard disk drives, networked and/or local storage units or systems, etc., may be coupled to one or more of the networks 304, 306, 308. It should be noted that databases and/or additional components may be utilized with, or integrated into, any type of network element coupled to the networks 304, 306, 308. In the context of the present description, a network element may refer to any component of a network.

[0043] According to some approaches, methods and systems described herein may be implemented with and/or on virtual systems and/or systems, which emulate one or more other systems, such as a UNIX system that emulates an IBM z/OS environment, a UNIX system that virtually hosts a MICROSOFT WINDOWS environment, a MICROSOFT WINDOWS system that emulates an IBM z/OS environment, etc. This virtualization and/or emulation may be implemented through the use of VMWARE software in some embodiments.

[0044] FIG. 4 shows a representative hardware system 400 environment associated with a user device 416 and/or server 314 of FIG. 3, in accordance with one embodiment. In one example, a hardware configuration includes a workstation having a central processing unit 410, such as a microprocessor, and a number of other units interconnected via a system bus 412. The workstation shown in FIG. 4 may include a Random Access Memory (RAM) 414, Read Only Memory (ROM) 416, an I/O adapter 418 for connecting peripheral devices, such as disk storage units 420 to the bus 412, a user interface adapter 422 for connecting a keyboard 424, a mouse 426, a speaker 428, a microphone 432, and/or other user interface devices, such as a touch screen, a digital camera (not shown), etc., to the bus 412, communication adapter 434 for connecting the workstation to a communication network 435 (e.g., a data processing network) and a display adapter 436 for connecting the bus 412 to a display device 438.

[0045] In one example, the workstation may have resident thereon an operating system, such as the MICROSOFT WINDOWS Operating System (OS), a MAC OS, a UNIX OS, etc. In one embodiment, the system

400 employs a POSIX® based file system. It will be appreciated that other examples may also be implemented on platforms and operating systems other than those mentioned. Such other examples may include operating systems written using JAVA, XML, C, and/or C++ language, or other programming languages, along with an object oriented programming methodology. Object oriented programming (OOP), which has become increasingly used to develop complex applications, may also be used. Microsoft, and Windows, are trademarks of Microsoft Corporation in the United States, other countries, or both. IBM, and z/OS are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. UNIX is a registered trademark of The Open Group in the United States and other countries. Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

[0046] FIG. 5 is a block diagram illustrating a processing node 500 for fast recovery from failures in a chronologically ordered log-structured key-value store in log-structured storage systems, according to an embodiment. The term garbage collection (GC) refers to reclaiming “disk space” occupied by stale entries in the log. For example, when a record is inserted, an entry is added to the tail of the log. When the same record is deleted, a tombstone entry is added to the tail of the log. The tombstone entry refers to the original location of the data on disk as created by the insert. The disk space occupied by the original inserted record may be garbage collected (provided the system is not maintaining older versions). Stale data may be the result of records that have been deleted or updated. Updates result in stale data because older versions of the data that are maintained in the log are no longer needed. Note that in a log-structured store, every insert, update or delete operation results in a record being inserted at the tail of the log.

[0047] In one embodiment, the processing node 500 includes one or more processors 510, a checkpoint interface 530 and a memory 520. In one embodiment, each processor(s) 510 performs processing for fast recovery from failures in a chronologically ordered log-structured key-value store in log-structured storage systems. Log structured storage system: complete in-memory index to quickly access record given the key, where index entry is represented by: <key, log address>. In a log-structured storage system, there is a need to quickly reconstruct an index during restart after failure or after shutdown. A naive option may include replay of the log from the beginning. This option, however, is time-consuming and resource intensive. In one embodiment, the checkpoint interface 530 provides for processing including a checkpoint of the index. For recovery from checkpoints, the checkpoint interface 530 may perform processing where checkpoints may occur concurrently with garbage collection. As a result, a checkpoint operation may record the position of a key on the “garbage collection target slot.” “Rollback” of a GC operation will void the contents on the GC target slot. Reads will fail, which results in data loss. Subsequent recovery points to an invalid location resulting in data loss.

[0048] In one embodiment, another approach by the checkpoint interface 530 may include to disallow checkpoint operations. Checkpoints, if available, do not proceed concurrently with GC. This approach results in

longer elapsed time to complete a checkpoint operation, and hence longer recovery time. This approach also impacts availability negatively by resulting in longer restart times after a failure.

[0049] In one embodiment, GC is performed by the checkpoint interface 530 as a transaction. Transaction processing divides information processing into individual, indivisible operations known as transactions. Each transaction must succeed or fail as a complete unit, and can never be only partially complete. In one embodiment, the checkpoint interface 530 performs rollback of a GC transaction in case of failure. GC transaction records include: <target slot (TS), victim slot (VS), begin offset in target slot (BO) >. In one embodiment, the GC target region is the <BO, end of slot> region on the GC target slot. GC transaction records begin and end on the GC transaction.

[0050] In one embodiment, the checkpoint interface 530 performs GC rollback as follows. The checkpoint interface 530 reads the recovery log to identify an incomplete GC target slot and victim slot. The checkpoint interface 530 begins index reconstruction from the checkpoint such that: 1) if a slot is a GC target slot & offset \geq BO, skip until end of slot and continue to the next slot; 2) if a checkpoint record A points to a GC target region: a) read the record in the GC target region pointed to by A (referred to as record B); b) insert a tombstone record for the checkpoint record A recording the key, and the contents of checkpoint record A (this is to allow subsequent use of the GC target slot region containing record B); c) insert a new checkpoint record using B's back-pointer pointing to the corresponding GC victim slot record (referred to as record C).

[0051] Next, the checkpoint interface 530 ends index reconstruction, zeros-out a target slot TS from offset BO until the end of the slot, and inserts an "Abort record for the garbage collection transaction" into the recovery log. At the end of the recovery, the system state is identical to the state of the system before GC processing.

[0052] FIG. 6 illustrates an example 600 checkpoint record pointing to a GC slot, according to an embodiment. The index 610 of the example 600 includes keys in the left column and slot mapping in the right column. In one example, K1 refers to slot 1 620 and offset 2; K2 refers to slot 1 620 and offset 3, etc. In the example 600, the slots shown include slot 1 620, slot 51 630, slot 3 640 and slot 4 650. The victim slot VS is shown as slot 3 640, and the target slot TS is shown as slot 51 630. In this example 600, the index 610 includes an incorrect entry 611 that points to a zeroed-out location. The back-pointer (points to previous address) chain 660 shows the pointer from slot 4 650 (k8) to slot 51 630 and offset 6 (k8). The checkpoint record 670 refers to the checkpoint record pointing to the GC target slot.

[0053] FIG. 7 illustrates a block diagram for a process 700 for recovery from failures in a chronologically ordered log-structured key-value store, according to one embodiment. In one embodiment, block 710 in process 700 includes recording, by a processor (e.g., by a processor 510, FIG. 5), a system state prior to an aborted GC operation. In block 720, process 700 performs writing, by the processor, tombstone entries in a log structure for

dirty checkpoint records to point to data records in an aborted target slot. In block 730, process 700 inserts new checkpoint records for the dirty checkpoint records in the log structure.

[0054] In one embodiment, in process 700 the system may be a chronologically ordered log-structured key-value store system. In one embodiment, for process 700 the GC operation is a GC transaction. For process 700, the aborted target slot may be a target slot of the GC operation.

[0055] In one embodiment, process 700 may further include providing for checkpoint operations to proceed concurrently with GC processing while maintaining chronological order of data using processing to restore the system to a consistent state after failure during the GC operation (e.g., a failed transaction).

[0056] In one embodiment, process 700 may further include reading a recovery log after the aborted GC operation, and identifying an incomplete GC target slot and an incomplete GC victim slot.

[0057] In one embodiment, process 700 may further include recording the GC transaction on a separate recovery log. In one embodiment, process 700 may additionally include zeroing out a region of the target slot of the log structure from a begin offset in the target slot until an end of the target slot, and inserting an abort record for the GC transaction into the separate recovery log.

[0058] As will be appreciated by one skilled in the art, aspects of the present invention may be embodied as a system, method or computer program product. Accordingly, aspects of the present invention may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects that may all generally be referred to herein as a "circuit," "module" or "system." Furthermore, aspects of the present invention may take the form of a computer program product embodied in one or more computer readable medium(s) having computer readable program code embodied thereon.

[0059] Any combination of one or more computer readable medium(s) may be utilized. The computer readable medium may be a computer readable signal medium or a computer readable storage medium. A computer readable storage medium may be, for example, but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, or device, or any suitable combination of the foregoing. More specific examples (a non-exhaustive list) of the computer readable storage medium would include the following: an electrical connection having one or more wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a magnetic storage device, 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 for use by or in connection with an instruction execution system, apparatus, or device.

[0060] A computer readable signal medium may include a propagated data signal with computer readable program code embodied therein, for example, in baseband or as part of a carrier wave. Such a propagated signal may take any of a variety of forms, including, but not limited to, electro-magnetic, optical, or any suitable combination thereof. A computer readable signal medium may be any computer readable medium that is not a computer readable storage medium and that can communicate, propagate, or transport a program for use by or in connection with an instruction execution system, apparatus, or device.

[0061] Program code embodied on a computer readable medium may be transmitted using any appropriate medium, including but not limited to wireless, wireline, optical fiber cable, RF, etc., or any suitable combination of the foregoing.

[0062] Computer program code for carrying out operations for aspects of the present invention may be written in any combination of one or more programming languages, including an object-oriented programming language such as Java, Smalltalk, C++ or the like and conventional procedural programming languages, such as the "C" programming language or similar programming languages. The program code may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider).

[0063] Aspects of the present invention are described below with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

[0064] These computer program instructions may also be stored in a computer readable medium that can direct a computer, other programmable data processing apparatus, or other devices to function in a particular manner, such that the instructions stored in the computer readable medium produce an article of manufacture

including instructions which implement the function/act specified in the flowchart and/or block diagram block or blocks.

[0065] The computer program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other devices to cause a series of operational steps to be performed on the computer, other programmable apparatus or other devices to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide processes for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

[0066] The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of instructions, which comprises one or more executable instructions for implementing the specified logical function(s). In some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts or carry out combinations of special purpose hardware and computer instructions.

[0067] References in the claims to an element in the singular is not intended to mean "one and only" unless explicitly so stated, but rather "one or more." All structural and functional equivalents to the elements of the above-described exemplary embodiment that are currently known or later come to be known to those of ordinary skill in the art are intended to be encompassed by the present claims. No claim element herein is to be construed under the provisions of 35 U.S.C. section 112, sixth paragraph, unless the element is expressly recited using the phrase "means for" or "step for."

[0068] The terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting of the invention. As used herein, the singular forms "a", "an" and "the" are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will be further understood that the terms "comprises" and/or "comprising," when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof.

[0069] The corresponding structures, materials, acts, and equivalents of all means or step plus function elements in the claims below are intended to include any structure, material, or act for performing the function in

combination with other claimed elements as specifically claimed. The description of the present invention has been presented for purposes of illustration and description, but is not intended to be exhaustive or limited to the invention in the form disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope of the invention. The embodiment was chosen and described in order to best explain the principles of the invention and the practical application, and to enable others of ordinary skill in the art to understand the invention for various embodiments with various modifications as are suited to the particular use contemplated.

CLAIMS

1. A method for recovery after failure using a checkpoint in a chronologically ordered log-structured key-value store in a chronologically ordered log-structured key-value store system, the method comprising:
 - recording, by a processor, a system state prior to an aborted garbage collection operation;
 - writing, by the processor, tombstone entries in a log structure for dirty checkpoint records to point to data records in an aborted target slot; and
 - inserting new checkpoint records for the dirty checkpoint records in the log structure.
2. The method of claim 1, wherein the garbage collection operation comprises a garbage collection transaction.
3. The method of claim 2, wherein the aborted target slot is a target slot of the garbage collection operation.
4. The method of either of claims 2 or 3, further comprising:
 - providing for checkpoint operations to proceed concurrently with garbage collection processing while maintaining chronological order of data using processing to restore the system to a consistent state after failure during the garbage collection operation.
5. The method of any of the preceding claims, further comprising:
 - reading a recovery log after the aborted garbage collection operation; and
 - identifying an incomplete garbage collection target slot and an incomplete garbage collection victim slot.
6. The method of any of claims 2 to 5, further comprising:
 - recording the garbage collection transaction on a separate recovery log.
7. The method of claim 6, further comprising:
 - zeroing out a region of the target slot of the log structure from a begin offset in the target slot until an end of the target slot; and
 - inserting an abort record for the garbage collection transaction into the separate recovery log.
8. A computer program product for recovery after failure using a checkpoint in a chronological log-structured key-value store in a system, the computer program product comprising:
 - a computer readable storage medium readable by a processing circuit and storing instructions for execution by the processing circuit for performing a method according to any of claims 1 to 7.

9. A computer program stored on a computer readable medium and loadable into the internal memory of a digital computer, comprising software code portions, when said program is run on a computer, for performing the method of any of claims 1 to 7.

10. An apparatus for recovery after failure using a checkpoint in a chronologically ordered log-structured key-value store in a chronologically ordered log-structured key-value store system, the apparatus comprising:
 - a memory storing instructions; and
 - a processor executing the instructions to:
 - record a system state prior to an aborted garbage collection operation;
 - write tombstone entries in a log structure for dirty checkpoint records to point to data records in an aborted target slot; and
 - insert new checkpoint records for the dirty checkpoint records in the log structure.

11. The apparatus of claim 10, wherein the garbage collection operation comprises a garbage collection transaction, and the aborted target slot is a target slot of the garbage collection transaction.

12. The apparatus of claim 11, wherein the processor further executes instructions comprising:
 - providing for checkpoint operations to proceed concurrently with garbage collection processing while maintaining chronological order of data using processing to restore the system to a consistent state after failure during the garbage collection transaction;
 - reading a recovery log after the aborted garbage collection transaction; and
 - identifying an incomplete garbage collection target slot and an incomplete garbage collection victim slot.

13. The apparatus of either of claims 11 or 12, wherein the processor further executes instructions comprising:
 - recording the garbage collection transaction on a separate recovery log.

14. The apparatus of claim 13, wherein the processor further executes instructions comprising:
 - zeroing out a region of the target slot of the log structure from a begin offset in the target slot until an end of the target slot; and
 - inserting an abort record for the garbage collection transaction into the separate recovery log.