



US 20140282513A1

(19) **United States**(12) **Patent Application Publication**
Pacheco et al.(10) **Pub. No.: US 2014/0282513 A1**(43) **Pub. Date: Sep. 18, 2014**(54) **INSTRUCTION SET ARCHITECTURE FOR
COMPUTE-BASED OBJECT STORES**(52) **U.S. Cl.**CPC **G06F 9/455** (2013.01)USPC **718/1**(71) Applicants: **David Pacheco**, San Francisco, CA
(US); **Mark Cavage**, Seattle, WA (US);
Yunong Xiao, Seattle, WA (US); **Bryan**
Cantrill, Piedmont, CA (US)


(57)

ABSTRACT(72) Inventors: **David Pacheco**, San Francisco, CA
(US); **Mark Cavage**, Seattle, WA (US);
Yunong Xiao, Seattle, WA (US); **Bryan**
Cantrill, Piedmont, CA (US)

Instruction set architectures for compute-centric object stores. An exemplary method may include receiving a request from a user, the request identifying parameters of a compute operation that is to be executed against one or more objects in a distributed object store, generating a set of tasks from the request that comprise instructions for a daemon, locating the one or more objects within the distributed object store, the one or more objects being stored on a physical node. The method includes providing the set of tasks to a daemon, the daemon controlling execution of the compute operation by a virtual operating system container based upon the set of tasks, and storing an output of the virtual operating system container in the distributed object store.

(21) Appl. No.: **13/831,349**(22) Filed: **Mar. 14, 2013****Publication Classification**(51) **Int. Cl.****G06F 9/455**

(2006.01)






FIG. 1





FIG. 2







FIG. 4

*FIG. 5*




FIG. 6

INSTRUCTION SET ARCHITECTURE FOR COMPUTE-BASED OBJECT STORES

FIELD OF THE PRESENT TECHNOLOGY

[0001] The present technology relates generally to an instruction set architecture (ISA) for compute-centric object stores. ISAs of the present technology allow for efficient scheduling and management of compute operations across distributed object stores. The ISAs provide a means for expressing compute operations within the context of a distributed object store, as well as a mechanism for coordinating how data flows through a compute-centric object store system.

BACKGROUND

[0002] Various methods and systems for providing multi-tenant computing systems, such as cloud computing, have been attempted. In general, a cloud-based computing environment is a resource that typically combines the computational power of a large model of processors and/or that combines the storage capacity of a large model of computer memories or storage devices. For example, systems that provide a cloud resource may be utilized exclusively by their owners; or such systems may be accessible to outside users who deploy applications within the computing infrastructure to obtain the benefit of large computational or storage resources.

[0003] The cloud may be formed, for example, by a network of servers with each server (or at least a plurality thereof) providing processor and/or storage resources. These servers may manage workloads provided by multiple users (e.g., cloud resource consumers or other users). Typically, each user places workload demands upon the cloud that vary in real-time, sometimes dramatically. The nature and extent of these variations typically depend on the type of business associated with the user.

[0004] Oftentimes, these cloud computing systems leverage virtual machines for their users. A virtual machine ("VM") is an emulation of a real-world computing system. Often, the virtual machine provides a user with one or more different operating systems than the operating system of the local machine ("host") that is running the virtual machine. The VM provides a complete system platform that provides the one or more operating systems. The VM is typically managed by a hypervisor that mediates computing resources of the host machine for use by the VM via hardware emulation. The use of hardware emulation is often deleterious to VM performance and, in turn, reduces the number of VMs that may run on a given host machine. Additionally, as the number of VMs on a host machine increases and they begin to operate concurrently, the hypervisor must coordinate the varying workloads of the VMs to prevent instability.

[0005] In general, systems that provide data centered distributed applications, such as systems that run applications on large clusters of shared hardware are often programming intensive for users. That is, these systems require users to create complex programs for executing compute operations against objects or data stores. This is often caused by the complexity of the hardware and/or software frameworks required to manage data flow through and/or hardware resource virtualization within these systems.

SUMMARY OF THE PRESENT TECHNOLOGY

[0006] According to some embodiments, the present technology may be directed to systems that comprise: (a) one or more processors; and (b) logic encoded in one or more tangible media for execution by the one or more processors and when executed operable to perform operations comprising: (i) receiving a request from a user, the request identifying a compute operation that is to be executed against an object in a distributed object store; (ii) locating the object within the distributed object store, the object being stored on a physical node; (iii) assigning a virtual operating system container to the object; (iv) providing an instruction set to a daemon associated with the object, the daemon controlling execution of the compute operation by the virtual operating system container according to the instruction sets; and (v) storing output of the virtual operating system container in the distributed object store.

[0007] According to some embodiments, the present technology may be directed to a multitenant object storage system that comprises: (a) receiving a request from a user, the request identifying parameters of a compute operation that is to be executed against objects in a distributed object store, the request also comprising an identifier; (b) generating a set of tasks from the request that comprise an instruction set for a daemon; (c) locating the objects within the distributed object store, the objects being stored on a physical node; (d) providing the set of tasks to a daemon of the physical node, the daemon controlling execution of the compute operation by a virtual operating system container based upon the set of tasks; and (e) storing an output of the virtual operating system container in the distributed object store.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] Certain embodiments of the present technology are illustrated by the accompanying figures. It will be understood that the figures are not necessarily to scale and that details not necessary for an understanding of the technology or that render other details difficult to perceive may be omitted. It will be understood that the technology is not necessarily limited to the particular embodiments illustrated herein.

[0009] FIG. 1 is a block diagram of an exemplary architecture in which embodiments of the present technology may be practiced;

[0010] FIG. 2 is a schematic diagram of an exemplary guest virtual operating system container;

[0011] FIG. 3 is a schematic diagram illustrating the colocation of guest virtual operating system containers for multiple tenants on an object store;

[0012] FIG. 4 is a schematic diagram of a guest virtual operating system container applied onto an object store;

[0013] FIG. 5 is a flowchart of an exemplary method for executing a compute flow using a set of tasks, according to an instruction set architecture; and

[0014] FIG. 6 illustrates an exemplary computing system that may be used to implement embodiments according to the present technology.

DESCRIPTION OF EXEMPLARY EMBODIMENTS

[0015] While this technology is susceptible of embodiment in many different forms, there is shown in the drawings and will herein be described in detail several specific embodiments with the understanding that the present disclosure is to

be considered as an exemplification of the principles of the technology and is not intended to limit the technology to the embodiments illustrated.

[0016] The terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting of the present technology. 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.

[0017] It will be understood that like or analogous elements and/or components, referred to herein, may be identified throughout the drawings with like reference characters. It will be further understood that several of the figures are merely schematic representations of the present technology. As such, some of the components may have been distorted from their actual scale for pictorial clarity.

[0018] FIG. 1 is a block diagram of an exemplary architecture 100 in which embodiments of the present technology may be practiced. The architecture 100 comprises a plurality of client devices 105A-N that communicatively couple with a compute-centric object store system, hereinafter “system 110.” It will be understood that the architecture 100 may include a plurality of systems, such as system 110. For the sake of brevity and clarity, a detailed description of an exemplary system 110 will be provided below, although the features of the system 110 apply equally to all of the plurality of systems. The plurality of client devices 105A-N may communicatively couple with the system 110 via any one or combination of a number of private and/or public networks, such as the Internet. According to some embodiments, the client devices 105A-N may submit requests or jobs to a network service 110B, which is a constituent part of the system 110. In some instances, the network service 110E evaluates request received from users to determine one or more physical nodes that comprise objects that correspond to the request.

[0019] In general, the system 110 comprises an object store 110A that provides “compute” as a first class citizen of an object store 110A. More specifically, compute operations (e.g., instructing the system to compute on objects in the object store) of the present technology resemble a top-level API function, similar to processes like storing or fetching objects in the object store 110A.

[0020] It will be understood that the terms “object store” comprise a network service for storing unstructured, arbitrary-sized chunks of data (objects). It will be further understood that the object store may not support modifications to existing objects, but supports full object replacement operations, although systems that support both object modification and full object replacement operations may also utilize the features of the present technology to perform compute operations directly on (e.g., in-situ) objects within the object store.

[0021] In some embodiments, the system 110 may be configured to receive a request to perform a compute operation on at least a portion of an object store, from a first user. Again, the user may be associated with one of the client devices. The request identifies parameters of the compute operation as well as objects against which the compute operation is executed.

[0022] In some instances, the system 110 may assign virtual operating system containers to a user, based upon a request. The system 110 may map objects to the containers that are associated with the user. Typically, these objects are identified by the user in the request. A virtual operating system container performs the compute operation on an object according to the identified parameters of the request. The system 110 may then clear the virtual operating system containers and return the virtual operating system containers to a pool of virtual operating system containers. Additional aspects of the system 110 will be described in greater detail below.

[0023] It will be understood that a compute-centric object store may be created to operate without the user of virtual operating system (global kernel) or virtual operating system containers. While such an object store would provide advantages such as in-situ computation of data (where objects are processed directly on the object store), the object store may not isolate tenants in the similarly to systems that utilize a virtual operating system and/or virtual operating system containers.

[0024] In these instances, the compute-centric object store may be configured to receiving a request to perform a compute operation on at least a portion of an object store from a first user via a network service, the request identifying parameters of the compute operation. The object store may also execute an operating system process for the objects identified in the request. The operating system process may perform the compute operation on the object according to the identified parameters of the request. Additionally, once the compute operation has been executed, the operating system process may be terminated by the virtual operating system.

[0025] The terms in-situ computation will be understood to include the execution of compute operations against objects in an object store, where the objects not moved or copied from or within the object store.

[0026] In some embodiments, the system 110 is comprised of a hardware layer 115 that provides a logical interface with at least one or more processors and a memory which stores logic that is executed by the one or more processors. Generally, the hardware layer 115 controls one or more of the hardware components of a computing system, such as the computing system 600 of FIG. 6, which will be described in greater detail below. By way of non-limiting example, the hardware layer 115 may manage the hardware components of a server blade or another similar device. The hardware layer 115 provides access to the physical hardware that services a global operating system kernel 120 that cooperates with the hardware layer 115. The global operating system kernel 120 may also be referred to as a host operating system kernel.

[0027] Generally, the global operating system kernel 120 is configured to administer and manage a pool of guest virtual operating system containers, such as containers 125A-N. The containers 125A-N may operate on a distributed object store in a multitenant manner, where multiple containers can operate on the same object store simultaneously. It will be understood that each user is assigned container from the pool, on an as-needed basis. When a container is applied to an object store the container is referred to as a tenant.

[0028] According to some embodiments, the system kernel 120 may be utilized to setup the pool of guest virtual operating system containers. The system kernel 120 may also be configured to provide a command line interpreter interface that allows users to request jobs, execute other operating

system implemented applications, and interact with a virtual operating system in a manner that is substantially indistinguishable relative to an operating system executing on a bare metal device.

[0029] Generally, a job may be input by a user via a command line interpreter, such as a Unix shell terminal. More specifically, the user may express a computation using the same language as the language used by a Unix shell terminal. The actual request is submitted to the network service **110B**. Indeed, a request may be submitted as an HTTP request to the network service **110B**. The body of the request describes the computation to perform in terms of what commands are input into the command line interpreter, which is running within a container. Contrastingly systems that utilize multiple VMs that each comprises an operating system kernel, which are managed by a hypervisor, often require users to construct complex programs or scripts to perform compute operations. Compute operations for traditional VM systems require complex programming due to a complex framework that is used by the hypervisor to coordinate hardware emulation for each of the VMs.

[0030] Using the command line interpreter interface, the user may specify one or more desired compute operations that are to be executed against objects (such as object **130**) within an object store **110A** (see FIG. 3). It is noteworthy that the object store **110A** may include, for example, a local or distributed object store that maintains contiguous blobs, blocks, or chunks of data. It will be understood that the objects stored in the object store **110A** are complete objects, such as files or other similar data structures. Moreover, the compute operations executed against the object store **110A** may be performed in such a way that partial stores of data are avoided.

[0031] In order to perform compute operations on objects for multiple users, the system kernel **120** may collocate containers **125A-N** onto the object store **110A**, and execute the containers **125A-N** simultaneously. In FIG. 3, a plurality of containers, such as container **125A** has been placed onto each of a plurality of objects within the object store **110A**. Thus, a virtual operating system container is assigned to each of the plurality of objects specified in the user request. Most frequently, the assignment of a single container to a single object occurs when the system executes a “map” phase operation. The details of map and reduce phases provided by the system **110** will be described in greater detail below.

[0032] Broadly speaking, a virtual operating system container may be a lightweight virtualization solution offering a complete and secure user environment that operates on a single global kernel (system kernel **120**), providing performance characteristics that are similar to operating systems that operate on bare metal devices. That is, a virtual machine operates on emulated hardware and is subject to control by a hypervisor, which produces computing inefficiencies. A virtual operating system container may operate without the computing inefficiencies of a typical virtual machine.

[0033] In some instances, the system kernel **120** may utilize a KVM (Kernel Virtual Machine) that improves the efficiency of the a virtual operating system, such as the global operating system kernel, by leveraging CPU virtualization extensions to eliminate a substantial majority of the binary translation (i.e., hardware emulation) that are frequently required by VMs.

[0034] Turning to FIG. 2, an exemplary virtual operating system container **125A** (FIG. 1) is shown as comprising a quick emulation layer (QEMU) **135**, a virtual guest operating

system **140**, and a compute application **145** that is managed by the virtual guest operating system **140**. The QEMU **135** provides hardware emulation and is also VMM (virtual machine monitor). It is noteworthy that in some embodiments the QEMU **135** is not a strict hypervisor layer, but rather each QEMU **135** may be independent in some exemplary embodiments. That is, there may be one QEMU **135** one per container instead of a single QEMU **135** supporting several VMs. Advantageously, the operations of both a VM and a VMM may be combined into the QEMU **135**.

[0035] According to some embodiments, the compute application **145** that is executed may include a primitive O/S compute operation. Exemplary compute operations may include operating system primitive operations, such as query, word count, send, receive, and so forth. Additionally, the operations may comprise more sophisticated operations, such as operations that include audio or video transcoding. Additionally, in some instances, users may store programs or applications in the object store itself. Users may then execute the programs as a part of a compute operation.

[0036] In some instances the compute operations may include one or more phases such as a map phase, followed by a reduce phase. Generally, a map phase may include an operation that is executed against each of a plurality of objects individually, by a plurality of containers. In some instances, a unique container is assigned to each object that is to be processed.

[0037] In contrast, a reduce phase may be executed by a single container against a plurality of objects in a batch manner. Using an example such as word count, it will be assumed that the objects of the object store **135** may comprise text files. The application **145** may execute a map phase to count the words in each of the text files. The output of the application **145** may be stored in a plurality of output objects that are stored in the object store **135**. A compute application **145** of another container may execute a reduce phase that sums the output objects of the map phase and generates a word count for all objects within the object store **135**.

[0038] It will be understood that the system kernel **120** may schedule and coordinate various compute operations (and phases) performed by the compute applications **145** of all containers. In sum, the system kernel **120** may act similarly to a hypervisor that manages the compute operations of the various active containers. Based upon the request input by the user, the system kernel **120** may instruct the containers to perform a series of map functions, as well as a reduce functions. The map and reduce functions may be coordinated to produce the desired output specified in the request.

[0039] Turning to FIG. 3, after receiving a request from a user, the system kernel **120** may select a first set of containers, which includes container **125A** from the pool of containers. This container **125A** is assigned to a user. In response to receiving a request from a second user, the system kernel **120** may also select a second set of containers from the pool of containers.

[0040] Based upon the request received from the first tenant, the system kernel **120** may map the first set of containers to a plurality of objects, such as object **130**, stored in the object store **110A**. Likewise, the system kernel **120** may map a second set of containers to a plurality of different objects stored in the object store **110A** for the second user. The objects and containers for the first user may be referred to as a compute zone of the first user, while the objects mapped to the container **125N** may be referred to as a compute zone of

the second user. The maintenance of compute zones allows the system kernel **120** to provide multitenant access to the object store **110A**, even when the first and second users are potentially adversarial. For example, the first and second users may be commercial competitors. For security, the system kernel **120** maintains compute zones in order to balkanize object storage and prevent access to objects of other users. Additionally, the balkanization of object storage also ensures fair distribution of resources between users.

[0041] It will be understood that the system kernel **120** may maintain as many containers and compute zones as allowed by the processor(s) of the hardware layer **115**. Additionally, the system kernel **120** assigns a container to a user on an as-needed basis, meaning that containers may not be assigned permanently to a user, which would result in a monopolization of resources when the user is not performing compute operations.

[0042] FIG. 4 illustrates the placement of the container **125A** onto the data store **110A**. It is understood that the container **125A** encircles a plurality of objects in the data store **110A**. This mapping of multiple object to a single container would be commonly seen in a reduce phase, where the container is performing a concatenating or summation process on the outputs of individual containers, such as the containers shown in FIG. 3.

[0043] Additionally, because the container is placed onto the object store, the system kernel **120** need not transfer objects from the object store **110A** into the container for processing in some exemplary embodiments. Advantageously, the container operates directly on the objects of the object store **110A**.

[0044] According to some embodiments, the containers **125A-N** managed by the system kernel **120** are empty when the containers **125A-N** are in the pool. After objects are mapped to the container, compute operations may be executed by the container on the objects, and a desired output is generated, the system kernel **120** may clear the container and return the container to the pool.

[0045] In some instances, the system kernel **120** may not generate containers until a request is received from a user. That is, the system kernel **120** may “spin up” or launch containers when a request is received from the user. This allows for minimum impact to the bare metal resources, such as the CPU, as the system kernel **120** need not even maintain a pool of virtual operating system containers, which are awaiting user requests. That is, maintaining a pool of containers requires CPU and memory resources. When the compute operations have been completed, the system kernel **120** may terminate the containers, rather than clearing the containers and returning the containers to a pool.

[0046] In accordance with the present disclosure, an instruction set architecture may be implemented within the system **110**. In some embodiments, the instruction set architecture may specify an application programming interface that allows the system **110** to interact with the distributed object store.

[0047] According to some embodiments, the system **110** communicatively couples with the object store **110A** using a services related application programming interface (SAPI) **155**, which provides features such as automatic discovery of object stores, dynamic configuration of object stores, and an API for a user portal. In sum, the SAPI allows users to configure, deploy, and upgrade applications using a set of loosely-coupled, federated services. In some embodiments,

the SAPI may include an underlying API and an autoconfig agent, also referred to as a daemon **150**. A SAPI client may also be disseminated to clients. It will be understood that the daemon **150** may be associated with a physical node **160** of the object store **110A**.

[0048] In accordance with some embodiments according to the present disclosure, various object stores, such as object store **110A** of FIGS. 3 and 4, comprise a single SAPI zone. It will be understood that the SAPI zone may be stateless and the SAPI zone may be configured to write objects into the object store **110A**. In addition to storing objects, the SAPI zone may also communicatively couple with a VM API to provision zones and a network API (NAPI) to reserve network interface controllers (NIC) and lookup network universal unique identifiers (UUID).

[0049] It will be understood that the SAPI **155** may comprise three main object types such as applications, services, and instances. It is noteworthy that an application may comprise one or more services, and each service may comprise one or more instances. Moreover, instances may represent actual object store zones, and such zones inherit zone parameters and metadata from their associated applications and services.

[0050] Also, the application, service, and instance information may be used by the compute application of a virtual operating system container that is placed onto an object store. The daemon **150** may control the operation of the containers operating on the daemon's object store.

[0051] Each application, service and instance may include three sets of properties. For example, “params” may comprise zone parameters like a zone's RAM size, disk quota, image UUID, and so forth. These parameters are evaluated when a zone is provisioned. Another property comprises “metadata”, which defines metadata available to the daemon **150**. These metadata keys and values form the input of a script template in a configuration manifest (described below). As these values are updated, the daemon **150** may rewrite any configuration and make reference to changed metadata values. Yet another property comprises “manifests” that define a set of configuration manifests are indexed by name to facilitate inheriting manifest from parent objects.

[0052] It is noteworthy that creating applications and services have no effect on running zones. When an instance is created, a zone is provisioned using the above information from its associated application, service, and instance. Stated otherwise, applications and services (e.g., a job or request) may be defined separate from the objects that the applications and services are to be executed against. Thus, a job may be thought of abstractly as a workflow template. Advantageously, when the user requests the execution of a job, objects need only be defined by the user. The workflow template is then applied against the objects.

[0053] In some embodiments, the daemon **150** of a zone may be tasked with maintaining configuration inside that zone. The daemon **150** queries the SAPI **155** directly to determine which files to write and where to write them within the object store **110A**.

[0054] The daemon **150** uses objects called configuration manifests; those objects describe the contents, location, and semantics of configuration files for a zone. Those manifests contain a script template which is rendered using the metadata from the associated application, service, and instance.

[0055] When a user provides a request to the system **110**, the system kernel **120** may coordinate a compute flow of

compute operations which are managed by the daemon **150**. That is, the system kernel **120** may receive a request or “job” from a user, via a command line interpreter. The request identifies parameters of a compute operation that is to be executed against objects in a distributed object store. For example, a request may include performing a word count operation on a file.

[0056] To facilitate compute flow during the compute process, the system kernel **120** may assign an identifier for the request. This identifier provides a unique identifier that allows objects and outputs of compute operations to be correlated to the user. Objects previously stored in the object store may be correlated to the user utilizing a unique identifier. According to some embodiments, the identifier comprises the name of an input object or job name. This name may be specified by an end user submitting a job/request to the system or may be generated by the system from the request.

[0057] The user may also identify objects for the compute operation, using, for example, the command line interpreter. An exemplary find object command may include Find|User|Object Store Location; where the Object Store Location defines the object store that includes the object(s) which are necessary for execution of the compute operation.

[0058] In some instances, the system kernel **120** may query various daemons of object stores to locate the objects within the distributed object store. After the object have been located, the system kernel **120** may generate a set of tasks (e.g., an instruction set) that defines the various compute operations that are to be performed by the daemon of the located object store. In the example provided above, the set of tasks may include only one word count task that is provided to a single daemon of an object store (e.g., physical node). This relatively simple compute operation does not require coordination or scheduling of operations of multiple objects.

[0059] The daemon **150** may provide instructions to one or more virtual operating system containers that are placed onto the object store by the system kernel **120**. That is, the instruction sets provided to the containers is based upon the task assigned to the daemon **150** from the system kernel **120**.

[0060] In some instances, the set of tasks may include a more complex arrangement of operations that are executed against a plurality of objects stores. The system kernel **120** may interact with the daemon to coordinate processing of these objects in a specified order.

[0061] Additionally, the set of tasks may define various map phases that are to be executed on the objects of the object store, as well as various reduce phases that are executed on the outputs of the map phases. It will be understood that objects within the workflow may be tracked and correlated together using the identifier. For example, if an instruction set passed to a daemon requires performing a word count compute operation on 100 text files, each of the objects of the compute operation would be correlated using the identifier. Thus, the objects of the compute operation would comprise 100 objects that each includes a word count value for their corresponding text file. The identifier may be appended to the object as metadata.

[0062] It will also be understood that a map phase may result in multiple outputs, which are generated from a single input object. For example, assume that usage logs for a computing device are stored for a 24 hour time period. To determine hourly usage rates, the 24 hour log object may be separated into 24 distinct objects. Thus, the map phase may

receive the 24 hour log object and may split the same into constituent output objects to complete the map phase.

[0063] It will be understood that a more complex request may require a more complicated set of tasks (e.g., phases). For example, if the user desires to look at all 11 p.m. to 12 p.m. user logs for a plurality of computing devices, the set of tasks may require not only the map task where a single input object is processed into multiple objects, but also a reduce phase that sums a plurality of 11 p.m. to 12 p.m. user logs for a plurality of devices.

[0064] In sum, the system kernel **120** will provide a daemon with tasks that include a map phase for generating the hour increment logs from various input objects. Additionally, the tasks also inform the daemon to return the output objects, which may be stored as an aggregate 11 p.m. to 12 p.m. log object within the object store.

[0065] It will be understood that the daemon of a physical node (e.g., object store) may control execution of compute operations by the one or more virtual operating system containers that are placed onto the object store via the system kernel **120**.

[0066] Thus, it is appreciated with that intermediate output objects may not be output to the user directly, but may be fed back into the system for additional processing, such as with the map and reduce phases described above. Moreover, the set of tasks generated by the system kernel **120** may include any number of map phases and/or reduce phases, which vary according to the steps required to produce the desired output.

[0067] FIG. 5 is a flowchart of an exemplary method **500** for executing a compute flow using a set of tasks, according to an instruction set architecture. The method **500** may include a step **505** of receiving a request from a user, the request identifying parameters of a compute operation that is to be executed against objects in a distributed object store.

[0068] According to some embodiments, the method may include a step **510** of locating one or more objects within the distributed object store, the one or more objects being stored on one or more physical nodes. Once objects have been located, the method may include a step **515** of generating a set of tasks from the request that comprise instructions for a daemon. Again, the set of tasks may define applications, services and instances that are that can be utilized by virtual operating system containers. It will be understood that the set of tasks comprises an instruction set that is a translation of the user request into meaningful input that can be executed by one or more virtual operating system containers.

[0069] Upon generating the set of tasks, the method includes a step **520** of providing the set of tasks to a daemon of the physical node. Again, the daemon controls execution of the compute operation by one or more virtual operating system containers based upon the set of tasks. The daemon functions as a virtual operating system hypervisor that coordinates the operation and execution of the containers.

[0070] Finally, the method includes a step **525** of storing an output of the virtual operating system container in the distributed object store.

[0071] FIG. 6 illustrates an exemplary computing system **600** that may be used to implement an embodiment of the present systems and methods. The system **600** of FIG. 6 may be implemented in the contexts of the likes of computing systems, networks, servers, or combinations thereof. The computing system **600** of FIG. 6 includes one or more processors **610** and main memory **620**. Main memory **620** stores, in part, instructions and data for execution by processor **610**.

Main memory **620** may store the executable code when in operation. The system **600** of FIG. **6** further includes a mass storage device **630**, portable storage device **640**, output devices **650**, user input devices **660**, a display system **670**, and peripheral devices **680**.

[0072] The components shown in FIG. **6** are depicted as being connected via a single bus **690**. The components may be connected through one or more data transport means. Processor unit **610** and main memory **620** may be connected via a local microprocessor bus, and the mass storage device **630**, peripheral device(s) **680**, portable storage device **640**, and display system **670** may be connected via one or more input/output (I/O) buses.

[0073] Mass storage device **630**, which may be implemented with a magnetic disk drive or an optical disk drive, is a non-volatile storage device for storing data and instructions for use by processor unit **610**. Mass storage device **630** may store the system software for implementing embodiments of the present technology for purposes of loading that software into main memory **620**.

[0074] Portable storage device **640** operates in conjunction with a portable non-volatile storage medium, such as a floppy disk, compact disk, digital video disc, or USB storage device, to input and output data and code to and from the computer system **600** of FIG. **6**. The system software for implementing embodiments of the present technology may be stored on such a portable medium and input to the computer system **600** via the portable storage device **640**.

[0075] User input devices **660** provide a portion of a user interface. User input devices **660** may include an alphanumeric keypad, such as a keyboard, for inputting alpha-numeric and other information, or a pointing device, such as a mouse, a trackball, stylus, or cursor direction keys. Additional user input devices **660** may comprise, but are not limited to, devices such as speech recognition systems, facial recognition systems, motion-based input systems, gesture-based systems, and so forth. For example, user input devices **660** may include a touchscreen. Additionally, the system **600** as shown in FIG. **6** includes output devices **650**. Suitable output devices include speakers, printers, network interfaces, and monitors.

[0076] Display system **670** may include a liquid crystal display (LCD) or other suitable display device. Display system **670** receives textual and graphical information, and processes the information for output to the display device.

[0077] Peripheral device(s) **680** may include any type of computer support device to add additional functionality to the computer system. Peripheral device(s) **680** may include a modem or a router.

[0078] The components provided in the computer system **600** of FIG. **6** are those typically found in computer systems that may be suitable for use with embodiments of the present technology and are intended to represent a broad category of such computer components that are well known in the art. Thus, the computer system **600** of FIG. **6** may be a personal computer, hand held computing system, telephone, mobile computing system, workstation, server, minicomputer, main-frame computer, or any other computing system. The computer may also include different bus configurations, networked platforms, multi-processor platforms, etc. Various operating systems may be used including Unix, Linux, Windows, Mac OS, Palm OS, Android, iOS (known as iPhone OS before June 2010), QNX, and other suitable operating systems.

[0079] It is noteworthy that any hardware platform suitable for performing the processing described herein is suitable for use with the systems and methods provided herein. Computer-readable storage media refer to any medium or media that participate in providing instructions to a central processing unit (CPU), a processor, a microcontroller, or the like. Such media may take forms including, but not limited to, non-volatile and volatile media such as optical or magnetic disks and dynamic memory, respectively. Common forms of computer-readable storage media include a floppy disk, a flexible disk, a hard disk, magnetic tape, any other magnetic storage medium, a CD-ROM disk, digital video disk (DVD), any other optical storage medium, RAM, PROM, EPROM, a FLASHEPROM, any other memory chip or cartridge.

[0080] Computer program code for carrying out operations for aspects of the present technology 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 coupled with 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).

[0081] 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 technology has been presented for purposes of illustration and description, but is not intended to be exhaustive or limited to the present technology in the form disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the present technology. Exemplary embodiments were chosen and described in order to best explain the principles of the present technology and its practical application, and to enable others of ordinary skill in the art to understand the present technology for various embodiments with various modifications as are suited to the particular use contemplated.

[0082] Aspects of the present technology are described above with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the present technology. 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.

[0083] These computer program instructions may also be stored in a computer readable medium that can direct a com-

puter, 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.

[0084] 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.

[0085] 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 technology. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of code, which comprises one or more executable instructions for implementing the specified logical function(s). It should also be noted that, 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 combinations of special purpose hardware and computer instructions.

[0086] While various embodiments have been described above, it should be understood that they have been presented by way of example only, and not limitation. The descriptions are not intended to limit the scope of the technology to the particular forms set forth herein. Thus, the breadth and scope of a preferred embodiment should not be limited by any of the above-described exemplary embodiments. It should be understood that the above description is illustrative and not restrictive. To the contrary, the present descriptions are intended to cover such alternatives, modifications, and equivalents as may be included within the spirit and scope of the technology as defined by the appended claims and otherwise appreciated by one of ordinary skill in the art. The scope of the technology should, therefore, be determined not with reference to the above description, but instead should be determined with reference to the appended claims along with their full scope of equivalents.

1. A system, comprising:
one or more processors; and
logic encoded in one or more tangible media for execution by the one or more processors and when executed operable to perform operations comprising:
receiving a request from a user, the request identifying a compute operation that is to be executed against an object in a distributed object store;
locating the object within the distributed object store, the object being stored on a physical node;
assigning a virtual operating system container to the object;

- providing an instruction set to a daemon associated with the object, the daemon controlling execution of the compute operation by the virtual operating system container according to the instruction set; and
storing output of the virtual operating system container in the distributed object store.
2. The system according to claim 1, wherein the instruction set causes the virtual operating system container to generate a plurality of outputs using the object.
3. The system according to claim 1, wherein the request is received via a command line interpreter.
4. The system according to claim 1, wherein the virtual operating system container performs the compute operation on the object in such a way that the object remains in-situ within the object store.
5. The system according to claim 1, wherein the virtual operating system container is selected from a pool of virtual operating system containers managed by a global operating kernel system.
6. The system according to claim 5, wherein the virtual operating system container is returned to the pool of virtual operating system containers managed by a global operating system kernel after the virtual operating system container has been cleared.
7. The system according to claim 1, wherein assigning comprises making a read-only instantiation of the object available to applications executing inside the virtual operating system container.
8. The system according to claim 1, further comprising assigning an identifier to the request, wherein the daemon correlates the output of the container to the request using the identifier.
9. The system according to claim 8, wherein the identifier is associated with a compute zone for a user, allowing a global operating system kernel to correlate the output, the request, and the user.
10. A method, comprising:
receiving a request from a user, the request identifying parameters of a compute operation that is to be executed against one or more objects in a distributed object store, the request also comprising an identifier;
generating a set of tasks from the request that comprise instructions for a daemon;
locating the one or more objects within the distributed object store, the one or more objects being stored on a physical node;
providing the set of tasks to the daemon, the daemon controlling execution of the compute operation by a virtual operating system container based upon the set of tasks; and
storing an output of the virtual operating system container in the distributed object store.
11. The method according to claim 10, wherein a plurality of outputs of the virtual operating system container are correlated to the identifier.
12. The method according to claim 10, wherein the set of tasks define a compute flow that defines compute operations that are executed according to a schedule.
13. The method according to claim 12, wherein the compute operations comprise compute phases that include any of:
a map phase where virtual operating system containers perform a first process on a plurality of single objects of the one or more physical nodes;

a reduce phase where the virtual operating system container performs a second process on a plurality of output objects of the map phase; and

any combination of map phases and reduce phases.

14. The method according to claim **13**, further comprising correlating the plurality of output objects of the map phase to the request using the identifier, wherein the virtual operating system container of the reduce phase locates the plurality of output objects stored in the one or more physical nodes using the identifier.

15. The method according to claim **10**, wherein the instructions define applications, services, and instances for virtual operating system containers that operate on the objects.

16. The method according to claim **10**, wherein the instructions cause the virtual operating system containers to generate a plurality of output objects from a single input object.

17. The method according to claim **16**, wherein the plurality of output objects are correlated to one another using the identifier.

18. The method according to claim **10**, wherein the virtual operating system container performs a compute operation on an object in such a way that the object remains in-situ within the object store.

19. The method according to claim **10**, wherein the daemon controls operation of a plurality of virtual operating system containers assigned to a plurality of objects, using the set of tasks.

* * * * *