



- (51) International Patent Classification:  
G06F 21/62 (2013.01) G06F 9/50 (2006.01)  
G06F 9/455 (2018.01) G06F 16/182 (2019.01)
- (21) International Application Number:  
PCT/US2020/052280
- (22) International Filing Date:  
23 September 2020 (23.09.2020)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
16/586,816 27 September 2019 (27.09.2019) US  
16/586,818 27 September 2019 (27.09.2019) US  
16/586,825 27 September 2019 (27.09.2019) US
- (63) Related by continuation (CON) or continuation-in-part (CIP) to earlier applications:  
US 16/586,816 (CON)  
Filed on 27 September 2019 (27.09.2019)  
US 16/586,818 (CON)  
Filed on 27 September 2019 (27.09.2019)  
US 16/586,825 (CON)  
Filed on 27 September 2019 (27.09.2019)

- (71) Applicant: AMAZON TECHNOLOGIES, INC. [US/US]; PO Box 81226, Seattle, Washington 98108-1226 (US).
- (72) Inventors: DATTA, Ramyanshu; c/o Amazon Technologies, Inc., 410 Terry Avenue North, Seattle, Washington 98109-5210 (US). HARRIS, Timothy Lawrence; c/o Amazon Technologies, Inc., 410 Terry Avenue North, Seattle, Washington 98109-5210 (US). MILLER, Kevin C.; c/o Amazon Technologies, Inc., 410 Terry Avenue North, Seattle, Washington 98109-5210 (US).
- (74) Agent: SEELIG, Melanie J.; Knobbe Martens Olson & Bear LLP, 2040 Main Street, Fourteenth Floor, Irvine, California 92614 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, IT, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

(54) Title: ON-DEMAND CODE OBFUSCATION OF DATA IN INPUT PATH OF OBJECT STORAGE SERVICE

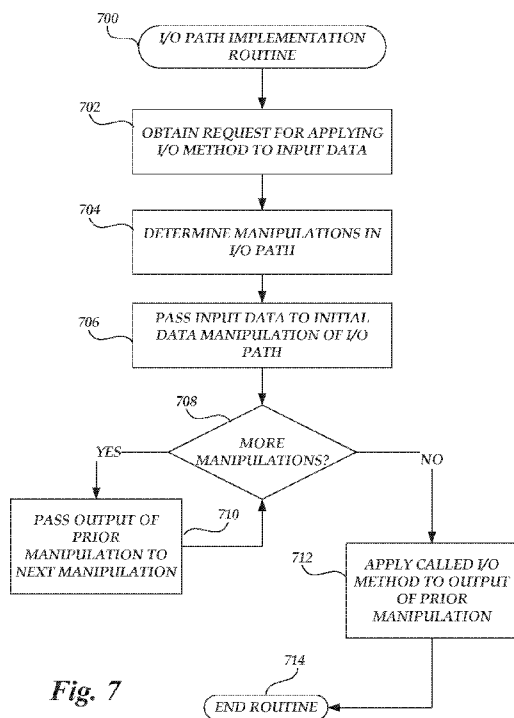


Fig. 7

(57) Abstract: Input and output (I/O) to an object storage service are modified by implementing one or more owner-specified functions to I/O requests. A function can implement a data manipulation, such as filtering out sensitive data before reading or writing the data. The functions can be applied prior to implementing a request method (e.g., GET or PUT) specified within the I/O request, such that the data to which the method is applied may not match the object specified within the request. For example, a user may request to obtain (e.g., GET) a data set. The data set may be passed to a function that filters sensitive data to the data set, and the GET request method may then be applied to the output of the function. In this manner, owners of objects on an object storage service are provided with greater control of objects stored or retrieved from the service.



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

**Declarations under Rule 4.17:**

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

**Published:**

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

## ON-DEMAND CODE OBFUSCATION OF DATA IN INPUT PATH OF OBJECT STORAGE SERVICE

### BACKGROUND

[0001] Computing devices can utilize communication networks to exchange data. Companies and organizations operate computer networks that interconnect a number of computing devices to support operations or to provide services to third parties. The computing devices can be located in a single geographic location or located in multiple, distinct geographic locations (e.g., interconnected via private or public communication networks). Specifically, data centers or data processing centers, herein generally referred to as a “data center,” may include a number of interconnected computing systems to provide computing resources to users of the data center. The data centers may be private data centers operated on behalf of an organization or public data centers operated on behalf, or for the benefit of, the general public.

[0002] To facilitate increased utilization of data center resources, virtualization technologies allow a single physical computing device to host one or more instances of virtual machines that appear and operate as independent computing devices to users of a data center. With virtualization, the single physical computing device can create, maintain, delete, or otherwise manage virtual machines in a dynamic manner. In turn, users can request computer resources from a data center, including single computing devices or a configuration of networked computing devices, and be provided with varying numbers of virtual machine resources.

[0003] In addition to computational resources, data centers provide a number of other beneficial services to client devices. For example, data centers may provide data storage services configured to store data submitted by client devices, and enable retrieval of that data over a network. A variety of types of data storage services can be provided, often varying according to their input/output (I/O) mechanisms. For example, database services may allow I/O based on a database query language, such as the Structured Query Language (SQL). Block storage services may allow I/O based on modification to one or more defined-length blocks, in a manner similar to how an operating system interacts with local storage, and may thus facilitate virtualized disk drives usable, for example, to store an operating system of a virtual

machine. Object storage services may allow I/O at the level of individual objects or resources, such as individual files, which may vary in content and length. For example, an object storage service may provide an interface compliant with the Representational State Transfer (REST) architectural style, such as by allowing I/O based on calls designating input data and a hypertext transport protocol request method (e.g., GET, PUT, POST, DELETE, etc.) to be applied to that data. By transmitting a call designating input data and a request method, a client can thus retrieve the data from an object storage service, write the data to the object storage service as a new object, modify an existing object, etc.

### BRIEF DESCRIPTION OF THE DRAWINGS

[0004] FIG. 1 is a block diagram depicting an illustrative environment in which an object storage service can operate in conjunction with an on-demand code execution system to implement functions in connection with input/output (I/O) requests to the object storage service;

[0005] FIG. 2 depicts a general architecture of a computing device providing a frontend of the object storage service of FIG. 1;

[0006] FIG. 3 is a flow diagram depicting illustrative interactions for enabling a client device to modify an I/O path for the object storage service by insertion of a function implemented by execution of a task on the on-demand code execution system;

[0007] FIG. 4 is an illustrative visualization of a pipeline of functions to be applied to an I/O path for the object storage service of FIG. 1;

[0008] FIGS. 5A-5B show a flow diagram depicting illustrative interactions for handling a request to store input data as an object on the object storage service of FIG. 1, including execution of an owner-specified task to the input data and storage of output of the task as the object;

[0009] FIGS. 6A-6B show a flow diagram depicting illustrative interactions for handling a request to retrieve data of an object on the object storage service of FIG. 1, including execution of an owner-specified task to the object and transmission of an output of the task to a requesting device as the object;

[0010] FIG. 7 is a flow chart depicting an illustrative routine for implementing owner-defined functions in connection with an I/O request obtained at the object storage service of FIG. 1 over an I/O path; and

[0011] FIG. 8 is a flow chart depicting an illustrative routine for executing a task on the on-demand code execution system of FIG. 1 to enable data manipulations during implementation of an owner-defined function.

[0012] FIG. 9 is a flow chart depicting an illustrative routine for executing a task on the on-demand code execution system of FIG. 1 to execute first and second functions in response to store a data object provided in multiple portions.

[0013] FIG. 10 is a system diagram of illustrative data flows and interactions between various components of the service provider system in connection with the routine illustrated in FIG. 9.

[0014] FIG. 11 is a flow chart depicting an illustrative routine for executing a task on the on-demand code execution system of FIG. 1 to dynamically obfuscate portions of input data in response to store the input data.

[0015] FIG. 12 is a system diagram of illustrative data flows and interactions between various components of the service provider system in connection with the routine illustrated in FIG. 11.

[0016] FIG. 13 is a flow chart depicting an illustrative routine for executing a task on the on-demand code execution system of FIG. 1 to dynamically determine and store an index of the contents of input data in response to a request to store the input data.

[0017] FIG. 14 is a system diagram of illustrative data flows and interactions between various components of the service provider system in connection with the routine illustrated in FIG. 13.

#### DETAILED DESCRIPTION

[0018] Generally described, aspects of the present disclosure relate to handling requests to read or write to data objects on an object storage system. More specifically, aspects of the present disclosure relate to modification of an input/output (I/O) path for an object storage service, such that one or more data manipulations can be inserted into the I/O path to modify the data to which a called request method is applied, without requiring a calling client

device to specify such data manipulations. In one embodiment, data manipulations occur through execution of user-submitted code, which may be provided for example by an owner of a collection of data objects on an object storage system in order to control interactions with that data object. For example, in cases where an owner of an object collection wishes to ensure that end users do not submit objects to the collection including any personally identifying information (to ensure end user's privacy), the owner may submit code executable to strip such information from a data input. The owner may further specify that such code should be executed during each write of a data object to the collection. Accordingly, when an end user attempts to write input data to the collection as a data object (e.g., via an HTTP PUT method), the code may be first executed against the input data, and resulting output data may be written to the collection as the data object. Notably, this may result in the operation requested by the end user—such as a write operation—being applied not to the end user's input data, but instead to the data output by the data manipulation (e.g., owner-submitted) code. In this way, owners of data collections control I/O to those collections without relying on end users to comply with owner requirements. Indeed, end users (or any other client device) may be unaware that modifications to I/O are occurring. As such, embodiments of the present disclosure enable modification of I/O to an object storage service without modification of an interface to the service, ensuring inter-compatibility with other pre-existing software utilizing the service.

[0019] In some embodiments of the present disclosure, data manipulations may occur on an on-demand code execution system, sometimes referred to as a serverless execution system. Generally described, on-demand code execution systems enable execution of arbitrary user-designated code, without requiring the user to create, maintain, or configure an execution environment (e.g., a physical or virtual machine) in which the code is executed. For example, whereas conventional computing services often require a user to provision a specific device (virtual or physical), install an operating system on the device, configure application, define network interfaces, and the like, an on-demand code execution system may enable a user to submit code and may provide to the user an application programming interface (API) that, when used, enables the user to request execution of the code. On receiving a call through the API, the on-demand code execution system may generate an execution environment for the code, provision the environment with the code, execute the code, and provide a result. Thus, an on-demand code execution system can remove a need for a user to handle configuration and

management of environments for code execution. Example techniques for implementing an on-demand code execution system are disclosed, for example, within U.S. Patent No. 9,323,556, entitled “PROGRAMMATIC EVENT DETECTION AND MESSAGE GENERATION FOR REQUESTS TO EXECUTE PROGRAM CODE,” and filed September 30, 2014 (the “556 Patent”), the entirety of which is hereby incorporated by reference.

[0020] Due to the flexibility of on-demand code execution system to execute arbitrary code, such a system can be used to create a variety of network services. For example, such a system could be used to create a “micro-service,” a network service that implements a small number of functions (or only one function), and that interacts with other services to provide an application. In the context of on-demand code execution systems, the code executed to create such a service is often referred to as a “function” or a “task,” which can be executed to implement the service. Accordingly, one technique for performing data manipulations within the I/O path of an object storage service may be to create a task on an on-demand code execution system that, when executed, performs the required data manipulation. Illustratively, the task could provide an interface similar or identical to that of the object storage service, and be operable to obtain input data in response to a request method call (e.g., HTTP PUT or GET calls), execute the code of the task against the input data, and perform a call to the object storage service for implementation of the request method on resulting output data. A downside of this technique is a complexity. For example, end users might be required under this scenario to submit I/O requests to the on-demand code execution system, rather than the object storage service, to ensure execution of the task. Should an end user submit a call directly to the object storage service, task execution may not occur, and thus an owner would not be enabled to enforce a desired data manipulation for an object collection. In addition, this technique may require that code of a task be authored to both provide an interface to end users that enables handling of calls to implement request methods on input data, and an interface that enables performance of calls from the task execution to the object storage service. Implementation of these network interfaces may significantly increase the complexity of the required code, thus disincentivizing owners of data collections from using this technique. Moreover, where user-submitted code directly implements network communication, that code may need to be varied according to the request method handled. For example, a first set of code may be required to support GET operations, a second set of code may be required to support PUT operations, etc.

Because embodiments of the present disclosure relieve the user-submitted code of the requirement of handling network communications, one set of code may in some cases be enabled to handle multiple request methods.

[0021] To address the above-noted problems, embodiments of the present disclosure can enable strong integration of serverless task executions with interfaces of an object storage service, such that the service itself is configured to invoke a task execution on receiving an I/O request to a data collection. Moreover, generation of code to perform data manipulations may be simplified by configuring the object storage service to facilitate data input and output from a task execution, without requiring the task execution to itself implement network communications for I/O operations. Specifically, an object storage service and on-demand code execution system can be configured in one embodiment to “stage” input data to a task execution in the form of a handle (e.g., a POSIX-compliant descriptor) to an operating-system-level input/output stream, such that code of a task can manipulate the input data via defined-stream operations (e.g., as if the data existed within a local file system). This stream-level access to input data can be contrasted, for example, with network-level access of input data, which generally requires that code implement network communication to retrieve the input data. Similarly, the object storage service and on-demand code execution system can be configured to provide an output stream handle representing an output stream to which a task execution may write output. On detecting writes to the output stream, the object storage service and on-demand code execution system may handle such writes as output data of the task execution, and apply a called request method to the output data. By enabling a task to manipulate data based on input and output streams passed to the task, as opposed to requiring the code to handle data communications over a network, the code of the task can be greatly simplified.

[0022] Another benefit of enabling a task to manipulate data based on input and output handles is increased security. A general-use on-demand code execution system may operate permissively with respect to network communications from a task execution, enabling any network communication from the execution unless such communication is explicitly denied. This permissive model is reflective of the use of task executions as micro-services, which often require interaction with a variety of other network services. However, this permissive model also decreases security of the function, since potentially malicious network communications



can also reach the execution. In contrast to a permissive model, task executions used to perform data manipulations on an object storage system's I/O path can utilize a restrictive model, whereby only explicitly-allowed network communications can occur from an environment executing a task. Illustratively, because data manipulation can occur via input and output handles, it is envisioned that many or most tasks used to perform data manipulation in embodiments of the present disclosure would require no network communications to occur at all, greatly increasing security of such an execution. Where a task execution does require some network communications, such as to contact an external service to assist with a data manipulation, such communications can be explicitly allowed, or "whitelisted," thus exposing the execution in only a strictly limited manner.

[0023] In some embodiments, a data collection owner may require only a single data manipulation to occur with respect to I/O to the collection. Accordingly, the object storage service may detect I/O to the collection, implement the data manipulation (e.g., by executing a serverless task within an environment provisioned with input and output handles), and apply the called request method to the resulting output data. In other embodiments, an owner may request multiple data manipulations occur with respect to an I/O path. For example, to increase portability and reusability, an owner may author multiple serverless tasks, which may be combined in different manners on different I/O paths. Thus, for each path, the owner may define a series of serverless tasks to be executed on I/O to the path. Moreover, in some configurations, an object storage system may natively provide one or more data manipulations. For example, an object storage system may natively accept requests for only portions of an object (e.g., of a defined byte range), or may natively enable execution of queries against data of an object (e.g., SQL queries). In some embodiments, any combination of various native manipulations and serverless task-based manipulations may be specified for a given I/O path. For example, an owner may specify that, for a particular request to read an object, a given SQL query be executed against the object, the output of which is processed via a first task execution, the output of which is processed via a second task execution, etc. The collection of data manipulations (e.g., native manipulations, serverless task-based manipulations, or a combination thereof) applied to an I/O path is generally referred to herein as a data processing "pipeline" applied to the I/O path.

[0024] In accordance with aspects of the present disclosure, a particular path modification (e.g., the addition of a pipeline) applied to an I/O path may vary according to attributes of the path, such as a client device from which an I/O request originates or an object or collection of objects within the request. For example, pipelines may be applied to individual objects, such that the pipeline is applied to all I/O requests for the object, or a pipeline may be selectively applied only when certain client devices access the object. In some instances, an object storage service may provide multiple I/O paths for an object or collection. For example, the same object or collection may be associated with multiple resource identifiers on the object storage service, such that the object or collection can be accessed through the multiple identifiers (e.g., uniform resource identifiers, or URIs), which illustratively correspond to different network-accessible endpoints. In one embodiment, different pipelines may be applied to each I/O path for a given object. For example, a first I/O path may be associated with unprivileged access to a data set, and thus be subject to data manipulations that remove confidential information from the data set prior during retrieval. A second I/O path may be associated with privileged access, and thus not be subject to those data manipulations. In some instances, pipelines may be selectively applied based on other criteria. For example, whether a pipeline is applied may be based on time of day, a number or rate of accesses to an object or collection, etc.

[0025] As will be appreciated by one of skill in the art in light of the present disclosure, the embodiments disclosed herein improve the ability of computing systems, such as object storage systems, to provide and enforce data manipulation functions against data objects. Whereas prior techniques generally depend on external enforcement of data manipulation functions (e.g., requesting that users strip personal information before uploading it), embodiments of the present disclosure enable direct insertion of data manipulation into an I/O path for the object storage system. Moreover, embodiments of the present disclosure provide a secure mechanism for implementing data manipulations, by providing for serverless execution of manipulation functions within an isolated execution environment. Embodiments of the present disclosure further improve operation of serverless functions, by enabling such functions to operate on the basis of local stream (e.g., "file") handles, rather than requiring that functions act as network-accessible services. The presently disclosed embodiments therefore address technical problems inherent within computing systems, such as the difficulty of

enforcing data manipulations at storage systems and the complexity of creating external services to enforce such data manipulations. These technical problems are addressed by the various technical solutions described herein, including the insertion of data processing pipelines into an I/O path for an object or object collection, potentially without knowledge of a requesting user, the use of serverless functions to perform aspects of such pipelines, and the use of local stream handles to enable simplified creation of serverless functions. Thus, the present disclosure represents an improvement on existing data processing systems and computing systems in general.

[0026] The general execution of tasks on the on-demand code execution system will now be discussed. As described in detail herein, the on-demand code execution system may provide a network-accessible service enabling users to submit or designate computer-executable source code to be executed by virtual machine instances on the on-demand code execution system. Each set of code on the on-demand code execution system may define a “task,” and implement specific functionality corresponding to that task when executed on a virtual machine instance of the on-demand code execution system. Individual implementations of the task on the on-demand code execution system may be referred to as an “execution” of the task (or a “task execution”). In some cases, the on-demand code execution system may enable users to directly trigger execution of a task based on a variety of potential events, such as transmission of an application programming interface (“API”) call to the on-demand code execution system, or transmission of a specially formatted hypertext transport protocol (“HTTP”) packet to the on-demand code execution system. In accordance with embodiments of the present disclosure, the on-demand code execution system may further interact with an object storage system, in order to execute tasks during application of a data manipulation pipeline to an I/O path. The on-demand code execution system can therefore execute any specified executable code “on-demand,” without requiring configuration or maintenance of the underlying hardware or infrastructure on which the code is executed. Further, the on-demand code execution system may be configured to execute tasks in a rapid manner (e.g., in under 100 milliseconds [ms]), thus enabling execution of tasks in “real-time” (e.g., with little or no perceptible delay to an end user). To enable this rapid execution, the on-demand code execution system can include one or more virtual machine instances that are “pre-warmed” or pre-initialized (e.g., booted into an operating system and executing a complete or substantially

complete runtime environment) and configured to enable execution of user-defined code, such that the code may be rapidly executed in response to a request to execute the code, without delay caused by initializing the virtual machine instance. Thus, when an execution of a task is triggered, the code corresponding to that task can be executed within a pre-initialized virtual machine in a very short amount of time.

[0027] Specifically, to execute tasks, the on-demand code execution system described herein may maintain a pool of executing virtual machine instances that are ready for use as soon as a request to execute a task is received. Due to the pre-initialized nature of these virtual machines, delay (sometimes referred to as latency) associated with executing the task code (e.g., instance and language runtime startup time) can be significantly reduced, often to sub-100 millisecond levels. Illustratively, the on-demand code execution system may maintain a pool of virtual machine instances on one or more physical computing devices, where each virtual machine instance has one or more software components (e.g., operating systems, language runtimes, libraries, etc.) loaded thereon. When the on-demand code execution system receives a request to execute program code (a “task”), the on-demand code execution system may select a virtual machine instance for executing the program code of the user based on the one or more computing constraints related to the task (e.g., a required operating system or runtime) and cause the task to be executed on the selected virtual machine instance. The tasks can be executed in isolated containers that are created on the virtual machine instances, or may be executed within a virtual machine instance isolated from other virtual machine instances acting as environments for other tasks. Since the virtual machine instances in the pool have already been booted and loaded with particular operating systems and language runtimes by the time the requests are received, the delay associated with finding compute capacity that can handle the requests (e.g., by executing the user code in one or more containers created on the virtual machine instances) can be significantly reduced.

[0028] As used herein, the term “virtual machine instance” is intended to refer to an execution of software or other executable code that emulates hardware to provide an environment or platform on which software may execute (an example “execution environment”). Virtual machine instances are generally executed by hardware devices, which may differ from the physical hardware emulated by the virtual machine instance. For example, a virtual machine may emulate a first type of processor and memory while being executed on

a second type of processor and memory. Thus, virtual machines can be utilized to execute software intended for a first execution environment (e.g., a first operating system) on a physical device that is executing a second execution environment (e.g., a second operating system). In some instances, hardware emulated by a virtual machine instance may be the same or similar to hardware of an underlying device. For example, a device with a first type of processor may implement a plurality of virtual machine instances, each emulating an instance of that first type of processor. Thus, virtual machine instances can be used to divide a device into a number of logical sub-devices (each referred to as a “virtual machine instance”). While virtual machine instances can generally provide a level of abstraction away from the hardware of an underlying physical device, this abstraction is not required. For example, assume a device implements a plurality of virtual machine instances, each of which emulate hardware identical to that provided by the device. Under such a scenario, each virtual machine instance may allow a software application to execute code on the underlying hardware without translation, while maintaining a logical separation between software applications running on other virtual machine instances. This process, which is generally referred to as “native execution,” may be utilized to increase the speed or performance of virtual machine instances. Other techniques that allow direct utilization of underlying hardware, such as hardware pass-through techniques, may be used, as well.

[0029] While a virtual machine executing an operating system is described herein as one example of an execution environment, other execution environments are also possible. For example, tasks or other processes may be executed within a software “container,” which provides a runtime environment without itself providing virtualization of hardware. Containers may be implemented within virtual machines to provide additional security, or may be run outside of a virtual machine instance.

[0030] The foregoing aspects and many of the attendant advantages of this disclosure will become more readily appreciated as the same become better understood by reference to the following description, when taken in conjunction with the accompanying drawings.

[0031] FIG. 1 is a block diagram of an illustrative operating environment 100 in which a service provider system 110 operates to enable client devices 102 to perform I/O operations on objects stored within an object storage service 160 and to apply path modifications to such

I/O operations, which modifications may include execution of user-defined code on an on-demand code execution system 120.

[0032] By way of illustration, various example client devices 102 are shown in communication with the service provider system 110, including a desktop computer, laptop, and a mobile phone. In general, the client devices 102 can be any computing device such as a desktop, laptop or tablet computer, personal computer, wearable computer, server, personal digital assistant (PDA), hybrid PDA/mobile phone, mobile phone, electronic book reader, set-top box, voice command device, camera, digital media player, and the like.

[0033] Generally described, the object storage service 160 can operate to enable clients to read, write, modify, and delete data objects, each of which represents a set of data associated with an identifier (an “object identifier” or “resource identifier”) that can be interacted with as an individual resource. For example, an object may represent a single file submitted by a client device 102 (though the object storage service 160 may or may not store such an object as a single file). This object-level interaction can be contrasted with other types of storage services, such as block-based storage services providing data manipulation at the level of individual blocks or database storage services providing data manipulation at the level of tables (or parts thereof) or the like.

[0034] The object storage service 160 illustratively includes one or more frontends 162, which provide an interface (a command-line interface (CLIs), application programming interface (APIs), or other programmatic interface) through which client devices 102 can interface with the service 160 to configure the service 160 on their behalf and to perform I/O operations on the service 160. For example, a client device 102 may interact with a frontend 162 to create a collection of data objects on the service 160 (e.g., a “bucket” of objects) and to configure permissions for that collection. Client devices 102 may thereafter create, read, update, or delete objects within the collection based on the interfaces of the frontends 162. In one embodiment, the frontend 162 provides a REST-compliant HTTP interface supporting a variety of request methods, each of which corresponds to a requested I/O operation on the service 160. By way of non-limiting example, request methods may include:

- a GET operation requesting retrieval of an object stored on the service 160 by reference to an identifier of the object;

- a PUT operation requesting storage of an object to be stored on the service 160, including an identifier of the object and input data to be stored as the object;
- a DELETE operation requesting deletion of an object stored on the service 160 by reference to an identifier of the object; and
- a LIST operation requesting listing of objects within an object collection stored on the service 160 by reference to an identifier of the collection.

A variety of other operations may also be supported. For example, the service 160 may provide a POST operation similar to a PUT operation but associated with a different upload mechanism (e.g., a browser-based HTML upload), or a HEAD operation enabling retrieval of metadata for an object without retrieving the object itself. In some embodiments, the service 160 may enable operations that combine one or more of the above operations, or combining an operation with a native data manipulation. For example, the service 160 may provide a COPY operation enabling copying of an object stored on the service 160 to another object, which operation combines a GET operation with a PUT operation. As another example, the service 160 may provide a SELECT operation enabling specification of an SQL query to be applied to an object prior to returning the contents of that object, which combines an application of an SQL query to a data object (a native data manipulation) with a GET operation. As yet another example, the service 160 may provide a “byte range” GET, which enables a GET operation on only a portion of a data object. In some instances, the operation requested by a client device 102 on the service 160 may be transmitted to the service via an HTTP request, which itself may include an HTTP method. In some cases, such as in the case of a GET operation, the HTTP method specified within the request may match the operation requested at the service 160. However, in other cases, the HTTP method of a request may not match the operation requested at the service 160. For example, a request may utilize an HTTP POST method to transmit a request to implement a SELECT operation at the service 160.

[0035] During general operation, frontends 162 may be configured to obtain a call to a request method, and apply that request method to input data for the method. For example, a frontend 162 can respond to a request to PUT input data into the service 160 as an object by storing that input data as the object on the service 160. Objects may be stored, for example, on object data stores 168, which correspond to any persistent or substantially persistent storage (including hard disk drives (HDDs), solid state drives (SSDs), network accessible storage

(NAS), storage area networks (SANs), non-volatile random access memory (NVRAM), or any of a variety of storage devices known in the art). As a further example, the frontend 162 can respond to a request to GET an object from the service 160 by retrieving the object from the stores 168 (the object representing input data to the GET resource request), and returning the object to a requesting client device 102.

[0036] In some cases, calls to a request method may invoke one or more native data manipulations provided by the service 160. For example, a SELECT operation may provide an SQL-formatted query to be applied to an object (also identified within the request), or a GET operation may provide a specific range of bytes of an object to be returned. The service 160 illustratively includes an object manipulation engine 170 configured to perform native data manipulations, which illustratively corresponds to a device configured with software executable to implement native data manipulations on the service 160 (e.g., by stripping non-selected bytes from an object for a byte-range GET, by applying an SQL query to an object and returning results of the query, etc.).

[0037] In accordance with embodiments of the present disclosure, the service 160 can further be configured to enable modification of an I/O path for a given object or collection of objects, such that a called request method is applied to an output of a data manipulation function, rather than the resource identified within the call. For example, the service 160 may enable a client device 102 to specify that GET operations for a given object should be subject to execution of a user-defined task on the on-demand code execution system 120, such that the data returned in response to the operation is the output of a task execution rather than the requested object. Similarly, the service 160 may enable a client device 102 to specify that PUT operations to store a given object should be subject to execution of a user-defined task on the on-demand code execution system 120, such that the data stored in response to the operation is the output of a task execution rather than the data provided for storage by a client device 102. As will be discussed in more detail below, path modifications may include specification of a pipeline of data manipulations, including native data manipulations, task-based manipulations, or combinations thereof. Illustratively, a client device 102 may specify a pipeline or other data manipulation for an object or object collection through the frontend 162, which may store a record of the pipeline or manipulation in the I/O path modification data store 164, which store 164, like the object data stores 168, can represent any persistent or substantially persistent



storage. While shown as distinct in FIG. 1, in some instances the data stores 164 and 168 may represent a single collection of data stores. For example, data modifications to objects or collections may themselves be stored as objects on the service 160.

[0038] To enable data manipulation via execution of user-defined code, the system further includes an on-demand code execution system 120. In one embodiment, the system 120 is solely usable by the object storage service 160 in connection with data manipulations of an I/O path. In another embodiment, the system 120 is additionally accessible by client devices 102 to directly implement serverless task executions. For example, the on-demand code execution system 120 may provide the service 160 (and potentially client devices 102) with one or more user interfaces, command-line interfaces (CLIs), application programming interfaces (APIs), or other programmatic interfaces for generating and uploading user-executable code (e.g., including metadata identifying dependency code objects for the uploaded code), invoking the user-provided code (e.g., submitting a request to execute the user codes on the on-demand code execution system 120), scheduling event-based jobs or timed jobs, tracking the user-provided code, or viewing other logging or monitoring information related to their requests or user codes. Although one or more embodiments may be described herein as using a user interface, it should be appreciated that such embodiments may, additionally or alternatively, use any CLIs, APIs, or other programmatic interfaces.

[0039] The client devices 102, object storage service 160, and on-demand code execution system 120 may communicate via a network 104, which may include any wired network, wireless network, or combination thereof. For example, the network 104 may be a personal area network, local area network, wide area network, over-the-air broadcast network (e.g., for radio or television), cable network, satellite network, cellular telephone network, or combination thereof. As a further example, the network 104 may be a publicly accessible network of linked networks, possibly operated by various distinct parties, such as the Internet. In some embodiments, the network 104 may be a private or semi-private network, such as a corporate or university intranet. The network 104 may include one or more wireless networks, such as a Global System for Mobile Communications (GSM) network, a Code Division Multiple Access (CDMA) network, a Long Term Evolution (LTE) network, or any other type of wireless network. The network 104 can use protocols and components for communicating via the Internet or any of the other aforementioned types of networks. For example, the

protocols used by the network 104 may include Hypertext Transfer Protocol (HTTP), HTTP Secure (HTTPS), Message Queue Telemetry Transport (MQTT), Constrained Application Protocol (CoAP), and the like. Protocols and components for communicating via the Internet or any of the other aforementioned types of communication networks are well known to those skilled in the art and, thus, are not described in more detail herein.

[0040] To enable interaction with the on-demand code execution system 120, the system 120 includes one or more frontends 130, which enable interaction with the on-demand code execution system 120. In an illustrative embodiment, the frontends 130 serve as a “front door” to the other services provided by the on-demand code execution system 120, enabling users (via client devices 102) or the service 160 to provide, request execution of, and view results of computer executable code. The frontends 130 include a variety of components to enable interaction between the on-demand code execution system 120 and other computing devices. For example, each frontend 130 may include a request interface providing client devices 102 and the service 160 with the ability to upload or otherwise communicate user-specified code to the on-demand code execution system 120 and to thereafter request execution of that code. In one embodiment, the request interface communicates with external computing devices (e.g., client devices 102, frontend 162, etc.) via a graphical user interface (GUI), CLI, or API. The frontends 130 process the requests and make sure that the requests are properly authorized. For example, the frontends 130 may determine whether the user associated with the request is authorized to access the user code specified in the request.

[0041] References to user code as used herein may refer to any program code (e.g., a program, routine, subroutine, thread, etc.) written in a specific program language. In the present disclosure, the terms “code,” “user code,” and “program code,” may be used interchangeably. Such user code may be executed to achieve a specific function, for example, in connection with a particular data transformation developed by the user. As noted above, individual collections of user code (e.g., to achieve a specific function) are referred to herein as “tasks,” while specific executions of that code (including, e.g., compiling code, interpreting code, or otherwise making the code executable) are referred to as “task executions” or simply “executions.” Tasks may be written, by way of non-limiting example, in JavaScript (e.g., node.js), Java, Python, or Ruby (or another programming language).

[0042] To manage requests for code execution, the frontend 130 can include an execution queue, which can maintain a record of requested task executions. Illustratively, the number of simultaneous task executions by the on-demand code execution system 120 is limited, and as such, new task executions initiated at the on-demand code execution system 120 (e.g., via an API call, via a call from an executed or executing task, etc.) may be placed on the execution queue and processed, e.g., in a first-in-first-out order. In some embodiments, the on-demand code execution system 120 may include multiple execution queues, such as individual execution queues for each user account. For example, users of the service provider system 110 may desire to limit the rate of task executions on the on-demand code execution system 120 (e.g., for cost reasons). Thus, the on-demand code execution system 120 may utilize an account-specific execution queue to throttle the rate of simultaneous task executions by a specific user account. In some instances, the on-demand code execution system 120 may prioritize task executions, such that task executions of specific accounts or of specified priorities bypass or are prioritized within the execution queue. In other instances, the on-demand code execution system 120 may execute tasks immediately or substantially immediately after receiving a call for that task, and thus, the execution queue may be omitted.

[0043] The frontend 130 can further include an output interface configured to output information regarding the execution of tasks on the on-demand code execution system 120. Illustratively, the output interface may transmit data regarding task executions (e.g., results of a task, errors related to the task execution, or details of the task execution, such as total time required to complete the execution, total data processed via the execution, etc.) to the client devices 102 or the object storage service 160.

[0044] In some embodiments, the on-demand code execution system 120 may include multiple frontends 130. In such embodiments, a load balancer may be provided to distribute the incoming calls to the multiple frontends 130, for example, in a round-robin fashion. In some embodiments, the manner in which the load balancer distributes incoming calls to the multiple frontends 130 may be based on the location or state of other components of the on-demand code execution system 120. For example, a load balancer may distribute calls to a geographically nearby frontend 130, or to a frontend with capacity to service the call. In instances where each frontend 130 corresponds to an individual instance of another component of the on-demand code execution system 120, such as the active pool 148 described below, the

load balancer may distribute calls according to the capacities or loads on those other components. Calls may in some instances be distributed between frontends 130 deterministically, such that a given call to execute a task will always (or almost always) be routed to the same frontend 130. This may, for example, assist in maintaining an accurate execution record for a task, to ensure that the task executes only a desired number of times. For example, calls may be distributed to load balance between frontends 130. Other distribution techniques, such as anycast routing, will be apparent to those of skill in the art.

[0045] The on-demand code execution system 120 further includes one or more worker managers 140 that manage the execution environments, such as virtual machine instances 150 (shown as VM instance 150A and 150B, generally referred to as a “VM”), used for servicing incoming calls to execute tasks. While the following will be described with reference to virtual machine instances 150 as examples of such environments, embodiments of the present disclosure may utilize other environments, such as software containers. In the example illustrated in FIG. 1, each worker manager 140 manages an active pool 148, which is a group (sometimes referred to as a pool) of virtual machine instances 150 executing on one or more physical host computing devices that are initialized to execute a given task (e.g., by having the code of the task and any dependency data objects loaded into the instance).

[0046] Although the virtual machine instances 150 are described here as being assigned to a particular task, in some embodiments, the instances may be assigned to a group of tasks, such that the instance is tied to the group of tasks and any tasks of the group can be executed within the instance. For example, the tasks in the same group may belong to the same security group (e.g., based on their security credentials) such that executing one task in a container on a particular instance 150 after another task has been executed in another container on the same instance does not pose security risks. As discussed below, a task may be associated with permissions encompassing a variety of aspects controlling how a task may execute. For example, permissions of a task may define what network connections (if any) can be initiated by an execution environment of the task. As another example, permissions of a task may define what authentication information is passed to a task, controlling what network-accessible resources are accessible to execution of a task (e.g., objects on the service 160). In one embodiment, a security group of a task is based on one or more such permissions. For example, a security group may be defined based on a combination of permissions to initiate network

connections and permissions to access network resources. As another example, the tasks of the group may share common dependencies, such that an environment used to execute one task of the group can be rapidly modified to support execution of another task within the group.

[0047] Once a triggering event to execute a task has been successfully processed by a frontend 130, the frontend 130 passes a request to a worker manager 140 to execute the task. In one embodiment, each frontend 130 may be associated with a corresponding worker manager 140 (e.g., a worker manager 140 co-located or geographically nearby to the frontend 130) and thus, the frontend 130 may pass most or all requests to that worker manager 140. In another embodiment, a frontend 130 may include a location selector configured to determine a worker manager 140 to which to pass the execution request. In one embodiment, the location selector may determine the worker manager 140 to receive a call based on hashing the call, and distributing the call to a worker manager 140 selected based on the hashed value (e.g., via a hash ring). Various other mechanisms for distributing calls between worker managers 140 will be apparent to one of skill in the art.

[0048] Thereafter, the worker manager 140 may modify a virtual machine instance 150 (if necessary) and execute the code of the task within the instance 150. As shown in FIG. 1, respective instances 150 may have operating systems (OS) 152 (shown as OS 152A and 152B), language runtimes 154 (shown as runtime 154A and 154B), and user code 156 (shown as user code 156A and 156B). The OS 152, runtime 154, and user code 156 may collectively enable execution of the user code to implement the task. Thus, via operation of the on-demand code execution system 120, tasks may be rapidly executed within an execution environment.

[0049] In accordance with aspects of the present disclosure, each VM 150 additionally includes staging code 157 executable to facilitate staging of input data on the VM 150 and handling of output data written on the VM 150, as well as a VM data store 158 accessible through a local file system of the VM 150. Illustratively, the staging code 157 represents a process executing on the VM 150 (or potentially a host device of the VM 150) and configured to obtain data from the object storage service 160 and place that data into the VM data store 158. The staging code 157 can further be configured to obtain data written to a file within the VM data store 158, and to transmit that data to the object storage service 160. Because such data is available at the VM data store 158, user code 156 is not required to obtain data over a network, simplifying user code 156 and enabling further restriction of network

communications by the user code 156, thus increasing security. Rather, as discussed above, user code 156 may interact with input data and output data as files on the VM data store 158, by use of file handles passed to the code 156 during an execution. In some embodiments, input and output data may be stored as files within a kernel-space file system of the data store 158. In other instances, the staging code 157 may provide a virtual file system, such as a filesystem in userspace (FUSE) interface, which provides an isolated file system accessible to the user code 156, such that the user code's access to the VM data store 158 is restricted.

[0050] As used herein, the term "local file system" generally refers to a file system as maintained within an execution environment, such that software executing within the environment can access data as file, rather than via a network connection. In accordance with aspects of the present disclosure, the data storage accessible via a local file system may itself be local (e.g., local physical storage), or may be remote (e.g., accessed via a network protocol, like NFS, or represented as a virtualized block device provided by a network-accessible service). Thus, the term "local file system" is intended to describe a mechanism for software to access data, rather than physical location of the data.

[0051] The VM data store 158 can include any persistent or non-persistent data storage device. In one embodiment, the VM data store 158 is physical storage of the host device, or a virtual disk drive hosted on physical storage of the host device. In another embodiment, the VM data store 158 is represented as local storage, but is in fact a virtualized storage device provided by a network accessible service. For example, the VM data store 158 may be a virtualized disk drive provided by a network-accessible block storage service. In some embodiments, the object storage service 160 may be configured to provide file-level access to objects stored on the data stores 168, thus enabling the VM data store 158 to be virtualized based on communications between the staging code 157 and the service 160. For example, the object storage service 160 can include a file-level interface 166 providing network access to objects within the data stores 168 as files. The file-level interface 166 may, for example, represent a network-based file system server (e.g., a network file system (NFS)) providing access to objects as files, and the staging code 157 may implement a client of that server, thus providing file-level access to objects of the service 160.

[0052] In some instances, the VM data store 158 may represent virtualized access to another data store executing on the same host device of a VM instance 150. For example, an

active pool 148 may include one or more data staging VM instances (not shown in FIG. 1), which may be co-tenanted with VM instances 150 on the same host device. A data staging VM instance may be configured to support retrieval and storage of data from the service 160 (e.g., data objects or portions thereof, input data passed by client devices 102, etc.), and storage of that data on a data store of the data staging VM instance. The data staging VM instance may, for example, be designated as unavailable to support execution of user code 156, and thus be associated with elevated permissions relative to instances 150 supporting execution of user code. The data staging VM instance may make this data accessible to other VM instances 150 within its host device (or, potentially, on nearby host devices), such as by use of a network-based file protocol, like NFS. Other VM instances 150 may then act as clients to the data staging VM instance, enabling creation of virtualized VM data stores 158 that, from the point of view of user code 156A, appear as local data stores. Beneficially, network-based access to data stored at a data staging VM can be expected to occur very quickly, given the co-location of a data staging VM and a VM instance 150 within a host device or on nearby host devices.

[0053] While some examples are provided herein with respect to use of IO stream handles to read from or write to a VM data store 158, IO streams may additionally be used to read from or write to other interfaces of a VM instance 150 (while still removing a need for user code 156 to conduct operations other than stream-level operations, such as creating network connections). For example, staging code 157 may “pipe” input data to an execution of user code 156 as an input stream, the output of which may be “piped” to the staging code 157 as an output stream. As another example, a staging VM instance or a hypervisor to a VM instance 150 may pass input data to a network port of the VM instance 150, which may be read-from by staging code 157 and passed as an input stream to the user code 157. Similarly, data written to an output stream by the task code 156 may be written to a second network port of the instance 150A for retrieval by the staging VM instance or hypervisor. In yet another example, a hypervisor to the instance 150 may pass input data as data written to a virtualized hardware input device (e.g., a keyboard) and staging code 157 may pass to the user code 156 a handle to the IO stream corresponding to that input device. The hypervisor may similarly pass to the user code 156 a handle for an IO stream corresponding to a virtualized hardware output device, and read data written to that stream as output data. Thus, the examples provided herein with respect to file streams may generally be modified to relate to any IO stream.

[0054] The object storage service 160 and on-demand code execution system 120 are depicted in FIG. 1 as operating in a distributed computing environment including several computer systems that are interconnected using one or more computer networks (not shown in FIG. 1). The object storage service 160 and on-demand code execution system 120 could also operate within a computing environment having a fewer or greater number of devices than are illustrated in FIG. 1. Thus, the depiction of the object storage service 160 and on-demand code execution system 120 in FIG. 1 should be taken as illustrative and not limiting to the present disclosure. For example, the on-demand code execution system 120 or various constituents thereof could implement various Web services components, hosted or “cloud” computing environments, or peer to peer network configurations to implement at least a portion of the processes described herein. In some instances, the object storage service 160 and on-demand code execution system 120 may be combined into a single service. Further, the object storage service 160 and on-demand code execution system 120 may be implemented directly in hardware or software executed by hardware devices and may, for instance, include one or more physical or virtual servers implemented on physical computer hardware configured to execute computer executable instructions for performing various features that will be described herein. The one or more servers may be geographically dispersed or geographically co-located, for instance, in one or more data centers. In some instances, the one or more servers may operate as part of a system of rapidly provisioned and released computing resources, often referred to as a “cloud computing environment.”

[0055] In the example of FIG. 1, the object storage service 160 and on-demand code execution system 120 are illustrated as connected to the network 104. In some embodiments, any of the components within the object storage service 160 and on-demand code execution system 120 can communicate with other components of the on-demand code execution system 120 via the network 104. In other embodiments, not all components of the object storage service 160 and on-demand code execution system 120 are capable of communicating with other components of the virtual environment 100. In one example, only the frontends 130 and 162 (which may in some instances represent multiple frontends) may be connected to the network 104, and other components of the object storage service 160 and on-demand code execution system 120 may communicate with other components of the environment 100 via the respective frontends 130 and 162.



[0056] While some functionalities are generally described herein with reference to an individual component of the object storage service 160 and on-demand code execution system 120, other components or a combination of components may additionally or alternatively implement such functionalities. For example, while the object storage service 160 is depicted in FIG. 1 as including an object manipulation engine 170, functions of that engine 170 may additionally or alternatively be implemented as tasks on the on-demand code execution system 120. Moreover, while the on-demand code execution system 120 is described as an example system to apply data manipulation tasks, other compute systems may be used to execute user-defined tasks, which compute systems may include more, fewer or different components than depicted as part of the on-demand code execution system 120. In a simplified example, the object storage service 160 may include a physical computing device configured to execute user-defined tasks on demand, thus representing a compute system usable in accordance with embodiments of the present disclosure. Thus, the specific configuration of elements within FIG. 1 is intended to be illustrative.

[0057] FIG. 2 depicts a general architecture of a frontend server 200 computing device implementing a frontend 162 of FIG. 1. The general architecture of the frontend server 200 depicted in FIG. 2 includes an arrangement of computer hardware and software that may be used to implement aspects of the present disclosure. The hardware may be implemented on physical electronic devices, as discussed in greater detail below. The frontend server 200 may include many more (or fewer) elements than those shown in FIG. 2. It is not necessary, however, that all of these generally conventional elements be shown in order to provide an enabling disclosure. Additionally, the general architecture illustrated in FIG. 2 may be used to implement one or more of the other components illustrated in FIG. 1.

[0058] As illustrated, the frontend server 200 includes a processing unit 290, a network interface 292, a computer readable medium drive 294, and an input/output device interface 296, all of which may communicate with one another by way of a communication bus. The network interface 292 may provide connectivity to one or more networks or computing systems. The processing unit 290 may thus receive information and instructions from other computing systems or services via the network 104. The processing unit 290 may also communicate to and from primary memory 280 or secondary memory 298 and further provide output information for an optional display (not shown) via the input/output device interface

296. The input/output device interface 296 may also accept input from an optional input device (not shown).

[0059] The primary memory 280 or secondary memory 298 may contain computer program instructions (grouped as units in some embodiments) that the processing unit 290 executes in order to implement one or more aspects of the present disclosure. These program instructions are shown in FIG. 2 as included within the primary memory 280, but may additionally or alternatively be stored within secondary memory 298. The primary memory 280 and secondary memory 298 correspond to one or more tiers of memory devices, including (but not limited to) RAM, 3D XPOINT memory, flash memory, magnetic storage, and the like. The primary memory 280 is assumed for the purposes of description to represent a main working memory of the worker manager 140, with a higher speed but lower total capacity than secondary memory 298.

[0060] The primary memory 280 may store an operating system 284 that provides computer program instructions for use by the processing unit 290 in the general administration and operation of the frontend server 200. The memory 280 may further include computer program instructions and other information for implementing aspects of the present disclosure. For example, in one embodiment, the memory 280 includes a user interface unit 282 that generates user interfaces (or instructions therefor) for display upon a computing device, e.g., via a navigation or browsing interface such as a browser or application installed on the computing device.

[0061] In addition to or in combination with the user interface unit 282, the memory 280 may include a control plane unit 286 and data plane unit 288 each executable to implement aspects of the present disclosure. Illustratively, the control plane unit 286 may include code executable to enable owners of data objects or collections of objects to attach manipulations, serverless functions, or data processing pipelines to an I/O path, in accordance with embodiments of the present disclosure. For example, the control plane unit 286 may enable the frontend 162 to implement the interactions of FIG. 3. The data plane unit 288 may illustratively include code enabling handling of I/O operations on the object storage service 160, including implementation of manipulations, serverless functions, or data processing pipelines attached to an I/O path (e.g., via the interactions of FIGS. 5A-6B, implementation of the routines of FIGS. 7-8, etc.).

[0062] The frontend server 200 of FIG. 2 is one illustrative configuration of such a device, of which others are possible. For example, while shown as a single device, a frontend server 200 may in some embodiments be implemented as multiple physical host devices. Illustratively, a first device of such a frontend server 200 may implement the control plane unit 286, while a second device may implement the data plane unit 288.

[0063] While described in FIG. 2 as a frontend server 200, similar components may be utilized in some embodiments to implement other devices shown in the environment 100 of FIG. 1. For example, a similar device may implement a worker manager 140, as described in more detail in U.S. Patent No. 9,323,556, entitled “PROGRAMMATIC EVENT DETECTION AND MESSAGE GENERATION FOR REQUESTS TO EXECUTE PROGRAM CODE,” and filed September 30, 2014 (the “556 Patent”), the entirety of which is hereby incorporated by reference.

[0064] With reference to FIG. 3, illustrative interactions are depicted for enabling a client device 102A to modify an I/O path for one or more objects on an object storage service 160 by inserting a data manipulation into the I/O path, which manipulation is implemented within a task executable on the on-demand code execution system 120.

[0065] The interactions of FIG. 3 begin at (1), where the client device 102A authors the stream manipulation code. The code can illustratively function to access an input file handle provided on execution of the program (which may, for example, be represented by the standard input stream for a program, commonly “stdin”), perform manipulations on data obtained from that file handle, and write data to an output file handle provided on execution of the program (which may, for example, be represented by the standard output stream for a program, commonly “stdout”).

[0066] While examples are discussed herein with respect to a “file” handle, embodiments of the present disclosure may utilize handles providing access to any operating-system-level input/output (IO) stream, examples of which include byte streams, character streams, file streams, and the like. As used herein, the term operating-system-level input/output stream (or simply an “IO stream”) is intended to refer to a stream of data for which an operating system provides a defined set of functions, such as seeking within the stream, reading from a stream, and writing to a stream. Streams may be created in various manners. For example, a programming language may generate a stream by use of a function library to

open a file on a local operating system, or a stream may be created by use of a “pipe” operator (e.g., within an operating system shell command language). As will be appreciated by one skilled in the art, most general purpose programming languages include, as basic functionality of the code, the ability to interact with streams.

[0067] In accordance with embodiments of the present disclosure, task code may be authored to accept, as a parameter of the code, an input handle and an output handle, both representing IO streams (e.g., an input stream and an output stream, respectively). The code may then manipulate data of the input stream, and write an output to the output stream. Given use of a general purpose programming language, any of a variety of functions may be implemented according to the desires of the user. For example, a function may search for and remove confidential information from the input stream. While some code may utilize only input and output handles, other code may implement additional interfaces, such as network communication interfaces. However, by providing the code with access to input and output streams (via respective handles) created outside of the code, the need for the code to create such streams is removed. Moreover, because streams may be created outside of the code, and potentially outside of an execution environment of the code, stream manipulation code need not necessarily be trusted to conduct certain operations that may be necessary to create a stream. For example, a stream may represent information transmitted over a network connection, without the code being provided with access to that network connection. Thus, use of IO streams to pass data into and out of code executions can simplify code while increasing security.

[0068] As noted above, the code may be authored in a variety of programming languages. Authoring tools for such languages are known in the art and thus will not be described herein. While authoring is described in FIG. 3 as occurring on the client device 102A, the service 160 may in some instances provide interfaces (e.g., web GUIs) through which to author or select code.

[0069] At (2), the client device 102A submits the stream manipulation code to the frontend 162 of the service 160, and requests that an execution of the code be inserted into an I/O path for one or more objects. Illustratively, the frontends 162 may provide one or more interfaces to the device 102A enabling submission of the code (e.g., as a compressed file). The frontends 162 may further provide interfaces enabling designation of one or more I/O paths to

which an execution of the code should be applied. Each I/O path may correspond, for example, to an object or collection of objects (e.g., a “bucket” of objects). In some instances, an I/O path may further corresponding to a given way of accessing such object or collection (e.g., a URI through which the object is created), to one or more accounts attempting to access the object or collection, or to other path criteria. Designation of the path modification is then stored in the I/O path modification data store 164, at (3). Additionally, the stream manipulation code is stored within the object data stores 166 at (4).

[0070] As such, when an I/O request is received via the specified I/O path, the service 160 is configured to execute the stream manipulation code against input data for the request (e.g., data provided by the client device 102A or an object of the service 160, depending on the I/O request), before then applying the request to the output of the code execution. In this manner, a client device 102A (which in FIG. 3 illustratively represents an owner of an object or object collection) can obtain greater control over data stored on and retrieved from the object storage service 160.

[0071] The interactions of FIG. 3 generally relate to insertion of a single data manipulation into the I/O path of an object or collection on the service 160. However, in some embodiments of the present disclosure an owner of an object or collection is enabled to insert multiple data manipulations into such an I/O path. Each data manipulation may correspond, for example, to a serverless code-based manipulation or a native manipulation of the service 160. For example, assume an owner has submitted a data set to the service 160 as an object, and that the owner wishes to provide an end user with a filtered view of a portion of that data set. While the owner could store that filtered view of the portion as a separate object and provide the end user with access to that separate object, this results in data duplication on the service 160. In the case that the owner wishes to provide multiple end users with different portions of the data set, potentially with customized filters, that data duplication grows, resulting in significant inefficiencies. In accordance with the present disclosure, another option may be for the owner to author or obtain custom code to implement different filters on different portions of the object, and to insert that code into the I/O path for the object. However, this approach may require the owner to duplicate some native functionality of the service 160 (e.g., an ability to retrieve a portion of a data set). Moreover, this approach would inhibit modularity

and reusability of code, since a single set of code would be required to conduct two functions (e.g., selecting a portion of the data and filtering that portion).

[0072] To address these shortcomings, embodiments of the present disclosure enable an owner to create a pipeline of data manipulations to be applied to an I/O path, linking together multiple data manipulations, each of which may also be inserted into other I/O paths. An illustrative visualization of such a pipeline is shown in FIG. 4 as pipeline 400. Specifically, the pipeline 400 illustrates a series of data manipulations that an owner specifies are to occur on calling of a request method against an object or object collection. As shown in FIG. 4, the pipeline begins with input data, specified within the call according to a called request method. For example, a PUT call may generally include the input data as the data to be stored, while a GET call may generally include the input data by reference to a stored object. A LIST call may specify a directory, a manifest of which is the input data to the LIST request method.

[0073] Contrary to typical implementations of request methods, in the illustrative pipeline 400, the called request method is not initially applied to the input data. Rather, the input data is initially passed to an execution of “code A” 404, where code A represents a first set of user-authored code. The output of that execution is then passed to “native function A” 406, which illustratively represents a native function of the service 160, such as a “SELECT” or byte-range function implemented by the object manipulation engine 170. The output of that native function 406 is then passed to an execution of “code B” 408, which represents a second set of user-authored code. Thereafter, the output of that execution 408 is passed to the called request method 410 (e.g., GET, PUT, LIST, etc.). Accordingly, rather than the request method being applied to the input data as in conventional techniques, in the illustration of FIG. 4, the request method is applied to the output of the execution 408, which illustratively represents a transformation of the input data according to one or more owner-specified manipulations 412. Notably, implementation of the pipeline 400 may not require any action or imply any knowledge of the pipeline 400 on the part of a calling client device 102. As such, implementation of pipelines can be expected not to impact existing mechanisms of interacting with the service 160 (other than altering the data stored on or retrieved from the service 160 in accordance with the pipeline). For example, implementation of a pipeline can be expected not to require reconfiguration of existing programs utilizing an API of the service 160.

[0074] While the pipeline 400 of FIG. 4 is linear, in some embodiments the service 160 may enable an owner to configure non-linear pipelines, such as by include conditional or branching nodes within the pipeline. Illustratively, as described in more detail below, data manipulations (e.g., serverless-based functions) can be configured to include a return value, such as an indication of successful execution, encountering an error, etc. In one example, the return value of a data manipulation may be used to select a conditional branch within a branched pipeline, such that a first return value causes the pipeline to proceed on a first branch, while a second return value causes the pipeline to proceed on a second branch. In some instances, pipelines may include parallel branches, such that data is copied or divided to multiple data manipulations, the outputs of which are passed to a single data manipulation for merging prior to executing the called method. The service 160 may illustratively provide a graphical user interface through which owners can create pipelines, such as by specifying nodes within the pipeline and linking those nodes together via logical connections. A variety of flow-based development interfaces are known and may be utilized in conjunction with aspects of the present disclosure.

[0075] Furthermore, in some embodiments, a pipeline applied to a particular I/O path may be generated on-the-fly, at the time of a request, based on data manipulations applied to the path according to different criteria. For example, an owner of a data collection may apply a first data manipulation to all interactions with objects within a collection, and a second data manipulation to all interactions obtained via a given URI. Thus, when a request is received to interact with an object within the collection and via the given URI, the service 160 may generate a pipeline combining the first and second data manipulations. The service 160 may illustratively implement a hierarchy of criteria, such that manipulations applied to objects are placed within the pipeline prior to manipulations applied to a URI, etc.

[0076] In some embodiments, client devices 102 may be enabled to request inclusion of a data manipulation within a pipeline. For example, within parameters of a GET request, a client device 102 may specify a particular data manipulation to be included within a pipeline applied in connection with the request. Illustratively, a collection owner may specify one or more data manipulations allowed for the collection, and further specify identifiers for those manipulations (e.g., function names). Thus, when requesting to interact with the collection, a client device 102 may specify the identifier to cause the manipulation to be included within a

pipeline applied to the I/O path. In one embodiment, client-requested manipulations are appended to the end of a pipeline subsequent to owner-specified data manipulations and prior to implementing the requested request method. For example, where a client device 102 requests to GET a data set, and requests that a search function be applied to the data set before the GET method is implemented, the search function can receive as input data the output of an owner-specified data manipulations for the data set (e.g., manipulations to remove confidential information from the data set). In addition, requests may in some embodiments specify parameters to be passed to one or more data manipulations (whether specified within the request or not). Accordingly, while embodiments of the present disclosure can enable data manipulations without knowledge of those manipulations on the part of client devices 102, other embodiments may enable client devices 102 to pass information within an I/O request for use in implementing data manipulations.

[0077] Moreover, while example embodiments of the present disclosure are discussed with respect to manipulation of input data to a called method, embodiments of the present disclosure may further be utilized to modify aspects of a request, including a called method. For example, a serverless task execution may be passed the content of a request (including, e.g., a called method and parameters) and be configured to modify and return, as a return value to a frontend 162, a modified version of the method or parameters. Illustratively, where a client device 102 is authenticated as a user with access to only a portion of a data object, a serverless task execution may be passed a call to “GET” that data object, and may transform parameters of the GET request such that it applies only to a specific byte range of the data object corresponding to the portion that the user may access. As a further example, tasks may be utilized to implement customized parsing or restrictions on called methods, such as by limiting the methods a user may call, the parameters to those methods, or the like. In some instances, application of one or more functions to a request (e.g., to modify the method called or method parameters) may be viewed as a “pre-data processing” pipeline, and may thus be implemented prior to obtaining the input data within the pipeline 400 (which input data may change due to changes in the request), or may be implemented independently of a data manipulation pipeline 400.

[0078] Similarly, while example embodiments of the present disclosure are discussed with respect to application of a called method to output data of one or more data manipulations,



in some embodiments manipulations can additionally or alternatively occur after application of a called method. For example, a data object may contain sensitive data that a data owner desires to remove prior to providing the data to a client. The owner may further enable a client to specify native manipulations to the data set, such as conducting a database query on the dataset (e.g., via a SELECT resource method). While the owner may specify a pipeline for the data set to cause filtering of sensitive data to be conducted prior to application of the SELECT method, such an order of operations may be undesirable, as filtering may occur with respect to the entire data object rather than solely the portion returned by the SELECT query. Accordingly, additionally or alternatively to specifying manipulations that occur prior to satisfying a request method, embodiments of the present disclosure can enable an owner to specify manipulations to occur subsequent to application of a called method but prior to conducting a final operation to satisfy a request. For example, in the case of a SELECT operation, the service 160 may first conduct the SELECT operation against specified input data (e.g., a data object), and then pass the output of that SELECT operation to a data manipulation, such as a serverless task execution. The output of that execution can then be returned to a client device 102 to satisfy the request.

[0079] While FIG. 3 and FIG. 4 are generally described with reference to serverless tasks authored by an owner of an object or collection, in some instances the service 160 may enable code authors to share their tasks with other users of the service 160, such that code of a first user is executed in the I/O path of an object owned by a second user. The service 160 may also provide a library of tasks for use by each user. In some cases, the code of a shared task may be provided to other users. In other cases, the code of the shared task may be hidden from other users, such that the other users can execute the task but not view code of the task. In these cases, other users may illustratively be enabled to modify specific aspects of code execution, such as the permissions under which the code will execute.

[0080] With reference to FIGS. 5A and 5B, illustrative interactions will be discussed for applying a modification to an I/O path for a request to store an object on the service 160, which request is referred to in connection with these figures as a “PUT” request or “PUT object call.” While shown in two figures, numbering of interactions is maintained across FIGS. 5A and 5B.

[0081] The interactions begin at (1), where a client device 102A submits a PUT object call to the storage service 160, corresponding to a request to store input data (e.g., included or specified within the call) on the service 160. The input data may correspond, for example, to a file stored on the client device 102A. As shown in FIG. 5A, the call is directed to a frontend 162 of the service 162 that, at (2), retrieves from the I/O path modification data store 164 an indication of modifications to the I/O path for the call. The indication may reflect, for example, a pipeline to be applied to calls received on the I/O path. The I/O path for a call may generally be specified with respect to a request method included within a call, an object or collection of objects indicated within the call, a specific mechanism of reaching the service 160 (e.g., protocol, URI used, etc.), an identity or authentication status of the client device 102A, or a combination thereof. For example, in FIG. 5A, the I/O path used can correspond to use of a PUT request method directed to a particular URI (e.g., associated with the frontend 162) to store an object in a particular logical location on the service 160 (e.g., a specific bucket). In FIGS. 5A and 5B, it is assumed that an owner of that logical location has previously specified a modification to the I/O path, and specifically, has specified that a serverless function should be applied to the input data before a result of that function is stored in the service 160.

[0082] Accordingly, at (3), the frontend 162 detects within the modifications for the I/O path inclusion of a serverless task execution. Thus, at (4), the frontend 162 submits a call to the on-demand code execution system 120 to execute the task specified within the modifications against the input data specified within the call.

[0083] The on-demand code execution system 120, at (5), therefore generates an execution environment 502 in which to execute code corresponding to the task. Illustratively, the call may be directed to a frontend 130 of the system, which may distribute instructions to a worker manager 140 to select or generate a VM instance 150 in which to execute the task, which VM instance 150 illustratively represents the execution environment 502. During generation of the execution environment 502, the system 120 further provisions the environment with code 504 of the task indicated within the I/O path modification (which may be retrieved, for example, from the object data stores 166). While not shown in FIG. 5A, the environment 502 further includes other dependencies of the code, such as access to an operating system, a runtime required to execute the code, etc.

[0084] In some embodiments, generation of the execution environment 502 can include configuring the environment 502 with security constraints limiting access to network resources. Illustratively, where a task is intended to conduct data manipulation without reference to network resources, the environment 502 can be configured with no ability to send or receive information via a network. Where a task is intended to utilize network resources, access to such resources can be provided on a “whitelist” basis, such that network communications from the environment 502 are allowed only for specified domains, network addresses, or the like. Network restrictions may be implemented, for example, by a host device hosting the environment 502 (e.g., by a hypervisor or host operating system). In some instances, network access requirements may be utilized to assist in placement of the environment 502, either logically or physically. For example, where a task requires no access to network resources, the environment 502 for the task may be placed on a host device that is distant from other network-accessible services of the service provider system 110, such as an “edge” device with a lower-quality communication channel to those services. Where a task requires access to otherwise private network services, such as services implemented within a virtual private cloud (e.g., a local-area-network-like environment implemented on the service 160 on behalf of a given user), the environment 502 may be created to exist logically within that cloud, such that a task execution 502 accesses resources within the cloud. In some instances, a task may be configured to execute within a private cloud of a client device 102 that submits an I/O request. In other instances, a task may be configured to execute within a private cloud of an owner of the object or collection referenced within the request.

[0085] In addition to generating the environment 502, at (6), the system 120 provisions the environment with stream-level access to an input file handle 506 and an output file handle 508, usable to read from and write to the input data and output data of the task execution, respectively. In one embodiment, files handle 506 and 508 may point to a (physical or virtual) block storage device (e.g., disk drive) attached to the environment 502, such that the task can interact with a local file system to read input data and write output data. For example, the environment 502 may represent a virtual machine with a virtual disk drive, and the system 120 may obtain the input data from the service 160 and store the input data on the virtual disk drive. Thereafter, on execution of the code, the system 120 may pass to the code a handle of the input data as stored on the virtual disk drive, and a handle of a file on the drive

to which to write output data. In another embodiment, files handle 506 and 508 may point to a network file system, such as an NFS-compatible file system, on which the input data has been stored. For example, the frontend 162 during processing of the call may store the input data as an object on the object data stores 166, and the file-level interface 166 may provide file-level access to the input data and to a file representing output data. In some cases, the file handles 506 and 508 may point to files on a virtual file system, such as a file system in user space. By providing handles 506 and 508, the task code 504 is enabled to read the input data and write output data using stream manipulations, as opposed to being required to implement network transmissions. Creation of the handles 506 and 508 (or streams corresponding to the handles) may illustratively be achieved by execution of staging code 157 within or associated with the environment 502.

[0086] The interactions of FIG. 5A are continued in FIG. 5B, where the system 120 executes the task code 504. As the task code 504 may be user-authored, any number of functionalities may be implemented within the code 504. However, for the purposes of description of FIGS. 5A and 5B, it will be assumed that the code 504, when executed, reads input data from the input file handle 506 (which may be passed as a commonly used input stream, such as stdin), manipulates the input data, and writes output data to the output file handle 508 (which may be passed as a commonly used output stream, such as stdout). Accordingly, at (8), the system 120 obtains data written to the output file (e.g., the file referenced in the output file handle) as output data of the execution. In addition, at (9), the system 120 obtains a return value of the code execution (e.g., a value passed in a final call of the function). For the purposes of description of FIGS. 5A and 5B, it will be assumed that the return value indicates success of the execution. At (10), the output data and the success return value are then passed to the frontend 162.

[0087] While shown as a single interaction in FIG. 5B, in some embodiments output data of a task execution and a return value of that execution may be returned separately. For example, during execution, task code 504 may write to an output file through the handle 508, and this data may be periodically or iteratively returned to the service 160. Illustratively, where the output file exists on a file system in user space implemented by staging code, the staging code may detect and forward each write to the output file to the frontend 162. Where the output file exists on a network file system, writes to the file may directly cause the written data to be

transmitted to the interface 166 and thus the service 160. In some instances, transmitting written data iteratively may reduce the amount of storage required locally to the environment 502, since written data can, according to some embodiments, be deleted from local storage of the environment 502.

[0088] In addition, while a success return value is assumed in FIGS. 5A and 5B, other types of return value are possible and contemplated. For example, an error return value may be used to indicate to the frontend 162 that an error occurred during execution of task code 504. As another example, user-defined return values may be used to control how conditional branching within a pipeline proceeds. In some cases, the return value may indicate to the frontend 162 a request for further processing. For example, a task execution may return to the frontend 162 a call to execute another serverless task (potentially not specified within a path modification for the current I/O path). Moreover, return values may specify to the frontend 162 what return value is to be returned to the client device 102A. For example, a typical PUT request method called at the service 160 may be expected to return an HTTP 200 code (“OK”). As such, a success return value from the task code may further indicate that the frontend 162 should return an HTTP 200 code to the client device 102A. An error return value may, for example, indicate that the frontend 162 should return a 3XX HTTP redirection or 4XX HTTP error code to the client device 102A. Still further, in some cases, return values may specify to the frontend 162 content of a return message to the client device 102A other than a return value. For example, the frontend 162 may be configured to return a given HTTP code (e.g., 200) for any request from the client device 102A that is successfully retrieved at the frontend 162 and invokes a data processing pipeline. A task execution may then be configured to specify, within its return value, data to be passed to the client device 102A in addition to that HTTP code. Such data may illustratively include structured data (e.g., extensible markup language (XML) data) providing information generated by the task execution, such as data indicating success or failure of the task. This approach may beneficially enable the frontend 162 to quickly respond to requests (e.g., without awaiting execution of a task) while still enabling a task execution to pass information to the client device 102.

[0089] For purposes of the present illustration, it will be assumed that the success return value of the task indicates that an HTTP 2XX success response should be passed to the device 102A. Accordingly, on receiving output data, the frontend 162 stores the output data

as an object within the object data stores 166, (11). Interaction (11) illustratively corresponds to implementation of the PUT request method, initially called for by the client device 102A, albeit by storing the output of the task execution rather than the provided input data. After implementing the called PUT request method, the frontend 162, at (12), returns to the client device 102A the success indicator indicated by the success return value of the task (e.g., an HTTP 200 response code). Thus, from the perspective of the client device 102A, a call to PUT an object on the storage service 160 resulted in creation of that object on the service 160. However, rather than storing the input data provided by the device 102A, the object stored on the service 160 corresponds to output data of an owner-specified task, thus enabling the owner of the object greater control over the contents of that object. In some use cases, the service 160 may additionally store the input data as an object (e.g., where the owner-specified task corresponds to code executable to provide output data usable in conjunction with the input data, such as checksum generated from the input data).

[0090] With reference to FIGS. 6A and 6B, illustrative interactions will be discussed for applying a modification to an I/O path for a request to retrieve an object on the service 160, which request is referred to in connection with these figures as a “GET” request or “GET call.” While shown in two figures, numbering of interactions is maintained across FIGS. 6A and 6B.

[0091] The interactions begin at (1), where a client device 102A submits a GET call to the storage service 160, corresponding to a request to obtain data of an object (identified within the call) stored on the service 160. As shown in FIG. 6A, the call is directed to a frontend 162 of the service 160 that, at (2), retrieves from the I/O path modification data store 164 an indication of modifications to the I/O path for the call. For example, in FIG. 6A, the I/O path used can correspond to use of a GET request method directed to a particular URI (e.g., associated with the frontend 162) to retrieve an object in a particular logical location on the service 160 (e.g., a specific bucket). In FIGS. 6A and 6B, it is assumed that an owner of that logical location has previously specified a modification to the I/O path, and specifically, has specified that a serverless function should be applied to the object before a result of that function is returned to the device 102A as the requested object.

[0092] Accordingly, at (3), the frontend 162 detects within the modifications for the I/O path inclusion of a serverless task execution. Thus, at (4), the frontend 162 submits a call to the on-demand code execution system 120 to execute the task specified within the

modifications against the object specified within the call. The on-demand code execution system 120, at (5), therefore generates an execution environment 502 in which to execute code corresponding to the task. Illustratively, the call may be directed to a frontend 130 of the system, which may distribute instructions to a worker manager 140 to select or generate a VM instance 150 in which to execute the task, which VM instance 150 illustratively represents the execution environment 502. During generation of the execution environment 502, the system 120 further provisions the environment with code 504 of the task indicated within the I/O path modification (which may be retrieved, for example, from the object data stores 166). While not shown in FIG. 6A, the environment 502 further includes other dependencies of the code, such as access to an operating system, a runtime required to execute the code, etc.

[0093] In addition, at (6), the system 120 provisions the environment with file-level access to an input file handle 506 and an output file handle 508, usable to read from and write to the input data (the object) and output data of the task execution, respectively. As discussed above, files handle 506 and 508 may point to a (physical or virtual) block storage device (e.g., disk drive) attached to the environment 502, such that the task can interact with a local file system to read input data and write output data. For example, the environment 502 may represent a virtual machine with a virtual disk drive, and the system 120 may obtain the object referenced within the call from the service 160, at (6'), and store the object on the virtual disk drive. Thereafter, on execution of the code, the system 120 may pass to the code a handle of the object as stored on the virtual disk drive, and a handle of a file on the drive to which to write output data. In another embodiment, files handle 506 and 508 may point to a network file system, such as an NFS-compatible file system, on which the object has been stored. For example, the file-level interface 166 may provide file-level access to the object as stored within the object data stores, as well as to a file representing output data. By providing handles 506 and 508, the task code 504 is enabled to read the input data and write output data using stream manipulations, as opposed to being required to implement network transmissions. Creation of the handles 506 and 508 may illustratively be achieved by execution of staging code 157 within or associated with the environment 502.

[0094] The interactions of FIG. 6A are continued in FIG. 6B, where the system 120 executes the task code 504 at (7). As the task code 504 may be user-authored, any number of functionalities may be implemented within the code 504. However, for the purposes of

description of FIGS. 6A and 6B, it will be assumed that the code 504, when executed, reads input data (corresponding to the object identified within the call) from the input file handle 506 (which may be passed as a commonly used input stream, such as stdin), manipulates the input data, and writes output data to the output file handle 508 (which may be passed as a commonly used output stream, such as stdout). Accordingly, at (8), the system 120 obtains data written to the output file (e.g., the file referenced in the output file handle) as output data of the execution. In addition, at (9), the system 120 obtains a return value of the code execution (e.g., a value passed in a final call of the function). For the purposes of description of FIGS. 6A and 6B, it will be assumed that the return value indicates success of the execution. At (10), the output data and the success return value are then passed to the frontend 162.

[0095] On receiving output data and the return value, the frontend 162 returns the output data of the task execution as the requested object. Interaction (11) thus illustratively corresponds to implementation of the GET request method, initially called for by the client device 102A, albeit by returning the output of the task execution rather than the object specified within the call. From the perspective of the client device 102A, a call to GET an object from the storage service 160 therefore results in return of data to the client device 102A as the object. However, rather than returning the object as stored on the service 160, the data provided to the client device 102A corresponds to output data of an owner-specified task, thus enabling the owner of the object greater control over the data returned to the client device 102A.

[0096] Similarly to as discussed above with respect to FIGS. 5A and 5B, while shown as a single interaction in FIG. 6B, in some embodiments output data of a task execution and a return value of that execution may be returned separately. In addition, while a success return value is assumed in FIGS. 6A and 6B, other types of return value are possible and contemplated, such as error values, pipeline-control values, or calls to execute other data manipulations. Moreover, return values may indicate what return value is to be returned to the client device 102A (e.g., as an HTTP status code). In some instances, where output data is iteratively returned from a task execution, the output data may also be iteratively provided by the frontend 162 to the client device 102A. Where output data is large (e.g., on the order of hundreds of megabytes, gigabytes, etc.), iteratively returning output data to the client device 102A can enable that data to be provided as a stream, thus speeding delivery of the



content to the device 102A relative to delaying return of the data until execution of the task completes.

[0097] While illustrative interactions are described above with reference to FIGS. 5A-6B, various modifications to these interactions are possible and contemplated herein. For example, while the interactions described above relate to manipulation of input data, in some embodiments a serverless task may be inserted into the I/O path of the service 160 to perform functions other than data manipulation. Illustratively, a serverless task may be utilized to perform validation or authorization with respect to a called request method, to verify that a client device 102A is authorized to perform the method. Task-based validation or authorization may enable functions not provided natively by the service 160. For example, consider a collection owner who wishes to limit certain client devices 102 to accessing only objects in the collection created during a certain time range (e.g., the last 30 days, any time excluding the last 30 days, etc.). While the service 160 may natively provide authorization on a per-object or per-collection basis, the service 160 may in some cases not natively provide authorization on a duration-since-creation basis. Accordingly, embodiments of the present disclosure enable the owner to insert into an I/O path to the collection (e.g., a GET path using a given URI to the collection) a serverless task that determines whether the client is authorized to retrieve a requested object based on a creation time of that object. Illustratively, the return value provided by an execution of the task may correspond to an “authorized” or “unauthorized” response. In instances where a task does not perform data manipulation, it may be unnecessary to provision an environment of the task execution with input and output stream handles. Accordingly, the service 160 and system 120 can be configured to forego provisioning the environment with such handles in these cases. Whether a task implements data manipulation may be specified, for example, on creation of the task and stored as metadata for the task (e.g., within the object data stores 166). The service 160 may thus determine from that metadata whether data manipulation within the task should be supported by provisioning of appropriate stream handles.

[0098] While some embodiments may utilize return values without use of stream handles, other embodiments may instead utilize stream handles without use of return values. For example, while the interactions described above relate to providing a return value of a task execution to the storage service 160, in some instances the system 120 may be configured to

detect completion of a function based on interaction with an output stream handle. Illustratively, staging code within an environment (e.g., providing a file system in user space or network-based file system) may detect a call to deallocate the stream handle (e.g., by calling a “file.close()” function or the like). The staging code may interpret such a call as successful completion of the function, and notify the service 160 of successful completion without requiring the task execution to explicitly provide return value.

[0099] While the interactions described above generally relate to passing of input data to a task execution, additional or alternative information may be passed to the execution. By way of non-limiting example, such information may include the content of the request from the client device 102 (e.g., the HTTP data transmitted), metadata regarding the request (e.g., a network address from which the request was received or a time of the request), metadata regarding the client device 102 (e.g., an authentication status of the device, account time, or request history), or metadata regarding the requested object or collection (e.g., size, storage location, permissions, or time created, modified, or accessed). Moreover, in addition or as an alternative to manipulation of input data, task executions may be configured to modify metadata regarding input data, which may be stored together with the input data (e.g., within the object) and thus written by way of an output stream handle, or which may be separately stored and thus modified by way of a metadata stream handle, inclusion of metadata in a return value, or separate network transmission to the service 160.

[0100] With reference to FIG. 7, an illustrative routine 700 for implementing owner-defined functions in connection with an I/O request obtained at the object storage service of FIG. 1 over an I/O path will be described. The routine 700 may illustratively be implemented subsequent to association of an I/O path (e.g., defined in terms of an object or collection, a mechanism of access to the object or collection, such as a URI, an account transmitting an IO request, etc.) with a pipeline of data manipulations. For example, the routine 700 may be implemented prior to the interactions of FIG. 3, discussed above. The routine 700 is illustratively implemented by a frontend 162.

[0101] The routine 700 begins at block 702, where the frontend 162 obtains a request to apply an I/O method to input data. The request illustratively corresponds to a client device (e.g., an end user device). The I/O method may correspond, for example, to an HTTP request method, such as GET, PUT, LIST, DELETE, etc. The input data may be included within the

request (e.g., within a PUT request), or referenced in the request (e.g., as an existing object on the object storage service 160).

[0102] At block 704, the frontend 162 determines one or more data manipulations in the I/O path for the request. As noted above, the I/O path may be defined based on a variety of criteria (or combinations thereof), such as the object or collection referenced in the request, a URI through which the request was transmitted, an account associated with the request, etc. Manipulations for each defined I/O path may illustratively be stored at the object storage service 160. Accordingly, at block 704, the frontend 162 may compare parameters of the I/O path for the request to stored data manipulations at the object storage service 160 to determine data manipulations inserted into the I/O path. In one embodiment, the manipulations form a pipeline, such as the pipeline 400 of FIG. 4, which may be previously stored or constructed by the frontend 162 at block 704 (e.g., by combining multiple manipulations that apply to the I/O path). In some instances, an additional data manipulation may be specified within the request, which data manipulation may be inserted, for example, prior to pre-specified data manipulations (e.g., not specified within the request). In other instances, the request may exclude reference to any data manipulation.

[0103] At block 706, the frontend 162 passes input data of the I/O request to an initial data manipulation for the I/O path. The initial data manipulation may include, for example, a native manipulation of the object storage service 160 or a serverless task defined by an owner of the object or collection referenced in the call. Illustratively, where the initial data manipulation is a native manipulation, the frontend 162 may pass the input to the object manipulation engine 170 of FIG. 1. Where the initial data manipulation is a serverless task, the frontend 162 can pass the input to the on-demand code execution system 120 of FIG. 1 for processing via an execution of the task. An illustrative routine for implementing a serverless task is described below with reference to FIG. 8.

[0104] While FIG. 7 illustratively describes data manipulations, in some instances other processing may be applied to an I/O path by an owner. For example, an owner may insert into an I/O path for an object or collection a serverless task that provides authentication independent of data manipulation. Accordingly, in some embodiments block 706 may be modified such that other data, such as metadata regarding a request or an object specified in the request, is passed to an authentication function or other path manipulation.

[0105] Thereafter, the routine 700 proceeds to block 708, where the implementation of the routine 700 varies according to whether additional data manipulations have been associated with the I/O path. If so, the routine 700 proceeds to block 710, where an output of a prior manipulation is passed to a next manipulation associated with the I/O path (e.g., a subsequent stage of a pipeline).

[0106] Subsequent to block 710, the routine 700 then returns to block 708, until no additional manipulations exist to be implemented. The routine 700 then proceeds to block 712, where the frontend 162 applies the called I/O method (e.g., GET, PUT, POST, LIST, DELETE, etc.) to the output of the prior manipulation. For example, the frontend 162 may provide the output as a result of a GET or LIST request, or may store the output as a new object as a result of a PUT or POST request. The frontend 162 may further provide a response to the request to a requesting device, such as an indication of success of the routine 700 (or, in cases of failure, failure of the routine). In one embodiment, the response may be determined by a return value provided by a data manipulation implemented at blocks 706 or 710 (e.g., the final manipulation implemented before error or success). For example, a manipulation that indicates an error (e.g., lack of authorization) may specify an HTTP code indicating that error, while a manipulation that proceeds successfully may instruct the frontend 162 to return an HTTP code indicating success, or may instruct the frontend 162 to return a code otherwise associated with application of the I/O method (e.g., in the absence of data manipulations). The routine 700 thereafter ends at block 714.

[0107] Notably, application of the called method to that output, as opposed to input specified in an initial request, may alter data stored in or retrieved from the object storage service 160. For example, data stored on the service 160 as an object may differ from the data submitted within a request to store such data. Similarly, data retrieved from the system as an object may not match the object as stored on the system. Accordingly, implementation of routine 700 enables an owner of data objects to assert greater control over I/O to an object or collection stored on the object storage service 160 on behalf of the owner.

[0108] In some instances, additional or alternative blocks may be included within the routine 700, or implementation of such blocks may include additional or alternative operations. For example, as discussed above, in addition to or as an alternative to providing output data, serverless task executions may provide a return value. In some instances, this return value may

instruct a frontend 162 as to further actions to take in implementing the manipulation. For example, an error return value may instruct the frontend 162 to halt implementation of manipulations, and provide a specified error value (e.g., an HTTP error code) to a requesting device. Another return value may instruct the frontend 162 to implement an additional serverless task or manipulation. Thus, the routine 700 may in some cases be modified to include, subsequent to blocks 706 and 710 for example, handling of the return value of a prior manipulation (or block 708 may be modified to include handling of such a value). Thus, the routine 700 is intended to be illustrative in nature.

[0109] With reference to FIG. 8, an illustrative routine 800 will be described for executing a task on the on-demand code execution system of FIG. 1 to enable data manipulations during implementation of an owner-defined function. The routine 800 is illustratively implemented by the on-demand code execution system 120 of FIG. 1.

[0110] The routine 800 begins at block 802, where the system 120 obtains a call to implement a stream manipulation task (e.g., a task that manipulations data provided as an input IO stream handle). The call may be obtained, for example, in conjunction with blocks 706 or 710 of the routine 700 of FIG. 7. The call may include input data for the task, as well as other metadata, such as metadata of a request that preceded the call, metadata of an object referenced within the call, or the like.

[0111] At block 804, the system 120 generates an execution environment for the task. Generation of an environment may include, for example, generation of a container or virtual machine instance in which the task may execute and provisioning of the environment with code of the task, as well as any dependencies of the code (e.g., runtimes, libraries, etc.). In one embodiment, the environment is generated with network permissions corresponding to permissions specified for the task. As discussed above, such permissions may be restrictively (as opposed to permissively) set, according to a whitelist for example. As such, absent specification of permissions by an owner of an I/O path, the environment may lack network access. Because the task operates to manipulate streams, rather than network data, this restrictive model can increase security without detrimental effect on functionality. In some embodiments, the environment may be generated at a logical network location providing access to otherwise restricted network resources. For example, the environment may be

generated within a virtual private local area network (e.g., a virtual private cloud environment) associated with a calling device.

[0112] At block 806, the system 120 stages the environment with an IO stream representing to input data. Illustratively, the system 120 may configure the environment with a file system that includes the input data, and pass to the task code a handle enabling access of the input data as a file stream. For example, the system 120 may configure the environment with a network file system, providing network-based access to the input data (e.g., as stored on the object storage system). In another example, the system 120 may configure the environment with a “local” file system (e.g., from the point of view of an operating system providing the file system), and copy the input data to the local file system. The local file system may, for example, be a filesystem in user space (FUSE). In some instances, the local file system may be implemented on a virtualized disk drive, provided by the host device of the environment or by a network-based device (e.g., as a network-accessible block storage device). In other embodiments, the system 120 may provide the IO stream by “piping” the input data to the execution environment, by writing the input data to a network socket of the environment (which may not provide access to an external network), etc. The system 120 further configures the environment with stream-level access to an output stream, such as by creating a file on the file system for the output data, enabling an execution of the task to create such a file, piping a handle of the environment (e.g., stdout) to a location on another VM instance colocated with the environment or a hypervisor of the environment, etc.

[0113] At block 808, the task is executed within the environment. Execution of the task may include executing code of the task, and passing to the execution handles or handles of the input stream and output stream. For example, the system 120 may pass to the execution a handle for the input data, as stored on the file system, as a “stdin” variable. The system may further pass to the execution a handle for the output data stream, e.g., as a “stdout” variable. In addition, the system 120 may pass other information, such as metadata of the request or an object or collection specified within the request, as parameters to the execution. The code of the task may thus execute to conduct stream manipulations on the input data according to functions of the code, and to write an output of the execution to the output stream using OS-level stream operations.

[0114] The routine 800 then proceeds to block 810, where the system 120 returns data written to the output stream as output data of the task (e.g., to the frontend 162 of the object storage system). In one embodiment, block 810 may occur subsequent to the execution of the task completing, and as such, the system 120 may return the data written as the complete output data of the task. In other instances, block 810 may occur during execution of the task. For example, the system 120 may detect new data written to the output stream and return that data immediately, without awaiting execution of the task. Illustratively, where the output stream is written to an output file, the system 120 may delete data of the output file after writing, such that sending of new data immediately obviates a need for the file system to maintain sufficient storage to store all output data of the task execution. Still further, in some embodiments, block 810 may occur on detecting a close of the output stream handle describing the output stream.

[0115] In addition, at block 812, subsequent to the execution completing, the system 120 returns a return value provided by the execution (e.g., to the frontend 162 of the object storage system). The return value may specify an outcome of the execution, such as success or failure. In some instances, the return value may specify a next action to be undertaken, such as implementation an additional data manipulation. Moreover, the return value may specify data to be provided to a calling device requesting an I/O operation on a data object, such as an HTTP code to be returned. As discussed above, the frontend 162 may obtain such return value and undertake appropriate action, such as returning an error or HTTP code to a calling device, implementing an additional data manipulation, performing an I/O operation on output data, etc. In some instances, a return value may be explicitly specified within code of the task. In other instances, such as where no return value is specified within the code, a default return value may be returned (e.g., a '1' indicating success). The routine 800 then ends at block 814.

[0116] Customers typically desire the ability to determine process data (such as determining a checksum value of a file, or perform some other function) once it has been uploaded to an object storage service in order to confirm the integrity of the uploaded data. However, current techniques often require waiting until the complete file is uploaded, even when the file is split into separate portions and the individual portions are uploaded in parallel (e.g., using a multi-part upload procedure, which is a term used to refer to any procedure where

multiple parts or sub-objects are individually uploaded and later combined into a complete, reassembled, or sometimes referred to as unified, file or object), before processing of the reassembled (or sometimes referred to as unified), complete file can be determined. Where multi-part upload is supported, embodiments enable insertion of a processing function into the input/output path of each portion, such that individual intermediate or initial (or first) functions can be executed on each portion. In addition, embodiments also enable insertion of a processing function that combines the individual intermediate or initial function outputs (e.g., the checksum values of each portion of the input file, etc.) to determine a final (or second) function output associated with the reassembled input file (e.g., such as determining a checksum value of the reassembled file, or determining some other function output based on the reassembled file). Where multi-part upload enables parallel upload, intermediate function outputs can also be calculated in parallel. Pre-calculation of an intermediate function output (such as a checksum), either in parallel or iteratively during upload of portions, enables the function output (e.g., the checksum) for a complete file to be calculated much more rapidly after uploading is complete, as compared to calculating the function output of the complete, reassembled file only after the uploading and reassembling of the input file is complete. The term “reassembled” may also be referred to as “unified.” For example, a reassembled file, object, or data may also be referred to as a unified file, object, or data.

[0117] Multi-part upload enables a client to split a file into separate portions and the upload the separate portions in parallel. Once all portions have been successfully uploaded, the client may submit a call to merge, or reassemble the separate portions to form the original file. The client may also submit a manifest with the call that indicates which portions are to be merged, and the order in which the portions are to be merged.

[0118] One particularly useful application of such processing is to determine a checksum of a large file based upon individual checksum values of file parts, each of which may be uploaded in parallel. A checksum value is an error-detecting code determined from a set of data and used to detect changes to the set of data. One such checksum value is determined using a cyclical redundancy check (e.g., CRC-32, which is a 32-bit cyclical redundancy check). A checksum algorithm enables calculation of a value, or a checksum, for an object, where the value is smaller than the object, but is would significantly change if even minor changes to the object occur. Therefore, checksums can be used to detect errors associated with the transfer



of the object from one location to another. The routine illustrated in FIG. 9 may be used to compute the checksum (or other value) of an input file from individual checksums (or other values) determined from individual portions of the input file.

[0119] FIG. 9 is a flow diagram of an illustrative routine 900 that may be executed by an object storage service, such as object storage service 160, a code execution service (or a function running within the code execution service), such as the on-demand code execution system 120, or both. The routine 900 may be used to dynamically process input data portions (sometimes referred to as chunks, parts, or data sub-objects) of input data at run time (“on-the-fly”). Such processing may occur as the input data portions are uploaded to and stored as data object portions in an object storage service (such as, for example, object storage service 160) and in response to a request to merge the data object portions into a data object stored on the object storage service. Although routine 900 is described with respect to calculating a checksum value of an input file based upon individual checksum values of certain independently-uploaded input file portions, the routine may be used to determine a function output based upon any initial, or intermediate function outputs.

[0120] In some embodiments, the routine 900 may be used to automatically determine a checksum value of (or perform a first function on) each individual input data portion as it is uploaded and prior to reassembling the individual data object portions into the data object representing the complete input data. Determining a checksum value of each individual input data portion as it is uploaded and prior to reassembling the input data can advantageously reduce the amount of time before the stored input data is ready for further processing or retrieval. For example, if errors occur during input data portion upload, the error may be detected as soon as the input data portion upload is completed, instead of after the complete input data is reassembled. Such error detection can result in the re-uploading of just the input data portion having such errors. Alternatively, first values determined from each input data portion as they uploaded may be used to detect a first condition, instead of determining the first condition after the complete input data is reassembled. Additional processing may be performed with respect to each input data portion based on its corresponding first value, as well. In addition, a checksum value of the complete input data may be determined from the checksum values of each of the individual input data portion checksum values instead of from the reassembled input data (e.g., after reassembling the portions into the data object).

Similarly, a second or final value associated with the complete input data may be determined from the first values of the individual data portions instead of from the reassembled input data by applying a second function to the first values. Determining the checksum (or second value) of the input data from the checksums (or first values) of its input data portions advantageously reduces latency and computing resource requirements. Aspects of the routine 900 will be described with additional reference to FIG. 10, which is a system diagram of illustrative data flows and interactions between various components of the service provider system 110.

[0121] The routine 900 may begin in response to an event, such as submission of a request from a client device 102 to upload input data to the object storage service 160. Illustratively, an owner of a collection of data objects to which the input data is to be added as a new data object may have previously specified that, on uploading of an object to the collection using multi-part upload, a first task should be executed to process each portion of the data object uploaded, and that a second task should be executed on a request to reassemble the portions into the data object. In some embodiments, the routine 900 or portions thereof may be implemented on multiple processors, serially or in parallel.

[0122] At block 902, the object storage service 160 can receive a request to store input data submitted via multi-part upload. FIG. 10 illustrates the object storage service 160 receiving the request at (1). The request illustratively includes parameters, such as an identifier of the input data to be stored by the object storage service 160 as a data object, a location to store the data object; context data regarding the request; other data; or some combination thereof. For example, the request may be a resource request, such as a PUT request, for particular input data to be stored in the object data store 166 of the object storage service 160, which input data is to be provided via multi-part upload.

[0123] At block 902, the object storage service 160 can also determine that function output is to be generated using portions of the input data to be stored in the object storage service 160. In some embodiments, the determination may be based on context data and/or the input data itself. For example, the object storage service 160 may receive an indication that the client will transfer the input data to the storage service 160 using a multi-part file transfer protocol, or the input data may be required to be uploaded to the object storage service 160 using a multi-part file transfer protocol. In such case, the object storage service 160 will determine an object identifier (e.g., an object ID) for the multi-part input data to be transferred.

The object storage service 160 will provide the object ID to the client. In some embodiments, the input data is not transferred using a multi-part file transfer protocol. Instead, the input data is transferred in portions (e.g., objects, sub-objects, files, delineated elements, etc.), but not necessarily according to a multi-part file transfer protocol. A manifest or list may be provided to identify the portions that are to be subsequently joined together, and the order in which they are to be joined together, to reassemble the complete input data from its portions.

[0124] At block 904, the object storage service 160 may receive a portion of the input data from the client. In one specific, non-limiting embodiment, the input data may be a file, a composite file (e.g., a compressed file, such as a file compressed according to a .zip, .tar or other compressed file format), a composable object, or a super-composable object composed of individual objects or sub-objects. Each input data portion is received with associated metadata, which can include the object ID and an indication of one or more functions to be performed on the input data portion, the complete input data, or both. For example, the metadata can include a checksum value associated with the data object portion (the “received CV”). The received input data portion, object ID, and metadata (e.g., the received CV) may be stored by the object storage service 160 in one or more staging areas. Staging areas are data storage locations, and include data storage accessible via a block storage service, a local disk, the object data store 166 of the object storage service 160, or other data storage location. The received input data portion, object ID, and metadata may be stored in the same or different staging areas. In addition, multiple input data portions may be received by the object storage system 160 in parallel, during at least partially overlapping time periods. Furthermore, the input data portions may be received in a different order than the order in which the input data portions are to be assembled into the complete data object. Therefore, the metadata can include an input data portion identifier (input data portion ID) that can be used to designate the input data portions to be used, and the order in which the input data portions are to be arranged, to assemble the complete input data. Furthermore, the input data portions may be the same size or have different sizes than one another.

[0125] The indications of one or more functions to be performed on the input data portion, the complete input data, or both, can include an indication to manipulate and/or validate the input data portion, the input data, or both, prior to storing the input data within the object storage service 160 object data store 166. For example, the indication can indicate that

the input data portion, the complete input data, or both, are to be compressed, decompressed, encrypted, decrypted, or a combination thereof, prior to being stored within the object storage service 160 object data store 166. In addition, the indication can indicate that the input data portion, the complete input data, or both, are to be error checked prior to subsequent manipulation. For example, the input data portions may be individually error checked, or checksum checked prior to being reassembled to the complete input data. Additionally, the reassembled input data may be checksum checked prior to being stored in the object data store storage 166. In some embodiments, the object storage service 160 may automatically error check each input data portion and/or the complete input data without receiving an indication instructing the object storage service 160 to do so. The object storage service 160 may initiate error detection of each input data portion as soon as it is completely received, without waiting to reassemble the complete input data. FIG. 10 illustrates the object storage service 160 receiving and storing portion(s) of the input data at (2).

[0126] At block 906, the object storage service 160 can make a call to the execution environment 502 to execute a function (e.g., a first function) to determine a checksum value of (or perform a different calculation or determination using) the input data portion. FIG. 10 illustrates the object storage service 160 making a call to the execution environment 502, and the execution environment 502 (or the function running within the execution environment 502) returning a result at (3). In response to the call, the VM instance 150 or other execution environment 502 can execute the function by using the input data portion. For example, the VM instance 150 or other execution environment 502 (or the function running within the execution environment 502) may determine a checksum value (a “determined CV”) associated with the input data portion. The execution environment 502 (or the function running within the execution environment 502) may perform any of a variety of error detection operations on the input data portion, including a cyclical redundancy check (e.g., CRC-32) or any other parallelizable error detection operation. A parallelizable error detection operation is an error detection operation that may be performed on portions of input data, and the individual outputs of the error detection operation may be combined or otherwise used to determine a checksum or other data integrity indication associated with the complete input data. Each determined CV may be stored in any of a variety of ways, including storing it as metadata with the input data

portion, with a relational or non-relational database service, using a relational or non-relational database management system, or storing it with the object storage service.

[0127] At block 908, the object storage service 160 can process the output of the function received from the execution environment 502 (or the function running within the execution environment 502). For example, the object storage service can perform error detection for (or perform some other calculation or determination using) the input data portion using the output data, such as the determined CV, received from the execution environment 502 (or the function running within the execution environment 502). Error detection may include comparing the determined CV to the stored, received CV. If the two values are different, the object storage service 160 may determine that an error has occurred during upload of the input data portion, and the client may be requested to re-send the associated input data portion. FIG. 10 illustrates the object storage service 160 processing the output(s) of the function(s) at (4).

[0128] In some embodiments, the object storage service 160 may provide the determined CV (or first values) associated with the input data portion to the client. The client may receive the determined CV and compare it to a client-determined checksum value of the input data portion (or otherwise process the first values). If the two values are different, the client may determine that it needs to re-send the associated input data portion to the object storage service 160. In such case, the client will instruct the object storage service 160 that the input data portion is being re-uploaded.

[0129] In some embodiments, instead of performing a checksum determination function on each input data portion that is received by the object storage service 160, the execution environment 502 (or the function running within the execution environment 502) is configured to perform the function on a fixed-sized portion of the input data (or input data portion). The size of the fixed-sized portion may be configured by the client. For example, the size may be designated using a parameter send to the object storage service 160 in connection with the initiation of the input data multi-part upload process. In some embodiments, the size is predetermined by the object storage service 160 or execution environment 502 (or the function running within the execution environment 502).

[0130] For example, the client may wish to upload a 10 GB file as input data using a multi-part upload process. The client may upload the input data in multiple portions, each

having the same or different size. For example, the client may upload the input data in ten 1 GB data object portions. The execution environment 502 may process each portion as it is received (as discussed above), or it may process a fixed-sized portion of each portion, instead. For example, the execution environment 502 may process each 100 MB (or other predetermined, fixed size) of each 1 GB data object as it is received.

[0131] Such fixed-sized portion processing can advantageously enable the execution environment 502 to operate on a known fixed sized input. Such configuration would greatly simplify and improve the efficiency of the provisioning the staging area storage used to process each fixed-sized portion of the input data portion. In some embodiments, fixed-sized portion processing is used automatically if the complete input data size, or if an input data portion size exceeds a threshold value.

[0132] Blocks 904 to 908 define a parallelizable block 909 that may be iterated multiple times in parallel or sequentially, or both. For example, the blocks of block 909 may be performed for each input data portion received from the client, and in parallel (e.g., during at least partially overlapping time periods).

[0133] At block 910, the object storage service 160 (or VM instance 150, other execution environment 502, or the function running within the execution environment 502) can receive a request to perform a second function based on at least a portion of the first outputs. For example, the object storage service 160 can receive a request to determine a checksum of the reassembled input data from the stored input data portions, submitted via multi-part upload, or a request to reassemble the input data from the stored input data portions. FIG. 10 illustrates the object storage service 160 receiving the request at (5). The request illustratively includes parameters, such as an identifier of the input data portions to be reassembled and stored by the object storage service 160 as a data object, a location to store the data object; context data regarding the request; other data; or some combination thereof. For example, the request may be a resource request, such as a PUT request.

[0134] At block 910, the object storage service 160 can also determine that function output is to be generated using portions of the input data stored in the object storage service 160. In some embodiments, the determination may be based on context data and/or the input data itself. For example, the object storage service 160 may receive an indication that previously received input data portions are to be combined together. A manifest or list may be

provided to identify the portions that are to be joined together, and the order in which they are to be joined together, to reassemble the complete input data from previously uploaded portions.

[0135] At block 912, the object storage service 160 may execute a call to the execution environment 502 (or the function running within the execution environment 502) to determine a checksum of the reassembled input data by (or perform a second function) using the individual checksums (or first values) of each of the input data portion checksums. FIG. 10 illustrates object storage service 160 executing the call to the execution environment 502 (or the function running within the execution environment 502) to determine a checksum of the reassembled input data at (6). In one embodiment, the object storage service 160 receives a manifest from the client that identifies the input data portions that are to be reassembled into the complete input data. In addition, the manifest also identifies the order in which the input data portions are to be reassembled into the complete input data. For example, individual input data portions may have been received out of order and the manifest may be used to determine the correct ordering of the input data portions within the complete input data. The checksum values of each of the input data portions identified in the manifest are provided to the execution environment 502 with the call to determine the checksum of the reassembled input data. The execution environment 502 may execute the function to determine the checksum of the reassembled input data by combining the individual checksums, or by determining a checksum of the individual checksum values. The execution environment 502 (or the function running within the execution environment 502) may return the checksum of the reassembled input data to the object storage service, as shown in FIG. 10 at (6).

[0136] At decision block 914, the object storage service 160 processes the output of the function. For example, the object storage service 160 may perform error detection using the checksum of the reassembled input data, or it may provide the output to the client to enable the client to perform error detection. In some embodiments, the object storage service 160 may process the output of the function by storing the output as an object within the object data store 166. If an error is detected, the client may re-upload one or more portions of the input data. If no error is detected, the object storage service 160 reassembles the complete input data from the stored input data portions based upon the contents of the manifest. FIG. 10 illustrates the object storage service processing the output of the function and reassembling the complete input data from the stored input data portions at (7).

[0137] At block 916, the object storage service 160 can store the reassembled input data as a data object in the object data store 166. FIG. 10 illustrates the object storage service storing the reassembled input data as a data object at (8).

[0138] Blocks 912 through 916 are illustrated as occurring in sequence. However, the order in which these blocks occur may vary. In some embodiments, the ordering may be different, or two or more block may be performed at the same time, or during at least partially overlapping time periods. For example, in some embodiments, block 912 can be performed concurrently (or partially concurrently) with block 914 and/or block 916. In some embodiments, blocks 914 and 916 may occur before block 912, as well.

[0139] The routine may terminate at block 918.

[0140] In some embodiments, a client sends a request to an object storage service (such as object storage service 160) to write input data or a file as a data object to a storage location, such as an object data store (including object data store 166). For example, the client may wish to store a collection of customer records that include personal customer information (e.g., customer government-issued identification numbers, social security numbers, etc.). The client may wish to obfuscate the customer records prior to storage so that users may only retrieve versions of the customer records where the personal customer information has been obfuscated. The client may wish to allow only a small number of users with superior security credentials to have access to the un-obfuscated personal customer information. In another example, the input data may include medical images (e.g., photograph, x-ray, sonogram, ultrasound images, etc.), where a portion of the image includes personally identifiable information, such as the patients' names. The client may wish to obfuscate the personally-identifiable information from the medical image. The client request may include the input data, or information usable by the object storage service 160 to obtain the input data. In response to the request, the object storage service 160 may stage the input data in a staging area, such as any of the staging areas discussed above. Once the input data has been staged, a routine to obfuscate the input data may be initiated, such as routine 1100 of FIG. 11.

[0141] FIG. 11 is a flow diagram of an illustrative routine 1100 that may be executed by a code execution service (or a function running within the execution environment 502), such as the on-demand code execution system 120, to dynamically tokenize, mask, scramble, obscure, encrypt, or otherwise render unintelligible (collectively referred to herein as



“obfuscate” for convenience) portions of input data at run time in response to a request to store or write the input data. Obfuscation also includes replacement (e.g., selective replacement) of one or more portions of input data with different, unique data, such as a token. The token for each instance of replaced data (e.g., each instance of private information) is different from every other token. In other words, there may be provided a one-to-one mapping of tokens to each instance of private information. The routine may be implemented as a function of the on-demand code execution system 120, and a user may attach the function to, or insert the function within, an input-output path for a given collection of objects. A request to write input data includes request to write or store the input data as a data object in a storage location, such as an object data store, including object data store 166. A client may wish to store a data set that includes both private and non-private information. However, the client may wish to store the data set in a manner in which the private information is separated from the non-private information, and in which the private information, and a mapping between tokens and private information, are stored in a secure location with access to only a limited number of authorized individuals or resources. The client may also wish to provide access to the non-private information to a larger group of individual or resource, or store the non-private information in a less secure location. Aspects of the routine 1100 will be described with reference to FIG. 12, which is a system diagram of illustrative data flows and interactions between various components of the service provider system 110.

[0142] The routine 1100 may begin in response to an event, such as when the routine illustrated in FIG. 8 reaches block 808. For example, the routine 1100 may be an owner-defined function, also referred to as a user-defined task, that is performed by a VM instance 150 or other execution environment 502 generated during the routine illustrated in FIG. 8. In some embodiments, the routine 1100 or portions thereof may be implemented on multiple processors, serially or in parallel.

[0143] At block 1102, the VM instance 150 or other execution environment 502 (or the function running within the execution environment 502) can receive parameters associated with a request to write input data. FIG. 12 illustrates the execution environment 502 (or the function running within the execution environment 502) receiving the parameters associated with the request at (1). In some embodiments, the parameters may include: reference data comprising a reference to input data to be stored as a data object; a reference to an output

location of the data object; context data regarding the request; other data or metadata; or some combination thereof. For example, the request may be a resource request, such as a PUT request, to store input data as a particular data object in the object storage service 160. The reference to the input data may be data that can be used by the execution environment 502 (or the function running within the execution environment 502) to access the input data, such as: a file descriptor; a file handle; a pointer; or some other data representing an address or identifier of the input data. The reference to the output location may be data that can be used by the execution environment 502 (or the function running within the execution environment 502) to write, store, or otherwise persist output data, such as: a file descriptor; a file handle; a pointer; or some other data representing an address or identifier of a location for providing output of the function. The context data or metadata may include data or metadata regarding the context of the request, such as: an identifier of a user, account or other source of the request; an identifier of an access or security profile under which the request is being make; data representing the access or security rights under which the request is to be processed; an identifier of a location associated with the request; an identifier of a language associated with the request; or data representing preferences or tendencies of a source of the request. While FIG. 12 depicts the object storage service providing parameters such as a reference to a requested data object or reference to an output location to the execution environment 502 (or the function running within the execution environment 502), in other instances these references may be provided by elements of the execution system 120, such as staging code 157.

[0144] At block 1104, the VM instance 150 or other execution environment 502 (or the function running within the execution environment 502) can obtain the input data to be stored as a data object using the reference data. The input data may be obtained in un-obfuscated or substantially un-obfuscated form. FIG. 12 illustrates the execution environment 502 (or the function running within the execution environment 502) obtaining the input data at (2). In some embodiments, the input data may not be obtained from the object storage service 160 at block 1104, but may be provided to the execution environment 502 (or the function running within the execution environment 502) previously. For example, during staging of the execution environment, the input data may be obtained and stored on a computing device of the execution environment 502 at a location indicated by the reference data.

[0145] At block 1106, the VM instance 150 or other execution environment 502 (or the function running within the execution environment 502) can determine that one or more portions of the input data are to be obfuscated. In some embodiments, the determination may be based on context data and/or the input data. For example, if a portion of the input data looks like, or is determined to be or to likely be a form of private or personally-identifiable information, the execution environment 502 (or the function running within the execution environment 502) can determine that such portion is to be obfuscated. The execution environment 502 (or the function running within the execution environment 502) may test one or more items of context data against one or more criteria to determine whether to perform an obfuscation and which portion(s) of the input data to obfuscate. If an item of context data satisfies one or more criteria, then the execution environment 502 (or the function running within the execution environment 502) can determine that one or more portions of the input data are to be obfuscated such that the obfuscated portion(s) render the portions of the input data as unable to be understood by a recipient. FIG. 12 illustrates the execution environment 502 (or the function running within the execution environment 502) determining to obfuscate portions of the input data at (3).

[0146] Testing the context data against the criteria may include: determining that the input data includes private, or personally identifiable information (including, but not limited to: an individual's name, address, age, government-issued identification number, social security number, date of birth, place of birth, mother's maiden name, biometric information, health information, a vehicle identification number (VIN), etc.); or determining that the input data includes information that has been designated confidential.

[0147] In one specific, non-limiting embodiment, the input data may be a data file, such as a spreadsheet, delimited file, or other collection of data records. Some portions of the data file, such as collections of records, collections of columns or data fields, or the like are to be stored in obfuscated form if the request satisfies one or more criteria. The execution environment 502 (or the function running within the execution environment 502) may determine that properties of the request indicated by the context data or otherwise associated with the request satisfy the criteria for particular records, columns, and/or fields of the requested data object. The execution environment 502 (or the function running within the execution environment 502) may determine, based on this criteria that the particular records,

columns, and/or fields of the requested input data are to be obfuscated prior to being output by the function (e.g., for storage as a data object).

[0148] At block 1108, the VM instance 150 or other execution environment 502 can selectively apply obfuscation to portions of the input data determined above. FIG. 12 illustrates the execution environment 502 (or the function running within the execution environment 502) obfuscating portions of the input data at (4). Obfuscating the content of a portion of the input data may involve the use of one or more obfuscation methods, such as scrambling the content in a pseudo random method, generating a hash of the content, replacing the content with a token mapped to the content in a data store (such as the object storage service 160), encrypting the portion, or the like. In some embodiments, encryption is performed using a key under the control of the data object owner and is managed using a key management service. In some embodiments, different obfuscation methods may be used for different portions of a data object, different data objects, different context data criteria, or the like.

[0149] For example, in one embodiment, the obfuscation method may include replacing a portion of the input data with a token that is mapped to a key-value pair secured in a secure location, such as an external database. For example, a social security number “909-09-0909” may be replaced with a globally unique identifier, such as “001,” and a distinct database may store a key-value pair mapping key “001” to “909-09-0909.”

[0150] In some embodiments, the obfuscation method may be specified by an entity that owns or is responsible for the data object requested to be stored (e.g., as part of the request to store the input data as the data object). For example, an entity may specify that particular type of obfuscation (e.g., an industry standard obfuscation method in the medical field) is to be used for a data object or bucket of data objects, while another entity may specify that a different type of obfuscation (e.g., tokenization using a mapping of tokens to data) is to be used for a different data object or bucket of data objects. If no obfuscation method is specified, the execution environment 502 (or the function running within the execution environment 502) may apply a default obfuscation method.

[0151] At block 1110, the VM instance 150 or other execution environment 502 (or the function running within the execution environment 502) can provide the selectively-obfuscated input data as output of the function. For example, the execution environment 502 (or the function running within the execution environment 502) can place the selectively-

obfuscated input data at the output location indicated by the reference data, and finalize the output. Finalizing output of the function may include closing the output stream or file identified by the reference to the output location and/or providing a return value (e.g., indicating success, failure, or some other characteristics of function execution) to the object storage service 160. In addition, at block 1110, the VM instance 150 or other execution environment 502, or the function running within the execution environment 502, can also provide an index as second output data. The index may include a mapping between tokens and obfuscated private information. The index may be subsequently stored using the object storage service, a different object storage service, or a different storage service, such as a database storage service, or any other storage service. FIG. 12 illustrates the execution environment 502 (or the function running within the execution environment 502) providing the selectively-obfuscated input data as output at (5). The routine 1100 may terminate at block 1112.

[0152] Obfuscation of data object portions at write provides certain data management advantages. For example, if input data includes customer records, such as purchase history, personally identifiable information, and other private and non-private information, a data object including obfuscated versions of that information may be more easily updated if a particular customer deletes her account. For example, instead of having to scan through an entire data object to locate and remove all of the deleted customer's private information, the system can instead delete the mapping of tokens associated with the deleted customer from the token mapping table (or mapping of tokens to key-value pairs, as discussed above) or delete the customer's private information from the location in which such private information is stored.

[0153] FIG. 13 is a flow diagram of an illustrative routine 1300 that may be executed by a code execution service, such as the on-demand code execution system 120, to dynamically determine and store, at run time, an index of the contents of input data in response to a request to store the input data as a data object. A client may wish to retrieve only a portion of a composite file stored as a data object in an object storage service. By providing an index that identifies the different files or data sets or items within the composite file and their locations, the object storage service is able to retrieve and provide to the client only the desired portions. Aspects of the routine 1300 will be described with reference to FIG. 14, which is a system

diagram of illustrative data flows and interactions between various components of the service provider system 110.

[0154] In some embodiments, a client sends a request to an object storage service (such as object storage service 160) to write input data as a data object at a storage location, such as an object data store (including object data store 166). For example, the client may wish to store input data that includes a composite file, such as a compressed file, sometimes referred to as a .zip archive, a .tar archive, or a compressed file, or other file made up of a collection of individual data elements. The composite file may include one or more individual files, each of which is compressed. The composite file may also include an index of the contents of the composite file. The index may include the names of each of the individual files within the composite file, as well as other metadata regarding the composite file's contents. The index may also provide a mapping between the contents of the composite file and the byte-range location of each of the contents. This index, therefore, enables a user to use a "byte-range GET" to request only the bytes for a certain desired file, or other content of the composite file. In other examples, the composite file does not include an index of the composite file's contents. In yet other examples, the input data is not a composite file, but the object storage service 160 is configured to generate a storable data object that corresponds to a compressed version of the input data, and to store the compressed version within the object storage service. The client request may include the input data, or information usable by the object storage service 160 to obtain the input data. In response to the request, the object storage service 160 may stage the input data in a staging area, such as any of the staging areas discussed above. Once the object has been staged, a routine to index the input data may be initiated, such as routine 1100 of FIG. 13.

[0155] The routine 1300 may begin in response to an event, such as when the routine illustrated in FIG. 8 reaches block 808. For example, the routine 1300 may be an owner-defined function, also referred to as a user-defined task, that is performed by a VM instance 150 or other execution environment 502 generated during the routine illustrated in FIG. 8. In some embodiments, the routine 1300 or portions thereof may be implemented on multiple processors, serially or in parallel.

[0156] At block 1302, the VM instance 150 or other execution environment 502 (or a function running within the execution environment 502) can receive parameters associated

with a request to store input data as a data object. FIG. 14 illustrates the execution environment 502 (or the function running within the execution environment 502) receiving the parameters associated with the request at (1). In some embodiments, the parameters may include: reference data comprising a reference input data to be stored as a data object; a reference to an output location of the data object; context data regarding the request; other data or metadata; or some combination thereof. For example, the request may be a resource request, such as a PUT request, to store input data as a particular data object in the object storage service 160. The reference to the input data may be data that can be used by the execution environment 502 (or the function running within the execution environment 502) to access the input data, such as: a file descriptor; a file handle; a pointer; or some other data representing an address or identifier of the input data. The reference to the output location may be data that can be used by the execution environment 502 (or the function running within the execution environment 502) to write, store, or otherwise persist output data, such as: a file descriptor; a file handle; a pointer; or some other data representing an address or identifier of a location for providing output of the function. The context data or metadata may include data or metadata regarding the context of the request, such as: an identifier of a user, account or other source of the request; an identifier of an access or security profile under which the request is being make; data representing the access or security rights under which the request is to be processed; an identifier of a location associated with the request; an identifier of a language associated with the request; or data representing preferences or tendencies of a source of the request. While FIG. 14 depicts the object storage service providing parameters such as a reference to a requested data object or reference to an output location to the execution environment 502 (or the function running within the execution environment 502), in other instances these references may be provided by elements of the execution system 120, such as staging code 157.

[0157] At block 1304, the VM instance 150 or other execution environment 502 (or the function running within the execution environment 502) can obtain the input data using the reference data. FIG. 14 illustrates the execution environment 502 obtaining the input data at (2). In some embodiments, the input data may not be obtained from the object storage service 160 at block 1104, but may be provided to the execution environment 502 (or the function running within the execution environment 502) previously. For example, during staging of the

execution environment, the input data may be obtained and stored on a computing device of the execution environment 502 at a location indicated by the reference data.

[0158] At block 1306, the VM instance 150 or other execution environment 502 (or the function running within the execution environment 502) generates an index of the input data's contents. In some embodiments, the index is generated by obtaining the names of the individual files stored within the input data. For example, the input data may include an index of the data object's contents. If not, the execution environment 502 (or the function running within the execution environment 502) can read and store the names of each file within the input data. In some embodiments, the files within the input data are extracted or decompressed so the file names and/or file contents may be determined. In some embodiments, the execution environment 502 (or the function running within the execution environment 502) generates an index of the input data's contents using metadata or headers stored within the input data. In some embodiments, the VM instance 150 or other execution environment 502 (or the function running within the execution environment 502) is configured to un-pack, or recursively un-pack the input data to determine its contents (e.g., identifiers of delineated elements within the input data, and the byte-range locations of the delineated elements within the input data, the delineated elements being files, or any other delineated element described herein). Recursive unpacking can include analyzing a second composite file that is located within a first file. The VM instance 150 or other execution environment 502 (or the function running within the execution environment 502) can unpack the first file to identify the second file (or second files), and then unpack the second file to determine identifiers of delineated elements and byte-range (or other) locations within the second file. In some embodiments, the execution environment 502 (or the function running within the execution environment 502) generates an index of the input data's contents by analyzing the text within the input data. The index includes content identifiers (e.g., file names, text fields, header information, metadata, etc.) as well as location information associated with each identifier. For example, the index can include a list of all files within the input data, as well as the location (e.g., byte range, etc.) of each file within the input data. In another example, the index can include a list of all the headers of the data sets within the input data (e.g., the sales data for various geographic regions), as well as the location of each data set within the input data (e.g., byte range, etc.). In addition, when the input file comprises a composite file, the VM instance 150 or other



execution environment 502 (or the function running within the execution environment 502) can determine a file aggregation technique used to form the composite file. For example, the VM instance 150 or other execution environment 502 (or the function running within the execution environment 502) may determine whether the composite file is a .zip, .tar, or other format by analyzing bytes within the file. For example, some aggregations techniques generate files having known header formats. Therefore, the VM instance 150 or other execution environment 502 (or the function running within the execution environment 502) can dynamically evaluate the input data based upon the bytes (sometimes referred to as file aggregation technique information), and use that information to determine how to further read and interpret the rest of the input data. For example, the file aggregation technique information may be used to determine whether to perform recursive unpacking of a file, such as discussed above. FIG. 14 illustrates the execution environment 502 (or the function running within the execution environment 502) determining an index of the input data contents at (3).

[0159] At block 1308, the VM instance 150 or other execution environment 502 (or the function running within the execution environment 502) can provide the index as output of the function. For example, the execution environment 502 (or the function running within the execution environment 502) can may return the index to the object storage service 160. In some embodiments, the VM instance 150 or other execution environment 502 (or the function running within the execution environment 502) can return the input data instead of, in addition to (as second output data), or combined with the index. FIG. 14 illustrates the execution environment 502 (or the function running within the execution environment 502) providing the index as output at (4).

[0160] At block 1310, the object storage service 160 can process the function output. For example, the object storage service 160 (or a different service) may store the index. FIG. 14 illustrates the object storage service 160 storing the index at (5). The index may be stored at any of a variety of locations. For example, the object storage service 160 may store the index in an object data store, such as the object data store 166 of the object storage service 160. In another example, the object storage service 160 may store the index as a table using a relational or non-relational data storage service or database management system. In yet another example, the index may be appended to or otherwise added to the input data, and the updated input data (with index) may be stored by the object storage service 160.

[0161] In some embodiments, at block 1310, the object storage service 160 may create a data object corresponding to the input data and to add metadata to the data object that includes a reference to the index. The reference can include an indication that there is an index associated with the data object. In another embodiment, the data object corresponding to the input data and the index may be associated with each other via a naming convention. For example, the data object and the index may have similar identifier or name portions, such as a prefix, suffix, or other identifier. The reference may be used by a subsequent user of the data object to obtain a desired portion of the data object. For example, a user may retrieve the index and select a desired portion of the data object. The object storage service 160 and execution environment 502 (or the function running within the execution environment 502) may use the desired portion indicated by the user and the index to identify the location within the data object (e.g., byte range) of the desired portion of the data object. The object storage service 160 and execution environment 502 (or the function running within the execution environment 502) may use the location to retrieve, e.g., extract or decompress, the desired portion of the data object (e.g., via executing a byte-range query or GET, etc. on the stored data object) and provide it to the user.

[0162] The routine may terminate at block 1312.

[0163] FIG. 14 illustrates the execution environment 502 (or the function running within the execution environment 502) indexing a data object for storage in the object storage service 160 in response to receiving a request to store the data object. Although the just-in-time transformation is shown as occurring in connection with operations of routine 1300 for indexing the data object, just-in-time transformations may be performed in connection with any of the other routines described herein, with any other owner-defined function or user-defined task, in a pipeline with multiple functions, etc.

Example Embodiments

[0164] Examples of the embodiments of the present disclosure can be described in view of the following clauses:

Clause 1. A system comprising:

an object storage service comprising one or more computing devices, wherein the object storage service is configured to store a plurality of data objects within an object data store; and

a code execution service comprising one or more computing devices for on-demand execution of functions in an input/output (I/O) path of the object storage service;

wherein the object storage service is configured to at least:

receive input data and a request to store the input data as a data object within the object data store;

determine that a function to obfuscate a portion of the input data, associated with the request to store the input data, is to be executed prior to storing the input data as the data object; and

transmit a call to the code execution service to execute the function against the input data; and

wherein the code execution service is configured to at least:

receive, from the object storage service, the call to execute the function, the call including the input data; and

execute the function, wherein executing the function causes the service to:

identify, within the input data, one or more instances of private information that is to be obfuscated prior to storing the input data as the data object within the object data store;

generate output data comprising the one or more instances of the private information of the input data in obfuscated form and a remaining portion of the input data in un-obfuscated form; and

return the output data to the object storage service;

wherein the object storage service is further configured to store the output data in the object data store as the data object.

Clause 2. The system of Clause 1, wherein the input data does not remain stored in the object storage service in un-obfuscated form after the object storage service stores the output data.

Clause 3. The system of Clause 1, wherein the code execution service is configured to generate the output data by:

- determining a unique token for each instance of the private information, wherein each unique token is different from every other unique token;

- storing the instances of the private information and a mapping of the unique tokens to the instances of the private information; and

- replacing each instance of the private information with the corresponding unique token.

Clause 4. The system of Clause 1, wherein the code execution service is configured to generate the obfuscated form of the private information by encrypting the private information.

Clause 5. A computer-implemented method comprising:

- under control of a computing system comprising one or more computer processors configured to execute specific instructions,

- receiving a request to store input data as a data object within a data store;

- determining, based at least partly on the request, to execute a function to obfuscate a portion of the input data prior to storing the input data as the data object;

- configuring a code execution system to execute the function, wherein the code execution system is configured to provide on-demand execution of functions in an input/output (I/O) path of the data store;

- executing the function using the code execution system prior to storing the input data as the data object, wherein executing the function comprises:

- obtaining the input data;

- determining to obfuscate a first portion of the input data;

- generating a first obfuscated portion comprising the first portion in obfuscated form; and

generating a output data comprising the first obfuscated portion, wherein the output data does not include the first portion in un-obfuscated form; and

storing the output data as the data object in the data store.

Clause 6. The computer-implemented method of Clause 5, wherein determining to obfuscate a first portion of the input data comprises determining that the first portion includes private information.

Clause 7. The computer-implemented method of Clause 6, wherein the private information represents one or more of: personally identifiable information, a name, an address, an age, a government-issued identification number, a date of birth, a place of birth, a mother's maiden name, an account number, or a biometric record.

Clause 8. The computer-implemented method of Clause 5, wherein generating the first obfuscated portion comprises:

determining a unique token corresponding to the first portion;

storing a mapping of the token to the first portion; and

replacing the first portion with the unique token.

Clause 9. The computer-implemented method of Clause 8, wherein storing the mapping of the token to the first portion comprises storing the mapping of the token to the first portion in a storage location having different access permissions than a location in the data store in which the data object is stored.

Clause 10. The computer-implemented method of Clause 5, wherein generating the first obfuscated portion comprises encrypting the first portion of the input data using an encryption key.

Clause 11. The computer-implemented method of Clause 10, further comprising storing the encryption key and a mapping of the encryption key to the first obfuscated portion.

Clause 12. The computer-implemented method of Clause 5, wherein determining to obfuscate the first portion of the input data is based at least in part upon a portion of the input data.

Clause 13. A system comprising:

a data store storing plurality of data objects; and

one or more computing devices in communication with the data store and configured to at least:

- receive a request to store input data as a data object within a data store;
- determine, based at least partly on the request, to execute a function to obfuscate a portion of the input data prior to storing the input data as the data object;

- configure a code execution system to execute the function, wherein the code execution system is configured to provide on-demand execution of functions in an input/output (I/O) path of the data store;

- execute the function using the code execution system prior to storing the input data as the data object, wherein executing the function comprises:

- obtaining the input data;

- determining to obfuscate a first portion of the input data;

- generating a first obfuscated portion comprising the first portion in obfuscated form; and

- generating an output data comprising the first obfuscated portion, wherein the output data does not include the first portion in un-obfuscated form; and

- store the output data as the data object in the data store.

Clause 14. The system of Clause 13, wherein the one or more computing devices are further configured to determine to obfuscate the first portion of the input data by determining that the first portion includes private information.

Clause 15. The system of Clause 14, wherein the private information represents one or more of: personally identifiable information, a name, an address, an age, a government-issued identification number, a date of birth, a place of birth, a mother's maiden name, an account number, or a biometric record.

Clause 16. The system of Clause 13, wherein the one or more computing devices are further configured to generate the first obfuscated portion by:

- determining a unique token corresponding to the first portion;

- storing a mapping of the token to the first portion; and

- replacing the first portion with the unique token.

Clause 17. The system of Clause 16, wherein the one or more computing devices are further configured to generate the first obfuscated portion by storing the mapping of the token to the first portion in a storage location having different access permissions than a location in the data store in which the data object is stored.

Clause 18. The system of Clause 13, wherein the one or more computing devices are further configured to generate the first obfuscated portion by encrypting the first portion of the input data using an encryption key.

Clause 19. The system of Clause 18, wherein the one or more computing devices are further configured to store the encryption key and a mapping of the encryption key to the first obfuscated portion.

Clause 20. The system of Clause 13, wherein the one or more computing devices are further configured to determine to obfuscate the first portion of the input data based at least in part upon a portion of the input data.

[0165] Further examples of the embodiments of the present disclosure can be described in view of the following clauses:

Clause 1. A system comprising:

an object storage service comprising one or more computing devices, wherein the object storage service is configured to store a plurality of data objects within an object data store; and

a code execution service comprising one or more computing devices for on-demand execution of functions in an input/output (I/O) path of the object storage service;

wherein the object storage service is configured to at least:

receive from a client: (1) input data as a plurality of input data portions, and (2) a request to store the input data portions as data object portions within the object data store;

determine that a first function to generate a checksum value for an input data portion, associated with the request to store the input data portion, is to be executed for each input data portion; and

transmit a first call to the code execution service to execute the first function against the input data portion for each of the input data portions; and

wherein the code execution service is configured to at least:

- receive, from the object storage service, the first call to execute the first function on the input data portion; and

- execute the first function, wherein executing the first function causes the code execution service to:

- generate an individual checksum value for the input data portion; and

- return the individual checksum value as first output data; and

wherein the object storage service is further configured to:

- store the first output data for each of the input data portions as individual checksum data objects;

- reassemble the input data using at least some of the input data portions;

- determine that a second function to generate a checksum value for the reassembled input data is to be executed; and

- transmit a second call to the code execution service to execute the second function against the individual checksum data objects; and

wherein the code execution service is configured to at least:

- receive, from the object storage service, the second call to execute the second function; and

- execute the second function, wherein executing the second function causes the code execution service to:

- generate a checksum value for the reassembled input data based on the individual checksum data objects; and

- return the checksum value for the reassembled input data as second output data; and

wherein the object storage service is further configured to:

- store the second output data as an input data checksum data object;

- perform error detection using the second output data to determine whether the input data has been received without error;

- reassemble the input data using the input data portions; and



store the reassembled input data as a data object within the object data store.

Clause 2. The system of Clause 1, wherein executing the first function causes the code execution service to generate the individual checksum value for the input data portion by performing a cyclical redundancy check using the input data portion, and wherein executing the second function causes the code execution service to generate the checksum value for the reassembled input data by performing a cyclical redundancy check using the individual checksum data objects.

Clause 3. The system of Clause 1, wherein the object storage service is further configured to store the second output data as metadata to the data object.

Clause 4. The system of Clause 1, wherein the object storage service is further configured to determine that one or more input data portions have been received with errors, based on the individual checksum data objects and provide information to the client regarding whether one or more input data portions have been received with errors.

Clause 5. A computer-implemented method comprising:

under control of a computing system comprising one or more computer processors configured to execute specific instructions,

receiving input data from a client as a plurality of input data portions via individual upload processes, and a request to store the input data portions as data object portions within a data store;

determining, based at least partly on the request, to execute a first function for each input data portion prior to confirming storage of the input data portion as the data object portion;

executing the first function using a code execution system for each of the input data portions, wherein the code execution system provides on-demand execution of functions designated in an input/output (I/O) path of the data store, wherein executing the first function comprises:

obtaining the input data portion;

performing the first function to generate a first function value for the input data portion; and

returning the first function value;

storing the first function value;  
receiving a request to assemble at least some of the input data portions into reassembled input data;  
determining, based at least partly on the request, to execute a second function for the reassembled input data;  
executing the second function using the code execution system, wherein executing the second function comprises:  
obtaining the individual first function values for the at least some of the input data portions;  
performing the second function to generate a second function value for the reassembled input data using the individual first function values; and  
returning the second function value for the reassembled input data as second output data; and  
storing the second output data.

Clause 6. The computer-implemented method of Clause 5, further comprising determining that a particular input data portion was received with errors by using the first function value, and providing information to the client regarding whether one or more input data portions have been received with errors.

Clause 7. The computer-implemented method of Clause 5, further comprising providing the first function value to the client to enable the client to perform error detection on the input data portions.

Clause 8. The computer-implemented method of Clause 5, wherein performing the first function comprises generating an individual checksum value by performing a cyclical redundancy check using the input data portion, and wherein performing the second function comprises generating a checksum value for the reassembled input data by performing a cyclical redundancy check using at least some of the individual checksum values.

Clause 9. The computer-implemented method of Clause 5, further comprising reassembling the input data using the input data portions and storing the reassembled input data in the data store as a data object.

Clause 10. The computer-implemented method of Clause 9, wherein storing the second output data comprises storing the second output data as metadata to the data object.

Clause 11. The computer-implemented method of Clause 9, wherein determining, based at least partly on the request, to execute the second function for the reassembled input data portion comprises determining, based at least partly on the request, to execute the second function for the reassembled data portion prior to reassembling the data input.

Clause 12. The computer-implemented method of Clause 5, wherein each input data portion comprises a plurality of input data parts, each input data part having a fixed size, and wherein executing the first function for each of the input data portions comprises executing the first function for each input data part of each input data portion.

Clause 13. A system comprising:

- a data store configured to store a plurality of data objects; and

- one or more computing devices in communication with the data store and configured to at least:

- receive input data from a client as a plurality of input data portions via individual upload processes, and a request to store the input data portions as data object portions within a data store;

- determine, based at least partly on the request, to execute a first function for each input data portion prior to storing the input data portion as the data object portion;

- execute the first function using a code execution system for each of the input data portions, wherein the code execution system provides on-demand execution of functions designated in an input/output (I/O) path of the data store, wherein executing the first function comprises:

- obtaining the input data portion;

- performing the first function to generate a first function value for the input data portion; and

- returning the first function value;

- storing the first function value;

- receive a request to assemble at least some of the input data portions into reassembled input data;

determine, based at least partly on the request, to execute a second function for the reassembled input data;

execute the second function using the code execution system, wherein executing the second function comprises:

obtaining the individual first function values for the at least some of the input data portions;

performing the second function to generate a second function value for the reassembled input data using the individual first function values; and

returning the second function value for the reassembled input data as second output data; and

store the second output data.

Clause 14. The system of Clause 13, wherein the code execution system is further configured to determine that a particular input data portion was received with errors by using the first function value, and provide information to the client regarding whether one or more input data portions have been received with errors.

Clause 15. The system of Clause 13, wherein the code execution service is configured to provide the first function value to the client to enable the client to perform error detection on the input data portions.

Clause 16. The system of Clause 13, wherein the code execution service is further configured to perform the first function to generate an individual checksum value by performing a cyclical redundancy check using the input data portion, and perform the second function to generate a checksum value for the reassembled input data by performing a cyclical redundancy check using at least some of the individual checksum values.

Clause 17. The system of Clause 13, wherein the one or more computing devices is further configured to reassemble the input data using the input data portions and store the reassembled input data in the data store as a data object.

Clause 18. The system of Clause 17, wherein the one or more computing devices is further configured to store the second output data as metadata to the data object.

Clause 19. The system of Clause 17, wherein the code execution service is further configured to determine, based at least partly on the request, to execute the second function for

the reassembled input data portion by determining, based at least partly on the request, to execute the second function for the reassembled data portion prior to reassembling the data input.

Clause 20. The system of Clause 13, wherein each input data portion comprises a plurality of input data parts, each input data part having a fixed size, and wherein the code execution service is configured to execute the first function for each of the input data portions by executing the first function for each input data part of each input data portion.

[0166] Further examples of the embodiments of the present disclosure can be described in view of the following clauses:

Clause 1. A system comprising:

an object storage service comprising one or more computing devices, wherein the object storage service is configured to store a plurality of data objects; and

a code execution service comprising one or more computing devices for on-demand execution of functions in an input/output (I/O) path of the object storage service;

wherein the object storage service is configured to at least:

receive input data and a request to store the input data as a data object within the object data store, the input data comprising a composite file, wherein the composite file comprises a plurality of individual files and, for each of the individual files, a file identifier and byte range location information identifying a byte range location of the individual file within the composite file;

determine that a function to create an index by extracting the file identifiers and byte range location information from the input data, associated with the request to store the input data, is to be executed prior to storing the input data as the data object; and

transmit a call to the code execution service to execute the function against the input data; and

wherein the code execution service is configured to at least:

receive, from the object storage service, the call to execute the function, the call including the input data; and

execute the function, wherein executing the function causes the code execution service to:

generate an index by extracting the file identifiers and byte range location information from the input data, the index mapping the file identifiers to corresponding byte range location information; and  
return the index as output data; and

wherein the object storage service is further configured to store the input data in the object data store as a first data object and wherein the output data may be stored in an index data store associated with the first data object.

Clause 2. The system of Clause 1, wherein the code execution service is further configured to determine file aggregation technique information associated with the composite file, and extract the file identifiers and byte range location information from the input data using the file aggregation technique information.

Clause 3. The system of Clause 1, wherein the first data object is associated with the output data via a naming convention, or by including metadata with the first data object that references the second data object.

Clause 4. The system of Clause 1, wherein the index enables the client to retrieve a desired portion of the composite file from the data store without having to retrieve the entire data object from the object data store.

Clause 5. A computer-implemented method comprising:

under control of a computing system comprising one or more computer processors configured to execute specific instructions,

receiving a request to store input data as a data object within in a data store, the input data comprising a collection of delineated elements;

determining, based at least partly on the request, to execute a function to generate an index mapping element identifiers and element locations for each delineated element prior to storing the input data as the data object;

configuring a code execution system to execute the function, wherein the code execution system provides on-demand execution of functions in an input/output (I/O) path of the data store;

executing the function using the code execution system prior to storing the input data as the data object, wherein executing the function comprises:

obtaining the input data;

generating an index mapping the element identifiers and element locations within the input data; and

returning the index as output data; and

storing the output data separately from and associated with the data object.

Clause 6. The computer-implemented method of Clause 5, further comprising decompressing the input data prior to generating the index.

Clause 7. The computer-implemented method of Clause 5, wherein generating the index comprises extracting the element identifiers and element locations from the input data.

Clause 8. The computer-implemented method of Clause 5, further comprising generating the element identifiers, the element locations, or both, using the delineated elements.

Clause 9. The computer-implemented method of Clause 5, wherein storing the output data separately from the data object comprises storing the output data as a second data object that may be independently accessed from within the data store.

Clause 10. The computer-implemented method of Clause 5, wherein storing the output data comprises storing the output data within a database using a data storage service.

Clause 11. The computer-implemented method of Clause 5, wherein the delineated elements include one or more of rows, lines, files, comma separated values, or columns of data.

Clause 12. The computer-implemented method of Clause 5, further comprising compressing the input data and storing the compressed input data as the data object.

Clause 13. A system comprising:

a data store configured to store a plurality of data objects; and

one or more computing devices in communication with the data store and configured to at least:

receive a request to store input data as a data object within the data store, the input data comprising a collection of delineated elements;

determine, based at least partly on the request, to execute a function to generate an index mapping element identifiers and element locations for each delineated element prior to storing the input data as the data object;

configure a code execution service to execute the function, wherein the code execution service is thereby configured to:

obtain the input data;

generate an index mapping the element identifiers and element locations within the input data; and

return the index as output data; and

store the output data separately from and associated with the data object.

Clause 14. The system of Clause 13, wherein the code execution system is further configured to decompress the input data prior to generating the index.

Clause 15. The system of Clause 13, wherein the code execution service is configured to generate the index by extracting the element identifiers and element locations from the input data.

Clause 16. The system of Clause 13, wherein the code execution service is further configured to generate the element identifiers, the element locations, or both, using the delineated elements.

Clause 17. The system of Clause 13, wherein the one or more computing devices is further configured to store the output data separately from the data object by storing the output data as a second data object that may be independently accessed from within the data store.

Clause 18. The system of Clause 13, wherein the one or more computing devices is further configured to store the output data by storing the output data within a database using a data storage service.

Clause 19. The system of Clause 13, wherein the delineated elements include one or more of rows, lines, files, comma separated values, or columns of data.

Clause 20. The system of Clause 13, wherein the one or more computing devices is further configured to compress the input data and store the compressed input data as the data object.



[0167] All of the methods and processes described above may be embodied in, and fully automated via, software code modules executed by one or more computers or processors. The code modules may be stored in any type of non-transitory computer-readable medium or other computer storage device. Some or all of the methods may alternatively be embodied in specialized computer hardware.

[0168] Conditional language such as, among others, "can," "could," "might" or "may," unless specifically stated otherwise, are otherwise understood within the context as used in general to present that certain embodiments include, while other embodiments do not include, certain features, elements or steps. Thus, such conditional language is not generally intended to imply that features, elements or steps are in any way required for one or more embodiments or that one or more embodiments necessarily include logic for deciding, with or without user input or prompting, whether these features, elements or steps are included or are to be performed in any particular embodiment.

[0169] Disjunctive language such as the phrase "at least one of X, Y or Z," unless specifically stated otherwise, is otherwise understood with the context as used in general to present that an item, term, etc., may be either X, Y or Z, or any combination thereof (e.g., X, Y or Z). Thus, such disjunctive language is not generally intended to, and should not, imply that certain embodiments require at least one of X, at least one of Y or at least one of Z to each be present.

[0170] Unless otherwise explicitly stated, articles such as 'a' or 'an' should generally be interpreted to include one or more described items. Accordingly, phrases such as "a device configured to" are intended to include one or more recited devices. Such one or more recited devices can also be collectively configured to carry out the stated recitations. For example, "a processor configured to carry out recitations A, B and C" can include a first processor configured to carry out recitation A working in conjunction with a second processor configured to carry out recitations B and C.

[0171] The term "or" should generally be understood to be inclusive, rather than exclusive. Accordingly, a set containing "a, b, or c" should be construed to encompass a set including a combination of a, b, and c.

[0172] Any routine descriptions, elements or blocks in the flow diagrams described herein or depicted in the attached figures should be understood as potentially representing

modules, segments, or portions of code which include one or more executable instructions for implementing specific logical functions or elements in the routine. Alternate implementations are included within the scope of the embodiments described herein in which elements or functions may be deleted, or executed out of order from that shown or discussed, including substantially synchronously or in reverse order, depending on the functionality involved as would be understood by those skilled in the art.

[0173] It should be emphasized that many variations and modifications may be made to the above-described embodiments, the elements of which are to be understood as being among other acceptable examples. All such modifications and variations are intended to be included herein within the scope of this disclosure and protected by the following claims.

## CLAIMS

WHAT IS CLAIMED IS:

1. A computer-implemented method comprising:
  - under control of a computing system comprising one or more computer processors configured to execute specific instructions,
    - receiving a request to store input data as a data object within a data store;
    - determining, based at least partly on the request, to execute a function to obfuscate a portion of the input data prior to storing the input data as the data object;
    - configuring a code execution system to execute the function, wherein the code execution system is configured to provide on-demand execution of functions in an input/output (I/O) path of the data store;
    - executing the function using the code execution system prior to storing the input data as the data object, wherein executing the function comprises:
      - obtaining the input data;
      - determining to obfuscate a first portion of the input data;
      - generating a first obfuscated portion comprising the first portion in obfuscated form; and
      - generating a output data comprising the first obfuscated portion, wherein the output data does not include the first portion in un-obfuscated form; and
      - storing the output data as the data object in the data store.
2. The computer-implemented method of claim 1, wherein determining to obfuscate a first portion of the input data comprises determining that the first portion includes private information.
3. The computer-implemented method of claim 2, wherein the private information represents one or more of: personally identifiable information, a name, an address, an age, a government-issued identification number, a date of birth, a place of birth, a mother's maiden name, an account number, or a biometric record.

4. The computer-implemented method of claim 1, wherein generating the first obfuscated portion comprises:

determining a unique token corresponding to the first portion;  
storing a mapping of the token to the first portion; and  
replacing the first portion with the unique token.

5. The computer-implemented method of claim 4, wherein storing the mapping of the token to the first portion comprises storing the mapping of the token to the first portion in a storage location having different access permissions than a location in the data store in which the data object is stored.

6. The computer-implemented method of claim 1, wherein generating the first obfuscated portion comprises encrypting the first portion of the input data using an encryption key.

7. The computer-implemented method of claim 6, further comprising storing the encryption key and a mapping of the encryption key to the first obfuscated portion.

8. The computer-implemented method of claim 1, wherein determining to obfuscate the first portion of the input data is based at least in part upon a portion of the input data.

9. A system comprising:

a data store storing plurality of data objects; and

one or more computing devices in communication with the data store and configured to at least:

receive a request to store input data as a data object within a data store;

determine, based at least partly on the request, to execute a function to obfuscate a portion of the input data prior to storing the input data as the data object;

configure a code execution system to execute the function, wherein the code execution system is configured to provide on-demand execution of functions in an input/output (I/O) path of the data store;

execute the function using the code execution system prior to storing the input data as the data object, wherein executing the function comprises:

obtaining the input data;

determining to obfuscate a first portion of the input data;  
generating a first obfuscated portion comprising the first portion in obfuscated form; and  
generating an output data comprising the first obfuscated portion, wherein the output data does not include the first portion in un-obfuscated form; and  
store the output data as the data object in the data store.

10. The system of claim 9, wherein the one or more computing devices are further configured to determine to obfuscate the first portion of the input data by determining that the first portion includes private information.

11. The system of claim 10, wherein the private information represents one or more of: personally identifiable information, a name, an address, an age, a government-issued identification number, a date of birth, a place of birth, a mother's maiden name, an account number, or a biometric record.

12. The system of claim 9, wherein the one or more computing devices are further configured to generate the first obfuscated portion by:

determining a unique token corresponding to the first portion;  
storing a mapping of the token to the first portion; and  
replacing the first portion with the unique token.

13. The system of claim 12, wherein the one or more computing devices are further configured to generate the first obfuscated portion by storing the mapping of the token to the first portion in a storage location having different access permissions than a location in the data store in which the data object is stored.

14. The system of claim 9, wherein the one or more computing devices are further configured to generate the first obfuscated portion by encrypting the first portion of the input data using an encryption key.

15. The system of claim 9, wherein the one or more computing devices are further configured to determine to obfuscate the first portion of the input data based at least in part upon a portion of the input data.

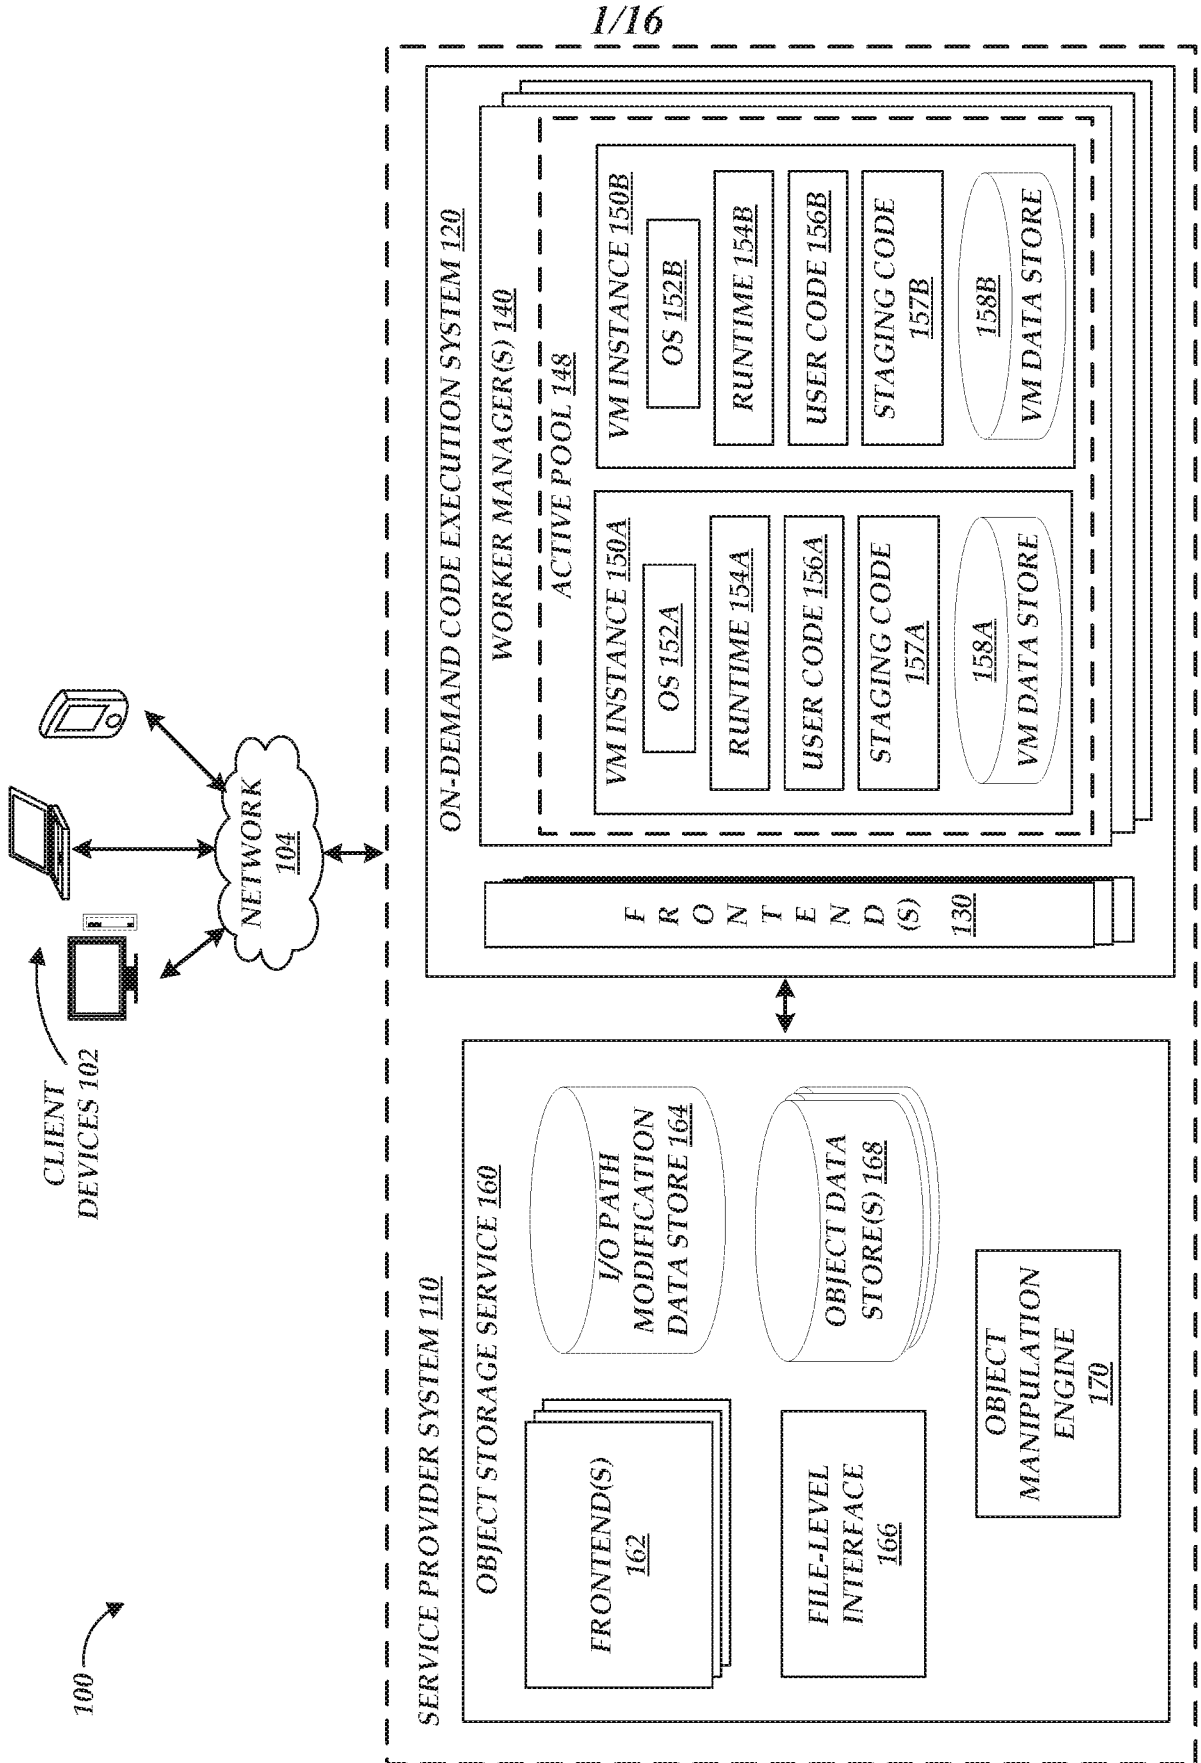


Fig. 1

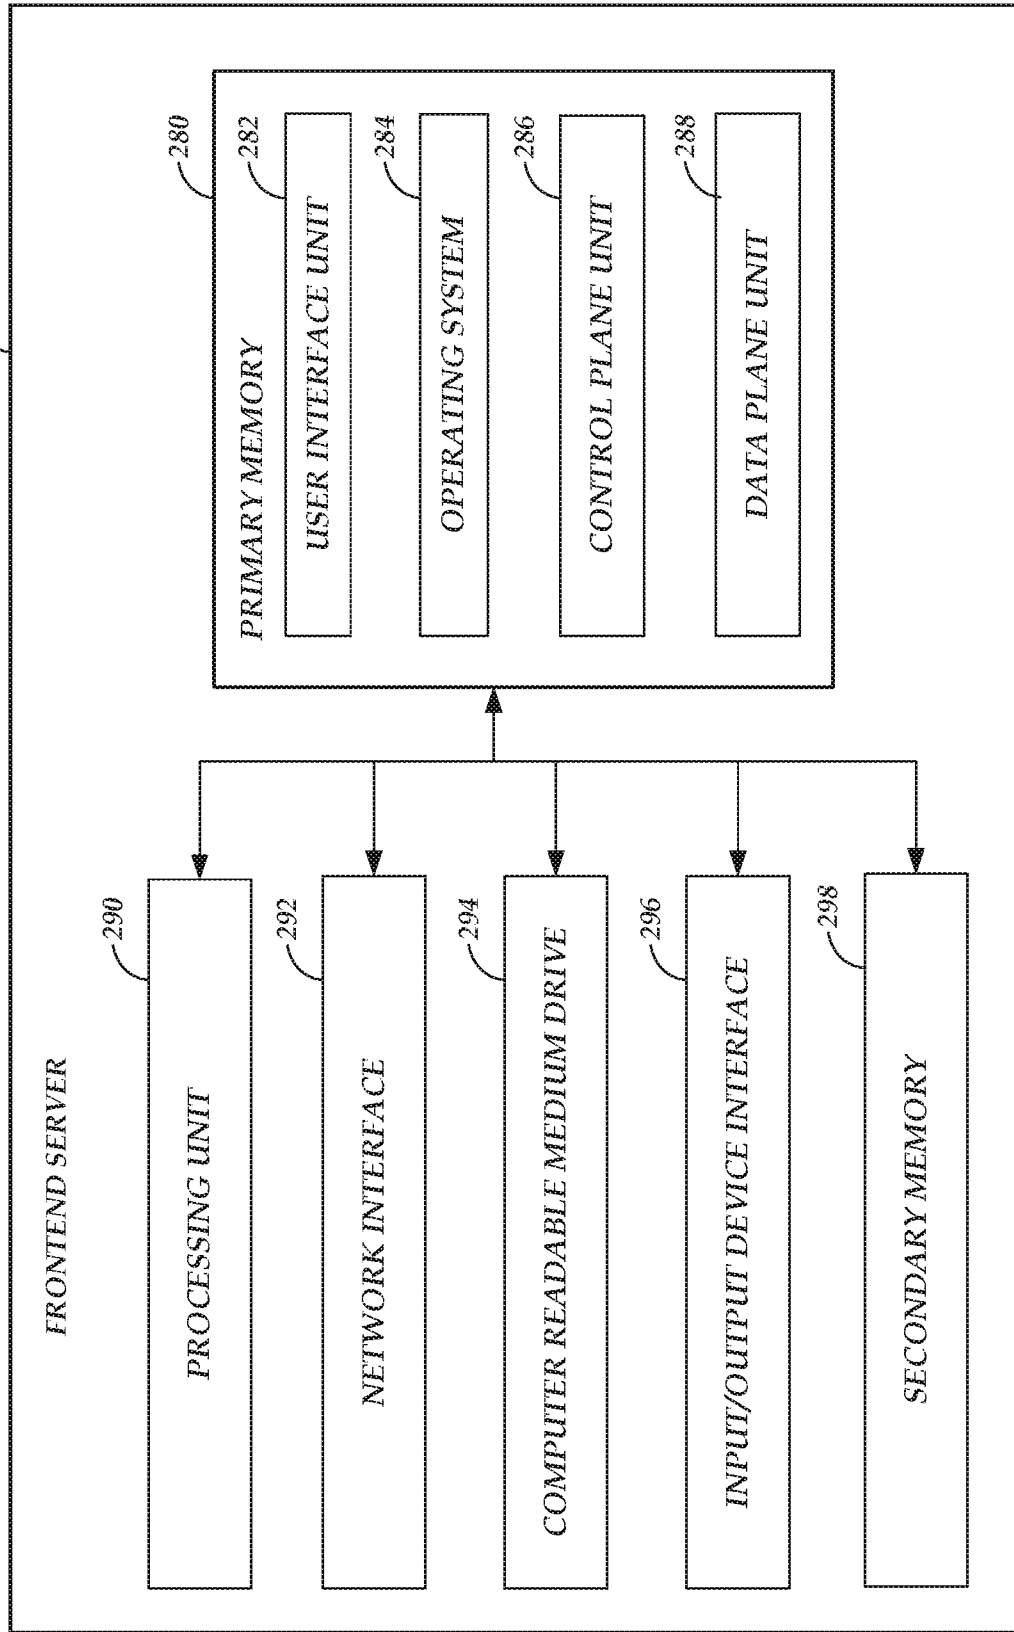


Fig. 2

3/16

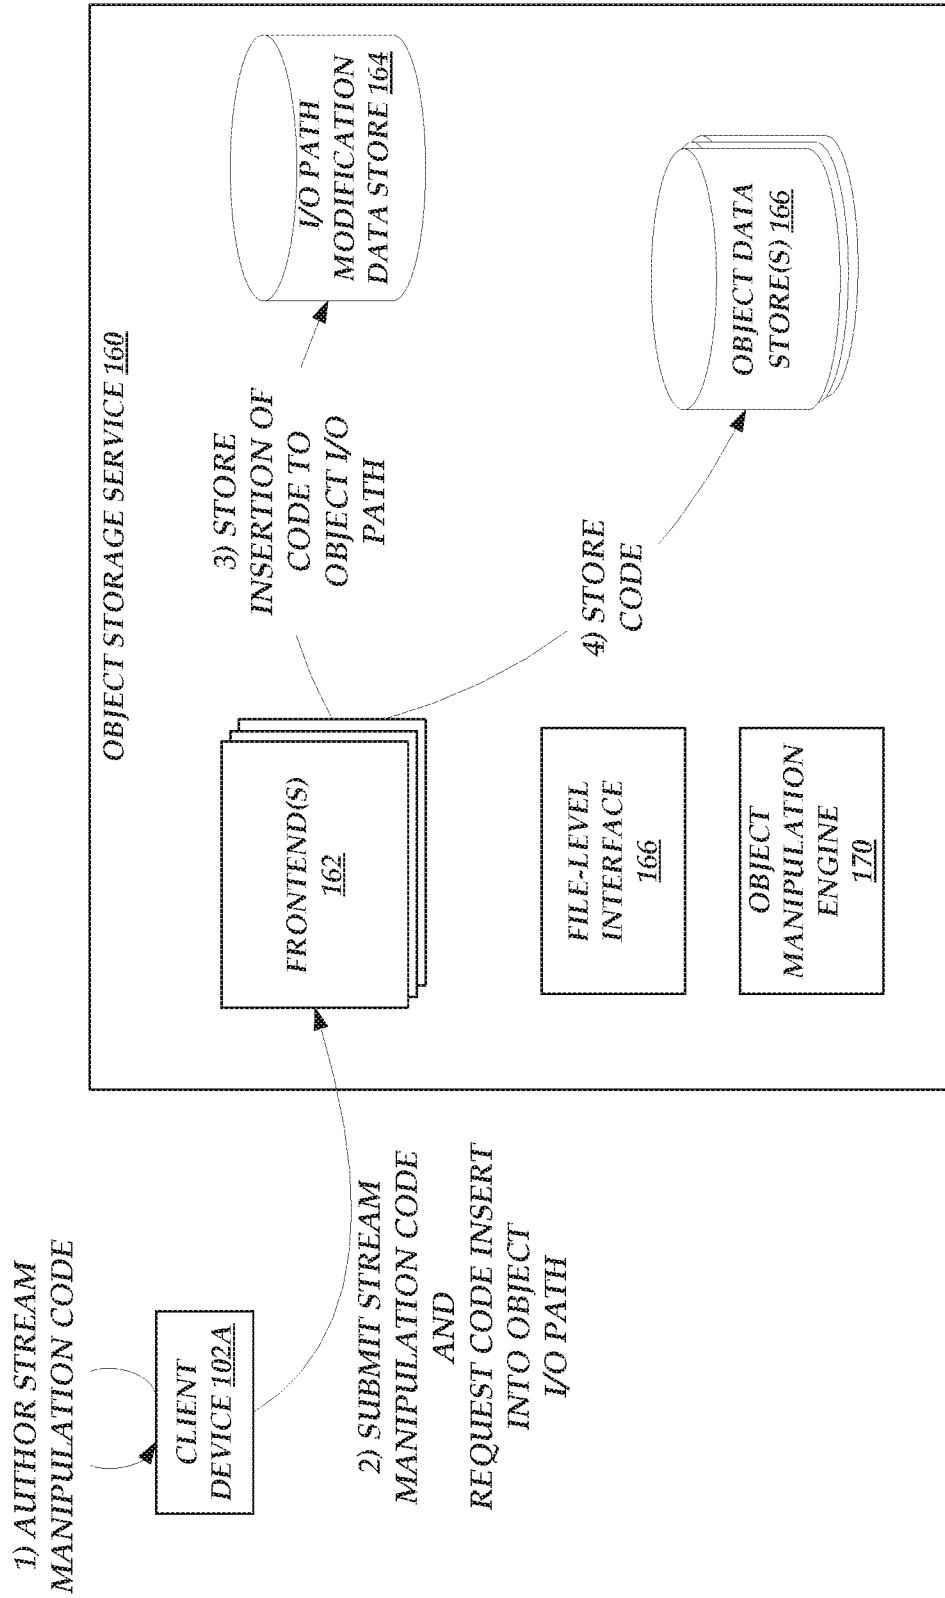


Fig. 3



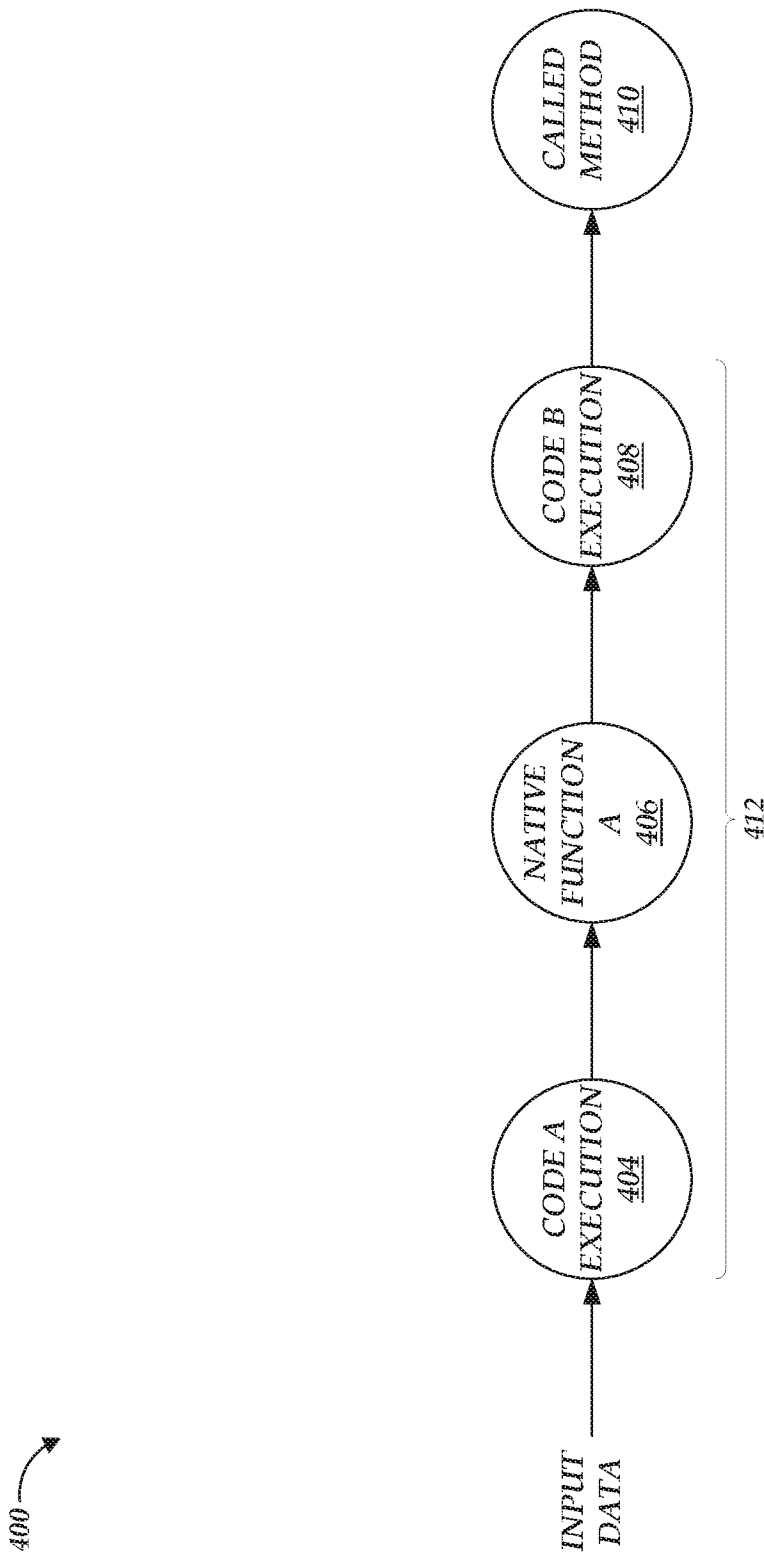


Fig. 4

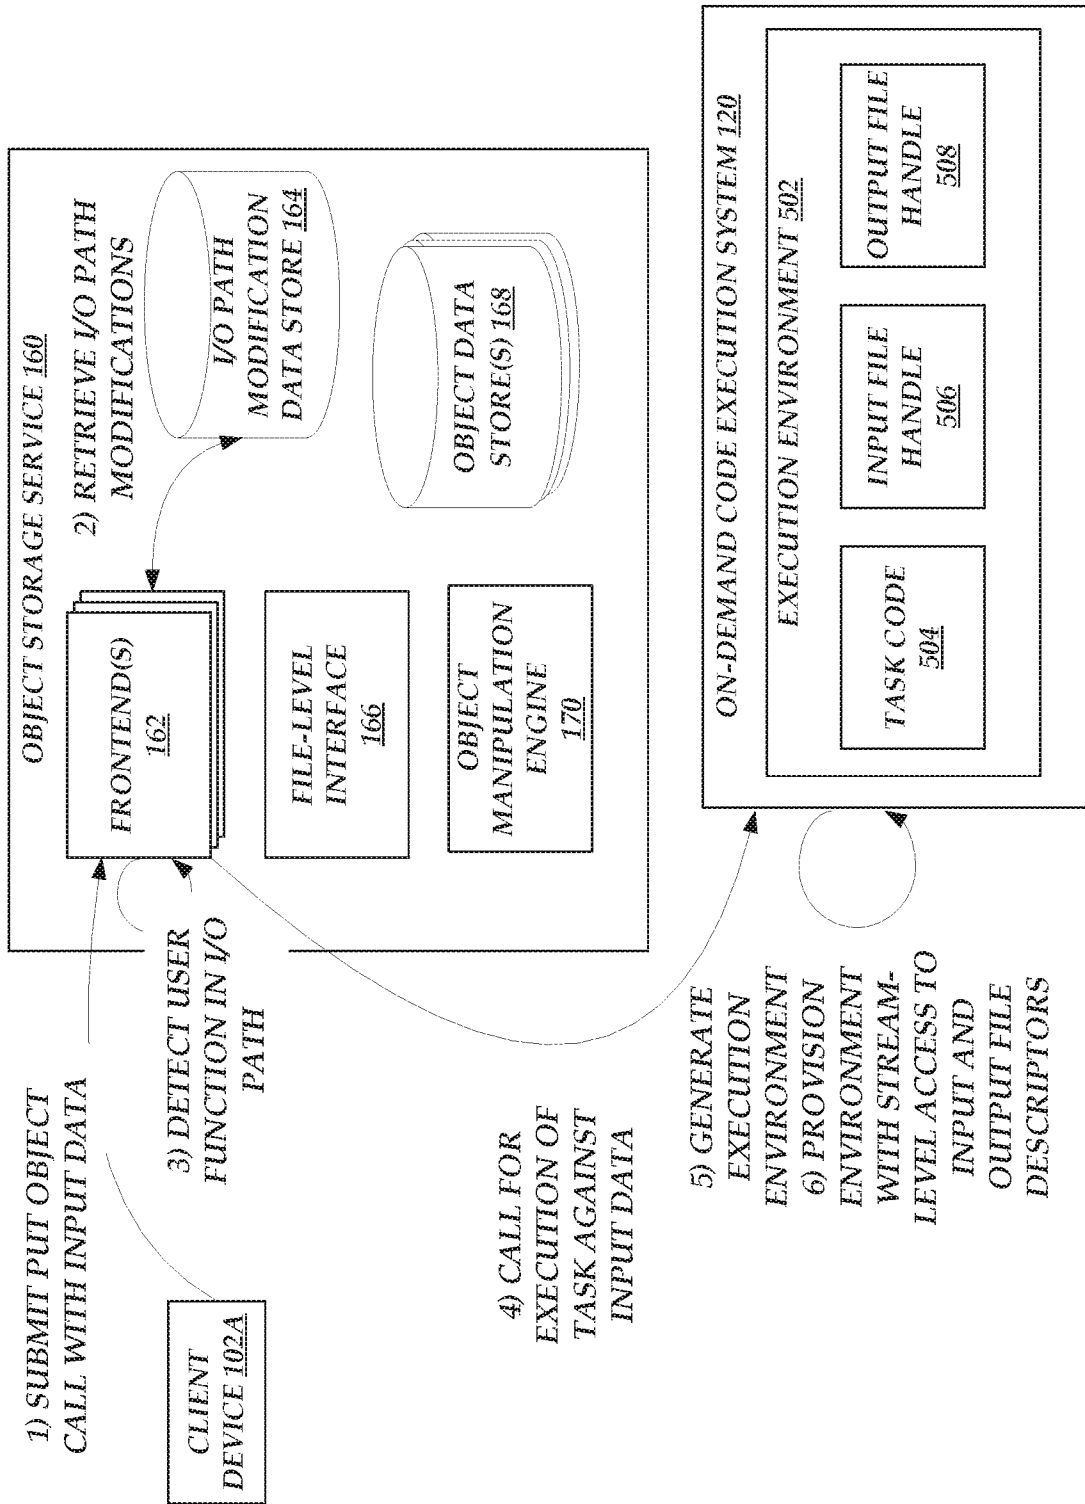


Fig. 5A

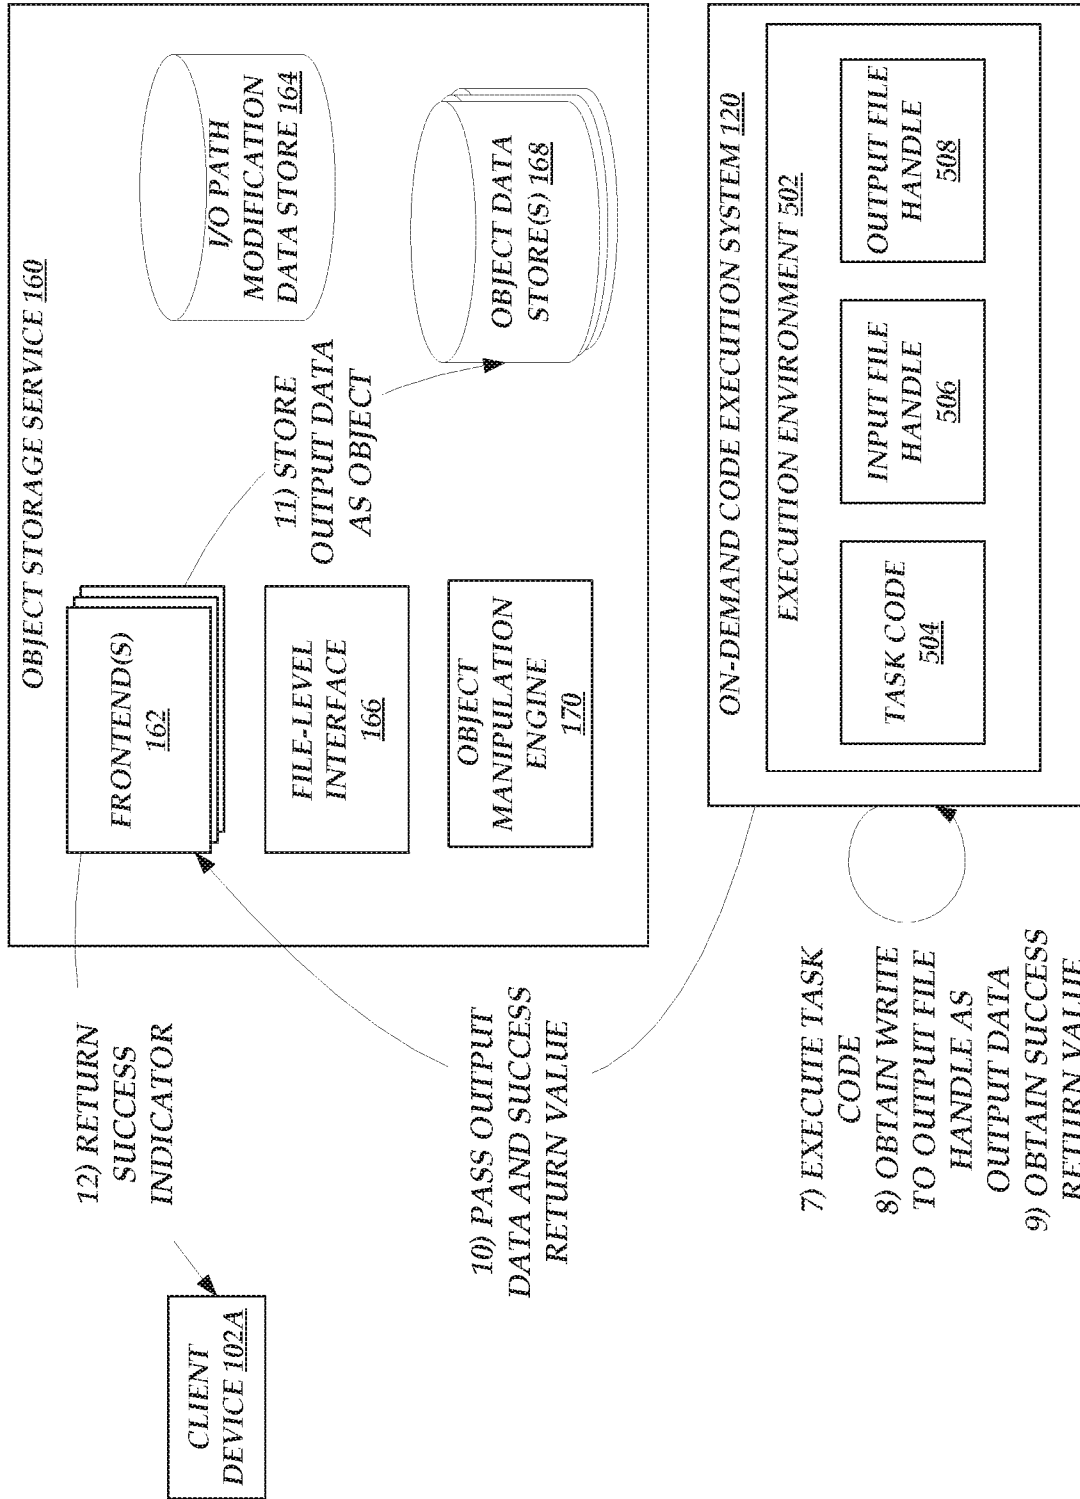


Fig. 5B

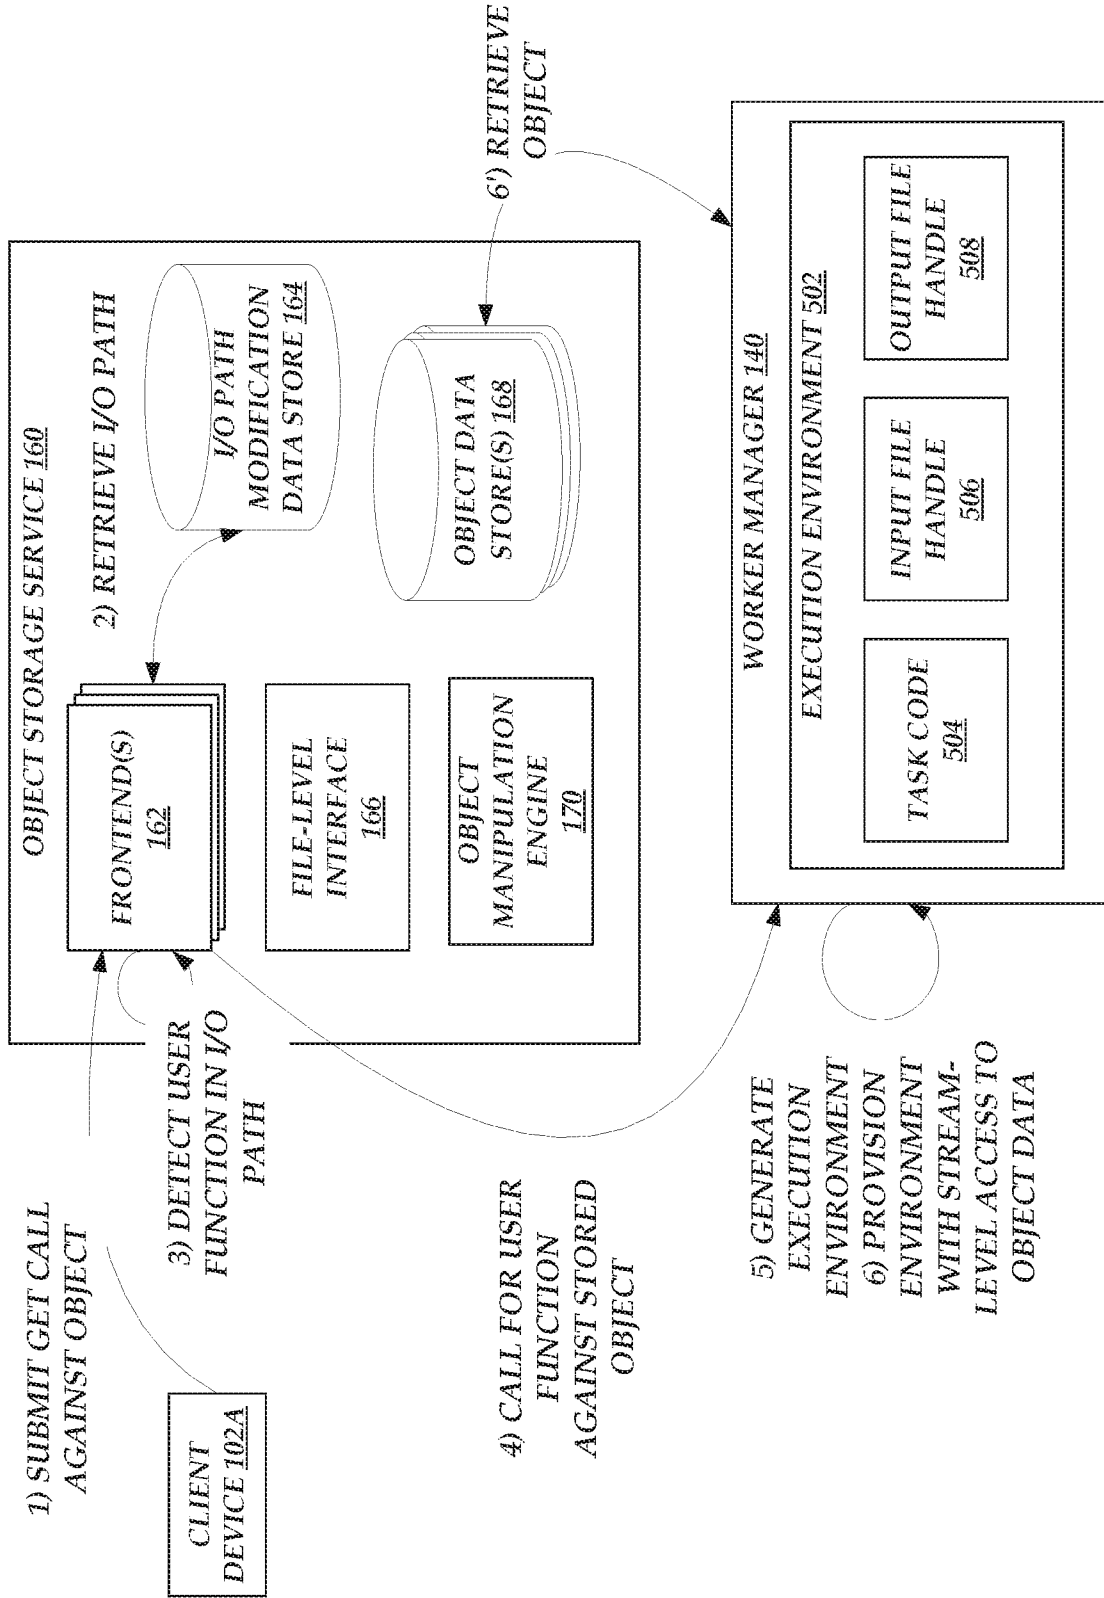


Fig. 6A

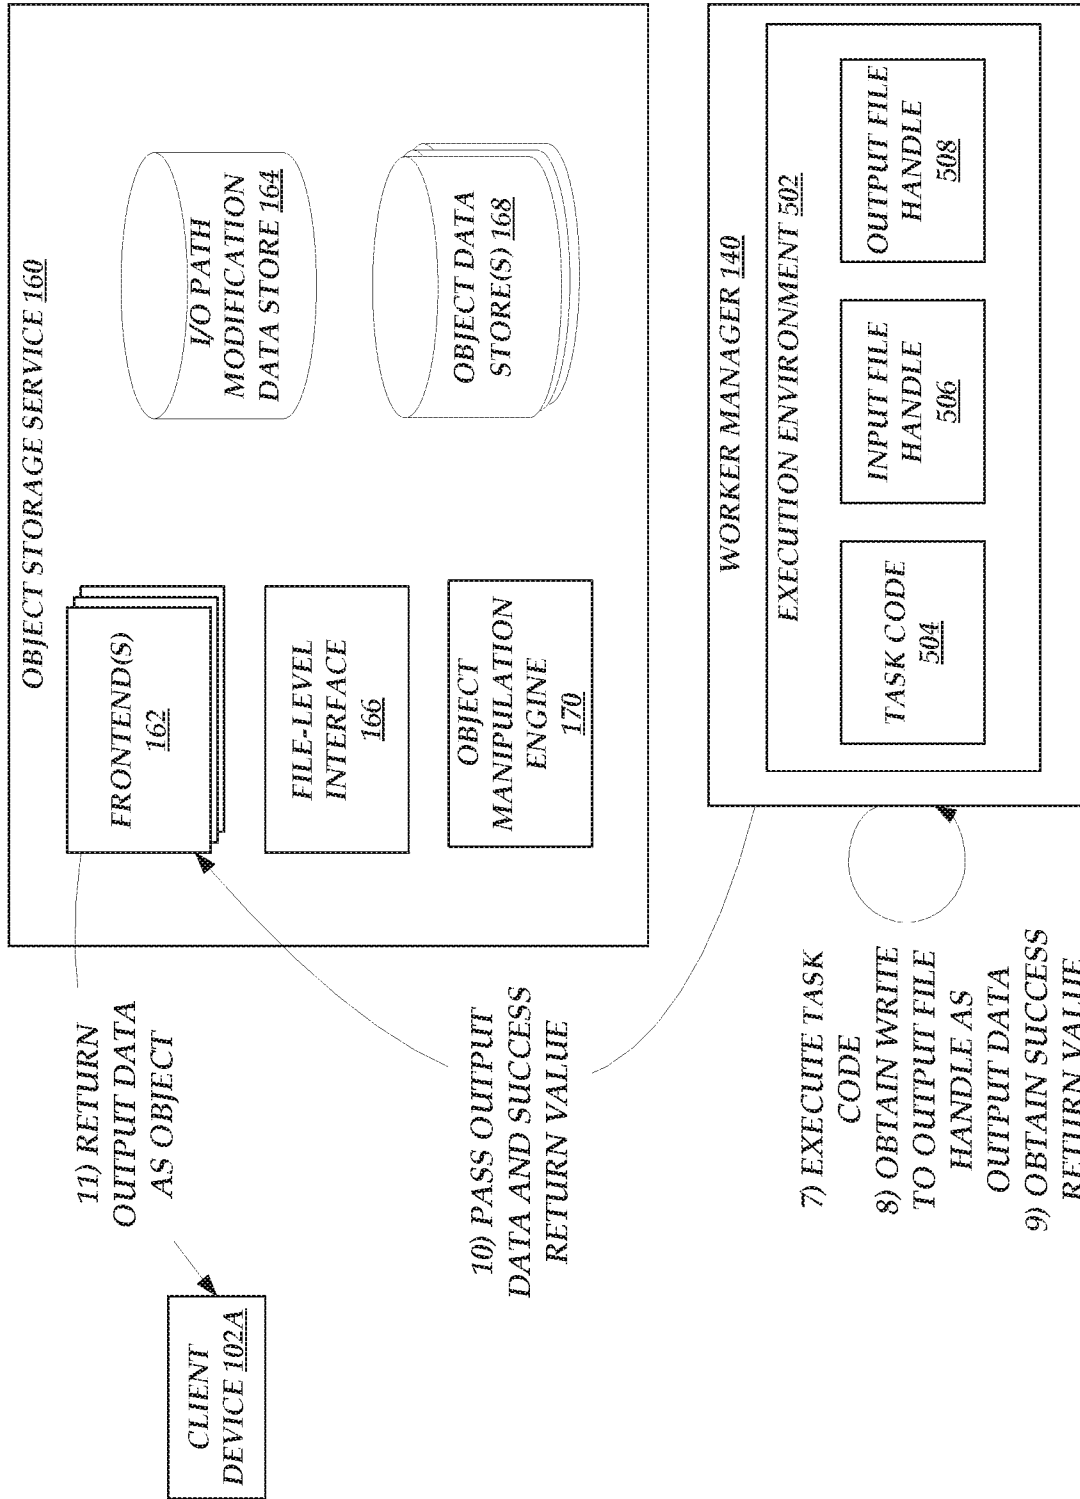


Fig. 6B

9/16

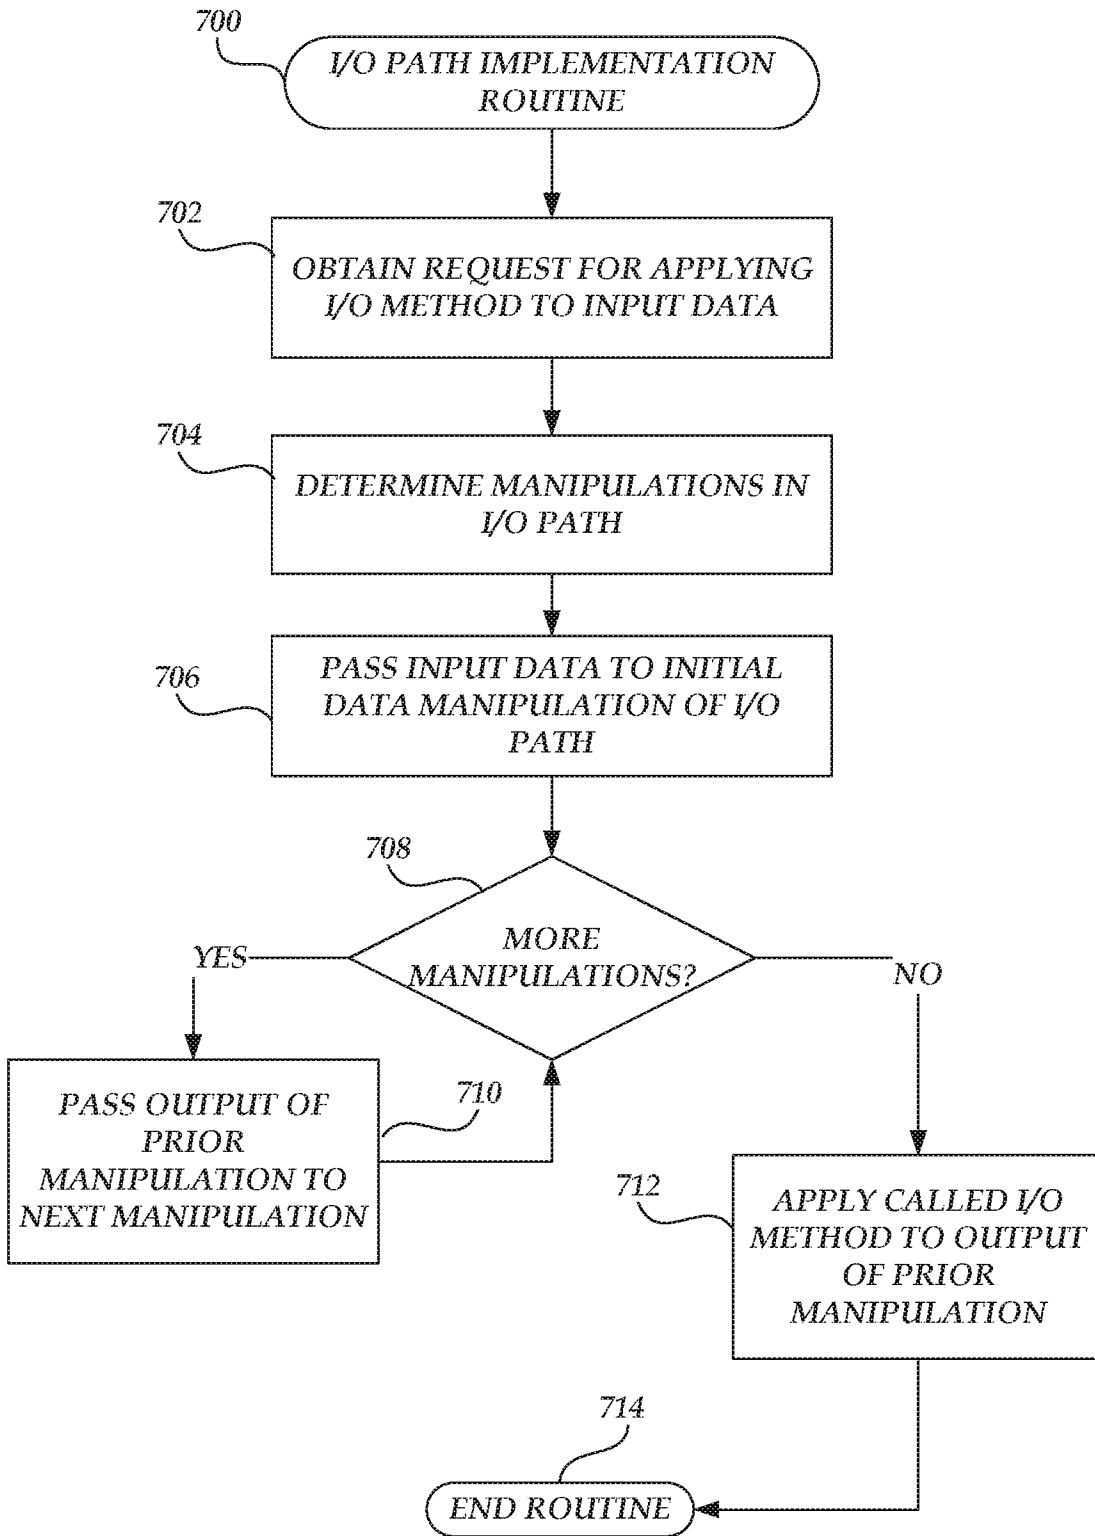


Fig. 7

10/16

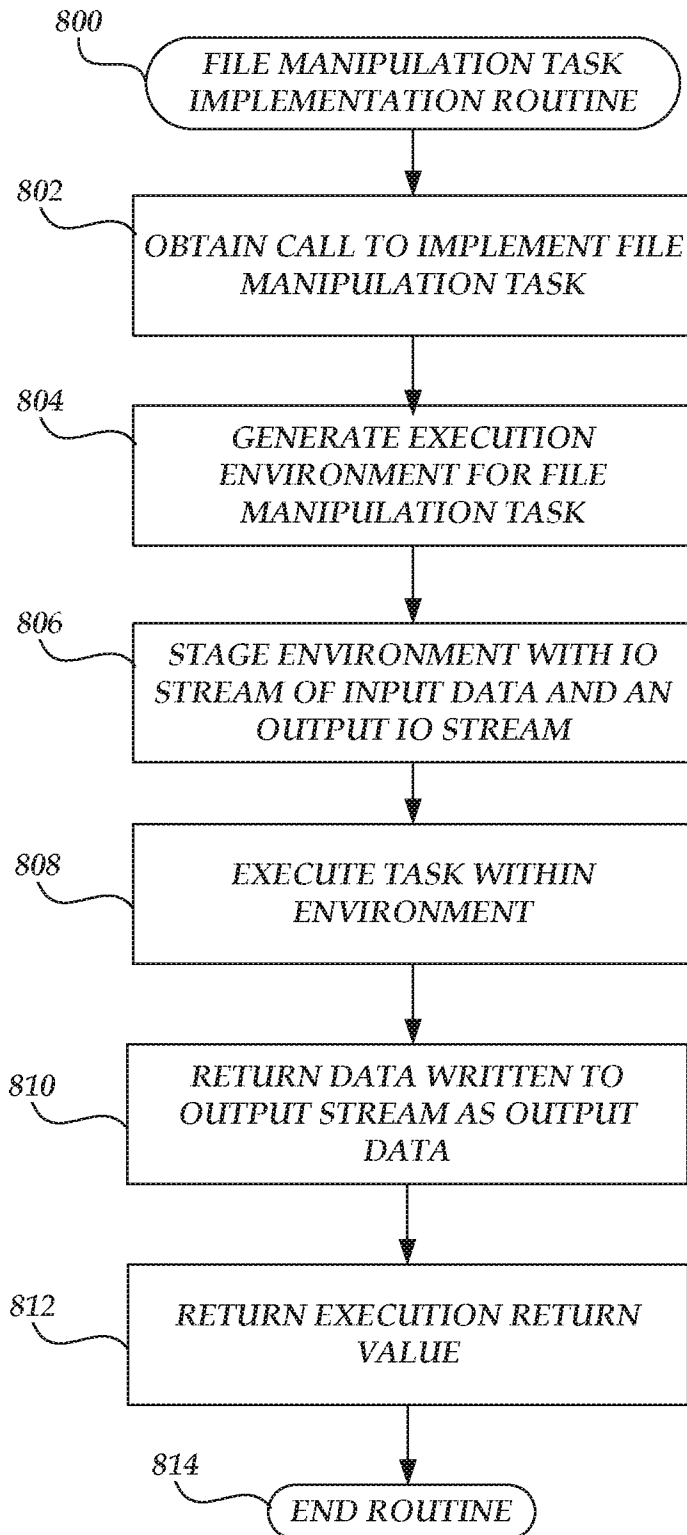


Fig. 8

11/16

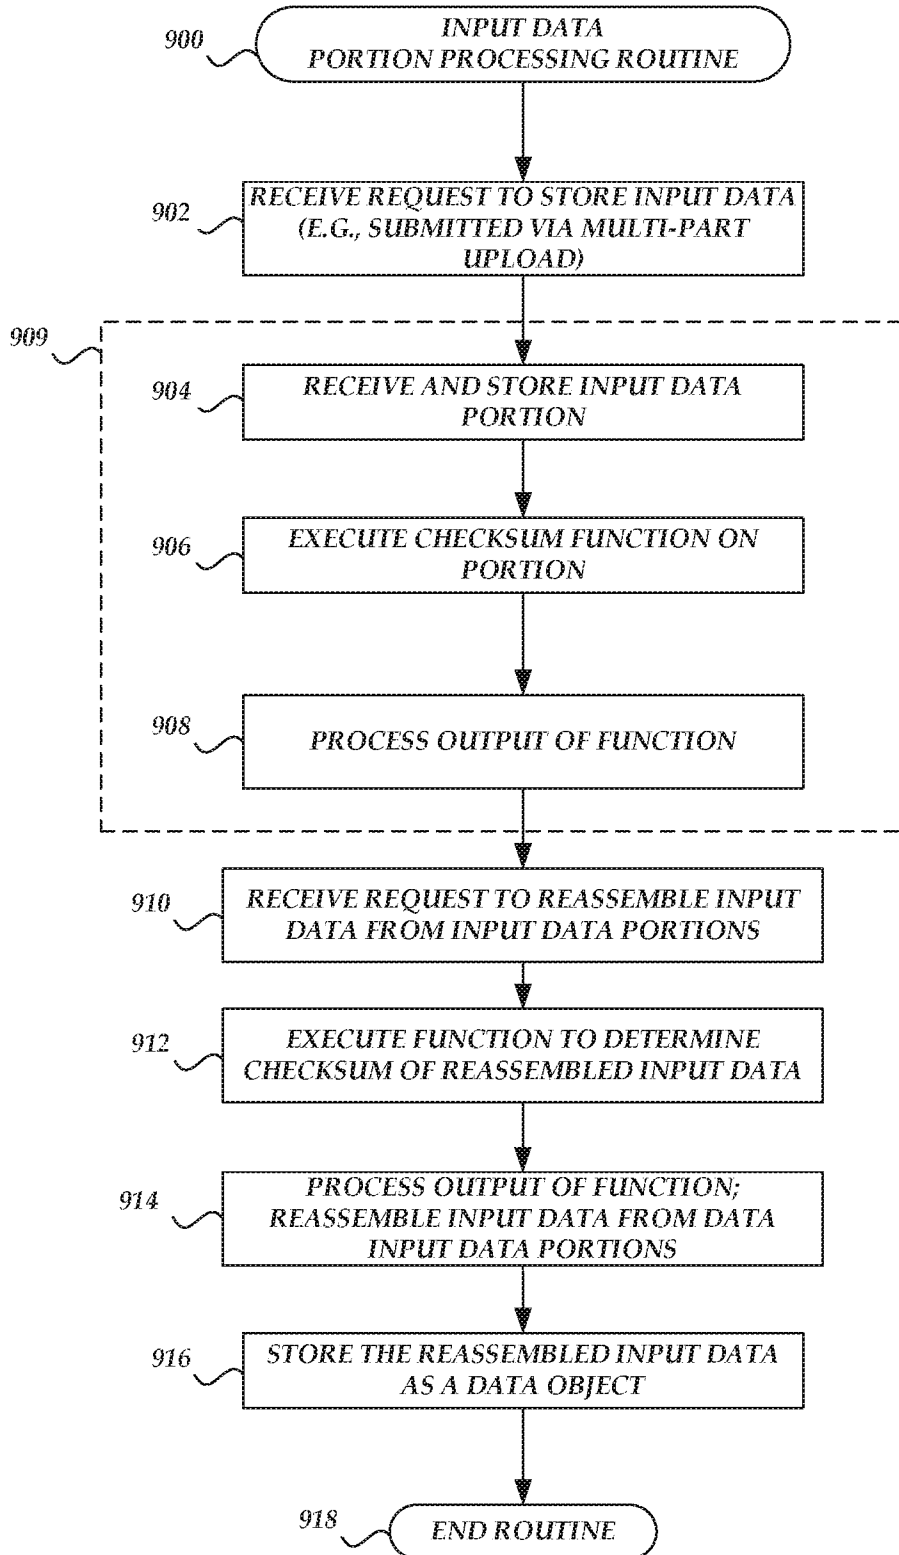


Fig. 9



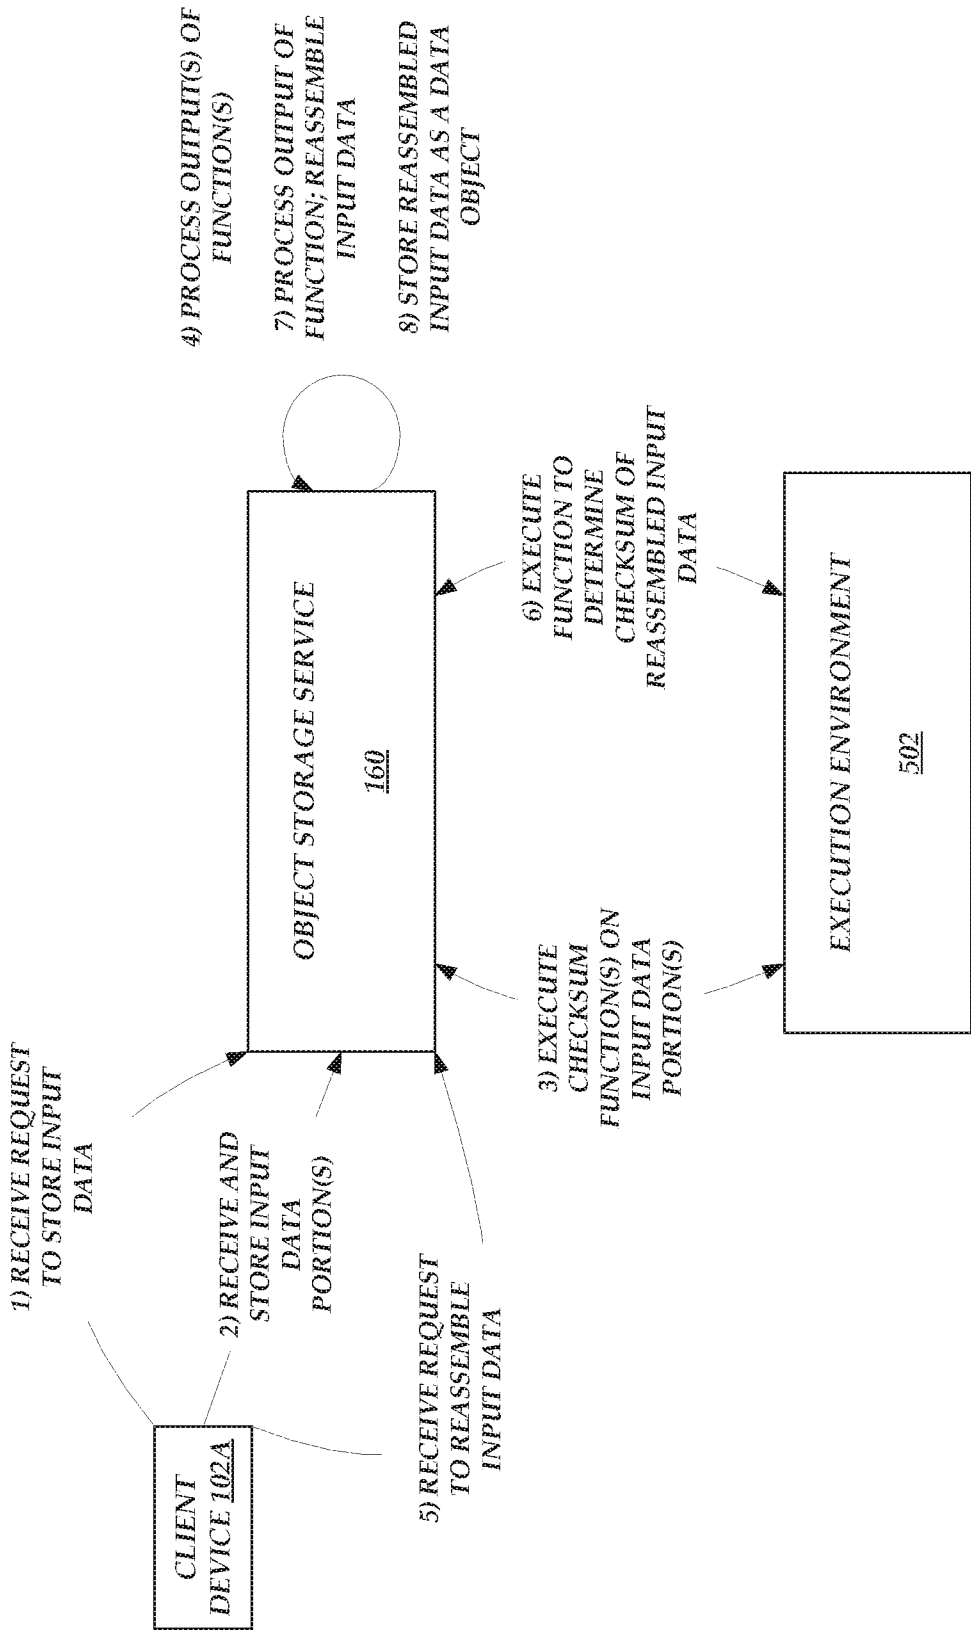
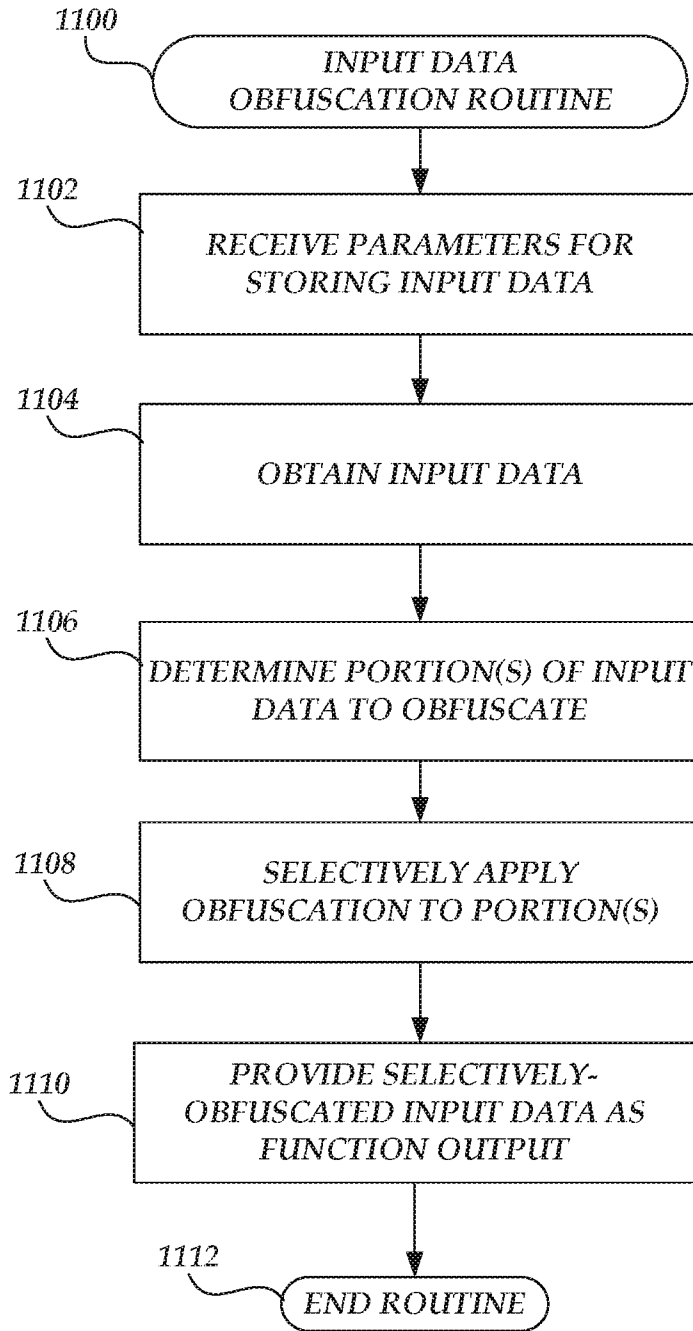


Fig. 10

13/16



*Fig. 11*

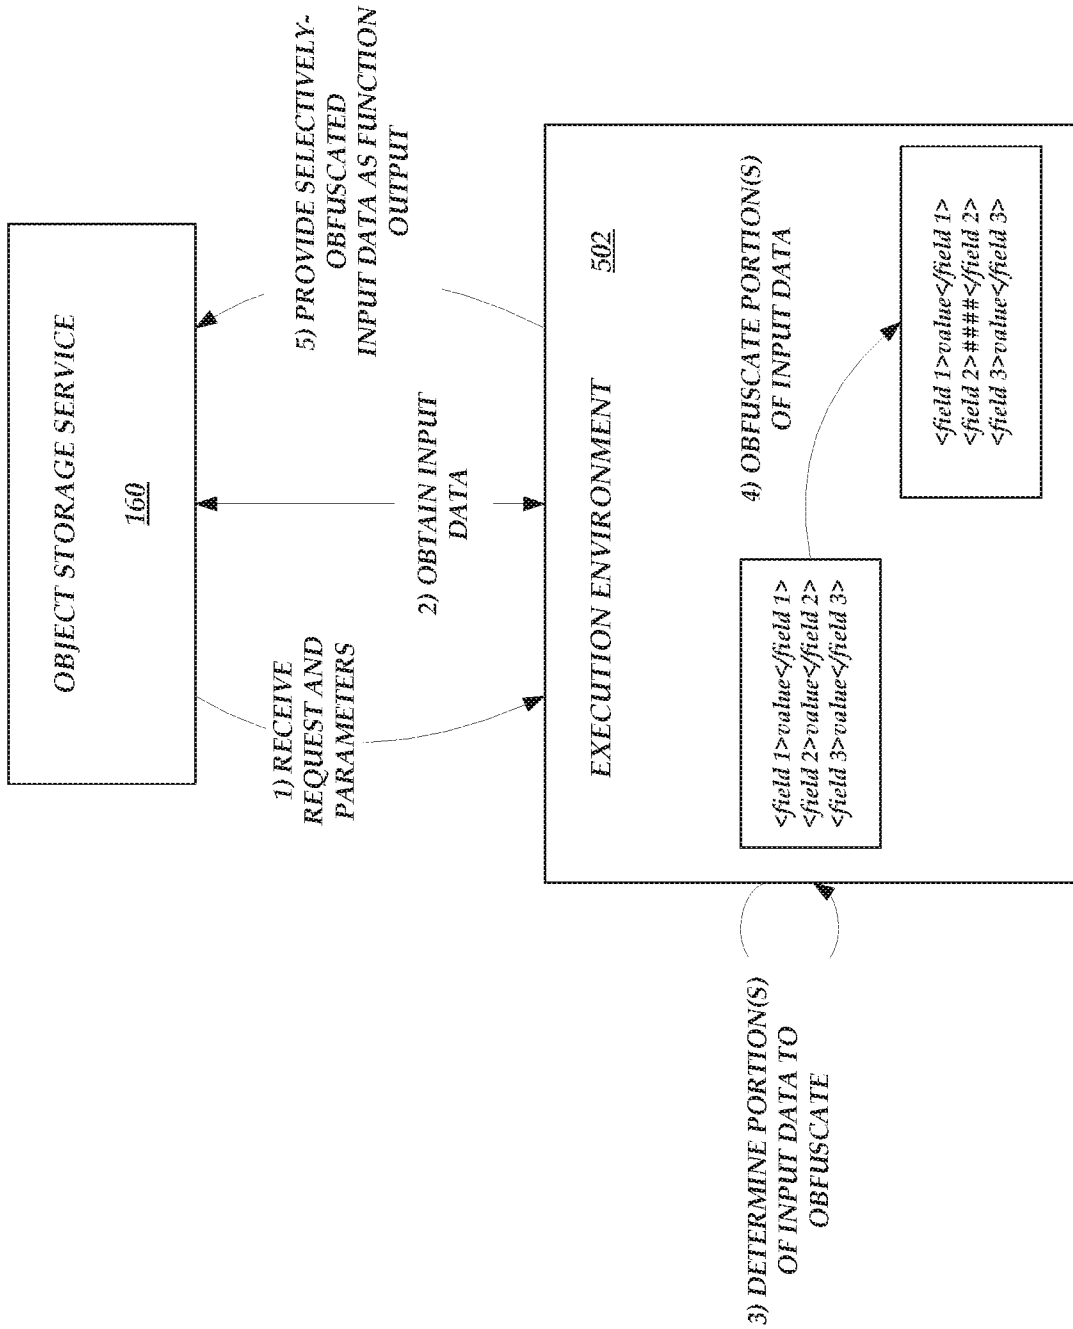


Fig. 12

15/16

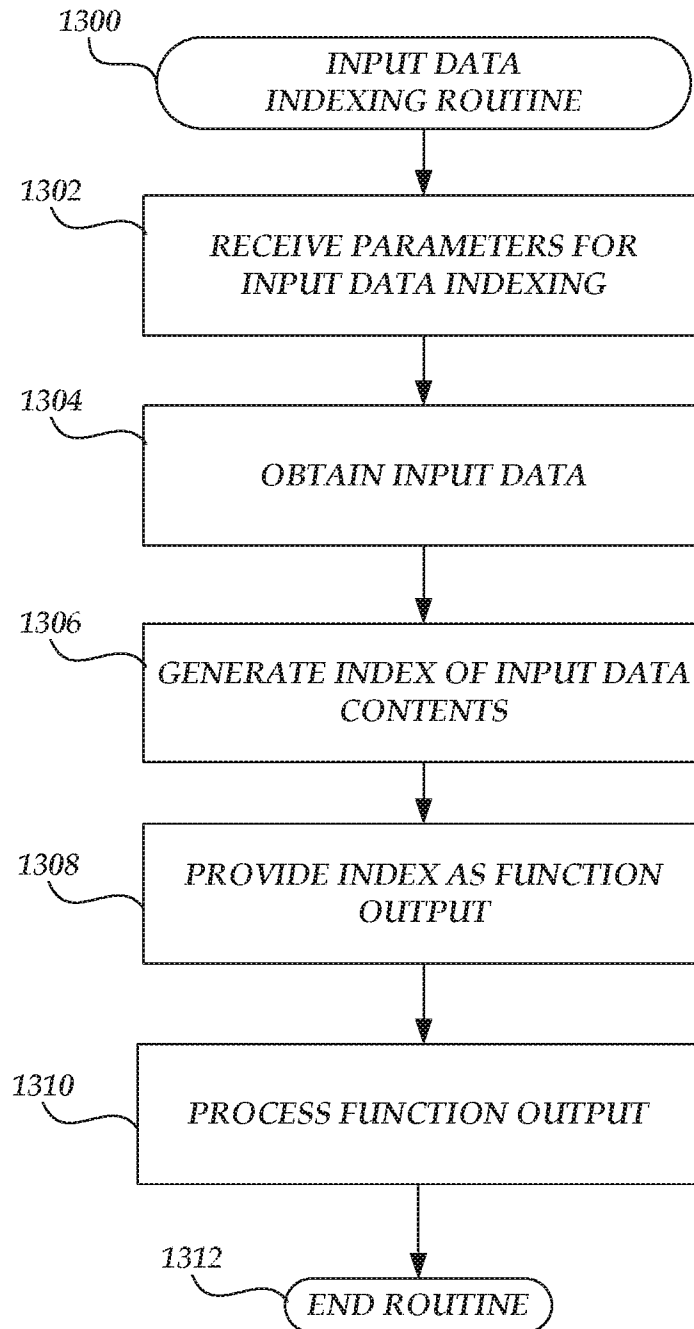
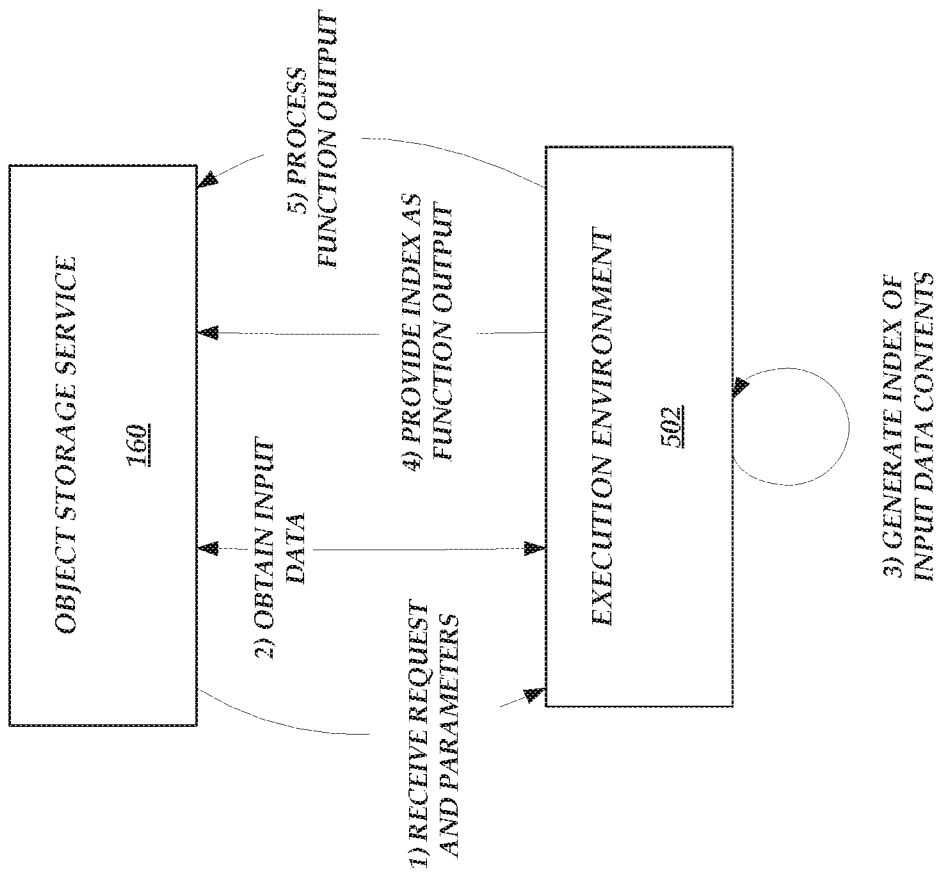


Fig. 13



**Fig. 14**

INTERNATIONAL SEARCH REPORT

International application No  
PCT/US2020/052280

A. CLASSIFICATION OF SUBJECT MATTER  
 INV. G06F21/62 G06F9/455 G06F9/50 G06F16/182  
 ADD.  
 According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED  
 Minimum documentation searched (classification system followed by classification symbols)  
 G06F

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

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)  
 EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2012/278621 A1 (WOLOSZYN TERRENCE PETER [CA]) 1 November 2012 (2012-11-01) paragraphs [0002], [0093] - [0096]; figures 1-4 paragraph [0150] - paragraph [0157]	1-15
Y	US 2016/092251 A1 (WAGNER TIMOTHY ALLEN [US]) 31 March 2016 (2016-03-31) paragraphs [0017], [0064]	1-15

Further documents are listed in the continuation of Box C.

See patent family annex.

\* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier application or patent but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
- "&" document member of the same patent family

Date of the actual completion of the international search

10 December 2020

Date of mailing of the international search report

17/12/2020

Name and mailing address of the ISA/  
 European Patent Office, P.B. 5818 Patentlaan 2  
 NL - 2280 HV Rijswijk  
 Tel. (+31-70) 340-2040,  
 Fax: (+31-70) 340-3016

Authorized officer

Mäenpää, Jari

## INTERNATIONAL SEARCH REPORT

International application No  
PCT/US2020/052280

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>Shri Rk Bigensana Singh ET AL: "Enhancing Cloud Data Security with Data Encryption &amp; Tokenization", International Journal of Current Trends in Engineering &amp; Research Scientific Journal Impact Factor, 1 May 2016 (2016-05-01), pages 191-196, XP055334669, Retrieved from the Internet: URL:<a href="http://www.ijcter.com/papers/volume-2/issue-5/eancing-cloud-data-security-with-d-ata-encryption-tokenization.pdf">http://www.ijcter.com/papers/volume-2/issue-5/eancing-cloud-data-security-with-d-ata-encryption-tokenization.pdf</a> [retrieved on 2017-01-12] the whole document -----</p>	1-15

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2020/052280

Patent document cited in search report	Publication date	Patent family member(s)	Publication date	
US 2012278621	A1	01-11-2012	CA 2775206 A1	27-10-2012
			CA 2775237 A1	27-10-2012
			CA 2775245 A1	27-10-2012
			CA 2775247 A1	27-10-2012
			CA 2775427 A1	27-10-2012
			EP 2702726 A1	05-03-2014
			US 2012278487 A1	01-11-2012
			US 2012278504 A1	01-11-2012
			US 2012278621 A1	01-11-2012
			US 2012278872 A1	01-11-2012
			US 2012278897 A1	01-11-2012
			WO 2012145827 A1	01-11-2012
			-----	
US 2016092251	A1	31-03-2016	CA 2962633 A1	07-04-2016
			CN 107111508 A	29-08-2017
			EP 3201768 A1	09-08-2017
			EP 3633506 A1	08-04-2020
			JP 6352535 B2	04-07-2018
			JP 2017534967 A	24-11-2017
			US 9652306 B1	16-05-2017
			US 2016092251 A1	31-03-2016
			US 2016239318 A1	18-08-2016
			US 2017371724 A1	28-12-2017
			WO 2016053973 A1	07-04-2016
-----				