



- (51) International Patent Classification:
G06F 21/33 (2013.01) *H04L 29/06* (2006.01)
G06F 21/62 (2013.01)
- (21) International Application Number:
PCT/US2016/065853
- (22) International Filing Date:
9 December 2016 (09.12.2016)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
14/965,055 10 December 2015 (10.12.2015) US
- (71) Applicant: BLUE TALON, INC. [US/US]; 541 Jefferson St, Ste. 202, Redwood City, California 94063 (US).
- (72) Inventors: ARUMUGAM, Dilli Dorai Minnal; 541 Jefferson St, Ste. 202, Redwood City, California 94063 (US).
MUJUMDAR, Prasad; 541 Jefferson St, Ste. 202, Redwood City, California 94063 (US).

- (74) Agent: GOREN, David J.; Fish & Richardson P.C., P.O. Box 1022, Minneapolis, Minnesota 55440-1022 (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, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK,

[Continued on next page]

(54) Title: POLICY ENFORCEMENT FOR COMPUTE NODES

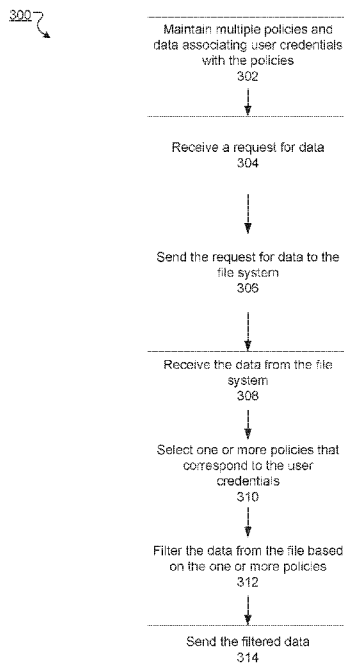


FIG. 3

(57) Abstract: Methods, systems, and apparatus, including computer programs encoded on computer storage media, for maintaining, by a policy enforcement system in a first compute node, a plurality of policies and data associating a plurality of user credentials with the plurality of policies. A request is received from a compute process for data from a file system in the first compute node. The request includes user credentials. The request for data is sent to the file system, and the data is received from the file system. Based on the received user credentials and the data associating the plurality of user credentials with the plurality of policies, one or more policies that correspond to the received user credentials is selected from the plurality of policies. The policy enforcement system filters the data from the file system based on the one or more policies, and sends the filtered data to the compute process.

WO 2017/100592 A1

SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG). **Published:**

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

POLICY ENFORCEMENT FOR COMPUTE NODES

BACKGROUND

A framework, e.g., Apache Hadoop, can be deployed to manage distributed storage and distributed processing of large data sets on clusters of many computers, i.e., nodes, which may be physical or virtual. The framework can include multiple components to be run on different nodes in the cluster. Each component can be responsible for a different task.

5 For example, a first component, e.g., Hadoop Distributed File System (HDFS), can implement a file system, and a second component, e.g., Hive, can implement a database access layer. The components work together to distribute processing of a workload of files among the nodes in the cluster.

10 A third component, e.g., YARN, of the framework can break up the workload into multiple tasks. In particular, the third component is a resource manager that assigns each task to a respective compute node in the cluster that performs computations of the task. Each compute node can retrieve a portion of data required for the task from the file system before executing a process that uses the portion of data to complete the task.

SUMMARY

In general, this specification describes a policy enforcement system.

15 In one aspect, a method includes maintaining, by a policy enforcement system in a first compute node, a plurality of policies and data associating a plurality of user credentials with the plurality of policies. A request is received from a compute process for data from a file system in the first compute node for use in a first task assigned to the first compute node by a resource manager managing a plurality of tasks comprising the first task and a plurality
20 of compute nodes comprising the first compute node. The request includes user credentials. The request for data is sent to the file system, and the data is received from the file system. Based on the received user credentials and the data associating the plurality of user credentials with the plurality of policies, one or more policies that correspond to the received user credentials is selected from the plurality of policies. The policy enforcement system
25 filters the data from the file system based on the one or more policies, and sends the filtered data to the compute process.

In another aspect, a computer-readable medium having instructions stored thereon, which, when executed by a processor, cause the processor to perform operations to perform method.

5 Features of either aspect may include one or more of the following. The one or more policies may cause the policy enforcement system to only include, in the filtered data, data accessible using the user credentials. The data from the file system may be organized by a table of columns and rows. Filtering the data from the file system may include determining, based on the one or more policies associated with the user credentials, that one or more columns of data in the table are restricted, and masking at least a portion of one or more columns of data. The masking may include encrypting or tokenizing the data. The masking may mask some data in a column and leave some data unmasked. Filtering the data from the file system may include determining, based on the one or more policies associated with the user credentials, that one or more rows of data in the table are restricted, and masking the one or more rows of data. The masking may include not returning the one or more masked rows to the compute process. The user credentials may include a user id and a group id, and the user id may associated with a first collection of policies, and the group id may be associated with a second collection of policies. Prior to the storing, user input may be received at a user interface associated with the policy enforcement system, the user input specifying one or more policies associated with the user credentials.

20 In another aspect, a system includes a client device and a file system that includes a master node and a plurality of compute nodes. The master node is configured to receive from the client device a request for data from the file system and divide a workload from the request into a plurality of tasks and assign the tasks to the plurality of compute nodes. Each compute node includes a processor and a computer-readable medium coupled to the processor and having instructions stored thereon, which, when executed by the processor, cause the processor to operating a compute process and a policy enforcement system. The instructions cause the processor to perform operations as follows. A first task is received at the compute process from the master node. A plurality of policies and data associating a plurality of user credentials with the plurality of policies are maintained by the policy enforcement system. A request for data from a file system in the compute node for use in the first task is received in the policy enforcement system from the compute process. The

request includes user credentials. The request for data is sent from the policy enforcement system to the file system. The data from the file system is received in the policy enforcement system. Based on the received user credentials and the data associating the plurality of user credentials with the plurality of policies, one or more policies that correspond to the received
5 user credentials is selected from the plurality of policies. The data from the file system is filtered by the policy enforcement system based on the one or more policies. The filtered data is sent from the policy enforcement system to the compute process, and a result is returned from the compute process in the compute node to the master node.

Features may include one or more of the following. The data from the file system
10 may be organized by a table of columns and rows. Filtering the data from the file system may include determining, based on the one or more policies associated with the user credentials, that one or more columns or rows of data in the table are restricted, and masking at least a portion of one or more columns or rows of data. The masking may include encrypting or tokenizing data in one or more columns. The masking may include not
15 returning one or more rows to the compute process.

Particular embodiments of the subject matter described in this specification can be implemented so as to realize one or more of the following advantages. Traditionally, a compute node is able to retrieve, or in some cases modify, data from a file system as long as the compute node has proper read or write access. With proper access permissions, the file
20 system returns the data without any filters or modifications. Using policies, the policy enforcement system can mask particular portions of the returned data that may be sensitive, e.g., personally identifiable information such as social security numbers or dates of birth. In addition, the policy enforcement system can filter out certain records based on the data in the record, e.g., a policy could allow only access to records record with field State having value
25 'CA'. This can provide an additional layer of privacy for user data stored in the file system. Also, when communicating with the file system, the system operates without interfering with existing communication protocols between a client and nodes within the file system.

The details of one or more embodiments of the subject matter of this specification are set forth in the accompanying drawings and the description below. Other features, aspects,
30 and advantages of the subject matter will become apparent from the description, the drawings, and the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 illustrates an example architecture for distributing tasks to compute nodes.

FIG. 2 illustrates an example architecture for integrating a policy enforcement system in a compute node.

FIG. 3 is a flow diagram of an example method for completing a task using the policy enforcement system.

Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION

Sensitive data, such as personally identifiable data, can be stored in a file system, e.g., HDFS, in the cloud. Compute nodes in the cluster can operate on the sensitive data retrieved from the file system to complete tasks within a workload. There can be restrictions to read or write access for some compute nodes due to the sensitivity of the data. Some database systems can, in response to a data request from a client without authorized user credentials, send a response indicating unauthorized access to the data; such gatekeeping is performed at the central server that receives the request from the client. However, for a system that has distributed storage and processing, e.g., HDFS, it can be desirable for policies to be enforced at each compute node. In addition, it can be desirable for a compute node to receive a modified version of the data instead of a response indicating unauthorized access. For example, it can be desirable for the compute node to receive a response in which some (but not all) of the data has been masked, e.g., characters replaced with a default characters or values.

Thus, a policy enforcement system integrated in a compute node can provide a masked version of the sensitive data. This can enable the compute node to access portions of the sensitive data without having the sensitive data be entirely exposed. The system is operable even where data is stored in a framework, e.g., HDFS, that includes name nodes and data nodes.

FIG. 1 illustrates an example architecture of a file system 100 for distributing tasks to compute nodes 110-116. One or more client devices 50 can communicate with the file system 100. The client device 50 can include a memory, e.g., a random access memory (RAM), for storing instructions and data and a processor for executing stored instructions.

The memory can include both read only and writable memory. For example, the client device 50 can be a computer coupled to the file system 100 through a data communication network, e.g., local area network (LAN) or wide area network (WAN), e.g., the Internet, or a combination of networks. Communication over the network can occur using TCP/IP protocols.

The client device 102 can submit a job to the file system 100. The job can involve various tasks that are either spread across multiple compute nodes 110-116.

An application 102 can be running on one or more nodes within a cluster. Each node can include a memory, e.g., a random access memory (RAM), for storing instructions and data and a processor for executing stored instructions. The memory can include both read only and writable memory. The application 102 can be configured to process workloads. For example, the application 102 can be configured to perform an aggregation operation on a large dataset.

A resource manager 104, which can serve as a master node, can analyze the workload provided to the application 102 and divide the workload into multiple tasks. The resource manager 104 assigns each task to a respective compute node, e.g., one of compute nodes 110-116. For example, the resource manager 104 can receive the job from the client 50, and divide the job into multiple individual tasks. The resource manager 104 can allocate the tasks so that each task operates on an independent portion of data from a file system. Each independent portion of the data can be located in a different compute node. The compute nodes complete their assigned tasks, and in some implementations, return a result to the application 102, e.g., to the resource manager 104. The application 102 can communicate with each compute node through a data communication network, e.g., local area network (LAN) or wide area network (WAN), e.g., the Internet, or a combination of networks. Communication over the network can occur using TCP/IP protocols.

The resource manager 104 can return the result of the job to the client device 50.

To perform an assigned task, a compute node communicates with a distributed file system 106. The distributed file system 106 provides a portion of data required for completing the assigned task. The compute node can then perform operations on the portion of data as specified by the application 102.

Each compute node can include a policy enforcement system. The policy enforcement system can communicate with a policy store 108, which maintains policies within the cluster. The compute node and the policy enforcement system will be described further below with reference to FIG. 2.

5 Policies limit access to portions of data stored in the distributed file system 106. By way of illustration, data in the distributed file system 106 can be stored in a table of rows and columns. Some columns in the table may include data about, e.g., names, dates of birth, phone numbers, credit card information, social security numbers, or other personal information. The rows can include data about the individuals, e.g. sorted by unique
10 identifier. Some policies allow user credentials of one particular person to access data from a column, or can allow user credentials of a group of people to access data from a column. Some other policies prevent any user from accessing complete data from a column. Policies can restrict access to these columns in the data to users with approved user credentials. This will be described in further detail with reference to FIG. 3. In some implementations,
15 policies are applicable to other formats for the data. The other formats can include unstructured data, key-value stores, document stores, or graphs.

Thus, each policy can be associated with user credentials of one or more users. Multiple user credentials can also be associated with one policy. User credentials can include a user identification, a group identification, or both. In some implementations, the
20 user identification is associated with a first collection of policies and the group identification is associated with a second collection of policies. These policies can be stored in the policy store 108. An example policy can be a data structure that specifies a policy identifier, a column, and a function to be executed on the column, e.g., mask everything except the last four characters of data in the column. An example association can be stored in a table with a
25 column for a user credential identifier and a column for a policy identifier with a user credential identifier associated with the policy identifier in the same row.

In some implementations, an administrator can generate a policy using a user interface that communicates with the policy store 108. The administrator can provide, at the user interface, user input specifying one or more policies associated with particular user
30 credentials, and the policy store 108 can store the one or more policies and the associations.

FIG. 2 illustrates an example architecture 200 for integrating a policy enforcement system 208 in a compute node 202. The compute node 202 can be an example of any of the compute nodes 110-116.

A user can send a request to an application 204 to process a workload divided into multiple tasks. The request can include user credentials. The application 204, or alternatively, a resource manager, can assign a task of the workload to the compute node 202. The application 204 can provide the user credentials to the compute node 202. In some implementations, a policy enforcement system 208 stores the user credentials in an internal database for later use.

In some implementations, the user credentials stored in the internal database are independent of or in addition to user credentials in the distributed file system 214 and the application 204. That is, the policy enforcement system 208 can manage, e.g., through the user interface, user credentials that are not managed by the distributed file system 214 or the application 204. In some implementations, the user credentials at the policy enforcement system 208 have group identifiers or category identifiers. Policies can apply a masking function, which will be described below, to data associated with user credentials of particular group or category identifiers.

The policy enforcement system 208 communicates with the policy store 108 to obtain policies associated with the user credentials stored in the internal database. The policy enforcement system 208 can store the policies in the internal database and associations of the one or more of the policies with the user credentials.

The application 204 can communicate with a compute process 206. The compute process 206 executes on the compute node 202 to complete the assigned task from the application 204. To complete the task, the compute process 206 obtains data from a policy enforcement system 208, which intercepts data from a file driver 212.

The file driver 212 communicates with the distributed file system 106. In some implementations, the file driver 212 obtains a portion of data from “local” storage 210, i.e., storage for the distributed file system 106 that is logically associated with the compute node 202. The local storage can be a physical or virtual drive associated with the particular compute node 202. In some implementations, the file driver 212 obtains a portion of data from the non-local storage 214, e.g., a physical or virtual drive associated with another node.

To complete the task, the compute process 206 requests the portion of data from the distributed file system 106. The policy enforcement system 208 can perform operations on the portion of data before providing the portion of data to the compute process 206. The operations performed by the policy enforcement system 208 will be described further below with reference to FIG. 3.

In some implementations, the policy enforcement system 203 operates between the distributed file system 214 and the compute node 202.

FIG. 3 is a flow diagram of an example method 300 for completing a task using the policy enforcement system. For convenience, the method 300 will be described with respect to a policy enforcement system, e.g., the policy enforcement system 208 of FIG. 2, executing software to implement the method 300.

The policy enforcement system maintains multiple policies and data associating user credentials with the multiple policies (step 302). The policies can be stored in an internal database as described above or obtained from a policy store, e.g., the policy store 108 of FIG. 2. The resource manager, can receive a job from a client. The application, e.g., the resource manager, sends a request for data to be retrieved from a file system, e.g., a distributed file system 214 of FIG. 2, for use in a task assigned to a compute node. The compute node can receive the request at a compute process, e.g., the compute process 206 of FIG. 2. The compute process can provide the request to the policy enforcement system, e.g., through a file system interface command, and the policy enforcement system receives the request (step 304). The request includes user credentials.

The policy enforcement system sends the request for data to the file system (step 306). In some implementations, the policy enforcement system sends the request for data to a file driver, as described above with reference to FIG. 2, and the file driver sends the request to the file system. In particular, the policy enforcement system sends the request for data to the local file system 210 within the node. For example, the tasks may be divided by the resource manager according to which compute node contains data to be access by the task; thus the policy enforcement system of a particular node need not access the local storage data of any other node.

In some implementations, the file system authenticates the request using the user credentials, and if the user credentials are valid, provides the data to the policy enforcement

a programming language like JavaScript. If the predicate evaluates to true, the policy enforcement system can mask the data. In some implementations, the policy enforcement system also encrypts or tokenizes the data before sending the data to the compute process, e.g., using a custom encryption or tokenization function.

5 In some implementations, the policy enforcement system can filter the data such that certain records are not returned to the client. Thus, if the predicate evaluates to true for a row, that is removed from the data returned to the client. For example a policy could allow only access to records record with field State having value 'CA'.

10 By way of illustration, the policy can specify for a "social-security-number" column that only a last four characters can be provided to the compute process and that for a "date-of-birth" column only the year can be provided to the compute process. Therefore, the policy enforcement system can replace any character that is not in the last four characters with an 'x'. See Table 2.

15

name	social-security-number	date-of-birth
John Smith	xxx-xx-6789	xx-xx-56
Jane Smith	xxx-xx-3333	xx-xx-33

Table 2

The policy enforcement system can also filter the data by simply not returning one or more columns. For example, a policy can specify that "social-security-number" column will not be provided to the compute process.

20 In some implementations, the policy enforcement system can also enforce file level operations such as read, write, create, delete, etc., based on the policies. If a particular user does not have authorization to execute the operation on entire file, e.g., the policy does not authorized the operation for the provided the credentials, then the request could be rejected immediately. In this case, the policy enforcement system can return an error to the compute node. This can enable an administrator to store all access control policies at a variety of
25 degrees of granularity (coarse grain file level access, as well as fine grain filter/masking of data) in single policy enforcement system.

The policy enforcement system then sends the filtered data to the compute process (step 314). In this way, the policy enforcement system, as integrated with a compute node, does not interfere with existing file retrieval protocols between compute nodes and file systems. The compute process can then return the filtered data to the application.

5 As described above, in the response returned to the compute, some (but not all) of the data has been masked, e.g., characters replaced with a default characters or values. In addition, because the policies can be associated with individual users, different users can receive responses with different masking. For example, a first user might not be allowed to see any portion of a credit card number (so all digit of the credit card number would be masked), a second user might be allowed to see the last 4 digits (so the remaining digits
10 would be masked), and a third user might be permitted to see the whole credit card number. By masking the data differently based on the policies, all users can run the same applications that read the full data records. The data each user sees will depend on their access privileges.

Also, the policy enforcement system can provide an additional level of security
15 without requiring the compute process to handle a different data structure, which minimizes cost and complexity of integrating the policy enforcement system. By way of illustration, the compute process expects a data structure with three columns based on request for data from “name”, “social-security-number”, and “date-of-birth” columns. The policy enforcement system can provide a data structure with the “name”, “social-security-number”, and “date-of-
20 birth” columns, although data within the columns are masked.

Embodiments of the subject matter and the functional operations described in this specification can be implemented in digital electronic circuitry, in tangibly-embodied computer software or firmware, in computer hardware, including the structures disclosed in this specification and their structural equivalents, or in combinations of one or more of them.
25 Embodiments of the subject matter described in this specification can be implemented as one or more computer programs, i.e., one or more modules of computer program instructions encoded on a tangible non-transitory program carrier for execution by, or to control the operation of, data processing apparatus. Alternatively or in addition, the program instructions can be encoded on an artificially-generated propagated signal, e.g., a machine-
30 generated electrical, optical, or electromagnetic signal, that is generated to encode information for transmission to suitable receiver apparatus for execution by a data processing

apparatus. The computer storage medium can be a machine-readable storage device, a machine-readable storage substrate, a random or serial access memory device, or a combination of one or more of them. The computer storage medium is not, however, a propagated signal.

5 The term “data processing apparatus” encompasses all kinds of apparatus, devices, and machines for processing data, including by way of example a programmable processor, a computer, or multiple processors or computers. The apparatus can include special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application-specific integrated circuit). The apparatus can also include, in addition to
10 hardware, code that creates an execution environment for the computer program in question, e.g., code that constitutes processor firmware, a protocol stack, a database management system, an operating system, or a combination of one or more of them.

 A computer program (which may also be referred to or described as a program, software, a software application, a module, a software module, a script, or code) can be
15 written in any form of programming language, including compiled or interpreted languages, or declarative or procedural languages, and it can be deployed in any form, including as a stand-alone program or as a module, component, subroutine, or other unit suitable for use in a computing environment. A computer program may, but need not, correspond to a file in a file system. A program can be stored in a portion of a file that holds other programs or data,
20 e.g., one or more scripts stored in a markup language document, in a single file dedicated to the program in question, or in multiple coordinated files, e.g., files that store one or more modules, sub-programs, or portions of code. A computer program can be deployed to be executed on one computer or on multiple computers that are located at one site or distributed across multiple sites and interconnected by a communication network.

25 As used in this specification, an “engine,” or “software engine,” refers to a software implemented input/output system that provides an output that is different from the input. An engine can be an encoded block of functionality, such as a library, a platform, a software development kit (“SDK”), or an object. Each engine can be implemented on any appropriate type of computing device, e.g., servers, mobile phones, tablet computers, notebook
30 computers, music players, e-book readers, laptop or desktop computers, PDAs, smart phones, or other stationary or portable devices, that includes one or more processors and computer

readable media. Additionally, two or more of the engines may be implemented on the same computing device, or on different computing devices.

The processes and logic flows described in this specification can be performed by one or more programmable computers executing one or more computer programs to perform
5 functions by operating on input data and generating output. The processes and logic flows can also be performed by, and apparatus can also be implemented as, special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application-specific integrated circuit).

Computers suitable for the execution of a computer program include, by way of
10 example, can be based on general or special purpose microprocessors or both, or any other kind of central processing unit. Generally, a central processing unit will receive instructions and data from a read-only memory or a random access memory or both. The essential elements of a computer are a central processing unit for performing or executing instructions and one or more memory devices for storing instructions and data. Generally, a computer
15 will also include, or be operatively coupled to receive data from or transfer data to, or both, one or more mass storage devices for storing data, e.g., magnetic, magneto-optical disks, or optical disks. However, a computer need not have such devices. Moreover, a computer can be embedded in another device, e.g., a mobile telephone, a personal digital assistant (PDA), a mobile audio or video player, a game console, a Global Positioning System (GPS) receiver,
20 or a portable storage device, e.g., a universal serial bus (USB) flash drive, to name just a few.

Computer-readable media suitable for storing computer program instructions and data include all forms of non-volatile memory, media and memory devices, including by way of example semiconductor memory devices, e.g., EPROM, EEPROM, and flash memory
25 devices; magnetic disks, e.g., internal hard disks or removable disks; magneto-optical disks; and CD-ROM and DVD-ROM disks. The processor and the memory can be supplemented by, or incorporated in, special purpose logic circuitry.

To provide for interaction with a user, embodiments of the subject matter described in this specification can be implemented on a computer having a display device, e.g., a CRT
(cathode ray tube) monitor, an LCD (liquid crystal display) monitor, or an OLED display, for
30 displaying information to the user, as well as input devices for providing input to the computer, e.g., a keyboard, a mouse, or a presence sensitive display or other surface. Other

kinds of devices can be used to provide for interaction with a user as well; for example, feedback provided to the user can be any form of sensory feedback, e.g., visual feedback, auditory feedback, or tactile feedback; and input from the user can be received in any form, including acoustic, speech, or tactile input. In addition, a computer can interact with a user
5 by sending resources to and receiving resources from a device that is used by the user; for example, by sending web pages to a web browser on a user's client device in response to requests received from the web browser.

Embodiments of the subject matter described in this specification can be implemented in a computing system that includes a back end component, e.g., as a data server, or that
10 includes a middleware component, e.g., an application server, or that includes a front end component, e.g., a client computer having a graphical user interface or a Web browser through which a user can interact with an implementation of the subject matter described in this specification, or any combination of one or more such back end, middleware, or front end components. The components of the system can be interconnected by any form or
15 medium of digital data communication, e.g., a communication network. Examples of communication networks include a local area network ("LAN") and a wide area network ("WAN"), e.g., the Internet.

The computing system can include clients and servers. A client and server are generally remote from each other and typically interact through a communication network.
20 The relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other.

While this specification contains many specific implementation details, these should not be construed as limitations on the scope of any invention or of what may be claimed, but rather as descriptions of features that may be specific to particular embodiments of particular
25 inventions. Certain features that are described in this specification in the context of separate embodiments can also be implemented in combination in a single embodiment. Conversely, various features that are described in the context of a single embodiment can also be implemented in multiple embodiments separately or in any suitable subcombination. Moreover, although features may be described above as acting in certain combinations and
30 even initially claimed as such, one or more features from a claimed combination can in some

cases be excised from the combination, and the claimed combination may be directed to a subcombination or variation of a subcombination.

Similarly, while operations are depicted in the drawings in a particular order, this should not be understood as requiring that such operations be performed in the particular
5 order shown or in sequential order, or that all illustrated operations be performed, to achieve desirable results. In certain circumstances, multitasking and parallel processing may be advantageous. Moreover, the separation of various system modules and components in the
10 embodiments described above should not be understood as requiring such separation in all embodiments, and it should be understood that the described program components and systems can generally be integrated together in a single software product or packaged into multiple software products.

Particular embodiments of the subject matter have been described. Other
embodiments are within the scope of the following claims. For example, the actions recited
15 in the claims can be performed in a different order and still achieve desirable results. As one example, the processes depicted in the accompanying figures do not necessarily require the particular order shown, or sequential order, to achieve desirable results. In certain implementations, multitasking and parallel processing may be advantageous.

WHAT IS CLAIMED IS:

1. A method, comprising:
 - maintaining, by a policy enforcement system in a first compute node, a plurality of policies and data associating a plurality of user credentials with the plurality of policies;
 - receiving, from a compute process, a request for data from a file system in the first
 - 5 compute node for use in a first task, the request further comprising user credentials, the first task assigned to the first compute node by a resource manager managing a plurality of tasks comprising the first task and a plurality of compute nodes comprising the first compute node;
 - sending the request for data to the file system;
 - receiving the data from the file system;
 - 10 selecting from the plurality of policies, based on the received user credentials and the data associating the plurality of user credentials with the plurality of policies, one or more policies that correspond to the received user credentials;
 - filtering, by the policy enforcement system, the data from the file system based on the one or more policies; and
 - 15 sending the filtered data to the compute process.
2. The method of claim 1, wherein the one or more policies cause the policy enforcement system to only include, in the filtered data, data accessible using the user credentials.
3. The method of claim 1, wherein the data from the file system is organized by a table
- 20 of columns and rows, wherein filtering the data from the file system comprises:
 - determining, based on the one or more policies associated with the user credentials, that one or more columns of data in the table are restricted; and
 - masking at least a portion of one or more columns of data.
4. The method of claim 3, wherein the masking comprises encrypting or tokenizing the
- 25 data.
5. The method of claim 3, wherein the masking masks some data in a column and leaves some data unmasked.

6. The method of claim 1, wherein the data from the file system is organized by a table of columns and rows, wherein filtering the data from the file system comprises:

determining, based on the one or more policies associated with the user credentials, that one or more rows of data in the table are restricted; and

5 masking the one or more rows of data.

7. The method of claim 6, wherein the masking comprises not returning the one or more masked rows to the compute process.

8. The method of claim 1, wherein the user credentials comprise a user id and a group id, and wherein the user id is associated with a first collection of policies, and wherein the group id is associated with a second collection of policies.

10

9. The method of claim 1, further comprising:

prior to the storing, receiving user input at a user interface associated with the policy enforcement system, the user input specifying one or more policies associated with the user credentials.

15

10. A computer-readable medium having instructions stored thereon, which, when executed by a processor, cause the processor to perform operations comprising:

maintaining, by a policy enforcement system in a first compute node, a plurality of policies and data associating a plurality of user credentials with the plurality of policies;

20

receiving, from a compute process, a request for data from a file system in the first compute node for use in a first task, the request further comprising user credentials, the first task assigned to the first compute node by a resource manager managing a plurality of tasks comprising the first task and a plurality of compute nodes comprising the first compute node;

sending the request for data to the file system;

receiving the data from the file system;

25

selecting from the plurality of policies, based on the received user credentials and the data associating the plurality of user credentials with the plurality of policies, one or more policies that correspond to the received user credentials;

filtering, by the policy enforcement system, the data from the file system based on the

one or more policies; and

 sending the filtered data to the compute process.

11. The computer program product of claim 10, wherein the one or more policies cause the policy enforcement system to only include, in the filtered data, data accessible using the user credentials.

12. The computer program product of claim 10, wherein the data from the file system is organized by a table of columns and rows, wherein filtering the data from the file system comprises:

 determining, based on the one or more policies associated with the user credentials, that one or more columns of data in the table are restricted; and
 masking at least a portion of one or more columns of data.

13. The computer program product of claim 12, wherein the masking comprises encrypting or tokenizing the data.

14. The computer program product of claim 12, wherein the masking masks some data in a column and leaves some data unmasked.

15. The computer program product of claim 10, wherein the data from the file system is organized by a table of columns and rows, wherein filtering the data from the file system comprises:

 determining, based on the one or more policies associated with the user credentials, that one or more rows of data in the table are restricted; and
 masking the one or more rows of data.

16. The computer program product of claim 15, wherein the masking comprises not returning the one or more masked rows to the compute process.

17. The computer program product of claim 10, wherein the user credentials comprise a user id and a group id, and wherein the user id is associated with a first collection of policies, and wherein the group id is associated with a second collection of policies.

18. The computer program product of claim 10, wherein the operations further comprise:
prior to the storing, receiving user input at a user interface associated with the policy enforcement system, the user input specifying one or more policies associated with the user credentials.

- 5 19. A system comprising:
a client device; and
a file system comprising a master node and a plurality of compute nodes,
wherein the master node is configured to receive from the client device a request for data from the file system and divide a workload from the request into a plurality of tasks and
10 assign the tasks to the plurality of compute nodes; and
wherein each compute node includes a processor and a computer-readable medium coupled to the processor and having instructions stored thereon, which, when executed by the processor, cause the processor operate a compute process and a policy enforcement system, including causing the processor to perform operations comprising
15 receiving at the compute process a first task from the master node;
maintaining, by the policy enforcement system a plurality of policies and data associating a plurality of user credentials with the plurality of policies;
receiving, in the policy enforcement system from the compute process, a request for data from a file system in the compute node for use in the first task, the request
20 further comprising user credentials;
sending the request for data from the policy enforcement system to the file system;
receiving the data from the file system in the policy enforcement system;
selecting from the plurality of policies, based on the received user credentials
25 and the data associating the plurality of user credentials with the plurality of policies, one or more policies that correspond to the received user credentials;
filtering, by the policy enforcement system, the data from the file system based on the one or more policies;
sending the filtered data from the policy enforcement system to the compute
30 process; and

returning a result from the compute process in the compute node to the master node.

20. The system of claim 19, wherein the data from the file system is organized by a table of columns and rows, wherein filtering the data from the file system comprises:

- 5 determining, based on the one or more policies associated with the user credentials, that one or more columns or rows of data in the table are restricted; and
 masking at least a portion of one or more columns or rows of data.

21. The system of claim 20, wherein the masking comprises encrypting or tokenizing data in one or more columns.

- 10 22. The system of claim 20, wherein the masking comprises not returning one or more rows to the compute process.

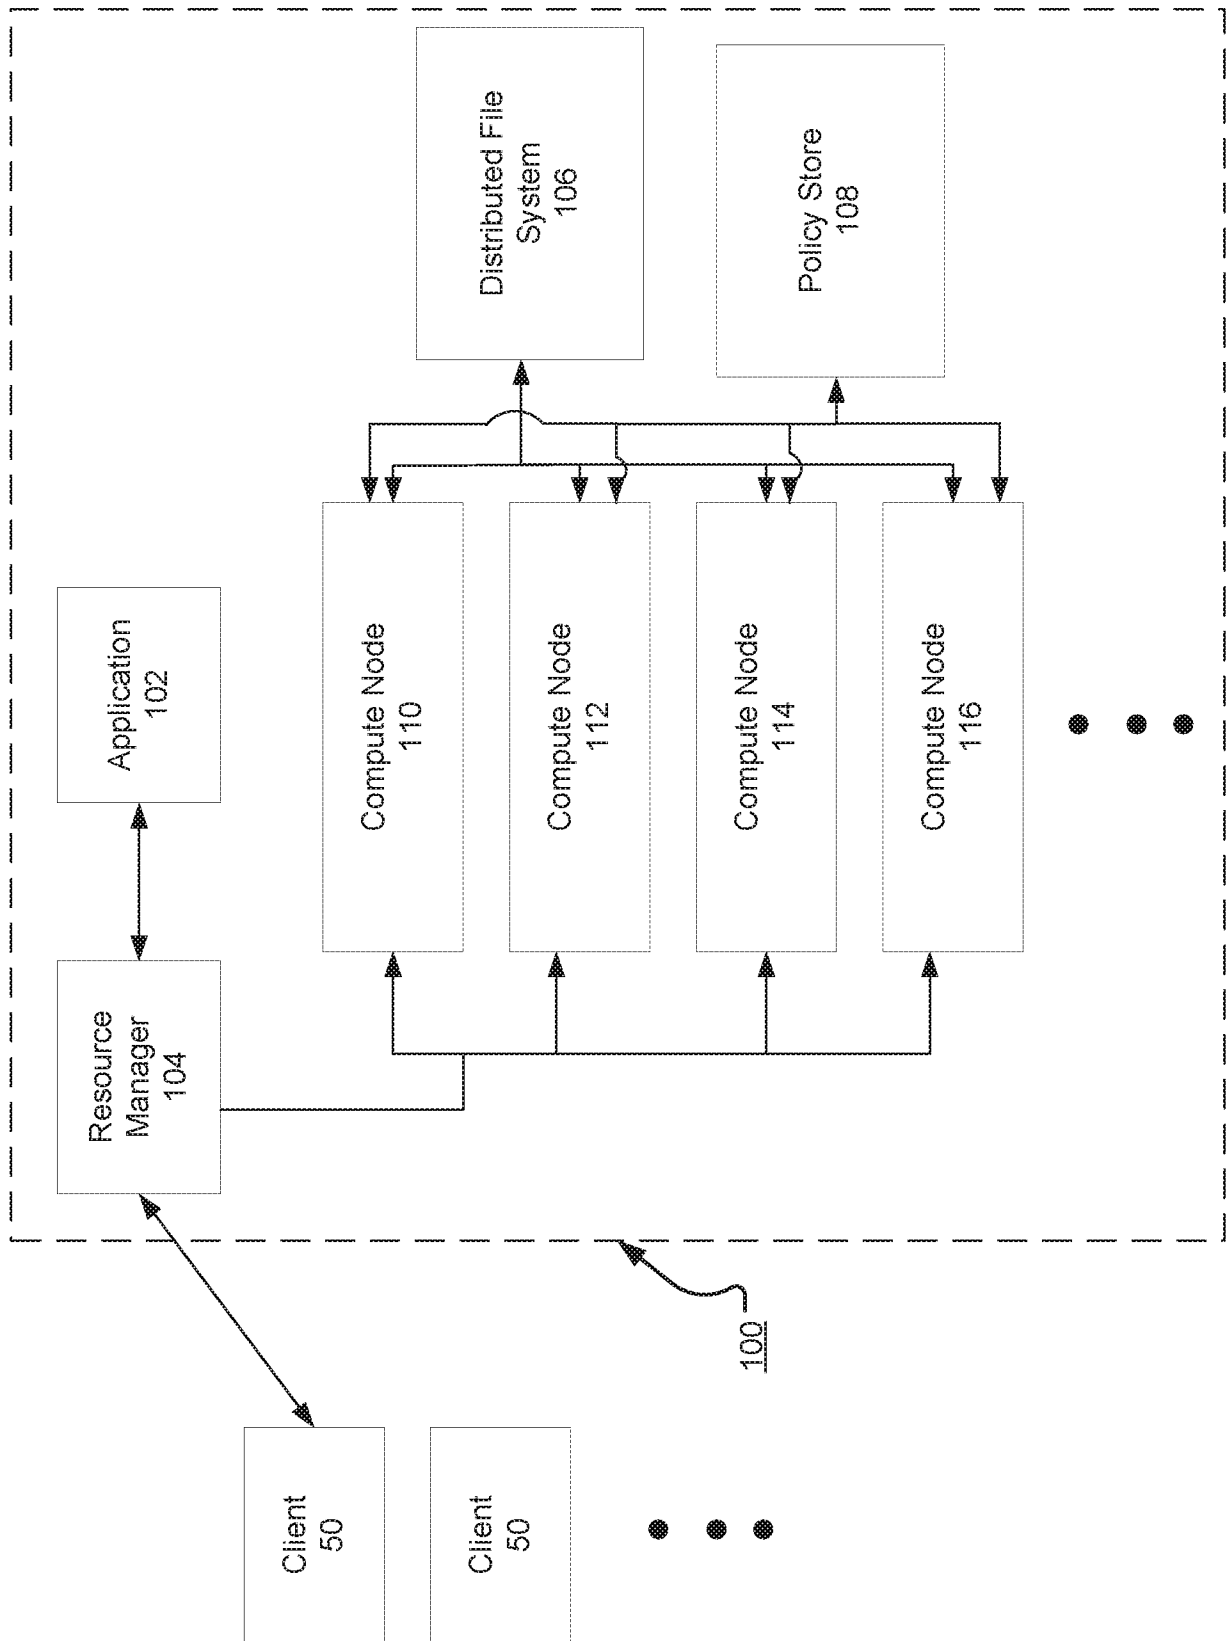


FIG. 1

3/3

300

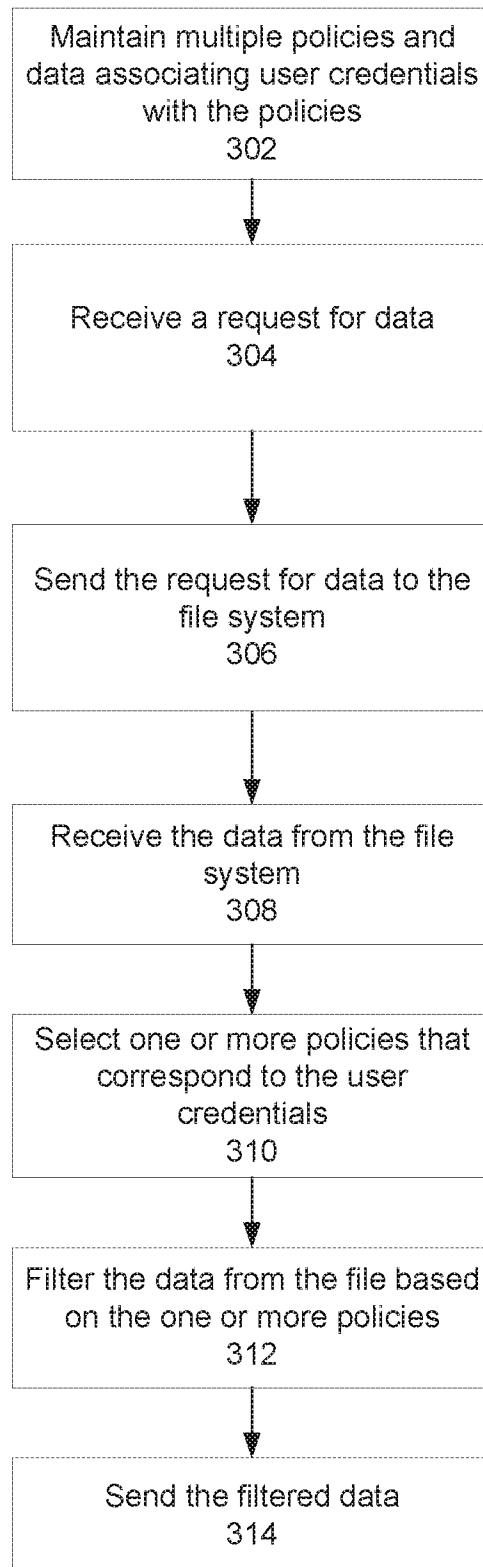


FIG. 3

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2016/065853**A. CLASSIFICATION OF SUBJECT MATTER****G06F 21/33(2013.01)i, G06F 21/62(2013.01)i, H04L 29/06(2006.01)i**

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 21/33; G06F 15/16; H04L 29/08; G06F 17/00; H04L 29/06; G06F 21/00; G06F 9/50; G06F 21/62

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

Korean utility models and applications for utility models
Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: policy enforcement, user credentials, selecting, masking, filtering, tokenizing

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 8677447 B1 (NIR ZUK et al.) 18 March 2014 See column 2, lines 61-62; column 4, lines 16-48; column 5, lines 19-21; claims 1, 15; and figures 1, 4.	1-22
Y	US 9071594 B2 (SALESFORCE.COM, INC.) 30 June 2015 See column 7, lines 18-20; column 12, lines 49-53; claim 9; and figure 1.	1-22
Y	US 8181221 B2 (S. A. VETHA MANICKAM et al.) 15 May 2012 See column 5, lines 28-34; and figure 2.	3-7, 12-16, 20-22
A	US 2011-0209194 A1 (KEVIN J. KENNEDY) 25 August 2011 See paragraphs [0098]-[0112]; and figure 1.	1-22
A	US 2012-0198467 A1 (DAVID B. JACKSON) 02 August 2012 See paragraphs [0062]-[0069]; and figure 3.	1-22

 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

23 March 2017 (23.03.2017)

Date of mailing of the international search report

24 March 2017 (24.03.2017)

Name and mailing address of the ISA/KR


International Application Division
Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

LEE, Seoung Young

Telephone No. +82-42-481-3535



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2016/065853

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 8677447 B1	18/03/2014	None	
US 9071594 B2	30/06/2015	US 2006-0075475 A1 US 2010-0192204 A1 US 2012-0096533 A1 US 2012-0096534 A1 US 2013-0014211 A1 US 2013-0014230 A1 US 2013-0247139 A1 US 2013-0247155 A1 US 2015-0304308 A1 US 2017-0006041 A1 US 7721328 B2 US 8108919 B2 US 8595802 B2 US 8667558 B2 US 8707411 B2 US 8707412 B2 US 8949963 B2 US 9450946 B2	06/04/2006 29/07/2010 19/04/2012 19/04/2012 10/01/2013 10/01/2013 19/09/2013 19/09/2013 22/10/2015 05/01/2017 18/05/2010 31/01/2012 26/11/2013 04/03/2014 22/04/2014 22/04/2014 03/02/2015 20/09/2016
US 8181221 B2	15/05/2012	US 2009-0049512 A1	19/02/2009
US 2011-0209194 A1	25/08/2011	DE 112011100626 T5 GB 2507941 A US 2011-0209193 A1 US 2011-0209195 A1 US 2011-0209196 A1 US 8434128 B2 US 8607325 B2 US 9215236 B2 WO 2011-103385 A1	24/01/2013 21/05/2014 25/08/2011 25/08/2011 25/08/2011 30/04/2013 10/12/2013 15/12/2015 25/08/2011
US 2012-0198467 A1	02/08/2012	US 2008-0216081 A1 US 2013-0346995 A1 US 8108869 B2 US 8464266 B2 US 9298514 B2 WO 2006-098725 A2 WO 2006-098725 A3	04/09/2008 26/12/2013 31/01/2012 11/06/2013 29/03/2016 21/09/2006 13/12/2007