



(51) International Patent Classification:

H04L 29/06 (2006.01) *G06F 21/62* (2013.01)

(21) International Application Number:

PCT/US2017/068680

(22) International Filing Date:

28 December 2017 (28.12.2017)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

15/394,148 29 December 2016 (29.12.2016) US

(71) Applicant: **T-MOBILE USA, INC.** [US/US]; 12920 SE 38th Street, Bellevue, Washington 98006-1350 (US).(72) Inventors: **PEPPE, Brett**; c/o T-Mobile USA, Inc., 12920 SE 38th Street, Bellevue, Washington 98006-1350 (US).
REITH, Greg; c/o T-Mobile USA, Inc., 12920 SE 38th Street, Bellevue, Washington 98006-1350 (US).(74) Agent: **CHEN, Elliott Y.**; Han Santos, PLLC, 500 Union Street, Suite 800, Seattle, Washington 98101 (US).(81) Designated States (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,

(54) Title: SECURE DATA CORRIDORS




FIG. 1

(57) **Abstract:** A system and method of providing a secure data corridor are provided. A request is received from a subject for a data feed comprising one or more data elements. A use-case is identified and a security level is determined for the data feed. For each data element of the data feed, a security level and one or more security controls that are mapped to the corresponding data element, are determined. A data sensitivity rating is assigned to the use-case. Upon determining that a clearance of the subject is at or above the data sensitivity rating of the use-case, the subject is granted privilege to the data feed via the secure data corridor.



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))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

SECURE DATA CORRIDORS

RELATED APPLICATIONS

[0001] This application is a Continuation in Part of U.S. Application No. 15/276,567 filed on September 26, 2016, which claims benefit under 35 USC §119 of U.S. Provisional Patent Application No. 62/233,153 filed on September 25, 2015, and titled “Distributed Big Data Security Architecture,” which are herein incorporated by reference in their entirety.

BACKGROUND

[0002] Traditional information technology (IT) architecture has relied on centralized computing structures, whereby all or most of the processing and computing is performed on a central server. Centralized computing can enable the deployment of central server computing resources, administration, and management. Typically, computational processing involves the extraction, transformation and loading of a resultant data feeds by the centralized computing structure. Security for data feeds is typically provided by securing the computational nodes

[0003] However, cost and performance inefficiencies of centralized computing structures have steered IT architectures towards decentralized computing structures, whereby data feeds are extracted, uploaded to a decentralized computing platform and then transformed. Doing so provides cost and performance efficiencies since the infrastructure to perform the computational processing can be outsourced to an entity that specializes in such.

[0004] As decentralized computing develops, both the velocity and diversity of data can cause some security measures to have an impact on latency, scalability, and recovery performance. Security risk has evolved beyond solely the computational node and now

includes network connections between those nodes. Thus, security can be viewed as a service constraint.

[0005] Security permissions are usually mapped to an object (system, file, database, etc.) and checked on subject request for operation via an Access Control Lists (ACLs). These ACL's are intended to support "authentication," "authorization," and "auditing." However, in most modern enterprise systems these capabilities have become separated. For example, authentication is deferred to an "authoritative master;" in authorization, the permissions depend on the "objects in scope;" and in auditing, the logging specifics depend on the "events in scope."

[0006] Because these security tasks have become functionally separated, the cost to keep them reasonably synchronous in a large enterprise often becomes so high that critical permissions maintenance and audit may not be performed. For example, the basic security on a file system may be statically mapped to folders and files. One either has the necessary permissions to list, read, write (accept) or one does not (deny). The subject is authenticated (e.g., mapped to a known principal) and authorized. Auditing occurs at the object level. However, the file system may be blind to the data that participates. One strategy to deal with this fundamental weakness is a multilevel classifications system (like Bell-LaPadula). However, such system still depends on static specifications.

[0007] Data may be classified by business sensitivity and risk controls, which in turn are factored into the permissions granted to an object. In this regard, subjects derive permissions to data based on their assigned clearance to the underlying secure data corridor. The data layer is now sensitive to the import, transformation, viewing and export of such data. A more modern form of the ACL would integrate risk and auditing. Such integration implies that metadata associated with risk controls and data sensitivity is stored

in the security hive. As such, these decisions no longer reside in the application layer, but are verified and enforced at the system level.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] The detailed description is described with reference to the accompanying figures, in which the left-most digit(s) of a reference number identifies the figure in which the reference number first appears. The use of the same reference numbers in different figures indicates similar or identical items.

[0009] FIG. 1 illustrates a schematic view of an architecture that facilitates establishing a secure session between a subject and a secure data container via a secure data corridor.

[0010] FIG. 2A illustrates a block diagram of an informational graph that describes the relationship between data objects, data elements, and data feeds within a use-case and/or secure data container, consistent with an exemplary embodiment.

[0011] FIG. 2B illustrates a block diagram of a diagraph that represents subject permissions to specific control parameters and a specific data feed, consistent with an exemplary embodiment.

[0012] FIG. 3A illustrates a block diagram of an exemplary permission lattice associated with granting access permissions to a subject.

[0013] FIG. 3B illustrates a data sensitivity lattice that associates control parameters with data feeds.

[0014] FIG. 3C illustrates a block diagram of an exemplary permission lattice having several sensitivity layers.

[0015] FIG. 3D illustrates a data sensitivity lattice that associates control parameters with data feeds. The control parameters are tied to data feeds.

[0016] FIG. 4 illustrates the creation of a third data feed from two existing data feeds.

[0017] FIG. 5 illustrates data movement within a secure data channel as a digraph.

[0018] FIG. 6 illustrates a block diagram of an informational graph that includes keys between data objects, data elements, and data feeds within a use-case and/or secure data container.

[0019] FIG. 7 illustrates a block diagram of a data digraph with different controls, consistent with an exemplary embodiment.

[0020] FIG. 8 illustrates a block diagram of a resource management platform that can establish a communicative connection between a subject and a secure data container via a secure data corridor.

[0021] FIG. 9 illustrates an example processes for establishing a session between a requestor subject and a secure data container using a secure data corridor.

DETAILED DESCRIPTION

Overview

[0022] In the following detailed description, numerous specific details are set forth by way of examples in order to provide a thorough understanding of the relevant teachings. However, it should be apparent to those skilled in the art that the present teachings may be practiced without such details. In other instances, well known methods, procedures, components, and/or circuitry have been described at a relatively high-level, without detail, in order to avoid unnecessarily obscuring aspects of the present teachings.

[0023] This disclosure generally relates to data security, and more particularly, to access control in a distributed computing environment. Techniques are described for defining security measures of a “secure data corridor” that enables data feeds to transmit from an ingress point to an egress point, while maintaining a desired level of security protection.

[0024] The term “secure data corridor” as used herein describes a communication channel used to transmit data feeds from a source location (i.e., ingress point) to a destination location (i.e., egress point) within a distributed computing environment. The techniques described herein further associate a secure data corridor with security measures that mandate access and/or confidentiality protection and/or data integrity protection that is commensurate with the nature of the data feed. For example, a data feed that relates to proprietary vendor data that is preferentially withheld from public dissemination may be associated with a less sensitive security measure when compared with a data feed that relates to personal consumer information that should be withheld from public dissemination. In various examples, security measures may be defined, and enforced by explicit, read, write, import and export provisions expressly associated with data feeds. Further, any other regime that is definable by a developer or an administrator, may be used to define security measures.

[0025] In various examples, data items may be classified by use-case, and risk controls may be applied based on the sensitivity of data that is maintained within each use-case. Each use-case may be classified and assigned to a “security label,” which defines the minimum security clearance required for a subject to access any data feed associated therewith. The term “data sensitivity rating” is used herein to describe a rating, or score, given to a data feed that quantifies the security measure that mandates confidentiality protection and data integrity protection. Using the above mentioned example, a data feed that is associated with personal consumer information to be withheld from public dissemination is likely to be assigned a more sensitive, that is higher, data sensitivity rating, relative to a data feed that relates to proprietary vendor data that is preferentially withheld from public dissemination.

[0026] Each data sensitivity rating may prompt one or more control parameters that protect the confidentiality and integrity of the data feed. The one or more control parameters may include principles such as no read-up, no write-down, tranquility, explicit read-in, and explicit write-out. The ‘no read-up’ principle refers to a condition whereby a subject at a given security level may not read a data feed at a high security level. A ‘no write-down’ principle refers to a condition whereby a subject at a given security level may not write to a data feed to a less sensitive security level. A ‘tranquility’ level refers to a condition where an object that is once classified to a level, remains at that level. The term ‘explicit read-in’ refers to a condition where landing permissions are explicitly granted (i.e., they are not default). In contrast, the term ‘explicit write-out’ refers to a condition where export permissions are explicitly granted (i.e., they are not default). The term “subject” as used herein describes a user, a non-user entity, a secure data container, or any other data structure requesting access permissions to data feeds within a secure data container, and via a secure data corridor.

[0027] In various examples, the classification of a “data sensitivity rating” may be associated directly with data elements and set by predetermined rules. In some examples, a default data sensitivity rating may be assigned to a secure data corridor. In other examples, the data sensitivity rating may be based on access privileges associated with the subject. Further, the data sensitivity rating of the secure data corridor may be dynamically set by interrogating the nature of the data feed that is transmitting through the secure data corridor. In some examples, the data sensitivity rating may be set in substantially real-time.

[0028] The term “data feed” as used herein represents groups of potentially streamed data elements used in an import or export operation. A data feed may include one or more data elements, which in turn are associated with data objects. A data feed may be assigned

a default control parameter that applies to all data elements in the feed. A data feed may have default controls and acquired controls. Acquired controls may supersede default controls. The ability to set default controls on a data feed may facilitate loading data into the cluster that requires transformation after loading, but still requires protection. A data feed may have zero to many associated data elements.

[0029] The term “data element” as used herein describes the physical and logical characteristics of an associated data type. It is a datum or a collection thereof. Further, a data element may define a data type as a string, integer, float, or object. Data elements may be associated with one to several data feeds as driven by the needs of a business use case. The data elements map to data feeds. Data objects map to a data element. Data objects, store the data value at a particular storage location.

[0030] In some examples, the control parameters are assigned to the data objects. The data elements that are associated with the data objects inherit control parameters from the data objects. Further, the data feed that is associated with the data elements inherit the control parameters from the data elements, and ultimately from the data objects. The level of the control parameters is dependent on the use-case. Thus, in one embodiment, the security model may be customized for different use-cases. In other words, a use-case may be associated with a specific security label, and all data feeds and objects associated with the use-case may be assigned a data sensitivity rating, specific to the security label. Thus, a subject with access to a use-case may have access to all data and services that are assigned to the use-case, provided that the subject retains the requisite access permissions.

[0031] Each use-case, may have its own namespace, which can be enforced by a resource management platform. Privileges and permission are attached to the namespace. The term ‘namespace’ is used herein to describe a set of symbols used to identify objects that are used to organize those objects (which may be of various kinds), so that those

objects may be referred to by that identifying set of symbols. In this way, the identifying set of symbols is the name of that object. The act of informing a system that the name of the object corresponds to that object is called “binding.” Because names may be desirable to use in multiple use-cases, a namespace itself may have a name to identify that namespace, and used to qualify names for that use-case. For example, consider there are two use-cases, called Use-Case1 and Use-Case2. Say that both use-cases have an object called “MainObject.” Use of names spaces can distinguish MainObject for Use-Case1 and for Use-Case2 by referring to the former as Use-Case1.MainObject and Use-Case2.MainObject. In this way, object names need not be unique across all use-cases, and namespace names allow a logical way to group objects for that use-case.

[0032] In a non-limiting example, a namespace may include identifiers within the context of file systems, programming languages, and computer networks and distributed systems. Further, a namespace may also include a secure data container having one or more data feeds.

[0033] In various examples, a subject may be added to an Access Control List (ACL) on the authentication service principals in order to gain access to a secure data corridor. Generally, a subject assigned to a given data sensitivity rating may access all data within that same classification.

[0034] This disclosure further describes techniques that illustrate multiple layers of authentication, authorization, and auditing of separate platform levels. As used herein, authentication is the verification and validation of credentials presented by a subject to an authoritative master (e.g., an authentication server) that either accepts or rejects the subject. Acceptance of the credentials assigns an identity to the subject. Authorization is the granting of system resources, privileges and/or permissions to a subject. Auditing is the collection, maintenance, and reporting of Access Control Management (ACM) records.

[0035] By virtue of the multiple layers of authentication, a single point of attack (SPOA) in the security trust chain can be avoided as any successful attack would need to be of a synchronized multi-point nature. Elements of a defense may include controlling access to the resource management platform by an authentication server entry. Authentication may be performed by an Active Directory (represented later herein as an authentication server), whereby a subject may provide a set of credentials that confirm their identity. In a non-limiting example, the credentials may include a username, password, or other biometric identifier. Further, authorization to access a secure data container, via a secure data corridor is dependent on the assigned access privileges and the data sensitivity rating of an associated use-case and data feed within the secure data container.

[0036] A secure data corridor (SDC) may comprise a secure data container (having one or more data feeds) and trust chain vertices. A trust chain is created by a directed graph of vertices between SDC's. Each vertex may be a physical computational node, but more commonly would be a software module. The traversal of a given trust chain not only verifies the integrity of the chain, but also enables an audit mechanism. The secure data container is, in essence, embedded within the data corridor. The trust chain vertices define logical input and output pathways to the embedded data container. In other words, a secure data corridor connotes a transport mechanism and in its basic pathway operates as a migration pathway of physical data.

[0037] In various examples, the secure data corridor may correspond to a link layer, a network layer, a transport layer, a session layer, or an application layer of a network protocol. For example, for the network, transport and session layers, the corresponding network protocol may be transmission control protocol (TCP) and internet protocol (IP). For the application layer, the network protocol may be hypertext transfer protocol (HTTP).

[0038] In a non-limiting example, a secure data corridor may include one input, one output, and one container. In a second non-limiting example, a secure data corridor may include one container with one input and no outputs. In this example, the SDC can act as a secure sink where data of an unknown sensitivity rating can be staged prior to being class-rated into other containers, and other secure data corridors. In a third non-limiting example, the secure data corridor may include a secure data container with multiple outputs and no input. In this example, the secure data corridor input may be disconnected once data is migrated. The multiple outputs, that is to say the vertices, may connect to several other secure data corridors.

[0039] Each SDC may include a base set of applicable security controls having a security label. The data elements used in an SDC are mapped to these controls. Different controls may have different confidentiality, integrity, and auditing requirements. A given data item may be tagged by as many controls as appropriate.

[0040] All Data elements (DE) used in a SDC may be registered in a SDC data dictionary (DD). Each data element can participate in one or more Data Feeds (DF), which represents groups of data elements used in an IMPORT or EXPORT operation. Each use-case has one or more data feeds and represents the complete set of data used to solve a particular problem domain. An SDC may include at least one use-case. In one embodiment, an SDC includes the following information: which use-case(s) are active for a data feed; what data feed(s) participate in each use case; the active data dictionary; which security labels are active and the data elements mapped to the security labels; which data elements are associated with a data feed; and which subject used, viewed or modified which use-case, data feed, and/or data element.

[0041] In one embodiment, there is a data linkage map set (DLMS), which is a table used by the secure data channel verifier to determine permissions on data objects via their

security label. For example, the DLMS defines the valid data triples (e.g., data feed, data element, and security control) for an authoritative master (e.g., authentication server) within the secure data channel as a linear map-set. A given data element may be tagged by as many controls as necessary. For example, controls may relate to the Payment Card Industry Data Security Standard (PCI DSS), customer proprietary network information (CPNI), Sarbanes-Oxley (SOX), etc.

[0042] In one embodiment, the DLMS includes a date/time stamp that captures the data/time of when a write operation created the entry of a data object, as illustrated in table 1 below.

Data Feed	Data Element	Control	Date/Time
DF ₁	DE ₁	C ₁	09/26/2014; 6:08PM
DF ₂		C ₂	09/26/2014; 6:18PM

Table 1: Example Data Linkage Map Set

[0043] The DLMS may be used as a reference to determine whether a requesting subject has clearance to a data object by having clearance to the data feed, data element(s), and controls within. In the example of Table 1, the security label for the first use-case is {DF₁, DE₁, C₁}. Accordingly, the requesting subject must have clearance to {DF₁, DE₁, C₁} for a data object.

[0044] In one embodiment, there is a subject linkage map set (SLMS), which is a table used by the secure data channel verifier to determine subject clearances to data objects. It defines the valid subject triples for the authoritative master within the secure data channel. In this regard, Table 2 below provides three different use-cases.

Data Feed	Subject	Control	Date/Time
DF ₃	S ₃	C ₃	09/26/2014; 6:08PM
DF ₄	S ₄		09/26/2014; 6:18PM
	S ₅	C ₅	09/26/2014; 6:39PM

Table 2: Example Subject Linkage Map Set

[0045] The empty fields in Table 2 above indicate that there are no subject linkages that have been defined within the secure data channel. In the example of Table 2, for the first use-case, the subject clearance level need be at or above the security label {DF₃, DE₃, C₃}. Accordingly, the DLMS permits both the DF₃ and C₃ to DE₃. For the second use-case, the subject clearance level need be at or above the security label {DF₄, DE₄, null}. Accordingly, the DLMS permits both the DF₄ and C₄ to DE₄. For the third use-case, the subject clearance level need be at or above the security label {null, DE₅, C₅}. Accordingly, the DLMS permits all DF and C₅ to DE under all DF.

[0046] The term “data container,” as used herein, describes a data structure that includes data items associated with the operation of a particular application. For example, a data container of a particular application and its dependencies may include an operating system, one or more data item libraries, configuration files, application binaries, and other parts of a technology stack used for execution of the particular application. Unlike a traditional virtual machine, a data container may hold its application and any dependencies. However, a same operating system kernel and libraries can be shared between multiple containers running in a host environment.

[0047] The term “attribute based access control” (ABAC) relates to a security model that maps attributes against access controls. It is an access control method, where a subject request to perform operations on objects are granted based on assigned attributes of the subject, assigned attributes of the object, environment conditions, and a set of policies that re specified in terms of those attributes and conditions. In ABAC, both the subject and object agree synchronously on the mediated attribute sets for access control. If the attributes are not consistently defined, then a formal ABAC is not feasible. In one aspect, ABAC can avoid a static (i.e., subject/object operation) pair lookup.

[0048] The principles described herein may be applicable to any distributed computing environment. Historical protection techniques include network (e.g., IP address), firewall (e.g., Ports, Services), platform (e.g., Folder, Share) and data (e.g., Schema, Table) segmentation.

[0049] The term “techniques,” as used herein, may refer to system(s), method(s), computer-readable instructions, module(s), algorithms, hardware logic, and/or operation(s) as permitted by the context described above and throughout the document.

Example Architecture

[0050] FIG. 1 illustrates a schematic view of an architecture 100 that facilitates establishing a secure session 108 between a subject 102 and a secure data container via a secure data corridor. In the example of FIG. 1, the subject 102 may transmit an access request 104 to an authentication server 106. In some examples, the authentication server 106 may perform authentication by an Active Directory whereby the subject 102 may provide a set of credentials that confirm their identity. In a non-limiting example, the credentials may include a username, password, or other biometric identifier associated with the subject. In various examples, the subject 102 may correspond to a user, a non-user entity, a secure data corridor, or any other data structure that is capable of requesting access permissions to data feeds within a secure data container.

[0051] The subject 102 may interface with the authentication server 106 via any sort of electronic device, such as a cellular phone, a smart phone, a tablet computer, an electronic reader, a media player, a gaming device, a personal computer (PC), a laptop computer, etc. Further, the electronic device of the subject 102 may include a subscriber identity module (SIM), such as an eSIM, to identify the electronic device of the subject 102 to a telecommunication service provider network (also referred to herein as “telecommunication network”).

[0052] In the illustrated example, in response to authenticating the subject 102, a session 106 is initiated via the authentication server 106 between the subject 102, through the staging server(s) 109. The decentralized cluster security policy balances the need to provide scalable aggregation while still maintaining isolation based on security classifications. Isolation may be based on authorization and via the staging server(s) 109.

[0053] Further, the staging server(s) 109 are communicatively coupled to a resource management platform 110. The resource management platform 110 may establish a communicative connection between a subject 102 and one or more secure data container(s) 112(1), 112(2), or 112(N), via a secure data corridor 114(1), 114(2), and 114(N), respectively. In various examples, the secure data corridor 114(1), 114(2), and 114(N) correspond to a link layer, session layer, a network layer, a transport layer, or an application layer of a transmission control protocol (TCP) and internet protocol (IP). In a non-limiting example, a secure data corridor may include one input, one output, and one container. In a second non-limiting example, a secure data corridor may include one container with one input and no outputs. The one or more secure data containers 112(1), 112(2), or 112(N) may be stored within a decentralized computing platform 116. The one or more secure data containers 112(1), 112(2), or 112(N) may include a data structure that includes data items associated with the operation of a particular application. For example, a data container of a particular application and its dependencies may include an operating system, one or more data item libraries, configuration files, application binaries, and other parts of a technology stack used for execution of the particular application.

[0054] The staging server(s) 109 and the decentralized computing platform 116 may operate on one or more distributed computing resource(s), respectively. The one or more distributed computing resource(s) may include one or more computing device(s) 118(1) to 118(N) and 120(1) to 120(N), respectively, which operate in a cluster or other

configuration to share resources, balance load, increase performance, provide fail-over support or redundancy, or for other purposes. The one or more computing device(s) 118 and 120 may include one or more interfaces that enable communications with electronic device(s), such as an electronic device associated with the subject 102, via one or more network(s) 122.

[0055] In the illustrated example, the one or more network(s) 122 may include public networks such as the Internet, private networks such as an institutional and/or personal intranet, or some combination of private and public networks. The one or more network(s) 122 can also include any type of wired and/or wireless network, including but not limited to local area network (LANs), wide area networks (WANs), satellite networks, cable networks, Wi-Fi networks, Wi-Max networks, mobile communications networks (e.g. 3G, 4G, and so forth), or any combination thereof.

Example Data Feeds

[0056] FIG. 2A illustrates a block diagram of an informational graph 200 that describes the relationship between data objects, data elements, and data feeds within a use-case and/or secure data container, consistent with an exemplary embodiment. In the illustrated example, control parameters 202, 204, and 206 may be assigned to data objects 208, 210, and 212, respectively. The control parameters may include principles such as no read-up, no write-down, tranquility, explicit read-in, and explicit write-out. The ‘no read-up’ principle refers to a condition whereby a subject at a given security level may not read a data feed at a high security level. A ‘no write-down’ principle refers to a condition whereby a subject at a given security level may not write to a data feed to a less sensitive security level.

[0057] In the illustrated example of FIG. 2A, the data objects 208, 210, and 212 inherit the respective control parameters 202, 204, and 206 from the data feeds 220 and 222, or

the mapping data elements 214, 215, 216, and 218 that define them. A data element may define the physical and logical characteristics of an associated data type. For example, a data element may define a data type as a string, integer, float, or object. Further, a data element may further define a risk rating or a generic sensitivity rating (collectively referred to herein as a security level) of the associated data. In some examples, the security level may change over time. The security level may also change to reflect a directed or explicit sensitivity policy. Further, a data element may define a data type as a string, integer, float, or object. The data objects store the data value at a particular storage location. The data feeds 220 and 222 that are associated with the data elements 214, 216, and 218 inherit the control parameters from the data elements 214, 216, and 218, and ultimately from the data objects 208, 210, and 212. Each data element is not limited to be associated with one data object. For example, data element 216 may be associated not only with data object 210, but also with one or more additional data objects (not shown). Thus, several data objects may be associated with a single data element. Further, one data object (e.g., 210) may be associated with several data elements (e.g., 214 and 216).

[0058] FIG. 2B illustrates a block diagram of a diagram that represents subject permissions to specific control parameters 224 and a specific data feed 226, consistent with an exemplary embodiment. A subject 228 is mapped to controls 224 and data feeds 226. In the illustrated example, permissions associated with the subject 228 are bound to specific control parameters 224 and stored in a list of triples 230 (i.e., subject, control, and data feed) that tie the subject 228 to specific controls, and correspondingly, the subject to the data feeds that are tagged by the specific controls. In some examples, the subject 228 may be granted access permissions to data feeds, whereby the access is governed by the specific control parameters specified in the list of triples 230. The data feeds 226 and specific control parameters 224 may be explicitly assigned. Alternatively, or additionally,

the subject 228 may be granted access to all data elements that are qualified by a specific control parameter.

[0059] Further, the resource management platform may interrogate the list of triples 230 associated with the subject to determine privileged data feeds as well as the controls that determine the access rights to those data feeds. The control parameters 224 may be expressed as a data sensitivity lattice of varying levels, as further described in FIGS. 3A and 3B.

Example Permission Lattice

[0060] FIG. 3A illustrates a block diagram of an exemplary permission lattice 300 associated with granting access permissions to a subject. In various examples, a subject may correspond to a user, a non-user entity, a secure data container, or any other data structure requesting access privileges. Particularly, FIG. 3A illustrates a block diagram of a data sensitivity lattice 300. In some examples, the data sensitivity lattice 300 may include a first data sensitivity layer 302 and a second data sensitivity layer 304. In the illustrated example, the first data sensitivity layer 302 may be associated with a first set of data items that mandate low security measures for data confidentiality and data integrity protection. In a non-limiting example, the first set of data items may correspond to proprietary vendor data that is preferentially withheld from public distribution. Comparatively, the second data sensitivity layer 304 may be associated with a second set of data items that mandate more data security measures relative to the first data sensitivity layer 302, for data confidentiality and data integrity protection. In a non-limiting example, the second set of data items may correspond to personal consumer information that must be withheld from public distribution. For illustrative purposes only, FIGS. 3A and 3B present a two-layer data sensitivity lattice 300. It is noteworthy that the data sensitivity lattice 300 may be defined to include any number of layers favored by a developer or

administrator, as discussed in more detail later in the context of the description of FIGS. 3C and 3D.

[0061] FIG. 3A illustrates the relative flow of read and write permissions between the first data sensitivity layer 302 and the second data sensitivity layer 304. It is noteworthy, that although FIG. 3A describes read and write permissions, permissions may further include import permissions, export permissions, or any other permission regime definable by a developer or administrator. In the illustrated example, a subject with access permission to the first data sensitivity layer 302 may write data items to the second data sensitivity layer 304, however the subject may not read data items within the second data sensitivity layer 304. In other words, the subject may add (i.e., write) data items to the more sensitive, second data sensitivity layer 304, however the subject may not access, that is read, data items within the more sensitive, second data sensitivity layer 304.

[0062] FIG. 3A further illustrates that a subject with access permission to the second data sensitivity layer 304 may read data items within the first data sensitivity layer 302, however the subject may not write data items within the first data sensitivity layer 302. In other words, the subject may read data items within the less sensitive, first data sensitivity layer 302, however the subject may not add (i.e., write) data items to the less sensitive, first data sensitivity layer 302.

[0063] A subject with access permission to the first data sensitivity layer 302 may read data items and write data items within the first data sensitivity layer 302. Similarly, a subject with access permission to the second data sensitivity layer 304 may read data items and write data items within the second sensitivity layer 304. Note that the terms data sensitivity layer and data sensitivity rating may be used interchangeably to describe the data sensitivity of a data item. Further the term data item may be used interchangeably to describe data feed.

[0064] FIG. 3B illustrates a data sensitivity lattice 306 that associates control parameters with data feeds. The control parameters are tied to data feeds. The control parameters may include read, write, export, and import parameters. For example, consider a subject with access permission to the second data sensitivity layer 304. As illustrated in FIG. 3A, the subject may read data items within the first data sensitivity layer 302, but may not add (i.e., write) data items to the first data sensitivity layer 302. The write limitation ensures that the subject may avoid inadvertently having more sensitive data items from the second data sensitivity layer 304 being added to the less sensitive first data sensitivity layer 302. The data sensitivity lattice 306 of FIG. 3B reflects the same, except access permissions are presented from the perspective of the data feed, rather than the subject.

[0065] For example, FIG. 3A illustrates the second data sensitivity layer 304 as the most data sensitivity layer. Thus, the second data sensitivity layer 304 can only interact downward with the first data sensitivity layer 302. In FIG. 3B, the interaction of the second data sensitivity layer 304 with the first data sensitivity layer 302 is reflected in the ‘down’ column of the data sensitivity lattice 306. The ‘down’ column indicates a capability to read down (i.e., ‘1’ in the read – down cell), and an inability to write down (i.e., ‘0’ in the write – down cell).

[0066] Similarly, as shown FIG. 3A, a subject with access permissions to first data sensitivity layer 302 may not read data contained within the second data sensitivity layer 304, because the higher sensitivity of the second data sensitivity layer 304 represents more sensitively classified data. Using the previous example of proprietary vendor data and consumer personal data, a subject with permission access to proprietary vendor data (i.e., classified as first data sensitivity level layer 302) cannot read the consumer personal data (i.e., classified as the second data sensitivity level layer 304).

[0067] It should be noted that permissions are not limited to a first and second sensitivity layer. In various embodiments, there may be any number of sensitivity layers that can be controlled for a subject. In this regard, FIG. 3C illustrates a block diagram of an exemplary permission lattice 330 having several sensitivity layers. By way of example, FIG. 3C illustrates three sensitivity layers 332, 334 and 336, while it will be understood that other number of sensitivity layers are supported as well. Accordingly, the data sensitivity lattice 330 may include a first data sensitivity layer 332, a second data sensitivity layer 334, and a third data sensitivity layer 336. In the illustrated example, the first data sensitivity layer 332 may be associated with a first set of data items that mandate low security measures for data confidentiality and data integrity protection. Comparatively, the second data sensitivity layer 334 may be associated with a second set of data items that mandate more data security measures relative to the first data sensitivity layer 332, for data confidentiality and data integrity protection. The third sensitivity layer 336 may be associated with a third set of data items that mandate more data security measures relative to the second data sensitivity layer 334.

[0068] In the illustrated example of FIG. 3C, a subject with access permission to the first data sensitivity layer 332 may write data items to the second data sensitivity layer 334 and the third sensitivity layer 336. However, the subject may not read data items within the second data sensitivity layer 334 or the third data sensitivity layer 336. In other words, the subject may add (i.e., write) data items to the more sensitive, second data sensitivity layer 334 and third sensitivity layer 336, but the subject may not access, that is read, data items within the more sensitive, second data sensitivity layer 334 and third data sensitivity layer 336.

[0069] FIG. 3A further illustrates that a subject with access permission to the third data sensitivity layer 336 may read data items within the second data sensitivity layer 334

and the first data sensitivity layer 332. However, the subject may not write data items within the second data sensitivity layer 334 and the first data sensitivity layer 332. In other words, the subject may read data items within the less sensitive, first and second data sensitivity layers 332 and 334, but the subject may not add (i.e., write) data items to the less sensitive, first and second data sensitivity layers 332 and 334. Note that the terms data sensitivity layer and data sensitivity rating may be used interchangeably to describe the data sensitivity of a data item. Further the term data item may be used interchangeably to describe data feed.

[0070] FIG. 3D illustrates a data sensitivity lattice 340 that associates control parameters with data feeds. The control parameters are tied to data feeds. The control parameters may include read, write, export, and import parameters. The controls in the lattice 340 need not be limited to “0” and “1” entries. Rather, in various embodiments, for a “read” operation, a “0” may indicate a permit for the subject to read at the same level; a “1” may indicate a permit for the subject to read one level down; a “2” may indicate a permit for the subject to read 2 levels down; and an “n” may indicate a permit for the subject to read any level down. A blank (i.e., “null”) may explicitly deny a read or write, respectively.

[0071] Similarly, for a write operation, a “0” may indicate a permit for the subject to write at the same level; a “1” may indicate a permit for the subject to write one level up; a “2” may indicate a permit for the subject to write 2 levels up; and an “n” may indicate a permit for the subject to write to any level up.

[0072] Clearances are not explicitly assigned to subjects; rather, a subject acquires a clearance when permission to a control is granted. Each subject may have a singular clearance level for each granted control. If a subject is granted access to a control, then

the subjects is permitted all data elements (and objects) referenced tagged by that control and data feed, unless it is explicitly denied.

Example Creation of a New Data Feed

[0073] FIG. 4 illustrates the creation of a third data feed from two existing data feeds. For example, a subject 406 may have read access to several data feeds, represented by data feeds 402 and 404. The first data feed $C_{C(1)}$ includes data elements DE_{11} to DE_{1n} , each data element having a corresponding security label, sometimes referred to herein as a data sensitivity rating. Similarly, the second data feed $C_{E(1)}$ includes data elements DE_{21} to DE_{2n} , each with a corresponding security label. The subject 406 can create a new data feed 408 from the two existing data feeds 402 and 404 that includes the data elements of each data feed and its corresponding data sensitivity ratings. Accordingly, the subject 406 has read access to the first and second data feeds 402 and 404, and write access to the new third data feed 408.

[0074] In one aspect, data is classified by business use-case and applied risk controls (sensitivity). Subjects derive permissions to data objects based on their assigned clearance to the respective data corridor. This data layer is sensitive to the import, transformation, viewing and export of any data. Such architecture provides a unified platform for authorization and auditing in a business enterprise.

[0075] Applicants have developed the Secure Data Corridor (SDC) upon the realization that data ingestion, transformation, processing and extraction may not be adequately protected on large Big Data clusters using traditional Access Control Management (ACM) techniques. In a traditional approach, any failure in the contractual trust chain may make the entire cluster vulnerable. In one aspect, a security model (based on SDC's) is used that is secure by design via a Container Based Access Control (CBAC) technique with integrated Risk Controls, where: (i) containers are classified by data

sensitivity and define processing corridors; (ii) containers may securely host multiple data pools at different risk levels; (iii) containers may dynamically re-rate themselves based on their contents; (iv) data feeds imported into containers are controlled by a risk rating (e.g., corridor on write); (v) work products extracted from containers are controlled by a risk rating (e.g., corridor on read); and (vi) container data pools can participate in overlays.

[0076] If an attack vector compromises a single SDC, the attack cannot be extended to other SDC's by way of architecture. In various scenarios, a Big Data cluster may include as many SDC's as appropriate. The concepts discussed herein may be applicable to any distributed computing environment.

Example Operations on a Data Feed

[0077] Since data feeds are a digraph, then any transformation of one or more data feeds is also a digraph. In this regard, FIG. 5 illustrates data movement within a secure data channel as a digraph. Data is imported from an external source 502 into the secure data channel as a first data feed 504. The second data feed 506 is created by transforming the first data feed 504. The fourth data feed 510 is created by joining the second and third data feeds 506 and 508. Data is then exported to a subject 512 outside of the secure data channel.

Mapping of Data Objects to Data Elements

[0078] FIG. 6 illustrates a block diagram of an informational graph that includes keys between data objects, data elements, and data feeds within a use-case and/or secure data container. In various embodiments, the same data element may map to more than one data object within the same data feed, as illustrated by data element 608 with respect to data objects 610 and 612. The DLMS maintains a lookup map-set of its mapped data feeds (DF), data elements (DE), and controls (C). Each data object includes a reference key back to its owning data element. Accordingly, data objects that belong to a set of data

elements can be identified by searching the secure data channel for the referenced keyset. Every data feed is now an information graph of source, elements, and values. Such a graph can track the transformation history by adding a temporal dimension.

Auditing Via Integrated Risk Controls

[0079] In one embodiment, a data element can be interpreted as a signature rather than a value. Secure data channel data objects store values that are redefined by data elements. In various scenarios, auditing may occur at different levels for different use-cases. As used in this context, auditing is the collection, maintenance, and reporting of Access Control Management (ACM) records. For example, such auditing may be performed at a system level, database level, application level, or data feed level. In various embodiments, at least one of the following aspects may be determined by way of audit: (i) which use-cases are active, (ii) what data feed(s) participate in each use-case, (iii) the active data dictionary for the secure data channel, (iv) which controls are active and the data elements mapped thereto, (iv) and which subject looked at which use-case data feed, and/or data element.

[0080] FIG. 7 illustrates a block diagram of a data diagram with different controls, consistent with an exemplary embodiment. The data element 706 may be attached to a plurality of controls 702. Each control may have an associated security level 704. The data element 706 is a metadata definition linked to controls 2302 and data feeds 2306. The data element 706 is assigned a clearance label (L_k) for each control 702. The data element 706 can be assigned a temporal or scoped nature such that its assigned label is modified.

Ingesting and Extracting Data Feeds into a Container

[0081] In one embodiment, since each container has a default classification (which may be dynamically assigned), these containers are sensitive to what kind of data that they

can import. In this regard, data feeds are assigned a risk rating. One cannot write data at a higher risk level into a container with a lower risk rating. One can write data at a lower risk level into a container with a higher risk rating.

[0082] A secure corridor logs all inputs, processing, and outputs. For example, in Hadoop®, a YARN (a resource-management platform responsible for managing computing resources in clusters and using them for scheduling of users' applications) plugin may facilitate a “corridor on write,” where data feeds are imported into a container. The data feed may be accepted, denied, imported, or overlaid.

[0083] In one embodiment, data may be extracted from a container. Since each container has a default classification and data pools with risk ratings, each container is sensitive to what kind of data it can export. In one example, each data sub-pool in a container may be marked to indicate whether it is “releasable” or “not-releasable.” A data sub-pool marked at a lower level can export into a container rated at a higher level (i.e., write-up). However, such sub-pool cannot export to a container rated at a lower level (i.e., write-down). Any data sub-pool may participate in an overlay with another container as long as these rules are followed. For example, in Hadoop®, a YARN plugin may facilitate a “corridor on read,” where work products extracted from containers are controlled by a risk rating. The data feed may be accepted, denied, exported, or overlaid.

Example Resource Management Platform

[0084] FIG. 8 illustrates a block diagram of a resource management platform 802 that can establish a communicative connection between a subject and a secure data container having at least one data feed, via a secure data corridor. The secure data corridor may be configured with a security label that is substantially similar to a security label of an adjoining secure data container. In some examples, the subject may correspond to a user device, a non-user entity device, or another adjoining secure data container.

[0085] In the illustrated example, the resource management platform 802 may correspond to the resource management platform 110 of FIG. 1. Further, the resource management platform 802 may include one or more processor(s) 804 that are operably connected to memory 806. In at least one example, the one or more processor(s) 804 may be a central processing unit(s) (CPU), graphics processing unit(s) (GPU), both a CPU and GPU, or any other sort of processing unit(s). Each of the one or more processor(s) 804 may have numerous arithmetic logic units (ALUs) that perform arithmetic and logical operations as well as one or more control units (CUs) that extract instructions and stored content from the processor cache memory, and then executes these instructions by calling on the ALUs, as appropriate during program execution. The one or more processor(s) 804 may also be responsible for executing all computer applications stored in the memory, which can be associated with common types of volatile (RAM) and/or nonvolatile (ROM) memory.

[0086] In some examples, memory 806 may include a system memory, which may be volatile (such as RAM), non-volatile (such as ROM, flash memory, etc.,) or some combination of the two. The memory may also include additional data storage devices (removable and/or non-removable) such as magnetic disks, optical disks, tape, etc.

[0087] The memory 806 may further include non-transitory computer-readable media, such as volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information, such as computer readable instructions, data structures, program modules, or other data. System memory, removable storage and non-removable storage are all examples of non-transitory computer-readable media. Examples of non-transitory computer-readable media include, but are not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic tape,

magnetic disk storage or other magnetic storage devices, or any other non-transitory medium which can be used to store the desired information.

[0088] In the illustrated example, the memory 806 may include an operating system 808 and various applications, such as a security application 810 that may be used to provide a secure data corridor between a subject and a data feed of interest. Execution of the security application 810 by the processor(s) 802 configures the resource management platform 802 to perform various functions. In one embodiment, these functions may be controlled by different modules, such as a data sensitivity module 812, a secure data corridor module 814, an audit module 816, and a secure data dictionary module 818. The operating system 808 may be used to implement the data sensitivity module 812 and the secure data corridor module 814. The operating system 808 may be any operating system capable of managing computer hardware and software resources. The data sensitivity module 812, the secure data corridor module 814, the audit module 816, and the secure data dictionary module 818 may include routines, program instructions, objects, and/or data structures that perform particular tasks or implement particular abstract data types.

[0089] In the illustrated example, the data sensitivity module 812 may determine the data sensitivity rating to assign to a secure data corridor. In one example, the data sensitivity module 812 may detect a data sensitivity rating at an ingress and egress point of a secure data corridor. The ingress and egress points may interchangeably correspond to a subject and a secure data container, depending on the flow of data feeds. Thus, the data sensitivity module 812 may further assign a data sensitivity rating to the secure data corridor that corresponds with the determined data sensitivity rating at the ingress and egress points.

[0090] In a non-limiting example, the data sensitivity module 812 may quantify a data sensitivity rating of a secure data corridor based on access permissions associated with a

subject. For example, the data sensitivity module 812 may detect that access permissions associated with the subject may limit the data sensitivity rating of data feeds that the subject may access. For example, if a subject is permitted to access data feeds with a first data sensitivity rating, whereby the first data sensitivity rating is less sensitive (i.e., mandates less data security measures for data confidentiality and data integrity protection) than a second data sensitivity rating, the data sensitivity module 812 may associate the first data sensitivity rating to a secure data corridor that establishes a session link between a secure data container and the subject. Doing so ensures that data feeds that egress from the secure container to the subject comply with access permissions of the subject and do not include data feeds that may be more sensitive (i.e., mandate more data security measures for data confidentiality and data integrity protection) than permitted by the subject access permissions.

[0091] In another non-limiting example, the data sensitivity module 812 may quantify a data sensitivity rating of the secure data corridor based on the data sensitivity rating of incoming data feeds that ingress from the subject to a secure data container. In other words, the data sensitivity rating of the secure data corridor may be based on the data sensitivity rating of incoming data feeds that are transmitting into the secure data container, rather than data feeds being accessed by the subject from the secure data container. In this example, the data sensitivity module 812 may quantify a data sensitivity rating for incoming data feeds and associate that data sensitivity rating to the secure data corridor that establishes a session link between the subject and the secure data container. Doing so ensures that incoming data feeds transmitting from the subject to the secure data container are protected while transmitting through the secure data corridor.

[0092] Further, in another non-limiting example, the data sensitivity module 812 may quantify a data sensitivity rating of the secure data corridor based on the data sensitivity

rating of data feeds that transmit through the secure data corridor. In other words, the data sensitivity module 812 may detect the data sensitivity rating of data feeds transmitting through the secure data corridor, and dynamically set a substantially similar data sensitivity rating for the secure data corridor. A benefit of this approach is that in some instances, data feeds may have an unknown data sensitivity rating. As a precautionary measure, these data feeds may be assigned a high data sensitivity rating until determined otherwise. Thus to avoid assigning a secure data corridor with an arbitrary, and unnecessarily high data sensitivity rating that is based on a data feed with unknown data sensitivity ratings transmitting through the secure data corridors, the data sensitivity module 812 may progressively adjust the data sensitivity rating of the secure data corridor in response to progressively quantifying the unknown data sensitivity ratings of the data feeds.

[0093] Additionally, the data sensitivity module 814 may lock the data sensitivity rating of a secure data corridor based at least in part on the data sensitivity rating of the adjoining secure data container. For example, a secure data container may have assigned a particular data sensitivity rating based at least in part on an associated use-case. Thus, by locking the data sensitivity rating of the secure data corridor, the data sensitivity module 814 may ensure that data feeds that transmit through the secure data corridor comply with the data sensitivity rating of the use-case.

[0094] In the illustrated example, the secure data corridor module 814 may establish a secure data corridor between a subject device and a secure data container. In various examples, the secure data corridor may correspond to a link layer, a network layer, a transport layer, or an application layer of a transmission control protocol (TCP) and internet protocol (IP). Further, secure data corridor module 814 may assign the secure

data corridor with a data sensitivity rating that is determined by the data sensitivity module 812.

[0095] In the illustrated example, the audit module 816 facilitates a developer or administrator to audit an operation of the resource management platform. The audit module 816 may identify access permissions associated with a subject and compare the access permissions with control parameters of the data feeds. Further, the audit module 816 may interrogate data sensitivity ratings associated with established secure data corridors, secure data containers, and use-cases associated with secure data containers. Additionally, the audit module may ensure that each of the secure data corridors and secure data containers comply with the control parameters of the transmitting data feeds.

[0096] There may be a secure data dictionary module 818 that registers all data elements that are used in each secure data channel. Each of the data elements that are registered in the data dictionary module 818 map to data objects, which in turn store the data value at a particular storage location.

[0097] In the illustrated example, the resource management platform 802 may further include input/output interface(s) 820. The input/output interface(s) 820 may include any type of output interface known in the art, such as a display (e.g., a liquid crystal display), speakers, a vibrating mechanism, or a tactile feedback mechanism. Input/output interface(s) 820 also include ports for one or more peripheral devices, such as headphones, peripheral speakers, or a peripheral display. Further, the input/output interface(s) 820 may further include a camera, a microphone, a keyboard/keypad, or a touch-sensitive display. A keyboard/keypad may be a push button numerical dialing pad (such as on a typical telecommunication device), a multi-key keyboard (such as a conventional QWERTY keyboard), or one or more other types of keys or buttons, and may also include a joystick-like controller and/or designated navigation buttons, or the like.

[0098] In the illustrated example, the resource management platform 802 may further include one or more network interface(s) 822. The one or more network interface(s) 822 may include any sort of transceiver known in the art. For example, the one or more network interface(s) 822 may include a radio transceiver that performs the function of transmitting and receiving radio frequency communications via an antenna. In addition, the one or more network interface(s) 822 may also include a wireless communication transceiver and a near field antenna for communicating over unlicensed wireless Internet Protocol (IP) networks, such as local wireless data networks and personal area networks (e.g., Bluetooth or near field communication (NFC) networks). Further, the one or more network interface(s) 822 may include wired communication components, such as an Ethernet port or a Universal Serial Bus (USB).

Example Process

[0099] With the foregoing overview of an example architecture 100 that facilitates establishing a secure session between a subject and a secure data container via a secure data corridor of FIG. 1 and the various block diagrams of data access relationships of FIGS. 2 to 7, it may be helpful to provide an example process. To that end, FIG. 9 illustrates an example process 900 for establishing a session between a subject and a data feed, which is part of a secure data container, by way of a secure data corridor.

[00100] Process 900 illustrates a collection of blocks in a logical flow chart, which represents a sequence of operations that can be implemented in hardware, software, or a combination thereof. In the context of software, the blocks represent computer-executable instructions that, when executed by one or more processors, perform the recited operations. Generally, computer-executable instructions may include routines, programs, objects, components, data structures, and the like that perform particular functions or implement particular abstract data types. The order in which the operations are described

is not intended to be construed as a limitation, and any number of the described blocks can be combined in any order and/or in parallel to implement the process. For discussion purposes, the process 900 is described with reference to the architecture 100 of FIG. 1.

[00101] Process 900 describes establishing a session between a subject and a target data feed, which may be part of a secure data container, using a secure data corridor. In various examples, the subject may correspond to a user, a non-user entity, or another secure data corridor that requests permission to access data feeds within the particular secure data corridor.

[00102] At block 902, a resource management platform 110 may receive a request from a subject to access a data feed that includes at least one data element. This data feed may be part of a secure data container. In one embodiment, each data feed comprises one or more data objects. Each data feed object stores values defined by the corresponding data element.

[00103] At block 904, the resource management platform 110 identifies a use-case for the data feed.

[00104] At block 906, the resource management platform 110 determines, for each data element, a security level of the data feed.

[00105] At block 908, the resource management platform 110 determines a security level of each data element in the data feed.

[00106] At block 910, the resource management platform 110 determines, for each data element, one or more security controls that are mapped to the corresponding data element.

[00107] At block 912, the resource management platform 110 assigns a security label for the use-case. In various embodiments, the security label of the use-case may be based on at least one of: (i) the data sensitivity rating of the data feed, (ii) the security level of

each data element in the data feed, and (iii) the one or more security controls that are mapped to each data element.

[00108] At block 914, the resource management platform 110 compares a clearance of the subject to a security label of the use case having the data feed of interest. In this regard, the request from the subject may include the access permissions associated with the subject. In one embodiment, the authentication server 106 may perform authentication by way of referencing an active directory to confirm the identity and the credentials (e.g., clearance) of the subject.

[00109] Upon determining that the clearance of the subject is at or above the data security label of the use-case (i.e., “Yes” at decision block 916), the process continues with block 918, where the resource management platform 110 allows the subject to access the data feed via the secure data corridor.

[00110] However, upon determining that the clearance of the subject is below the security label of the use-case (i.e., “No” at decision block 916), the process continues with block 920, where the resource management platform 110 denies the subject to access the data feed via the secure data corridor.

Conclusion

[00111] While the foregoing has described what are considered to be the best mode and/or other examples, it is understood that various modifications may be made therein and that the subject matter disclosed herein may be implemented in various forms and examples, and that the teachings may be applied in numerous applications, only some of which have been described herein. It is intended by the following claims to claim any and all applications, modifications and variations that fall within the true scope of the present teachings.

[00112] It is understood that any specific order or hierarchy of steps in the process disclosed in FIG. 9 are illustrations of exemplary approaches. Based upon design preferences, it is understood that the specific order or hierarchy of steps in the processes may be rearranged, expanded, and some steps omitted. Some of the blocks may be performed simultaneously. For example, not all of the actions of determining a security level of the data feed, determining a security level of each data element in the data feed, or determining one or more security controls need to be performed (i.e., only one of the actions may be sufficient). In some embodiments, the actions of determining a security level of the data feed, determining a security level of each data element in the data feed, or determining one or more security controls may be performed concurrently or one after another (in any order).

[00113] Except as stated immediately above, nothing that has been stated or illustrated is intended or should be interpreted to cause a dedication of any component, step, feature, object, benefit, advantage, or equivalent to the public, regardless of whether it is or is not recited in the claims.

[00114] It will be understood that the terms and expressions used herein have the ordinary meaning as is accorded to such terms and expressions with respect to their corresponding respective areas of inquiry and study except where specific meanings have otherwise been set forth herein. Relational terms such as first and second and the like may be used solely to distinguish one entity or action from another without necessarily requiring or implying any actual such relationship or order between such entities or actions. The terms “comprises,” “comprising,” or any other variation thereof, are intended to cover a non-exclusive inclusion, such that a process, method, article, or apparatus that comprises a list of elements does not include only those elements but may include other elements not expressly listed or inherent to such process, method, article, or apparatus. An

element proceeded by “a” or “an” does not, without further constraints, preclude the existence of additional identical elements in the process, method, article, or apparatus that comprises the element.

[00115] The Abstract of the Disclosure is provided to allow the reader to quickly ascertain the nature of the technical disclosure. It is submitted with the understanding that it will not be used to interpret or limit the scope or meaning of the claims. In addition, in the foregoing Detailed Description, it can be seen that various features are grouped together in various embodiments for the purpose of streamlining the disclosure. This method of disclosure is not to be interpreted as reflecting an intention that the claimed embodiments require more features than are expressly recited in each claim. Rather, as the following claims reflect, inventive subject matter lies in less than all features of a single disclosed embodiment. Thus the following claims are hereby incorporated into the Detailed Description, with each claim standing on its own as a separately claimed subject matter.

CLAIMS

1. A computing device configured to provide a secure data corridor, the computing device comprising:
 - a processor;
 - a network interface communicatively coupled to the processor and configured to enable communications with a mobile traffic network;
 - a storage device for content and programming;
 - a security application stored in the storage device, wherein execution of the security application by the processor configures the computing device to perform acts comprising:
 - receiving a request from a subject for a data feed comprising at least one data element;
 - identifying a use-case for the data feed;
 - determining a data sensitivity rating of the data feed;
 - determining a security level of each data element of the data feed;
 - determining, for each data element of the data feed, one or more security controls that are mapped to the data element;
 - assigning a security label to the use-case;
 - comparing a clearance of the subject to the security label of the use-case;
 - upon determining that a clearance of the subject is at or above the data sensitivity rating of the use-case, allowing the subject privilege to the data feed via the secure data corridor; and
 - upon determining that the clearance of the subject is below the data sensitivity rating of the use-case, denying the subject privilege to the data feed via the secure data corridor.
2. The computing device of claim 1, wherein the security label is assigned to the use-case based on a combination of the data sensitivity rating of the data feed, the security

level of each data element, and the one or more security controls that are mapped to each data element.

3. The computing device of claim 1, wherein:
each data element defines its data type as a string, integer, float, or object; and
each data element maps to one or more data objects.
4. The computing device of claim 1, wherein execution of the security application further configures the computing device to perform acts comprising: upon determining that a subject has access to a use-case, allowing the subject to access all data elements and data objects related to the use-case.
5. The computing device of claim 1, wherein the secure data corridor corresponds to at least one of: (i) a link layer, (ii) a network layer, (iii) a transport layer, (iv) a session layer, and (v) an application layer of transmission control protocol and internet protocol (TCP/IP).
6. The computing device of claim 1, wherein the secure data corridor includes a container having an input and no output.
7. The computing device of claim 1, wherein the secure data corridor includes an input, an output, and a secure data container.
8. The computing device of claim 1, wherein the secure data corridor includes a secure data container with one or more outputs and no input.
9. A non-transitory computer-readable medium having stored thereon a plurality of sequences of instructions which, when executed by a processor, cause the processor to perform actions, the actions comprising:
receiving a request from a subject for a data feed comprising at least one data element;
identifying a use-case for the data feed;

determining a data sensitivity rating of the data feed;
determining a security level of a data element of the data feed;
determining, for the element of the data feed, one or more security controls that are mapped to the data element;
assigning a security label to the use-case;
comparing a clearance of the subject to the security label of the use-case;
upon determining that a clearance of the subject is at or above the data sensitivity rating of the use-case, allowing the subject privilege to the data feed via a secure data corridor; and
upon determining that the clearance of the subject is below the data sensitivity rating of the use-case, denying the subject privilege to the data feed via the secure data corridor.

10. The non-transitory computer-readable medium of claim 10, wherein the security label is assigned to the use-case based on a combination of the data sensitivity rating of the data feed, the security level of the data element, and the one or more security controls that are mapped to the data element.

11. The non-transitory computer-readable medium of claim 10, wherein the data sensitivity rating of the secure data corridor is dynamically set in substantially real time upon receiving the request from the subject for the data feed.

12. A computer-implemented method, the method comprising:
receiving a request from a subject for a data feed comprising at least one data element;
identifying a use-case for the data feed;
determining a data sensitivity rating of the data feed;
determining a security level of a data element of the data feed;
determining, for the element of the data feed, one or more security controls that are mapped to the data element;

assigning a security label to the use-case;
comparing a clearance of the subject to the security label of the use-case;
upon determining that a clearance of the subject is at or above the data sensitivity rating of the use-case, allowing the subject privilege to the data feed via a secure data corridor; and

upon determining that the clearance of the subject is below the data sensitivity rating of the use-case, denying the subject privilege to the data feed via the secure data corridor.

13. The computer-implemented method of claim 13, wherein:

the data element defines its data type as a string, integer, float, or object;
the data element maps to one or more data objects; and
each data object is configured to store one or more data elements mapped to the corresponding data object, at a storage location.

14. The computer-implemented method of claim 13, further comprising, upon determining that a subject has access to a use-case, allowing the subject to access all data elements and data objects related to the use-case.

15. The computer-implemented method of claim 13, wherein the control parameters include at least one of the following limitations: (i) no read-up, (ii) no write-down, (iii) tranquility, (iv) explicit read-in, and (v) explicit write-out.




FIG. 1




FIG. 2A




FIG. 2B




FIG. 3A

Control	Up	Down
Read	0	1
Write	1	0

306

FIG. 3B





FIG. 3C

Control	Up	Down
Read	0	2
Write	1	0

FIG. 3D






FIG. 5




FIG. 6





FIG. 7

**FIG. 8**

11/11




FIG. 9

A. CLASSIFICATION OF SUBJECT MATTER**H04L 29/06(2006.01)i, G06F 21/62(2013.01)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)

H04L 29/06; G06F 17/00; G06F 21/78; G06F 17/30; E05G 1/00; G06F 21/60; 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: secure data corridor, computing device, UE, network interface, security control, data element**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 8938775 B1 (GREGORY B. ROTH et al.) 20 January 2015 See column 1, lines 13-16; column 3, lines 39-41; column 6, lines 41-42; column 7, lines 3-32; column 9, lines 51-59; column 10, lines 1-48; column 11, lines 17-49; column 12, lines 44-46; claims 1, 3, 5; and figure 5.	1-15
Y	US 2015-0339497 A1 (BANK OF AMERICA CORPORATION) 26 November 2015 See paragraphs [0032], [0083], [0144]; and figure 7.	1-15
Y	US 8321387 B2 (RICHARD D. DETTINGER et al.) 27 November 2012 See column 4, lines 15-19.	11
A	US 2015-0256559 A1 (SRI INTERNATIONAL) 10 September 2015 See paragraphs [0099]-[0109] and figure 5.	1-15
A	WO 2010-106552 A2 (BHARATHRAJ, KRISHNA DHARMA) 23 September 2010 See claim 1 and figure 1.	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

26 April 2018 (26.04.2018)

Date of mailing of the international search report

26 April 2018 (26.04.2018)

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

NHO, Ji Myong

Telephone No. +82-42-481-8528



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2017/068680

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 8938775 B1	20/01/2015	None	
US 2015-0339497 A1	26/11/2015	US 9443112 B2	13/09/2016
US 8321387 B2	27/11/2012	US 2007-0027880 A1	01/02/2007
US 2015-0256559 A1	10/09/2015	US 2014-0007184 A1	02/01/2014
		US 2016-0072844 A1	10/03/2016
		US 9047463 B2	02/06/2015
		US 9210194 B2	08/12/2015
		US 9680876 B2	13/06/2017
WO 2010-106552 A2	23/09/2010	WO 2010-106552 A3	04/11/2010