



US012367294B2

(12) **United States Patent**
Yannuzzi et al.

(10) **Patent No.:** **US 12,367,294 B2**
(45) **Date of Patent:** **Jul. 22, 2025**

(54) **AUTOMATED DATA COMPLIANCE AND OBSERVABILITY**

11,507,408 B1 * 11/2022 Gabrielson H04L 9/3247
11,982,993 B2 * 5/2024 Cella G06F 18/23
2003/0086422 A1 * 5/2003 Klinker H04L 45/38
370/235

(71) Applicant: **Cisco Technology, Inc.**, San Jose, CA (US)

2005/0021689 A1 1/2005 Marvin et al.
2009/0099860 A1 4/2009 Karabulut et al.
2014/0098671 A1 * 4/2014 Raleigh H04M 15/61
370/235

(72) Inventors: **Marcelo Yannuzzi**, Vufflens-la-Ville (CH); **Hervé Muyal**, Gland (CH); **Jean Andrei Diaconu**, Haute-Savoie (FR); **Frank Brockners**, Cologne (DE); **Carlos Goncalves Pereira**, Carlsbad, CA (US)

(Continued)

(73) Assignee: **Cisco Technology, Inc.**, San Jose, CA (US)

OTHER PUBLICATIONS

Williamson et al., "Throttling viruses: restricting propagation to defeat malicious mobile code", 18th Annual Computer Security Applications Conference, 2002. Proceedings, Date of Conference: Dec. 9-13 (Year: 2002).*

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 257 days.

(Continued)

(21) Appl. No.: **17/859,707**

Primary Examiner — Morshed Mehedi

(22) Filed: **Jul. 7, 2022**

(74) *Attorney, Agent, or Firm* — BEHMKE INNOVATION GROUP LLC; James M. Behmke

(65) **Prior Publication Data**

(57) **ABSTRACT**

US 2024/0012911 A1 Jan. 11, 2024

In one embodiment, an observability and assurance service, associated with various clusters of application services for an application that are executed in a data mesh, may configure a data compliance filter for a particular application service in one of the clusters of application services according to a data compliance policy. The observability and assurance service may monitor the data and traffic associated with the particular application service, wherein the data compliance filter is applied to the traffic to restrict sensitive data in the traffic from being processed by the particular application service. The observability and assurance service may make a determination that the data compliance policy has been violated by the particular application service. The observability and assurance service may modify, based on the determination, the data compliance filter for the particular application service.

(51) **Int. Cl.**

G06F 21/00 (2013.01)
G06F 21/60 (2013.01)
G06F 21/62 (2013.01)

(52) **U.S. Cl.**

CPC **G06F 21/602** (2013.01); **G06F 21/6245** (2013.01)

(58) **Field of Classification Search**

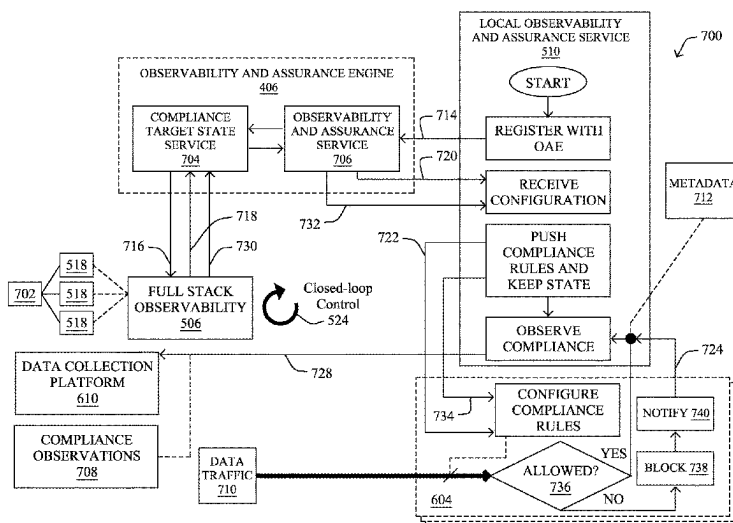
CPC G06F 21/602; G06F 21/6245
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

9,949,129 B1 * 4/2018 Henry H04W 12/08
10,713,664 B1 7/2020 Alagappan et al.

18 Claims, 10 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

2014/0280961 A1 9/2014 Martinez et al.
 2015/0254456 A1* 9/2015 Jacquin G06F 21/6245
 726/1
 2015/0281287 A1 10/2015 Gill et al.
 2016/0344736 A1 11/2016 Khait et al.
 2017/0149737 A1* 5/2017 Betzler H04W 12/02
 2017/0170970 A1 6/2017 Leighton et al.
 2017/0201569 A1 7/2017 Fu et al.
 2017/0300309 A1 10/2017 Berger et al.
 2018/0027022 A1 1/2018 Nagaratnam et al.
 2018/0115586 A1* 4/2018 Chou G06F 8/656
 2018/0124066 A1 5/2018 Minkovich et al.
 2018/0124113 A1 5/2018 Lock et al.
 2018/0260566 A1 9/2018 Chaganti et al.
 2018/0357226 A1* 12/2018 Su G06F 16/9535
 2019/0014123 A1 1/2019 Akireddy et al.
 2019/0228171 A1 7/2019 Mathur
 2020/0159947 A1* 5/2020 Shenefiel H04L 63/1425
 2020/0364351 A1 11/2020 Sanchez et al.
 2021/0006972 A1 1/2021 Guim Bernat et al.
 2021/0152561 A1 5/2021 Shelton et al.
 2021/0286638 A1 9/2021 Fan et al.
 2021/0329001 A1 10/2021 Barton et al.
 2021/0360037 A1 11/2021 Beckman et al.
 2022/0070189 A1* 3/2022 Liu H04L 63/1425
 2023/0155984 A1* 5/2023 Adam H04L 63/0281
 726/26
 2023/0195515 A1* 6/2023 Zhang G06F 9/5027
 718/102

OTHER PUBLICATIONS

“Global Apps, Local Compliance”, online: <https://incountry.com/>, accessed May 24, 2022, 10 pages.
 “Global Cloud Service Provider”, online: <https://us.ovhcloud.com/>, accessed May 24, 2022, 11 pages.
 “Google Distributed Cloud”, online: <https://cloud.google.com/distributed-cloud>, accessed May 24, 2022, 8 pages.
 Kurian, Thomas, “How Google Cloud is addressing the need for data sovereignty in Europe in 2020”, online: <https://cloud.google.com/blog/products/identity-security/how-google-cloud-is-addressing-data-sovereignty-in-europe-2020>, accessed May 24, 2022, 2 pages.
 “Gaia-X: A Federated Secure Data Infrastructure”, online: <https://www.gaia-x.eu/>, accessed May 24, 2022, 6 pages.
 “RegTech 100”, online: <https://fintech.global/regtech100/>, accessed May 24, 2022, 14 pages.
 “OneTrust Cloud Solutions”, online: <https://www.onetrust.com/>, accessed May 24, 2022, 5 pages.
 “Collibra—The Data Intelligence Cloud”, online: <https://www.collibra.com/us/en>, accessed May 24, 2022, 4 pages.
 “LogicGate Risk Cloud”, online: <https://www.logicgate.com/>, accessed May 24, 2022, 5 pages.
 “Governance and Security for Low-Code/No-Code Applications”, online: <https://www.zenity.io/>, accessed May 24, 2022, 6 pages.
 Zacks, et al., “Network Data Objectivization, Classification, Verification and Privacy via Ring-Oriented Metadata”, Defensive Publication Series, Jul. 2021, 11 pages, Technical Disclosure Commons.

* cited by examiner

100 ↙




FIG. 1A

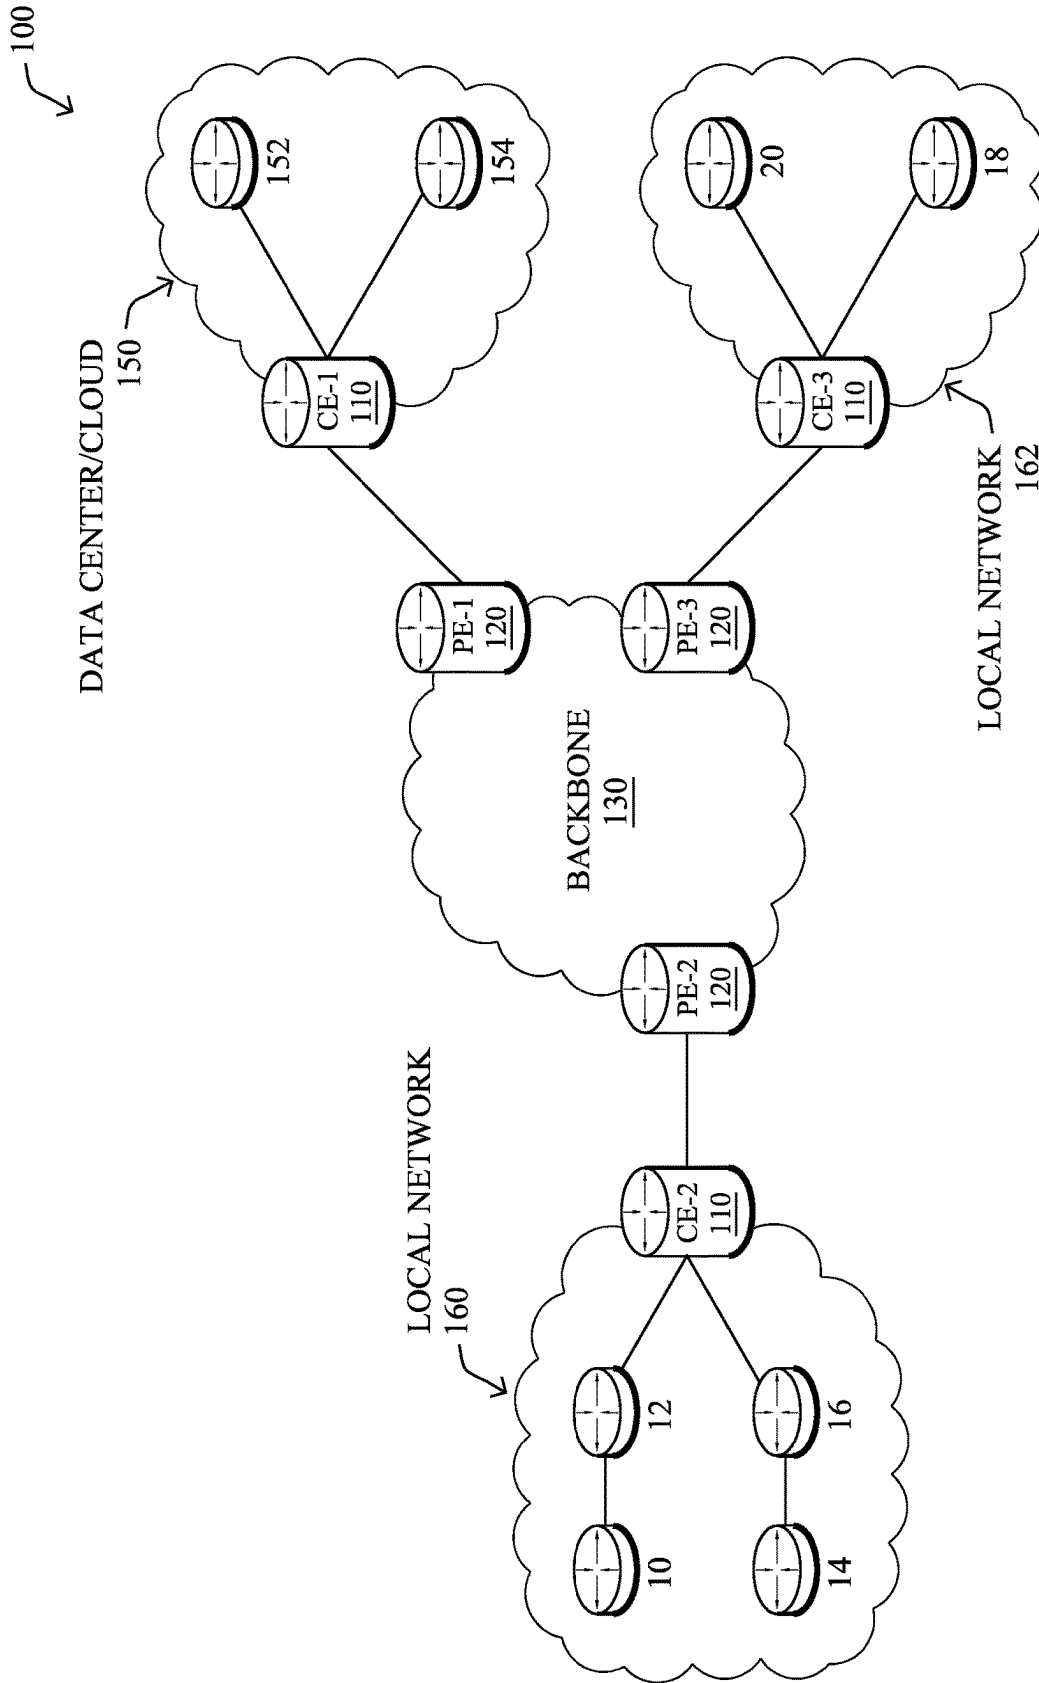


FIG. 1B

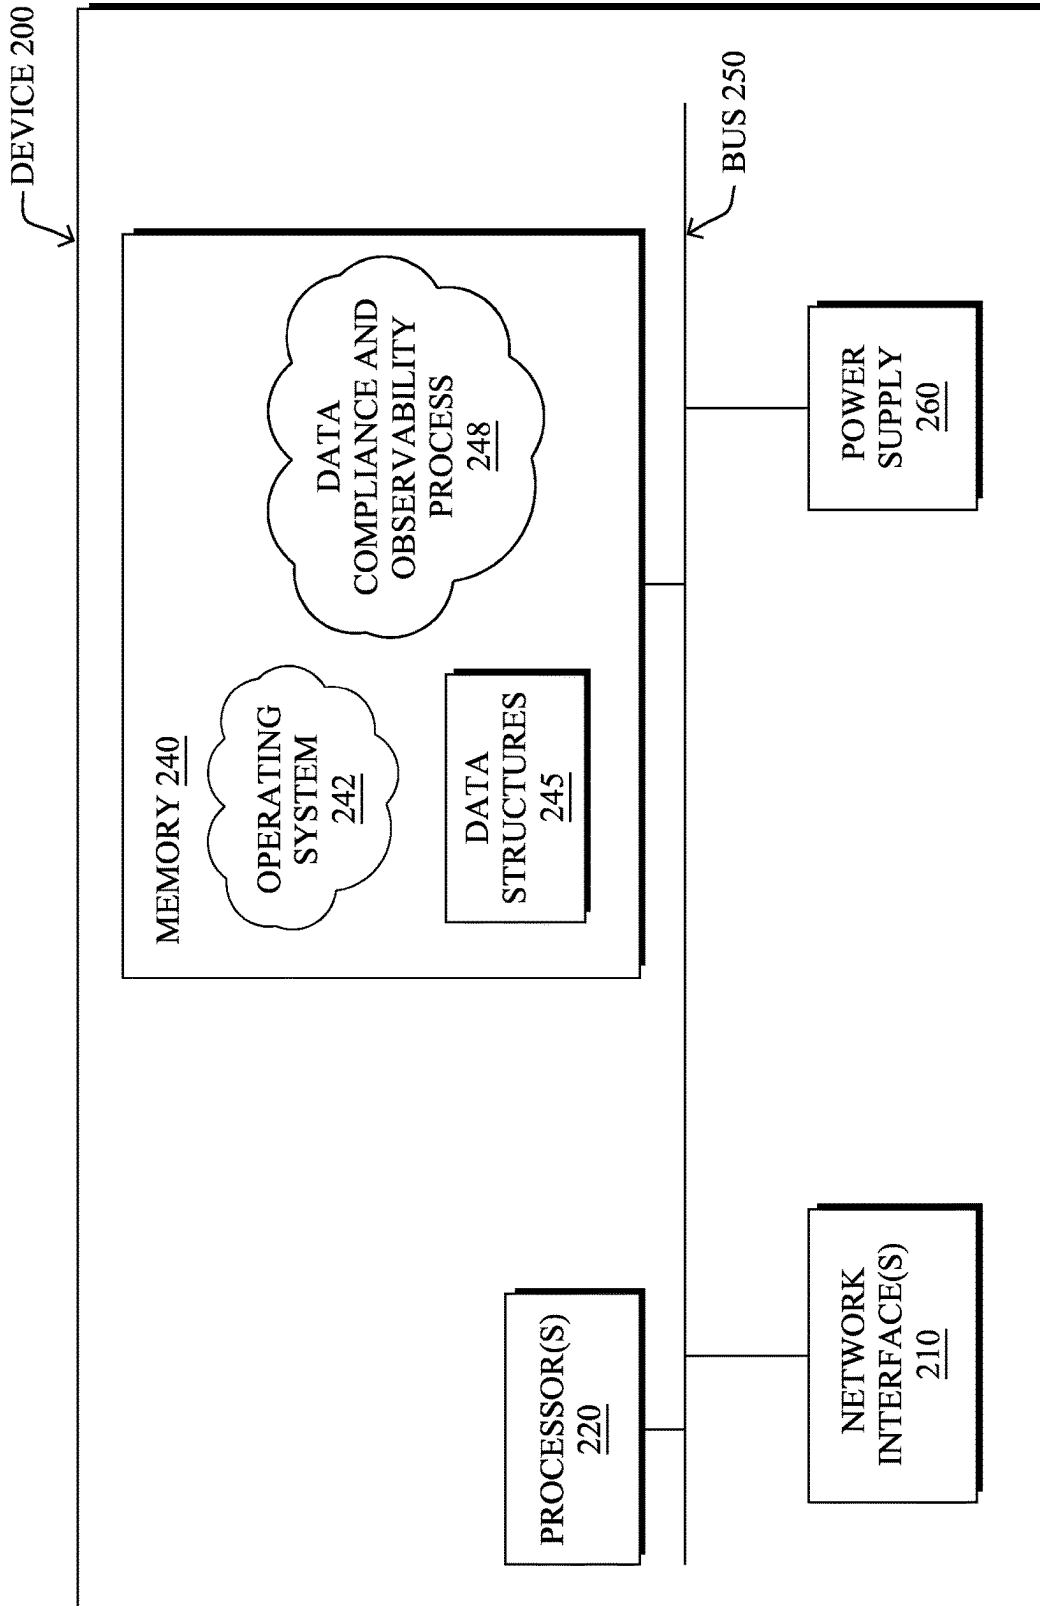


FIG. 2




FIG. 3




FIG. 4A

400

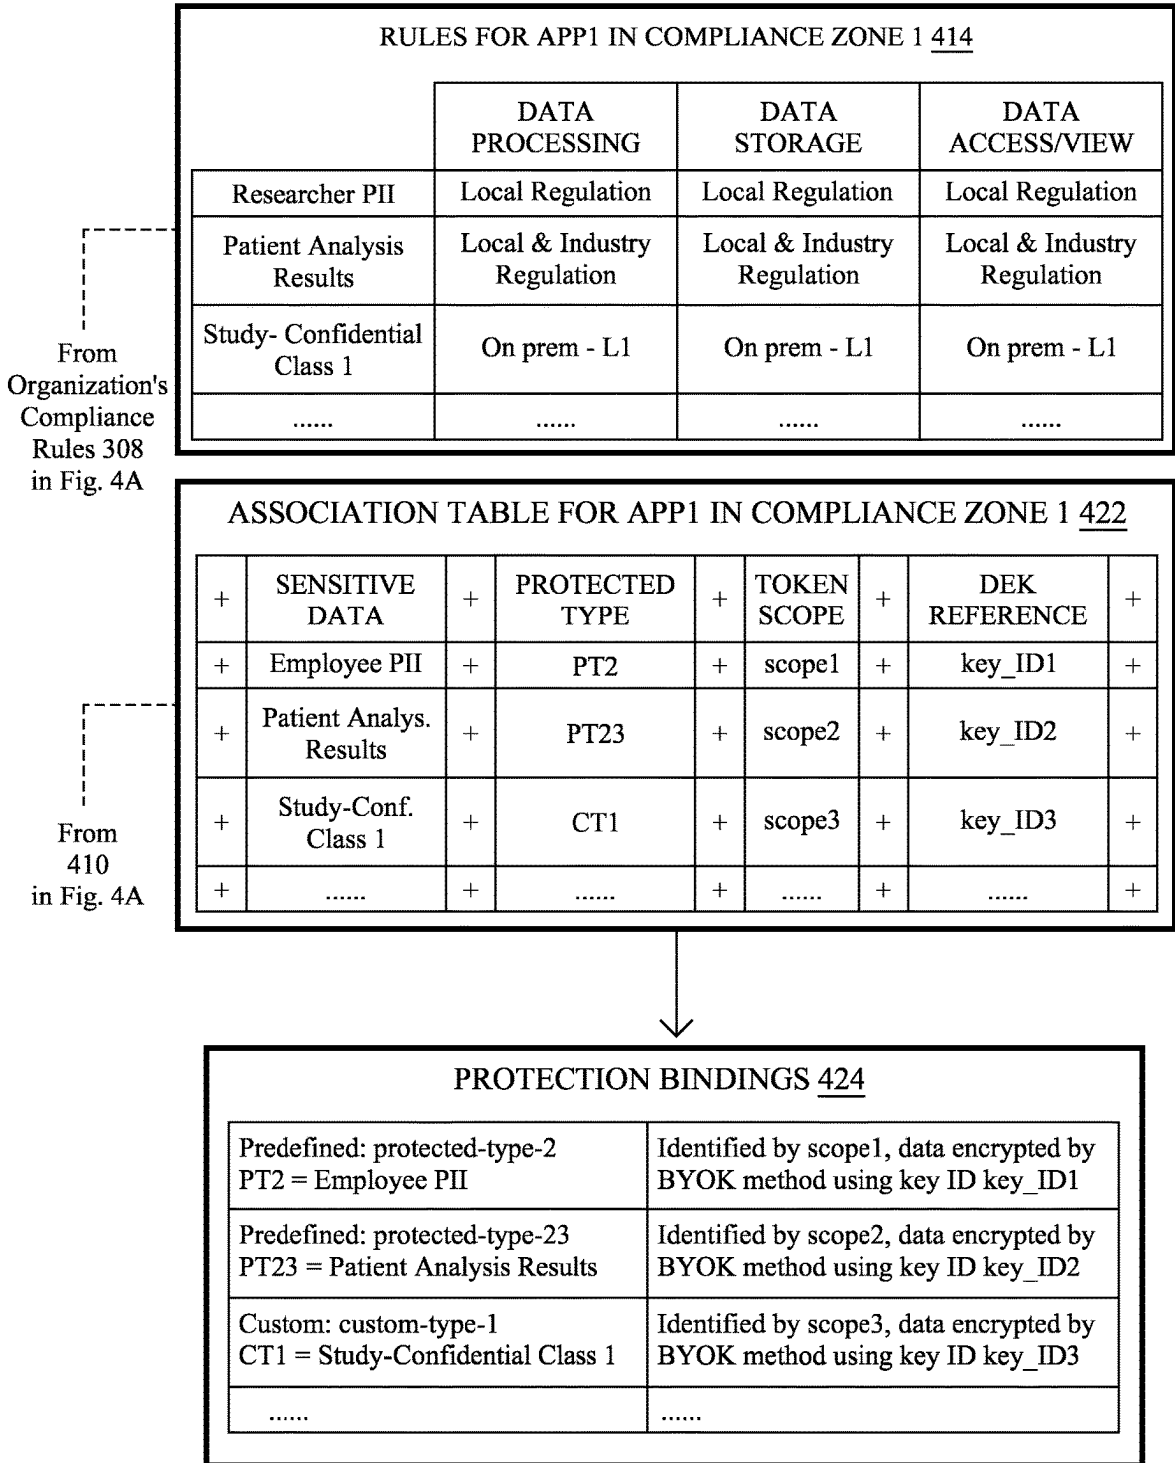


FIG. 4B




FIG. 5




FIG. 6




FIG. 7

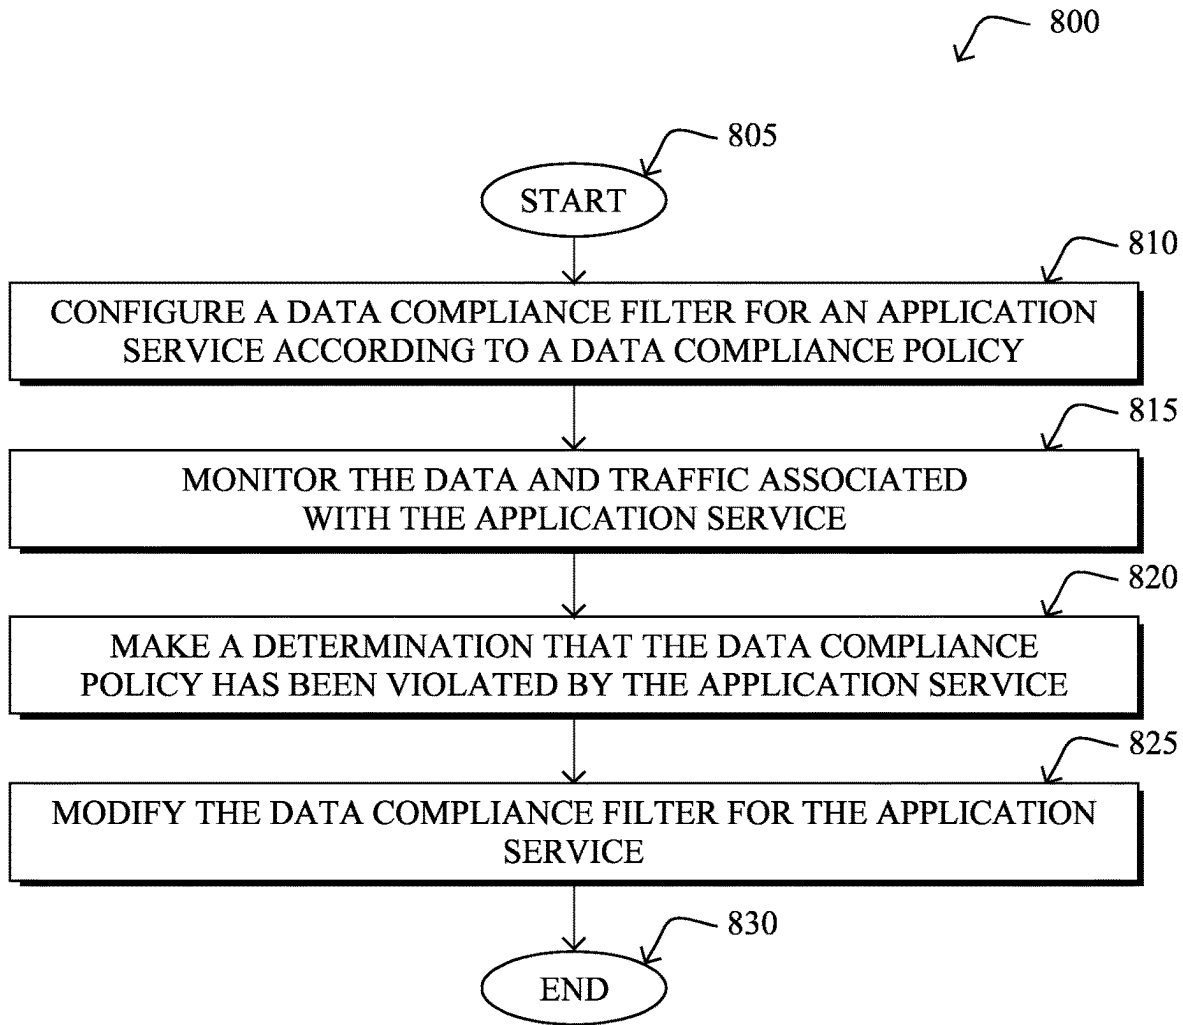


FIG. 8

1

AUTOMATED DATA COMPLIANCE AND OBSERVABILITY

TECHNICAL FIELD

The present disclosure relates generally to automated data compliance and observability.

BACKGROUND

Applications operate by handling data. For instance, executing an application can involve the storage, communication, processing, etc. of various types of data. The various types of data may include data whose handling is subject to various regulations. For example, data handling regulations at national, federal, state, industry, and/or organizational levels may be applicable to the data handled by an application.

The enforcement of data compliance has only been made more complex and potential violations made more likely as applications are increasingly being developed as a set of distributed services running across a mix of multi-cloud and edge infrastructures and handling a mix of data types differentially subject to various regulations. Unfortunately, following deployment of an application, enforcement of data compliance requirements largely occurs non-specifically and in a programmatic blind-spot. Given the current regulatory environment and trends, continuing to treat data compliance as an afterthought in this manner will likely yield increased violations of data compliance regulations which may result in substantial fines, penalties, and/or other negative impacts to data handlers.

BRIEF DESCRIPTION OF THE DRAWINGS

The embodiments herein may be better understood by referring to the following description in conjunction with the accompanying drawings in which like reference numerals indicate identically or functionally similar elements, of which:

FIGS. 1A-1B illustrate an example communication network;

FIG. 2 illustrates an example network device/node;

FIG. 3 illustrates an example architecture for data compliance;

FIGS. 4A-4B illustrates an example architecture for data compliance and protection bindings;

FIG. 5 illustrates an example architecture for data compliance and observability;

FIG. 6 illustrates an example architecture for data compliance and observability;

FIG. 7 illustrates an example process for data compliance and observability; and

FIG. 8 illustrates an example simplified procedure for data compliance and observability.

DESCRIPTION OF EXAMPLE EMBODIMENTS

Overview

According to one or more embodiments of the disclosure an observability and assurance service, associated with various clusters of application services for an application that are executed in a data mesh, may configure a data compliance filter for a particular application service in one of the clusters of application services according to a data compliance policy. The observability and assurance service

2

may monitor the data and traffic associated with the particular application service, wherein the data compliance filter is applied to the traffic to restrict sensitive data in the traffic from being processed by the particular application service.

5 The observability and assurance service may make a determination that the data compliance policy has been violated by the particular application service. The observability and assurance service may modify, based on the determination, the data compliance filter for the particular application service.

Description

A computer network is a geographically distributed collection of nodes interconnected by communication links and segments for transporting data between end nodes, such as personal computers and workstations, or other devices, such as sensors, etc. Many types of networks are available, with the types ranging from local area networks (LANs) to wide area networks (WANs). LANs typically connect the nodes over dedicated private communications links located in the same general physical location, such as a building or campus. WANs, on the other hand, typically connect geographically dispersed nodes over long-distance communications links, such as common carrier telephone lines, optical light-paths, synchronous optical networks (SONET), or synchronous digital hierarchy (SDH) links, or Powerline Communications (PLC) such as IEEE 61334, IEEE P1901.2, and others. The Internet is an example of a WAN that connects disparate networks throughout the world, providing global communication between nodes on various networks. The nodes typically communicate over the network by exchanging discrete frames or packets of data according to pre-defined protocols, such as the Transmission Control Protocol/Internet Protocol (TCP/IP). In this context, a protocol consists of a set of rules defining how the nodes interact with each other. Computer networks may be further interconnected by an intermediate network node, such as a router, to extend the effective "size" of each network.

40 Smart object networks, such as sensor networks, in particular, are a specific type of network having spatially distributed autonomous devices such as sensors, actuators, etc., that cooperatively monitor physical or environmental conditions at different locations, such as, e.g., energy/power consumption, resource consumption (e.g., water/gas/etc. for advanced metering infrastructure or "AMI" applications) temperature, pressure, vibration, sound, radiation, motion, pollutants, etc. Other types of smart objects include actuators, e.g., responsible for turning on/off an engine or perform any other actions. Sensor networks, a type of smart object network, are typically shared-media networks, such as wireless or PLC networks. That is, in addition to one or more sensors, each sensor device (node) in a sensor network may generally be equipped with a radio transceiver or other communication port such as PLC, a microcontroller, and an energy source, such as a battery. Often, smart object networks are considered field area networks (FANs), neighborhood area networks (NANs), personal area networks (PANs), etc. Generally, size and cost constraints on smart object nodes (e.g., sensors) result in corresponding constraints on resources such as energy, memory, computational speed and bandwidth.

FIG. 1A is a schematic block diagram of an example computer network 100 illustratively comprising nodes/devices, such as a plurality of routers/devices interconnected by links or networks, as shown. For example, customer edge (CE) routers 110 may be interconnected with provider edge

(PE) routers **120** (e.g., PE-1, PE-2, and PE-3) in order to communicate across a core network, such as an illustrative network backbone **130**. For example, routers **110**, **120** may be interconnected by the public Internet, a multiprotocol label switching (MPLS) virtual private network (VPN), or the like. Data packets **140** (e.g., traffic/messages) may be exchanged among the nodes/devices of the computer network **100** over links using predefined network communication protocols such as the Transmission Control Protocol/Internet Protocol (TCP/IP), User Datagram Protocol (UDP), Asynchronous Transfer Mode (ATM) protocol, Frame Relay protocol, or any other suitable protocol. Those skilled in the art will understand that any number of nodes, devices, links, etc. may be used in the computer network, and that the view shown herein is for simplicity.

In some implementations, a router or a set of routers may be connected to a private network (e.g., dedicated leased lines, an optical network, etc.) or a virtual private network (VPN), such as an MPLS VPN thanks to a carrier network, via one or more links exhibiting very different network and service level agreement characteristics. For the sake of illustration, a given customer site may fall under any of the following categories:

- 1.) Site Type A: a site connected to the network (e.g., via a private or VPN link) using a single CE router and a single link, with potentially a backup link (e.g., a 3G/4G/5G/LTE backup connection). For example, a particular CE router **110** shown in network **100** may support a given customer site, potentially also with a backup link, such as a wireless connection.
 - 2.) Site Type B: a site connected to the network by the CE router via two primary links (e.g., from different Service Providers), with potentially a backup link (e.g., a 3G/4G/5G/LTE connection). A site of type B may itself be of different types:
 - 2a.) Site Type B1: a site connected to the network using two MPLS VPN links (e.g., from different Service Providers), with potentially a backup link (e.g., a 3G/4G/5G/LTE connection).
 - 2b.) Site Type B2: a site connected to the network using one MPLS VPN link and one link connected to the public Internet, with potentially a backup link (e.g., a 3G/4G/5G/LTE connection). For example, a particular customer site may be connected to network **100** via PE-3 and via a separate Internet connection, potentially also with a wireless backup link.
 - 2c.) Site Type B3: a site connected to the network using two links connected to the public Internet, with potentially a backup link (e.g., a 3G/4G/5G/LTE connection).
- Notably, MPLS VPN links are usually tied to a committed service level agreement, whereas Internet links may either have no service level agreement at all or a loose service level agreement (e.g., a “Gold Package” Internet service connection that guarantees a certain level of performance to a customer site).
- 3.) Site Type C: a site of type B (e.g., types B1, B2 or B3) but with more than one CE router (e.g., a first CE router connected to one link while a second CE router is connected to the other link), and potentially a backup link (e.g., a wireless 3G/4G/5G/LTE backup link). For example, a particular customer site may include a first CE router **110** connected to PE-2 and a second CE router **110** connected to PE-3.

FIG. 1B illustrates an example of network **100** in greater detail, according to various embodiments. As shown, network backbone **130** may provide connectivity between

devices located in different geographical areas and/or different types of local networks. For example, network **100** may comprise local/branch networks **160**, **162** that include devices/nodes **10-16** and devices/nodes **18-20**, respectively, as well as a data center/cloud environment **150** that includes servers **152-154**. Notably, local networks **160-162** and data center/cloud environment **150** may be located in different geographic locations.

Servers **152-154** may include, in various embodiments, a network management server (NMS), a dynamic host configuration protocol (DHCP) server, a constrained application protocol (CoAP) server, an outage management system (OMS), an application policy infrastructure controller (APIC), an application server, etc. As would be appreciated, network **100** may include any number of local networks, data centers, cloud environments, devices/nodes, servers, etc.

In some embodiments, the techniques herein may be applied to other network topologies and configurations. For example, the techniques herein may be applied to peering points with high-speed links, data centers, etc.

According to various embodiments, a software-defined WAN (SD-WAN) may be used in network **100** to connect local network **160**, local network **162**, and data center/cloud environment **150**. In general, an SD-WAN uses a software defined networking (SDN)-based approach to instantiate tunnels on top of the physical network and control routing decisions, accordingly. For example, as noted above, one tunnel may connect router CE-2 at the edge of local network **160** to router CE-1 at the edge of data center/cloud environment **150** over an MPLS or Internet-based service provider network in backbone **130**. Similarly, a second tunnel may also connect these routers over a 4G/5G/LTE cellular service provider network. SD-WAN techniques allow the WAN functions to be virtualized, essentially forming a virtual connection between local network **160** and data center/cloud environment **150** on top of the various underlying connections. Another feature of SD-WAN is centralized management by a supervisory service that can monitor and adjust the various connections, as needed.

FIG. 2 is a schematic block diagram of an example node/device **200** (e.g., an apparatus) that may be used with one or more embodiments described herein, e.g., as any of the computing devices shown in FIGS. 1A-1B, particularly the PE routers **120**, CE routers **110**, nodes/device **10-20**, servers **152-154** (e.g., a network controller/supervisory service located in a data center, etc.), any other computing device that supports the operations of network **100** (e.g., switches, etc.), or any of the other devices referenced below. The device **200** may also be any other suitable type of device depending upon the type of network architecture in place, such as IoT nodes, etc. Device **200** comprises one or more network interfaces **210**, one or more processors **220**, and a memory **240** interconnected by a system bus **250**, and is powered by a power supply **260**.

The network interfaces **210** include the mechanical, electrical, and signaling circuitry for communicating data over physical links coupled to the network **100**. The network interfaces may be configured to transmit and/or receive data using a variety of different communication protocols. Notably, a physical network interface **210** may also be used to implement one or more virtual network interfaces, such as for virtual private network (VPN) access, known to those skilled in the art.

The memory **240** comprises a plurality of storage locations that are addressable by the processor(s) **220** and the network interfaces **210** for storing software programs and

data structures associated with the embodiments described herein. The processor 220 may comprise necessary elements or logic adapted to execute the software programs and manipulate the data structures 245. An operating system 242 (e.g., the Internetworking Operating System, or IOS®, of Cisco Systems, Inc., another operating system, etc.), portions of which are typically resident in memory 240 and executed by the processor(s), functionally organizes the node by, inter alia, invoking network operations in support of software processors and/or services executing on the device. These software processors and/or services may comprise data compliance and observability process 248, as described herein, any of which may alternatively be located within individual network interfaces.

FIG. 3 illustrates an example architecture 300 for data compliance, according to various embodiments. The architecture 300 may include a data compliance process 328. Data compliance process 328 may be utilized to provide configuration, observability, and enforcement of data compliance rules. Data compliance process 328 may accomplish these functions utilizing Data Compliance as Code (DCaC).

DCaC may include integrating a data compliance mechanism into the program code of the application. For example, data compliance process 328 may be utilized to build data compliance into the application development process, supported by automated code annotations, bindings between such annotations and categories of sensitive data, and controls at code, build, and pre-deploy time. Data compliance process 328 may provide a mechanism whereby application developers proactively assist data teams, application managers, and legal departments with data compliance, while ensuring that developers may remain oblivious to specific regulations, data related obligations, or compliance requirements that organizations might have across different regions.

For example, data compliance process 328 may include data annotating process 302. Data annotating process 302 may facilitate application developer 312 automatically adding metadata to program code of an application 316 during the development of the application 316. In various embodiments, this may be performed by automated annotations of data fields in the program code and by the creation of references to such annotations at code-build time. These references to annotated code may be automatically rendered in the form of machine-readable data manifest 314.

More specifically, data annotating process 302 may provide a mechanism for automated annotations of the program code of application 316, including classes, application programming interfaces (APIs), and the resulting data at code/build time (e.g., by implementing a Low-Code/No-Code approach supported by software development kits (SDKs) 320 and tooling 318). Application developers may utilize SDKs 320 and tooling 318 to automatically label data topics, data producers, data consumers, data processors, data holders, etc. For instance, developers may label certain data by annotating it with a data type identifier. For example, a developer may annotate certain data as “protected-type-1,” or other data as “protected-type-2,” and so on.

SDKs 320 in data annotating process 302 may provide a set of predefined data types out-of-the-box, including associations by default to specific categories of sensitive data. Sensitive data may include a type of data that may be considered personal, private, sensitive, confidential, protected, secret, restricted, personally identifiable information (PII), etc. In some examples, sensitive data may include data that is subject to regulation. For example, Table 1 lists

examples of predefined protected data types and default associations to some examples of categories of sensitive data.

TABLE 1

PROTECTED DATA TYPE	DEFAULT ASSOCIATION
protected-type-1	Customer PII
protected-type-2	Employee PII
...	...
protected-type-23	Patient Analysis Results
...	...
protected-type-41	Sales Confidential
...	...
protected-type-56	Restricted HR
...	...
unprotected	NA

A list of the associations, such as the example illustrated in Table 1, may provide associations by default to several categories of sensitive data, including but not limited to PII, confidential, restricted, and unprotected data. In some embodiments, the set of predefined protected data types might be standardized or rely on an existing taxonomy.

SDKs 320 in data annotating process 302 may also provide a mechanism to define and use custom data types in annotating program data of the application 316. For example, custom data types may be utilized, which identify protected data types that are not covered by any of those available by default in SDKs 320. For example, “custom-type-1” might be a custom data type associated to a category of sensitive data such as “Restricted Employee Poll.” In various embodiments, the generation and/or insertion of the annotations into the program code of the application 316 may be accomplished by an automated process (e.g., a programmatic identification of data of a particular data type triggering an automated insertion of an annotation of the data as the particular data type, etc.), a partially automated process (e.g., a programmatic flagging of data of a particular data type with a supervised or manual annotation of the data as the particular data type, etc.), and/or a manual process (e.g., a manual flagging of data of a particular data type and/or a manual annotation of the data as the particular data type, etc.).

In various embodiments, associations between protected data types and categories of sensitive data may be assigned and/or instrumented by different organizations and at different moments in time. In some cases, the association between protected data types and categories of sensitive data may be assigned by application developers 312 at code/build time. This might be the case when the team of application developers 312 is part of, or develops for, the organization that may use or manage the application 316. In such cases, the team of application developers 312 might have sufficient knowledge about the data and their use, so that they may either use the associations provided by default or create custom ones.

In additional instances, application developers 312 of application 316 and/or the users of the application 316 might belong to different organizations. For example, this may be the case when application developers 312 are a DevSecOps team that develops an application 316 that may be used across different organizations, industries, etc. In such cases, application developers 312 may be unaware of the categories of data that should be assigned by a data team and/or application manager 304 in another organization (e.g., precisely what data is confidential and what data is not with respect to that organization and its use of the application

316). In these instances, application developers 312 may leverage SDKs 320 and tooling 318 to approach data labeling and association in a manner that sidesteps the knowledge deficit while still instilling the functionality. For example, the application developers 312 may leverage SDKs 320 and tooling 318 to automatically add the different classes of protected data type at code and build time (e.g., utilizing predefined and custom protected data types). Additionally, or alternatively, the application developers 312 may leverage SDKs 320 and tooling 318 to automatically insert references in the form of machine-readable descriptions for the protected data types that may be used to generate data manifest 314 bound to application 316 at build time.

The protected data type annotations and their corresponding references may be utilized by a data team and/or application manager 304 in another organization to select and/or create automated associations 326 between categories of annotated data in the application 316 (e.g., metadata provided by application developers 312) and specific categories of sensitive data (e.g., personal data, private data, sensitive data, confidential data, protected data, secret data, restricted data, PII, etc.). For instance, each protected data type might be bonded to a class of tokens (e.g., JSON Web Tokens with a specific scope), which in turn might represent different categories of sensitive data for a data team and/or application manager 304.

In a specific example, an API call for application 316 may be labeled by application developers 312 with a data type identifier such as “custom-type-7” at code/build time. The “custom-type-7” labeled API call may attempt to access certain data using its bound token (e.g., “Token 7”) with a scope defined by, for example, a data team and/or application manager 304 before application 316 was deployed. From the data team and/or application manager 304 perspective, the attempt to access this data may translate to a request to access, for instance, “Confidential Partner” data. As such, the data type labels, and their associations may be utilized as an automated data mapping between the programmatic operations of application 316 and the sensitive data implicated in those operations. In various embodiments, these associations and functionalities may be supported by compliance engine 306 based on the selection, configuration, and automation of data compliance rules before application 316 is deployed and/or post-deployment.

In some examples, application developers 312, which again may be a DevSecOps team, might opt for a hybrid approach to generating these associations. For example, this may be the case when making some custom associations between data types and categories of sensitive data or using those predefined in the system (e.g., “protected-type-1” to “Customer PII”) might not only be trivial for the application developers 312 but also may facilitate the task of a data team and/or application manager 304 in defining associations. However, other associations might not be apparent to application developers 312. Hence, certain data in application 316 may be labeled as “protected types” along with their corresponding machine-readable descriptions in data manifest 314, though they may remain unassigned to a specific category of sensitive data, so they can be associated later by a data team and/or application manager 304 before the application is deployed, or by an automated data lineage, classification, and tagging process at run time (e.g., during the testing phase, that is, before the application is deployed in production).

In some embodiments, a data team and/or application manager 304 may be provided with a mechanism to change the associations created by application developers 312 or

even associate more than one category of sensitive data to a given data type (e.g., a data team and/or application manager 304 may associate certain data with both “Employee PIT” and “Confidential Data”). Hence, two categories of data compliance policies (e.g., one for “Employee PIT” and another for “Confidential Data”) may apply and restrict even further the access to this category of data. In general, a data team and/or application manager 304 may be able to Create, Read, Update, or Delete (CRUD) any association between the metadata provided by application developers 312 and categories of sensitive data.

In various embodiments, a data team and/or application manager 304 may proactively create a set of custom data types. A data team and/or application manager 304 may provide the set of custom data types to application developers 312. Application developers 312 may then utilize the set of custom data types so that application 316 is annotated at development based on guidelines (e.g., the set of custom data types, etc.) provided beforehand by the data team and/or application manager 304.

In additional embodiments, application developers 312 and a data team and/or application manager 304 may collaborate to annotate application 316. For example, application developers 312 and a data team and/or application manager 304 may iterate in the annotation and association processing in an agile manner. For example, the iteration may be performed as part of a continuous integration/continuous delivery (CI/CD) pipeline (e.g., at testing, staging, and production).

In some examples, application 316 may be composed of several services developed with different programming languages. Therefore, application 316 may utilize different SDKs 320. In some instances, the annotation methods and terminology applied to application 316 may vary depending on the programming language (e.g., usually referred to as attributes in C#, decorators in Python, annotations in Golang, etc.). In such cases, tooling 318 of data annotating process 302 may examine the different predefined and custom data types used with different SDKs 320, perform checks, and ensure consistency in the annotations and enumeration across the different services at build time. For example, these consistency checks may ensure that a given “custom-type-X” data type identifier represents the same type of data across services programmed using different programming languages even if they were programmed by different developers. Overall, the data annotating process 302 may provide different degrees of freedom to application developers 312, data teams and/or application managers 304, and the number of protected data types used, and their corresponding associations may vary depending on the type of application 316.

Data annotating process 302 may, as described above, be utilized in generating automated data references. Specifically, data annotating process 302 may automatically render a data manifest 314 bonded to application 316 at build time. Data manifest 314 may provide machine-readable descriptions of the predefined and/or custom data types used by application 316. A combination of SDKs 320 and tooling 318 may facilitate the instrumentation and automation of the program code at build time, including the automated rendering of data manifest 314. In some cases, application 316 may be composed of various containers. Each container may be built and packaged with its own data manifest, such that the final data manifest rendered for application 316 may be a composition of the individual data manifests. In some cases, application 316 may include dependencies on external services, such as a MySQL database. Such dependencies

may be captured as a dependency manifest. Data fed, processed, retained, or retrieved from these external services may also be annotated and automatically captured in application 316 data manifest 314.

Data annotating process 302 may, as described above, be utilized for decoupling data compliance from the business logic of application 316. For example, SDKs 320 and tooling 318 of data annotating process 302 may provide automated mechanisms for decoupling the configuration, observability, and enforcement of data compliance rules from the business logic of application 316. In some instances, application 316 may be a cloud/edge native application, which may be implemented as a set of workloads composing a service mesh. The decoupling of data compliance from the business logic may be especially relevant for applications of this type, as geographically dispersed and/or variably deployed workloads may implicate increased data compliance complexity.

Various possible embodiments for decoupling data compliance from the business logic of application 316 may be utilized. For instance, a sidecar model, where the services that implement the business logic of application 316 are deployed together with sidecar proxies associated to each of those services, may be utilized. The sidecar proxies may be utilized to enforce horizontal functions that are independent of the business logic, such as routing, security, load balancing, telemetry, retries, etc. As such, the sidecars may be well-positioned to decouple, observe, and control data compliance. For example, a combination of distributed data compliance controllers and sidecar proxies may be used to configure, observe, and enforce data compliance rules across different geographies, and distributed multi-cloud and edge infrastructures 334.

Instead of, or in addition to, using sidecars, various embodiments may use client libraries, daemons working in tandem with the application-specific services, or sandboxed programs in the OS kernel, e.g., using the Extended Berkeley Packet Filter (eBPF). Further embodiments may use an agentless approach or embed such functionality in Kubernetes itself. In any case, the functionality introduced herein may enable the portability and reuse of observability and enforcement of data compliance functions across not only different applications but also cloud and edge environments.

The above-described data annotating process 302 may yield a portable annotated application 316 that is geared with built-in annotations for different types of protected data. In addition, the yielded annotated application 316 may be structured to operate while remaining agnostic of any state, country, industry, organization-specific regulation and/or data policy requirements that a data team and/or application manager 304 might have. As a result, data annotating process 302 may be leveraged as a new model of building applications including DCaC by not only data teams and/or application managers 304, but also software as a service (SaaS) providers and others.

Data compliance process 328 may provide configuration, observability, and enforcement of data compliance rules. As described above, associations 326 between categories of annotated data in application 316 and specific categories of sensitive data may be instrumented prior to a deployment of application 316. The associations 326 may be used to control the processing and use of data during and after the deployment of application 316. More specifically, compliance engine 306 may utilize associations 326 together with current data compliance regulations governing data handling in each region where application 316 may be used, as well as a specific organization's compliance rules 308 for/while

using application 316, to enforce compliance with them. Such controls may apply to data access requests, data storage and retention policies, data processing requirements, etc. of application 316 both at deploy and execution time, etc.

To this end, data compliance process 328 may include data compliance regulation repository 322. Data compliance regulation repository 322 may provide a repository of data compliance rules. For example, data compliance regulation repository 322 may include a repository of industry regulations 324 which may be applicable to the use of application 316. For example, with respect to instances where application 316 is used by a healthcare provider, data compliance regulation repository 322 may include industry regulations 324 such as Health Insurance Portability and Accountability Act of 1996 (HIPAA) regulations applicable to handling of data in the healthcare industry. In other examples, data compliance regulation repository 322 may include a repository of national regulations 330 which may be applicable to the use of application 316. For example, with respect to instances where application 316 is based in a member state of the E.U., data compliance regulation repository 322 may include national regulations 330 such as the GDPR applicable to handling of data in the E.U.

The data compliance regulations included in data compliance regulation repository 322 may be consumed by a data team and/or application manager 304 as a service (aaS). Data compliance regulation repository 322 may support input, expression, collection, approval, visualization, and/or use of data compliance policies covering multiple categories of rules. For example, data compliance regulation repository 322 may store data compliance policies that are specific to an industry, those that may apply at a national, multi-national, federal, state, and industry levels, etc. For instance, an organization (e.g., a multi-national company) may leverage a data compliance regulation repository 322 service of a data compliance process 328 and utilize the regulations already available in data compliance regulation repository 322, which may cover regulations across several industries and countries out-of-the-box. An organization may select the target state, country or region, the industry if needed, and select the data compliance regulations that may be applicable at the organizational level (e.g., organization's compliance rules 308).

Compliance engine 306 may offer APIs and a user-friendly user interface (UI) through which a data team and/or application manager 304 may select and define data compliance requirements. For instance, if application 316, which handles Customer PII data, needs to be deployed in British Columbia, Canada, a data team and/or application manager 304 may simply select "Customer PII→Apply Local Regulation" to constrain the processing, storage, retention, and access to Customer PII data according to the regulations in British Columbia as retrieved from data compliance regulation repository 322. To this end, compliance engine 306 may compute and handle the resulting constraints that apply to Customer PII data in British Columbia transparently to data teams and/or application managers 304. More specifically, the set of data compliance constraints may be captured in a machine-readable format from data compliance regulation repository 322, and therefore, used by compliance engine 306 programmatically.

In some examples, compliance engine 306 may be utilized as a pluggable module working in tandem with one or more workload engines 332, such as Cisco Intersight or any automation tool offered by a hyperscaler, or other cloud and edge providers. Workload engines 332 may manage the

deployment of application 316, subject to the rules and constraints provided by compliance engine 306.

In various embodiments, compliance engine 306 may operate either in a push or a pull model. For instance, in a pull model, a workload engine 332 may receive a request to deploy application 316 in a given region (e.g., a request from a site reliability engineering (SRE) and/or information technology (IT) team 310). In such a case, workload engine 332 may issue a request to compliance engine 306, to compute and return data compliance rules and constraints that must be applied for their specific deployment. Alternatively, in a push model, a data team and/or application manager 304 may select the compliance rules required and a declarative intent for application deployment may be issued from compliance engine 306 to one or more workload engines 332. Such deployments may involve multi-cluster service meshes, which may run across multi-cloud and edge infrastructures 334. In various embodiments, the above-described sidecar proxies in the service mesh may not only provide monitoring and observability of data compliance to compliance engine 306 but also may receive configuration and compliance updates in real-time 336. In additional embodiments, the same functionality may be implemented utilizing client libraries, daemons, eBPF, an agentless approach, or Kubernetes itself. In addition, some embodiments may support the techniques described herein without utilizing a service mesh.

FIGS. 4A-4B illustrate an example architecture 400 for data compliance and protection bindings according to various embodiments. Architecture 400 may be utilized to deliver DCaC. For example, architecture 400 may include compliance engine 306. Compliance engine 306 may include data compliance rules module 402, compliance intent engine 404, and/or observability and assurance engine 406.

Data compliance rules module 402 may compute compliance constraints based on a combination of inputs. For example, the constraints may be computed based on a combination of inputs including the target state, country, or multi-country region for an application, the industry the application is being utilized within, and/or the compliance rules required by an organization using the application.

Compliance intent engine 404 may include an association service 408. Association service 408 may manage a set of association tables 410. Association service 408 may include functionality to allow a data team and/or application manager 304 to select, configure, and create the associations (e.g., associations 326 from FIG. 3) and store them in the form of protection bindings 424 in association tables 410 (e.g., populated association table 422 and its associated protection bindings 424 provide a non-limiting specific example of one such association table). The protection bindings 424 may define a data handling scope bonded to the association between a data type and its associated category of sensitive data. The data handling scope may be an indication of how data will be handled by the application (e.g., applicable tokens, token scopes, encryption keys, etc.). The protection bindings 424, stored in populated association table 422 may be created and/or maintained by compliance engine 306. The protection bindings 424 may not be reinserted into the program code but rather maintained by compliance engine 306 since protection bindings 424 may be subject to changes over time (e.g., the scopes might change, encryption keys may be rotated, etc.) and keeping the protection bindings 424 outside of the program code may prevent these changes from affecting the program code.

Additionally, compliance intent engine 404 may include a service intent engine 412. Service intent engine 412 may interface with workload engines 332. Service intent engine 412 may either receive and/or process data compliance requests from workload engines 332 (e.g., pull model) or create and/or issue declarative intents encoding a deployment request to workload engines 332 (e.g., push model).

Observability and assurance engine 406 may receive telemetry data from services deployed in the field (e.g., from a service mesh deployed across multi-cloud and edge infrastructures 334). In addition, observability and assurance engine 406 may push data compliance configurations and data traffic filters in real-time 336 out to workloads deployed in the multi-cloud and edge infrastructures 334.

To populate association tables 410 and create protection bindings 424, association service 408 may obtain inputs defining the associations of the protected data types and/or their data handling scopes. For example, for each annotated application 316, association service 408 may obtain data manifests 314 describing the set of predefined and custom protected data types handled by annotated applications 316 and potential associations already made by application developers (e.g., application developer 312 in FIG. 3).

Additionally, association service 408 may obtain data handling scopes to be bonded to individual associations between given data types and their associated category of sensitive data. For example, association service 408 may obtain, as an input, categories of tokens and corresponding scopes that may be associated with each of the protected data types used in the annotated classes and methods that compose the annotated application 316. Such tokens may be defined and/or obtained from external token management service 416 (e.g., from OKTA).

Further, association service 408 may obtain, as an input, identifiers of encryption keys. The encryption keys may be those keys associated with each of the protected data types used in the annotated classes and methods that compose annotated applications 316. For example, different categories of sensitive data may be encrypted utilizing techniques such as bring your own key (BYOK) or hold your own key (HYOK). The encryption key identifiers (ID) may be defined or obtained from external encryption key service 418, which may extract encryption key IDs. External encryption key service 418 may interface with key management service (KMS) 420 and may create references (e.g., key IDs) to encryption keys stored and managed by KMS 420. In this manner, the keys may not be managed by association service 408, but instead may remain secure with KMS 420.

This set of inputs may be utilized by association service 408 to populate association tables 410. In some examples, there may be one association table populated per annotated application 316 and/or per data compliance zone (e.g., a geographical area where the application is deployed, etc.). A populated association table 422 may include the automatically associated annotated data types (e.g., protected data type labels), with categories of sensitive data (e.g., encoded in the form of tokens with specific scopes as illustrated in populated association table 422), along with pointers to the encryption keys used for each category of protected data (e.g., key IDs).

Compliance engine 306 may handle one populated association table 422 per compliance zone for each annotated application 316. In some embodiments, populated association table 422 may be extended to include columns denoting API paths to access the data resources. In some cases, the paths may explicitly embed the protected data types used by the application developers. In addition, the tokens, scopes,

and key IDs may be renewed and dynamically updated in populated association table 422 depending on their validity and/or expiration time. Populated association table 422 may also manage more advanced associations, including m:n bindings (e.g., two or more protected data types might be associated to a single token/scope class).

Once association tables 410 are defined and/or populated, a data team and/or application manager 304 may select the organization's compliance rules 308. Organization's compliance rules 308 may include data compliance rules (e.g., data compliance rules 414) selected to be required for a specific annotated application 316 (e.g., "App 1") that should be deployed in a specific geographical region (e.g., "Compliance Zone 1").

For example, data compliance rules 414 may include constraints selected to be applied to each category of sensitive data (e.g., "Researcher PII," "Patient Analysis Results," "Study-Confidential Class 1," etc.) within annotated application 316 "App 1". Data compliance rules 414 may further specify the selected constraints associated with processing the category of sensitive data, storage, and retention of the category of sensitive data, and/or accessing or viewing of the category of sensitive data by the annotated application 316.

Data compliance rules 414 may be rules that may apply to a new annotated application 316 (e.g., "App 1"). In this example, application "App 1" may handle data about clinical trials of a new drug developed by a pharmaceutical company. Application "App 1" may collect and analyze data and provide insights about the new drug. In this example, the application "App 1" may support several categories of sensitive data.

For example, the application "App 1" may support the category of sensitive data "Researcher PII" data, which may include PII of the employee conducting one of the trials. In this example, the rule chosen may be one to restrict the processing, storage, retention, and access to the data according to a "local regulation." For example, if the compliance zone where the application is going to be deployed is British Columbia, Canada, then this rule automatically constrains the location of workloads, the storage, and any attempt to access or view such data according to the regulation on PII in British Columbia, Canada. If the trials were instead to be conducted in New Delhi, India, the rule might be the same (e.g., applying a local regulation), but the resulting constraints yielded by compliance engine 306 will usually differ from those in British Columbia, Canada by virtue of the two locations having different data handling regulations.

Application "App 1" may also support the category of sensitive data "Patient Analysis Results" data. In this example, the rule chosen by a data team and/or application manager 304 may adhere to both the local and industry-specific regulation. In the examples of application "App 1," the industry-specific regulation may include, for example, specific legislation constraining the processing, storage, retention, and access to patients' data with respect to clinical trials.

Application "App 1" may additionally support the category of sensitive data "Study-Confidential Class 1" data. This data type may be a custom protected data type that may be implemented to enable researchers to keep a specific category of data related to the clinical trials as highly confidential. In this example, the processing, retention, and access to the data may be constrained to a specific facility. For example, perhaps the processing, storage, retention, and access of the "Study-Confidential Class 1" data is con-

strained to the premises of a "Laboratory 1" associated with the clinical trial (e.g., "On prem-L1").

Based on data compliance rules 414 input along with the corresponding protection bindings 424, data compliance rules module 402 may identify the implicated categories of sensitive data. For example, the three categories of sensitive data (e.g., "Employee PII," "Patient Analysis Results," and "Study-Confidential Class 1") listed in data compliance rules 414, along with corresponding protection bindings 424 in populated association table 422 may be identified by data compliance rules module 402. These categories of sensitive data may be defined and/or used by a data team and/or application manager 304 and/or may have already been associated to specific predefined and custom protected data types, such as the ones shown in populated association table 422 and/or protection binding 424.

In addition, data compliance rules module 402 may identify the selected compliance requirements listed in data compliance rules 414. For example, data compliance rules module 402 may identify the compliance requirements specified for processing, storage, retention, and access for each of the categories of sensitive data as defined in data compliance rules 414.

Data compliance rules module 402 may compute the set of compliance constraints that apply to application "App 1" based on data compliance rules 414 and/or a compliance zone selected (e.g., a target country and industry for "App 1"). In some examples, the set of compliance constraints may be computed from, for example, industry regulations 324, national regulations 330, etc. obtained from data compliance regulation repository 322.

The output of data compliance rules module 402 (e.g., the computed set of compliance constraints for a category of sensitive data) may be processed by compliance intent engine 404. Compliance intent engine 404 may link the resulting constraints to the corresponding populated association table 422 and send this output to both observability and assurance engine 406 and service intent engine 412. As such, the compliance constraints may be linked to categories of sensitive data and/or their associated protected data types in the program code. Therefore, the compliance constraints may be linked to individual portions of the application code. For instance, the constraints may be linked to control a data transfer through an API call that was previously annotated by application developers 312 using the protected data types referenced in protection bindings 424, populated in association table 422, and constrained by data compliance rule module 402 according to data compliance rules 414.

In various embodiments, once a service mesh is deployed, a data consumer process may request access to a data resource through an API. This may be implemented using a GET method including a path containing the field "custom-type-1" (CT1), which, according to populated association table 422, represents the custom protected type "Study-Confidential Class 1." The HTTP request may be transported and forwarded over mTLS across the sidecar proxies in the service mesh. The authorization header in the service mesh may carry "token 3" with a specific scope "scope 3," as defined in populated association table 422. In this example, "scope 3" represents the category of sensitive data "Study-Confidential Class 1," and the constraint in this case is that the data of that type must be retained on "prem Lab1." To that end, service intent engine 412 may have requested and/or instructed a workload engine 332 to deploy the workloads handling "CT1" "on prem L1." In turn, observability and assurance engine 406 may have configured data filters in the sidecar proxies to enforce access control. For

instance, API calls using an authorization token with “scope 3” may be restricted to data consumers located “on prem L1.”

Compliance engine **306** may process more elaborate data compliance rules than simply those illustrated, including the selection of specific locations for processing, storage, retention, and access for each category of sensitive data. For instance, a data team and/or application manager **304** may choose a specific data center (e.g., where their data warehouse is hosted), a compliant public or edge zone, a compliant private cloud or edge site, combinations of these, etc. While some of these selections may be very specific, others might remain openly declarative, which a member of SRE/IT team **310** may translate into a specific infrastructure request for deploying the application, or some of the services that comprise the application.

Some of the tokenization mechanisms described herein may be externally handled, such as by an authorization server, which may potentially work in concert with a delegated authorization solution (e.g., OAuth 2.0/OpenID Connect), a single sign-on (SSO) solution, etc. In such scenarios, the specific categories of tokens and scopes references in association tables **410** may be obtained from external systems.

In addition to the tokenization mechanisms, a data team and/or application manager **304** may also select the user and/or process groups that may have access to the different categories of sensitive data (e.g., read only, read and write, or no access). Such groups may be managed using internal tools or they may be externally handled by an authorization service (e.g., OKTA). In some embodiments, the selection of token scopes and the access rights applied to user and/or process groups may be made jointly. These additional constraints may also be part of the data compliance rules and state maintained by data compliance rule module **402**.

In an alternative embodiment, compliance engine **306** may also enable the use of third-party annotations and/or data catalogs (e.g., imported from external data classification and tagging systems, such as from Collibra, OneTrust, or others). In such cases, the protected types used in protection bindings **424**, and populated in association table **422**, may be comprised of a set of annotation labels (i.e., metadata) added by application developers **312** (e.g., PT2, PT23 and CT1 in FIGS. 4A-4B) as well as third-party labels provided by external systems.

As such, DCaC implemented through architecture **400** may provide two levels of decoupling. First, a decoupling between the annotations or metadata embedded in the program code of the application as provided by application developers **312** at code/build time and the categories of sensitive data that may be selected and associated by a data team and/or application manager **304** before the annotated application **316** is deployed. Second, a decoupling between the categories of sensitive data and the rules selected by a data team and/or application manager **304**, and the specific data compliance regulation and data compliance constraints that may apply to a given industry and/or region.

Such an approach may facilitate application developers **312** proactively assisting a data team and/or application manager **304**, while all of them are allowed to remain oblivious to the specificities and intricacies of the different data compliance regulations across the different industries and regions.

It will be apparent to those skilled in the art that other processor and memory types, including various computer-readable media, may be used to store and execute program instructions pertaining to the techniques described herein.

Also, while the description illustrates various processes, it is expressly contemplated that various processes may be embodied as modules configured to operate in accordance with the techniques herein (e.g., according to the functionality of a similar process). Further, while processes may be shown and/or described separately, those skilled in the art will appreciate that processes may be routines or modules within other processes.

As noted above, the number of laws, regulations, and rules regarding the storage and use of certain types of data are continually increasing across the globe. For instance, the General Data Protection Regulation (GDPR) in Europe places strict regulations on how a user’s personal data is collected and shared. These and other regulations have spawned independent efforts across several countries to ensure that online applications comply with specific data regulations at national, federal, or state level, and particularly, those that are cloud-based. This acceleration in data sovereignty regulations is posing complex challenges to the organizations that use or manage that data, since legal obligations and constraints vary from country to country. The challenge is even greater since data compliance requirements are often not limited to data sovereignty obligations. For example, depending on the type of an application, data compliance may demand the amalgamation of other regulations, such as industry-specific regulation (e.g., complying with HIPAA obligations in the healthcare industry in the United States), or organization-specific rules (e.g., on how to deal with confidential data). Regardless, of the source of the data compliance requirement, it’s observance may involve the handling of data in a particular way at runtime.

In addition, applications are increasingly being developed as a set of distributed services that run across a mix of multi-cloud and edge infrastructures. However, enforcing data compliance and ensuring data sovereignty across multi-cloud and edge is particularly challenging. For example, a certain law may require user data generated in a particular country to remain within the borders of that country. Ensuring that these obligations are met across the various clouds (e.g., public, private, etc.) and edges, as well as garnering insight into where the application data is retained and how it is used, are functionalities that are wholly absent in existing data compliance efforts.

For example, existing data compliance strategies lack integration in the application program code itself (e.g., they do not utilize a DCaC model) and instead rely on dynamically trying to determine what type of data is being shared and/or used. In addition, existing strategies may include restricting data residency to specific infrastructure providers to try to address data residency issues, but this strategy can lead to data lock-ins and utilization of sub-optimal infrastructure lacking a global footprint. These strategies are insufficient for cloud/edge native applications that are increasingly being developed as service meshes deployed across a mix of multi-cloud and edge infrastructures. Furthermore, these strategies leave an organization with poor data visibility and an inability to automatically geofence sensitive data across different cloud and edge infrastructure providers.

As such, existing data compliance strategies lack any ability to provide automated geofencing and observability of application data subject to the various data sovereignty and compliance rules that may apply across different geographical locations, particularly, when applications are deployed in distributed multi-cloud and edge environments.

Automated Data Compliance and Observability

The techniques herein introduce mechanisms for automated data geofencing and observability across multi-cloud

and edge infrastructures, in a manner that ensures data compliance and data sovereignty across different geographical regions. Indeed, the techniques herein provide portability and reuse of observability and enforcement of data compliance functions across not only different applications and regions but also cloud and edge environments.

Illustratively, the techniques described herein may be performed by hardware, software, and/or firmware, such as in accordance with data compliance and observability process 248, which may include computer executable instructions executed by the processor 220 (or independent processor of interfaces 210) to perform functions relating to the techniques described herein.

Specifically, according to various embodiments, an observability and assurance service, associated with various clusters of application services for an application that are executed in a data mesh, may configure a data compliance filter for a particular application service in one of the clusters of application services according to a data compliance policy. The observability and assurance service may monitor the data and traffic associated with the particular application service, wherein the data compliance filter is applied to the traffic to restrict sensitive data in the traffic from being processed by the particular application service. The observability and assurance service may make a determination that the data compliance policy has been violated by the particular application service. The observability and assurance service may modify, based on the determination, the data compliance filter for the particular application service.

Operationally, FIG. 5 illustrates an example architecture 500 for data compliance and observability, according to various embodiments. Architecture 500 may include data compliance and observability process 502. Data compliance and observability process 502 may be utilized to geofence data in heterogeneous cloud and edge environments at run time. Additionally, data compliance and observability process 502 may be utilized to provide observability into the access, processing, storage, and retention of the data, including data subject to sovereignty obligations and other compliance constraints. In various embodiments, data compliance and observability process 502 may execute these functions utilizing DCaC elements embedded in the program code of applications.

Architecture 500 may include compliance engine 306 and data compliance regulation repository 322. As described above, it may enable a data team and/or application manager 304 to select the target geographical area (e.g., state, country, or region, etc.), the target industry, and the data compliance rules that may be applicable at the organization level.

Different rules may apply to different types of data/categories of sensitive data. A mechanism may be utilized to differentiate between the different types of data/categories of sensitive data in order to selectively apply the different rules. As such, the data may be differentially classified and/or labeled.

For example, data discovery techniques, data categorization, classification, tagging may be utilized to differentiate between the various different types of data/categories of sensitive data. However, in various embodiments, the DCaC model may be leveraged in order to differentiate between the various different types of data/categories of sensitive data on the basis of annotations embedded in the program code of the application and/or their corresponding data associations.

Compliance engine 306 may use these annotations and the corresponding associations to enforce compliance rules on the data. In this manner, the DCaC model provides a

future-proof approach to compliance, since organizations using compliance engine 306 can quickly adjust to any ruleset, today and in the future, country by country, industry by industry, etc.

As described above, workload engines 332 may receive a request to deploy an application in a given region (e.g., from SRE/IT team 310). Workload engine 332 may issue a request to compliance engine 306, to compute and return the data compliance constraints that will be applied for that specific deployment. Alternatively, a data team and/or application manager 304 might select the compliance rules required, and a declarative intent for deploying the application might be issued from compliance engine 306 to one or more workload engines 332.

If the application deployment performed by the workload engines 332 is successful, several clusters 508 (e.g., a set of K8s or K3s clusters) may be spun up across various cloud and edge nodes. More specifically, such deployments might involve multi-cluster service meshes, which may run across multi-cloud and edge infrastructures 334. In various embodiments, clusters 508 may contain various pods 512 supporting the set of distributed services that compose the application (e.g., application services 514). Such services may comprise a service/data mesh 518, with instances running across various clouds and edges.

Clusters 508 and the application services 514 that they run may change over time. For instance, some workloads might be reallocated to another cloud or edge infrastructure (e.g., for cost or optimization reasons). Canary or alternative deployments might be instrumented to gradually transition to a new application or service version, which may change the traffic pattern of the data mesh 518. Security breaches might also produce deviations in traffic patterns. In addition, regulation might also evolve and change over time, creating new constraints to be enforced on the data mesh.

Overall, a data mesh 518 where application services are executed may change over time. These changes may impact where the data associated with these application services is processed, retained, and accessed. However, presently there are no existing mechanisms to dynamically adapt to changes in the data mesh 518 and/or changes in data compliance regulations that affect how data traffic in the data mesh should be processed (e.g., accessed, stored, forwarded, or otherwise used). In contrast, in various embodiments introduced herein, data compliance and observability process 502 is executable to provide visibility and/or automated mechanisms across multi-cloud and edge to dynamically adapt a data mesh 518 upon possible changes in regulation, the configuration of compliance rules, or the applications themselves.

These functionalities may be implemented, at least in part, using local observability and assurance service 510. In various embodiments, each cluster 508 may run an instance of local observability and assurance service 510. The set of local observability and assurance services 510 executing across the clusters 508 may work in concert both with the compliance engine 306 as well as with full stack observability tool 506. The interplay between local observability and assurance services 510, compliance engine 306, and full stack observability tool 506 may be accomplished through interface 522. This interplay may support the observability (e.g., data compliance observability 504) and dynamic data geofencing (e.g., geofenced data 520) functions on the data mesh 518, as described in greater detail below.

FIG. 6 illustrates an example architecture 600 for data compliance and observability according to various embodiments. The architecture 600 may include the set of clusters

508 that may run on multi-cloud and edge infrastructures **334**. The set of clusters **508** may include components constituting control and data planes. For instance, is each pod **512** within cluster **508** may be composed of a set of services **514** that implement the business logic of the application (e.g., application service **514** “Serv_1” in one pod **512** might implement part of this logic).

Architecture **600** may include an automated mechanism to decouple ancillary functions from the application business logic. This decoupling may be especially relevant for cloud/edge native applications, which are increasingly being implemented as a set of workloads composing a service mesh.

For example, a service mesh may include a set of application services **514** (e.g., Serv_1, . . . , Serv_n) that implement the business logic along with a set of interconnected sidecar proxies **604** associated to each one of those application services **514**. Sidecar proxies **604** (e.g., an Envoy proxy) may decouple ancillary functions from the application’s business logic, such as routing control, proxying and securing the communications between the services at the data plane level, collecting and reporting telemetry data on the service mesh, load balancing, managing retries, etc.

Control plane **606** may manage and configure sidecar proxies **604**, including, for example, the capacity to push routing rules to control and block traffic at run time. In some examples, control plane **606** may be an Istio control plane. The control plane **606** may be used for fine-grained traffic control for HTTP, gRPC, TCP, and other types of traffic. In some embodiments, the control plane **606** may also include controls for the gateways and/or other API gateways **602** that may support communications across the different clusters **508** composing the service mesh.

Control plane **606** and sidecar proxies **604** may also provide native functions to decouple the configuration, observability, and enforcement of data compliance rules from the application business logic. However, these capabilities are not limited to a service mesh nor to the use of sidecar proxies **604** and can be applied to any type of data mesh. For instance, the functions enabling observability and the enforcement of data compliance rules might be embedded in the sidecar proxies **604** themselves.

Alternatively, the decoupling of data compliance from the application’s business logic might be implemented without using sidecar proxies **604**. In such cases, this functionality might be embedded in other portable processes. For instance, the functionality may be achieved utilizing client libraries, daemons working in tandem with application-specific services, sandboxed programs in the OS kernel (e.g., using the Extended Berkeley Packet Filter or eBPF), or an agentless model, or by embedding such functionality in the Kubernetes control plane itself.

In various embodiments, the configuration, observability, and enforcement of data compliance rules, etc. functionalities may be decoupled from sidecar proxies **604** and implemented as local compliance controllers deployed within the service/data mesh **518** through a set of local observability and assurance services **510**. For example, each cluster **508** may include a local observability and assurance service **510** deployed together with the application pods **512**. This decoupled functionality may provide the portability and reuse of observability and enforcement of data compliance functions across not only different applications but also cloud and edge environments.

Local observability and assurance service **510** may act as a local control point enabling the enforcement of data geofencing and observability of data compliance at run time

(i.e., in real time). Such functionality may be implemented in various ways. For example, by leveraging filters (e.g., Envoy filters) and external logic interfacing with the sidecar proxies **604** based on an instruction format such as WebAssembly. This pluggable extension model may also facilitate the enforcement of data compliance rules and the collection of specific observations for traffic of interest in the service/data mesh **518**.

To this end, each local observability and assurance service **510** may interface both with the sidecar proxies **604** and control plane **606**. Sidecar proxies **604** may support the to collection of telemetry data with specific focus on data compliance, while control plane **606** may facilitate the configuration and enforcement of traffic policy on the sidecar proxies **604** and the API gateways **602**, thereby enabling data geofencing based on the data compliance rules defined. These controls on the service/data mesh **518** traffic may be exercised at run time. The controls may include hot reconfiguration and reloading of telemetry and data compliance rules.

In some embodiments, the traffic across the service/data mesh **518** may be encrypted beforehand, including the application-layer headers. In such cases, each sidecar proxy **604** may terminate, decrypt, enforce traffic controls, and either block or encrypt and forward the traffic to the next service in the mesh. Other embodiments may also use the local observability and assurance services **510**, but instead of interfacing with the sidecar proxies **604**, they may interface with client libraries, or daemons, or with filtering functionality supported by eBPF, or directly by the Kubernetes control plane, as described above.

The set of local observability and assurance services **510** in the clusters **508** may interface with observability and assurance engine **406**, which may be part of compliance engine **306**. While workload engine **332** may receive a set of data compliance requirements as an input to restrict (e.g., geofence) the placement of workloads at deploy time, the push and update of compliance-specific configuration to the local observability and assurance services **510** may be out of the scope of a workload engine **332**. Hence, each local observability and assurance service **510** may leverage interface **522** to register with observability and assurance engine **406** to receive compliance-specific configuration and updates.

Once a configuration phase (e.g., registration with observability and assurance engine **406**, etc.) is completed, local observability and assurance services **510** may process and send events to the observability and assurance engine **406** through interface **522** (e.g., capturing relevant observations locally, including data access requests that were blocked, potential attempts to deviate from the compliance rules configured, anomalies detected, etc.). Moreover, local observability and assurance services **510** may also operate and buffer events even when backhaul connectivity to observability and assurance engine **406** is lost. The continued operating and buffering functionalities of local observability and assurance services **510** may be utilized, for example, in many edge environments, where connectivity to the backends and/or the cloud might not be reliable or always available.

In various embodiments, local observability and assurance services **510** may interface with observability and assurance engine **406** to get configuration and updates, while it may interface with full stack observability tool **506**, such as Cisco’s Full Stack Observability solution, to push events yield by the observability function. This approach may be utilized to add data compliance implementation and/or

observation to application observability stacks. In additional embodiments, local observability and assurance services **510** may interface additionally or alternatively directly with observability and assurance engine **406**, which in turn may interface with full stack observability tool **506**.

The observability and assurance components may be implemented as a distributed function across multi-cloud and edge infrastructures **334**. For example, Local observability and assurance services **510** may interface both with a data collection platform **610** as well as with observability and assurance engine **406**. Data collection platform **610** may operate as a common element that gathers, formats, and feeds data to full stack observability tool **506**. Other data sources **608** may also supply data to data collection platform **610**. For instance, telemetry and events from other monitoring systems, such as Cisco's ThousandEyes, workload engines **332**, a service mesh manager, etc. may be supplied to data collection platform **610**.

Local observability and assurance services **510** may send events, logs, and other metrics related to data compliance to data collection platform **610**. As a result, the full stack observability tool **506** may add data compliance to the observability stack. Observability and assurance engine **406** may also interface with the full stack observability tool **506** to provide this functionality. Therefore, observability and assurance engine **406** may not only communicate the targeted state of data compliance policies for a given application to the full stack observability tool **506**, but it may also allow the full stack observability tool **506** to notify the observability and assurance engine **406** upon any deviations or potential issues based on observations and analysis made across the full stack.

Such notifications may be used by the observability and assurance engine **406** to dynamically update the configuration of data compliance rules in one, or several, of the local observability and assurance services **510** in the clusters **508**. Local observability and assurance services **510**, full stack observability tool **506**, and observability and assurance engine **406** may operate in closed-loop **524** mode, thereby enabling dynamic control of data compliance rules and data geofencing (e.g., geofenced data **520**) in real time.

FIG. 7 illustrates an example process **700** for data compliance and observability according to various embodiments. Process **700** may be executed utilizing the above-described observability and assurance components including the observability and assurance engine **406**, local observability and assurance service **510**, sidecar proxies **604**, data collection platform **610**, and full stack observability tool **506**, which may be cooperatively operated to provide data handling (e.g., data geofencing) and observability thereof according to predefined data compliance rules in real time both at intra- and inter-cluster level across hybrid cloud and edge infrastructures.

For example, local observability and assurance service **510** may be deployed in a service mesh. Once local observability and assurance service **510** is deployed, it may, at **714**, register with observability and assurance service **706**. Observability and assurance service **706** may be part of observability and assurance engine **406**.

Observability and assurance engine **406** may include a compliance target state service **704**, which may manage the data compliance rules and desired states for various organizations across different regions, concurrently. The targeted compliance state may include a target compliance configuration determined from the data compliance rules selected by data team and/or application managers and/or their

corresponding data handling constraints computed by the compliance engine (e.g., compliance engine **306** in FIG. 6).

At **716**, observability and assurance engine **406** may push a targeted compliance state to full stack observability tool **506**. For example, Observability and assurance engine **406** may push the targeted compliance state for each category of sensitive data that may be subject to the data compliance rules defined by a data team and/or application manager **304**, and the corresponding data handling constraints computed by the compliance engine **306** and sent to full stack observability tool **506**.

Full stack observability tool **506** may be configured to monitor resulting data meshes **702** of different categories of sensitive data for potential variations from their corresponding targeted compliance state. Additionally, at **718**, full stack observability tool **506** may confirm the reception and preparation to observe the resulting data meshes **702** and potential variations from their corresponding desired states.

At **720**, observability and assurance engine **406** may send the targeted compliance configuration to local observability and assurance service **510**. In turn, local observability and assurance service **510** may retain the targeted compliance states and, at **722**, push the corresponding targeted compliance configuration to a set of local services that may filter the data/service mesh for each category of sensitive data. For example, local observability and assurance service **510** may push targeted compliance states to sidecar proxies **604** or alternative models such as client libraries, daemons, etc.

Additionally, local observability and assurance service **510** may prepare and configure its observability module to send any relevant notifications (e.g., observations of data compliance policy violations, deviations from established patterns of behavior, suspicious activity, etc.) to data collection platform **610**. In various embodiments, local observability and assurance service **510** may monitor traffic associated with the application services and/or determine whether the observed data traffic violates a data compliance policy (e.g., deviates from an established pattern of behavior, violates the targeted compliance configuration for a category of sensitive data, exhibits suspicious traffic patterns, etc.).

The targeted compliance configurations pushed to the set of local services, at **722**, may be utilized to configure corresponding data compliance filters for an application service. Specifically, a data compliance filter may be configured to enforce the targeted compliance configurations (e.g., compliance rules) by filtering data traffic **710** corresponding to the category of sensitive data in the data mesh where application services for the application are executing. In various embodiments, the data compliance filter, once configured, may be applied to the data traffic **710** to restrict sensitive data in the data traffic **710** from being processed (e.g., accessed, stored, forwarded, or otherwise used) by particular application services (e.g., located in a particular geographic location).

In some examples, the data compliance filters may be Envoy filters. For example, a filter may be implemented in a filter chain. The filter may be associated with a listener that listens on a port associated with an application service. The filter may detect data traffic **710**. Data traffic **710** may include connections, requests, other data, etc. directed to the application services in the mesh.

The filter may then determine whether the detected data traffic **710** will violate one or more of the target data compliance configurations configured into the filter and/or if the detected data traffic **710** is allowed by the target data

compliance configurations. The filter may then make a routing decision based on the determination.

For example, if at **736** it is determined that data traffic **710** is allowed (e.g., “Yes”), then the filter may pass data traffic **710** on (e.g., to a next filter in the filter chain, to its destination application service, etc.) and/or report the allowance of data traffic **710** to local observability and assurance service **510**. In contrast, if at **736** it is determined that data traffic **710** is not allowed (e.g., “No”) then the filter may block, at **738**, data traffic **710** and/or may notify, at **740**, local observability and assurance service **510** that data traffic **710** was blocked.

At **724**, once the data compliance filters are configured, local observability and assurance service **510** may observe the traffic and data exchanged across the different application services in the data/service mesh. In various embodiments such observations may be received in the form of metadata **712**. For example, the metadata **712** may be collected by and/or received from tools, APIs, SDKs, etc. used to instrument, generate, collect, export, etc. telemetry data (metrics, logs, traces, etc.), such as OpenTelemetry/OpenTracing.

At **728**, upon detection of any suspicious behavior, variations from the targeted compliance states for any category of sensitive data, etc., local observability and assurance service **510** may notify full stack observability tool **506**, via data collection platform **610**, of the detection. In some examples, the notification may include compliance observations **708**. Compliance observations **708** may include relevant portions of the observed metadata **712**. Compliance observations **708** may also include notifications of blocked data traffic **710** that was reported to local observability and assurance service **510**.

At **730**, the notifications received by full stack observability tool **506** through data collection platform **610** may be sent to observability and assurance service **706**. At **732**, receiving the notification at observability and assurance service **706** may trigger observability and assurance service **706** to push new and/or modified data compliance rules to local observability and assurance service **510**.

The new and/or modified data compliance rules may be a new and/or modified set of targeted compliance states for a category of sensitive data. In various embodiments, the new and/or modified data compliance rules may include changes to the previous set of targeted compliance states that are designed to address and/or remedy a detected issue identified in the notification received at observability and assurance service **706**. This may include modifications to how data is handled (e.g., modifications to what type of data is included in a category of sensitive data, modification to how categories of sensitive data are handled by filters, modifications to which data is geofenced and/or how it is geofenced, modifications to which application services are allowed to handle specific types of data, etc.).

At **734**, the new and/or modified rules pushed to the local observability and assurance service **510** may be utilized to reconfigure the data compliance filters in the service/data mesh. For example, the filters may be modified so that they filter data traffic **710** according to the new and/or modified data compliance rules.

The techniques described with respect to **724**, **728**, **730**, **732**, and **734** may be utilized for performing data compliance observability and geofencing functions in real time while providing continuous assessment and a closed-loop **524** control of the data compliance state.

While the examples shown in FIG. **7** are mainly focused on observability and policy enforcement for data in transit

710, those skilled in the art would understand that the techniques described herein may also be applied to data at rest (e.g., to control data storage and retention policies for specific categories of data) as well as to data in use (e.g., for controlling data processing by an application service **514**, using processor **220** and memory **240** within device **200**).

FIG. **8** illustrates an example simplified procedure (e.g., a method) for automated data compliance and observability, in accordance with one or more embodiments described herein. For example, a non-generic, specifically configured device (e.g., device **200**), may perform procedure **800** by executing stored instructions (e.g., data compliance and observability process **248**). The procedure **800** may start at step **805**, and continues to step **810**, where, as described in greater detail above, an observability and assurance service, associated with various clusters of application services for an application that are executed in a data mesh, may configure a data compliance filter for a particular application service in one of the clusters of application services according to a data compliance policy. The data compliance policy may indicate a targeted data compliance state. For example, the data compliance policy may indicate the data handling constraints (e.g., targeted compliance states) that apply to a category of sensitive data for the associated application service.

In various embodiments, the observability and assurance service may take the form of one or more physical devices (e.g., device **200** in FIG. **2**) that are specially configured to perform the functions described herein. As would be appreciated by one of ordinary skill in the art, when executed in a distributed manner across multiple such devices, the collection of devices can be viewed as a singular device for performing these functions.

At step **815**, as detailed above, the observability and assurance service may monitor the data and traffic associated with the particular application service. In various embodiments, the data compliance filter is applied to the traffic to restrict sensitive data in the traffic from being processed (e.g., accessed, stored, forwarded, or otherwise used) by the particular application service. The data compliance filter may include and/or filter data traffic according to one or more data compliance rules that restrict a portion of the sensitive data from being processed by the particular application service based on a geographic location associated with the sensitive data and a geographic location associated with the particular application service (e.g., geofencing filters to restrict handling of portions of sensitive data to infrastructure located in geographical locations that comply with data compliance constraints for the portions of sensitive data). The data compliance filter may include a data traffic filter, utilizable within a data traffic filter chain, that detects the traffic associated with the particular application service arriving at a port associated with the particular application service and makes a routing decision for the traffic based on one or more data compliance rules.

The data traffic may be traffic from another application service within the cluster of application services. Additionally, or alternatively, the data traffic may be traffic sent from outside the cluster of application services (e.g., from another cluster in multi-cluster settings or from an external process with respect to the application).

At step **820**, the observability and assurance service may make a determination that the data compliance policy has been violated by the particular application service. Observability and assurance service may report the determination that the data compliance policy has been violated. For example, the violation may be reported to a data collection

platform, a full stack observability tool, and/or an observability and assurance engine. In various embodiments, the determination may indicate that a specific type of sensitive data in the traffic is being processed by the particular application service. In some examples, processing of this specific type of sensitive data by the particular application service may represent a violation of a data compliance policy.

At step **825**, as detailed above, observability and assurance service may modify, based on the determination, the data compliance filter for the particular application service. Modifying the data compliance filter may cause the specific type of sensitive data to be prevented from being processed by the particular application service. For example, where processing of this specific type of sensitive data by the particular application service may represent a violation of a data compliance policy, the data compliance filter may be modified to cause the specific type of sensitive data to be prevented from being processed by the particular application service. Procedure **800** then ends at step **830**.

It should be noted that while certain steps within procedure **800** may be optional as described above, the steps shown in FIG. **8** are merely examples for illustration, and certain other steps may be included or excluded as desired. Further, while a particular order of the steps is shown, this ordering is merely illustrative, and any suitable arrangement of the steps may be utilized without departing from the scope of the embodiments herein.

The techniques described herein, therefore, may be utilized to achieve data compliance in multi-cloud and edge infrastructures and may provide the capability to observe and geofence data both at the intra-cluster and inter-cluster level in real time. As such, data compliance and data sovereignty across geographical regions may be ensured utilizing these techniques. In addition, these techniques provide portability and reuse of observability and enforcement of data compliances functions across not only different applications and regions but also cloud and edge environments.

For example, the techniques described herein allow management of data compliance and its observability across multiple clusters by running a local observability and assurance system inside of a duster that registers with an observability and assurance agent outside of the cluster. In addition, utilizing the DCaC model, the techniques described herein allows the receipt of various data compliance rules associated to an application programmatically and automated configuration of a set of local observability and assurance services once they are deployed in the dusters. These techniques allow for the expression, communication, processing, and consumption of the desired state of data compliance rules for a given application.

In addition, the local observability and assurance service facilitates the automated extraction of telemetry from the application services and observation of the state and compliance of the workloads and their associated data. This feature is utilizable to provide notifications of deviations from the desired data compliance/residency rules to trusted entities, such as an observability and assurance engine or other data collection systems. In combination, these techniques allow for dynamic processing and enforcement configuration locally in a service/data mesh upon detection of deviations from the desired data compliance and/or data residency rules. For example, the techniques allow for observation of compliance deviations, updates on regulation, and/or identification of potential issues based on observations and analysis made across the full application stack.

As such, these techniques may be utilized to instrument a set of local observability and assurance services to enforce data geofencing both at intra- and inter-cluster level across hybrid cloud and edge infrastructures. Accordingly, these techniques allow for dynamic control of data compliance and enforcement of data geofencing rules in real time and across different applications, different geographical regions, and/or different cloud and edge environments

While there have been shown and described illustrative embodiments that provide automated data compliance and observability, it is to be understood that various other adaptations and modifications may be made within the spirit and scope of the embodiments herein. For example, while certain embodiments are described herein with respect to using the techniques herein for certain purposes, the techniques herein may be applicable to any number of other use cases, as well.

The foregoing description has been directed to specific embodiments. It will be apparent, however, that other variations and modifications may be made to the described embodiments, with the attainment of some or all of their advantages. For instance, it is expressly contemplated that the components and/or elements described herein can be implemented as software being stored on a tangible (non-transitory) computer-readable medium (e.g., disks/CDs/RAM/EEPROM/etc.) having program instructions executing on a computer, hardware, firmware, or a combination thereof. Accordingly, this description is to be taken only by way of example and not to otherwise limit the scope of the embodiments herein. Therefore, it is the object of the appended claims to cover all such variations and modifications as come within the true spirit and scope of the embodiments herein.

What is claimed is:

1. A method comprising:

configuring, by an observability and assurance service associated with various clusters of application services for an application that are executed in a data mesh, a data compliance filter for a particular application service in one of the various clusters of application services according to a data compliance policy that indicates a targeted compliance state for the various clusters of application services according by region;

monitoring, by the observability and assurance service, data and traffic associated with the particular application service, wherein the data compliance filter is applied to the traffic to restrict sensitive data in the traffic from being processed by the particular application service;

making, by the observability and assurance service, a determination that the data compliance policy has been violated by the particular application service; and

modifying, by the observability and assurance service and based on the determination, the data compliance filter for the particular application service.

2. The method as in claim 1, wherein the traffic is from another application service within the various clusters of application services.

3. The method as in claim 1, wherein the traffic is sent from outside the various clusters of application services.

4. The method as in claim 1, wherein the data compliance filter comprises one or more data compliance rules that restrict a portion of the sensitive data from being processed by the particular application service based on a geographic location associated with the sensitive data and a geographic location associated with the particular application service.

- 5. The method as in claim 1, further comprising: reporting, by the observability and assurance service, the determination that the data compliance policy has been violated.
- 6. The method as in claim 1, wherein the data compliance filter is based in part on an annotation in program code for the application that indicates a category of sensitive data.
- 7. The method as in claim 1, wherein the determination indicates that a specific type of sensitive data in the traffic is being processed by the particular application service.
- 8. The method as in claim 7, wherein modifying the data compliance filter causes the specific type of sensitive data to be prevented from being processed by the particular application service.
- 9. The method as in claim 1, wherein the data compliance filter is a data traffic filter that:
 - detects the traffic associated with the particular application service arriving at a port associated with the particular application service; and
 - makes a routing decision for the traffic based on one or more data compliance rules.
- 10. An apparatus, comprising:
 - one or more network interfaces;
 - a processor coupled to the one or more network interfaces and configured to execute one or more processes; and
 - a memory configured to store a process that is executable by the processor, the process when executed configured to:
 - configure, by an observability and assurance service associated with various clusters of application services for an application that are executed in a data mesh, a data compliance filter for a particular application service in one of the various clusters of application services according to a data compliance policy that indicates a targeted compliance state for the various clusters of application services according by region;
 - monitor data and traffic associated with the particular application service, wherein the data compliance filter is applied to the traffic to restrict sensitive data in the traffic from being processed by the particular application service;
 - make a determination that the data compliance policy has been violated by the particular application service; and
 - modify, based on the determination, the data compliance filter for the particular application service.
- 11. The apparatus as in claim 10, wherein the traffic is from another application service within the various clusters of application services.

- 12. The apparatus as in claim 10, wherein the traffic is sent from outside the various clusters of application services.
- 13. The apparatus as in claim 10, wherein the data compliance filter comprises one or more data compliance rules that restrict a portion of the sensitive data from being processed by the particular application service based on a geographic location associated with the sensitive data and a geographic location associated with the particular application service.
- 14. The apparatus as in claim 10, wherein the process when executed is further configured to:
 - report, by the observability and assurance service, the determination that the data compliance policy has been violated.
- 15. The apparatus as in claim 10, wherein the data compliance filter is based in part on an annotation in program code for the application that indicates a category of sensitive data.
- 16. The apparatus as in claim 10, wherein the determination indicated that a specific type of sensitive data in the traffic is being processed by the particular application service.
- 17. The apparatus as in claim 16, wherein the process when executed is configured to modify the data compliance filter by causing the specific type of sensitive data to be prevented from being processed by the particular application service.
- 18. A tangible, non-transitory, computer-readable medium storing program instructions that cause a device to execute a process comprising:
 - configuring, by an observability and assurance service associated with various clusters of application services for an application that are executed in a data mesh, a data compliance filter for a particular application service in one of the various clusters of application services according to a data compliance policy that indicates a targeted compliance state for the various clusters of application services according by region;
 - monitoring, by the observability and assurance service, data and traffic associated with the particular application service, wherein the data compliance filter is applied to the traffic to restrict sensitive data in the traffic from being processed by the particular application service;
 - making, by the observability and assurance service, a determination that the data compliance policy has been violated by the particular application service; and
 - modifying, by the observability and assurance service and based on the determination, the data compliance filter for the particular application service.

* * * * *