(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
08 June 2023 (08.06.2023)

WIPO | PCT

(10) International Publication Number
WO 2023/101736 A1

(51) International Patent Classification:
G06F 21/60 (2013.01) *H04L 9/40* (2022.01)
G06F 21/62 (2013.01)

(21) International Application Number:
PCT/US2022/042323

(22) International Filing Date:
01 September 2022 (01.09.2022)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
63/264,829 02 December 2021 (02.12.2021) US
17/652,461 24 February 2022 (24.02.2022) US

(71) Applicant: **MICROSOFT TECHNOLOGY LICENSING, LLC.** [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).

(72) Inventors: **GARG, Shikha**; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **SCHOLL, Daniel, Keith**; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **UMAY, Mehmet Kadri**; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **SIDDIQUE, Imran**; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **PATEL, Nayana Singh**; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US).

(54) Title: HYBRID CLOUD DATA CONTROL

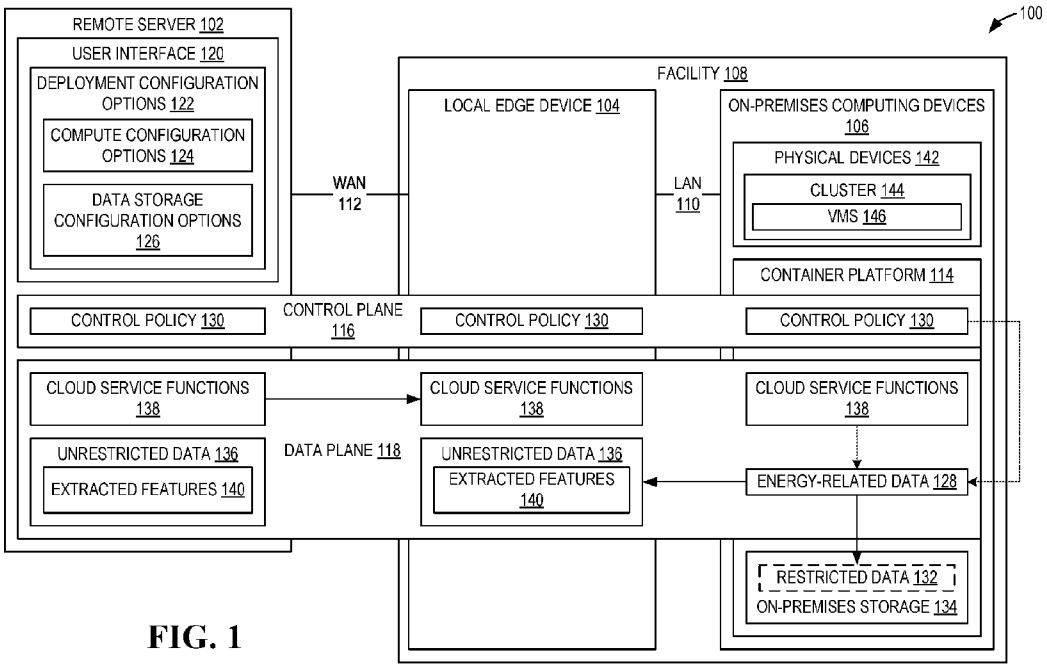

FIG. 1

(57) **Abstract:** A remote server computing system is configured to present a user interface with a plurality of deployment configuration options including compute configuration options and data storage configuration options for energy-related data within a hybrid cloud environment. The hybrid cloud environment comprises a cloud-service-managed control plane and a data plane utilizing local compute resources and storage. A data control policy is generated that provides cloud-service-managed governance over at least a portion of the data plane. The control plane is configured to enforce the data control policy by subjecting at least a portion of the energy-related data to a data transmission restriction or a local storage restriction. The data plane is used to deploy one or more cloud service functions configured to process at least the portion of the energy-related data and output one or more extracted features from at least the portion of the energy-related data to the data plane.

WO 2023/101736 A1

[Continued on next page]

(74) **Agent: CHATTERJEE, Aaron, C.** et al.; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US).

(81) **Designated States** *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CV, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IQ, IR, IS, IT, JM, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

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

Declarations under Rule 4.17:

— *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

— *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

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

HYBRID CLOUD DATA CONTROL

BACKGROUND

Energy companies can generate large amounts of data from such activities as energy exploration, production, transport, and usage. Such data is often processed and stored using hardware and software that is specialized to applications within the energy industry.

SUMMARY

Examples are disclosed that relate to computing devices and methods for providing cloud-service-managed governance in a hybrid cloud environment. One example provides a remote server computing system, comprising a processor and a memory storing instructions executable by the processor. The instructions are executable to present a user interface with a plurality of deployment configuration options including compute configuration options and data storage configuration options for energy-related data within a hybrid cloud environment. The hybrid cloud environment comprises a cloud-service-managed control plane and a data plane utilizing local compute resources and storage located on-premises at an energy production or distribution facility. The cloud-service-managed control plane and the data plane span the remote server computing system, a local edge computing device, and the local compute resources and storage. A user input of one or more of the deployment configuration options is received. A data control policy is generated that provides cloud-service-managed governance over at least a portion of the data plane using the one or more user-input deployment configuration options. The data control policy is provided to the local edge computing device and the local compute resources via the cloud-service-managed control plane. The cloud-service-managed control plane is configured to enforce the data control policy by subjecting at least a portion of the energy-related data to a data transmission restriction or a local storage restriction. The data plane is used to deploy one or more cloud service functions to the local compute resources. The one or more cloud service functions are configured to process at least the portion of the energy-related data and output one or more extracted features from at least the portion of the energy-related data to the data plane. The remote server computing system is further configured to receive, via the data plane, the one or more extracted features from at least the portion of the energy-related data.

Another aspect provides a computing device comprising a processor and a memory storing instructions executable by the processor. The instructions are executable to receive a data control policy from a remote server computing system via a cloud-service-managed control plane. The data control policy provides cloud-service-managed governance over at least a portion of a data plane utilizing local compute resources and storage located on-premises at an energy production or distribution facility. Energy-related data is received. The data control policy is enforced by

subjecting at least a portion of the energy-related data to a data transmission restriction or a local storage restriction. One or more cloud service functions are received from the remote server computing system via the data plane. At least a portion of the energy-related data is processed using the one or more cloud service functions to extract one or more features from at least the portion of the energy-related data. The one or more extracted features are output from at least the portion of the energy-related data to the remote server computing system via the data plane.

This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used to limit the scope of the claimed subject matter. Furthermore, the claimed subject matter is not limited to implementations that solve any or all disadvantages noted in any part of this disclosure.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 shows a schematic view of an example of a hybrid cloud environment.

FIG. 2 shows an example of a graphical user interface (GUI) that can be implemented to configure the hybrid cloud environment of FIG. 1.

FIG. 3 shows additional example features of the GUI of FIG. 2.

FIG. 4 shows an example policy template menu in the GUI of FIG. 2.

FIG. 5 shows an example of a regulatory data control policy template in the GUI of FIG. 2.

FIG. 6 shows another example of a regulatory data control policy template in the GUI of FIG. 2.

FIG. 7 shows an example of a high-latency network data control policy template in the GUI of FIG. 2.

FIG. 8 shows an example of a high-data-volume data control policy template in the GUI of FIG. 2.

FIG. 9 shows an example of an intermittent connection data control policy template in the GUI of FIG. 2.

FIG. 10 shows another example of a hybrid cloud environment including an example deployment of a machine learning function.

FIG. 11 shows another example of a hybrid cloud environment including an example search engine deployment.

FIGS. 12A-12B show a block diagram of an example method for providing cloud-service-managed governance in a hybrid cloud environment.

FIG. 13 shows a block diagram of an example computing system.

DETAILED DESCRIPTION

As introduced above, energy companies can produce large volumes of data during activities such as energy exploration, energy production, energy transport, and/or usage. Such data is often

processed and stored using hardware and software that is domain specific. For example, the hardware and software may be custom-built or specialized for applications within the energy industry. However, it can be challenging to provide this data to other computing devices (e.g., a server hosting a cloud service), and to integrate services provided by such computing devices (e.g., data processing, storage, and policy management) with a company's existing domain-specific computing infrastructure.

In some examples, energy-related data that is produced at an energy production facility can be considered sensitive and subject to export controls. For example, seismic data collected as part of oil exploration may be subject to data residency rules in a jurisdiction where the data is collected. Accordingly, per the rules and regulations of the host jurisdiction, the seismic data may not be transmitted to a cloud server that is not located within the borders of that jurisdiction.

Additionally, and as introduced above, energy-related datasets can be large. For example, seismic data can occupy petabytes of storage. However, delays in accessing such data (e.g., due to network latency) can be detrimental for visualization and other compute workloads. In some areas, such as remote locations with little communication infrastructure, it can be challenging to transmit and receive data with low latency (e.g., less than 50 ms) from a cloud data center. Similarly, it can be challenging to transmit data in real time to the cloud due to the large volume of data. For example, an energy facility may generate 10-1000 GB of data per second, which can surpass bandwidth availability, thereby preventing real-time data transmission. Further, some energy facilities operate in remote locations, which may be far offshore, deep underwater, or a deserted region of land. In some examples, these facilities are not connected to the internet or the cloud, or connect intermittently (e.g., once per month).

In some examples, existing energy-related facilities attempt to address some of the above-identified issues by providing cloud-independent data storage and computing resources on premises. However, such cloud-independent on-premises deployments of computing resources has led to the development of a fragmented ecosystem of data types, data formats, tools and services, each being siloed at the facility of the energy producer. This fragmentation tends to increase the cost and complexity of storing and processing energy data, increase the latency of transmitting and processing energy data, and potentially force companies to create custom tools for converting and interfacing between different data types, formats, applications, and devices, increasing the cost and development time for projects. On-premises, cloud-independent computer systems are especially challenging for integrated energy companies that engage in a combination of upstream, midstream, and downstream activities, and/or energy companies that utilize a variety of energy sources.

The present disclosure addresses the dual needs for on-premises storage and efficient central management. The example computing systems described herein relate to providing cloud-service-managed governance in a hybrid cloud environment comprising a cloud-service-managed control plane and a data plane utilizing local compute resources and storage located on-premises at an energy production or distribution facility. Briefly, a user interface is presented with a plurality of deployment configuration options including compute configuration options and data storage configuration options for energy-related data within the hybrid cloud environment. A data control policy is generated that provides cloud-service-managed governance over at least a portion of the data plane using the one or more user-input deployment configuration options. The data control policy is provided to a local edge computing device and the local compute resources via the cloud-service-managed control plane, and the cloud-service-managed control plane is configured to enforce the data control policy by subjecting at least a portion of the energy-related data to a data transmission restriction or a local storage restriction. The data plane is used to deploy one or more cloud service functions to the local compute resources. The one or more cloud service functions are configured to process at least the portion of the energy-related data and output one or more extracted features from at least the portion of the energy-related data to the data plane.

In this manner, local compute and storage resources (e.g., hardware that was purchased and previously managed by an energy company) are managed by a cloud service provider, providing uniformity across the data plane. Cloud service functions implemented at a remote cloud server can also be deployed to the network edge and local resources, allowing clients to specify what workloads are run at the cloud, the edge, or locally, based on latency, bandwidth, sovereignty, connectivity, or other constraints. New configurations, applications, and updates are also easily deployable in a uniform manner across the hybrid cloud environment.

FIG. 1 shows one example of a hybrid cloud environment 100. The hybrid cloud environment 100 comprises a remote server computing system 102, a local edge computing device 104, and local compute resources and storage in the form of one or more computing devices 106 located on-premises at an energy production or distribution facility 108. In some examples, the facility 108 comprises an exploratory oil well site. Other examples of facilities include, but are not limited to, an offshore platform, a pipeline, a wind farm, a solar farm, a terminal, a power station, and a facility where energy is consumed (e.g., a data center, a manufacturing facility, or a residential development).

In the example depicted in FIG. 1, the local edge computing device 104 and the one or more computing devices 106 are located on-premises at the facility 108, and the remote server 102 is located at a different facility. As one non-limiting example, the local edge computing device 104 and the one or more computing devices 106 may be located at a first facility, such as an oil

exploration facility, in a first jurisdiction, such as Nigeria, and the remote server 102 may be located at a data center in a second jurisdiction, such as the United States. In other examples, the local edge computing device 104 and the one or more computing devices 106 are located at different facilities. For example, the one or more computing devices 106 may be located at an offshore platform, and the local edge computing device 104 is located on another vessel (e.g., another platform or a supply/transport ship).

The local edge computing device 104 and the one or more computing devices 106 are connected by a local area network (LAN) 110. The local edge computing device 104 is connected to the remote server 102 by a wide area network (WAN) 112, such as the internet. The local edge computing device 104 is a computing device having a position on a network topology between the LAN 110 and the WAN 112. Advantageously, and as described in more detail below, the local edge computing device 104 provides at least a portion of the functionality of the remote server 102 at the local network level, such as providing cloud-service-managed data governance, storage, and processing.

It will also be appreciated that, in some examples, aspects of the local edge computing device 104, the on-premises computing devices 106, and/or the remote server 102 are provided by two or more devices. For example, aspects of the remote server 102 may be implemented at a plurality of servers at one or more data centers. The on-premises computing devices 106 may comprise a cluster of a plurality of connected computing devices. It will also be appreciated that, in some examples, the local edge computing device 104, the on-premises computing devices 106, and/or the remote server 102 comprise real devices, virtual devices, or a combination of real and virtual devices. In some examples, the local compute resources and storage comprise one or more physical computing devices 142 running a cluster 144 comprising a plurality of virtual machines (VMs) 146.

In some examples, aspects of the remote server 102 and/or the local edge computing device 104 are provided on a software-as-a-service (SaaS) basis or a hardware-as-a-service (HaaS) basis, in which the underlying computing hardware is housed and managed by a service provider. In contrast, the one or more on-premises computing devices 106 comprise hardware that was originally independent of the cloud service ecosystem, such as legacy hardware owned by an energy company, housed and managed at the facility 108. Advantageously, such hardware is integrated into the hybrid cloud environment 100, providing uniform management and preventing fragmentation of the computing ecosystem.

The one or more on-premises computing devices 106 comprise a container platform 114, such as KUBERNETES. The container platform 114 comprises a system for deploying, scaling, and managing containerized applications. The remote server 102 and the local edge device 104 are

compatible with the container platform 114, thus the container platform 114 is a foundation for applications and services to be deployed across the hybrid cloud environment 100.

The remote server computing system 102, the local edge computing device 104, and the one or more on-premises computing devices 106 each have a processor and associated memory storing instructions executable by the processor. Additional details regarding the components of the remote server computing system 102, the local edge computing device 104, the one or more on-premises computing devices 106, the LAN 110, and the WAN 112 are provided below with reference to FIG. 13.

The hybrid cloud environment 100 also comprises a cloud-service-managed control plane 116 and a data plane 118 utilizing the local compute resources and storage provided by the one or more on-premises computing devices 106. The control plane 116 controls how data is routed within the hybrid cloud environment 100. For example, data routing table(s) are created in the control plane 116 that define how data is routed between the server 102, the edge device 104, and/or the on-premises computing device(s) 106. The data plane 118 comprises the data that is routed between the server 102, the edge device 104, and/or the on-premises computing device(s) 106, following one or more paths established at the control plane 116. The cloud-service-managed control plane 116 and the data plane 118 span the remote server computing system 102, the local edge computing device 104, and the local compute resources and storage provided by the one or more on-premises computing devices 106. In this manner, the cloud-service-managed control plane 116 and the data plane 118 facilitate data transmission between the server 102, the edge device 104, and/or the on-premises computing device(s) 106.

The remote server computing system 102 is configured to present a user interface 120 with a plurality of deployment configuration options 122. In some examples, and as described in more detail below with reference to FIGS. 2-8, the user interface is a graphical user interface (GUI). In other examples, the user interface comprises any other suitable type of interface. Another example of a suitable interface is a command line/text interface.

FIGS. 2-9 show an example of a GUI 200 for configuring a hybrid cloud environment, such as the hybrid cloud environment 100 of FIG. 1. In some examples, the GUI 200 serves as the user interface 120 of FIG. 1.

The GUI 200 comprises a first “RESOURCES” tab 202, a second “EXTENSIONS” tab 204, and a third “POLICY” tab 206. With reference now to FIG. 2, the “RESOURCES” tab 202 is selected, as indicated by dashed box 208. Accordingly, the GUI 200 displays a resource page 210 comprising a table 212. The table 212 comprises a “RESOURCE” column 214 listing a plurality of computing resources making up the hybrid cloud environment, and a “STATUS” column 216 indicating a status of each computing resource. In the example depicted in FIG. 2, the hybrid cloud

environment comprises a cloud server 218, an edge device 220, and a first-party computing device 222 located on-premises at an energy-related facility.

The table 212 further comprises a status indicator 224 for each of the resources in the column 214 configured to provide a graphical representation of a status of each resource. For example, the status indicator 224 may be colored green to indicate that the resource is connected to a cloud service platform, amber to indicate that the resource is configured, but not presently connected, and red to indicate a configuration error. The column 214 further comprises a text indicator 226 of the status (e.g., "CONNECTED"). In the example depicted in FIG. 2, the cloud server 218, the edge device 220, and the first-party computing device 222 are each connected to the cloud service platform.

The GUI 200 further comprises an "ADD" button 228 that is user-selectable to add a new resource to the hybrid cloud environment. An "EDIT" button 230 is user-selectable to edit a name or one or more configuration parameters of a computing resource. In addition, the GUI 200 comprises a "DELETE" button 232 that is user-selectable to remove one or more computing resources from the hybrid cloud environment. Each resource may have an associated IP address, and in one example, a routing table of IP addresses may be generated for use as the control policy 130 discussed below, based on the user input into in GUI 200.

FIG. 3 shows the GUI 200 with the "EXTENSIONS" tab 204 selected, as indicated by the dashed box 208. The GUI 200 is displaying an extensions page 234 comprising an extension table 236. The extension table 236 comprises an "EXTENSION" column 238 listing one or more applications that can be deployed in the hybrid cloud environment, and a "STATUS" column 240 indicating a status of each computing resource. In the example depicted in FIG. 3, the extensions include a containerized application 242 implemented at the first-party computing device 222 using a cloud-service-compatible container platform (e.g., the container platform 114 of FIG. 1). The extensions further include a cloud service application 244 implemented at the cloud server 218 and the edge device 220. By integrating the containerized application 242 and the cloud service application 244 in the extensions table 236, the extensions are deployable at any location in the hybrid cloud environment (e.g., at the cloud server 218, the edge device 220, and/or the first-party computing device 222).

The extension table 236 further comprises a status indicator 246 for each of the extensions in the column 238 configured to provide a graphical representation of a status of each resource. The column 214 further comprises a text indicator 248 of the status (e.g., "ENABLED"). In the example depicted in FIG. 3, the containerized application 242 and the cloud service application 244 are enabled for use in the hybrid cloud environment.

The GUI 200 further comprises, in the extensions page 234, an "ADD" button 250 that is user-selectable to add a new extension to the hybrid cloud environment. An "EDIT" button 252 is user-selectable to edit a name or one or more configuration parameters of an extension. In addition, the GUI 200 comprises a "DELETE" button 254 that is user-selectable to remove one or more extensions from the hybrid cloud environment.

With reference again to FIG. 1, the remote server computing system 102 is configured to receive a user input of one or more of the deployment configuration options 122. As described in more detail below with reference to FIGS. 4-8, the deployment configuration options 122 include compute configuration options 124 and data storage configuration options 126 for energy-related data 128 within the hybrid cloud environment 100. A data control policy 130 is generated that provides cloud-service-managed governance over at least a portion of the data plane 118 using the one or more user-input deployment configuration options. The data control policy 130 is provided to the local edge computing device 104 and/or the local compute resources via the cloud-service-managed control plane 116, and may include a routing table of IP addresses of the local edge device 104, on-premises computing devices 106, and remote server 102, for example. The cloud-service-managed control plane 116 is configured to enforce the data control policy 130 by subjecting at least a portion 132 of the energy-related data 128 to a data transmission restriction or a local storage restriction.

As described in more detail below, the data transmission restriction or local storage restriction may be implemented as a separate routing table in the control policy 130 that is different than a routing table that is applied to unrestricted data. In this manner, the routing tables are configured to specify allowed data paths for restricted data and unrestricted data in accordance with the user-input deployment configuration options 122. For example, an unrestricted data routing table may allow the unrestricted data 136 to be transmitted between any of the remote server 102, the edge device 104, and the one or more on-premises computing devices 106. However, a restricted data routing table may allow the restricted data 132 to be transmitted between any of the one or more on-premises computing devices 106, but not to the remote server 102 and/or the local edge device 104.

With reference now to FIGS. 4-9, the "POLICY" tab 206 in the GUI 200 is selected, as indicated by the dashed box 208. Accordingly, the GUI 200 displays a plurality of deployment configuration options, including compute configuration options and data storage configuration options that can be used to formulate the data control policy 130 of FIG. 1.

In some examples, the deployment configuration options are provided in the form of a policy template. As depicted in FIG. 4, the GUI 200 comprises a policy template menu 256. The policy template menu 256 comprises a plurality of policy templates, including a regulatory data control

template 258, a high-latency template 260, a high-volume template 262, and an intermittent connection template 264. The menu 256 also comprises a blank template 266 labeled "NEW CONFIGURATION" for a user to provide a new, custom set of configuration options.

In some examples, and as introduced above, some energy-related data (e.g., seismic data) is considered sensitive data and is subjected to data residency rules in certain jurisdictions. Accordingly, the regulatory data control template 258 is populated with deployment configuration options that comply with such data residency and control rules. These may, for example, be implemented via a restricted routing table with only on-premises IP address entries as discussed above.

FIG. 5 shows the regulatory data control template 258. As shown in FIG. 5, the GUI 200 comprises a menu of user-selectable storage and transmission restriction options 268, and a graphical representation of the cloud device 218, edge device 220, and first-party device 222. The cloud device 218 can serve as the remote server 102 of FIG. 1, the edge device 220 can serve as the local edge device 104 of FIG. 1, and the first-party device 222 can serve as the one or more on-premises computing devices 106 of FIG. 1.

The menu 268 includes a user-selectable storage icon 270, which is configured to be selected and placed inside one or more of the graphical representations of the cloud device 218, the edge device 220, and/or the first-party device 222 to specify one or more locations at which data storage is permitted. A user-selectable storage restriction icon 272 is configured to be selected and placed inside one or more of the graphical representations of the cloud device 218, the edge device 220, and/or the first-party device 222 to specify one or more locations at which data is not permitted to be stored. Data transmission icons 274A and 274B are configured to be placed between two of the graphical representations of the cloud device 218, the edge device 220, and/or the first-party device 222 to specify between which devices, and in which direction, data is permitted to be transmitted. Transmission restriction icons 276A and 276B are configured to be placed between two of the graphical representations of the cloud device 218, the edge device 220, and/or the first-party device 222 to specify between which devices, and in which direction, data transmission is not permitted. The menu 268 further comprises a processing icon 278 configured to be selected and placed inside one or more of the graphical representations of the cloud device 218, the edge device 220, and/or the first-party device 222 to specify one or more locations at which data processing occurs.

In the regulatory template 258, energy-related data is stored on one or more storage devices located within a geographic area 280 (e.g., within a country's borders). For example, as shown in FIG. 5, data is permitted to be stored on the edge device 220 and the first-party device 222, which are located within the geographic area 280. The energy-related data is also subject to a data

transmission restriction within the geographic area. For example, as shown in FIG. 5, data is permitted to be transmitted between the edge device 220 and the first-party device 222, but is not permitted to be transmitted out of the geographic area 280 to the cloud server 218. In some examples, data is permitted to be transmitted from the cloud server 218 to the edge device 220, for example to provide core platform and policy management services to devices within the geographic area 280. The geographic area restriction may be implemented by restricting transfer to IP addresses within the geographic area using a routing table, for example.

With reference now to FIG. 6, the cloud server 218 is optionally located in a sovereign cloud. The regulatory compliance template 258 permits access to the energy-related data via the sovereign cloud 282. For example, as shown in FIG. 5, data is permitted to be transmitted between the edge device 220 and the cloud device 218, and the data is permitted to be stored on the cloud device 218 when the cloud device 218 is in the sovereign cloud. For example, and with reference again to FIG. 1, based on determining that the remote server computing system 102 is located in a sovereign cloud, the remote server 102 is configured to receive, via the data plane 118, at least the portion 132 of the energy-related data 128 subject to the data transmission restriction or the local storage restriction.

With reference again to FIG. 5, the edge device 220 comprises a high-performance edge server that is connected to the first party device 222 via a high-throughput network. Accordingly, the regulatory template 258 indicates that data processing occurs at the edge device 220. In this manner, the edge device 220 provides data storage, API access, and processing services (e.g., artificial intelligence (AI) services) in the cloud environment.

In some examples, and as introduced above, network latency is detrimental for data visualization and other compute workloads. For example, seismic data can have a size on the order of petabytes, and visualization and other jobs lag when network latency is greater than a threshold latency (e.g., 50 ms). Accordingly, and as illustrated in the example of FIG. 7, the high-latency template 260 is populated with deployment configuration options that mitigate the effects of high-latency network(s).

The high-latency template 260 is implemented based upon determining that a network connection between the remote server computing system and the local edge computing device has a latency that is greater than or equal to a threshold latency. The threshold latency may comprise any suitable latency defined by the user. In some examples, the threshold latency is 50 ms.

As shown in FIG. 7, the high-latency template 260 specifies that the energy-related data is processed at the local edge computing device 220, as the latency of the network connection between the remote server computing system and the local edge computing device is too high for

applications to be executed in real time using the energy-related data. There are no data storage restrictions.

As indicated in FIG. 7, at least a portion of data processing optionally occurs at the cloud 218. For example, data can be streamed over the high-latency network and stored or cached at the cloud device 218 for non-real-time processing (e.g., AI training and inference).

With reference now to FIG. 8, and as introduced above, in some examples, large amounts of data are generated at the edge device in the range of 10-1000 GB/second, for each data collection site. This surpasses real-time network bandwidth between the edge device 220 and the cloud 218, such that the cloud 218 cannot process the data in real time. Accordingly, the user interface 200 is populated with the high-data-volume template 262 based on determining that the data volume of the energy-related data is greater than or equal to a threshold data volume. The threshold data volume may comprise any suitable latency. In some examples, the threshold data volume is in the range of 10-15 terabytes (TB) streaming data per day per site.

Accordingly, the high-data-volume template 262 specifies that the energy-related data is stored and processed at the local edge computing device, and is not stored at the cloud device 218. In other examples, the data is consolidated at the edge 220 before being sent to the cloud 218 in non-real time (e.g., over the network or via an offline databox transfer device). In this manner, some data processing and storage can take place at the cloud 218 (e.g., AI and other high-performance computing applications).

With reference now to FIG. 9, and as introduced above, some facilities operate in remote locations that are not connected to the Internet or the cloud 218 all the time. Based upon determining that a network connection between the remote server computing system and the local compute resources is intermittent, the user interface 200 is populated with intermittent connection template 264. The intermittent connection template 264 specifies that the energy-related data is cached and processed at the local compute resources and storage (e.g., at the first-party device 222). Data can be transmitted to the edge 220 and/or cloud 218 when a network connection becomes available, or via an offline databox transfer device. Other than caching, the data can be stored at any available location. It will also be appreciated that, in other examples, the energy-related data is cached and processed at the local edge device 220 when a high-throughput connection is available between the edge 220 and the first-party device 222, but not between the edge 220 and the cloud 218.

With reference again to FIG. 1, The local edge computing device 104 and/or the local compute resources are configured to receive the data control policy 130 from the remote server computing system 102 via the cloud-service-managed control plane 116. In some examples, the data control policy 130 comprises one or more files in a distributed version control repository (e.g., a Git repository). By housing the data control policy 130 in the distributed version control repository,

modifications to the data control policy 130, updates, and rollbacks can be implemented uniformly across the hybrid cloud environment 100.

In the example of FIG. 1, the one or more on-premises computing devices 106 receive the energy-related data 128. In some examples, the energy-related data 128 is received from one or more sensors, such as geophones, temperature sensors, or pressure sensors. In other examples, the energy-related data 128 is generated at the one or more on-premises computing devices 106.

The one or more on-premises computing devices 106 are configured to enforce the data control policy 130 by subjecting at least a portion of the energy-related data 128 to a data transmission restriction or a local storage restriction. It will also be appreciated that, in other examples, the energy-related data 128 is additionally or alternatively received by the local edge device 104. In such examples, enforcement of the data control policy 130 and following processing are implemented at the edge device 104.

When a subset of the energy-related data 128 is subject to the data transmission restriction or the local storage restriction, the restricted subset 132 of the energy-related data 128 is stored at an on-premises storage device 134, and is not transmitted to the remote server 102 or the local edge device 104. In contrast, when a subset of the energy-related data 128 is not subject to the data transmission restriction or the local storage restriction, the unrestricted portion 136 of the energy-related data 128 is provided to the remote server 102 and the local edge device 104 via the data plane 118.

The data plane 118 is also used to deploy one or more cloud service functions 138 to the local compute resources. The local edge device 104 and/or the one or more on-premises computing devices 106 are configured to receive the one or more cloud service functions 138 from the remote server computing system 102 via the data plane 118. In this manner, the one or more cloud service functions 138 can be executed at the local edge device 104 and/or the one or more on-premises computing devices 106.

The one or more cloud service functions 138 are configured to process at least the portion 132 of the energy-related data 128 that is subject to the data transmission restriction or the local storage restriction. In this manner, and as described in more detail below, the restricted data 132 is processed at on-premises at the facility 108, enabling compliance with regulatory data controls and/or preventing data processing delays related to network latency, bandwidth, and connectivity.

In the example of FIG. 1, the one or more on-premises computing devices 106 are configured to process at least the restricted portion 132 of the energy-related data 128 using the one or more cloud service functions 138 to extract one or more features 140 from at least the restricted portion 132 of the energy-related data 128. In some examples, the one or more cloud service functions 138 are configured to process the restricted portion 132 of the energy-related data 128 at the one

or more on-premises computing devices 106 and/or the local edge device 104, and any unrestricted data 136 is processed at the remote server 102. In this manner, at least a portion of the data processing can be offloaded to the remote server 102.

In other examples, all of the energy-related data 128 is processed locally at the one or more on-premises computing devices 106 and/or the local edge device 104. In this manner, the energy-related data can be processed using the cloud service functions 138 even when restricted to the local facility 108 by the control policy 130.

In some examples, the one or more features 140 extracted from the energy-related data 128 comprise inferences drawn from at least a portion of the energy-related data 128 (e.g., using AI). In other examples, the one or more features 140 comprise any other suitable data features. Some examples of suitable features include, but are not limited to, search results output in response to a search query, descriptive statistics, data visualizations, and a reduced-dimensionality version of the restricted data 132. Additional details regarding the data processing are provided in more detail below with reference to FIGS. 10-11.

FIG. 10 shows another example of a hybrid cloud environment 1000 for processing energy-related data 1002. The hybrid cloud environment 1000 includes a remote server 1004. The remote server 1004 can serve as the remote server 102 of FIG. 1. The hybrid cloud environment 1000 further comprises one or more on-premises computing devices 1006. The one or more on-premises computing devices 1006 can serve as the one or more on-premises computing devices 106 of FIG. 1.

The remote server 1004 is configured to receive energy-related training data 1008, including a plurality of energy-related training data pairs 1010. Each energy-related training data pair 1010 includes, as input, remote data 1012 accessible by the remote server computing system 1004. In some examples, the remote data 1012 is an unrestricted subset of the energy-related data 1002. In other examples, the remote data 1012 comprises a dataset that is not derived from the energy-related data 1002, but which shares similar data characteristics. For example, the energy-related data 1002 and the remote data 1012 may each comprise seismic imaging datasets, but each dataset may be obtained from a different geographic location.

The energy-related training data pairs 1010 also include a label 1014 serving as ground-truth output. The remote server computing system 1004 is configured to train a machine learning (ML) function 1016 using the plurality of training data pairs 1010. The machine learning function 1016 is trained to predict a classified label 1018 for restricted energy-related data 1020 that is restricted by the cloud service management control plane to one or more of the local edge computing device and the local on-premises computing device. While the present example refers to the one or more on-premises computing devices 1006, it will also be appreciated that the machine learning

function 1016 may be trained to classify data located at an edge computing device (e.g., the local edge device 104 of FIG. 1).

It will be appreciated that the particular set of features included in the energy-related training data pairs 1010 to train the machine learning function 1016 will be included for each and every training session, and will also be included in an input vector at run time, with each input indicated on a normalized scale of zero to one. When a particular feature is present in one session, but is not present in another session, it will be indicated as zero when it is not present.

In some examples, the machine learning function 1016 includes a neural network. The training may take place using any suitable method(s), such as by using backpropagation with gradient descent. As the neural network is trained, an input vector (e.g., a vector comprising a normalized form of the unrestricted data 1012) and matched ground truth labels (e.g., the label 1014), which can be classifications or scalar values, are applied to an input layer and an output layer respectively, and the weights in the network are computed through gradient descent and the backpropagation algorithm, for example, such that the trained neural network will properly classify (or properly value) the input vector to the matched ground truth classification or scalar value in the output layer. In other examples, another suitable AI model may be used, such as neural network of another structure, a support vector machine, etc.

The restricted energy-related data 1020 is not accessible to the remote server computing system 1004 for training due to transmission or storage restriction on the restricted data. For example, the restricted energy-related data 1020 may comprise the restricted data 132 of FIG. 1. As introduced above, in some examples, the restricted energy-related data 1020 comprises sensitive data that is not allowed to be transmitted or stored outside of a jurisdiction. In other examples, the restricted energy-related data 1020 comprises a large dataset 1022 that is challenging to transmit to the remote server 1004 due to its size or network issues between the on-premises computing device(s) 1006 and the remote server 1004. Initial training of the machine-learning function 1016 is also a computationally intensive process, which is advantageously offloaded to the remote server 1004 using the unrestricted energy-related training data 1008.

However, the trained machine learning function 1016 is deployed, via the cloud service data plane (e.g., the data plane 118) to the one or more of the local edge computing device and the local on-premises computing device. In this manner, the trained machine learning function 1016 is used to process the restricted data 1020 locally at the one or more local on-premises computing devices 1006. In some examples, adaptive learning (which may be less computationally intensive than initial training) takes place at the one or more on-premises computing devices 1006 to adapt the machine-learning function 1016 to the restricted data 1020.

In some examples, the classified label 1018 output by the machine learning function 1016 comprises a restricted label 1024 subject to a data transmission or storage restriction to the one or more on-premises computing devices 1006. Accordingly, the restricted label 1024 may not be transmitted to the remote server 1004.

In other examples, the classified label 1018 output by the machine learning function 1016 comprises an unrestricted label 1026 that is not subject to the data transmission or storage restriction. Accordingly, the unrestricted label 1026 is output to the cloud service data plane. In this manner, the remote server 1004 is configured to receive, via the cloud service data plane, the classified output of the trained machine learning function 1016.

In some examples, the one or more on-premises computing devices 1006 additionally or alternatively host a feature extractor 1028. In the example of FIG. 10, the feature extractor 1028 obtains outputs from a feature extraction layer within the machine learning function 1016. In this manner, the feature extractor 1028 can be applied to the large dataset 1022 to produce a reduced-dimensionality dataset 1030. In the example shown in FIG. 10, an input layer of the trained machine learning function 1016 comprises n neurons, which collectively receive an input vector of dimension n. The feature extractor 1028 obtains outputs from a downstream layer comprising $k < n$ neurons. Accordingly, the reduced-dimensionality dataset 1030 has a dimensionality of k, which is less than n, and thereby occupies less storage volume than the large dataset 1022.

In other examples, the one or more on-premises computing devices comprise a feature extractor in the form of a separate AI object from the machine learning function 1016. For example, the feature extractor may comprise a convolutional neural network or a recurrent neural network having a number of output neurons that is less than a number of input neurons, thereby reducing the dimensionality of an input dataset. Further, by using a separate AI object, the feature extractor can be built or trained independent of the machine learning function 1016 (e.g., using different neuron functions, weights, or biases) to provide different outputs.

Like the classified label 1018, in some examples, the reduced-dimensionality dataset 1030 comprises restricted output data 1032 subject to a data transmission or storage restriction to the one or more on-premises computing devices 1006. In other examples, the reduced-dimensionality dataset 1030 comprises unrestricted output data 1034 that is not subject to the data transmission or storage restriction. Accordingly, the unrestricted output data 1034 is output to the cloud service data plane. In this manner, the remote server 1004 is configured to receive, via the cloud service data plane, the unrestricted output data 1034.

FIG. 11 shows another example of a search engine deployment in a hybrid cloud environment 1100 for processing energy-related data 1102. Like the hybrid cloud environment 1000 of FIG. 10, the hybrid cloud environment 1100 includes a remote server 1104. The remote server 1104

can serve as the remote server 102 of FIG. 1. The hybrid cloud environment 1100 further comprises one or more on-premises computing devices 1106. The one or more on-premises computing devices 1106 can serve as the one or more on-premises computing devices 106 of FIG. 1. While the present example refers to the one or more on-premises computing devices 1106, it will also be appreciated that the search engine may be implemented at least in part using an edge computing device (e.g., the local edge device 104 of FIG. 1).

The search engine comprises a crawler 1108. The crawler 1108 is configured to generate an index 1110 of the energy-related data 1102, including a restricted portion 1112 of the energy-related data 1102. For example, the energy-related data 1102 may comprise the energy-related data 128 of FIG. 1, and the restricted data 1112 may comprise the restricted data 132 of FIG. 1.

As introduced above, the restricted energy-related data 1112 is not accessible to the remote server computing system 1104 for training of the search engine or indexing due to transmission or storage restriction on the restricted data. Accordingly, and in one potential advantage of the present disclosure, the remote server 1104 is configured to receive energy-related training data 1114, which is used to train a query engine 1116 to parse a query 1118 and to query the index 1110 at the one or more on-premises computing devices 1106. In some examples, the query 1118 originates at the one or more on-premises computing devices. In other examples, the query 1118 is provided by the remote server 1104.

Like the energy-related training data pairs 1010 of FIG. 10, in some examples, the energy-related training data 1114 is an unrestricted subset of the energy-related data 1102. In other examples, the energy-related training data 1114 comprises a dataset that is not derived from the energy-related data 1102, but which shares similar data characteristics.

The trained query engine 1116 is deployed, via the cloud service data plane (e.g., the data plane 118) to the one or more on-premises computing devices 1106. In this manner, the query function 1116 is used to process the query 1118 locally at the one or more local on-premises computing devices 1106. In some examples, adaptive learning (which may be less computationally intensive than initial training) takes place at the one or more on-premises computing devices 1106 to adapt the query function 1116 to the restricted data 1112.

In some examples, a result 1120 returned in response to the query 1118 is output to the cloud service data plane. In this manner, the remote server 1104 is configured to receive, via the cloud service data plane, the result 1120 of the query.

With reference again to FIG. 1, the one or more extracted features 140 from at least the portion 132 of the energy-related data 128 are output to the data plane 118. For example, the one or more extracted features 140 may comprise the classified label 1018 and/or the reduced-dimensionality dataset 1024 of FIG. 10, or the query result 1120 of FIG. 11. In this manner, the remote server

102 and/or the local edge device 104 receive the one or more extracted features 140, when the one or more extracted features 140 are not subject to the data transmission restriction or the local storage restriction.

With reference now to FIGS. 12A-12B, a flowchart is illustrated depicting an example method 1200 for providing cloud-service-managed governance in a hybrid cloud environment comprising a cloud-service-managed control plane and a data plane utilizing local compute resources and storage located on-premises at an energy production or distribution facility. The following description of method 1200 is provided with reference to the software and hardware components described above and shown in FIGS. 1-11 and 13, and the method steps in method 1200 will be described with reference to corresponding portions of FIGS. 1-11 and 13 below. For example, the method 1200 may be performed at the remote server 102 of FIG. 1. It will be appreciated that method 1200 also may be performed in other contexts using other suitable hardware and software components.

It will be appreciated that the following description of method 1200 is provided by way of example and is not meant to be limiting. It will be understood that various steps of method 1200 can be omitted or performed in a different order than described, and that the method 1200 can include additional and/or alternative steps relative to those illustrated in FIGS. 12A and 12B without departing from the scope of this disclosure.

With reference now to FIG. 12A, the method 1200 includes, at 1202, presenting a user interface with a plurality of deployment configuration options including compute configuration options and data storage configuration options for energy-related data within the hybrid cloud environment. For example, the user interface may comprise the user interface 120 of FIG. 1 and/or the user interface 200 of FIGS. 2-9.

As indicated at 1204, in some examples, the plurality of deployment configuration options includes a regulatory compliance template, such as the regulatory compliance template 258 of FIG. 5, specifying that the energy-related data is stored on one or more storage devices located within a geographic area, and the energy-related data is subject to a data transmission restriction within the geographic area. At 1206, in some examples, the method 1200 comprises determining that the remote server computing system is located in a sovereign cloud, wherein the regulatory compliance template permits access to the energy-related data via the sovereign cloud, and receiving via the data plane, at least the portion of the energy-related data subject to the data transmission restriction or the local storage restriction.

In some examples, at 1208, the method 1200 includes determining that a latency of a network connection between the remote server computing system and the local edge computing device is greater than or equal to a threshold latency, wherein the plurality of deployment configuration

options includes a high-latency network template (such as the high-latency network template 260 of FIG. 7) specifying that the energy-related data is processed at the local edge computing device; and based on determining that the latency of the network connection is greater than or equal to the threshold latency, populating the user interface with the high-latency network template.

At 1210, in some examples, the method 1200 comprises determining that a data volume of the energy-related data is greater than or equal to a threshold data volume, wherein the plurality of deployment configuration options includes a high-data-volume template (such as the high volume template 262 of FIG. 8) specifying that the energy-related data is stored and processed at the local edge computing device; and based on determining that the data volume of the energy-related data is greater than or equal to the threshold data volume, populating the user interface with the high-data-volume template.

In some examples, at 1212, the method 1200 comprises determining that a network connection between the local edge computing device and the local compute resources is intermittent, wherein the plurality of deployment configuration options includes an intermittent connection template (such as the intermittent connection template 264 of FIG. 9) specifying that the energy-related data is cached and processed on the local compute resources and storage; and based on determining that the network connection is intermittent, populating the user interface with the intermittent connection template.

With reference now to FIG. 12B, the method 1200 comprises receiving a user input of one or more of the deployment configuration options. For example, the remote server 102 may receive the user input via the user interface 120 of FIG. 1 or the user interface 200 of FIGS. 2-9.

The method 1200 further comprises, at 1216, generating a data control policy that provides cloud-service-managed governance over at least a portion of the data plane using the one or more user-input deployment configuration options. At 1218, the method 1200 comprises providing the data control policy to a local edge computing device and the local compute resources via the cloud-service-managed control plane, wherein the cloud-service-managed control plane is configured to enforce the data control policy by subjecting at least a portion of the energy-related data to a data transmission restriction or a local storage restriction.

At 1220, the method 1200 comprises using the data plane to deploy one or more cloud service functions to the local compute resources, the one or more cloud service functions configured to process at least the portion of the energy-related data and output one or more extracted features from at least the portion of the energy-related data to the data plane. For example, the data plane 118 of FIG. 1 is used to deploy the cloud service functions 138 to the local edge device 104 and/or the one or more on-premises computing devices 106. Similarly, the trained machine learning function 1016 of FIG. 10 is deployed to the one or more on-premises computing devices 1006 of

FIG. 10, and the query engine 1116 of FIG. 11 is deployed to the one or more on-premises computing devices 1106 of FIG. 11.

With reference again to FIG. 12B, the method 1200 comprises, at 1222, receiving, via the data plane, the one or more extracted features from at least the portion of the energy-related data. For example, the remote server 102 of FIG. 1 receives the extracted features 140 from the local edge device 104 and/or the one or more on-premises computing devices 106 via the data plane 118. Similarly, the remote server 1004 of FIG. 10 receives a classified label 1018 output by the trained machine learning function 1016, and the remote server 1104 can receive a result 1120 output in response to a query 1118. Advantageously, the hybrid cloud system governance disclosed herein enables data, configuration parameters, applications, and services to be deployed at various locations in a uniform manner, preventing the development of a fragmented ecosystem of data types, data formats, tools and services that can occur in other deployments.

In some embodiments, the methods and processes described herein may be tied to a computing system of one or more computing devices. In particular, such methods and processes may be implemented as a computer-application program or service, an application-programming interface (API), a library, and/or other computer-program product.

FIG. 13 schematically shows an example of a computing system 1300 that can enact one or more of the devices and methods described above. Computing system 1300 is shown in simplified form. Computing system 1300 may take the form of one or more personal computers, server computers, tablet computers, home-entertainment computers, network computing devices, gaming devices, mobile computing devices, mobile communication devices (e.g., smart phone), and/or other computing devices, and wearable computing devices such as smart wristwatches and head mounted augmented reality devices. In some examples, the computing system 1300 may embody the remote server 102 of FIG. 1, the local edge device 104 of FIG. 1, the one or more on-premises computing devices 106 of FIG. 1, the remote server 1004 of FIG. 10, the one or more on-premises computing devices 1006 of FIG. 10, the remote server 1104 of FIG. 11, or the one or more on-premises computing devices 1106 of FIG. 11.

The computing system 1300 includes a logic processor 1302 volatile memory 1304, and a non-volatile storage device 1306. The computing system 1300 may optionally include a display subsystem 1308, input subsystem 1310, communication subsystem 1312, and/or other components not shown in FIG. 13.

Logic processor 1302 includes one or more physical devices configured to execute instructions. For example, the logic processor may be configured to execute instructions that are part of one or more applications, programs, routines, libraries, objects, components, data structures, or other logical constructs. Such instructions may be implemented to perform a task, implement a data

type, transform the state of one or more components, achieve a technical effect, or otherwise arrive at a desired result.

The logic processor may include one or more physical processors (hardware) configured to execute software instructions. Additionally or alternatively, the logic processor may include one or more hardware logic circuits or firmware devices configured to execute hardware-implemented logic or firmware instructions. Processors of the logic processor 1302 may be single-core or multi-core, and the instructions executed thereon may be configured for sequential, parallel, and/or distributed processing. Individual components of the logic processor optionally may be distributed among two or more separate devices, which may be remotely located and/or configured for coordinated processing. Aspects of the logic processor may be virtualized and executed by remotely accessible, networked computing devices configured in a cloud-computing configuration. In such a case, these virtualized aspects are run on different physical logic processors of various different machines, it will be understood.

Non-volatile storage device 1306 includes one or more physical devices configured to hold instructions executable by the logic processors to implement the methods and processes described herein. When such methods and processes are implemented, the state of non-volatile storage device 1306 may be transformed—e.g., to hold different data.

Non-volatile storage device 1306 may include physical devices that are removable and/or built in. Non-volatile storage device 1306 may include optical memory (e.g., CD, DVD, HD-DVD, Blu-Ray Disc, etc.), semiconductor memory (e.g., ROM, EPROM, EEPROM, FLASH memory, etc.), and/or magnetic memory (e.g., hard-disk drive, floppy-disk drive, tape drive, MRAM, etc.), or other mass storage device technology. Non-volatile storage device 1306 may include nonvolatile, dynamic, static, read/write, read-only, sequential-access, location-addressable, file-addressable, and/or content-addressable devices. It will be appreciated that non-volatile storage device 1306 is configured to hold instructions even when power is cut to the non-volatile storage device 1306.

Volatile memory 1304 may include physical devices that include random access memory. Volatile memory 1304 is typically utilized by logic processor 1302 to temporarily store information during processing of software instructions. It will be appreciated that volatile memory 1304 typically does not continue to store instructions when power is cut to the volatile memory 1304.

Aspects of logic processor 1302, volatile memory 1304, and non-volatile storage device 1306 may be integrated together into one or more hardware-logic components. Such hardware-logic components may include field-programmable gate arrays (FPGAs), program- and application-specific integrated circuits (PASIC / ASICs), program- and application-specific standard products (PSSP / ASSPs), system-on-a-chip (SOC), and complex programmable logic devices (CPLDs), for example.

The terms "program" and "engine" may be used to describe an aspect of computing system 1300 typically implemented in software by a processor to perform a particular function using portions of volatile memory, which function involves transformative processing that specially configures the processor to perform the function. Thus, a program or engine may be instantiated via logic processor 1302 executing instructions held by non-volatile storage device 1306, using portions of volatile memory 1304. It will be understood that different programs and/or engines may be instantiated from the same application, service, code block, object, library, routine, API, function, etc. Likewise, the same program and/or engine may be instantiated by different applications, services, code blocks, objects, routines, APIs, functions, etc. The terms "program" and "engine" may encompass individual or groups of executable files, data files, libraries, drivers, scripts, database records, etc.

When included, display subsystem 1308 may be used to present a visual representation of data held by non-volatile storage device 1306. The visual representation may take the form of a GUI. As the herein described methods and processes change the data held by the non-volatile storage device, and thus transform the state of the non-volatile storage device, the state of display subsystem 1308 may likewise be transformed to visually represent changes in the underlying data. Display subsystem 1308 may include one or more display devices utilizing virtually any type of technology. Such display devices may be combined with logic processor 1302, volatile memory 1304, and/or non-volatile storage device 1306 in a shared enclosure, or such display devices may be peripheral display devices.

When included, input subsystem 1310 may comprise or interface with one or more user-input devices such as a keyboard, mouse, touch screen, or game controller. In some examples, the input subsystem may comprise or interface with selected natural user input (NUI) componentry. Such componentry may be integrated or peripheral, and the transduction and/or processing of input actions may be handled on- or off-board. Example NUI componentry may include a microphone for speech and/or voice recognition; an infrared, color, stereoscopic, and/or depth camera for machine vision and/or gesture recognition; a head tracker, eye tracker, accelerometer, and/or gyroscope for motion detection and/or intent recognition; as well as electric-field sensing componentry for assessing brain activity; and/or any other suitable sensor.

When included, communication subsystem 1312 may be configured to communicatively couple various computing devices described herein with each other, and with other devices. Communication subsystem 1312 may include wired and/or wireless communication devices compatible with one or more different communication protocols. For example, the communication subsystem may be configured for communication via a wireless telephone network, or a wired or wireless local- or wide-area network. In some examples, the communication subsystem may allow

computing system 1300 to send and/or receive messages to and/or from other devices via a network such as the Internet.

The following paragraphs discuss several aspects of the present disclosure. According to one aspect of the present disclosure, a remote server computing system comprises a processor, and a memory storing instructions executable by the processor. The instructions are executable to present a user interface with a plurality of deployment configuration options including compute configuration options and data storage configuration options for energy-related data within a hybrid cloud environment comprising a cloud-service-managed control plane and a data plane utilizing local compute resources and storage located on-premises at an energy production or distribution facility. The cloud-service-managed control plane and the data plane span the remote server computing system, a local edge computing device, and the local compute resources and storage. A user input of one or more of the deployment configuration options is received. A data control policy is generated using the one or more user-input deployment configuration options. The data control policy provides cloud-service-managed governance over at least a portion of the data plane. The data control policy is provided to the local edge computing device and the local compute resources via the cloud-service-managed control plane. The cloud-service-managed control plane is configured to enforce the data control policy by subjecting at least a portion of the energy-related data to a data transmission restriction or a local storage restriction. The data plane is used to deploy one or more cloud service functions to the local compute resources. The one or more cloud service functions are configured to process at least the portion of the energy-related data and output one or more extracted features from at least the portion of the energy-related data to the data plane. The instructions are further executable to receive, via the data plane, the one or more extracted features from at least the portion of the energy-related data.

The plurality of deployment configuration options may additionally or alternatively include a regulatory compliance template specifying that the energy-related data is stored on one or more storage devices located within a geographic area, and the energy-related data is subject to the data transmission restriction within the geographic area. The instructions may be additionally or alternatively executable to determine that the remote server computing system is located in a sovereign cloud; wherein the regulatory compliance template permits access to the energy-related data via the sovereign cloud; and receive, via the data plane, at least the portion of the energy-related data subject to the data transmission restriction or the local storage restriction.

The remote server computing system may additionally or alternatively include a network connection between the remote server computing system and the local edge computing device. The instructions may be additionally or alternatively executable to determine that a latency of the network connection is greater than or equal to a threshold latency; wherein the plurality of

deployment configuration options includes a high-latency network template specifying that the energy-related data is processed at the local edge computing device; and based on determining that the latency of the network connection is greater than or equal to the threshold latency, populate the user interface with the high-latency network template. The threshold latency may additionally or alternatively be 50 ms.

The instructions may be additionally or alternatively executable to determine that a data volume of the energy-related data is greater than or equal to a threshold data volume; wherein the plurality of deployment configuration options includes a high-data-volume template specifying that the energy-related data is stored and processed at the local edge computing device; and based on determining that the data volume of the energy-related data is greater than or equal to the threshold data volume, populate the user interface with the high-data-volume template. The threshold data volume may additionally or alternatively be in the range of 10-15 TB streaming data per day per site.

The remote server computing system may additionally or alternatively include a network connection between the local edge computing device and the local compute resources. The instructions may be additionally or alternatively executable to determine that the network connection is intermittent; wherein the plurality of deployment configuration options includes an intermittent connection template specifying that the energy-related data is cached and processed on the local compute resources and storage; and based on determining that the network connection is intermittent, populate the user interface with the intermittent connection template.

The local compute resources and storage may additionally or alternatively comprise one or more physical computing devices running a cluster comprising a plurality of virtual machines.

A container platform of the local compute resources and storage may additionally or alternatively be a foundation for applications and services to be deployed across the hybrid cloud environment.

The data control policy may additionally or alternatively comprise one or more files in a distributed version control repository.

Another aspect provides, at a computing device, a method for providing cloud-service-managed governance in a hybrid cloud environment comprising a cloud-service-managed control plane and a data plane utilizing local compute resources and storage located on-premises at an energy production or distribution facility. The cloud-service-managed control plane and the data plane span a remote server computing system, a local edge computing device, and the local compute resources and storage. The method comprises presenting a user interface with a plurality of deployment configuration options including compute configuration options and data storage configuration options for energy-related data within the hybrid cloud environment. A user input of one or more of the deployment configuration options is received. A data control policy is

generated that provides cloud-service-managed governance over at least a portion of the data plane using the one or more user-input deployment configuration options. The data control policy is provided to the local edge computing device and the local compute resources via the cloud-service-managed control plane. The cloud-service-managed control plane is configured to enforce the data control policy by subjecting at least a portion of the energy-related data to a data transmission restriction or a local storage restriction. The method further comprises using the data plane to deploy one or more cloud service functions to the local compute resources. The one or more cloud service functions are configured to process at least the portion of the energy-related data and output one or more extracted features from at least the portion of the energy-related data to the data plane. The method further comprises receiving, via the data plane, the one or more extracted features from at least the portion of the energy-related data.

The plurality of deployment configuration options may additionally or alternatively include a regulatory compliance template specifying that the energy-related data is stored on one or more storage devices located within a geographic area, and the energy-related data is subject to a data transmission restriction within the geographic area. The method may additionally or alternatively include determining that the remote server computing system is located in a sovereign cloud, wherein the regulatory compliance template permits access to the energy-related data via the sovereign cloud; and receiving via the data plane, at least the portion of the energy-related data subject to the data transmission restriction or the local storage restriction.

The method may additionally or alternatively include determining that a latency of a network connection between the remote server computing system and the local edge computing device is greater than or equal to a threshold latency, wherein the plurality of deployment configuration options includes a high-latency network template specifying that the energy-related data is processed at the local edge computing device; and based on determining that the latency of the network connection is greater than or equal to the threshold latency, populating the user interface with the high-latency network template.

The method may additionally or alternatively include determining, that a data volume of the energy-related data is greater than or equal to a threshold data volume, wherein the plurality of deployment configuration options includes a high-data-volume template specifying that the energy-related data is stored and processed at the local edge computing device; and based on determining that the data volume of the energy-related data is greater than or equal to the threshold data volume, populating the user interface with the high-data-volume template.

The method may additionally or alternatively include determining that a network connection between the local edge computing device and the local compute resources is intermittent, wherein the plurality of deployment configuration options includes an intermittent connection template

specifying that the energy-related data is cached and processed on the local compute resources and storage; and based on determining that the network connection is intermittent, populating the user interface with the intermittent connection template.

Another aspect provides a computing device comprising a processor, and a memory storing instructions executable by the processor. The instructions are executable to receive a data control policy from a remote server computing system via a cloud-service-managed control plane. The data control policy provides cloud-service-managed governance over at least a portion of a data plane utilizing local compute resources and storage located on-premises at an energy production or distribution facility. Energy-related data is received. The instructions are further executable to enforce the data control policy by subjecting at least a portion of the energy-related data to a data transmission restriction or a local storage restriction. One or more cloud service functions are received from the remote server computing system via the data plane. At least the portion of the energy-related data is processed using the one or more cloud service functions to extract one or more features from at least the portion of the energy-related data. The one or more extracted features from at least the portion of the energy-related data are output via the data plane.

The instructions may be additionally or alternatively executable to run a container platform that provides a foundation for applications and services to be deployed across a hybrid cloud environment comprising the computing device and a remote server.

The computing device may additionally or alternatively include a distributed version control repository housing the data control policy.

It will be understood that the configurations and/or approaches described herein are exemplary in nature, and that these specific embodiments or examples are not to be considered in a limiting sense, because numerous variations are possible. The specific routines or methods described herein may represent one or more of any number of processing strategies. As such, various acts illustrated and/or described may be performed in the sequence illustrated and/or described, in other sequences, in parallel, or omitted. Likewise, the order of the above-described processes may be changed.

The subject matter of the present disclosure includes all novel and non-obvious combinations and sub-combinations of the various processes, systems and configurations, and other features, functions, acts, and/or properties disclosed herein, as well as any and all equivalents thereof.

CLAIMS

1. A remote server computing system, comprising:

a processor, and

a memory storing instructions executable by the processor to:

present a user interface with a plurality of deployment configuration options including compute configuration options and data storage configuration options for energy-related data within a hybrid cloud environment comprising a cloud-service-managed control plane and a data plane utilizing local compute resources and storage located on-premises at an energy production or distribution facility, the cloud-service-managed control plane and the data plane spanning the remote server computing system, a local edge computing device, and the local compute resources and storage;

receive a user input of one or more of the deployment configuration options;

generate a data control policy that provides cloud-service-managed governance over at least a portion of the data plane using the one or more user-input deployment configuration options;

provide the data control policy to the local edge computing device and the local compute resources via the cloud-service-managed control plane, wherein the cloud-service-managed control plane is configured to enforce the data control policy by subjecting at least a portion of the energy-related data to a data transmission restriction or a local storage restriction;

use the data plane to deploy one or more cloud service functions to the local compute resources, the one or more cloud service functions configured to process at least the portion of the energy-related data and output one or more extracted features from at least the portion of the energy-related data to the data plane; and

receive, via the data plane, the one or more extracted features from at least the portion of the energy-related data.

2. The remote server computing system of claim 1, wherein the plurality of deployment configuration options includes a regulatory compliance template specifying that the energy-related data is stored on one or more storage devices located within a geographic area, and the energy-related data is subject to the data transmission restriction within the geographic area.

3. The remote server computing system of claim 2, wherein the instructions are further executable to:

determine that the remote server computing system is located in a sovereign cloud;

wherein the regulatory compliance template permits access to the energy-related data via the sovereign cloud; and

receive, via the data plane, at least the portion of the energy-related data subject to the data

transmission restriction or the local storage restriction.

4. The remote server computing system of claim 1, further comprising a network connection between the remote server computing system and the local edge computing device, wherein the instructions are further executable to:

determine that a latency of the network connection is greater than or equal to a threshold latency;

wherein the plurality of deployment configuration options includes a high-latency network template specifying that the energy-related data is processed at the local edge computing device; and

based on determining that the latency of the network connection is greater than or equal to the threshold latency, populate the user interface with the high-latency network template.

5. The remote server computing system of claim 1, wherein the instructions are further executable to:

determine that a data volume of the energy-related data is greater than or equal to a threshold data volume;

wherein the plurality of deployment configuration options includes a high-data-volume template specifying that the energy-related data is stored and processed at the local edge computing device; and

based on determining that the data volume of the energy-related data is greater than or equal to the threshold data volume, populate the user interface with the high-data-volume template.

6. The remote server computing system of claim 1, further comprising a network connection between the local edge computing device and the local compute resources, wherein the instructions are further executable to:

determine that the network connection is intermittent;

wherein the plurality of deployment configuration options includes an intermittent connection template specifying that the energy-related data is cached and processed on the local compute resources and storage; and

based on determining that the network connection is intermittent, populate the user interface with the intermittent connection template.

7. The remote server computing system of claim 1, wherein the local compute resources and storage comprise one or more physical computing devices running a cluster comprising a plurality of virtual machines.

8. The remote server computing system of claim 1, wherein a container platform of the local compute resources and storage is a foundation for applications and services to be deployed across

the hybrid cloud environment.

9. The remote server computing system of claim 1, wherein the data control policy comprises one or more files in a distributed version control repository.

10. At a computing device, a method for providing cloud-service-managed governance in a hybrid cloud environment comprising a cloud-service-managed control plane and a data plane utilizing local compute resources and storage located on-premises at an energy production or distribution facility, the cloud-service-managed control plane and the data plane spanning a remote server computing system, a local edge computing device, and the local compute resources and storage, the method comprising:

presenting a user interface with a plurality of deployment configuration options including compute configuration options and data storage configuration options for energy-related data within the hybrid cloud environment;

receiving a user input of one or more of the deployment configuration options;

generating a data control policy that provides cloud-service-managed governance over at least a portion of the data plane using the one or more user-input deployment configuration options;

providing the data control policy to the local edge computing device and the local compute resources via the cloud-service-managed control plane, wherein the cloud-service-managed control plane is configured to enforce the data control policy by subjecting at least a portion of the energy-related data to a data transmission restriction or a local storage restriction;

using the data plane to deploy one or more cloud service functions to the local compute resources, the one or more cloud service functions configured to process at least the portion of the energy-related data and output one or more extracted features from at least the portion of the energy-related data to the data plane; and

receiving, via the data plane, the one or more extracted features from at least the portion of the energy-related data.

11. The method of claim 10, wherein the plurality of deployment configuration options includes a regulatory compliance template specifying that the energy-related data is stored on one or more storage devices located within a geographic area, and the energy-related data is subject to a data transmission restriction within the geographic area.

12. The method of claim 11, further comprising:

determining that the remote server computing system is located in a sovereign cloud, wherein the regulatory compliance template permits access to the energy-related data via the sovereign cloud; and

receiving via the data plane, at least the portion of the energy-related data subject to the

data transmission restriction or the local storage restriction.

13. The method of claim 10, further comprising:

determining that a latency of a network connection between the remote server computing system and the local edge computing device is greater than or equal to a threshold latency, wherein the plurality of deployment configuration options includes a high-latency network template specifying that the energy-related data is processed at the local edge computing device; and

based on determining that the latency of the network connection is greater than or equal to the threshold latency, populating the user interface with the high-latency network template.

14. The method of claim 10, further comprising:

determining, that a data volume of the energy-related data is greater than or equal to a threshold data volume, wherein the plurality of deployment configuration options includes a high-data-volume template specifying that the energy-related data is stored and processed at the local edge computing device; and

based on determining that the data volume of the energy-related data is greater than or equal to the threshold data volume, populating the user interface with the high-data-volume template.

15. The method of claim 10, further comprising:

determining that a network connection between the local edge computing device and the local compute resources is intermittent, wherein the plurality of deployment configuration options includes an intermittent connection template specifying that the energy-related data is cached and processed on the local compute resources and storage; and

based on determining that the network connection is intermittent, populating the user interface with the intermittent connection template.

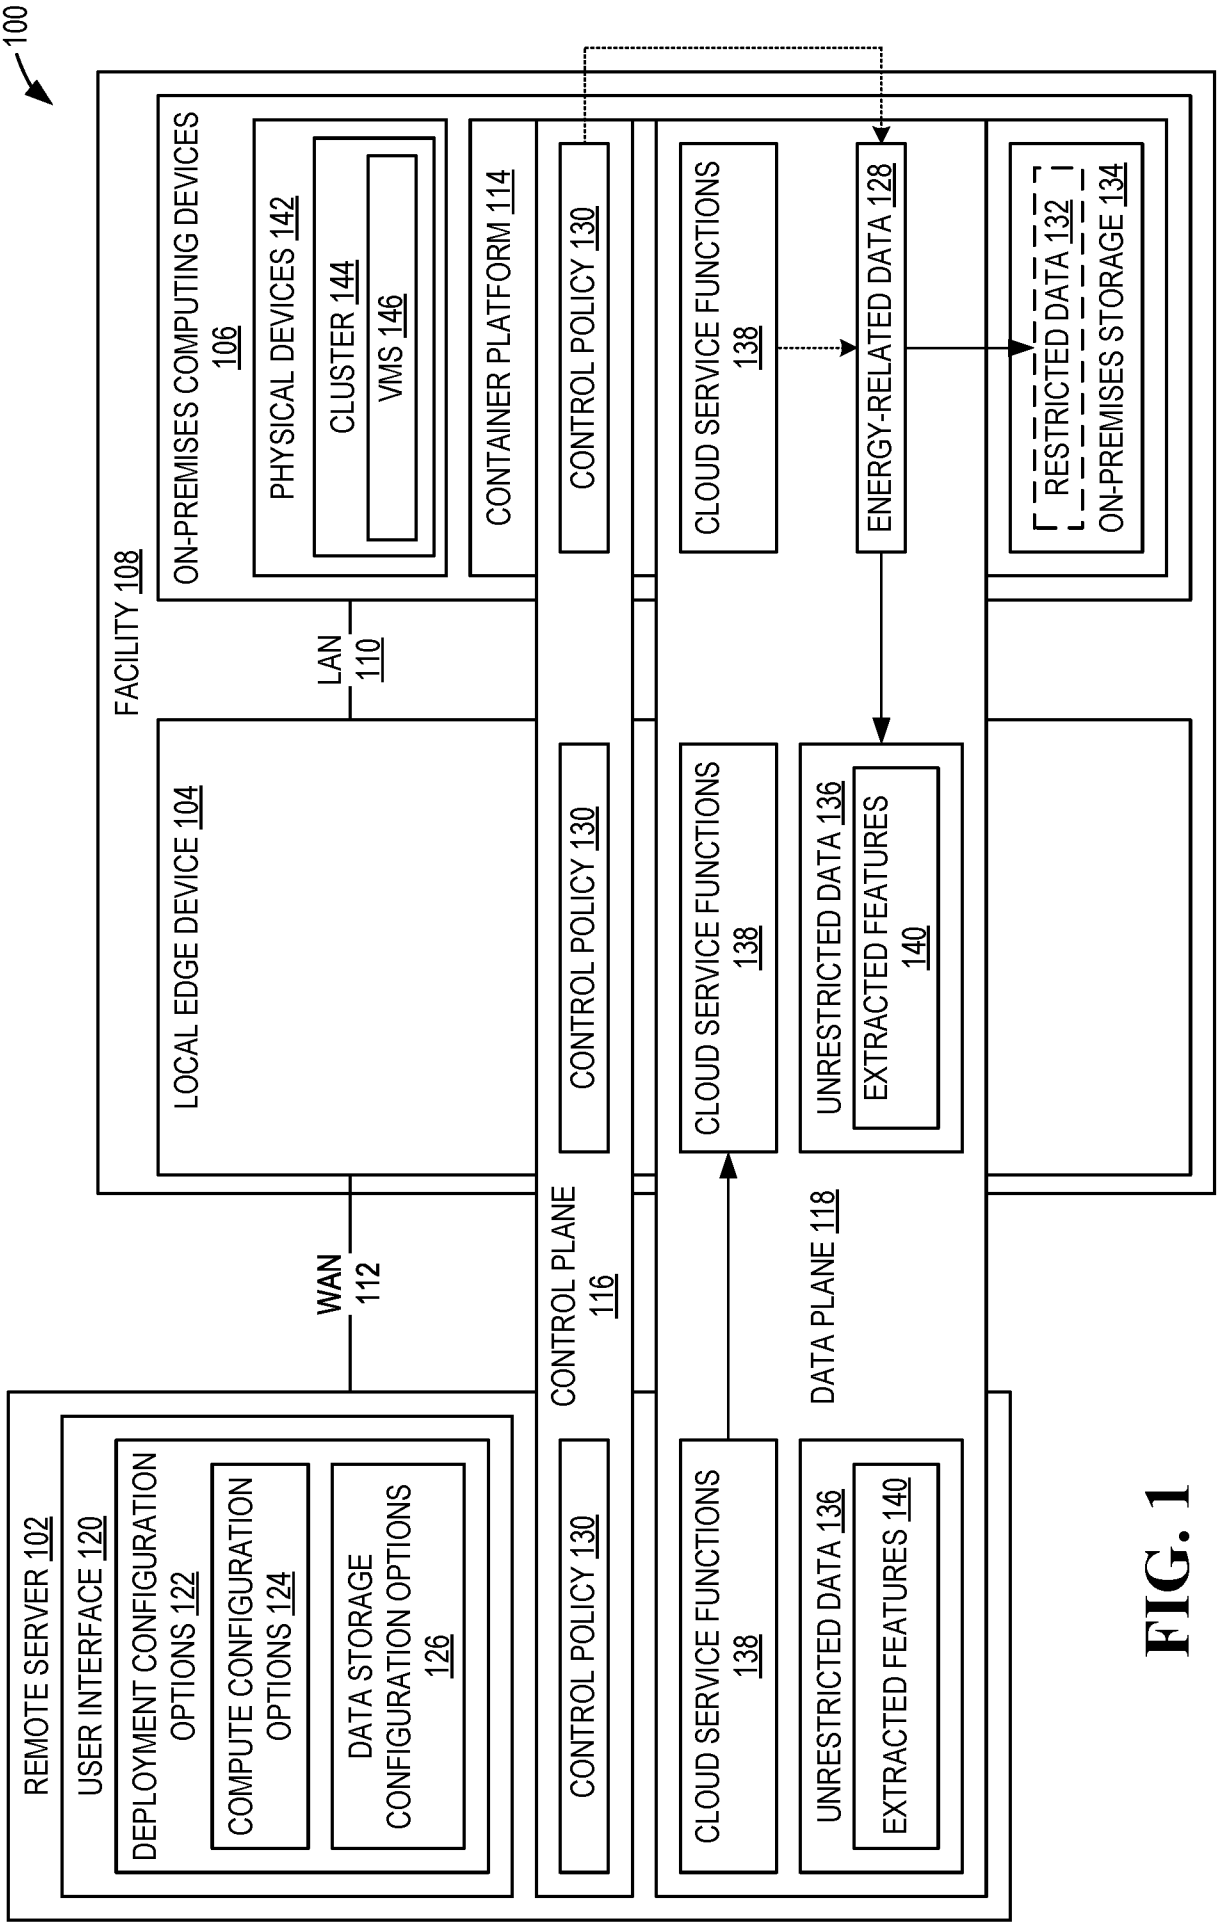
100
REMOTE SERVER 102
USER INTERFACE 120
DEPLOYMENT CONFIGURATION OPTIONS 122
COMPUTE CONFIGURATION OPTIONS 124
DATA STORAGE CONFIGURATION OPTIONS 126
WAN 112
FACILITY 108
LOCAL EDGE DEVICE 104
LAN 110
ON-PREMISES COMPUTING DEVICES 106
PHYSICAL DEVICES 142
CLUSTER 144
VMS 146
CONTAINER PLATFORM 114
CONTROL POLICY 130
CONTROL PLANE 116
CONTROL POLICY 130
CONTROL POLICY 130
CLOUD SERVICE FUNCTIONS 138
CLOUD SERVICE FUNCTIONS 138
CLOUD SERVICE FUNCTIONS 138
ENERGY-RELATED DATA 128
UNRESTRICTED DATA 136
EXTRACTED FEATURES 140
UNRESTRICTED DATA 136
EXTRACTED FEATURES 140
DATA PLANE 118
RESTRICTED DATA 132
ON-PREMISES STORAGE 134

FIG. 1

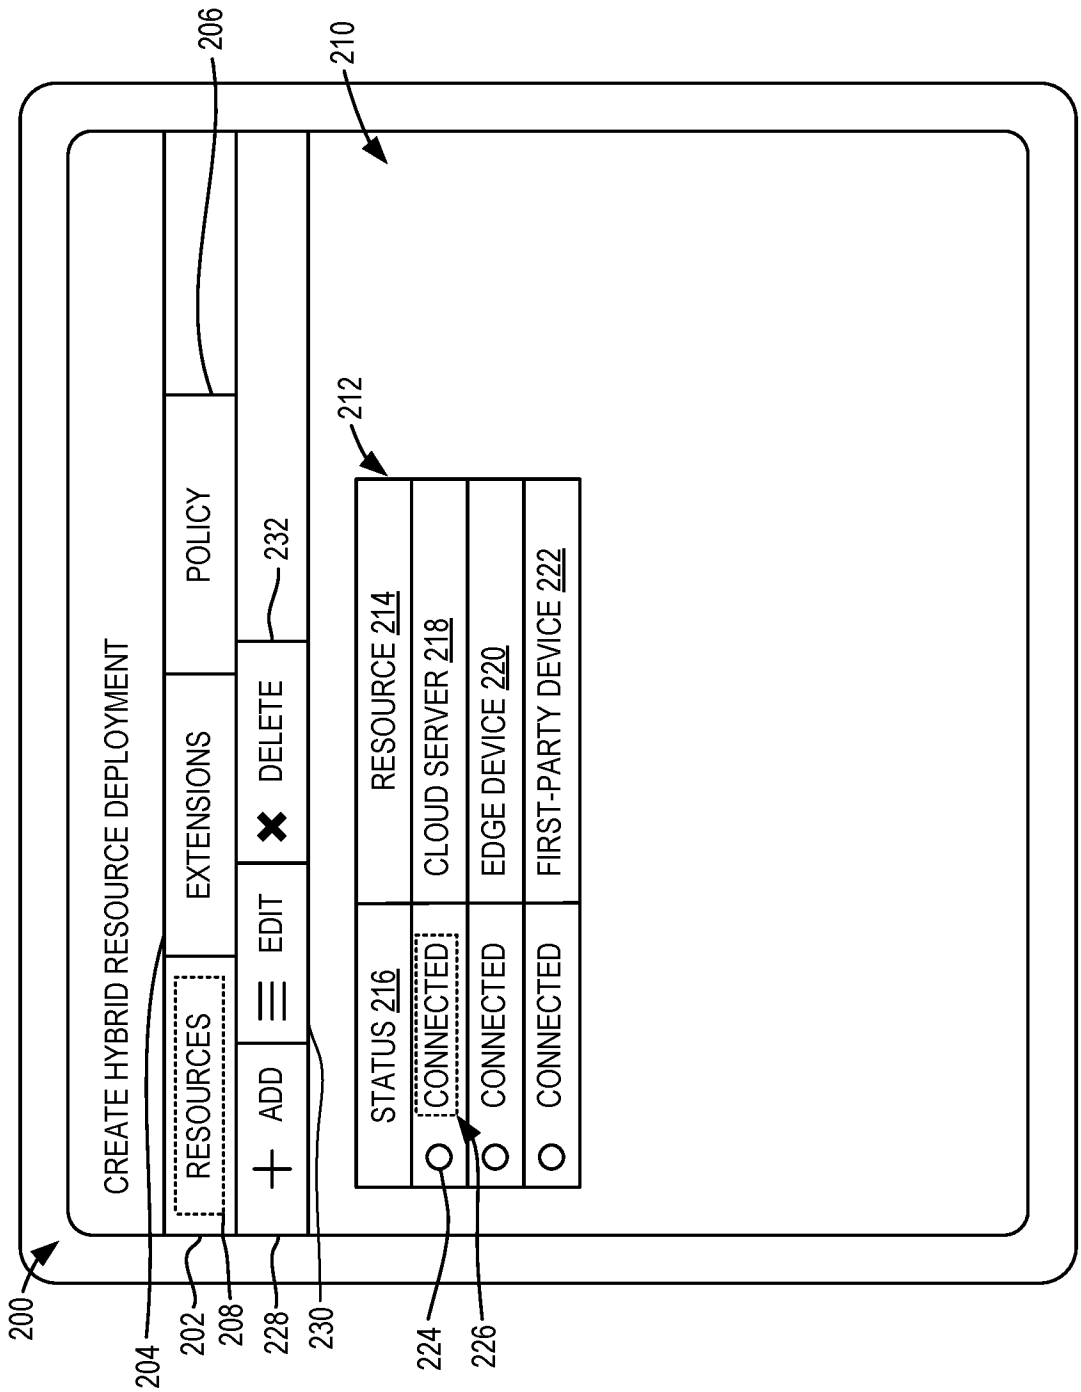
200
202
204
206
208
210
212
214
216
218
220
222
224
226
228
230
232
CREATE HYBRID RESOURCE DEPLOYMENT
RESOURCES
EXTENSIONS
POLICY
ADD
EDIT
DELETE
STATUS 216
RESOURCE 214
CONNECTED
CLOUD SERVER 218
CONNECTED
EDGE DEVICE 220
CONNECTED
FIRST-PARTY DEVICE 222

FIG. 2

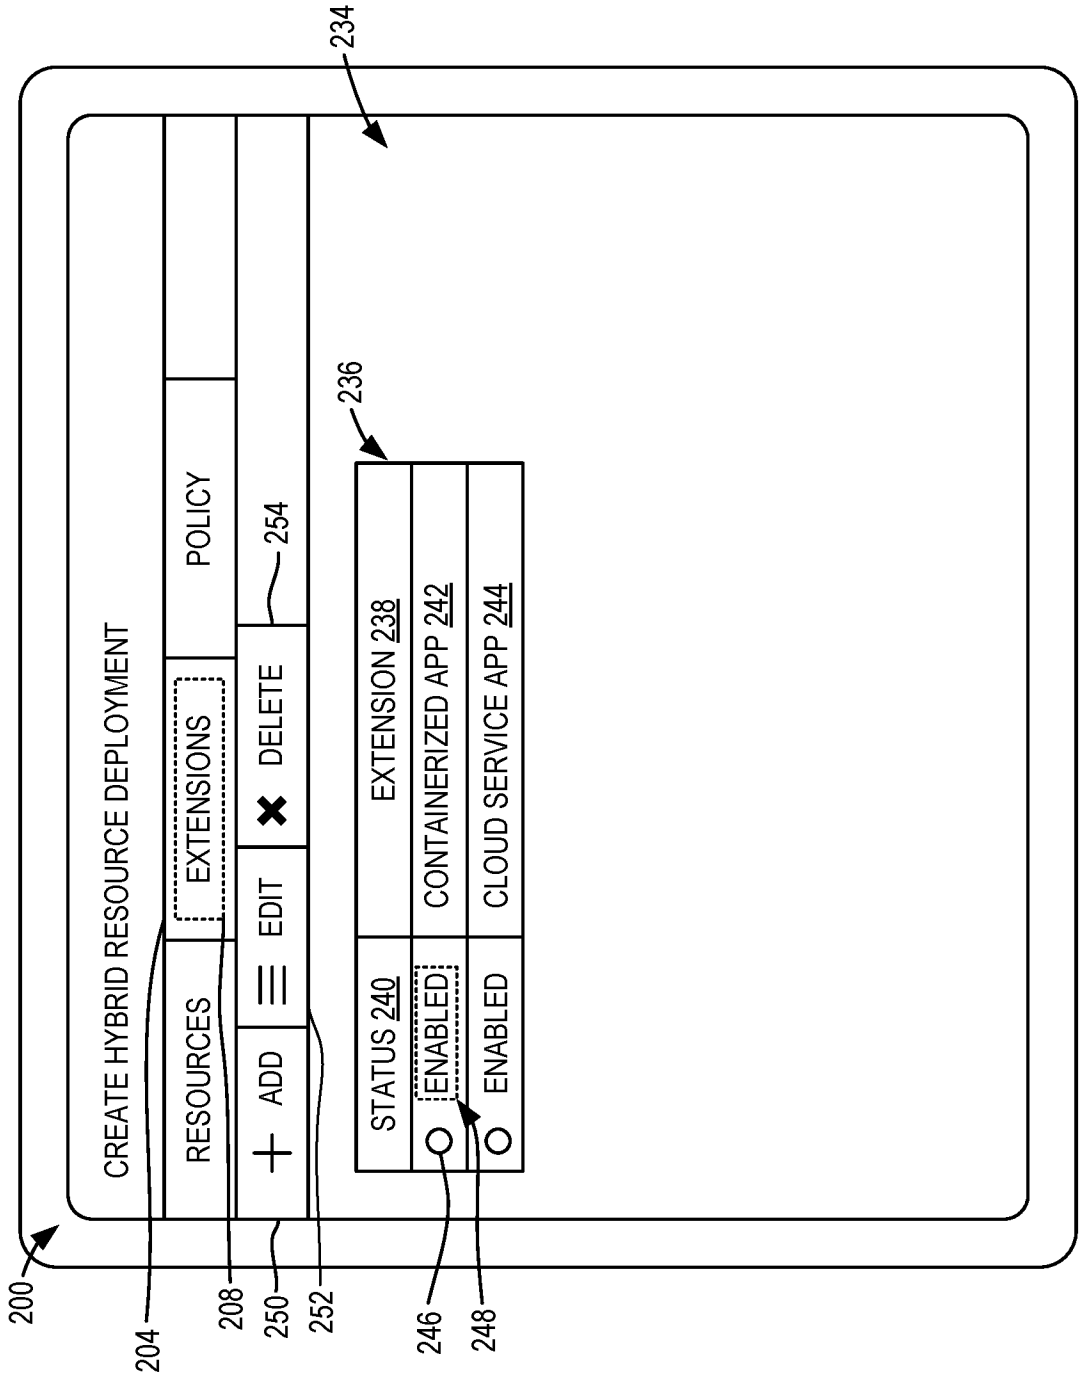
234
CREATE HYBRID RESOURCE DEPLOYMENT
RESOURCES
EXTENSIONS
POLICY
ADD
EDIT
DELETE
254
236
STATUS 240
EXTENSION 238
ENABLED
CONTAINERIZED APP 242
ENABLED
CLOUD SERVICE APP 244
200
204
208
250
252
246
248

FIG. 3

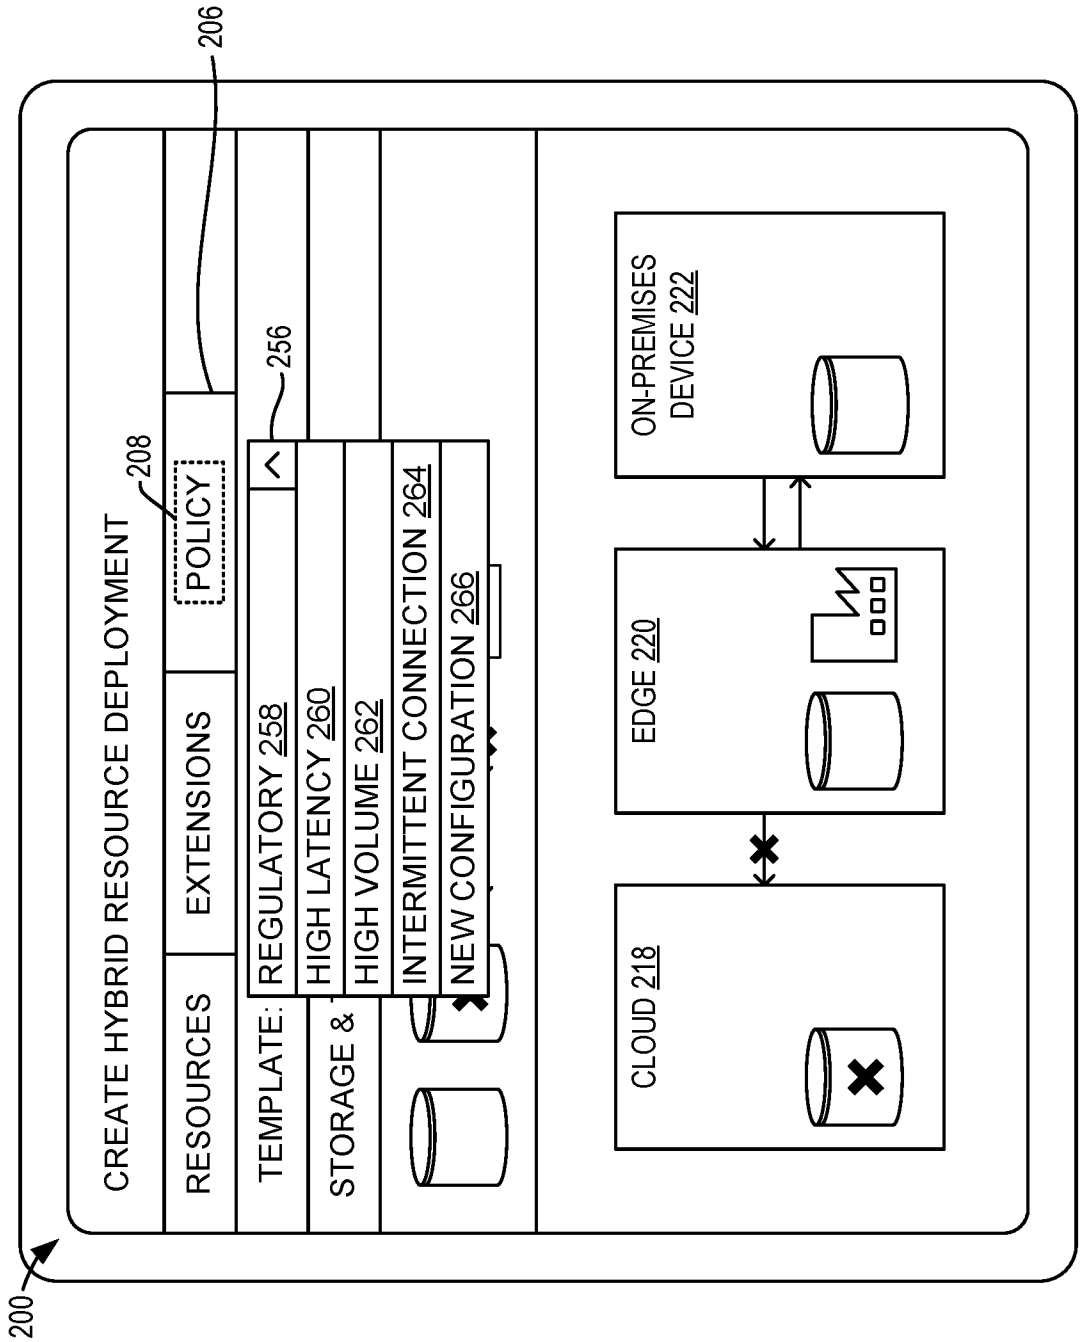
200
206
CREATE HYBRID RESOURCE DEPLOYMENT
208
RESOURCES
EXTENSIONS
POLICY
TEMPLATE:
256
REGULATORY 258
HIGH LATENCY 260
HIGH VOLUME 262
INTERMITTENT CONNECTION 264
NEW CONFIGURATION 266
STORAGE &
CLOUD 218
EDGE 220
ON-PREMISES DEVICE 222

FIG. 4

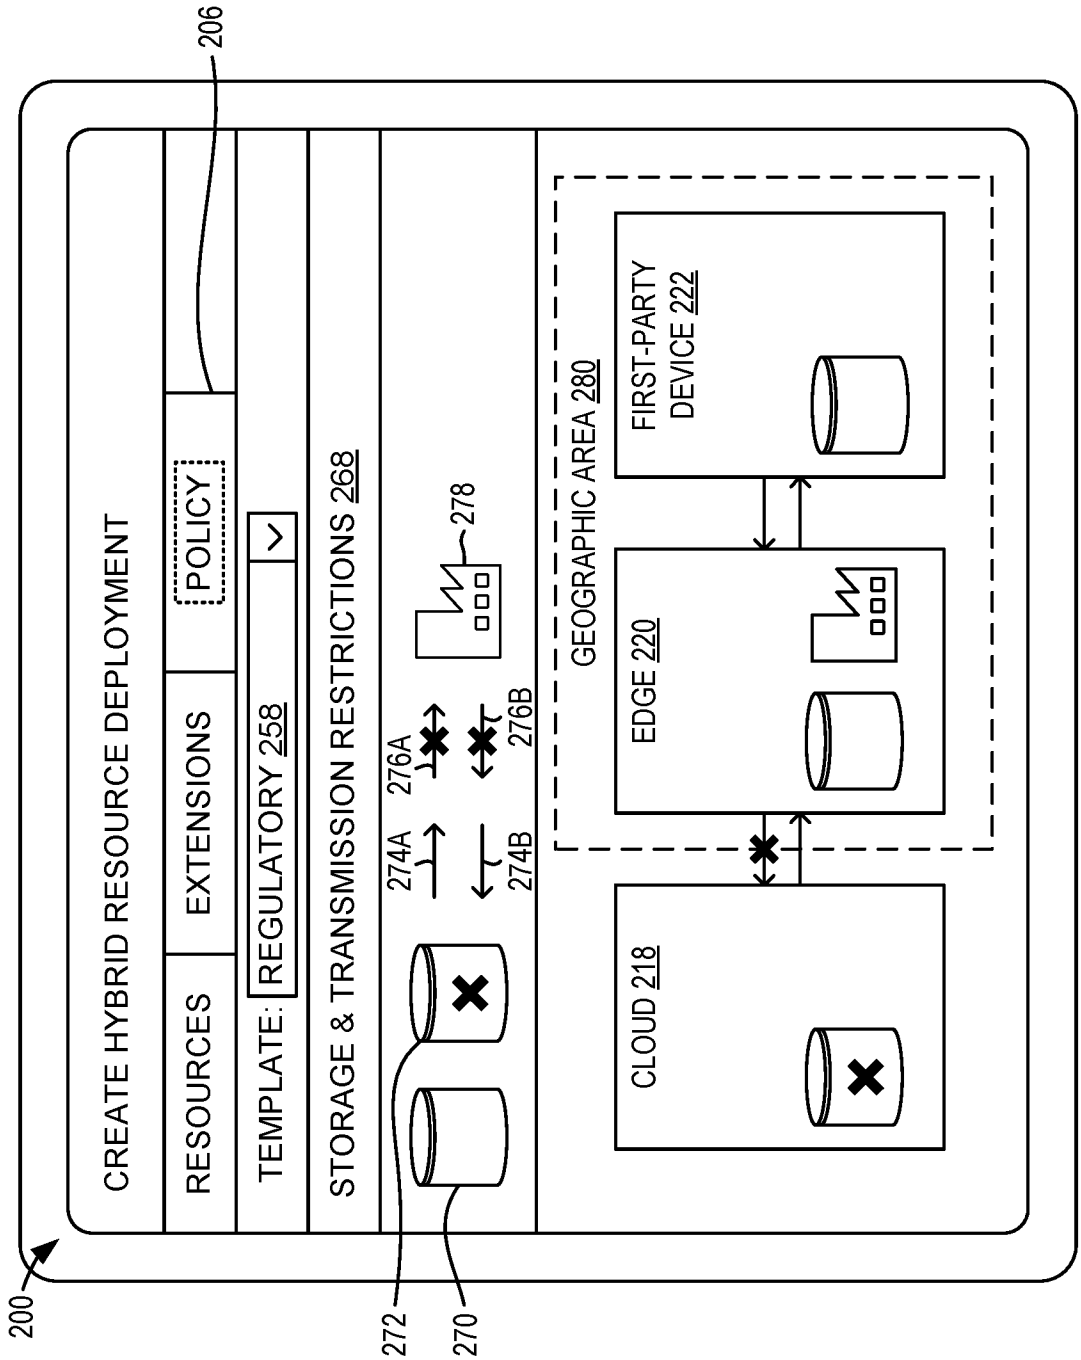
200
206
CREATE HYBRID RESOURCE DEPLOYMENT
RESOURCES
EXTENSIONS
POLICY
TEMPLATE: REGULATORY 258
STORAGE & TRANSMISSION RESTRICTIONS 268
272
270
274A
274B
276A
276B
278
GEOGRAPHIC AREA 280
CLOUD 218
EDGE 220
FIRST-PARTY DEVICE 222

FIG. 5

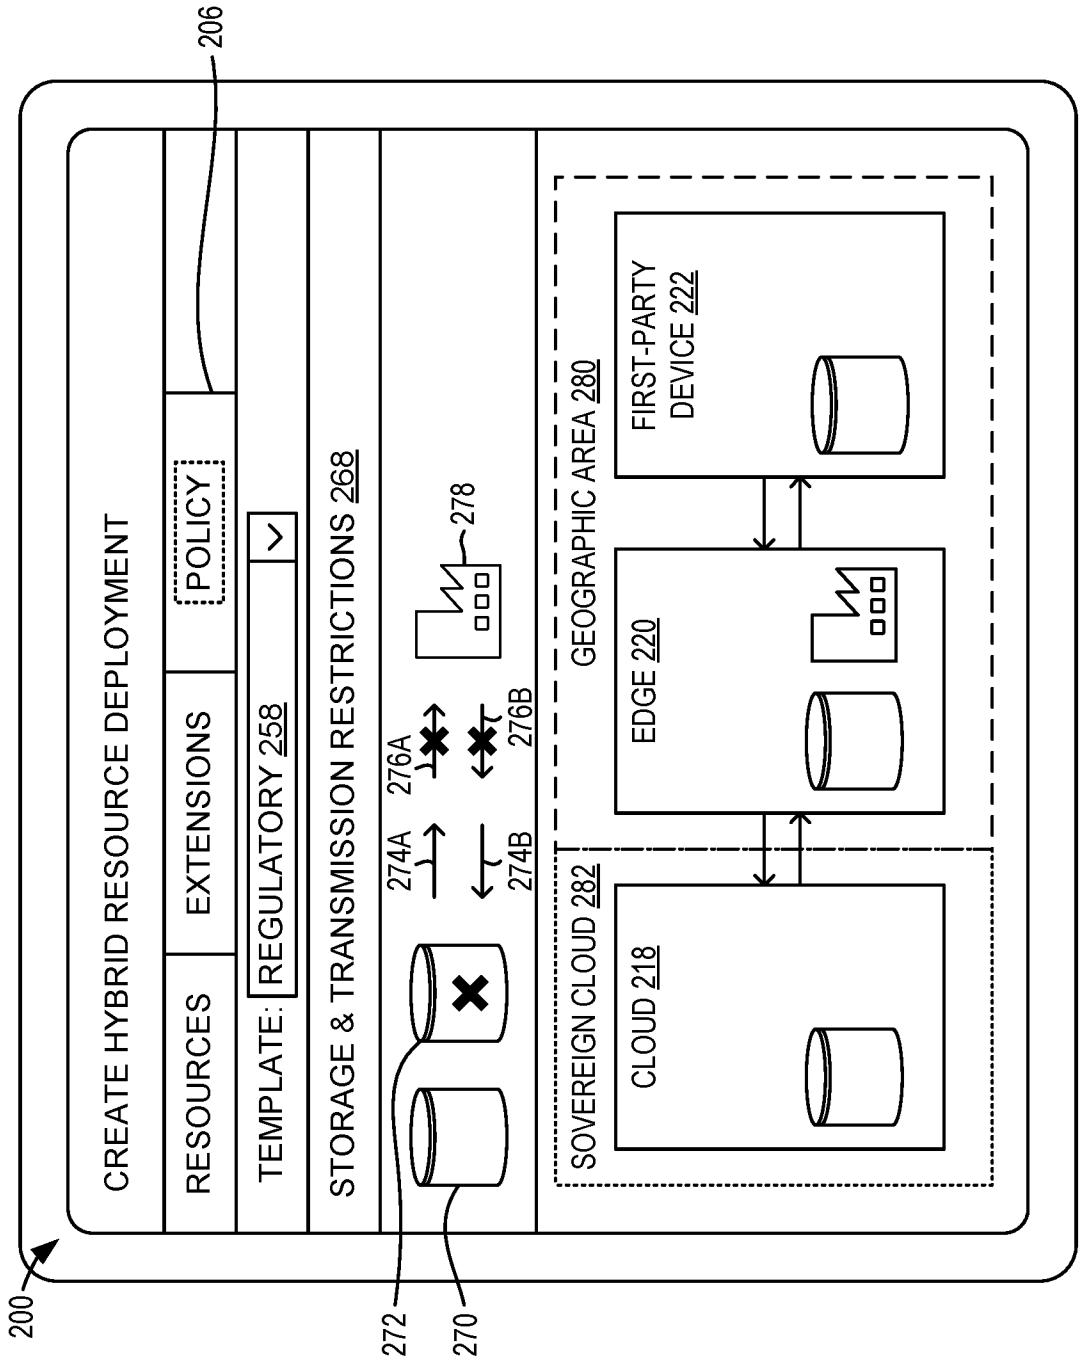
200
206
CREATE HYBRID RESOURCE DEPLOYMENT
RESOURCES
EXTENSIONS
POLICY
TEMPLATE: REGULATORY 258
STORAGE & TRANSMISSION RESTRICTIONS 268
272
270
274A
274B
276A
276B
278
GEOGRAPHIC AREA 280
SOVEREIGN CLOUD 282
CLOUD 218
EDGE 220
FIRST-PARTY DEVICE 222

FIG. 6

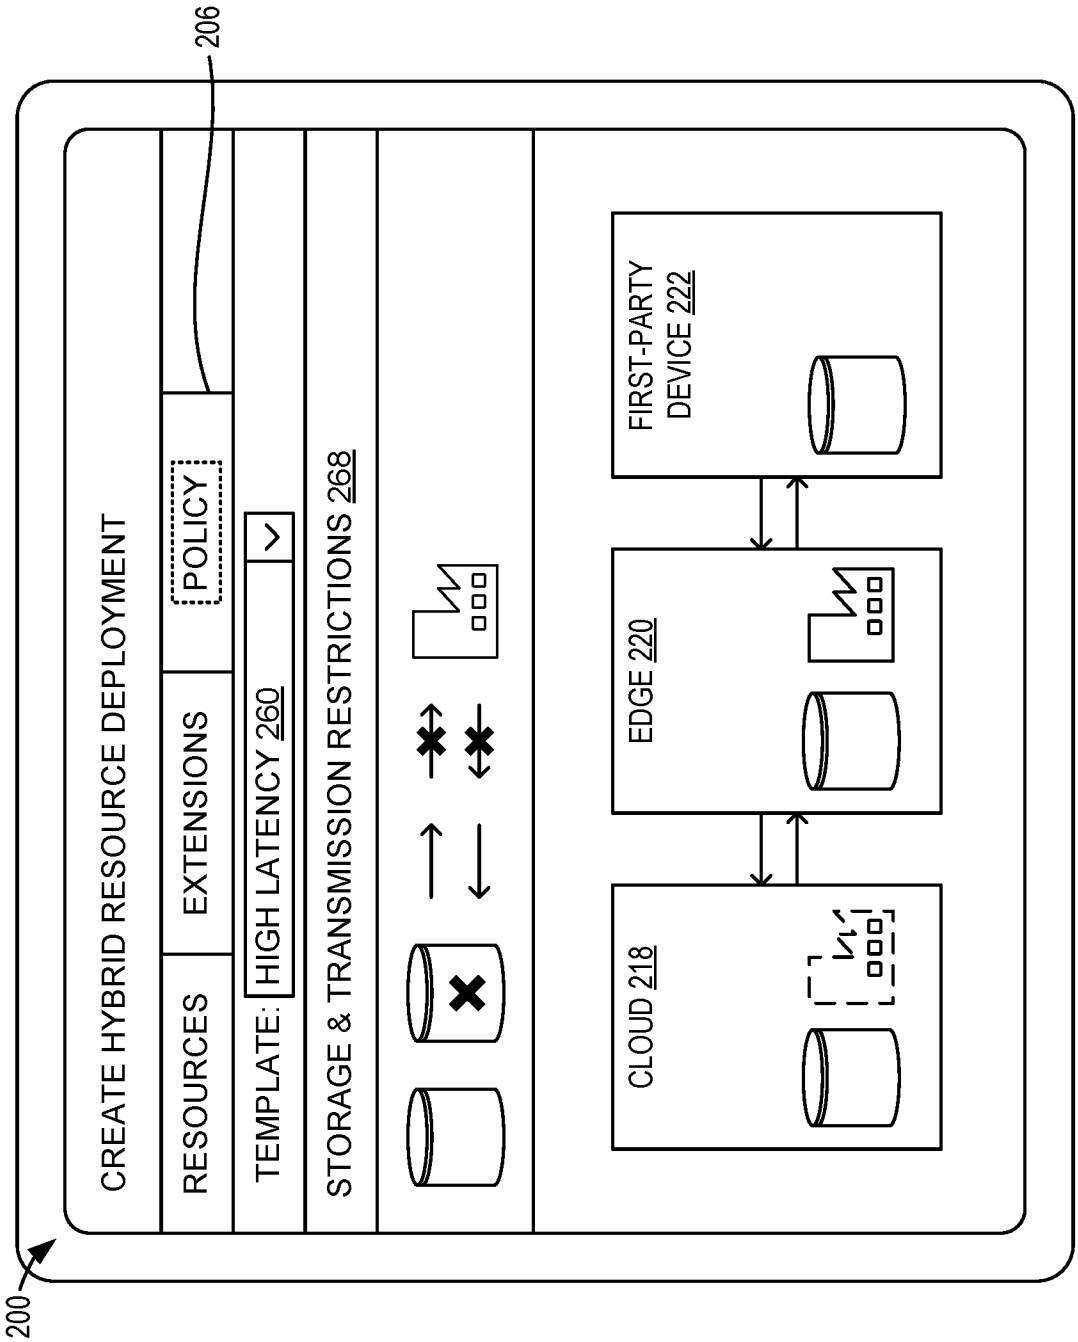
200
206
CREATE HYBRID RESOURCE DEPLOYMENT
RESOURCES
EXTENSIONS
POLICY
TEMPLATE: HIGH LATENCY 260
STORAGE & TRANSMISSION RESTRICTIONS 268
CLOUD 218
EDGE 220
FIRST-PARTY DEVICE 222

FIG. 7

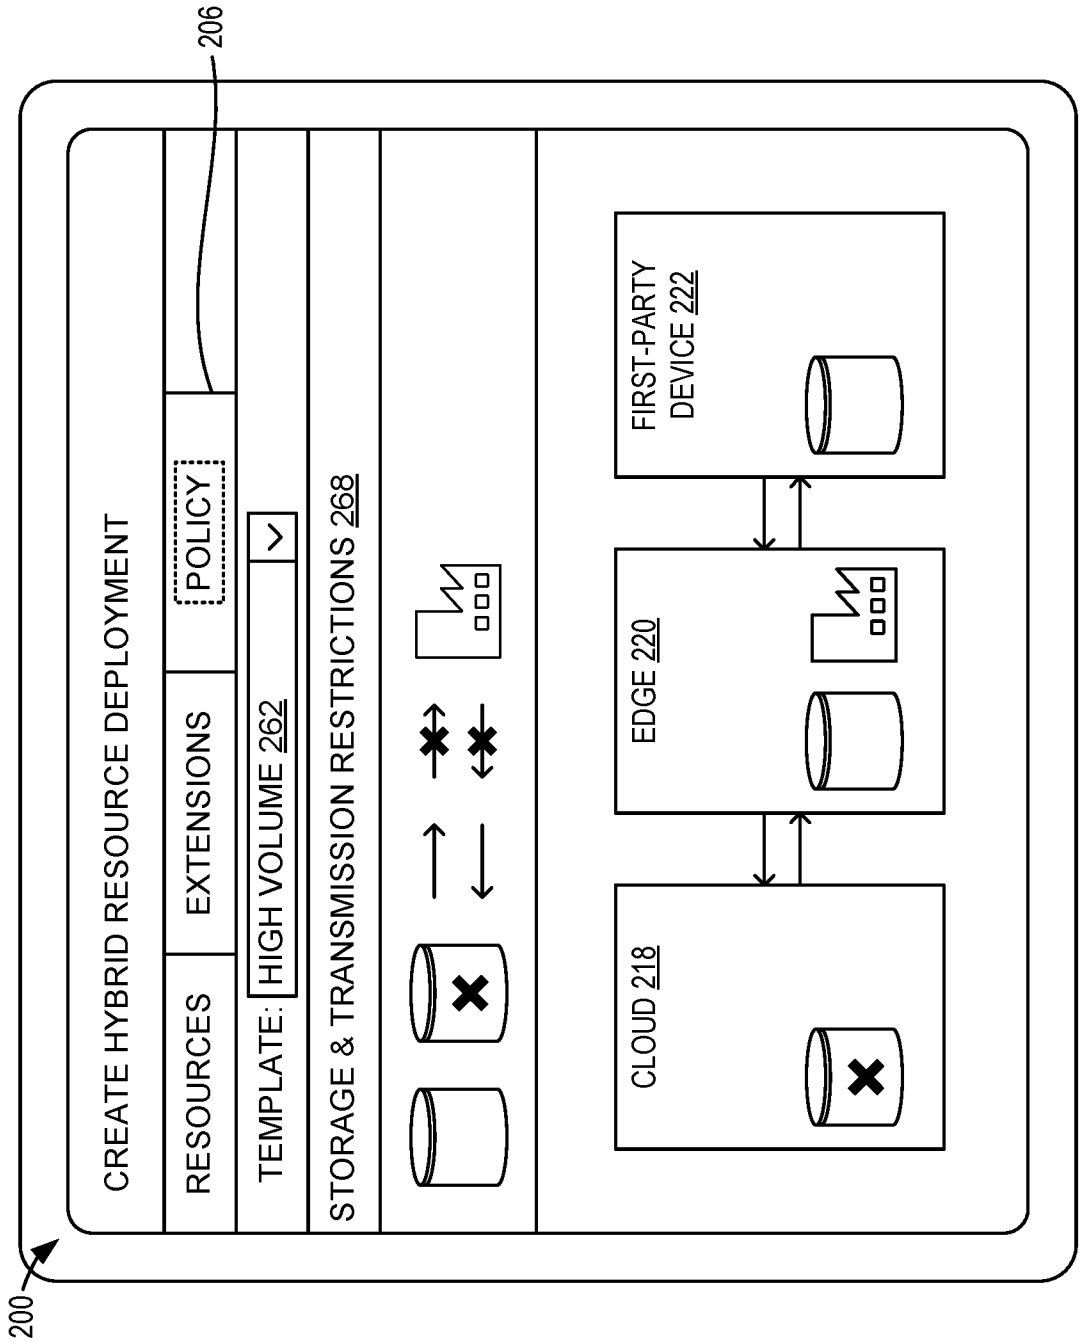
200
206
CREATE HYBRID RESOURCE DEPLOYMENT
RESOURCES
EXTENSIONS
POLICY
TEMPLATE: HIGH VOLUME 262
STORAGE & TRANSMISSION RESTRICTIONS 268
CLOUD 218
EDGE 220
FIRST-PARTY DEVICE 222

FIG. 8

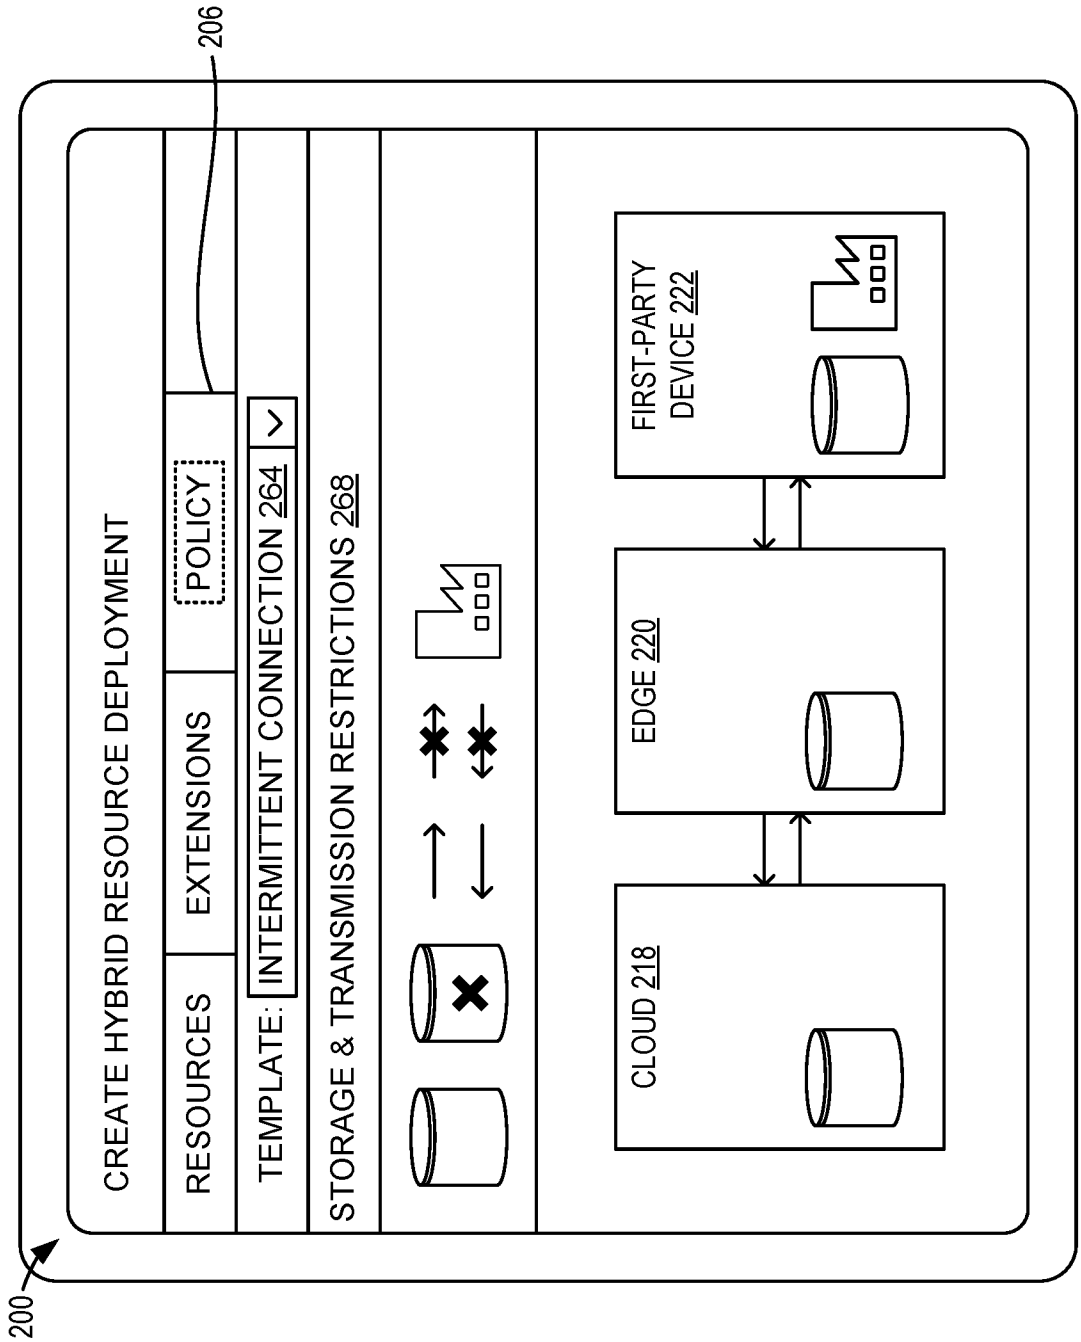
200
206
CREATE HYBRID RESOURCE DEPLOYMENT
RESOURCES
EXTENSIONS
POLICY
TEMPLATE: INTERMITTENT CONNECTION 264
STORAGE & TRANSMISSION RESTRICTIONS 268
CLOUD 218
EDGE 220
FIRST-PARTY DEVICE 222

FIG. 9

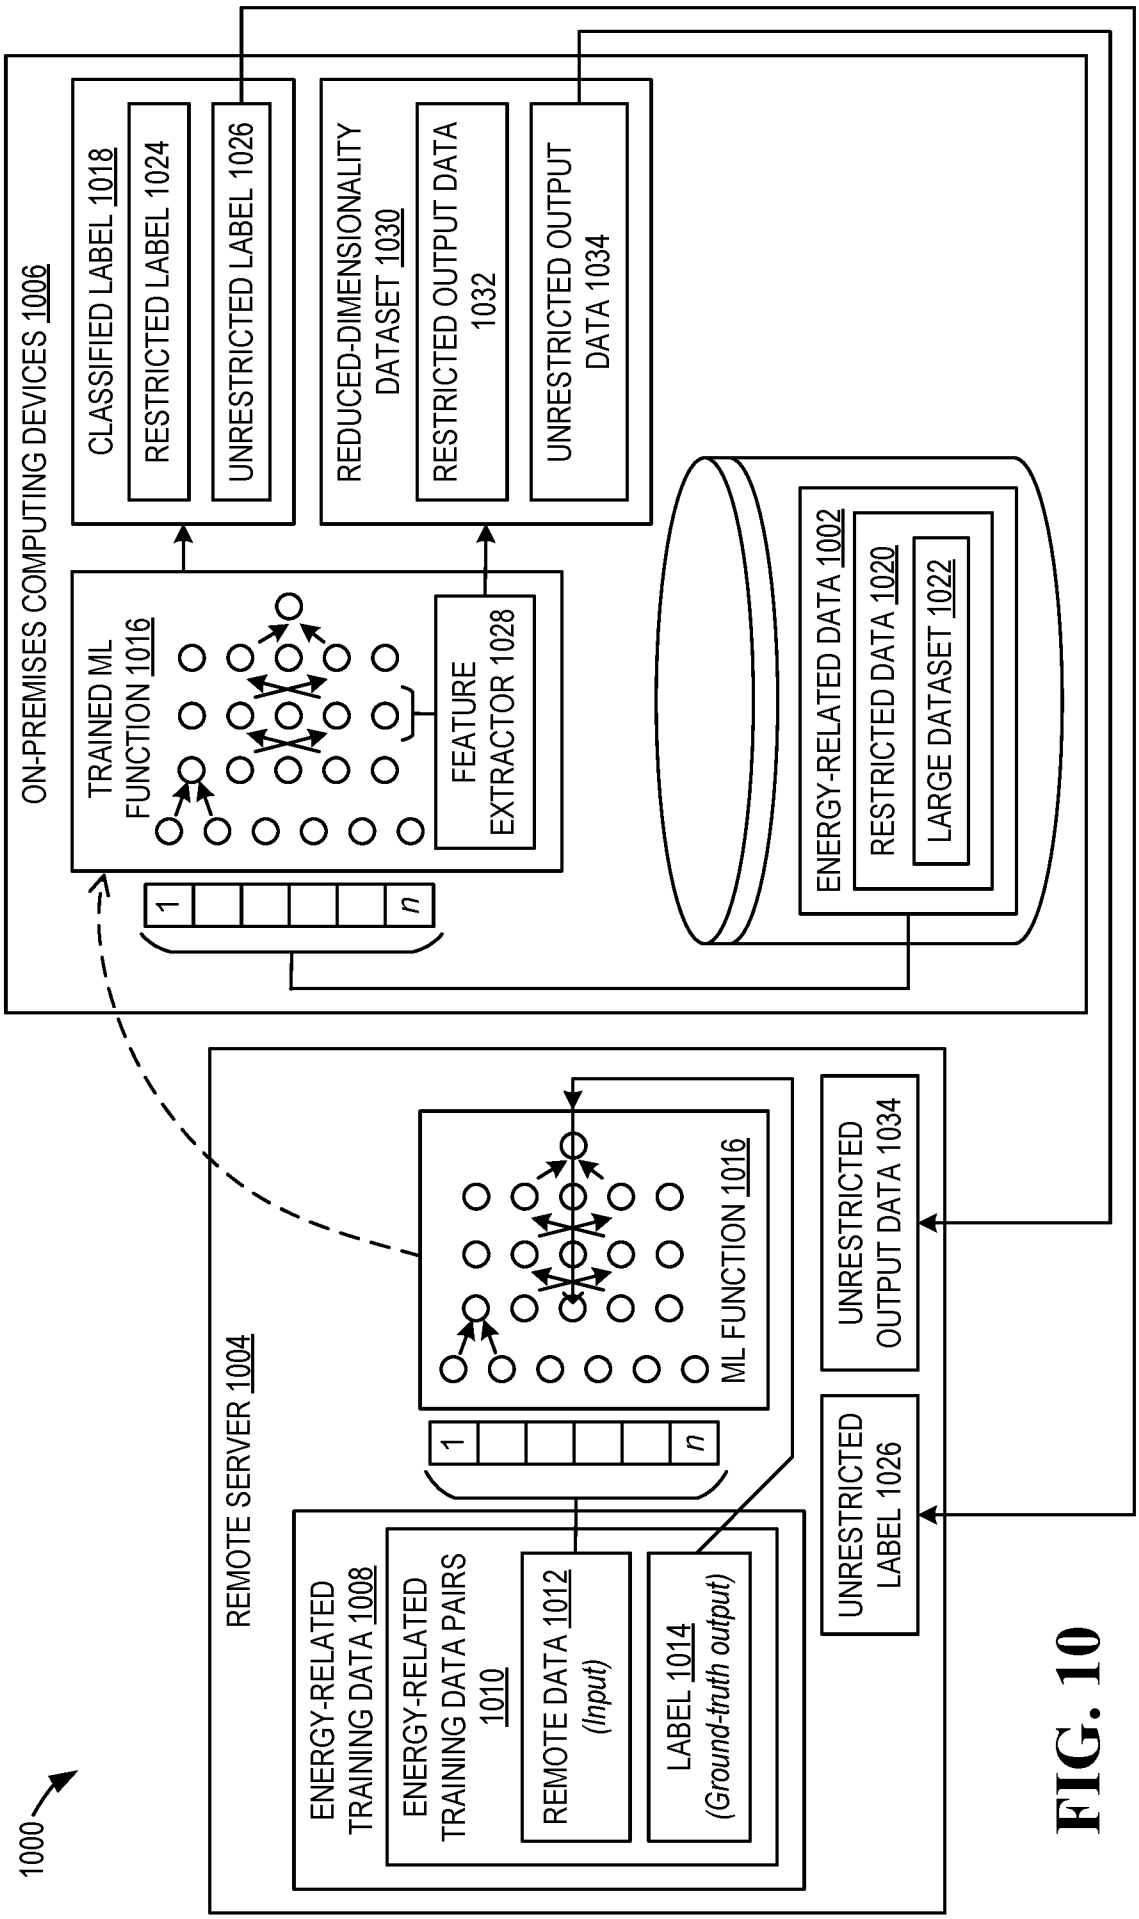
1000
ON-PREMISES COMPUTING DEVICES 1006
CLASSIFIED LABEL 1018
RESTRICTED LABEL 1024
UNRESTRICTED LABEL 1026
REDUCED-DIMENSIONALITY DATASET 1030
RESTRICTED OUTPUT DATA 1032
UNRESTRICTED OUTPUT DATA 1034
TRAINED ML FUNCTION 1016
FEATURE EXTRACTOR 1028
1
n
ENERGY-RELATED DATA 1002
RESTRICTED DATA 1020
LARGE DATASET 1022
REMOTE SERVER 1004
ML FUNCTION 1016
UNRESTRICTED OUTPUT DATA 1034
UNRESTRICTED LABEL 1026
ENERGY-RELATED TRAINING DATA 1008
ENERGY-RELATED TRAINING DATA PAIRS 1010
REMOTE DATA 1012 (Input)
LABEL 1014 (Ground-truth output)

FIG. 10

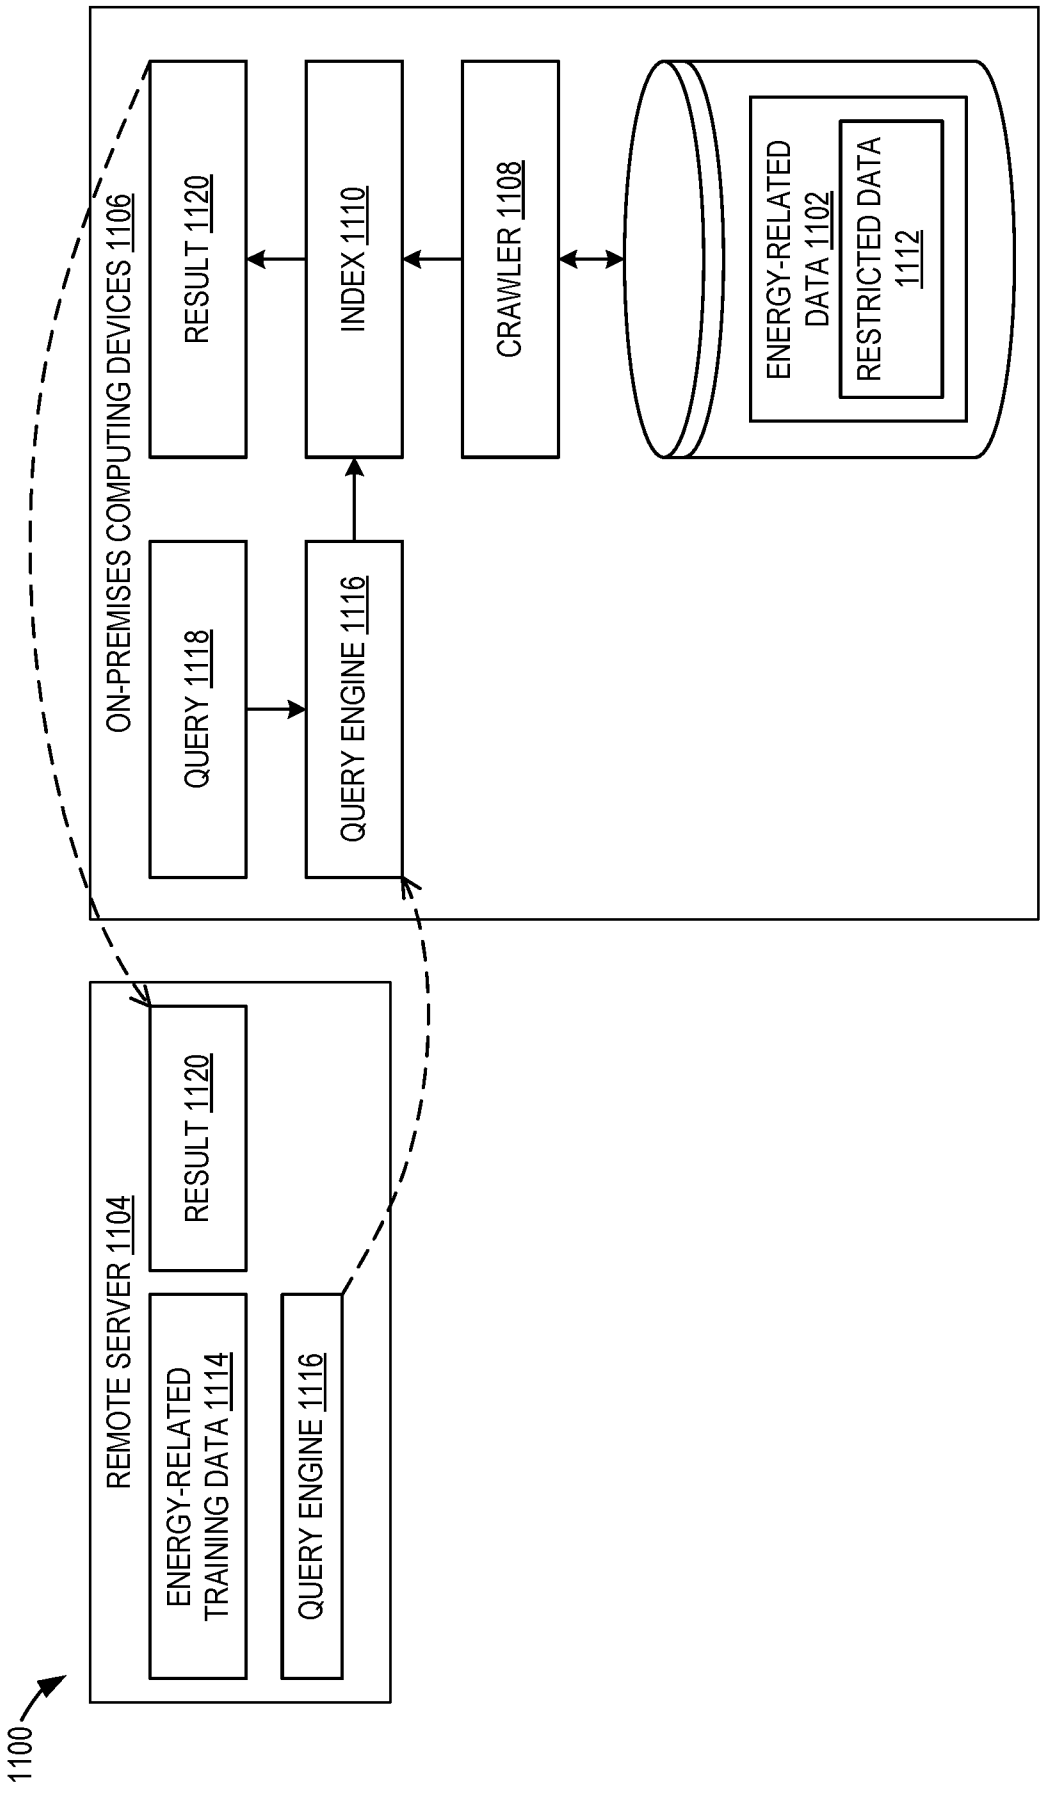
1100
REMOTE SERVER 1104
ENERGY-RELATED TRAINING DATA 1114
QUERY ENGINE 1116
RESULT 1120
ON-PREMISES COMPUTING DEVICES 1106
QUERY 1118
QUERY ENGINE 1116
RESULT 1120
INDEX 1110
CRAWLER 1108
ENERGY-RELATED DATA 1102
RESTRICTED DATA 1112

FIG. 11

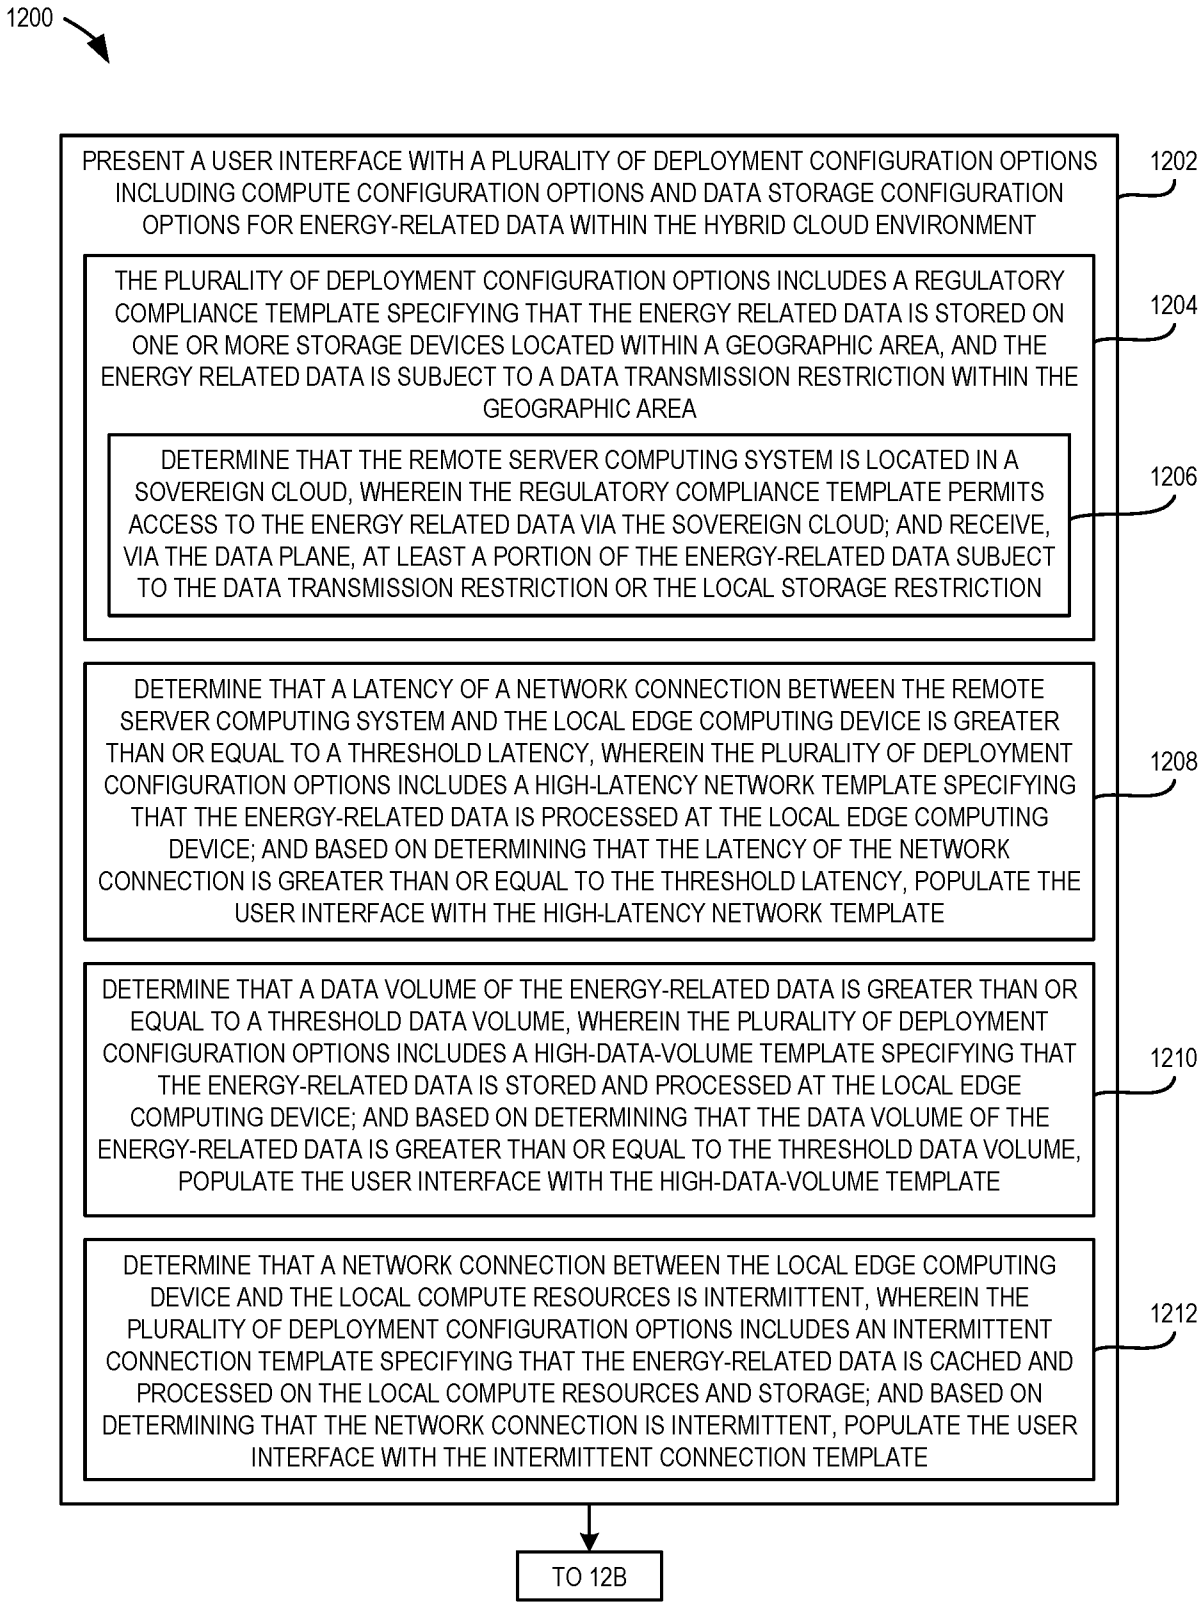
1200
PRESENT A USER INTERFACE WITH A PLURALITY OF DEPLOYMENT CONFIGURATION OPTIONS INCLUDING COMPUTE CONFIGURATION OPTIONS AND DATA STORAGE CONFIGURATION OPTIONS FOR ENERGY-RELATED DATA WITHIN THE HYBRID CLOUD ENVIRONMENT
1202
THE PLURALITY OF DEPLOYMENT CONFIGURATION OPTIONS INCLUDES A REGULATORY COMPLIANCE TEMPLATE SPECIFYING THAT THE ENERGY RELATED DATA IS STORED ON ONE OR MORE STORAGE DEVICES LOCATED WITHIN A GEOGRAPHIC AREA, AND THE ENERGY RELATED DATA IS SUBJECT TO A DATA TRANSMISSION RESTRICTION WITHIN THE GEOGRAPHIC AREA
1204
DETERMINE THAT THE REMOTE SERVER COMPUTING SYSTEM IS LOCATED IN A SOVEREIGN CLOUD, WHEREIN THE REGULATORY COMPLIANCE TEMPLATE PERMITS ACCESS TO THE ENERGY RELATED DATA VIA THE SOVEREIGN CLOUD; AND RECEIVE, VIA THE DATA PLANE, AT LEAST A PORTION OF THE ENERGY-RELATED DATA SUBJECT TO THE DATA TRANSMISSION RESTRICTION OR THE LOCAL STORAGE RESTRICTION
1206
DETERMINE THAT A LATENCY OF A NETWORK CONNECTION BETWEEN THE REMOTE SERVER COMPUTING SYSTEM AND THE LOCAL EDGE COMPUTING DEVICE IS GREATER THAN OR EQUAL TO A THRESHOLD LATENCY, WHEREIN THE PLURALITY OF DEPLOYMENT CONFIGURATION OPTIONS INCLUDES A HIGH-LATENCY NETWORK TEMPLATE SPECIFYING THAT THE ENERGY-RELATED DATA IS PROCESSED AT THE LOCAL EDGE COMPUTING DEVICE; AND BASED ON DETERMINING THAT THE LATENCY OF THE NETWORK CONNECTION IS GREATER THAN OR EQUAL TO THE THRESHOLD LATENCY, POPULATE THE USER INTERFACE WITH THE HIGH-LATENCY NETWORK TEMPLATE
1208
DETERMINE THAT A DATA VOLUME OF THE ENERGY-RELATED DATA IS GREATER THAN OR EQUAL TO A THRESHOLD DATA VOLUME, WHEREIN THE PLURALITY OF DEPLOYMENT CONFIGURATION OPTIONS INCLUDES A HIGH-DATA-VOLUME TEMPLATE SPECIFYING THAT THE ENERGY-RELATED DATA IS STORED AND PROCESSED AT THE LOCAL EDGE COMPUTING DEVICE; AND BASED ON DETERMINING THAT THE DATA VOLUME OF THE ENERGY-RELATED DATA IS GREATER THAN OR EQUAL TO THE THRESHOLD DATA VOLUME, POPULATE THE USER INTERFACE WITH THE HIGH-DATA-VOLUME TEMPLATE
1210
DETERMINE THAT A NETWORK CONNECTION BETWEEN THE LOCAL EDGE COMPUTING DEVICE AND THE LOCAL COMPUTE RESOURCES IS INTERMITTENT, WHEREIN THE PLURALITY OF DEPLOYMENT CONFIGURATION OPTIONS INCLUDES AN INTERMITTENT CONNECTION TEMPLATE SPECIFYING THAT THE ENERGY-RELATED DATA IS CACHED AND PROCESSED ON THE LOCAL COMPUTE RESOURCES AND STORAGE; AND BASED ON DETERMINING THAT THE NETWORK CONNECTION IS INTERMITTENT, POPULATE THE USER INTERFACE WITH THE INTERMITTENT CONNECTION TEMPLATE
1212
TO 12B

FIG. 12A

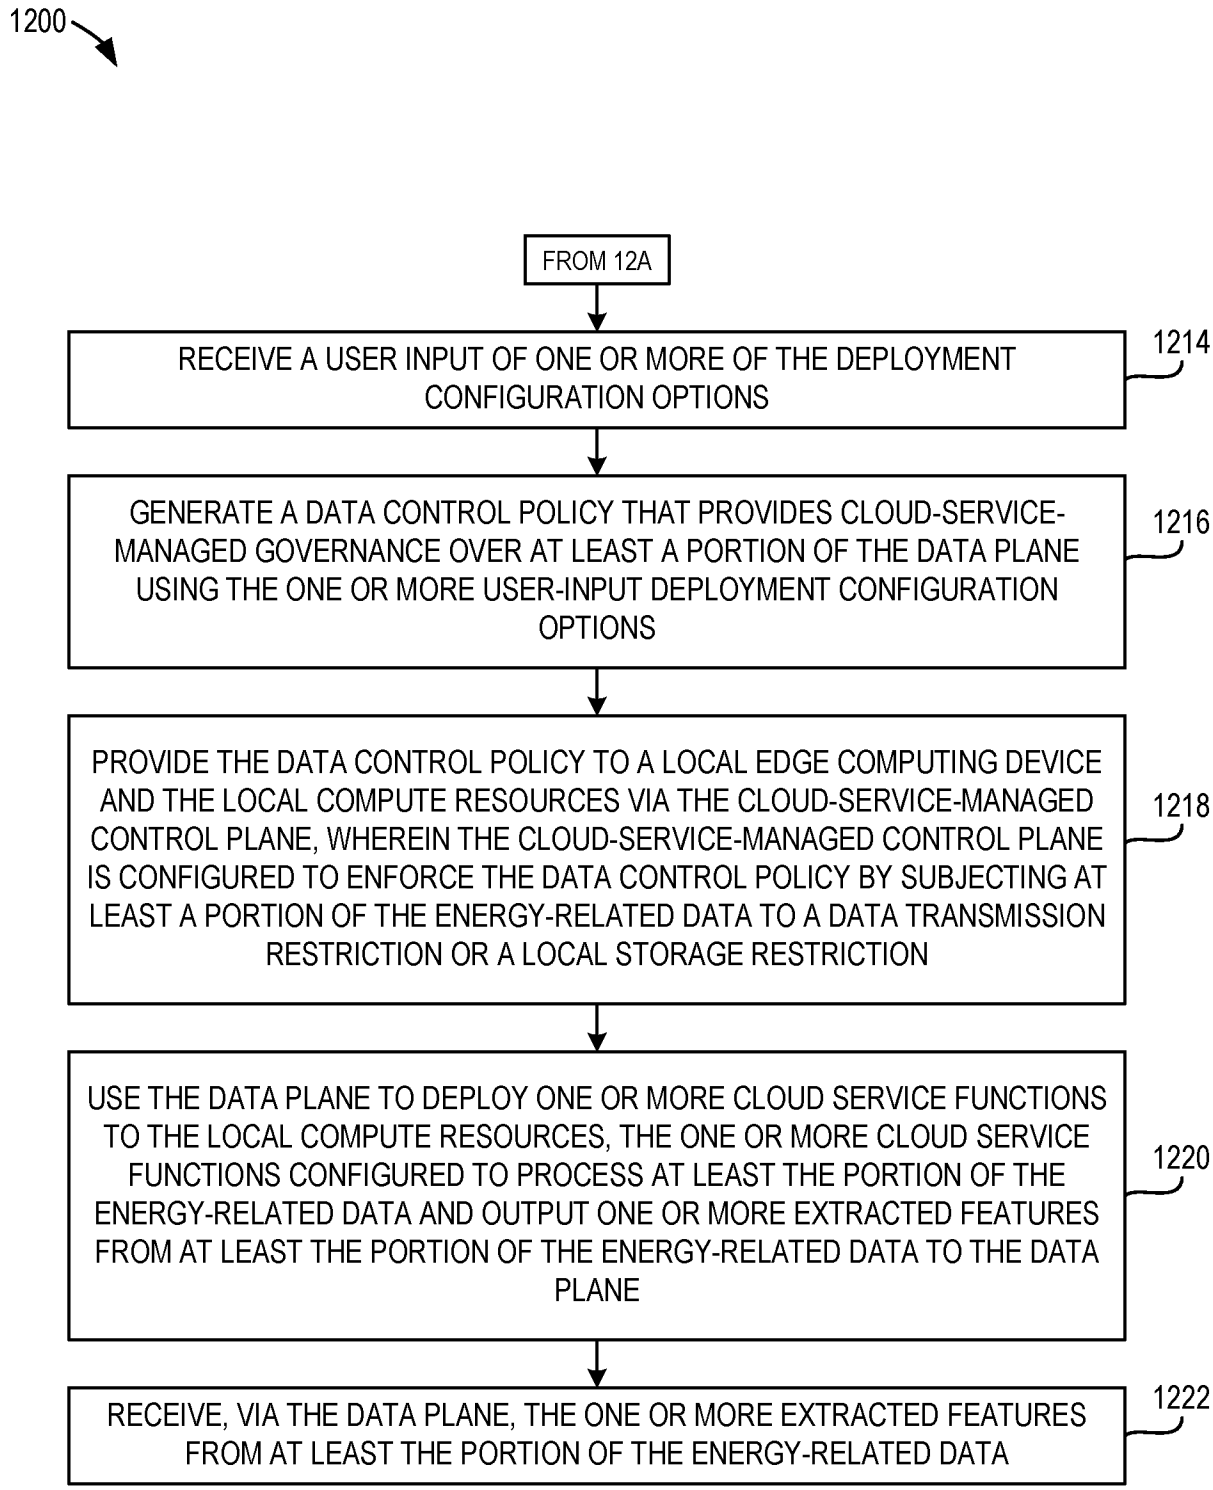
1200
FROM 12A
RECEIVE A USER INPUT OF ONE OR MORE OF THE DEPLOYMENT CONFIGURATION OPTIONS
1214
GENERATE A DATA CONTROL POLICY THAT PROVIDES CLOUD-SERVICE-MANAGED GOVERNANCE OVER AT LEAST A PORTION OF THE DATA PLANE USING THE ONE OR MORE USER-INPUT DEPLOYMENT CONFIGURATION OPTIONS
1216
PROVIDE THE DATA CONTROL POLICY TO A LOCAL EDGE COMPUTING DEVICE AND THE LOCAL COMPUTE RESOURCES VIA THE CLOUD-SERVICE-MANAGED CONTROL PLANE, WHEREIN THE CLOUD-SERVICE-MANAGED CONTROL PLANE IS CONFIGURED TO ENFORCE THE DATA CONTROL POLICY BY SUBJECTING AT LEAST A PORTION OF THE ENERGY-RELATED DATA TO A DATA TRANSMISSION RESTRICTION OR A LOCAL STORAGE RESTRICTION
1218
USE THE DATA PLANE TO DEPLOY ONE OR MORE CLOUD SERVICE FUNCTIONS TO THE LOCAL COMPUTE RESOURCES, THE ONE OR MORE CLOUD SERVICE FUNCTIONS CONFIGURED TO PROCESS AT LEAST THE PORTION OF THE ENERGY-RELATED DATA AND OUTPUT ONE OR MORE EXTRACTED FEATURES FROM AT LEAST THE PORTION OF THE ENERGY-RELATED DATA TO THE DATA PLANE
1220
RECEIVE, VIA THE DATA PLANE, THE ONE OR MORE EXTRACTED FEATURES FROM AT LEAST THE PORTION OF THE ENERGY-RELATED DATA
1222

FIG. 12B

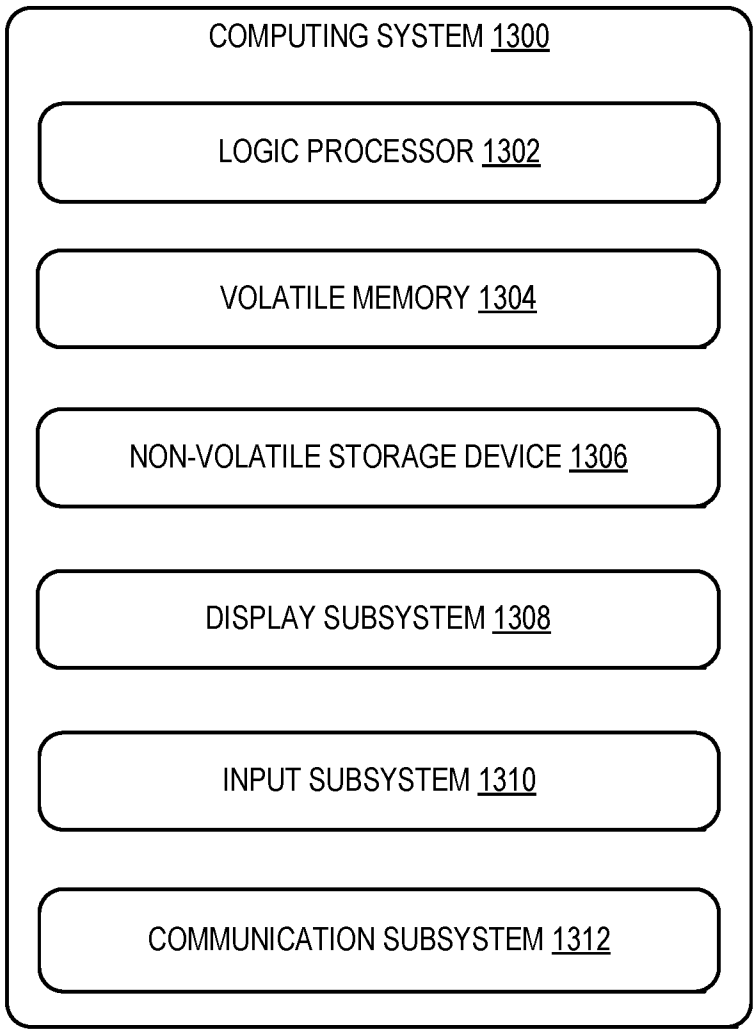
COMPUTING SYSTEM 1300
LOGIC PROCESSOR 1302
VOLATILE MEMORY 1304
NON-VOLATILE STORAGE DEVICE 1306
DISPLAY SUBSYSTEM 1308
INPUT SUBSYSTEM 1310
COMMUNICATION SUBSYSTEM 1312

FIG. 13

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2022/042323

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F21/60 G06F21/62 H04L9/40
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06F H04L

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

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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 10 897 468 B1 (CHHABRA JASMEET [US] ET AL) 19 January 2021 (2021-01-19) columns 4, 6, 9 - column 11; figure 1	1-15
A	ALHAZMI KHALED ET AL: "Optimized provisioning of SDN-enabled virtual networks in geo-distributed cloud computing datacenters", JOURNAL OF COMMUNICATIONS AND NETWORKS, NEW YORK, NY, USA,IEEE, US, vol. 19, no. 4, 1 August 2017 (2017-08-01) , pages 402-415, XP011660261, ISSN: 1229-2370, DOI: 10.1109/JCN.2017.000064 [retrieved on 2017-09-13] the whole document	1-15

☐ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance;; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance;; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
10 November 2022	18/11/2022

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

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2022/042323

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 10897468 B1	19-01-2021	NONE	