



(51) International Patent Classification:
 G06F 9/54 (2006.01) G06F 9/455 (2018.01)
 G06F 9/445 (2018.01)

(21) International Application Number:
 PCT/US2020/013631

(22) International Filing Date:
 15 January 2020 (15.01.2020)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
 62/792,580 15 January 2019 (15.01.2019) US

(71) Applicant: **AFFIRMED NETWORKS, INC.** [US/US];
 35 Nagog Park, 1st Floor, Acton, MA 01720 (US).

(72) Inventor: **PARKER, Ronald, M.**; 328 Joseph Dr.,
 Boxborough, MA 01719 (US).

(74) Agent: **HOBGOOD, John, V.** et al.; Wilmer Cutler Pickering
 Hale And Dorr LLP, 60 State Street, Boston, MA 02109
 (US).

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

(54) Title: DYNAMIC AUTO-CONFIGURATION OF MULTI-TENANT PaaS COMPONENTS

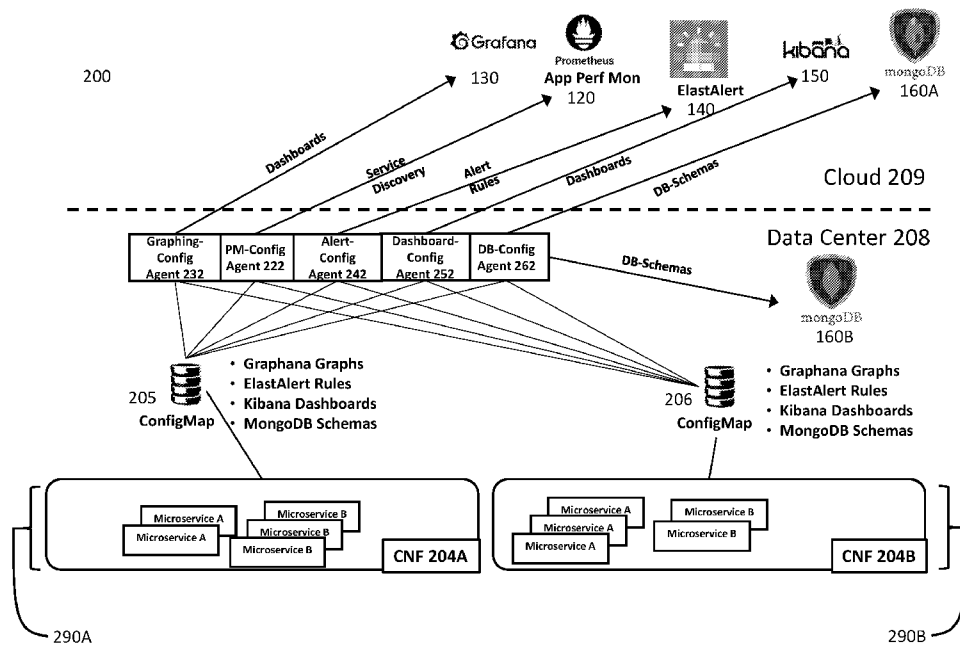


FIG. 2

(57) **Abstract:** Embodiments of the present disclosure describe configuration agents, or "config-agents," that can identify new PaaS component configuration information when a software system instance is deployed or updated without requiring a developer to manually configure each PaaS component instance whenever a new software system instance is installed or updated. In some examples, config-agents retrieve new configuration information from configuration files associated with the software system instances such as ConfigMaps and/or Custom Resource Objects of a Custom Resource Definition type operated on by the config-agents. The config-agents can keep track of each software system instance and PaaS component instance and ensure that necessary configurations are passed onto PaaS component instances.



(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).

Published:

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

DYNAMIC AUTO-CONFIGURATION OF MULTI-TENANT PAAS COMPONENTS**RELATED APPLICATIONS**

[0001] This application claims priority under 35 U.S.C. § 119(e) to U.S. Provisional Application No. 62/792,580, filed January 15, 2019, entitled “DYNAMIC AUTO-CONFIGURATION OF MULTI-TENANT PAAS COMPONENTS,” the contents of which is incorporated herein in its entirety.

TECHNICAL FIELD

[0002] The present disclosure relates to PaaS architecture, and more particularly to the process of updating PaaS components when new configurations are deployed.

BACKGROUND

[0003] In cloud computing, a user of a cloud computing service can outsource certain computing processes to third parties in order to streamline the computing experience of the user. Platform as a Service (“PaaS”) is one form of cloud computing, where a user can develop and run an application without having to create and manage each and every dependent system component of the application from the ground up. For example, a user can employ various PaaS components that are designed to perform specific tasks needed by the user’s application. Typical users of PaaS architecture include, but are not limited to, network operators that deploy Network Functions (“NFs”) with various configurations and functionalities. These NFs (e.g., Virtual Network Functions, or “VNFs”) interface with PaaS components to perform their various functions.

[0004] Example Platform as a Service (“PaaS”) components include databases such as MongoDB and observability components such as Prometheus™ (key performance indicator “KPI” management, threshold crossing alerts, and historical storage of KPIs), Grafana™ (graph visualization of KPIs), Elasticsearch™ (historical storage of log messages), ElastAlert (rule-based logic to raise alerts from log messages), Kibana™ (dashboard visualization of Elasticsearch and ElastAlert activity). These PaaS components are multi-tenant in that, when implemented to support networked systems, there may be multiple distinct systems (e.g., Network Functions) making use of each of the PaaS components simultaneously. When a new tenant (e.g., each Network Function) is deployed, configurations are applied at the PaaS

components. For example, a logical database name, list of document collections and their keys are provided to MongoDB™; lists of supported KPIs are uploaded to Prometheus™; identities of the endpoints from which the KPIs may be pulled are also be uploaded to Prometheus™ (sometimes referred to as “service discovery”); Dashboards that reference to the KPIs are uploaded to Grafana™; alerting rules are uploaded to ElastAlert; and dashboards to visualize log message flows are uploaded to Kibana™. However, coordinating this process for multi-tenant architectures is time-consuming and requires complicated manual execution or complex centralized tools like Ansible™ offered by Red Hat™ to automate changes across the platform and ensure that PaaS components receive the files they consume.

SUMMARY

[0005] In some embodiments, a method includes installing a service instance using a first descriptor chart, the installing comprising: installing a plurality of microservices declared in the first descriptor chart that comprise the service instance, at least one of the microservices configured to request services from an instance of a Platform-as-a-Service (“PaaS”) component of a PaaS component type, and creating at least one configuration file based on the first descriptor chart; parsing, by a configuration agent associated with the PaaS component type, the at least one configuration file for configuration data associated with the PaaS component; and instructing, by the configuration agent, the instance of the PaaS component to install the configuration data.

[0006] In some embodiments, the configuration file comprises descriptors of resources, and wherein the method further comprises mounting the configuration file to the configuration agent.

[0007] In some embodiments, the service instance comprises a network function (“NF”).

[0008] In some embodiments, the service instance is installed by a network operator, and the PaaS component is developed and managed by a third-party.

[0009] In some embodiments, the method further includes updating the service instance using a second descriptor chart, the updating comprising: updating at least one of the plurality of microservices declared in the second descriptor chart that comprise the service instance, the updated at least one of the plurality of microservices configured to request services from the instance of the PaaS component, and updating the at least one configuration file based on

the second descriptor chart; parsing, by the configuration agent, the at least one updated configuration file to identify the update to the configuration data associated with the PaaS component; and instructing, by the configuration agent, the instance of the PaaS component to install the updated configuration data.

[0010] In some embodiments, the updating the software instance is performed by a container orchestration engine (“COE”).

[0011] In some embodiments, a method includes: monitoring, by a configuration agent associated with a Platform-as-a-Service (“PaaS”) component type, for a notification identifying a configuration file change in at least one configuration file directory, the at least one configuration file in the configuration file directory containing configuration data for at least one instance of the PaaS component of the PaaS component type; monitoring, by the configuration agent, a configuration of the at least one instance of the PaaS component type to identify a lack of conformity with configuration data in the configuration file directory; automatically transmitting, by the configuration agent, an instruction to the at least one instance of the PaaS component when one or more of a configuration file change or a lack of conformity is identified, the instruction comprising an instruction to install a configuration on the at least one instance of the PaaS component based on the configuration data in the configuration file directory.

[0012] In some embodiments, the at least one instance of the PaaS component is configured to service at least one microservice of a plurality of microservices comprising a network function (“NF”).

[0013] In some embodiments, the change in configuration data comprises a change in data of a ConfigMap in the configuration file directory.

[0014] In some embodiments, the configuration data is associated with at least one service instance, and wherein the method further comprises updating the at least one service instance using a descriptor chart, the updating the at least one service instance causing the change in the configuration data.

[0015] In some embodiments, the at least one service instance is updated by a network operator, and the PaaS component is developed and managed by a third-party.

[0016] In some embodiments, a method includes: installing a service instance using a first descriptor chart, the installing comprising: installing a plurality of microservices declared in the first descriptor chart that comprise the service instance, at least one of the microservices configured to request services from an instance of a Platform-as-a-Service (“PaaS”) component of a PaaS component type, and creating at least one resource instance based on the first descriptor chart, the resource instance having a type associated with a resource definition, the resource instance containing configuration data associated with the instance of the PaaS component; receiving, by a configuration agent that operates resource instances of the type associated with the resource definition, a notification of the creation of the at least one resource instance, the notification further comprising the configuration data associated with the instance of the PaaS component; and instructing, by the configuration agent, the instance of the PaaS component to install the configuration data.

[0017] In some embodiments, the instructing the instance of the PaaS component to install the configuration data comprises performing, by the configuration agent, an API call to the instance of the PaaS component.

[0018] In some embodiments, the service instance comprises a network function (“NF”).

[0019] In some embodiments, the service instance is installed by a network operator, and the PaaS component is developed and managed by a third-party.

[0020] In some embodiments, the method further includes: updating the service instance using a second descriptor chart, the updating comprising: updating at least one of the plurality of microservices declared in the second descriptor chart that comprise the service instance, the updated at least one of the plurality of microservices configured to request services from the instance of the PaaS component, and updating the at least one resource instance based on the second descriptor chart by updating configuration data in the resource instance associated with the instance of the PaaS component; receiving, by the configuration agent that operates resource instances of the type associated with the resource definition, a notification of the update to the at least one resource instance, the notification further comprising the updated configuration data associated with the instance of the PaaS component; and instructing, by the configuration agent, the instance of the PaaS component to install the updated configuration data.

[0021] In some embodiments, the updating the software instance is performed by a container orchestration engine (“COE”).

[0022] In some embodiments, a method includes: monitoring, by a configuration agent associated with a resource definition type, for a notification identifying that a declaration in a resource instance of the resource definition type has changed, the changed declaration being associated with at least one instance of a Platform-as-a-Service (“PaaS”) component; monitoring, by the configuration agent, a configuration of the at least one instance of the PaaS component to identify a lack of conformity with the declaration in the resource instance of the resource definition type; automatically transmitting, by the configuration agent, an instruction to the at least one instance of the PaaS component when one or more of a declaration change or a lack of conformity is identified, the instruction comprising an instruction to install a configuration on the at least one instance of the PaaS component based on the declaration in the resource instance.

[0023] In some embodiments, the at least one instance of the PaaS component is configured to service at least one microservice of a plurality of microservices comprising an instance of a network function (“NF”).

[0024] In some embodiments, the method further includes updating the instance of the NF using a descriptor chart, the updating the instance of the NF causing the change in the declaration.

[0025] In some embodiments, the instance of the NF is updated by a network operator, and the PaaS component is developed and managed by a third-party

[0026] In some embodiments, the changed declaration comprises a change to a consumable of the at least one instance of the PaaS component.

BRIEF DESCRIPTION OF FIGURES

[0027] Various objectives, features, and advantages of the disclosed subject matter can be more fully appreciated with reference to the following detailed description of the disclosed subject matter when considered in connection with the following drawings, in which like reference numerals identify like elements.

[0028] FIG. 1 is a system diagram showing a configuration of a prior art multi-tenant PaaS component configuration using Ansible™.

[0029] FIG. 2 is a system diagram showing an exemplary auto-configuration architecture for multi-tenant PaaS components, according to some embodiments.

[0030] FIG. 3 shows a call flow diagram of an exemplary PaaS auto-configuration during cloud-native network function (“CNF”) deployment, according to some embodiments.

[0031] FIG. 4 shows a call flow diagram for a PaaS auto-configuration during a microservice upgrade, according to some embodiments.

[0032] FIG. 5 shows a call flow diagram of an exemplary PaaS auto-configuration during cloud-native network function (“CNF”) deployment, according to some embodiments.

[0033] FIG. 6 shows a call flow diagram for a PaaS auto-configuration during a microservice upgrade, according to some embodiments.

DETAILED DESCRIPTION

[0034] It is to be understood that the phraseology and terminology employed herein are for the purpose of description and should not be regarded as limiting.

[0035] In multi-tenant PaaS implementations, individual instances of software systems, such as NFs, can be installed as needed based on a descriptor file and input variables. Such software systems can comprise an amalgamation of microservices, each of which performs a particular service for the software system. Deployments of such software systems can request particular instances of PaaS components to provide services to their software systems based on custom configurations provided to the PaaS component instances using configuration files. As developers scale both the number of software system instances, the number of microservices in such instances, and number of instances of PaaS components that provide services to the software system instances, managing configuration information for each instance of software and PaaS components becomes cumbersome and error prone. While configuration managers such as Ansible™ can facilitate deployment and updates to such implementations, as described below, these solutions require manual scripting to keep track of configuration information for each PaaS component instance. Accordingly, embodiments of the present disclosure describe configuration agents, or “config-agents,” that

can identify new PaaS component configuration information when a software system instance is deployed or updated without requiring a developer to manually configure each PaaS component instance whenever a new software system instance is installed or updated. In some examples, config-agents retrieve new configuration information from configuration files associated with the software system instances such as ConfigMaps and/or Custom Resource Objects of a Custom Resource Definition type operated on by the config-agents. The config-agents can keep track of each software system instance and PaaS component instance, for example by using a namespace of each, and ensure that necessary configurations are passed onto PaaS component instances.

[0036] FIG. 1 is a system diagram showing a configuration of a prior art multi-tenant PaaS component configuration 100 using Ansible™. Configuration 100 includes Ansible™ 101, Kubernetes™ 170, and Helm™ 176, and PaaS components such as Prometheus™ 120, Grafana™ 130, ElastAlert 140, Kibana™ 150, MongoDB™ 160A/B. These components are described in more detail below. It should be appreciated that these components are merely exemplary, and can be replaced with other PaaS components that perform similar or different functions.

[0037] In general, a software system, such as a network function in a networked system, can communicate with PaaS components 120, 130, 140, 150, 160A/B to request the performance of particular services. Third-party deployers can manage PaaS component instances, which can operate in the public cloud. Users of the PaaS components, such as network operators, can customize PaaS components to serve their applications without having to create, manage, and maintain each PaaS component. Instead, specific instances of PaaS components can be deployed for use by particular software system instances (e.g., NFs), which are customized based on configuration files that are sometimes referred to by different names such as “dashboards,” “schemas,” etc. Thus, these PaaS components can be multi-tenant, in that multiple software systems (e.g., NFs) can use one or more instances of these PaaS components based on particular configurations fed to the PaaS components.

[0038] Helm™ 176: an open source system orchestrator. Helm™ 176 is a package manager that manages software deployments, such as deployments of NFs. Helm™ 176 takes descriptor files, referred to as Helm™ Charts, as input and decomposes them into individual Kubernetes™ commands. For example, a Helm™ Chart can be deployed with a command such as ‘helm install <chart.name> <variable values>...’, where <chart.name> is

the name of a particular chart, and the <variable values> are variables that are used to customize a particular instance of a software system generated from the Helm™ Chart. Multiple instances of software systems can be generated from the same Helm™ Chart using different variable values. According to some embodiments, a Helm™ Chart is structured like a Network Function (“NF”) descriptor. It can contain the instructions for Kubernetes™ (a container orchestration engine described below) to deploy the microservices that comprise the NF. For example, the Helm™ Chart can read as a set of declarations that declare various microservices, etc. that will make up a software system, such as a NF. This can include elements such as, but not limited to, versioned docker image file names, as well as metadata such as replica counts, auto-scale policies, service definitions (e.g., an abstraction to hide the individual members of the a microservice type’s replica set), and any Custom Resource Definitions (“CRDs,” which can be arbitrary and changeable data that can be injected by the COE (“Core Orchestration Engine,” e.g., Kubernetes™, Docker Swarm™, or Mesosphere™) to the running microservices). The docker images can be used to deploy microservices (e.g., pods) that make up the NF. The Helm™ Chart can also include additional data and files that are necessary to configure the microservices. In addition, a ConfigMap, which can be a dictionary of configuration settings that contains configuration information to be read from microservices (e.g., as a key-value pair of strings), can be declared in the Helm™ Chart itself. The Helm™ Chart can also indicate to which microservices the ConfigMaps should be mounted (as a file) such that the microservices can see the configuration information contained in the respective ConfigMap. As described in more detail below, unlike in FIG. 1, according to some embodiments, the metadata in a Helm™ Chart can be used to specify certain files to be consumed by various PaaS components. ConfigMaps allow for configuration files to be separated from image content in order to make microservice deployment more flexible. Furthermore, as described in more detail below with reference to FIGS. 2-6, by including these files in the Helm™ Chart, these files may be passed through to the various PaaS components without requiring coordination software like Ansible. Although Helm™ 176 and Helm™ Charts are discussed in the present disclosure, it should be appreciated that other types of orchestrators that allow for specification of docker images for deploying software and metadata associated with the same can be used.

[0039] Kubernetes™ 170: an open source container orchestration engine (“COE”) for automating deployment. Kubernetes™ can assist in deploying, scaling, and managing software systems. For example, Kubernetes™ can manage addressing, scheduling, and

scaling of instances of services, which can be made up of multiple pods containing containers.

[0040] Prometheus™ 120: an open source KPI /Statistics manager. Prometheus™ 120 includes a Time Series Database (“TSDB”) for intermediate term storage. Prometheus™ 120 can be regularly synchronized with list of endpoints from which KPIs are periodically gathered. Service discovery can refer to the process by which Prometheus™ 120 discovers various endpoints (e.g., nodes) from which to collect metrics. An instance of Prometheus™ 120 that serves a particular NF needs to be provided with this configuration information in order to learn about which contain KPIs to be gathered.

[0041] Kibana™ 150: an open source dashboard engine that allows visualization of log flows and alerts. Multiple instances of Kibana™ 150 can be deployed to serve various software systems, such as NFs. Each Kibana™ 150 consume a dashboard file containing information about how to generate a particular dashboard, which it can use to construct queries to ElasticSearch for data to visualize. According to some embodiments, dashboard files can be created manually via the Kibana tool in Kibana™ 150. A user can create a dashboard, add widgets to it, specify the log filters that can be applied to determine which logs match to the added widget, and finally save the dashboard. Alternatively, network operator can construct a set of these dashboards to accomplish particular desired functionalities and then add them to a source code control system such that they are released along with the software. In some embodiments, a release of software is comprised of at least two things: a set of versioned docker images and the Helm™ Chart for the particular NF. Thus, the dashboards can be declared within the Helm™ Chart. The packaging of the Helm™ Chart can be done during the build of the NF release.

[0042] ElastAlert 140: an open source plugin to Kibana™ that can receive configuration information in the form of alert rule files which it uses to construct appropriate queries to ElasticSearch™, which is an intermediate term storage for log messages. According to some embodiments, each microservice emits log messages, for example, by processes such as “printf” or “stdout” file descriptors.

[0043] MongoDB™ 160A/B: an open source NoSQL database. MongoDB™ 160A/B creates particular instances of logical databases based on configuration information contained in schema files. As explained below, schema files can delineate the database structure and

corresponding nodes from which data is to be collected. Each instance of MongoDB™ 160A/B can be for a particular software system, such as a NF. Instances of MongoDB™ 160A/B can collect key information via API calls to various nodes in the system, based on information contained in the schema. According to some embodiments, when a new NF is deployed (e.g., via a ‘helm install’ command), an instance of MongoDB™ 160A/B can be pre-instructed regarding a logical DB name that the NF expects to use, document collections (“repos”) within the logical DB, and the primary/secondary key information for each repo based on a schema file. APIs can be offered by MongoDB™ 160A/B within the set of TCP-based APIs defined by MongoDB™ 160A/B. After such steps have been accomplished, microservices within the newly orchestrated NF can write and read documents to MongoDB™ 160A/B using the proprietary TCP-based APIs that are defined by MongoDB™ 160A/B.

[0044] As shown in FIG. 1, in prior art systems, Ansible™ 101, as a higher-level configuration manager, is responsible for deploying new distinct software systems such as NFs and pushing all relevant configurations to instances of the PaaS components used by that new distinct software system. Ansible™ 101 does so via Playbooks 110, which are scripts that are written to conform to Ansible™ 101 requirements. When a new software system is deployed, Ansible™ 101 may need to push down information to each PaaS component instance so that each PaaS component instance is aware of what the new system requires of it. For example, Ansible™ 101 may push Helm™ Create 182 from a Helm™ chart to Helm™ 176, install new dashboards (e.g., files with information about graphing templates) 183 for Grafana™ 130, install new alerting rules 184 for ElastAlert 140, install dashboards 185 for Kibana™ 150, and configure database(s) 186 for MongoDB™ 160A/B, all in new instances of each PaaS component that are meant to serve a particular NF deployed via the Helm™ Chart. Some or all of these configurations are declared in playbooks 110. Playbooks take the form of scripts. These scripts can be prepared by an IT or other operational professional. Thus, in order to ensure that PaaS component instances receive the correct configuration to provide services to a new NF instance, someone must manually script in a Playbook that the information is to be provided to particular PaaS component instances. The construction and maintenance of these scripts in Playbooks can be elaborate and time consuming, and make the process of updating various components or adding new NFs/microservices more labor intensive.

[0045] Higher level automation software like Ansible™ 101 present a number of drawbacks. For example, when deployed at production scale, it becomes awkward, inefficient, and error-prone to utilize this approach. Deployment and automation solutions like Ansible™ 101 and its Playbooks 110 must capture significant software system-specific knowledge in order to know exactly how to coordinate the deployment of new software and the update of the respective PaaS components that the new software will utilize. For example, someone must specify a particular Helm™ Chart from which to deploy the software instance, and further specify, as discussed above, how each PaaS component instance is to learn of configuration to service that particular software instance. It should be appreciated that the PaaS component instances can be instantiated using any suitable methods including, but not limited to being deployed via a Helm™ Chart or being created by inputting data via a graphical user interface or via an API called from a script. Thus, when deploying a new software instance, someone must configure the Playbook to explicitly configure PaaS components to create and/or update PaaS component instances. Furthermore, each time an existing system is upgraded, the Ansible™ Playbooks 110 potentially need to be updated. For example, for new deployments or if a software update requires a new or updated PaaS component instance, the Playbooks 110 will need to be updated to specify that that PaaS component instance is to receive the correct configuration information. In addition, the Ansible Playbooks 110 do not sufficiently handle the needs of multiple versions of distinct systems. Accommodating different versions that may have different configuration needs with respect to the multi-tenant PaaS components is challenging. Versions can refer to either the version of a particular PaaS components or the version of a particular NF. According to some embodiments, a PaaS component instance can change versions and thus change its approach to configuration, for example using different APIs or different structure of the data it consumes. Such different data structures can relate, for example to the file format of a Kibana™ dashboard, for example changing from JSON to YAML, or simply adding new fields. For versioning of a NF, a new microservice version may add new log messages, which can require new versions of ElastAlert rules files and/or Kibana™ dashboards to accommodate such log messages. Another example can involve the introduction of new KPIs. According to some embodiments, a new microservice version can modify the way it interacts with a database, for example by introducing new collections or adding more secondary keys to existing collections. According to some embodiments, an upgrade to a NF can employ a new type of microservice that did not previously exist. While a NF can be

upgraded to a new version using a ‘helm upgrade’ command, someone must still add updated configuration information to be provided to PaaS components in the Playbook 110.

[0046] There are further drawbacks to this system. On a significant number of prior art deployments, there is no universal scripting environment. Instead, before or after deploying a particular software system (e.g., an NF), configuration or reconfiguration of the multi-tenant PaaS components is performed manually via a CLI or GUI interface. This creates significant disadvantages for production scale environments, where many manual configurations are time-consuming and difficult to coordinate. In addition, particular deployments can require additional debugging time.

[0047] FIG. 2 is a system diagram showing a multi-tenant PaaS component configuration 200 with config-agents, according to some embodiments. The configuration 200 includes similar components to the system 100, including Prometheus™ 120, Grafana™ 130, ElastAlert 140, Kibana™ 150, and MongoDB™ 160A and 160B. The system 200 also includes CNF 204A, CNF 204B, microservices 290A, microservices 290B, ConfigMap 205, ConfigMap 206 and graphing-config-agent 232, PM-config-agent 222, alert-config-agent 242, dashboard-config-agent 252, and database (“DB”)-config-agent 262. These elements are described in more detail below. It should be appreciated that these elements are merely exemplary, and can be replaced with other PaaS components that perform similar or different functions. It should also be appreciated that the config-agents can be structured as separate microservices on each CNF 204A, 204B.

[0048] As shown in FIG. 2, one or more NFs or other software systems (e.g., CNFs 204A, 204B) can be implemented using a microservices architecture. Accordingly, each of CNFs 204A, 204B comprise at least one microservice 290A, 290B, respectively, that provide one or services (e.g., functional logic) via API calls on behalf of the CNF. In some embodiments, a CNF can be considered to be a federation of microservices, each of which performs one service associated with the function of the CNF. Services can include, for example, functional logic such as call control logic or MME tracking area management logic. According to some embodiments, microservices 290A, 290B converse through APIs (for example, RESTful over HTTP or gRPC). According to some embodiments, the microservices 290A, 290B can be dynamically and multiply instantiated based on network load. As discussed above, a COE, such as Kubernetes™, can manage microservices for each CNF.

[0049] In some embodiments, the CNFs 204A, 204B can operate on a PaaS architecture, and can therefore each access instances of PaaS components, for example, Prometheus™ 120, Grafana™ 130, ElastAlert 140, Kibana™ 150, and MongoDB™ 160A/B, to accomplish particular functions associated with each PaaS component. According to some embodiments, PaaS components instances themselves can also be dynamically scalable based on the load which they receive.

[0050] In some embodiments, each CNF 204A, 204B includes at least one ConfigMap 205, 206, respectively. ConfigMaps 205, 206 can be used in implementations that employ Kubernetes. In such implementations, a ConfigMaps can be specified in a Helm™ Chart and implemented for each deployed CNF. According to some embodiments, each deployed CNF can have many ConfigMaps. A ConfigMap can be, in some embodiments, a resource that allows for injection of configuration data into a pod or other resource. ConfigMaps can be separated so as to be mounted separately to distinct microservice types. In an example, the type of data used by the DB-config-agent can be located in one ConfigMap which includes the collection of DB requirements files of each microservice type in the CNF that makes use of the DB. Similarly, according to some embodiments, the collection of Kibana™ dashboards that can be released with each microservice type can be orchestrated into their own ConfigMap for purposes of mounting it to the Dashboard-config-agent. According to some embodiments, mounting can denote that a ConfigMap is visible in a microservice's local files system at some specified coordinate, such as `/etc/config-maps/dashboard-config-map/<actual-config-map-file>`.

[0051] In some embodiments, the contents of the ConfigMaps 205, 206 can be mounted to each microservice 290A, 290B, respectively, as a file system containing the files that were present in a Helm™ Chart. The Helm™ Chart can include raw data that can be put into a ConfigMap, such as, for example, the Kibana™ dashboard files (i.e., configuration information for an instance of a Kibana™ PaaS component), and can also include instruction that such data be assembled into a ConfigMap and a directive that one or more microservice types should mount the ConfigMap at a given file system path. According to some embodiments, the ConfigMap is a general mechanism that is similar to an “environment variable” but, that can be changed during the lifetime of the running microservice. According to some embodiments, the ConfigMaps discussed throughout the present disclosure can instead be replaced by Custom Resource Definitions (“CRDs”) in an API-oriented solution

rather than a file-oriented solution, as described in more detail with reference to FIGS. 5 and 6.

[0052] According to some embodiments, CNFs 205, 206 can be associated config-agents associated with each PaaS component, such as Prometheus™ 120, Grafana™ 130, ElastAlert 140, Kibana™ 150, and MongoDB™ 160A/B, of the PaaS component configuration 200. Each config-agent can be managed by a separate entity from the PaaS component and can be singleton microservice that is deployed as a member of each CNF 204A, 204B that can make API calls to the PaaS component instances associated with the config-agent. As described in more detail with reference to FIGS. 3-4, each config-agent is configured to extract files from the mounted ConfigMap 205, 206 and recognizes dynamic changes in the structure thereof (e.g., additions, changes, or deletions of configuration information). The extracted files can be renamed to include the CNF ID. This can be performed, for example, to ensure that there are not any naming collisions. For example, two CNFs may have different life cycles (e.g., upgrades at different times), which means their metadata may diverge. Accordingly, by naming files based on the CNF instance name, the life cycle of these files can also be independently controlled separately from other CNF instances. The files can then be installed in each PaaS component using API calls to enable the operation of each respective component. For example, a config-agent can pass changes to configuration information for a specific instance of a PaaS component associated with a CNF via an API call to that PaaS component instance. Accordingly, when the ConfigMap is updated during an upgrade, the config-agents can identify the updated portion of the ConfigMap and synchronize the relevant PaaS component instances with the update.

[0053] According to some embodiments, CRDs can be used in addition to or instead of ConfigMaps. A CRD can be used to define a consumable of a PaaS component, such as a schema for MongoDB™. The CRD can be provided to a COE, such as Kubernetes™. Once defined, the CRD becomes one of the types of resources that the COE (e.g., Kubernetes™) can understand.

[0054] In some embodiments, particular resource instances, such as Custom Resource Objects (“CROs”) can be deployed via a Helm™ Chart. These CROs are of a type defined in the CRD, and can contain specific information therein, such as configuration information for a PaaS component instance. For example, when software or upgrades are deployed, the Helm™ Chart can include declarations of various CROs, which can include a CRD type

(already understood by Kubernetes™, name, and indication to Kubernetes™ 170 that it exists as a CRO. Furthermore, in the Helm™ Chart, certain files can be specified to be included in the CRO which include, for example, dashboard files to be consumed by a PaaS component instance such as Grafana™ 130. In another example, a particular schema CRO can be created of the type defined in a schema CRD, and can contain information about a database (e.g., KPIs, etc.) for a MongoDB™ instance. Each CRO can contain multiple files or a single file. CRDs operate using a publish and subscribe model where publishing occurs at deployment of the CROs, for example, when they are specified in a Helm™ Chart acted upon by Helm™. Although the present disclosure describes operation of Config-Agents based on information in CROs, ConfigMaps, or both, a person of skill in the art would understand from the present disclosure that Config-Agents can operate in other software environments using other types of resources intended to be bound to particular instances.

[0055] In some embodiments, config-agents can be created to operate CROs of certain CRD types in order to provide new, changed, or deleted configuration information to a PaaS component instance. In such embodiments, a microservice, such as a config-agent, can be instantiated and configured to inform the COE (e.g., Kubernetes™) that it operates resources of the type defined by a CRD (i.e., is subscribed to CROs of the CRD type). For example, a developer of the PaaS- config-agent microservice can document the type, name, and structure of the CRD that it consumes. When a new CRO of the CRD type is deployed, or when a CRO of that type is updated, in a Helm™ Chart, the COE (e.g., Kubernetes™ 170_ can make the CROs, including the files referenced therein, available to microservices subscribing to that particular CRD type. For example, the COE can inform that microservice (e.g., config-agent) that a CRO of that type has been created. The microservice can then read configuration information from the CRO via an API call, and then place another API call to the associated PaaS component instance to provide that information so that the PaaS component instance can properly supply services to an NF. Thus, microservice config-agents that subscribe to particular CRD types can listen for CROs to which it subscribes, which may be new or updated, such as those having dashboard files embedded therein, and then forward them to respective PaaS component instances. In this way, developers do not need to coordinate the process by which each respective PaaS component instance receives new or updated files. Config-agents and associated CNFs can operate using ConfigMaps, CROs, or both, in order to adequately provide configuration information to PaaS instances.

[0056] In some embodiments, both CROs and ConfigMaps can be used in the same system 200. For example, one Helm™ Chart can be deployed to instantiate CNFs that provide configuration information only in ConfigMaps, while a second Helm™ Chart can be deployed to instantiate CNFs that provide configuration information only via CROs. In another example, one Helm™ Chart can be deployed to instantiate a CNF that provides configuration information to one PaaS component instance using information read by config-agents read from ConfigMaps, while a second PaaS component instance can receive configuration information from a config-agent that receives such information from an API call to a CRO of the CRD type it operates. In any of these examples, a configuration manager, such as Ansible™, is not needed to ensure that configuration is provided to PaaS component instances. Instead, when an NF is instantiated or updated from a Helm™ Chart that may specify new, changed, or deleted configuration information for one or more PaaS components, config-agents associated with each PaaS component can pass updated configuration information to their respective PaaS components instances without requiring someone to manually code such an action in an Ansible™ playbook.

[0057] According to some embodiments, config-agents 232, 222, 242, 252, 262 can take the form of singleton microservices. Config-agents 232, 222, 242, 252, 262 can be declared in and installed via a separate Helm™ chart, which can provide the config-agents with information on how to discover updates to PaaS instance configurations (e.g., via monitoring for changes of a certain type to ConfigMaps or informing Kubernetes™ 170 that it operates CROs of particular CRD types. When a new piece of software containing one or more new types of microservices (e.g., via a Helm™ Chart) is instantiated or updated to implement new configuration on a PaaS component instance (e.g., dashboards in Grafana™ 130, endpoints to be accessed by Prometheus™ 120, alert-rules in an alert file to be implemented by ElastAlert 140, dashboards to be consumed by Kibana™, or schemas to be consumed by MongoDB™ 160A/B) for one of CNFs 204A, 204B, the respective Graphing-, PM-, Alert-, Dashboard-, and/or DB-config-agent 232, 222, 242, 252, and/or 262 can extract files from ConfigMaps 205 or 206, rename them to include the CNF ID, and install them in an instance of the PaaS component 130, 120, 140, 150, 160A/B associated with the CNF via API calls. Accordingly, one or more of config-agent 232, 222, 242, 252, 262 can also recognize new or dynamic changes to the ConfigMaps 205, 206 and synchronize with instances of PaaS components 120, 130, 140, 150, 160A/B via API calls without having to specify in a Playbook or the Helm™ Chart that changed the ConfigMap which particular PaaS component instances need

updated configuration information. For synchronization, in some examples, Microservices will utilize a Linux file system change notification mechanism (e.g., inotify) in order to get a callback from Linux when a change is detected. Kubernetes™ can be responsible for remounting the regenerated ConfigMap such that inotify can detect it and inform the microservice.

[0058] In some embodiments, files extracted from the ConfigMaps can include Grafana™-specific files that instruct it how to paint a web display in a browser and which KPIs to request from the data source (e.g., Prometheus™) in order to plot the lines on the web display

[0059] In some embodiments, files extracted from the ConfigMaps can include files that cause an instance of Prometheus™ 120 to perform remote service discovery (e.g., provide Prometheus™ 120 with endpoints (microservice instances) that are to data available to it. Since Prometheus™ 120 can be hierarchical – a parent Prometheus™ instance can need to know which child Prometheus™ instances make (aggregated) data available to it, PM-config-agent 222 can automate synchronization of a CNF microservice members list with Prometheus™ 120 to enable automated data collection by Prometheus™ 120 for hierarchical instances.

[0060] In some embodiments, files extracted from the ConfigMaps can include alert-rule files containing alert rules. An alert rule can be associated with a log message declaring that database connectivity has been lost – an operator may have created such a rule, declaring that this type of log message should result in a human detectable notification over email, SMS, slack, or many other communications schemes. Alert-rule files can be consumed by an instance of ElastAlert 140.

[0061] In some embodiments, files extracted from the ConfigMaps can include dashboard files to be consumed by an instance of Kibana™ 152. For example, an SMF-call-control microservice can have an associated dashboard constructed to view log activity emanating from the SMF-call-control, which can contain filters used to identify the specific log messages that are desired to be displayed in widgets on the dashboard (via a web browser).

[0062] In some embodiments, files extracted from the ConfigMaps can include DB-config-schema files to be consumed by an instance of MongoDB™ 160A/B to create or

modify a logical DB. In some examples, each associated set of schema files associated with a tenant can use the same file names to refer to different files for each tenant.

[0063] As described throughout the present disclosure, the information listed above can instead be retrieved by a config-agent 232, 222, 242, 252, or 262 via an API call to a CRO after receiving a notification that a CRO of the CRD type monitored by the config-agent has been instantiated.

[0064] As shown in FIG. 2, PaaS components can be multiple instantiated. For example, there are two instances of MongoDB™: 160A and 160B. In some embodiments, multiple network functions 204A, 204B can be served by the same PaaS component instance, such as Grafana™ 130. In some embodiments, each network function 204A, 204B can be serviced by a single PaaS component instance, such as MongoDB™ 160A and 160B, respectively. In some embodiments, PaaS components can be located in the cloud 209, while NFs 204A, 204B and configmaps 205, 206 are located in a data center 208. However, a person of skill in the art would understand that NFs 204A, 204B; ConfigMaps 205, 206; Config-Agents 232, 222, 242, 252, and 262; and/or PaaS components 130, 120, 140, 150, 160A, and 160B could be located in any configuration spread across multiple or a single data center 208 and/or with some or all elements residing in the cloud 209.

[0065] FIG. 3 is a call flow diagram showing an exemplary PaaS auto-configuration during CNF instantiation (e.g., installation), according to some embodiments. At step 302, a ‘helm install’ command can be called for a new CNF instance 204A, which is made up of at least one microservice, Microservice 290A, based on a Helm™ Chart. As discussed above, Microservice A 290A can assist with one portion of a NF, such as SMF call control, AMF call control, or AMF Tracking Area Management. Other microservices are contemplated. At step 304, as discussed above, a ConfigMap 205 can be created, which can include configuration information contained in the Helm™ Chart to be consumed by PaaS components (e.g., dashboards to be consumed by Grafana™ 130 and dashboards to be consumed by Kibana™ 150). The ConfigMap 205 allows for deployment of such information, which can take the form of files that are contained in the Helm™ Chart, to be passed through to PaaS component instances and to be mounted as a file system to the running microservice. At step 306 Helm™ 176 can provide, via the Helm™ Chart, instructions for Kubernetes™, which coordinates with the container runtime agent (e.g., Docker or a CRI-O) to create at least one instantiation of Microservice 290A based on the

docker images contained in the Helm™ Chart. The number of instantiations can also be declared in the Helm™ Chart.

[0066] Config-agents 232, 252 can continuously or periodically watch Linux for file changes at 314 and 354 (e.g., by issuing an inotify watchpoint to Linux for a file directory such that the software is notified with an event if there are any file deletes, adds, or modifies in the watched directory). Accordingly, when new CNFs are deployed or upgraded, for example via “helm install” or “helm upgrade” commands, config-agents simplify deployment automatically updating instances of PaaS components like Grafana™ 130 and Kibana™ 150. Thus, deploying new metadata to PaaS components instances does not require scripting commands that cause new PaaS component consumables to be installed to respective PaaS component instances, but rather config-agents continuously monitor for new files and pass them to respective PaaS components as needed. Once config-agents have been implemented, a developer need only declare such files in a Helm™ Chart using ConfigMaps, and the config-agents will take care of ensuring that they are properly installed on respective PaaS components instances.

[0067] In steps 330, 350, Graphing- and Dashboard-config-agents 232, 252 see changes to ConfigMap 205, including new dashboards. According to some embodiments, this involves gaining access to docker images that each respective config-agent runs and metadata, such as ConfigMaps, that should be mounted (such as that built by Helm™ and Kubernetes™ from data embedded in the Helm™ Chart). These steps can be accomplished in tandem or at different times. The config-agents 232, 252 can parse the changes new configurations that relate to the particular PaaS component which they serve. Upon identifying a new configuration, the config-agent 232, 252 can identify which PaaS component instance requires this new configuration (e.g., by using a namespace of the NF instance with which the configuration information is associated) and use POST API calls 332/352 to install the dashboards contained in the ConfigMap to Grafana™ 130 and Kibana™ 150, respectively. According to some embodiments where such an API call is not available, a separate config-agent can be deployed at the PaaS component in order to receive updated configurations (e.g., dashboards) and manipulate files local to the config-agent to accommodate the new configurations. According to some embodiments, each config-agent can be specifically designed to relay data to its respective PaaS component instance in the manner required by that PaaS component. For example, for ElastAlert™, a file that is locally

mounted to ElastAlert can be changed, where the file is not a ConfigMap mount. In such an example, the system can run a “sidecar” – a new type of container that runs alongside the ElastAlert container in a manner where their lifecycle is conjoined. The sidecar can have all the privileges of the ElastAlert container, including the ability to manipulate its files.

According to this example, the config-agent, as part of a CNF, can be the client side of the API relationship and the sidecar, as part of ElastAlert, can be considered to be on the server side. Accordingly, a remote API scheme can be implemented for this type of communication where none existed before. In an example, the sidecar and the ElastAlert Container can reside in the same Kubernetes™ pod.

[0068] FIG. 4 is a call flow diagram showing a PaaS auto-configuration during a microservice upgrade, according to some embodiments. At 402, a Helm™ Chart is updated, for example, by a developer delivering a new version thereof that indicates new or different docker image versions and/or sets of metadata. This can be, for example, a “child” Helm™ Chart associated with a particular CNF comprising microservices that is referenced in a top-level Helm™ Chart. Thus, according to some embodiments, the update may be, for example, to a particular microservice (Microservice A) or for a number of microservices in the CNF. Updates can also involve, for example, new, changed, or deleted configuration information for PaaS components, such as new, changed or deleted dashboards. When the top-level Helm™ Chart is updated to reference the new “child” Helm™ Chart, a “helm upgrade” command can be performed that associates an already-deployed CNF with the new top-level chart. Helm™ 173 can then decompose the Helm™ Chart and make calls to Kubernetes™ to effectuate the required changes to the already-deployed CNF. In such cases, Helm™ 176 can call Kubernetes™ to provide an update in step 404, which Kubernetes™ can use to replace files in ConfigMaps (or replace CRD objects for CRD implementations). Similarly, Helm™ 176 can update Microservice A 290A, for example on a rolling basis, by sending an update in step 406.

[0069] After ConfigMap 205 is updated, the Graphing- and Dashboard-config-agents 232A, 252A can identify changes in the ConfigMap, for examples based on file directory changes via receipt of notifications 432 and 452, respectively, as discussed above. In some embodiments, Linux may send a notification to config-agents 232, 252 of file directory changes as discussed above. For example, where files have been replaced in response to the updated Helm™ Chart, Linux can notify any microservice (e.g., config-agents 232 or 252) to

which the ConfigMap is mounted that a change in the files has occurred. After determining (e.g., being notified of) changes to the ConfigMap 205, the Graphing- and Dashboard-config-agents 232, 252 can send and/or install updates 434 and 454 including new dashboards to Grafana™ 130 and Kibana™ 150, respectively. Accordingly, higher level provisioning software like Ansible™ is not required in order to ensure that instances of PaaS components like Grafana™ and Kibana™ are kept up to date with new configurations across each CNF. Instead, config-agents become aware of updates to the ConfigMap and manage updates to individual components automatically. Such a system is scalable for larger deployments, and can significantly reduce costs for upgrading where a large number of open source components are used. Furthermore, the upgrade process does not require a single part of the system to both keep track of which CNFs use particular PaaS components that will need to be updated and coordinate the updating process. Thus, as a CNF is instantiated or updated to require different configurations of PaaS components, the respective PaaS component instances associated with that CNF are automatically updated to meet those requirements.

[0070] FIG. 5 is a call flow diagram showing an exemplary PaaS auto-configuration during CNF instantiation using CRDs, according to some embodiments. As discussed above, CRDs can be used by config-agents instead of, alongside, or in place of ConfigMaps. At step 502, a 'helm install' command can be called for a new CNF instance 204A, which is made up of at least one new microservice, Microservice 290A, based on a Helm™ Chart. As discussed above, Microservice A 290A can assist with one portion of a NF, such as SMF call control, AMF call control, or AMF Tracking Area Management. Other microservices are contemplated. At step 504, Helm™ 176 instructs Kubernetes™ 170 to instantiate a CROs of certain CRD types (which can be declared and deployed in a separate Helm™ Chart prior to step 504 such that Kubernetes™ understand such CRD types), such as CRO 505 having, for example, new graphs or dashboards contained therein as declared in the Helm™ Chart. As discussed above, CRO 505 can be constructed such that subscribing microservices can see their contents, for example by completing API calls thereto. In step 506, Microservices are created and replicated based on docker images and metadata contained in the Helm™ Chart. For example, Microservice A (290A) is created.

[0071] At step 530, the Graphing-config-agent 232 can subscribe to (e.g., informed Kubernetes™ 170 that it operates) CROs of a particular CRD type (e.g., for Grafana™ 130) that publish specified dashboard files for which the Graphing-config-agent 232 is configured

to watch. This can occur before the Helm™ Chart is installed in 502. In step 531, Kubernetes™ 170 notifies the Graphing-config-agent 232, which is subscribed to the CRO 505, that dashboard files are available and passes the files contained in the CRO 505 to the Graphing-config-agent 232. Graphing-config-agent 232 can then identify the correct PaaS component instance (e.g., via a namespace of the CNF with which the CRO 505 is associated) and make a POST API call to the particular instance of Grafana™ 130 to install these dashboards contained in the CROs by indicating that they should add a fully specified object, remove an object, and/or over write an object with a fully specified replacement value in step 532. The API call can also indicate the new CNF instance 204A which is to be associated with a new instance of the PaaS component Grafana™ 130. Accordingly, an instance of Grafana™ 130 is provided with new dashboards upon deployment of the CNF without requiring special scripting in the Helm™ Chart or an Ansible™ Playbook specifying that they be installed. Similarly, at 550, the Dashboard-config-agent 252 can subscribe to CROs of a particular CRD type (e.g., for Kibana™ 150) that publish specified dashboard files for which the Dashboard-config-agent 252 is configured to watch. This can occur before the Helm™ Chart is installed in 502. In step 551, Kubernetes™ 170 notifies the Dashboard-config-agent 252, which is subscribed to the CRO 505 of a particular CRD type that dashboard files are available along with the files contained in the CRO 505. Dashboard-config-agent 252 can then identify the correct PaaS component instance (e.g., via a namespace of the CNF with which the CRO 505 is associated) and make a POST API call to particular instance of Kibana™ 150 to install these dashboards contained in the CROs on a particular Kibana™ instance as described above. Accordingly, Kibana™ 150 is provided with new dashboards upon deployment of the CNF without requiring special scripting in the Helm™ Chart specifying that they be installed.

[0072] FIG. 6 is a call flow diagram showing a PaaS auto-configuration during a CNF upgrade using CRDs, according to some embodiments. At 602, a Helm™ Chart is updated, for example, by a developer delivering a new version thereof that indicates new docker image versions and/or metadata. This update can be similar to the one discussed with reference to FIG. 4, but can use CRDs in combination with config-agents. At 604, Helm™ 176 can call Kubernetes™ 170 to provide an update, which Kubernetes™ 170 can use to update object(s) in CRO 505. In response, at 605 and 607, Kubernetes™ 170 notifies all subscribing config-agents (e.g., those that operate CROs of the particular CRD type of CROs 505), such as Graphing-config-agent 232 and Dashboard-config-agent 252 that there are new or updated

dashboard files and passes the files in the CROs 505 to the config-agents 232, 252. Graphing-config-agent 232 and Dashboard-config-agent 252 then execute POST/PUT/DELETE API calls to install/update/delete any new or updated dashboards to instances Grafana™ 130 and Kibana™ 150, respectively, associated with the updated CNF in steps 634/654. In 606, Helm™ 176 may instruct Kubernetes™ 170 to upgrade microservices as specified in the Helm™ Chart (e.g., loading new software from docker images) in step 606.

[0073] In some embodiments, a Config-Agent can further monitor PaaS component instances to ensure that configurations thereof are consistent with declarations contained in the CROs and/or ConfigMaps monitored by the Config-Agent. For example, if a PaaS component instance fails and is re-instantiated without the necessary configuration information, the Config-Agent can identify that the configuration of the PaaS component instantiation is either incomplete or missing (e.g., via an API call) and push the necessary configuration information thereto (e.g., via an API call). Accordingly, a Config-Agent can ensure that adequate configuration is installed on PaaS component instances at all times whenever a difference occurs between CROs and/or ConfigMaps monitored by the Config-Agent and the associated PaaS component instance, regardless of where the difference occurred from.

[0074] The subject matter described herein can be implemented in digital electronic circuitry, or in computer software, firmware, or hardware, including the structural means disclosed in this specification and structural equivalents thereof, or in combinations of them. The subject matter described herein can be implemented as one or more computer program products, such as one or more computer programs tangibly embodied in an information carrier (e.g., in a machine readable storage device), or embodied in a propagated signal, for execution by, or to control the operation of, data processing apparatus (e.g., a programmable processor, a computer, or multiple computers). A computer program (also known as a program, software, software application, or code) can be written in any form of programming language, including compiled or interpreted languages, and it can be deployed in any form, including as a stand-alone program or as a module, component, subroutine, or other unit suitable for use in a computing environment. A computer program does not necessarily correspond to a file. A program can be stored in a portion of a file that holds other programs or data, in a single file dedicated to the program in question, or in multiple coordinated files

(e.g., files that store one or more modules, sub programs, or portions of code). A computer program can be deployed to be executed on one computer or on multiple computers at one site or distributed across multiple sites and interconnected by a communication network.

[0075] The processes and logic flows described in this specification, including the method steps of the subject matter described herein, can be performed by one or more programmable processors executing one or more computer programs to perform functions of the subject matter described herein by operating on input data and generating output. The processes and logic flows can also be performed by, and apparatus of the subject matter described herein can be implemented as, special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application specific integrated circuit).

[0076] Processors suitable for the execution of a computer program include, by way of example, both general and special purpose microprocessors, and any one or more processor of any kind of digital computer. Generally, a processor will receive instructions and data from a read only memory or a random access memory or both. The essential elements of a computer are a processor for executing instructions and one or more memory devices for storing instructions and data. Generally, a computer will also include, or be operatively coupled to receive data from or transfer data to, or both, one or more mass storage devices for storing data, e.g., magnetic, magneto optical disks, or optical disks. Information carriers suitable for embodying computer program instructions and data include all forms of nonvolatile memory, including by way of example semiconductor memory devices, (e.g., EPROM, EEPROM, and flash memory devices); magnetic disks, (e.g., internal hard disks or removable disks); magneto optical disks; and optical disks (e.g., CD and DVD disks). The processor and the memory can be supplemented by, or incorporated in, special purpose logic circuitry.

[0077] To provide for interaction with a user, the subject matter described herein can be implemented on a computer having a display device, e.g., a CRT (cathode ray tube) or LCD (liquid crystal display) monitor, for displaying information to the user and a keyboard and a pointing device, (e.g., a mouse or a trackball), by which the user can provide input to the computer. Other kinds of devices can be used to provide for interaction with a user as well. For example, feedback provided to the user can be any form of sensory feedback, (e.g., visual feedback, auditory feedback, or tactile feedback), and input from the user can be received in any form, including acoustic, speech, or tactile input. The subject matter described herein can

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

[0078] It is to be understood that the disclosed subject matter is not limited in its application to the details of construction and to the arrangements of the components set forth in the following description or illustrated in the drawings. The disclosed subject matter is capable of other embodiments and of being practiced and carried out in various ways. Also, it is to be understood that the phraseology and terminology employed herein are for the purpose of description and should not be regarded as limiting.

[0079] As such, those skilled in the art will appreciate that the conception, upon which this disclosure is based, may readily be utilized as a basis for the designing of other structures, methods, and systems for carrying out the several purposes of the disclosed subject matter. It is important, therefore, that the claims be regarded as including such equivalent constructions insofar as they do not depart from the spirit and scope of the disclosed subject matter.

[0080] Although the disclosed subject matter has been described and illustrated in the foregoing exemplary embodiments, it is understood that the present disclosure has been made only by way of example, and that numerous changes in the details of implementation of the disclosed subject matter may be made without departing from the spirit and scope of the disclosed subject matter, which is limited only by the claims which follow.

CLAIMS

1. A method comprising:

installing a service instance using a first descriptor chart, the installing comprising:

installing a plurality of microservices declared in the first descriptor chart that comprise the service instance, at least one of the microservices configured to request services from an instance of a Platform-as-a-Service (“PaaS”) component of a PaaS component type, and

creating at least one configuration file based on the first descriptor chart;

parsing, by a configuration agent associated with the PaaS component type, the at least one configuration file for configuration data associated with the PaaS component; and

instructing, by the configuration agent, the instance of the PaaS component to install the configuration data.

2. The method of claim 1, wherein the configuration file comprises descriptors of resources, and wherein the method further comprises mounting the configuration file to the configuration agent.

3. The method of claim 1, wherein the service instance comprises a network function (“NF”).

4. The method of claim 3, wherein the service instance is installed by a network operator, and the PaaS component is developed and managed by a third-party.

5. The method of claim 1, further comprising:

updating the service instance using a second descriptor chart, the updating comprising:

updating at least one of the plurality of microservices declared in the second descriptor chart that comprise the service instance, the updated at least one of the plurality of microservices configured to request services from the instance of the PaaS component, and

updating the at least one configuration file based on the second descriptor chart;

parsing, by the configuration agent, the at least one updated configuration file to identify the update to the configuration data associated with the PaaS component; and

instructing, by the configuration agent, the instance of the PaaS component to install the updated configuration data.

6. The method of claim 5, wherein the updating the software instance is performed by a container orchestration engine (“COE”).

7. A method comprising:

monitoring, by a configuration agent associated with a Platform-as-a-Service (“PaaS”) component type, for a notification identifying a configuration file change in at least one configuration file directory, the at least one configuration file in the configuration file directory containing configuration data for at least one instance of the PaaS component of the PaaS component type;

monitoring, by the configuration agent, a configuration of the at least one instance of the PaaS component type to identify a lack of conformity with configuration data in the configuration file directory;

automatically transmitting, by the configuration agent, an instruction to the at least one instance of the PaaS component when one or more of a configuration file change or a lack of conformity is identified, the instruction comprising an instruction to install a configuration on the at least one instance of the PaaS component based on the configuration data in the configuration file directory.

8. The method of claim 7, wherein the at least one instance of the PaaS component is configured to service at least one microservice of a plurality of microservices comprising a network function (“NF”).

9. The method of claim 7, wherein the change in configuration data comprises a change in data of a ConfigMap in the configuration file directory.

10. The method of claim 7, wherein the configuration data is associated with at least one service instance, and wherein the method further comprises updating the at least one service instance using a descriptor chart, the updating the at least one service instance causing the change in the configuration data.

11. The method of claim 10, wherein the at least one service instance is updated by a network operator, and the PaaS component is developed and managed by a third-party.

12. A method comprising:

installing a service instance using a first descriptor chart, the installing comprising:

installing a plurality of microservices declared in the first descriptor chart that comprise the service instance, at least one of the microservices configured to request services from an instance of a Platform-as-a-Service (“PaaS”) component of a PaaS component type, and

creating at least one resource instance based on the first descriptor chart, the resource instance having a type associated with a resource definition, the resource instance containing configuration data associated with the instance of the PaaS component;

receiving, by a configuration agent that operates resource instances of the type associated with the resource definition, a notification of the creation of the at least one resource instance, the notification further comprising the configuration data associated with the instance of the PaaS component; and

instructing, by the configuration agent, the instance of the PaaS component to install the configuration data.

13. The method of claim 12, wherein the instructing the instance of the PaaS component to install the configuration data comprises performing, by the configuration agent, an API call to the instance of the PaaS component.

14. The method of claim 12, wherein the service instance comprises a network function (“NF”).

15. The method of claim 14, wherein the service instance is installed by a network operator, and the PaaS component is developed and managed by a third-party.

16. The method of claim 12, further comprising:

updating the service instance using a second descriptor chart, the updating comprising:

updating at least one of the plurality of microservices declared in the second descriptor chart that comprise the service instance, the updated at least one of the plurality of microservices configured to request services from the instance of the PaaS component, and

updating the at least one resource instance based on the second descriptor chart by updating configuration data in the resource instance associated with the instance of the PaaS component;

receiving, by the configuration agent that operates resource instances of the type associated with the resource definition, a notification of the update to the at least one resource instance, the notification further comprising the updated configuration data associated with the instance of the PaaS component; and

instructing, by the configuration agent, the instance of the PaaS component to install the updated configuration data.

17. The method of claim 5, wherein the updating the software instance is performed by a container orchestration engine (“COE”).

18. A method comprising:

monitoring, by a configuration agent associated with a resource definition type, for a notification identifying that a declaration in a resource instance of the resource definition type has changed, the changed declaration being associated with at least one instance of a Platform-as-a-Service (“PaaS”) component;

monitoring, by the configuration agent, a configuration of the at least one instance of the PaaS component to identify a lack of conformity with the declaration in the resource instance of the resource definition type;

automatically transmitting, by the configuration agent, an instruction to the at least one instance of the PaaS component when one or more of a declaration change or a lack of conformity is identified, the instruction comprising an instruction to install a configuration on the at least one instance of the PaaS component based on the declaration in the resource instance.

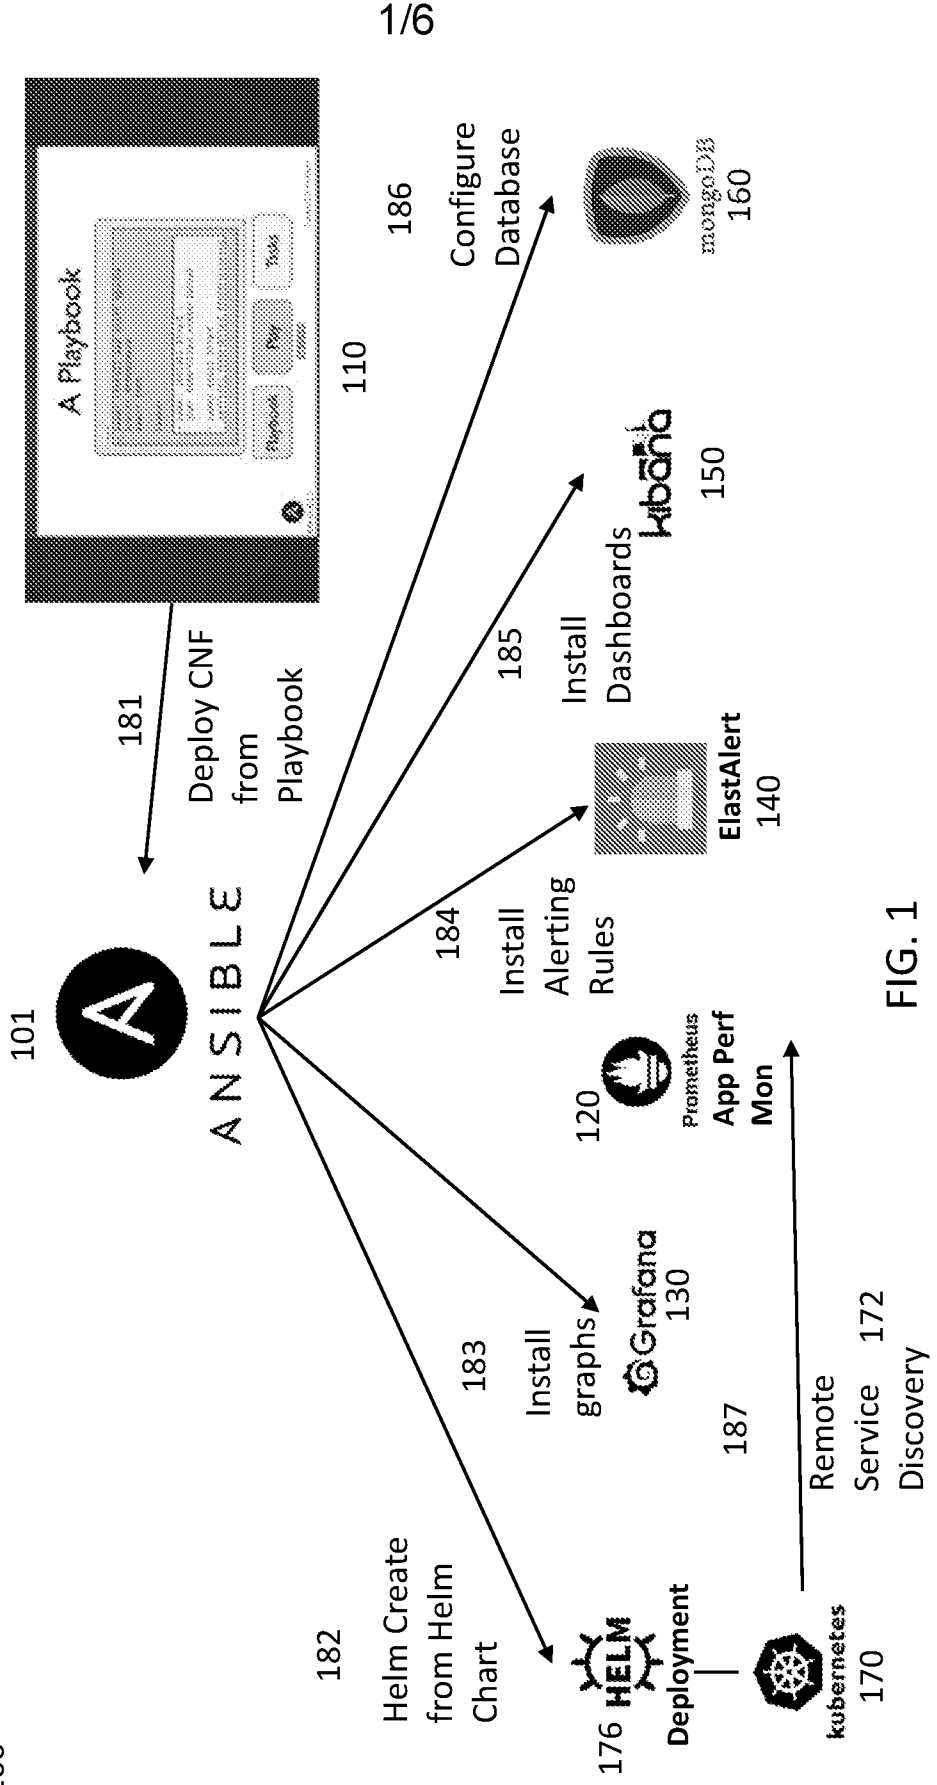
19. The method of claim 18, wherein the at least one instance of the PaaS component is configured to service at least one microservice of a plurality of microservices comprising an instance of a network function (“NF”).

20. The method of claim 19, further comprising updating the instance of the NF using a descriptor chart, the updating the instance of the NF causing the change in the declaration.

21. The method of claim 20, wherein the instance of the NF is updated by a network operator, and the PaaS component is developed and managed by a third-party

22. The method of claim 18, wherein the changed declaration comprises a change to a consumable of the at least one instance of the PaaS component.

100



1/6

FIG. 1
(PRIOR ART)

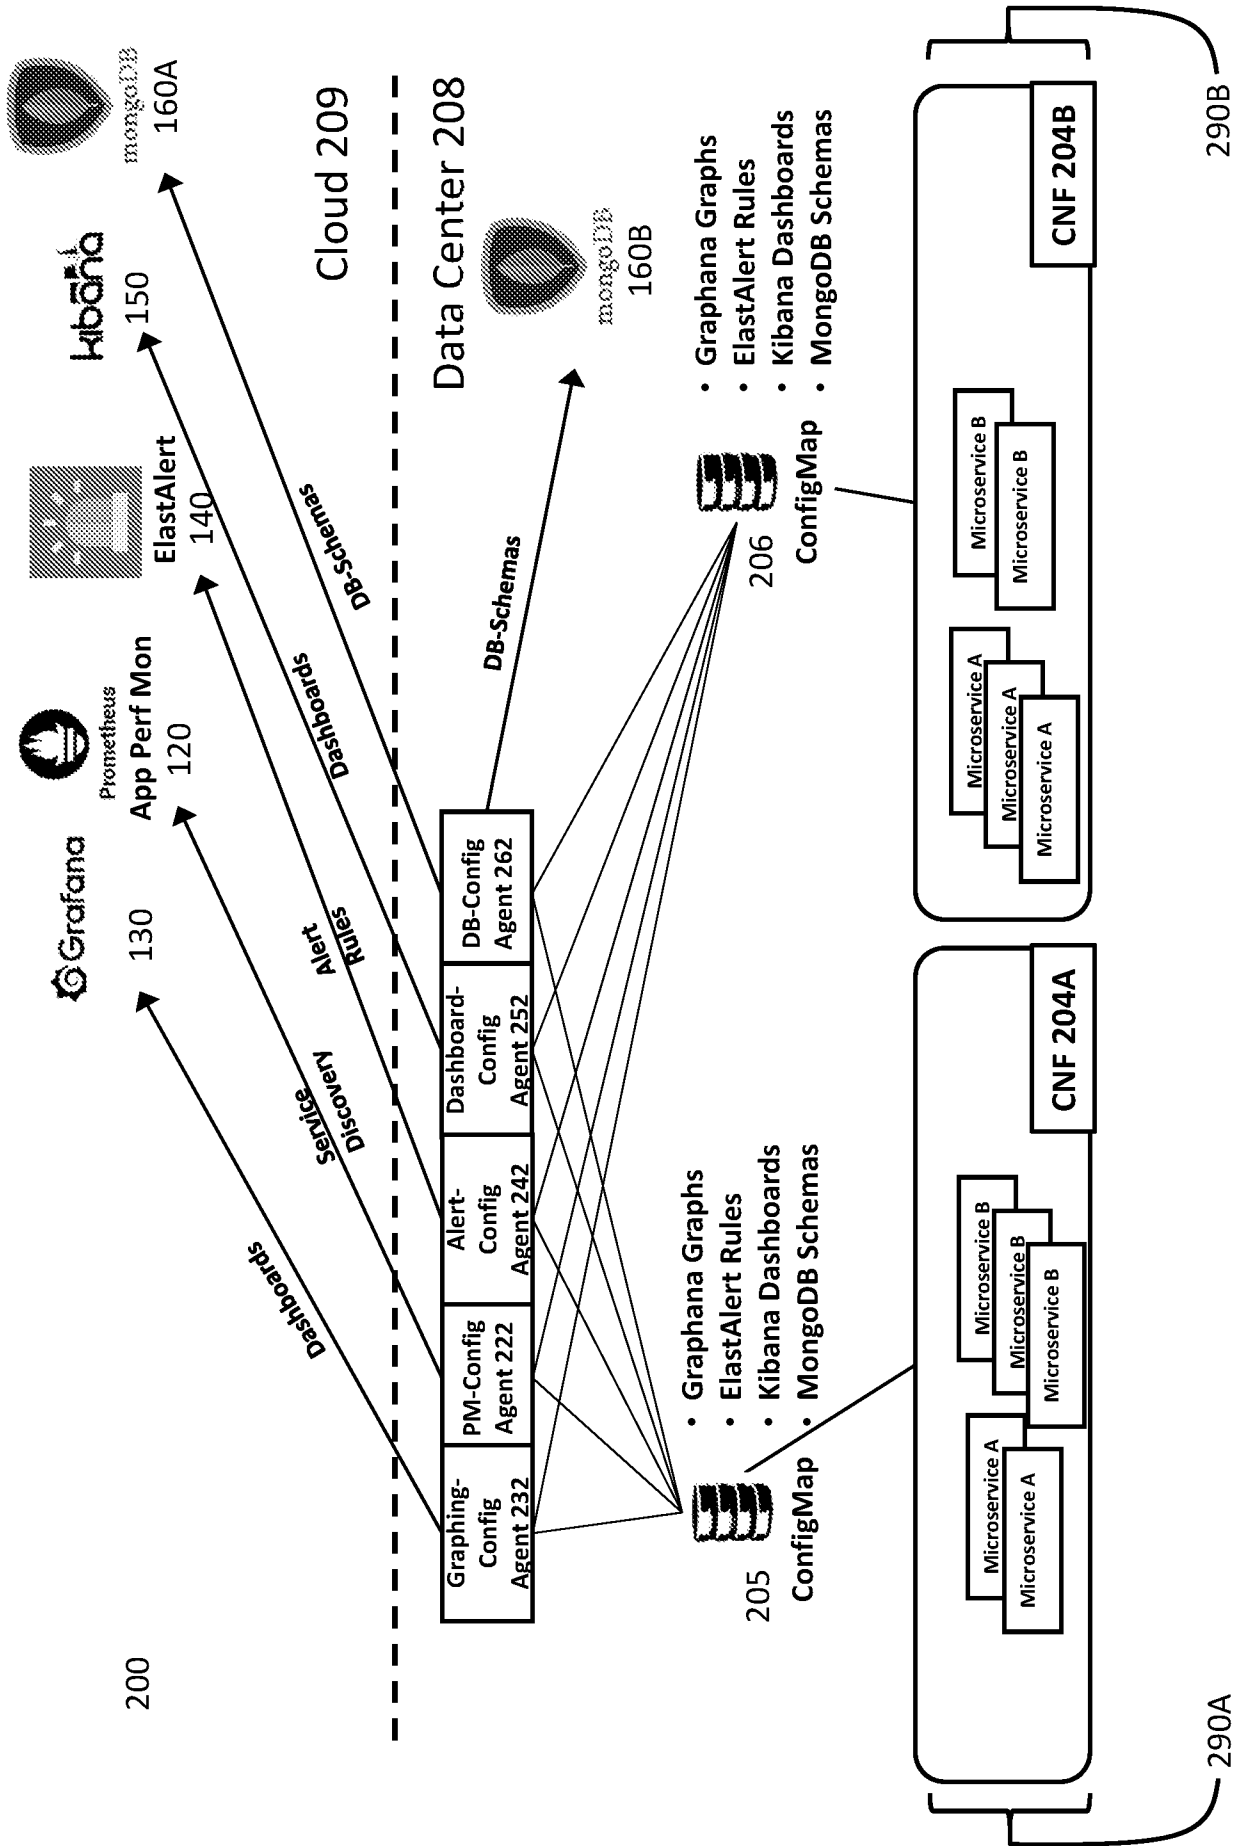


FIG. 2

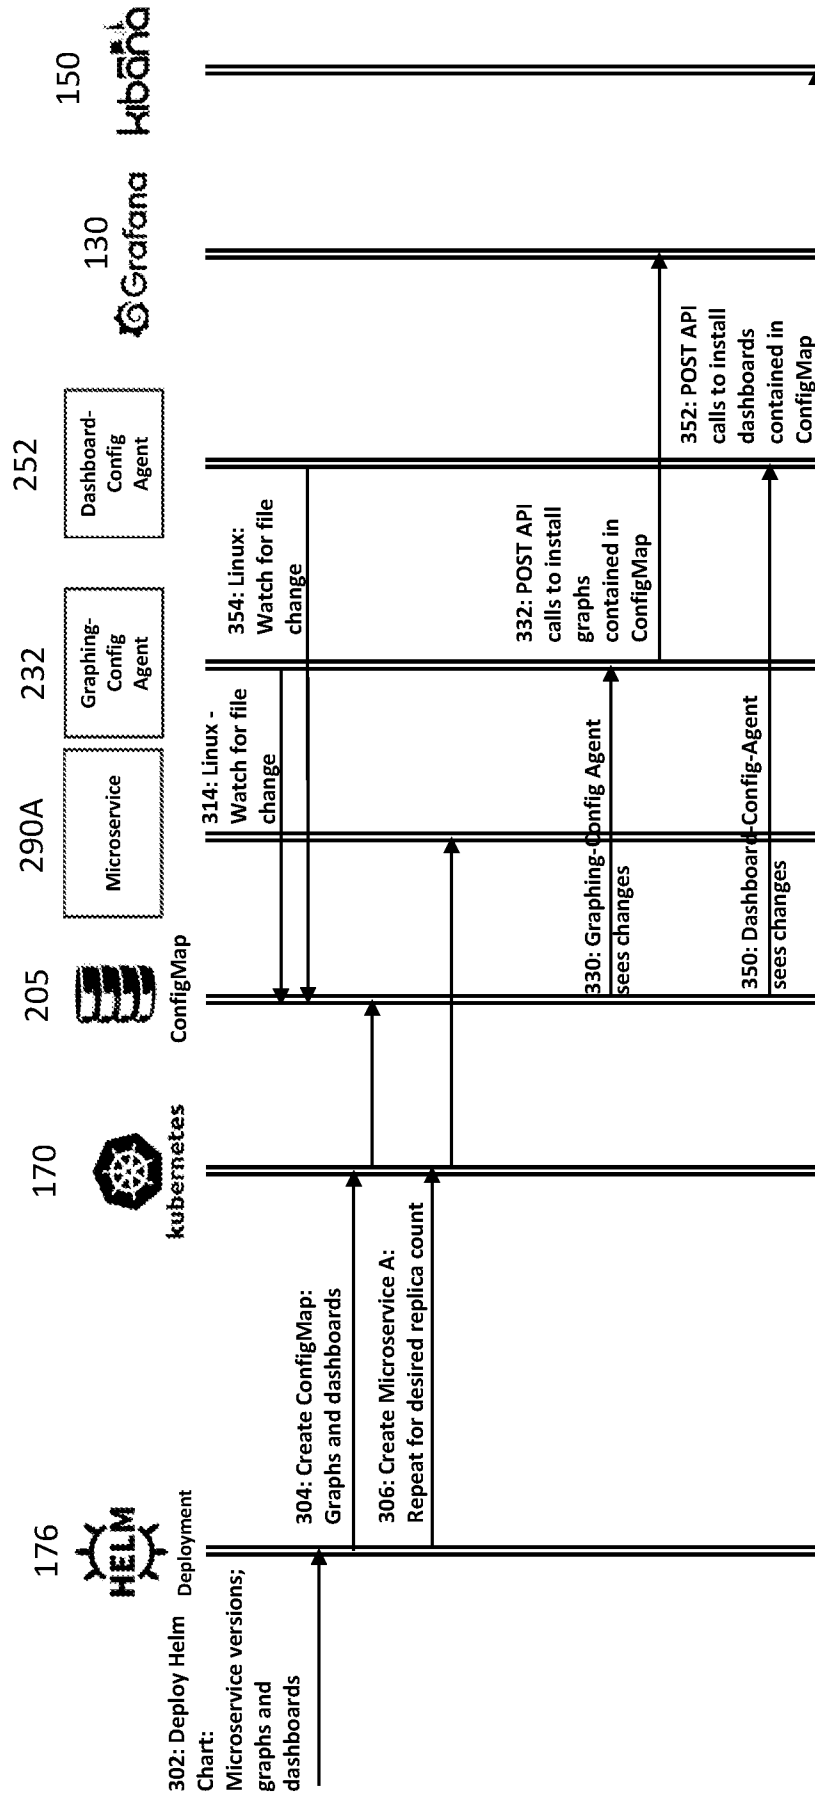


FIG. 3

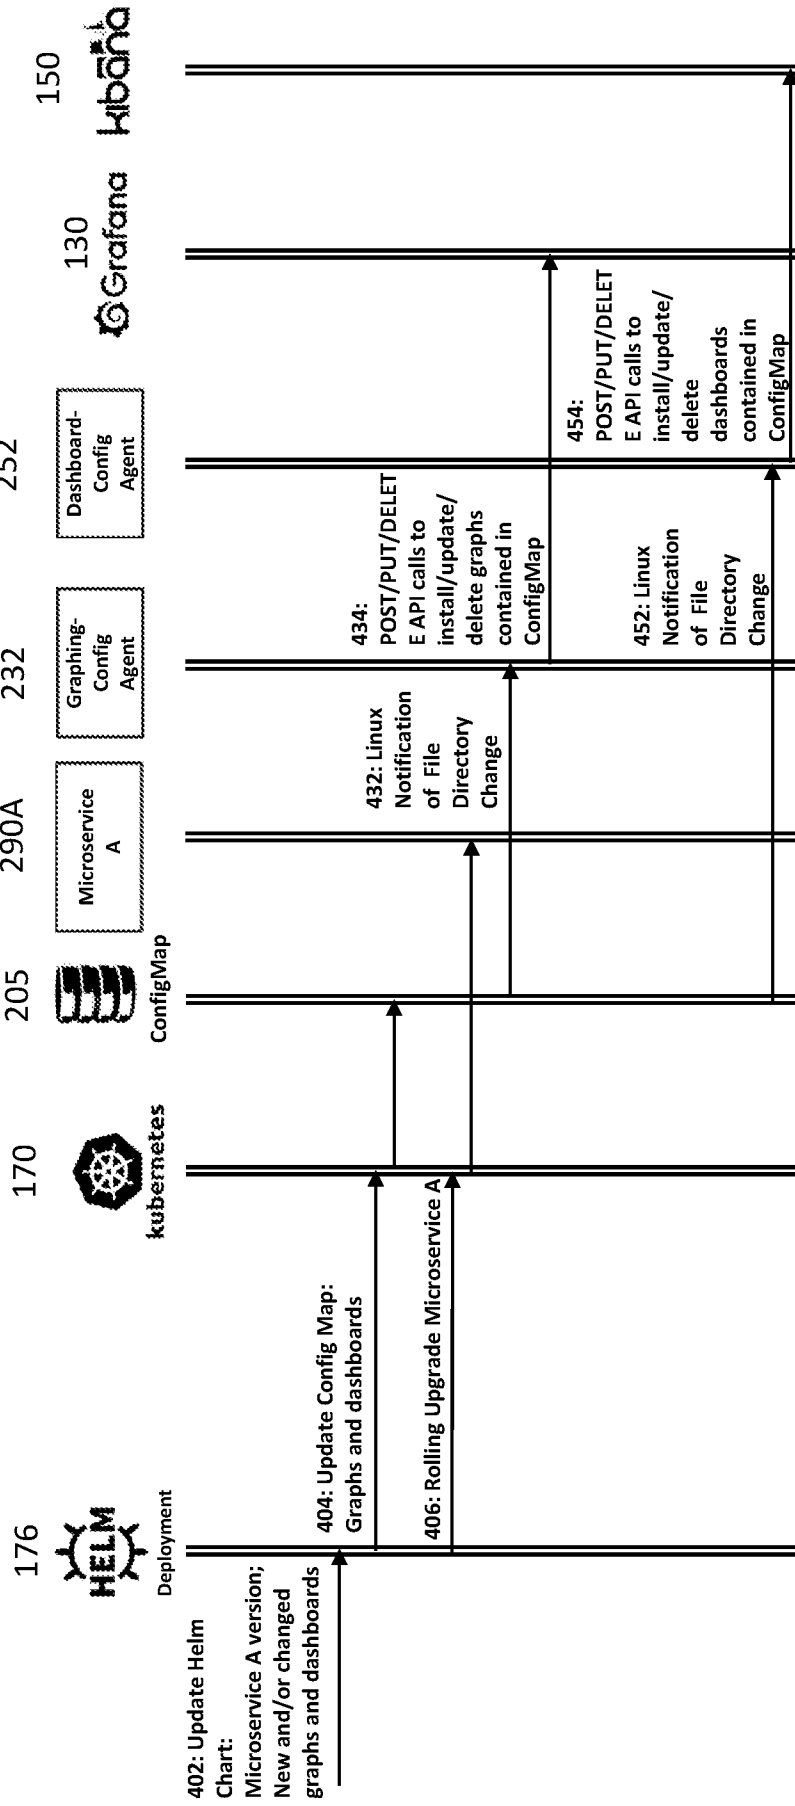


FIG. 4

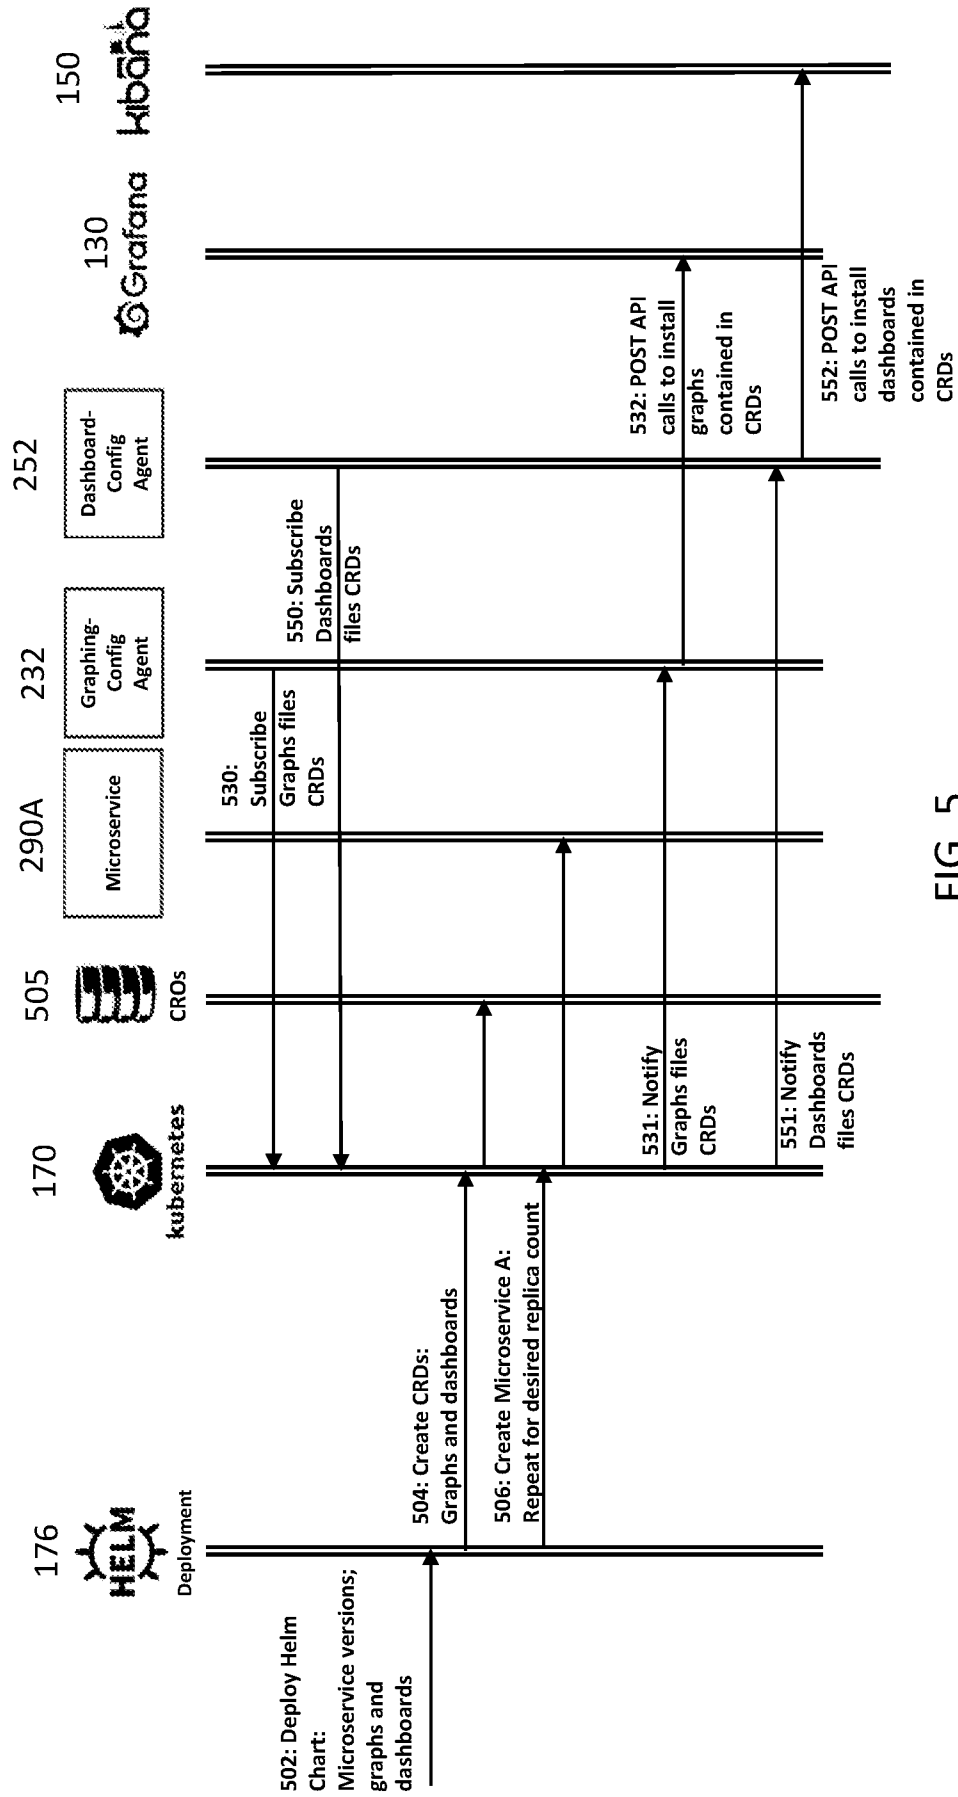


FIG. 5

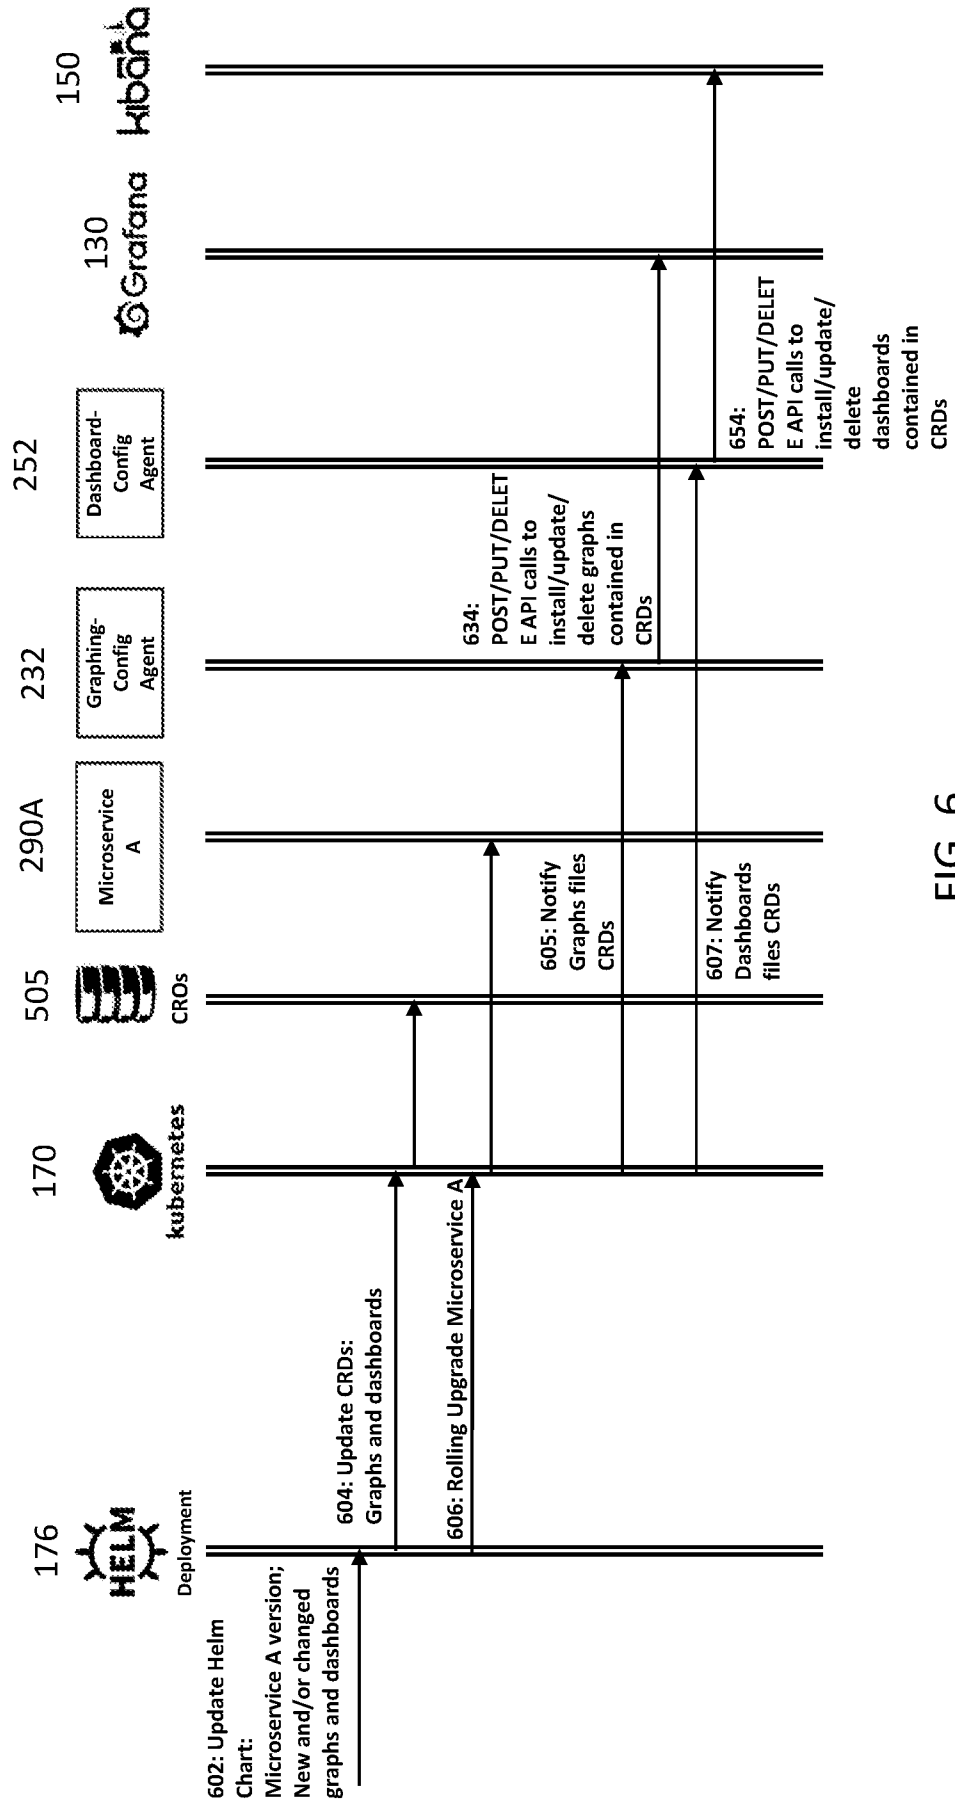


FIG. 6

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2020/013631

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F9/54 G06F9/445 G06F9/455
ADD.
According to International Patent Classification (IPC) or to both national classification and IPC

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

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

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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	Matt Butcher: "TechnoSophos: How Helm Uses ConfigMaps to Store Data", 1 November 2018 (2018-11-01), XP055688779, Retrieved from the Internet: URL:http://web.archive.org/web/20181101124257/http://technosophos.com/2017/03/23/how-helm-uses-configmaps-to-store-data.html [retrieved on 2020-04-23] page 1 - page 4	1-6, 12-17
A	US 2014/075412 A1 (KANNAN MAHESH [US] ET AL) 13 March 2014 (2014-03-13) paragraph [0045] - paragraph [0155]; figure 10	1-6, 12-17

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

24 April 2020

Date of mailing of the international search report

26/06/2020

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Kalejs, Eriks

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2020/013631

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:

3. Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

see additional sheet

1. As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.

2. As all searchable claims could be searched without effort justifying an additional fees, this Authority did not invite payment of additional fees.

3. As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:

4. No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

1-6, 12-17

Remark on Protest

- The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- No protest accompanied the payment of additional search fees.

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

This International Searching Authority found multiple (groups of) inventions in this international application, as follows:

1. claims: 1-6, 12-17

Installing microservices using descriptor charts

2. claims: 7-11, 18-22

Monitoring by an agent configuration files of a microservice for lack of conformity and updating microservices if nonconformity detected

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2020/013631

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2014075412 A1	13-03-2014	US 2014075412 A1	13-03-2014
		WO 2014039892 A1	13-03-2014
