



- (51) International Patent Classification:
G06F 17/30 (2006.01) *G06F 15/16* (2006.01)
- (21) International Application Number:
PCT/US2013/067513
- (22) International Filing Date:
30 October 2013 (30.10.2013)
- (25) Filing Language: English
- (26) Publication Language: English
- (71) Applicant: **HEWLETT-PACKARD DEVELOPMENT COMPANY, L.P.** [US/US]; 11445 Compaq Center Drive W., Houston, TX 77070 (US).
- (72) Inventor: **MAES, Stephane, Herman**; 1140 Enterprise Way, Sunnyvale, CA 94089 (US).
- (74) Agents: **MCKINNEY, Jack, H.** et al.; Hewlett-Packard Company, Intellectual Property Administration, 3404 E. Harmony Road, Mail Stop 35, Fort Collins, CO 80528 (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, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, 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).

[Continued on next page]

(54) Title: MONITORING A CLOUD SERVICE MODELED AS A TOPOLOGY

(57) Abstract: A method of monitoring a cloud service modeled as a topology with a processor includes establishing a number of monitoring policies to monitor an instantiated topology, detecting an event at a node based on the number of monitoring policies, and processing the detected event based on the number of monitoring policies.

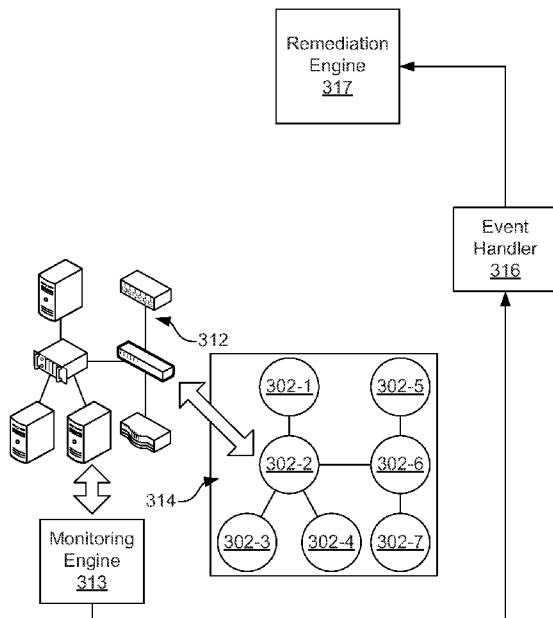


Fig. 10

WO 2015/065370 A1

Declarations under Rule 4.17:

- *as to the identity of the inventor (Rule 4.17(i))*
- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

Published:

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

MONITORING A CLOUD SERVICE MODELED AS A TOPOLOGY

BACKGROUND

[0001] An increasingly larger number of business entities and individuals are turning to cloud computing and the services provided through a cloud computing system in order to, for example, sell goods or services, maintain business records, and provide individuals with access to computing resources, among other cloud-related objectives. Cloud computing provides consumers of the cloud with scalable and pooled computing, storage, and networking capacity as a service or combinations of such services built on the above. A cloud may be designed, provisioned, deployed, and maintained by or for the entity for which the cloud computing system is created. Designing, provisioning, deploying, and maintaining a cloud computing system may be a difficult task.

BRIEF DESCRIPTION OF THE DRAWINGS

[0002] The accompanying drawings illustrate various examples of the principles described herein and are a part of the specification. The illustrated examples are given merely for illustration, and do not limit the scope of the claims.

[0003] Fig. 1 is a block diagram of a blueprint, according to one example of the principles described herein.

[0004] Fig. 2 is a block diagram showing an architecture derived topology, according to one example of the principles described herein.

[0005] Fig. 3 depicts a block diagram showing a functional overview of a topology-based management broker for designing, provisioning, deploying, monitoring, and managing a cloud service, according to one example of the principles described herein.

[0006] Fig. 4 depicts a block diagram showing a functional overview of a topology-based management broker for designing, provisioning, deploying, monitoring, and managing a cloud service, according to one example of the principles described herein.

[0007] Fig. 5 is a block diagram of an execution flow of the execution of a topology using provisioning policies, according to one example of the principles described herein.

[0008] Fig. 6 is a flowchart showing a method for brokering a cloud service, according to one example of the principles described herein.

[0009] Fig. 7 is a flowchart showing a method for brokering a cloud service, according to another example of the principles described herein.

[0010] Fig. 8 is a flowchart showing a method for remediating a number of incidents within a cloud service, according to one example of the principles described herein.

[0011] Fig. 9 is a flowchart showing a method of designing a topology, according to one example of the principles described herein.

[0012] Fig. 10 depicts a block diagram showing a functional overview of a monitoring engine for monitoring a cloud service modeled as a topology, according to one example of the principles described herein.

[0013] Fig. 11 is a flowchart showing a method for monitoring a cloud service modeled as a topology, according to one example of the principles described herein.

[0014] Fig. 12 is a flowchart showing another method for monitoring a cloud service modeled as a topology, according to another example of the principles described herein.

[0015] Fig. 13 depicts a block diagram of an instantiated service that includes monitoring policies, according to one example of the principles described herein.

[0016] Fig. 14 depicts a block diagram of an instantiated service and a monitoring system topology, according to one example of the principles described herein.

[0017] Throughout the drawings, identical reference numbers designate similar, but not necessarily identical, elements.

DETAILED DESCRIPTION

[0018] Cloud computing provides services for a user's data, software, and computation. Applications deployed on resources within the cloud service may be manually deployed. This manual deployment consumes considerable administrative time. Manual deployment of an application may include the provisioning and instantiation of the infrastructure. This may include linking the installation of an application or a platform such as middleware and DB+ applications or deployment of an image with or without the full knowledge of the deployed infrastructure. Manual deployment may further include numerous operations launched by a user who attempts to deploy the application. Thus, the manual linking of an application to a deployed infrastructure consumes large amounts of computing and personnel resources, and may lead to mistakes and irreconcilable issues between the application and the underlying infrastructure. Linking of an application to a deployed infrastructure may be automated with a number of tools, scripts, and executables, with orchestrators automating the sequence of execution of these processes. A number of devices used in the designing, provisioning, deploying, and maintaining of applications deployed on resources within the cloud service may include data centers, private clouds, public clouds, managed clouds, hybrid clouds, and combinations thereof.

[0019] More specifically, cloud services provided to users over a network may be designed, provisioned, deployed, and managed using a cloud service manager. The cloud service provider or other entity or individual designs, provisions, deploys, and manages such a cloud service that appropriately comprises of a number of services, applications, platforms or infrastructure capabilities deployed, executed, and managed in a cloud

environment. These designs may then be offered to user who may order, request, and subscribe to them from a catalog via a market place or via an API call, and then manage the lifecycles of a cloud service deployed based on the designs through the same mechanism. The service designs in a cloud service manager such as CLOUD SERVICE AUTOMATION (CSA 3.2) designed and distributed by Hewlett Packard Corporation, described in more detail below, are expressed with “blueprints.”

[0020] Blueprints describe services in terms of the collections of workflows that are to be executed to provision or manage all the components that make up the service in order to perform a particular lifecycle management action. Some of the functions of the workflows defined by blueprints are actual life cycle management actions that are then performed as calls to a resource provider. The resource provider converts the calls into well-formed and exchanged instructions specific to the particular resource or instance offered by a resource provider.

[0021] Fig. 1 is a block diagram of a blueprint (100), according to one example of the principles described herein. Each object (102-1, 102-2, 102-3, 102-4, 102-5, 102-6, 102-7, 102-8, 102-9, 102-10, 102-11, 102-12) in the blueprint may be associated with action workflows that call resource providers. A number of challenges exist with a blueprint (100) approach to designing, provisioning, deploying, and managing cloud services. The structure of a blueprint, while consisting of objects comprising properties and actions linked by relationships, do not identify relationships to physical topologies such as, for example, the actual physical architecture of the system that supports the cloud service. This renders it difficult to associate additional metadata with the blueprints (100) to describe, for example, policies associated with the system. Further, this association of policies with nodes in a blueprint (100) is not intuitive for a designer or administrator of the to-be-deployed cloud service.

[0022] Further, the structures of blueprints (100), for the same reason, are difficult to use as models of applications or templates of infrastructures as CONTINUOUS DELIVERY AUTOMATION (CDA) does. CDA is system tool utilized within a topology designer that independently models

infrastructure and application requirements while managing versions, configurations, and other application components. CDA 1.2 is also developed and distributed by Hewlett Packard Corporation. The structures of blueprints (100), for the same reason given above, are difficult to use as models of applications because blueprints do not describe the architecture of the application. Further, blueprints are difficult to use as templates of an infrastructure because they also do not describe the architecture of the infrastructure. As a result, systems aiming at modeling application models and infrastructure or platform templates, and mapping the application models and infrastructure or platform templates to each other are not easily reconciled with the blueprints because they are based on different methods of modeling these services.

[0023] The present systems and methods describe architecture-descriptive topologies that define the physical architecture of a system that constitutes a cloud service. Fig. 2 is a block diagram showing an architecture derived topology (200), according to one example of the principles described herein. As depicted in Fig. 2, the architecture derived topology (200) may comprise a number of nodes (201, 202, 203, 204, 205, 206, 207, 208, 209, 210, 211, 212, 213, 214, 215) associated with one another. Associations between nodes within the topology (200) are indicated by the open arrows. A number of nodes (201, 202, 203, 204, 205, 206, 207, 208, 209, 210, 211, 212, 213, 214, 215) within the topology (200) may also be aggregated with one another as designated by the filled arrows. Aggregation is a computing term used to describe combining (aggregating) multiple network connections in parallel to increase throughput beyond what a single connection could sustain, and to provide redundancy in case one of the links fails.

[0024] For example, the load balancer (201), web server service (202), application enterprise archive (203), and the database (204) are associated with one another. The web server service (202) is aggregated with a web virtual machine (205) and its security group (213) as well as a web virtual local area network (209). Similarly, the application enterprise archive (203) is aggregated with an application server service such as the JavaBeans Open

Source Software Application Server (JBoss) service (206), a JBoss virtual machine (208) and its associated security group (214), and a secure application virtual local area network (210). Again, similarly, the database (204) is associated with a database virtual machine (207) and its security group (215), and a secure database virtual local area network (211). The web virtual local area network (209), secure application virtual local area network (210), and secure database virtual local area network (211) are then associated with a router (212).

[0025] Thus, a cloud service based on an instantiation of the architecture derived topology (200) may be expressed as a topology of nodes with a number of relationships defined between a number of nodes within the topology. A number of properties and actions are associated with a number of the nodes, a number of groups of nodes, a portion of the topology, the topology as a whole, or combinations thereof. Further, a number of policies are associated with the number of the nodes, a number of groups of nodes, a portion of the topology, the topology as a whole, or combinations thereof. Still further, a number of lifecycle management actions (LCMAs) are associated with the number of the nodes, a number of groups of nodes, a portion of the topology, the topology as a whole, or combinations thereof.

[0026] Thus, the present systems and methods describe cloud service broker or manager that supports both topologies and blueprints while using the same lifecycle management engine. The lifecycle management engine supports lifecycle management of cloud services, and mapping of application models with infrastructure templates. The present systems and methods also describe a policy-based framework for managing the provisioning, deployment, monitoring, and remediation processes within a cloud service. Further, the present systems and methods provide support for usage models supported by CSA, CDA, and blueprints as will be described in more detail below.

[0027] Further, the present systems and methods also describe how a monitoring system may be described by a topology and may be guided by the

policies. As described herein, the monitoring may be of a realized service that has been provisioned using the topology.

[0028] The monitoring system may perform a number of operations. For example, the monitoring system may check to ensure that deployment of the code is compliant, that the realized service code has been checked for vulnerabilities, and that the correct version has been deployed. The monitoring system may then monitor metrics and other criteria as the instantiated service runs. For example, the monitoring system may check that performance is satisfactory and that a latest patch since instantiation has been applied. In other examples, the monitoring system may check that deployment is compliant or that no security attach has been detected. In other words, the monitoring of the instantiated service may be continuous, or performed otherwise as specified in the policies. In one example, the present systems and methods relies on monitoring policies that indicate how events resulting from the monitoring are to be handled and processed (if at all). For example, it may be the case that the monitoring system is set up, and the monitoring system may operate without more human interaction.

[0029] As described above, the monitoring system may be captured by a number of monitoring policies. In some examples, the monitoring system may be captured as an additional topology, or set of topologies, in addition to the service topology. This additional topology or set of topologies may be used to describe how the monitoring system is to be provisioned, configured or built. Similarly, while the above description has indicated that the event handler may be captured by a number of policies, the event handler may also be captured as an additional topology, or set of topologies. For example, handling policies, or an event handler topology may indicate how incidents are processed and handled. Other systems may also be represented as policies, or as additional topologies. For example, the remediation engine may be captured as a number of remediation policies, or may be captured as an additional topology, or as a set of topologies.

[0030] As used in the present specification and in the appended claims, the term "broker" is meant to be understood broadly as any computing

device or a collection of computing devices in a network of computing devices that manages the designing, provisioning, deployment of a topology within the cloud, and the maintenance and life cycle management of (an) instantiated service based on that topology.

[0031] As used in the present specification and in the appended claims, the term “cloud service” is meant to be understood broadly as any number of services provided over a number of computing devices that are connected through a real-time communication network. Cloud services may include services provided on a distributed system implementing distributed hardware and software resources. In one example, a cloud service may be any service offered on a private cloud, public cloud, managed cloud, hybrid cloud, or combinations thereof. In another example, a cloud service may be services provided on physically independent machines such as, for example, a data center.

[0032] Further, as used in the present specification and in the appended claims, the terms “node or “computing device” are meant to be understood broadly as any hardware device, virtual device, group of hardware devices, group of virtual devices, or combination thereof within a network. Nodes may include, for example, servers, switches, data processing devices, data storage devices, load balancers, routers, and virtual versions thereof, among many other types of hardware and virtual devices. Further, nodes may be representations of the above hardware and virtual devices before execution and instantiation of a topology of which the node is a part.

[0033] Still further, as used in the present specification and in the appended claims, the term “topology” is meant to be understood broadly as data representing a graph of nodes where branches between the nodes represent relationships between the nodes. The nodes may comprise any number of computing devices located within a network. Thus, the topology of the network may comprise the physical and logical layout of networked computing devices, and definitions of the relationships between the computing devices. A number of policies and lifecycle management actions (LCMA) may

be associated with the topologies, portions of the topologies, nodes within the topologies, groups of nodes within the topologies, and combinations thereof.

[0034] Still further, as used in the present specification and in the appended claims, the term “blueprint” is meant to be understood broadly as an execution flow for allowing automation of cloud service deployment and life cycle management of cloud services. A blue print may include a functional description of a number of hardware and/or virtualized components included within a service such as, for example, operating systems, application stacks, databases. A blueprint may further include a functional description of the configuration and connectivity between the hardware and virtualized components. The blueprints may also include a number of deployment models to enable the functional description to be deployed. The blueprints may further include a set of user-configurable options to allow a user to establish a number of aspects of the deployed service. Blueprints are an example of non-architecture derived executable topologies.

[0035] Still further, in addition to the blueprints described above, the present disclosure provides for the utilization of executable topologies. Thus, the present systems and methods provide for the execution and instantiation of both blueprint- and architecture-derived topologies. Both blueprint- and architecture-derived topologies are executable. Thus, as used in the present specification and in the appended claims, the term “topology” is meant to be understood broadly as any set of executable logic or interpretable logic that may be expressed as executable logic that defines the characteristics of the network to be instantiated. The topology may define a number of nodes. Further, the topology may define and a number of policies and lifecycle management actions associated with the nodes as a number of groups, individually, or a combination thereof. In one example, blueprints may be expressed as topologies. In this example, the blueprint-derived topologies may also define a number of nodes and a number of policies and lifecycle management actions associated with the nodes within the topologies, groups of nodes within the topologies, portions of the topologies, the topology as a whole, and combinations thereof.

[0036] Still further, as used in the present specification and in the appended claims, the term “policy” is meant to be understood broadly as any data or metadata used to assist in the management of the provisioning, deploying, monitoring, enforcement, and remediation within a cloud service. The policies may represent a number of rules or sets of rules that are applicable to the provisioning, deploying, monitoring, enforcement, and remediation tasks associated with a number of computing devices within a cloud service environment.

[0037] Still further, as used in the present specification and in the appended claims, the term “user” is meant to be understood broadly as any individual or entity for whom or by whom a cloud service is designed, provisioned, deployed, monitored, policy enforced, incident remediated, otherwise managed, or combinations thereof. In one example, the user may purchase use of the cloud service at a cost. For example, the user may pay a subscription to use the cloud resources and services, and, in this case, also be classified as a subscriber. In another example, a user may be a designer or administrator of the cloud service. In still another example, a user may be any individual who manages the cloud service.

[0038] Even still further, as used in the present specification and in the appended claims, the term “a number of” or similar language is meant to be understood broadly as any positive number comprising 1 to infinity; zero not being a number, but the absence of a number.

[0039] In the following description, for purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present systems and methods. It will be apparent, however, to one skilled in the art that the present apparatus, systems, and methods may be practiced without these specific details. Reference in the specification to “an example” or similar language means that a particular feature, structure, or characteristic described in connection with that example is included as described, but may not be included in other examples.

[0040] The present systems may be utilized in any data processing scenario including, for example, within a network including the design,

provisioning, deployment, and management of a number of computing devices within the network. For example, the present systems may be utilized in a cloud computing scenario where a number of computing devices, real or virtual, are designed, provisioned, deployed, and managed within a service-oriented network. In another example, the present systems may be utilized in a standalone data center or a data center within a cloud computing scenario. The service oriented network may comprise, for example, the following: a Software as a Service (SaaS) hosting a number of applications; a Platform as a Service (PaaS) hosting a computing platform comprising, for example, operating systems, hardware, and storage, among others; an Infrastructure as a Service (IaaS) hosting equipment such as, for example, servers, storage components, network, and components, among others; application program interface (API) as a service (APIaaS), other forms of cloud services, or combinations thereof. The present systems may be implemented on one or multiple hardware platforms, in which the modules in the system are executed on one or across multiple platforms. Such modules may run on various forms of cloud technologies and hybrid cloud technologies or offered as a SaaS (Software as a service) that may be implemented on or off the cloud.

[0041] Further, the present systems may be used in a public cloud network, a private cloud network, a hybrid cloud network, other forms of networks, or combinations thereof. In one example, the methods provided by the present systems are provided as a service over a network by, for example, a third party. In another example, the methods provided by the present systems are executed by a local administrator. In still another example, the present systems may be utilized within a single computing device. In this data processing scenario, a single computing device may utilize the devices and associated methods described herein to deploy cloud services and manage life cycles of the cloud services. In the above examples, the design of the cloud service, provisioning of a number of computing devices and associated software within the cloud service, deployment of the designed and provisioned cloud resources and services, management of the cloud resources and services, and combinations thereof may be provided as the service.

[0042] Aspects of the present disclosure may be embodied as a system, method, or computer program product, and may take the form of hardware, an element combining software and hardware aspects that may all generally be referred to herein as a "circuit," "module" or "system." Furthermore, aspects of the present disclosure may take the form of a computer program product embodied in a number of computer readable mediums comprising computer readable program code embodied thereon. Any combination of a number of computer readable mediums may be utilized.

[0043] A computer readable medium may be a computer readable storage medium in contrast to a computer readable signal medium. A computer readable storage medium may be, for example, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, or device, or any suitable combination of the foregoing. More specific examples of the computer readable storage medium may include the following: an electrical connection having a number of wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), a compact disc read-only memory (CD-ROM), an optical storage device, a magnetic storage device, or any suitable combination of the foregoing. In the context of this disclosure, a computer readable storage medium may be any tangible medium that can contain, or store a program for use by or in connection with an instruction execution system, apparatus, or device.

[0044] Throughout the present disclosure, various computing devices are described. The computing devices may comprise real or virtual computing elements including data processing devices, data storage devices, and data communication devices. Although these various devices may be described in connection with real and physical devices, any number of the devices may be virtual devices. The virtual devices, although describing a software-based computer that is based on specifications of emulated computer architecture and functions of a real world computer, the virtual devices comprise or are functionally connected to a number of associated hardware devices. Accordingly, aspects of the present disclosure may be implemented by

hardware elements, software elements (including firmware, resident software, micro-code, etc.), or a combination of hardware and software elements.

[0045] In some examples, the monitoring policies may data or metadata provided in the same file describing the nodes or topology. The metadata may be in a separate file associated with the nodes or topology. In one example, the monitoring policies may be accessed via uniform resource identifiers (URIs) of application programming interfaces (APIs) to perform calls in order to execute the APIs.

[0046] The monitoring policies may include the monitoring topology, metadata on how to provision a monitoring system. The monitoring policies may also setup a monitoring system, or configure an existent monitoring system. The monitoring policies may also include instructions to execute the monitoring (e.g. via actual workflows to perform). In one example, instead of the monitoring policies, the monitoring system (Fig. 3, 313) may be captured as an additional topology that describes how the monitoring system is deployed, managed, or configured.

[0047] Figs. 3 and 4 depict a block diagram of a topology-based management broker (300) along with a designing phase for provisioning, deploying, monitoring, protecting and remediating a cloud service, according to one example of the principles described herein. The system of Figs. 3 and 4 support both topologies and blueprints while using the same lifecycle management engine as will be described in more detail below.

[0048] As depicted in Figs. 3 and 4, one or a number of topology designers (301) contribute in designing various aspects of the cloud service topology. In one example, topology design is performed via a design tool that uses hardware devices and modules such as graphical user interfaces (GUI) and coding scripts. A human designer designs the topology with the use of a design tool (301). Thus, the design of the topology (302) is achieved through a combination of autonomous and human-provided design methods. In one example, the topology designer (301) may be an interface utilizing API's that enables separate creation of an application model (Fig. 4, 319) and its

associated components along with creation of an infrastructure template (Fig. 4, 320) which specifies infrastructure and lifecycle conditions for the infrastructure.

[0049] The subsystem depicted in Fig. 3 of the overall topology-based management broker (200) comprises a subsystem capable of provisioning, deploying, monitoring, enforcing policies within a cloud service, and remediating incidents within the cloud service. These tasks are all performed with the use of topologies with LCMAs and policies, whether the topologies are blueprint or architecture derived. Thus, the present systems and associated methods also support all the use cases that CSA 3.2 supports. As described above, CSA 3.2 is an automation system tool used to deploy and manage cloud computing applications, and is developed and distributed by Hewlett Packard Corporation. CSA 3.2 technologies are capable of supporting blueprints or architecture topologies. Further, CSA is described in International Patent App. Pub. No. PCT/US2012/045429, entitled "Managing a Hybrid Cloud Service," to Maes, which is hereby incorporated by reference in its entirety. As will be described in more detail below, the subsystem depicted in Fig. 3 uses a number of types of policies and lifecycle management actions (LCMAs) to provision, deploy, monitor, enforce policies within, and remediate incidents within a deployed cloud service.

[0050] Further, the subsystem depicted in Fig. 4 of the overall topology-based management broker (200) comprises a subsystem capable of independently modeling infrastructure and application requirements of a topology on the same stack as the subsystem depicted in Fig. 3. The present systems and associated methods also support all the use cases that a CDA subsystem such as those use cases of CDA 1.2 support. As described above, CDA is an automation system tool utilized within a topology designer that independently models infrastructure and application requirements while managing versions, configurations, and other application components. CDA 1.2 is also developed and distributed by Hewlett Packard Corporation. Further, CDA is described in International Patent App. Pub. No. PCT/US2012/041625, entitled "Cloud Application Deployment," to Maes, which is hereby incorporated by reference in its entirety.

[0051] In this manner, the subsystems of Figs. 3 and 4 work under a common stack and work together within the topology-based management broker (200) as a single computing system with common use of topologies, realized topologies, and policies to support all use cases of constructing topologies and supporting multiple providers' associated technologies. Thus, in one example, the present systems and methods reconcile the differing models, templates, and blueprints used respectively by CDA and CSA by utilizing, on the same stack, designed topologies (preferably architecture derived) of a cloud service, a number of policies, and a number of LCMAs associated with the topology nodes / subsets/full.

[0052] As depicted in Fig. 3, a topology designer (301) may design and present a lifecycle management (LCM) topology (302) to the topology-based management broker (200). In one example, the topology designers (301) described herein may be an integrated part of the topology-based management broker (200). In another example, the topology designers (301) may be separate from the topology-based management broker (200). In another example, a number of persons may use the topology designers (301) to design the topologies (302). These individuals may be service designers, infrastructure architects or administrators, system administrators, information technology operators, offer managers, or users, among other personnel with roles in the design of a topology. In still another example, the topology designers (301) may be operated by a third party.

[0053] The LCM topology (302) may define a number of nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7), and a number of relationships between the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7). Although in Fig. 3, seven nodes are depicted, any number of nodes may be designed into the topology (302) to achieve any data processing objectives. In one example, the topology-based management broker (200) may represent the topology (302) as an extensible markup language (XML) file. In another example, the topology-based management broker (200) may represent the topology (302) in JavaScript object notation (JSON) format; a text-based open standard designed for human-readable data interchange that is derived from

the JavaScript scripting language for representing objects. In still another example, the topology-based management broker (200) may represent the topology (302) in YAML syntax format; a human-readable data serialization format.

[0054] In Fig. 3, the relationships between nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) are depicted as lines connecting the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7). Each of the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7), the entire topology (302), a group of nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7), portions of the topology (302), or combinations thereof are associated with a number of policies (303). Policies (303) are data or metadata provided in the same file describing the nodes or topology, or in a file associated therewith. In one example, the association of the policies (303) within the topology (302) may be performed during the designing of the topology (302), by, for example, an administrator when offering the design. In another example, the association of the policies (303) within the topology (302) may be performed during the designing of the topology (302) when a user, for example, selects the design as a subscription or request.

[0055] Further, in one example, the addition of a policy (303) to the topology or portions thereof may cause the design of the topology to change. In this example, a policy (303) added to an element of the topology (302) may effect a number of other policies. For example, associating with a topology (302) a policy that indicates that a node be highly available may evolve the policies (303) and topology (302) as a whole to direct, for example, a cluster of nodes. In this manner, policies may drive the design of the topology (302).

[0056] Each of the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7), the entire topology (302), a group of nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7), portions of the topology (302), or combinations thereof are further associated with a number of lifecycle management actions (LCMAs) (304). In examples where LCMAs (304) are associated with the nodes, a single LCMA is associated with a given node. In examples where a number of LCMAs

are associated with portions of the topology (302) or the topology (302) as a whole, the LCMAs are subjected to an orchestration of resource providers.

[0057] LCMAs are expressed as a number of application programming interfaces (APIs), wherein the LCMAs are called during execution of the topology (302), and a number of computing resources are provisioned for purposes of managing the lifecycle of a given cloud capability. In one example, the LCMAs may be accessed via uniform resource identifiers (URIs) of application programming interfaces (APIs) to perform calls in order to execute the APIs. In one example, the LCMAs are provided by reference within the file comprising the data or metadata described above in connection with the policies (303).

[0058] In one example, the LCMAs are associated with the aspects of the topology by default by virtue of what computing device the node or nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) represent. In another example, the LCMAs are associated with the aspects of the topology by explicitly providing a number of functions, F_{Action} , that define how to select a resource provider to implement the action based on the policies associated with the aspects of the topology and the policies of the different relevant resource providers. These functions define how a resource provider is selected to implement the action based on the policies associated with the aspect of the topology and the policies of the different relevant resource providers.

[0059] The policies and LCMAs will be denoted herein by elements 303 and 304, respectively, to denote that the policies (303) and LCMAs (304) are associated with the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7), the entire topology (302), a group of nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7), portions of the topology (302), or combinations thereof. In one example, the association of the policies and LCMAs with aspects of the topology is performed via the topology designer (301).

[0060] In one example, although not depicted, a subset of nodes making up a group may also be associated with a number of policies (303) and a number of LCMAs (304). In this example, a number of nodes, for example, nodes (302-2, 302-3, 302-4, 302-6, 302-7), may be associated as a group with

a number of policies (303) and a number of LCMAs (304) associated therewith. Several groupings of the nodes may be present within the entire topology (302). In one example, the groups of nodes may overlap, in which a single node in a first group of nodes may also belong to a second group of nodes, and be subjected to both the first and second groups of nodes' policies (303) and LCMAs (304). Policies and their associations with individual nodes and groups of nodes will be described in more detail below.

[0061] The policies (303) associated with the nodes may be expressed and attached with the nodes in any manner (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7). In one example, the policies (303) are associated with the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) by defining properties of the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7). In another example, the policies (303) are associated with the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) by metalanguage expressions.

[0062] The policies (303) are a number of descriptions, metadata, workflows, scripts, rules, or sets of rules that are applicable to guiding the provisioning, monitoring, enforcement, governance, and remediation tasks associated with the lifecycle management of a number of nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) within a cloud service environment in which the topology (302) is to be or has been implemented. The policies (303) define the access control and usage control of the APIs of the topology-based management broker (200). Further, policies (303) define the access control and usage control of the APIs used to manage or use the instantiated services. For example, when a security threat is detected by a monitoring system (313), a remediation option may comprise making changes to a number of access control policies.

[0063] The policies (303) may be associated with and operable against a number of individual nodes, a number of groups of nodes, a number of nodes of a class of nodes, a subset of the nodes within the entire topology of the cloud service; the entire topology of the cloud service as a whole, or combinations thereof. If the policies (303) are initiated on the individual nodes, groups of nodes, or the entire topology of the cloud service as a whole, the

policies will guide how life cycle management actions are taken with respect to, or performed on the individual nodes, groups of nodes, nodes of a class of nodes, a subset of the nodes within the entire topology of the cloud service, or the entire topology of the cloud service as a whole.

[0064] One example of a type of policy is a provisioning policy. Provisioning policies may, if implemented, define the characteristics of the computing devices that comprise the cloud when the topology is provisioned, deployed, and executed. This provisioning can include the infrastructure and platform of the topology (302). The provisioning policies may include definitions of characteristics such as, for example, the physical location of a node. Provisioning policies may also include definitions of characteristics such as, for example, a geographical or deployment type location such as a network zone with or without access to an internet or behind or not behind a firewall, among other geographical or deployment type provisioning policies. In this example, a policy may have a provisioning policy component that may be associated with a server device that directs the server device to be located in a particular geographic area of a country, a particular region such as, for example, the east coast of the United States versus the west Coast, a particular server facility, or any other geographic location.

[0065] As to a provisioning policy that defines a physical location of the computing device, other characteristics may include, for example, the level of security of the location or access to the internet at which the node is located. Other provisioning policies may also include, for example, the speed in, for example, bandwidth of the network to which the node is coupled, whether the node is to be connected to an internet or intranet such as, for example, a demilitarized zone (DMZ) or perimeter network, whether the node is firewalled, whether the node has access to an internet, whether the node is to be located on top of another node, and whether the node is to be located on top of another node using a particular infrastructure element or platform, among other provisioning policies.

[0066] Provisioning policies may also, if implemented, rely on the requirements and capabilities of the nodes within the proposed cloud service

that is based on the topology (302). Requirements define the needs of nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) such as, for example, server or network needs in relation to processing, memory, and operating system (OS) needs, among other forms of needs. For example, the requirements policies may indicate that a node requires particular software or a particular software version associated with it such as a particular operating system. As another example, a requirements policy may also indicate that a particular node may require additional hardware devices associated with it such as, for example, a server device, a server group, or a high availability configuration, among others.

[0067] Capabilities such as the nature of the processors, memory, capacity, OS, middleware type and version, among others, define what each node (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) offers. Thus, in one example, capabilities policies may indicate that a node is capable of processing data at a certain rate. In another example, a capabilities policy may indicate that a memory device may have a terabyte (TB) of data storage space.

[0068] In still another example, the requirements policies may indicate that a node requires a particular computing platform. When designing a topology (302), the topology or association of metadata supports capturing data defining a number of hardware devices within the computing platform including hardware architecture and a software framework (including application frameworks). When the metadata is presented or associated, it is used to guide provisioning policies in order to better select appropriate elements within the computing platform such as, for example, a suitable data center. The metadata, when presented or associated, may also be used to guide matching fragments of topologies to other fragments as will be discussed in more detail below in connection with stitching of application models to infrastructure templates.

[0069] With regard to capability policies, the nodes may define what kind of device they are, what versions of software they capable of, or are executing, and what they can do. An example, of a capability policy may include a definition associated with the node that defines it as an application server, that it provides a Java Platform, Enterprise Edition (J2EE) environment,

that it runs a particular operating system, a version of an operating system, or a particular release of a version of the operating system, among many other capabilities. As described above, this may be used to determine, for example, what else may be deployed or what other devices may use the cloud services.

[0070] Another type of policy (303) that may be assigned includes monitoring policies. Monitoring policies are policies that, if implemented, define the operational monitoring of the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7), the security monitoring of the nodes, the compliance monitoring of the nodes, analytics among the nodes and groups of nodes, usage monitoring of the nodes, performance monitoring, and intelligence monitoring such as, for example, collection of metrics, business intelligence (BI) and business activity monitoring (BAM) and analytics/big data integration, among other types monitoring-related policies. In some examples, the monitoring policies (303) can provision or setup a big data system that collects data about a broker that deploys the topology, the associated systems (e.g. the monitoring systems, remediation system, events, among other associated systems), data produced by the instantiated cloud services, among other types of data. Applications may then be built based on a big data platform to generate events or derive other information. Applications may also behave like other monitoring systems (e.g. detecting operations anomalies, usage anomalies, security anomalies, among other functional anomalies).

[0071] The monitoring policies may also define what kind of monitoring is expected and how the monitoring is to be implemented. Examples of monitoring policies regarding node operations include performance, monitoring CPU levels and loads of the various nodes within the network, monitoring the speed at which data is processed through a node or a number of nodes or exchanged between nodes, and monitoring the operational state of applications running on a node or nodes at any level of the network, among many other operations parameters of the nodes, group of nodes, and the cloud service as a whole.

[0072] In another example, the monitoring policies also define how monitored events that occur in an instantiated topology are handled. In this

example, the monitoring policies assist an event handler (316) in receiving and processing the events, and in making decisions regarding remediation of incidents resulting from the events, and in sending notification messages regarding the incidents. The handling of events within the topology-based management broker (200) will be described in more detail below. As will be described in more detail below, the monitoring policies include a portion that defines what to do with the monitored events that result from the monitoring such as, for example, how to handle the events, where the events are sent, what devices or individuals address the events, how incidents resulting from the processing of the events are handled, how the events and incidents are processed (e.g., processed as aggregated, filtered, or correlated events, among other forms of processing), and how the resulting incidents are handled.

[0073] Monitoring policies also include monitoring policies regarding security. Security policies define how to monitor for abnormal behaviors or behaviors known as being associated with known or suspected security issues. Examples of monitoring policies regarding security include monitoring whether a node or a group of nodes is experiencing an attack, whether there is strange behavior occurring within the cloud service or interactions with the cloud service, and whether there is a virus or other anomaly with a node or group of nodes, among other security-related monitoring policies.

[0074] Monitoring policies also include monitoring policies regarding compliance. Examples of monitoring policies regarding compliance include, determinations as to whether the nodes or group of nodes are running an appropriate version of an operating system, determining whether the most recent patch associated with the release of a software program running on the nodes has been installed, determining if an installed patch is a correct patch, checking that a code or artifacts that have been used to provision the node and cloud service have been appropriately checked and vetted for any weakness or problem, if governance and access control to the node and cloud service or the node and cloud service management is appropriate, and if settings of a provisioned system match provisioning, security, or other compliance requirements such as correct logging levels, correct setup for access controls,

and correct setup for passwords, among other compliance-related monitoring policies.

[0075] Monitoring policies also include monitoring policies regarding usage. Examples of monitoring policies regarding usage include, determining how much a user has been using CPUs of a node or group of nodes, determining how much memory a user has utilized, determining how much money has been charged to the user, and determining whether the user has paid for the services provide through the designing, provisioning, deploying, and monitoring of the network topology, among other usage-related monitoring policies.

[0076] The policies (303) may further comprise governance policies that, if implemented, define access controls of nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) or groups of nodes within the topology (302) or the cloud service. For example, governance policies may include policies that define who may access the nodes within the topology (302) or the cloud service, and under what conditions may those individuals obtain such access.

[0077] The policies (303) may further comprise analytics policies that, if implemented, define what is needed to ensure analytics and big data monitoring within or among the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) or groups of nodes within the topology (302), and ensure that this is occurring as expected. For example, the analytics policies may define a number of workflows by which the monitoring system (313) may operate to configure the cloud service, provide analytics, collect big data, and process the data.

[0078] Still further, the policies (303) may comprise remediation policies that define what actions are to take place within the topology (302) should a problem arise or an incident be raised during deployment and execution of the topology (302). Remediation policies may include policies that define a number of actions taken by the topology-based management broker (200) during remediation processes, and include: (1) providing notifications to a user, consumer, or administrator; (2) obtaining instructions from the user, consumer, or administrator; (3) taking manual actions input by the user,

consumer, or administrator; (4) taking autonomous actions after receiving instructions from the user, consumer, or administrator; (5) taking autonomous actions without receiving instructions from the user, consumer, or administrator; (6) taking autonomous actions without notifying the user, consumer, or administrator or receiving instructions from the user, consumer, or administrator; (7) proposing a remediation action to a user or administrator for approval, and performing the proposed remediation action if approved by the user or administrator, or combinations thereof.

[0079] As an example, a failure, or poor performance, of the cloud service as instantiated or realized by the topology (302) may occur. The failure or degradation in performance may occur at any time after instantiation. A used herein the term failure may include not meeting the performance metrics, or performance expectations. The remediation policies may define how that failure may be handled based on the above potential scenarios. In addition, the remediation policies provide the actual rules and workflows of actions to perform to remediate the incidents under any number of conditions or indicate to whom or which device to delegate the decision making and orchestration and fulfillment of these remediation actions. Another remediation example may regard a potential need to maintain a level of service based on, for example, a service level agreement (SLA), or a quality of service (QoS) within the cloud service that is realized based on the topology (302). In this example, the addition of resources to support the increase in demand for resources may be handled based on the above potential scenarios. More details regarding monitoring of the deployed topology and event handling therein will be described in more detail below.

[0080] As described above, the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) may include a number of lifecycle management actions (LCMA) (304) associated with the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7). The LCMA (304) are a number of actions associated with the policies (303) that are executed by a processor when triggered by the policies (303) within a cloud service environment in which the topology (302) is implemented. The LCMA may be associated with and operable against a

number of individual nodes, a number of groups of nodes, a number of nodes of a class of nodes, a subset of the nodes within the entire topology of the cloud service; the entire topology of the cloud service as a whole, or combinations thereof. If the LCMA's are executed with respect to the individual nodes, groups of nodes, or the entire topology of the cloud services as a whole, the LCMA's will take an action with respect to the individual nodes, groups of nodes, the nodes of a class of nodes, a subset of the nodes within the entire topology of the cloud service, or the entire topology of the cloud service as a whole as defined within the LCMA's. LCMA's (304) include actions such as, for example, provisioning of computing resources within the topology, updating the topology, copying all or portions of the topology, modifying computing resources within the topology, moving computing resources within the topology, destroying or deleting resources within the topology, among other lifecycle management actions.

[0081] The various policies described herein define what actions are to be performed throughout the lifecycle of the topology (302) before, during, and after instantiation of a service based on the topology (302). Further, the various policies described herein define how these actions are to be performed. Still further, the various policies described herein define which device, individual, or combination thereof to which the actions are to be delegated. Even still further, the various policies described herein define combinations of the above. For example, any of the monitoring policies used in event handling and processing, or remediation may define what devices or portions of the cloud service are to be monitored or remediated, how to execute such monitoring and remediation, to whom or what devices to delegate the roles of monitoring and remediation, or a combination thereof.

[0082] Different policies play different roles at different times within the lifecycle of a topology. Further, the different policies may be executed at different times of the lifecycle of the cloud service and throughout the flows of the topology-based management broker (200). Fig. 5 is a block diagram of an execution flow of the execution of a topology (302) using provisioning policies, according to one example of the principles described herein. For example, this execution flow may be implemented when there are existent policies. In the

example of provisioning policies with their number of capabilities and requirements, a topology (302) may be designed with a number of associated policies (303) as described above. As depicted in Fig. 5, the topology (302) with its associated policies (303) may be an input (501) to a provisioning policy engine (502) that may evaluate provisioning policies including capabilities and requirements with context information and resource provider policies to determine which resource provider to use to perform LCMAs. After evaluating the provisioning policies and selecting the resource provider, an interpreter (503) may evaluate the other policies (e.g., monitoring policies, event handling policies, event processing policies, incident handling policies, among other policies). The interpreter (503) may also convert policies such that they may be implemented. For example, the interpreter (503) may convert policies into topologies associated with the topology (Fig. 3, 302), nodes of the topology (Fig. 3, 302) or a subset of nodes of a topology (Fig. 3, 302). The interpreter (503) may associate, delegate or configure resource providers, workflows, scripts or combinations thereof or provide actual instructions to implement a monitoring system. In some examples, if a monitoring system is already described, the monitoring and event handling topologies may be an input (501) to a provision policy engine (502) as described above. In one example, the topology (302) may be an architecture based topology. A policy provisioning engine (502) may be a standalone device or incorporated into a device of Fig. 1A such as, for example, the resource offering manager (308). The policy provisioning engine (502) may obtain a number of provisioning policies from a resource provider called resource provider policies (PR) (308-1), a number of provisioning policies as defined by a user, a number of policies as defined by the topology designer (301), or combinations thereof. In some examples, an instantiated monitoring system may be associated to a realized topology that covers many details of the node and monitoring system, which details may be stored in the RTSM (Fig. 3, 315). As described above, the monitoring system and event handling system may be stored as a realized topology. Similarly as described above, policies and LCMA may remain associated to the realized topology of the service. Accordingly, when a realized topology is updated for

any reason, 1) the monitoring system and event handling system are updated in a similar fashion and 2) the policies in the realized topology, for example to monitor additional or modified nodes or to monitor consistently if a service or portion of a service is moved, to another cloud for example

[0083] Resource provider policies (308-1) may be any policies that associated with a number of resource providers' offerings that guide the selection of a number of resources. In one example, the resource provider policies (308-1) may be dynamic functions that define the computing abilities of a computing resource. In this example, a computing resource that provides a defined level of computing resources such as, for example, processing power may be provisioned by the LCM engine (311) and resource offering manager (308) if the defined level of that computing resource meets a number of requirements within the topology (302).

[0084] Further, in one example, the addition of a policy (303, 308-1) to the topology or portions thereof may cause the design of the topology to change. In this example, a policy (303, 308-1) added to an element of the topology (302) may effect a number of other policies (303, 308-1). For example, associating with a topology (302) a policy that indicates that a node be highly available may evolve the policies (303) and topology (302) as a whole to direct, for example, a cluster of nodes. In this manner, policies may drive the design of the topology (302).

[0085] Accordingly, a designed topology such as, for example, an architecture topology generated, for example, by an automated or manual matching process with policies and LCMA's associated with the nodes of the topology (302) is modified at the time of provisioning. This may be performed by executing, with the provisioning policy engine (502) or the resource offering manager (308), the provisioning policies to determine a topology that satisfies the provisioning policies perfectly or in the best obtainable manner. Obtaining a best fit topology may involve a number of resource provider policies (308-1) provided by the resource offering manager (308) which describe the capabilities and selection criteria of a resource provider. The resource offering manager

(308) selects, for example, the resource provider from which the resource is to be obtained, and may also make other modifications to the topology (302).

[0086] The topology (302) is modified per the received provisioning policies (308-1) by the provisioning policy engine (502) as indicated by arrow 507, and sent to an interpreter (503). The interpreter (503) is any hardware device or a combination of hardware and software that interprets the provisioning policies to create an execution plan as indicated by arrow 508. The result is then interpreted and converted into an execution plan (508) that comprises a workflow or sequence of serial and/or parallel scripts in order to create an instance of the topology (Fig. 1A, 312). In one example, the interpreter (503) is a standalone device or is contained within the LCM engine (Fig. 1A, 311). The execution plan (508) comprises a number of workflows or sequences of serial and/or parallel scripts. The topology LCM engine (311) obtains the workflows or sequences of serial and/or parallel scripts, calls a resource provider via the resource offering manager (308) as indicated by arrow 509, and creates an instantiated service (312) at block 505. Assuming the workflow or sequence of serial and/or parallel scripts is executable, which it should be in the case of an architecture descriptive topology, the actions associated with the workflow or sequence of serial and/or parallel scripts are executed by the LCM engine (311).

[0087] With the above-described sequence based topology, an execution plan (508) may be represented as a blueprint. Conversely, a blueprint may be expressed as an execution plan (508). A blueprint with nodes expanded by policies (303) and LCMAs (304) may be similarly processed, instead, in a manner similar to the processing of an infrastructure topology. In this example, the blueprint in the form of a sequential service design (506) is input to the interpreter for processing as described above in connection with Fig. 5.

[0088] The execution of the execution plan (508) by the topology life cycle management engine (311) results in an instantiation of the cloud services including the provisioning of devices for monitoring, event handling, and processing and remediation of events and incidents as will be described in

more detail below. The result of the topology life cycle management engine (311) executing the workflow or sequence of serial and/or parallel scripts as defined by the execution plan (508) is an instantiated service (312) as indicated by block 505. Further, a realized topology (314) may be created based on the instantiated service (312), and stored as will also be described in more detail below.

[0089] As to the monitoring and remediation policies described herein, the same type of process may be applied, but with a number of realized policies defined within an instantiated service (312) and its realized topology (314) as input. In this process, additional LCMAs (304) may be created and used to assist in provisioning resources in an updated instantiated service (312). The explanation below across CSA/CDA use cases with architecture topologies or with blueprints shows the notion of common engine, pattern, and platform across all these cases.

[0090] The present systems and methods may be used in conjunction with any third party modeling such as, for example, HEAT command language interpreter that is an open source software developed and distributed by the OpenStack Foundation and released under the terms of the Apache License. Although HEAT may assist in the creation of a set of scripts fitting in the space of the execution plan, HEAT may provide support by interpreting or translating data, and converting the data into scripts. The present systems and methods may add the policies (303) and LCMAs (304) to the HEAT interpreter, and execute as described above.

[0091] Further, the present systems and methods may use topology and orchestration OASIS specification for cloud applications (TOSCA), a cloud computing standard to express topologies. In this example, the policies (303) and LCMAs (304) are added to a TOSCA-based topology.

[0092] Thus, the policies (303) and the LCMAs (304) may be implemented as function calls (305) or scripts in order to provision and deploy the topology (302) when the policies (303) and the LCMAs (304) trigger such provisioning and deployment. A resource offering manager (308) may be provided within the topology-based management broker (200) to manage and

provide computing resources within the topology (302) based on the policies (302) and LCMAs (304).

[0093] The resource offering manager (308) provides a number of plug-ins to execute the life cycle manager (311). As described above, the resource offering manager (308) associates a number of resource policies (308-1) to the plug-ins of a number of resource providers so that the resource providers may assist in guiding the selection process of a number of the resource providers. The non-resource provider policies such as policies (103) associated to the nodes are different in that they are associated with the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) during the designing of a topology (302).

[0094] The resource offering manager (308) may be operated by, for example, an administrator, or a service provider in order to provision the resources within the cloud service to be created via the deployment of the topology (302). As discussed above, the resource offering manager (308) comprises the hardware and software to define a number of resource provider policies (308-1), associate a number of those resource provider policies (308-1) with a number of the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7), the topology (302), or portions of the topology (302). When the topology (302) is deployed, the resource offering manager (308) provides the computing resources to the user that will implement the topology (302) based on the policies (303), the LCMAs (304), and the resource provider policies (308-1). As a result, the LCMAs are functions of the policies (303) associated with the topology (302), and the resource provider policies (308-1).

[0095] Thus, in one example, the resource offering manager (308) may implement a number of resource provider policies (308-1) that define under which conditions a computing resource from a number of service providers may be selected for deployment within the topology (302). In this example, the policies (303) associated with a node as well as the resource provider policies (308-1) define which resource offering from the resource offering manager (308) is selected for provisioning within the to-be-deployed instantiated topology (312). For example, if a policy associated with node (302-1) indicates that the

provisioned computing resource be located in a secure facility, and the policies of the resources offered by the resource offering manager (308) indicate that those available computing resources are not located in a secure facility, then that non-secure computing resource provided by that particular service provider will not be selected. In this manner, the policies associated with the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) and the policies associated with the resource offering manager (308) determine which computing resources may be provisioned and deployed within the topology (302).

[0096] The topology-based management broker (200) may store the topology (302) in a catalog (310). In one example, the topologies (302) designed and stored in the catalog (310) may be made available to any interested party via a self-service portal (309). In another example, an application program interface (API) may be provided instead of or in addition to the self-service portal (309). In this example, the API may be used by an application executing within the topology-based management broker (200) which may make a request from the catalog (310) for a number of topologies (302).

[0097] In another example, the user may be given the opportunity to view the catalog (310) of stored topologies to obtain a topology that was created for a first user or organization, and use a number of those topologies as the user's topology by ordering or subscribing to a topology (302). In still another example, the user may be given the opportunity to view the catalog (310) of stored topologies to obtain a topology that was created for a first user or organization, obtain a number of those topologies, and add a number of other topologies to it such as in an example where an application model is built on an infrastructure template using stitching processes described below.

[0098] In still another example, the user may be given the opportunity to view the catalog (310) of stored topologies to obtain topologies that were created for a first user or organization, obtain a number of those topologies, and add a number of other topologies to it such as topologies designed de novo or stored within the catalog (310). In still another example, the user may be given the opportunity to view the catalog (310) of stored topologies to obtain

topologies that were created for a first user or organization, obtain a number of those topologies, and build a new cloud service that comprises aspects of all the predefined topologies and the respective services described by the predefined topologies.

[0099] In another example, the user, a service designer, or a combination thereof may design the topology anew, design a topology based on a topology stored in the catalog (310), or design a topology based partially on a topology stored in the catalog (310). Design of a topology (302) may be split among a number of users, designers, and administrators. The designing of the topology (302) may include separating the design of the topology into a number of topologies and attaching to the separate pieces of the individual topologies and the topology as a whole a number of properties, LCMAAs, and policies. The user may also order a desired topology, be given an opportunity to approve of the chosen topology, and view and operate the topology by executing a number of applications on the resultant cloud service.

[00100] In another example, an application program interface (API) may be made available that invokes the call functions associated with the desired topology (302). In the self-service portal (309) example, the catalog (310) may be made available to the user, may identify to the user the item or items associated with the desired topology (302), may provide the ability for the user to order a number of services, and provide for the deployment of the selected topology (302). In one example, the deployment of the topology (302) may be approved by the user or a manager as defined by an approval workflow before deployment based on, for example, a service level agreement (SLA), cost of the cloud services, and the policies, among other considerations. In still another example, once the topologies (302) are designed and stored in the catalog (310), the topologies (302) may be identified by commercial terms and associated descriptions of how the topology (302) may be used.

[00101] When a topology (302) has been designed, the topology (302) may be provisioned on behalf of the user to create a subscription within the SLA. The topology lifecycle management (LCM) engine (311) is a data processing device that will execute the topology (302) to provision and deploy

computing resources to form the cloud service for use by the user. The topology LCM engine (311) analyzes the topology (302) created, and creates a set of scripts that form execution logic in the form of the execution plan to instantiate and realize the topology (302). In one example, the set of scripts define a sequence of provisioning and deployment of computing resources. The topology LCM engine (311) applies the policies associated with the topology (302) and the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) of the topology (302) as well as the policies set for the resources managed by the resource offering manager (308).

[00102] As a result of the above systems and methods, an instantiated service (312) is provided to the user for use. The instantiated service (312) comprises a number of computing devices that match the designed topology (302) and the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) within the topology (302). The instantiated service (312) functions based on the policies described above. The computing devices that make up the instantiated service (312) may comprise, for example, servers, switches, client devices, and databases, among many other computing devices. A realized topology (314) is derived by the LCM engine (311) or other device based on the instantiated service (312).

[00103] In addition to the instantiated service (312), a monitoring system (313) is also deployed if not already existent, or setup and configured if already available in order to monitor the instantiated service (312). With the inclusion of a monitoring system (313) within the topology-based management broker (200), the topology-based management broker (200) provides for a converged management and security (CM&S) environment. In one example, the CM&S environment may be the CM&S environment developed and distributed by Hewlett Packard Corporation. In another example, the CM&S environment may be the CM&S environment described in International Patent App. Pub. No. PCT/US2012/059209, entitled "Hybrid Cloud Environment" to Maes et al., which is hereby incorporated by reference in its entirety. The CM&S environment provides for template- and model-based approaches to application and service development and deployment, with the ability to bind

management and security capabilities to service models at deployment time in order to ensure common capabilities across hybrid cloud environments. CM&S also provides portability across private and public cloud environments, which may include heterogeneous infrastructures, management, and security tools. Further, CM&S provides efficient delivery and management of the application release, whether the infrastructure resources are on premise, in the public cloud or in a hybrid environment across public and private clouds. CM&S also provides role-based, predictive, and real-time performance and risk insights, and analytics such as, Business Intelligence (BI), Business Activity Monitoring (BAM), and big data analyses across heterogeneous systems, networks, and cloud environments.

[00104] In one example, the monitoring system (313) operates based on the monitoring policies associated with the topology (302) and the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) of the topology (302) as described above. In this example, the monitoring system (313) is used to monitor the operations, the security, the compliance, and the usage of the topology (302) and the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) of the topology (302), among other items to monitor within the instantiated service (312).

[00105] In one example, the monitoring system (313) is deployed to monitor the instantiated service (312) in cases where the monitoring system (313) already exists. In this example, a number of existing monitoring devices may be used to monitor the instantiated service (312) autonomously, through human intervention, or a combination thereof by configuring the existing monitoring system (313) to match the monitoring policies defined when designing the topology (302). In this example, the monitoring system (313) already existent may be configured based on the monitoring policies associated with the topology (302) and the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) of the topology (302) as described above, configured based on input from a user, or combinations thereof.

[00106] In another example, a previously non-existent monitoring system (313) may be provisioned and deployed based on the monitoring

policies defined when designing the topology (302). In this example, the monitoring system (313) is provisioned and deployed at the same time as the provisioning and deployment of the instantiated service (312). Further, the monitoring system (313), in this example, is deployed and managed based on the monitoring policies associated with the topology (302) and the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) of the topology (302) as described above. In any of the above examples, a complete service as outlined by the topology (302) is created, including the instantiated system (312) and the monitoring system (313).

[00107] The topology-based management broker (200) further comprises a realized topology system management (RTSM) database (315). The RTSM database (315) is a logical system repository of realized topologies (314), and may be any form of data repository. In one example, the RTSM database (315) comprises a database management system (DBMS). The DBMS is a combination of hardware devices and modules that interact with a user, other applications, and the database itself to capture and analyze data. In one example, the RTSM database (315) is a configuration management database (CMDB). A CMDB is a repository of information related to all the components of a realized topology (314).

[00108] The DBMS of the RTSM database (315) is designed to allow the definition, creation, querying, update, and administration of a database of realized topologies (314). Realized topologies are a model of the topologies (302) with the policies described above associated therewith. Thus, the realized topology (314) comprises a model of the topology (302), with the policies applied to the various nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7). A number of properties of the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) of the realized topology (314) are defined within the realized topology (314). These properties include any details of any instantiated service (312) that is created or updated via the topology-based management broker (200), and may include, for example, the internet protocol (IP) address of the nodes, and characteristics and computing parameters of the nodes, among many other properties.

[00109] The RTSM (315) is a repository that stores each instance of a realized topology (314). In this manner, every time a topology (302) is designed, provisioned, and deployed, the topology-based management broker (200) captures the realized topology (314) of that topology (302). Thus, the RTSM (315) contains a realized topology (314) of every topology (302) that has been instantiated within the topology-based management broker (200) or, through the below-described remediation processes, stores a modification of a realized topology or an instantiated service (312). Thus, in one example, in every instance of the modification of an existing topology (302), the realized topology (314) resulting from that modification is also stored within the RTSM (315). The remediation processes will now be described in more detail.

[00110] As may happen within the topology-based management broker (200), a number of events may occur within the topology-based management broker (200). These events may include, for example, a policy failure within a node of the instantiated service (312), a failure of a number of hardware or software components within the instantiated service (312), and an unauthorized access of the instantiated service (312), among many other computing-related events. Further, the monitoring system (313) monitors a number of performance- and utilization-related events that may occur within the instantiated service (312). These performance- and utilization-related events may include, for example, processor utilization within a number of the nodes, utilization of a number of the nodes by, for example, customers of the user's business, and levels of remaining data storage space within a data storage device, among many other performance- and utilization-related events.

[00111] In one example, the monitoring system (313) informs the event handler (316) of any events detected by the monitoring system (313). The event handler (316) is any computing device that receives data associated with detected events from the monitoring system (313), and processes the data in order to create a number of incidents that may arise from the detected events.

[00112] Thus, the topology-based management broker (200) processes the events that are detected by the monitoring system (313).

Processing of events detected by the monitoring system (313) may be performed by the event handler (316). The event handler (316) may receive any kind or amount of data from the monitoring system (313). As described above, the data received from the monitoring system (313) by the event handler (316) may include any data associated with the operation and usage of the instantiated service (312) as a whole, and the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) within the instantiated service (312) as groups of nodes and as individual nodes. In one example, the event handler (316) performs a number of requests for the event data. In this example, the event handler (316) may poll the monitoring system (313) for the event data after a predefined time period, randomly, when triggered by another event, or a combination thereof. As described above, event handling and processing may, in one example, be delegated to another system or third party service. For example, event handling such as correlation and filtering of events and incidents and incident identification may be delegated to HP BUSINESS SERVICE MANAGEMENT; a suite of service management software tools developed and distributed by the Hewlett Packard Corporation. Remediation processes may be delegated to HP OPERATIONS MANAGER I (HP OMI) or SITESCOPE; both comprising a suite of software tools developed and distributed by the Hewlett Packard Corporation. Security event notification, processing, and remediation may be delegated to HP ARCSIGHT; a suite of service management software tools developed and distributed by the Hewlett Packard Corporation. In one example, HP ARCSIGHT may reference the service agreement (SA) associated with the instantiated service (312) to comply with the SA.

[00113] The data received from the monitoring system (313) is processed by the event handler (316), and the event handler (316) determines whether an event indicates a remediation action, and whether and how to present a notification of the event to a user, administrator, third party, or other user of the topology-based management broker (200) or instantiated service (312). If the event handler (316) determines that a remediation action is to be taken in connection with an event, the event handler (316) generates an incident based on the event, and the data associated with the event is sent to a

remediation engine (317). In one example, the event handler (316) may process the events received from the monitoring system (313) using a number of processing types. Types of processing that the event handler (316) may perform include filtering, correlation, and aggregation of the events, among other forms of event processing, and combinations thereof. In one example, a number of events may collectively be subjected to a number of forms of event processing in order to create an incident. In this example, the events may individually not support the creation of an incident that indicates remediation, but a number of events, when analyzed by the event handler (316), may indicate that an issue within the instantiated topology (312) is not in agreement with the policies (303), or is otherwise in need of remediation.

[00114] In another example, incidents may be identified from a number of ticket support systems. For example, an information technology (IT) service management system (ITSM) (316-1) may also be a source of incidents. An ITSM system (316-1) implements and manages the quality of IT services that meet the needs of the user. In one example, the ITSM system (316-1) is managed by the user, a service provider, a third party, or combinations thereof, in which a service ticket is opened by one of these groups or individuals. In another example, the ITSM system (316-1) may automatically enter a service ticket based on the events detected by the monitoring system. If the ITSM system (316-1) determines that the instantiated system (312) or a number of nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) thereof are not appropriately provisioned, are wrongly provisioned, or are otherwise unfit for the instantiated system (312), the ITSM system (316-1) may, like the event handler (316), provide a remediation determination in the form of an incident sent to the remediation engine (317).

[00115] The incidents generated by the event handler (316) and the ITSM system (316-1) may be brought to the attention of a user, administrator, third party, or other user of the topology-based management broker (200) or instantiated service (312) in the form of a notification. As described above, the remediation policies define how a remediation action is to be performed, and may include: (1) providing notifications to a user, consumer, or administrator; (2)

obtaining instructions from the user, consumer, or administrator; (3) taking manual actions input by the user, consumer, or administrator; (4) taking autonomous actions after receiving instructions from the user, consumer, or administrator; (5) taking autonomous actions without receiving instructions from the user, consumer, or administrator; (6) taking autonomous actions without notifying the user, consumer, or administrator or receiving instructions from the user, consumer, or administrator; or combinations thereof. In this manner, the issuance of notifications within the system is defined by the remediation policies.

[00116] In some examples, once events are generated by the event handler (316) a remediation engine (317) may execute, via a processor, logic to correct the incidents reported by the event handler (316) and/or ITSM system (316-1). Remedies issued by the remediation engine (317) may include, for example, allocation of additional computing resources, allocation of different computing resources, and reallocation of computing resources from one geographical area to another, among many other remediation actions. In one example, the remediation actions taken by the remediation engine (317) are implemented to remedy a misallocation of computing resources that does not comply with the policies associated with the topology (302) designed. In another example, the remediation actions taken by the remediation engine (317) are implemented to remedy a failure of a number of computing resources within the instantiated service (312). In still another example, the remediation actions taken by the remediation engine (317) are implemented to adjust the security levels of the instantiated service (312) and the groups and individual computing resources therein. Any number of other remediation actions may be implemented by the remediation engine (317) for any number of reasons.

[00117] In one example, the remediation actions taken by the remediation engine (317) are implemented with or without notification to a user, administrator, third party, or other user as described above. Further, in another example, the remediation actions taken by the remediation engine (317) are implemented autonomously, without user interaction or confirmation from a user.

[00118] In still another example, the remediation actions taken by the remediation engine (317) are implemented with user interaction from the consumer, administrator, third party, or other user. In this example, the remediation engine (317) sends data to a self-service subscription management engine (318). The self-service subscription management engine (318) executes, via a processor, logic to present information to a user regarding the events detected by the monitoring system (313) and the incidents generated by the event handler (316) and ITSM system (316-1). The self-service subscription management engine (318) also executes, via a processor, logic to present to a user a number of recommendations for remediation of the events and incidents.

[00119] As depicted in Fig. 4, the topology-based management broker (200) further comprises a subsystem capable of independently modeling infrastructure and application requirements of a topology on the same stack as the subsystem depicted in Fig. 3. However, as described above, the present systems and associated methods also support all the use cases that a CDA supports such as those CDA 1.2 supports. As described above, CDA is a number of software tools utilized within a topology designer that independently model infrastructure and application requirements while managing versions, configurations, and other application components. CDA 1.2 is also developed and distributed by Hewlett Packard Corporation.

[00120] The subsystem of the topology-based management broker (200) depicted in Fig. 4 may be used to design a topology for a number of applications to be executed on the instantiated service (312). The subsystem of Fig. 4 assists in the provisioning, deploying, and maintaining of a topology that supports the applications, and provides application models that match appropriate infrastructure templates. In one example, the models of the applications executed on the deployed topology utilize designed topologies that are easily reconciled with the templates defining the infrastructure topologies of the topology.

[00121] A topology designer (301) may be used to design and create an application model (319). The application model (319) is defined by a

lifecycle management topology. As described above in connection with the LCM topology (302), the application model (319) comprises a number of nodes (319-1, 319-2, 319-3). A number of policies and lifecycle management actions (LCMA) are associated with each of the nodes (319-1, 319-2, 319-3) of the application model (319).

[00122] A topology designer (301) may also be used to create a number of infrastructure and/or platform templates (320). The templates (320) are defined by a lifecycle management topology. As described above in connection with the LCM topology (302), the templates (320) comprise a number of nodes (320-1, 320-2, 320-3, 320-4, 320-5). A number of policies and lifecycle management actions (LCMA) are also associated with each of the nodes (320-1, 320-2, 320-3, 320-4, 320-5) of the templates (320).

[00123] In one example, the topology designers (301), self-service portal (309), and resource offering manager (308), alone or in combination, may associate a number of policies (303) and LCMAs (304) with the nodes (319-1, 319-2, 319-3, 320-1, 320-2, 320-3, 320-4, 320-5) of the application model (319) and infrastructure template (320). In another example, a separate policy engine and LCMA engine may be provided to associate the nodes (319-1, 319-2, 319-3, 320-1, 320-2, 320-3, 320-4, 320-5) of the application model (319) and infrastructure template (320) with the policies and LCMAs as described above.

[00124] As depicted in Fig. 4, a number of models (319) may be presented as possible matches or near matches for a number of infrastructure templates (320). In one example, rather than using a topology designer (301), a number of application models (319) resources may be provided within the topology-based management broker (200). In this example, the topology-based management broker (200) may obtain application models (319) from, for example, the catalog (310), the RTSM (315), another model source, or combinations thereof. A user may browse through these model sources and obtain a number of application models (319) that may be reconciled with the infrastructure templates (320). In this manner, the topology designer (301) may design a number of application models (319) or a number of application models (319) may be obtained from the above-described resource. Thus, the

application models (319) may be application topologies designed by the topology designer (301), or realized application topologies as described above.

[00125] Similarly, as depicted in Fig. 4, a number of templates (320) are presented as possible matches or near matches for the application model (319). In one example, rather than using a topology designer (301), a number of template (320) resources may be provided within the topology-based management broker (200). In this example, the topology-based management broker (200) may obtain templates (320) from, for example, the catalog (310), the RTSM (315), another template source, or combinations thereof. A user may browse through these template sources and obtain a number of templates (320) that may be reconciled with the application model (319). In this manner, the topology designer (301) may design a number of templates (320) or a number of templates may be obtained from the above-described resource. Thus, the templates (320) may be infrastructure topologies designed by the topology designer (301), or realized infrastructure topologies as described above.

[00126] The CDA subsystem described in Fig. 4 comprises a stitching engine (321) to stitch or combine the application model (319) to the infrastructure template (320). The stitching engine (321) may use any type of method to stitch the application model (319) to the infrastructure template (320) based on the policies and LCMA associated with the application model (319) to the infrastructure template (320). In one example, the stitching engine (321) may use a matching process in which the stitching engine (321) matches the policies, requirements, and capabilities associated with the nodes (319-1, 319-2, 319-3) of a number of application models (319) with the policies, requirements, and capabilities of the nodes (320-1, 320-2, 320-3, 320-4, 320-5) of a number of infrastructure templates (320). In this example, the stitching engine (321) may browse through the template sources described above to find a match or near match. Once a match is found, the stitching engine (321) matches a number of nodes (319-1, 319-2, 319-3) of the application model (319) with a number of the nodes (320-1, 320-2, 320-3, 320-4, 320-5) of the matching infrastructure template (320).

[00127] Another method the stitching engine (321) may use to stitch the application model (319) to the infrastructure template (320) may comprise an algorithmic matching method. In this method, the stitching engine (321) determines a match mathematically via algorithms that employ the policies in performing the matching decisions. In one example, this may include inference methods in which requirements in the application level are tagged or otherwise associated with components that support them in a library of infrastructure topologies called a DSL database (323), wherein the overall infrastructure template (320) is aggregated first before the aggregation is extended to the application model (319).

[00128] A definitive software library (DSL) is a secure storage device, consisting of physical media or a software repository located on a network file server. Definitive authorized versions of all software configuration items (CIs) or artifacts that may be required to deploy the application designed in the application model (319) may be stored and protected in a DSL. In the present example, a number of infrastructure topologies (320) are stored in the DSL. Thus, the DSL contains master copies of a number of infrastructure topologies (320) developed using the present systems and methods or purchased from a third party. All related documentation related to the infrastructure topologies (320) is also stored in the DSL. The DSL database (323) of the present topology-based management broker (200) comprises a number of objects used in the deployment of the application after the application model (319) has been developed and is ready for deployment on the infrastructure template (320). In one example, a topology designer (301) may also provide additional design elements within the topology before, during, and/or after the stitching engine (321) processes the application model (319) and the infrastructure template (320) to create the topology (302) with a number of nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7).

[00129] Once the stitching engine (321) has completed the stitching process as described above, a complete topology (302) is created. The topology created by the subsystem of Fig. 4 may have additional policies and LCMAAs associated with the nodes as described above in connection with Fig. 3.

The topology (302) created via the subsystem of Fig. 4 may be stored in the catalog (310), the DSL database, or other storage device or system. The topology (302) created via the subsystem of Fig. 4 may be processed in a similar manner as described above in connection with the topology (302) developed in Fig. 1A. The LCM engine (311) obtains the artifacts required to deploy the application designed in the application model (319) from the DSL (323) and executes the topology (302).

[00130] In one example, an application lifecycle management (ALM) device (322) depicted in Fig. 3 is used to trigger the deployment of the topology developed on the subsystem depicted in Fig. 4 of the overall topology-based management broker (200). In one example, Hewlett Packard's Application Lifecycle Management (HP ALM) is used. HP ALM is a unified software platform developed and distributed by Hewlett Packard Company. HP ALM assists in accelerating the delivery of secure, reliable modern applications in a network.

[00131] Fig. 6 is a flowchart showing a method for brokering a cloud service, according to one example of the principles described herein. The method of Fig. 6 includes generating (block 601) a topology (Fig. 3 and 4, 102). As described above, in one example, a number of topology designers (Fig. 3, 301) including a number of topology design tools, GUIs, and coding scripts, may be used by a human designer to design the topology (Fig. 3 and 4, 302). The topology (Fig. 3 and 4, 302) may be designed using either or both of the subsystems depicted in Figs. 3 and 4. Further, in one example, topologies (Fig. 3 and 4, 302) designed and stored may be browsed or search for in a database of topologies (Fig. 3 and 4, 302) and used as a portion of the topology (Fig. 3 and 4, 302) to be instantiated.

[00132] In one example, topologies (302) may be generated by designing a topology (302) de novo via a number of topology designers (301). In another example, the topology may be generated (block 601) by stitching a number of applications models (Fig. 4, 319) and a number infrastructure templates (Fig. 4, 320) together using a number of stitching methods. As will be described in more detail below, the stitching engine (Fig. 4, 321) may obtain a

number of infrastructure topologies (Fig. 4, 320), and stitch (Fig. 9, block 903) a number of application models (Fig. 4, 319) to a number of appropriate infrastructure templates (Fig. 4, 320). In another example, the application models (Fig. 4, 319) and infrastructure templates (Fig. 4, 320) may be designed de novo by a number of topology designers (301). In one example, a number of persons may use the topology designers (301) to design the topologies (302) in accordance with the method of Fig. 6. These individuals may be service designers, infrastructure architects or administrators, system administrators, information technology operators, offer managers, or users, among other personnel with roles in the design of a topology. In still another example, the topology designers (301) may be operated by a third party.

[00133] The method may continue by associating (block 602) a number of LCMAs (304) with a number of nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) within the topology (302). In one example, block 602 may be performed with the resource offering manager (Fig. 3, 308). The LCMAs orchestrate a number of application programming interfaces (APIs) of a number of resources for purposes of managing the lifecycle of a given cloud service capability. In one example, the LCMAs are uniform resource identifiers (URIs) of application programming interfaces (APIs) that perform calls in order to execute the APIs.

[00134] In one example, policies (Fig. 3, 303) may also be associated with a number of nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) within the topology (302). In one example, association of policies (Fig. 3, 303) with a number of nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) within the topology (302) may be performed with the resource offering manager (Fig. 3, 308). A policy is any data or metadata used to assist in the management of the provisioning, deploying, monitoring, enforcement, and remediation within a cloud service. The policies may represent a number of rules or sets of rules that are applicable to the provisioning, deploying, monitoring, enforcement, and remediation tasks associated with a number of computing devices within a cloud service environment.

[00135] The topology (302) may be executed (block 603). In one example, the topology (302) is executed (block 603) based on the LCMAAs (304) associated (block 602) with a number of nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) within the topology (302). Further, in another example, the topology (302) is executed (block 603) based on the policies (303) associated with a number of nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) within the topology (302).

[00136] In still another example, a number of scripts may be created for execution (block 603) of the topology (302). The scripts define executable logic for instantiating a cloud service based on the topology (Fig. 3 and 4, 302) and policies (Fig. 3, 303). The method of Fig. 6 will be described in more detail in connection with Fig. 7.

[00137] Fig. 7 is a flowchart showing a method for brokering a cloud service, according to another example of the principles described herein. The method of Fig. 7 may begin by generating (block 701) a topology. As described above, in one example, a number of topology designers (Fig. 3, 301) including a number of topology design tools, GUIs, and coding scripts, may be used by a human designer to design the topology (Fig. 3 and 4, 302). The topology (Fig. 3 and 4, 302) may be designed using either or both of the subsystems depicted in Figs. 3 and 4. Further, in one example, topologies (Fig. 3 and 4, 302) designed and stored may be browsed or search for in a database of topologies (Fig. 3 and 4, 302) and used as a portion of the topology (Fig. 3 and 4, 302) to be instantiated.

[00138] In one example, topologies (302) may be generated by designing a topology (302) de novo via a number of topology designers (301). In another example, the topology may be generated (block 601) by stitching a number of applications models (Fig. 4, 319) and a number infrastructure templates (Fig. 4, 320) together using a number of stitching methods. As will be described in more detail below, the stitching engine (Fig. 4, 321) may obtain a number of infrastructure topologies (Fig. 4, 320), and stitch (block 903) a number of application models (Fig. 4, 319) to a number of appropriate infrastructure templates (Fig. 4, 320). In another example, the application

models (Fig. 4, 319) and infrastructure templates (Fig. 4, 320) may be designed de novo by a number of topology designers (301).

[00139] In one example, a number of persons may use the topology designers (301) to design the topologies (302) in accordance with the method of Fig. 6. These individuals may be service designers, infrastructure architects or administrators, system administrators, information technology operators, offer managers, or users, among other personnel with roles in the design of a topology. In still another example, the topology designers (301) may be operated by a third party.

[00140] The method may continue by associating (block 702) a number of policies (Fig. 3, 303) with a number of nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) within the topology (302). In one example, block 702 may be performed with the resource offering manager (Fig. 3, 308). A policy is any data or metadata used to assist in the management of the provisioning, deploying, monitoring, enforcement, and remediation within a cloud service. The policies may represent a number of rules or sets of rules that are applicable to the provisioning, deploying, monitoring, enforcement, and remediation tasks associated with a number of computing devices within a cloud service environment.

[00141] At block 703, a number of lifecycle management actions (LCMAs) (Fig. 3, 304) may be applied to a number of nodes within the topology. The LCMAs orchestrate a number of application programming interfaces (APIs) of a number of resources for purposes of managing the lifecycle of a given cloud service capability. In one example, the LCMAs are uniform resource identifiers (URIs) of application programming interfaces (APIs) that perform calls in order to execute the APIs.

[00142] In one example, the policies (Fig. 3, 303) and LCMAs (Fig. 3, 304) may be associated with the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) within the topology (302) via data or metadata describing the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) of the topology (Fig. 3, 302). The data or metadata may be provided in a number of files describing the nodes or topology, or in a file associated therewith. In another example, the

LCMAs are associated with the aspects of the topology by default by virtue of what computing device the node or nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) represent.

[00143] In another example, the LCMAs are associated with the aspects of the topology by explicitly providing a number of functions, F_{Action} , that define how to select a resource provider to implement the action based on the policies associated with the aspects of the topology and the policies of the different relevant resource providers. These functions define how a resource provider is selected to implement the action based on the policies associated with the aspect of the topology and the policies of the different relevant resource providers. In one example, the processes of blocks 702 and 703 may be performed in any order serially, or in parallel. Further, in one example, a number of persons may use the topology designers (301) to design the topologies (302) in accordance with the method of Fig. 6. These individuals may be service designers, infrastructure architects or administrators, system administrators, information technology operators, offer managers, or users, among other personnel with roles in the design of a topology. In still another example, the topology designers (301) may be operated by a third party.

[00144] A number of resource provider policies (308-1) may be associated (block 704) with a number of nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) within the topology (302). Resource provider policies (308-1) are any policies associated with a number of resource providers' offerings that guide the selection of a number of resources. In one example, the resource provider policies (308-1) may be dynamic functions that define the computing abilities of a computing resource. In this example, a computing resource that provides a defined level of computing resources such as, for example, processing power, may be provisioned by the LCM engine (311) and resource offering manager (308) if the defined level of that computing resource meets a number of requirements within the topology (302).

[00145] As described in connection with Fig. 5, after evaluating the provisioning policies and selecting the resource provider, an interpreter (503) may evaluate the other policies (e.g., monitoring policies, event handling

policies, event processing policies, incident handling policies, among other policies). The interpreter (Fig. 5, 503) may also convert policies such that they may be implemented. For example, the interpreter (Fig. 5, 503) may convert policies into topologies associated with the topology (Fig. 3, 302), nodes of the topology (Fig. 3, 302) or a subset of nodes of a topology (Fig. 3, 302). The interpreter (Fig. 5, 503) may then associate, delegate or configure resource providers, workflows, scripts or combinations thereof or provide actual instructions to implement a monitoring system. In some examples, if a monitoring system is already described, the monitoring and event handling topologies may be an input (Fig. 5, 501) to a provision policy engine (Fig. 5, 502) as described above.

[00146] The topology (302) may be executed (block 705). In one example, the topology (302) is executed (block 705) based on the policies, (303), LCMA's (304), resource provider policies (308-1), or combinations thereof. In one example, a number of scripts may be created for execution (block 705). The scripts define executable logic for instantiating a cloud service based on the topology (Fig. 3 and 4, 302), policies (Fig. 3, 303), LCMA's (Fig. 3, 304), resource provider policies (308-1), or combinations thereof.

[00147] A topology LCM engine (Fig. 3, 311) instantiates (block 706) the topology (Fig. 3 and 4, 302). In one example, instantiation (block 706) of the topology (302) is based on the policies (Fig. 3, 303), LCMA's (Fig. 3, 304) resource provider policies (308-1), executable scripts, or combinations thereof. In one example, the topology LCM engine (Fig. 3, 311) obtains the workflows or sequences of serial and/or parallel scripts created at block 705 during execution, calls a resource provider via the resource offering manager (Fig. 3, 308), and instantiates the topology (Fig. 3 and 4, 302) based on the policies (Fig. 3, 303), LCMA's (Fig. 3, 304) resource provider policies (308-1), and executable scripts to create an instantiated service (Fig. 3, 312).

[00148] A number of realized topologies (Fig. 3, 314) may be derived (block 707) from the instantiated service (Fig. 3, 312). In one example, the topology LCM engine (Fig. 3, 311) derives a realized topology (Fig. 3, 314) from each instantiated service (Fig. 3, 312). A number of the realized

topologies (Fig. 3, 314) may be stored (block 708) in a database of realized topologies. In one example, the LCM engine (Fig. 3, 311) stores the realized topologies (Fig. 3, 314) in the realized topology system management (RTSM) database (Fig. 3, 315); a logical system repository of realized topologies (Fig. 3, 314). In one example, the RTSM database (315) comprises a database management system (DBMS). The DBMS is a combination of hardware devices and software modules that interact with a user, other applications, and the database itself to capture and analyze data.

[00149] In one example, the RTSM database (Fig. 3, 315) is a configuration management database (CMDB); a repository of information related to all the components of a realize topology (Fig. 3, 314). The realized topology (Fig. 3, 314) comprises a model of the topology (Fig. 3, 302), with the policies applied to the various nodes (Fig. 3, 302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7). A number of properties of the nodes (Fig. 3, 302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) of the realized topology (Fig. 3, 314) are defined within the realized topology (Fig. 3, 314). These properties include any details of any instantiated topology (Fig. 3, 312) that is created or updated via the topology-based management broker (Fig. 3, 200), and may include, for example, the internet protocol (IP) address of the nodes, and characteristics and computing parameters of the nodes, among many other properties.

[00150] The RTSM (Fig. 3, 315) is a repository that stores each instance of a realized topology (Fig. 3, 314). In this manner, every time a topology (Fig. 3, 302) is designed, provisioned, and deployed, the topology-based management broker (Fig. 3, 200) captures the realized topology (Fig. 3, 314) of that instantiated topology (312). Thus, the RTSM (Fig. 3, 315) contains a realized topology (Fig. 3, 314) of every topology (Fig. 3, 302) that has been instantiated within the topology-based management broker (Fig. 3, 200). In one example, in every instance of the modification of an existing instantiated topology (312), the realized topology (Fig. 3, 314) resulting from that modification is also stored within the RTSM (Fig. 3, 315).

[00151] Fig. 8 is a flowchart showing a method for remediating a number of incidents within a cloud service, according to one example of the

principles described herein. The remediation method of Fig. 8 may be performed alone, or in combination with any number of additional process described herein such as those process described in Figs. 6, 7, and 8. Further, any block within the method of Fig. 8 may be performed alone or in combination with any number of other processes within Fig. 8. For example, a monitoring process described at block 801 may be performed alone without the remaining processes being performed, or less than all of the remaining processes being performed.

[00152] The remediation method of Fig. 8 may include monitoring (block 801) an instantiated topology (Fig. 3, 312) for a number of events. The monitoring system (313) monitors (block 801) an instantiated topology (Fig. 3, 312) based on the monitoring policies associated with the topology (302) and the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) of the topology (302) as described above. In one example, the monitoring system, based on the policies, monitors for a number or set of metrics. A number of events may be derived from the detected metrics,

[00153] The monitoring system (313) sends data representing a number of the events to the event handler (313) based on a number of the policies associated with the designed topology (302) and the instantiated service (312). For example, as described above, the monitoring policies include a portion that defines what to do with the monitored events that result from the monitoring such as, for example, how to handle the events, where the events are sent, what devices or individuals address the events, how incidents resulting from the processing of the events are handled, how the events and incidents are processed (e.g., processed as aggregated, filtered, or correlated events, among other forms of processing), and how the resulting incidents are handled.

[00154] A number of events detected by the monitoring system (313) may be processed by the event handler (316) based on a number of the policies described above. Handling (block 802) of events may include, for example, processing the events as aggregated, filtered, or correlated events, among other forms of processing. Further, based on the above-described

policies, the event handler (313) may handle (block 802) the events by determining whether the events should be processed into incidents, or whether to notify a number of users of the system (200), for example.

[00155] A number of incidents are generated (block 802). In one example, the incidents are created by the event handler (Fig. 3, 316) based on a number of policies including, for example, monitoring and remediation policies. Further, in one example, the incidents are generated (block 803) by the event handler (Fig. 3, 316) based on the events detected by the monitoring system (313). In another example, the incidents are generated (block 803) by obtaining a number of service tickets from an information technology (IT) service management system (ITSM), and, with the event handler, creating a number of incidents based on the service tickets. As described above, an ITSM (316-1) may also be a source of incidents. An ITSM system (316-1) implements and manages the quality of IT services that meet the needs of the user. In one example, the ITSM system (316-1) is managed by the user, a service provider, a third party, or combinations thereof, in which a service ticket is opened by one of these groups or individuals. In another example, the ITSM system (316-1) may automatically enter a service ticket based on the events detected by the monitoring system. If the ITSM system (316-1) determines that the instantiated system (312) or a number of nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) thereof are not appropriately provisioned, are wrongly provisioned, or are otherwise unfit for the instantiated system (312), the ITSM system (316-1) may, like the event handler (316), provide a remediation determination in the form of an incident sent to the remediation engine (317).

[00156] The incidents generated by the event handler (316) and the ITSM system (316-1) may be brought to the attention of a user, administrator, third party, or other user of the topology-based management broker (200) or instantiated service (312) in the form of a notification. A number of notifications are sent (block 804) regarding the incidents created by the event handler (313). These notifications may be sent (block 804) to a number of devices and users within the system (200). For example, a number of notifications may be sent to the self-service subscription management engine (318). The self-service

subscription management engine (318) may present the notifications to a user via, for example, the GUI (318-1) associated with the self-service subscription management engine (318). Thus, a number of notifications are presented (block 804) to a user regarding the incidents.

[00157] In one example, the process defined by block 804 is optional. As described above, the event handler (Fig. 3, 316) may or may not provide notifications to a user based on a number of policies associated with the instantiated topology (312). When the event handler (Fig. 3, 316) does dispatch notifications to a user, a varying level of user interaction may be allowed or requested including allowing a user to interact with, for example, a number of the GUIs (318-1) produced by the self-service subscription management engine (318) before a number of remediation actions are taken. As described above, remediation policies define whether a notification is to take place, how that notification is handled, and at what degree user input is allowed or requested. Thus, the remediation policies may include: (1) providing notifications to a user, consumer, or administrator; (2) obtaining instructions from the user, consumer, or administrator; (3) taking manual actions input by the user, consumer, or administrator; (4) taking autonomous actions after receiving instructions from the user, consumer, or administrator; (5) taking autonomous actions without receiving instructions from the user, consumer, or administrator; (6) taking autonomous actions without notifying the user, consumer, or administrator or receiving instructions from the user, consumer, or administrator; (7) proposing a remediation action to a user or administrator for approval, and performing the proposed remediation action if approved by the user or administrator, or combinations thereof.

[00158] At block 805, a number of function calls are generated. The function calls issued to the LCM engine (311) by the remediation engine (317) to remediate the incidents may be based on a number of LCMAAs associated with the elements of the instantiated topology (312), the incidents to be remediated, and the policies associated with the elements of the topology (302). In this manner, the remediation engine (317) executes, via a processor,

logic to correct the incidents reported by the event handler (316) and/or ITSM system (316-1) in order to generate (block 805) the function calls.

[00159] Using the function calls generated by the remediation engine (317), the topology LCM engine (Fig. 3, 311) modifies (block 806) an instantiated topology (Fig. 3, 312) based on the subsequent LCMA's created by the remediation engine (317). Modification of an instantiated topology (Fig. 3, 312) may include modifying the topology (312) or a portion thereof, modifying the a number of nodes or a group of nodes, addition of a number of nodes, groups of nodes, or topologies, deletion of a number of nodes, groups of nodes, or topologies, among many other types of changes that may be made to an instantiated service (312). Further, modification of the instantiated topology (312) may include re-instantiation of a previously instantiated topology (312).

[00160] A subsequent realized topology (Fig. 3, 314) may be derived (block 807) from the modified topology (Fig. 3, 312), and stored (block 808) in a database of realized topologies. In one example, the LCM engine (Fig. 3, 311) stores the realized topologies (Fig. 3, 314) in the realized topology system management (RTSM) database (Fig. 3, 315).

[00161] A determination (block 809) may be made as to whether monitoring of an instantiated topology (Fig. 3, 312) is to end. Reasons to end the monitoring of an instantiated topology (Fig. 3, 312) may include, for example, completion of a contract such as an SLA, ending of the cloud services provided by a number of service providers. If it is determined that monitoring of the instantiated topology (Fig. 3, 312) is to end (block 809, determination YES), then the process terminates. If, however, it is determined that monitoring of the instantiated topology (Fig. 3, 312) is not to end (block 809, determination NO), then the process loops back to block 801, and the process of remediation is repeated. In one example, the remediation process may be performed any number of iterations throughout the lifecycle of an originally instantiated topology (Fig. 3, 312). In this manner, events that may occur within the instantiated topology (Fig. 3, 312) may be addressed in order to maintain a working instantiated topology (Fig. 3, 312). Further, the remediation process

described in Fig. 8 allows for the instantiated topology (Fig. 3, 312) to be amended or adjusted to provide a scalable instantiated topology (Fig. 3, 312).

[00162] A number of the steps depicted in Fig. 8 may or may not or may not take place. Some of the steps may be performed without being prescribed by the topologies. In other words, the steps may be prescribed using other logic. The method (800) of Fig. 8 is an example model of a case where remediation is performed. In some examples, the events may not be modeled at the level of the topology.

[00163] Fig. 9 is a flowchart showing a method of designing a topology, according to one example of the principles described herein. The method of Fig. 9 may begin by generating (block 901) an application model (Fig. 4, 319). In one example, a topology designer (301) may be used to design and create the application model (Fig. 4, 319), and, in this manner, generate (701) an application model (Fig. 4, 319). In another example, the application model (Fig. 4, 319) may be obtained from a number of application model (Fig. 4, 319) sources such as, for example, the catalog (Fig. 1A, 310), the RTSM (Fig. 1A, 315), or the DSL database (Fig. 4, 323), among other application model (Fig. 4, 319) sources. The application model (Fig. 4, 319) is defined by a lifecycle management topology. As described above in connection with the LCM topology (Fig. 3, 302), the application model (Fig. 4, 319) comprises a number of nodes (Fig. 4, 319-1, 319-2, 319-3).

[00164] A number of infrastructure templates (Fig. 4, 320) may also be generated (block 902). In one example, a topology designer (301) may be used to design and create the infrastructure template (Fig. 4, 320). In another example, the infrastructure template (Fig. 4, 320) may be obtained from a number of infrastructure template (Fig. 4, 320) sources such as, for example, the catalog (Fig. 1A, 310), the RTSM (Fig. 1A, 315), or the DSL database (Fig. 4, 323), among other infrastructure template (Fig. 4, 320) sources. The infrastructure template (Fig. 4, 320) is defined by a lifecycle management topology. As described above in connection with the LCM topology (Fig. 3, 302), the infrastructure template (Fig. 4, 320) comprises a number of nodes (Fig. 4, 319-1, 319-2, 319-3). In one example, a number of persons may use

the topology designers (301) to design the application models (Fig. 4, 319) and infrastructure templates (Fig. 4, 320). These individuals may be service designers, infrastructure architects or administrators, system administrators, information technology operators, offer managers, or users, among other personnel with roles in the design of a topology.

[00165] A number of application models (Fig. 4, 319) are stitched (block 903) to a number of infrastructure templates (Fig. 4, 320). In one example, the stitching engine (Fig. 4, 321) may obtain a number of infrastructure topologies (Fig. 4, 320) stored in, for example, the DSL database (Fig. 4, 323) or other source of infrastructure templates (320), and stitch (block 902) a number of application models (Fig. 4, 319) to a number of appropriate infrastructure templates (Fig. 4, 320). In another example, the infrastructure templates (Fig. 4, 320) may be designed de novo by a number of topology designers (301).

[00166] The stitching engine (Fig. 4, 321) may use any type of method to stitch the application models (Fig. 4, 319) to the infrastructure templates (Fig. 4, 320) based on the policies and LCMA associated with the application models (Fig. 4, 319) to the infrastructure templates (Fig. 4, 320). In one example, the stitching engine (Fig. 4, 321) may use a matching process in which the stitching engine (Fig. 4, 321) matches the policies, requirements, and capabilities associated with the nodes (Fig. 4, 319-1, 319-2, 319-3) of the application models (Fig. 4, 319) with the policies, requirements, and capabilities of the nodes (Fig. 4, 320-1, 320-2, 320-3, 320-4, 320-5) of the infrastructure templates (Fig. 4, 320). In this example, the stitching engine (Fig. 4, 321) may browse through the template sources described above to find a match or near match. Once a match is found, the stitching engine (Fig. 4, 321) matches a number of nodes (Fig. 4, 319-1, 319-2, 319-3) of the application models (319) with a number of the nodes (Fig. 4, 320-1, 320-2, 320-3, 320-4, 320-5) of the matching infrastructure templates (Fig. 4, 320).

[00167] Another method the stitching engine (Fig. 4, 321) may use to stitch the application models (Fig. 4, 319) to the infrastructure templates (Fig. 4, 320) may comprise an algorithmic matching method. In this method, the

stitching engine (Fig. 4, 321) determines mathematically via algorithms that employ the policies in performing the matching decisions. In one example, this may include inference methods in which requirements in the application level are tagged or otherwise associated with components that support them in the DSL database (Fig. 4, 323), wherein the overall infrastructure topology (Fig. 4, 320) is aggregated first before the aggregation is extended to the application models (Fig. 4, 319).

[00168] A number of policies and lifecycle management actions (LCMAs) are associated (blocks 704 and 705) with each of the nodes (Fig. 4, 319-1, 319-2, 319-3) of the application model (Fig. 4, 319) and nodes of the infrastructure topology (Fig. 4, 320). In one example, the association (blocks 704 and 705) of the number of policies (303) and LCMAs (304) with the nodes (319-1, 319-2, 319-3, 320-1, 320-2, 320-3, 320-4, 320-5) of the application model (319) and infrastructure topology (320) may be performed by the topology designers (301), self-service portal (309), and resource offering manager (308), alone or in combination. In another example, a separate policy engine and LCMA engine may be provided to associate the nodes (319-1, 319-2, 319-3, 320-1, 320-2, 320-3, 320-4, 320-5) of the application model (319) and infrastructure topology (320) with the policies (303) and LCMAs (304) as described above.

[00169] In one example, the processes of blocks 704 and 705 of associating policies (303) and lifecycle management actions (LCMAs) (304) with each of the nodes (Fig. 4, 319-1, 319-2, 319-3) of the application model (319) and nodes of the infrastructure topology (Fig. 4, 320) may be performed before, during, or after the stitching process described in connection with block 903. In one example where policies and LCMAs are associated before the stitching process of block 902, the policies (303) and LCMAs (304) may be associated with a number of nodes or groups of nodes within the application model (319) and infrastructure topology (320), as well as with the application model (319) as a whole and infrastructure topology (320) as a whole. In this example, additional policies (303) and LCMAs (304) may be associated with the topology (302) created via the stitching process of block 902. In another

example, the processes of blocks 704 and 705 of associating policies (303) and lifecycle management actions (LCMAs) (304) with each of the nodes (Fig. 4, 319-1, 319-2, 319-3) of the application model (319) and nodes of the infrastructure topology (Fig. 4, 320) may be optional as to performance of these processes after the stitching process of block 902. In still another example, the processes of blocks 704 and 705 of associating policies (303) and lifecycle management actions (LCMAs) (304) with each of the nodes (Fig. 4, 319-1, 319-2, 319-3) of the application model (319) and nodes of the infrastructure topology (Fig. 4, 320) may be performed before and after stitching process of block 902.

[00170] The above processes described in Fig. 9 results in a completely designed topology (302) similar to the topology (302) described above in connection with Fig. 3. Thus, the method described in Fig. 9 may be further associated with the process described herein regarding Figs. 6, 7, and 8. For example, the topology (Fig. 4, 302) resulting from the method of Fig. 9 may be used as the input topology (Fig. 3, 302) for the method described in connection with Figs. 6 and 7 at, for example, blocks 601 and 701. Further, in another example, the topology (Fig. 4, 302) resulting from the method of Fig. 9 may be used as the input topology (Fig. 3, 302) for instantiation in the remediation method described in connection with Fig. 8. Further still, in one example, a number of persons participate in the method described in Fig. 9. These individuals may be service designers, infrastructure architects or administrators, system administrators, information technology operators, offer managers, or users, among other personnel with roles in the design, execution, monitoring, and remediation of a topology (302).

[00171] Fig. 10 depicts a block diagram showing a functional overview of a monitoring engine (313) for monitoring a cloud service modeled as a topology, according to one example of the principles described herein. As may happen within the topology-based management broker (Fig. 3, 300), a number of events may occur within the topology-based management broker (Fig. 3, 300). These events may include, for example, a policy failure within a node of the instantiated service (312), a failure of a number of hardware or

software components within the instantiated service (312), and an unauthorized access of the instantiated service (312), among many other computing-related events. A monitoring engine (313) may monitor the operation of the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) and detect these events. In other words, the monitoring engine (313) may operate based on the monitoring policies associated with the topology (302) and the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) of the instantiated service (312). In this example, the monitoring engine (313) is used to monitor the operations, the security, the compliance, the usage, and the analytics of the topology (302) and the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) of the topology (302), among other items to monitor within the instantiated service (312).

[00172] To establish these monitoring policies, the monitoring engine (313) may include an interface. The interface may prompt a user to select a number of monitoring policy elements to incorporate into a monitoring policy. Monitoring policies are policies that, if implemented, define the functionality of the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) and may include criteria against which actual node (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) functionality may be compared. Monitoring policies may include various types of monitoring functionality. For example, the monitoring policies may include operational monitoring, security monitoring, compliance monitoring, analytics monitoring, usage monitoring, or combinations thereof. Accordingly, the monitoring engine (313) may detect events relating to the results of the different monitoring parameters. For example, the monitoring engine (313) may detect an operational event, a security event, a compliance event, an analytics event, a usage event, or combinations thereof. Specific examples of the different monitoring policy elements are described as follows.

[00173] Examples of monitoring policies regarding node operations include performance, monitoring CPU levels and loads of the various nodes within the cloud, monitoring the speed at which data is processed through a node or a number of nodes or exchanged between nodes, and monitoring the operational state of applications running on a node or nodes at any level of the

network, among many other operations parameters of the nodes, group of nodes, and the cloud service as a whole.

[00174] Monitoring policies also include monitoring policies regarding security; monitoring for abnormal behaviors or behaviors known as being associated with known or suspected problems. Examples of monitoring policies regarding security include monitoring whether a node or a group of nodes is experiencing an attack, whether there is strange behavior occurring within the cloud or interactions with the cloud service, and whether there is a virus or other anomaly with a node or group of nodes, among other security-related monitoring policies.

[00175] Examples of monitoring policies regarding compliance include, determinations as to whether the nodes or group of nodes are running an appropriate version of an operating system, determining whether the most recent patch associated with the release of a software program running on the nodes has been installed, determining if an installed patch is a correct patch, checking that a code or artifacts that have been used to provision the node and cloud service have been appropriately checked and vetted for any weakness or problem, and if governance and access control to the node and cloud service or the node and cloud service management is appropriate, among other compliance-related monitoring policies.

[00176] Examples of monitoring policies regarding usage include, determining how much a consumer has been using CPUs of a node or group of nodes, determining how much memory a consumer has utilized, determining how much money has been charged to the consumer, and determining whether the consumer has paid for the services provide through the designing, provisioning, deploying, and monitoring of the network topology, among other usage-related monitoring policies.

[00177] Examples of monitoring policies regarding analytics include, defining what is needed to ensure analytics and big data monitoring within or among the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) or groups of nodes within the topology (302) is occurring as expected. For example, the analytics policies may define a number of workflows by which the monitoring

engine (313) may operate to configure the cloud service, provide analytics, collect big data, and process the data.

[00178] In addition to those policies described above, the monitoring policies may include other metrics. Other examples of metrics that may be used in monitoring policies to indicate functionality of nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) include log monitoring, performance monitoring, and intelligence monitoring such as, for example, collection of metrics, business intelligence (BI) and business activity monitoring (BAM) and analytics/big data integration, among other types monitoring-related policies.

[00179] In some examples, the monitoring policies (303) can provision or setup a big data system that collects data about a broker that deploys the topology, the associated systems (e.g. the monitoring systems, remediation system, events, among other associated systems), data of the instantiated cloud services, among other types of data. Applications may then be built based on a big data platform to generate events or derive other information. Applications may also behave like another monitoring systems (e.g. detecting operations anomalies, usage anomalies, security anomalies, among other functional anomalies).

[00180] Further, the monitoring engine (313) may monitor a number of performance- and utilization-related events that may occur within the instantiated service (312). These performance- and utilization-related events may include, for example, processor utilization within a number of the nodes, utilization of a number of the nodes by, for example, customers of the consumer's business, and levels of remaining data storage space within a data storage device, among many other performance- and utilization-related events.

[00181] In some examples, the monitoring policies used in node monitoring and event processing may define what devices or portions of the cloud service are to be monitored, how to execute such monitoring, and to whom or what devices to delegate the roles of monitoring.

[00182] In another example, the monitoring policies also define how detected events that occur in an instantiated topology (312) are processed. As will be described in more detail below, the monitoring policies may define what

to do with the detected events. For example, the monitoring policies may indicate how to process the events, where the events are sent, what devices or individuals address the events, how incidents resulting from the processing of the events are handled, how the events and incidents are processed, and how the resulting incidents are handled. For example, the monitoring policies may indicate a number of notification recipients, and may indicate what processes are to be performed on the detected events. More detail concerning the notification recipients and processes to be performed is given below in connection with Figs. 11 and 12.

[00183] As described above, a monitoring engine (313) may be deployed if not already existent, or setup and configured if already available in order to monitor the instantiated service (312). In one example, the monitoring engine (313) is deployed to monitor the instantiated service (312) in cases where the monitoring engine (313) already exists. In this example, a number of existing monitoring devices may be used to monitor the instantiated service (312) autonomously, through human intervention, or a combination thereof by configuring the existing monitoring engine (313) to match the monitoring policies defined when designing the topology (302). In this example, the monitoring engine (313) already existent may be configured based on the monitoring policies associated with the topology (302) and the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) of the topology (302) as described above, configured based on input from a user, or combinations thereof. In another example, a previously non-existent monitoring engine (313) may be provisioned and deployed based on the monitoring policies defined when designing the topology (302). In this example, the monitoring engine (313) is provisioned and deployed at the same time as the provisioning and deployment of the instantiated service (312). Further, the monitoring engine (313), in this example, is deployed and managed based on the monitoring policies associated with the topology (302) and the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) of the topology (302) as described above. In any of the above examples, a complete service as outlined by the topology (302) is

created, including the instantiated system (312) and the monitoring engine (313).

[00184] As described above, the system may include an event handler (316) that processes the detected events based on the monitoring policies. In other words, the monitoring policies that define the processes to be performed on detected events may assist an event handler (316) in receiving and processing the detected events, and in making decisions and sending notification messages regarding incidents arising out of the events.

[00185] The monitoring engine (313) may inform the event handler (316) of any events detected by the monitoring engine (313). The event handler (316) may be any computing device that receives data associated with detected events from the monitoring engine (313), and processes the data in order to create a remediation of any incidents that may arise from the detected events. Thus, the topology-based management broker (Fig. 3, 300) processes the events that are detected by the monitoring engine (313). Processing of events detected by the monitoring engine (313) may be performed by the event handler (316).

[00186] The event handler (316) may receive any kind or amount of data from the monitoring engine (313). As described above, the data received from the monitoring engine (313) by the event handler (316) may include any data associated with the operation, security, compliance, usage, and analytics of the instantiated service (312) as a whole, and the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) within the instantiated service (312) as groups of nodes and individually. In one example, the event handler (316) performs a number of requests for the event data. In this example, the event handler (316) may poll the monitoring engine (313) for the event data after a predefined time period, randomly, when triggered by another event, or a combination thereof. As described above, event processing may, in one example, be delegated to another system or third party service. For example, event processing such as correlation and filtering of events and incidents and incident identification may be delegated to HP BUSINESS SERVICE MANAGEMENT; a suite of service management software tools developed and distributed by the Hewlett Packard

Corporation. Remediation processes may be delegated to HP OPERATIONS MANAGER I (HP OMI) or SITESCOPE; both comprising a suite of software tools developed and distributed by the Hewlett Packard Corporation. Security event notification, processing, and remediation can be delegated to HP ARCSIGHT; a suite of service management software tools developed and distributed by the Hewlett Packard Corporation. In one example, HP ARCSIGHT may reference the service agreement (SA) associated with the instantiated service (312) to comply with the SA.

[00187] As described above, the event handler (316) may process the events detected from the monitoring engine (313). Types of processing that the event handler (316) may perform include filtering, correlation, aggregation of the events, among other forms of event processing, and combinations thereof. In one example, a number of events may collectively be subjected to a number of forms of event processing in order to create an incident. More detail concerning the processing of events is given below in connection with Fig. 11 and 12.

[00188] The event handler (316) may determine whether an event indicates a remediation action, and whether and how to present a notification of the event to a consumer, administrator, third party, or other user of the topology-based management broker (Fig. 3, 300) or instantiated service (312). If the event handler (316) determines that a remediation action is to be taken in connection with an event, the event handler (316) generates an incident based on the detected event. In this example, the events may individually not support the creation of an incident that indicates remediation, but a number of events, when analyzed by the event handler (316), may indicate that an issue within the instantiated topology (312) is not in agreement with the policies (303), or is otherwise in need of remediation.

[00189] The event handler (316) may send a notification of the incident to a number of notification recipients, which notification recipients are defined in the monitoring policies. The event handler (316) may also send the data associated with the detected event to a remediation engine (317) that remediates the incident. In other examples, the event handler (316) may not

send data to a remediation engine (317). In other words, the event handler (316) may not have a prescription on what to do with the events, and hence may not be present. Accordingly, the event handler (316) may be configured via an element other than the content of a topology, or being associated to the topologies.

[00190] The monitoring system (313) as described herein may be beneficial in that it provides a way to establish the monitoring system (313) and also provides a way to provision the monitoring system (313). Additionally, the monitoring system (313) provides a topology-based monitoring configuration. Yet further, the monitoring system (313) provides a separation of event processing from incident remediation.

[00191] Fig. 11 is a flowchart showing a method (1100) for monitoring a cloud service modeled as a topology, according to one example of the principles described herein. The method (1100) of Fig. 11 may begin by establishing (block 1101) a number of monitoring policies to monitor an instantiated service (Fig. 3, 312). As described above, the monitoring policies may define the functionality of the nodes (Fig. 3, 302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) and include criteria by which events in node (Fig. 3, 302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) functionality may be detected.

[00192] In some examples, the monitoring policies define the distribution of event notifications. For example, a monitoring policy may indicate a number of notification recipients. Such notification recipients may include service designers, infrastructure architects or administrators, system administrators, information technology operators, offer managers, or users, among other personnel with roles in the monitoring and remediation of a topology.

[00193] The monitoring policies may also indicate processes to be performed on detected events. For example, a monitoring policy may indicate whether a detected event should be filtered, correlated, aggregated, or combinations thereof, in relation to other detected events. More detail concerning the filtration, correlation and aggregation of detected events is given below in connection with block 1103. In addition to the specific elements

described in connection with Fig. 11, the monitoring policies may include any number of other elements described in connection with Fig. 10, including elements that indicate which, and how many, nodes in the topology are to be monitored.

[00194] Still further, the monitoring policies may indicate how to remediate incidents related to the detected events. As described above, in some examples detected events may result in incidents that are processed by a remediation engine (Fig. 3, 317). In this example, the monitoring policies may indicate how an incident based on a particular event is to be remediated. For example, the monitoring policy may indicate that an incident based on a first detected event may be manually remediated via a graphical user interface (GUI). In another example, the monitoring policy may indicate that an incident based on a second particular detected event may be passed to logic such that an autonomous remediation may be performed on the incident.

[00195] In some examples, the monitoring policies may be established based on a ticketing support system that may indicate instances of past detected events. The monitoring policies may also be established based on self-service identification of events, or based on the predetermined avoidance of particular events.

[00196] As described above, a user may use an interface of the monitoring system (Fig. 3, 313) to establish a number of monitoring policies. These users may be service designers, infrastructure architects or administrators, system administrators, information technology operators, offer managers, or users, among other personnel with roles in the design and/or monitoring of a topology. In still another example, the interface may be operated by a third party.

[00197] The monitoring system (Fig. 3, 313) may then be executed (block 1102) and instantiated based on the monitoring policies. In some examples, executing (block 1102) and instantiating the monitoring system may include setting up a monitoring system based on a topology, configuring (or reconfiguring) an existent monitoring system, implementing monitoring policies while stitching topologies, or combinations thereof.

[00198] The monitoring system (Fig. 3, 313) may collect metric data (block 1103) based on the number of monitoring policies. For example, the monitoring system (Fig. 3, 313) may detect instances when a node's or a group of nodes' functionality does not comply with the functionalities defined in the monitoring policies. For example, as described above, the monitoring policies may relate to operational monitoring, security monitoring, compliance monitoring, usage monitoring, and analytics monitoring. Accordingly, the monitoring system (Fig. 3, 313) may collect operational data, security data, compliance data, usage data, analytics data, big data, or combinations thereof.

[00199] An example of operational data may be that a node or a group of nodes is not operating at an established CPU level. An example of a security data may be that a node or group of nodes contain a virus. An example of a compliance data may be that a node or group of nodes is running an outdated version of an operating system. An example of a usage data may be that a consumer has utilized more memory than allotted per the policy allowance. An example of an analytics data may be that a node or group of nodes is not following a defined workflow.

[00200] Any block within the method of Fig. 11 may be performed alone or in combination with any number of other processes within Fig. 11. For example, monitoring policies may be established at block 1101 without the remaining processes being performed, or less than all of the remaining processes being performed.

[00201] Fig. 12 is a flowchart showing another method (1200) for monitoring a cloud service modeled as a topology, according to another example of the principles described herein. The method (1200) of Fig. 12 may begin by establishing (block 1201) a number of monitoring policies to monitor an instantiated topology. In some examples, this may be performed as described in connection with Fig. 11. In some examples, the monitoring system (Fig. 3, 313) may be established (block 1202) as an additional topology in addition to the topology (Fig. 3, 302). Similarly, an event handling system may be established as an additional topology in addition to the topology (Fig. 3, 302).

[00202] The method (1200) may allow dynamic addition of notification recipients. As described above, monitoring policies may indicate notification procedures of a detected event. In particular, the monitoring policy may indicate a number of notification recipients that receive a notification of an event. In some examples, notification recipients may be added to a monitoring policy, through a registration process for example. Such additional notification recipients may include service designers, infrastructure architects or administrators, system administrators, information technology operators, offer managers, or users, among other personnel with roles in the monitoring and remediation of a topology.

[00203] The monitoring system (Fig. 3, 313) may collect (block 1203) metric data based on the number of monitoring policies. In some examples, this may be performed as described in connection with Fig. 11.

[00204] A determination (block 1204) may be made as to whether a service is to be modified. For example, an instantiated service (Fig. 3, 312) may be monitored as was specified by the policies (Fig. 3, 303), or the monitoring/event handling, incident management topologies. If the instantiated service (Fig. 3, 312) is modified as part of a normal lifecycle (user updates, administrative updates, service deletion, or service retirement, for example), then the modification may be performed via LCMA orchestration that may also retire the monitoring system (Fig. 3, 313), the event handler (Fig. 3, 316) and an incident handling system. In some examples, this may be done based on the information stored in the RTSM and with the realized topology in the form of the associated monitoring policies (e.g. in the form or along with a topology of a monitoring system) and LCMAs and provisioning policies. For example, if an instantiated service (Fig. 3, 312) is retired, the system (Fig. 3, 300) may see also what monitoring service to retire or make available. In another example, if a node (Fig. 3, 302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) is moved, the system (Fig. 3, 300) knows how to update the monitoring system (Fig. 3, 313). If a node (Fig. 3, 302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) is added, the addition may include monitoring policies (Fig. 3, 303). The addition may be derived from a topology-wide monitoring policy.

[00205] In some examples, the event handler (Fig. 3, 316) may process (block 1205) the collected metric data based on user input. For example, in some cases, the monitoring topology or the monitoring policies (Fig. 3, 303) may not specify what to do with collected metric data. In this case, a user of the monitoring system (Fig. 3, 313) may decide what to do with the collected metric data.

[00206] In another example, the event handler (Fig. 3, 316) may process (block 1206) the collected metric data based on the number of monitoring policies. For example, the monitoring system (Fig. 3, 313) may monitor particular metrics and data that may indicate an event. The monitoring system (Fig. 3, 313) may send the events, per the policies and monitoring system (Fig. 3, 313) setup, to the event handler (Fig. 3, 316) based on a number of the policies associated with the designed topology (Fig. 3, 302) and the instantiated service (Fig. 3, 312). The monitoring policies include a portion that defines what to do with the data that result from the monitoring such as, for example, how to process (e.g. filter, correlate, aggregate, among other processes) the events, where the events are sent, what devices or individuals address the events, how incidents resulting from the processing of the events are handled, how the events (e.g., processed as aggregated, filtered, or correlated events, among other forms of processing), and how the resulting incidents are notified and handled.

[00207] Accordingly, the event handler (Fig. 3, 316) may process (block 1206) the collected metrics. In one example, the event handler (Fig. 3, 316) may filter a detected event. For example, some detected events while not in compliance with the policies (Fig. 3, 303) may not pose a threat to the cloud service. Accordingly, the event handler (Fig. 3, 316) may include thresholds, or other filtering elements that filter detected events based on a number of criteria. While a specific example depicting security monitoring has been presented, as described above, a number of different node functionalities may be monitored including operations, usage, compliance and analytics. In these examples, some of the metric data is just used for collection (and possible later processing such as usage analysis, analytics analysis via BAM, BI, or big data analysis).

For example, with the big data, the event handler (Fig. 3, 316) may send events to a big data system for storage and real-time processing with historical data to detect anomalies and generate reports. In this example, the big data system may be a feature of the monitoring system (Fig. 3, 313), or set up via a number of policies, or as an additional topology.

[00208] The event handler (Fig. 3, 316) may correlate the metric data. Correlating metric data may be beneficial in that it identifies a possible problem. Another way to indicate a potential problem is using big data analytics for real-time processing of many events. For example, the event handler (Fig. 3, 316) may indicate that a number of detected events relate to a particular issue. Accordingly, the event handler (Fig. 3, 316) may correlate the detected events or otherwise associate the detected events to one another. An incident that reflects the relationship between the detected events may then be generated if an issue that warrants notification is indicated (or identified based on policies).

[00209] The event handler (Fig. 3, 316) may aggregate a number of detected events. In one example, the event handler (Fig. 3, 316) may aggregate a number of instances of a particular type of detected event. For example, the event handler (Fig. 3, 316) may aggregate a number of instances of a node or group of nodes operating at a CPU level greater than the CPU level defined in a policy. In another example, the event handler (Fig. 3, 316) may wait until a certain number of detected events are received before generating an incident based on the detected events.

[00210] In some examples, a report may be generated (block 1206) based on the collected metric data. For example, log monitoring, performance monitoring, and intelligence monitoring such as, for example, collection of metrics, business intelligence (BI) and business activity monitoring (BAM) and analytics/big data integration, among other types monitoring-related policies may be collected and a report generated (block 1206) based on this information. By comparison, in some examples, a report is not generated, and no action beyond monitoring the event is performed. In other examples, a report may not specify what to do with an event.

[00211] In yet another example, as described above, the collected data may be used to generate an incident which may be processed in remediation. For example, policies (Fig. 3, 302) defining a monitoring system (Fig. 3, 313), or a monitoring system topology may prescribe that events relating to the collected metric data are to be processed and incidents relating to the events generated. Accordingly, in some examples, a remediate engine (Fig. 3, 317) may remediate (block 1208) incidents generated based on a number of detected events. In some examples, the remediate engine (Fig. 3, 317) may remediate the incidents as described above in connection with Fig. 6.

[00212] Any block within the method of Fig. 12 may be performed alone or in combination with any number of other processes within Fig. 12. For example, monitoring policies may be established at block 1201 without the remaining processes being performed, or less than all of the remaining processes being performed.

[00213] Fig. 13 depicts a block diagram of an instantiated service that includes monitoring policies, according to one example of the principles described herein. As described above, a topology (302) may define a number of nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7), and a number of relationships between the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7). In Fig. 13, the relationships between nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7) are depicted as lines connecting the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7). Each of the nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7), the entire topology (302), a group of nodes (302-1, 302-2, 302-3, 302-4, 302-5, 302-6, 302-7), portions of the topology (302), or combinations thereof are associated with a number of monitoring policies (303-1, 303-2, 303-3, 303-4, 303-5, 303-6, 303-7).

[00214] In some examples, the monitoring policies may be data or metadata provided in the same file describing the nodes or topology. The metadata may be in a separate file associated with the nodes or topology. In one example, the monitoring policies may be accessed via uniform resource identifiers (URIs) of application programming interfaces (APIs) to perform calls in order to execute the APIs.

[00215] The monitoring policies may include the monitoring topology, metadata on how to provision a monitoring system. The monitoring policies may also setup a monitoring system, or configure an existent monitoring system. The monitoring policies may also include instructions to execute the monitoring.

[00216] Fig. 14 depicts a block diagram of an instantiated service (314) and a monitoring system topology (302), according to one example of the principles described herein. As depicted in Fig. 14, in some examples, the monitoring policies (303) may be established in a monitoring system topology (302). In this example, the monitoring policies (303) may be implemented in the instantiated service (314) according to the monitoring system topology (302). In this example, the monitoring system topology (302) may indicate how events are to be processed. For example, the monitoring system topology (302) may indicate when and how events are to be reported, when and how events are to be remediated as incidents, and who receives notification of events and/or incidents, among other event related instructions. It may be a modeling choice to represent a monitoring topology as a particular case of monitoring policies. In some examples, when a realized topology is updated, the monitoring system (Fig. 3, 3113) may also be updated as a result.

[00217] Aspects of the present system and method are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to examples of the principles described herein. Each block of the flowchart illustrations and block diagrams, and combinations of blocks in the flowchart illustrations and block diagrams, may be implemented by computer usable program code. The computer usable program code may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the computer usable program code, when executed via, for example, the a number of processors within the devices comprising the topology-based management broker (200) or other programmable data processing apparatus, implement the functions or acts specified in the flowchart and/or block diagram block or blocks.

In one example, the computer usable program code may be embodied within a computer readable storage medium; the computer readable storage medium being part of the computer program product. In one example, the computer readable storage medium is a non-transitory computer readable medium.

[00218] The specification and figures describe methods and systems of managing the lifecycle of cloud service modeled as a topology. These systems and methods include, with a processor, generating a topology, the topology representing a cloud service, associating a number of lifecycle management actions (LCMAs) with a number of nodes within the topology, and with a lifecycle management engine, executing the topology.

[00219] This management of the lifecycle of a cloud service modeled as a topology may have a number of advantages, including: (1) providing a common stack along with common use of topologies, realized topologies, and policies may be used to support all use cases for both cloud service automation (CSA) and continued delivery automation (CDA) platforms and services to construct topologies while utilizing the same technology and supporting multiple providers' associated technologies; (2) providing a computing environment in which CSA and CDA use the same topology representations such as, for example, extensible mark-up language (XML) or JavaScript object mutation (JSON) or YAML; (3) providing a method of managing migration of content for CSA by reusing existing CSA content, creating a path to migrate resource providers, and reusing providers; (4) avoiding or alleviating the risk of perpetuating a CSA/CDA confusion, duplication of efforts and endangering future CSA opportunities; (5) complex applications may be automatically deployed on requested infrastructure without also requiring users to understand how to perform such operations, and (6) supports a CM&S environment, among many other advantages.

[00220] The preceding description has been presented to illustrate and describe examples of the principles described. This description is not intended to be exhaustive or to limit these principles to any precise form disclosed. Many modifications and variations are possible in light of the above teaching.

CLAIMS

WHAT IS CLAIMED IS:

1. A method of monitoring a cloud service modeled as a topology with a processor, comprising:
 - establishing a number of monitoring policies to monitor an instantiated topology;
 - executing and instantiating a monitoring system based on the monitoring policies; and
 - collecting metric data at a node based on the number of monitoring policies.
2. The method of claim 1, in which executing and instantiating a monitoring system comprises provisioning, configuring, delegating, or combinations thereof, a monitoring system based on the monitoring policies.
3. The method of claim 1, in which the monitoring policies are established based on a monitoring topology.
4. The method of claim 1, in which the monitoring policies comprise event handler instructions to related to filtering events, correlating events, aggregating events, generating incidents, generating incident notifications, or combinations thereof.
5. The method of claim 1, in which the number of monitoring policies indicate a number of notification recipients, a number of nodes to monitor, what metric data to collect, processes to perform on events, incident generation instructions, reporting instructions, or combinations thereof.

6. The method of claim 1, in which collecting metric data comprises collecting operation data, security data, compliance data, usage data, analytics data, big data, or combinations thereof,

in which the operation data includes performance data, service level agreement data, performance degradation data, service level agreement degradation data, trend data, or combinations thereof, and

in which the analytics data includes business activity monitoring data (BAM), business intelligence data (BI), or combinations thereof.

7. The method of claim 1, further comprising updating the monitoring policies based on remediation, called lifecycle management actions (LCMAs), or combinations thereof, in which updated monitoring policies remain associated with a realized topology such that an update to a monitoring policy may be promulgated to instances of a realized topology.

8. The method of claim 1, further comprising converting a number of policies into topologies.

9. The method of claim 1, further comprising implementing monitoring policies as defined in a monitoring system topology.

10. A system for monitoring a cloud service modeled as a topology, comprising:

an interface to assist in establishing a number of monitoring policies;

a broker to execute and instantiate a monitoring system based on the monitoring policies; and

a monitoring engine to collect metric data based on the number of monitoring policies.

11. The system of claim 10, further comprising an event handler to:
process the metric data based on the monitoring policies;
create an incident based on the detected events; and

send a notification of the number of incidents to a notification recipient, the notification comprising information regarding an incident.

12. The system of claim 10, in which the interface prompts a user to establish, configure, update or combinations thereof, an instantiated service via subscription management tools or application programming interfaces (APIs).

13. The system of claim 11, in which the event handler performs filtration, correlation, aggregation, or combinations thereof on a number of detected events.

14. A computer program product for monitoring a cloud service modeled as a topology, the computer program product comprising:

a computer readable storage medium comprising computer usable program code embodied therewith, the computer usable program code comprising:

computer usable program code to, when executed by a processor, establish a number of monitoring policies to monitor an instantiated topology;

computer usable program code to, when executed by a processor, determine a number of notification recipients, the notification recipients being indicated in the number of monitoring policies;


computer usable program code to, when executed by a processor, determine a number of processes to be performed on detected events, the processes being identified in the number of monitoring policies;

computer usable program code to, when executed by a processor, detect an event based on the number of monitoring policies; and

computer usable program code to, when executed by a processor, process the detected event based on the number of monitoring policies.

15. The computer program product of claim 14, further comprising computer usable program code to, when executed by a processor, generate a report based on the detected events.

1/14

100 




Fig. 1

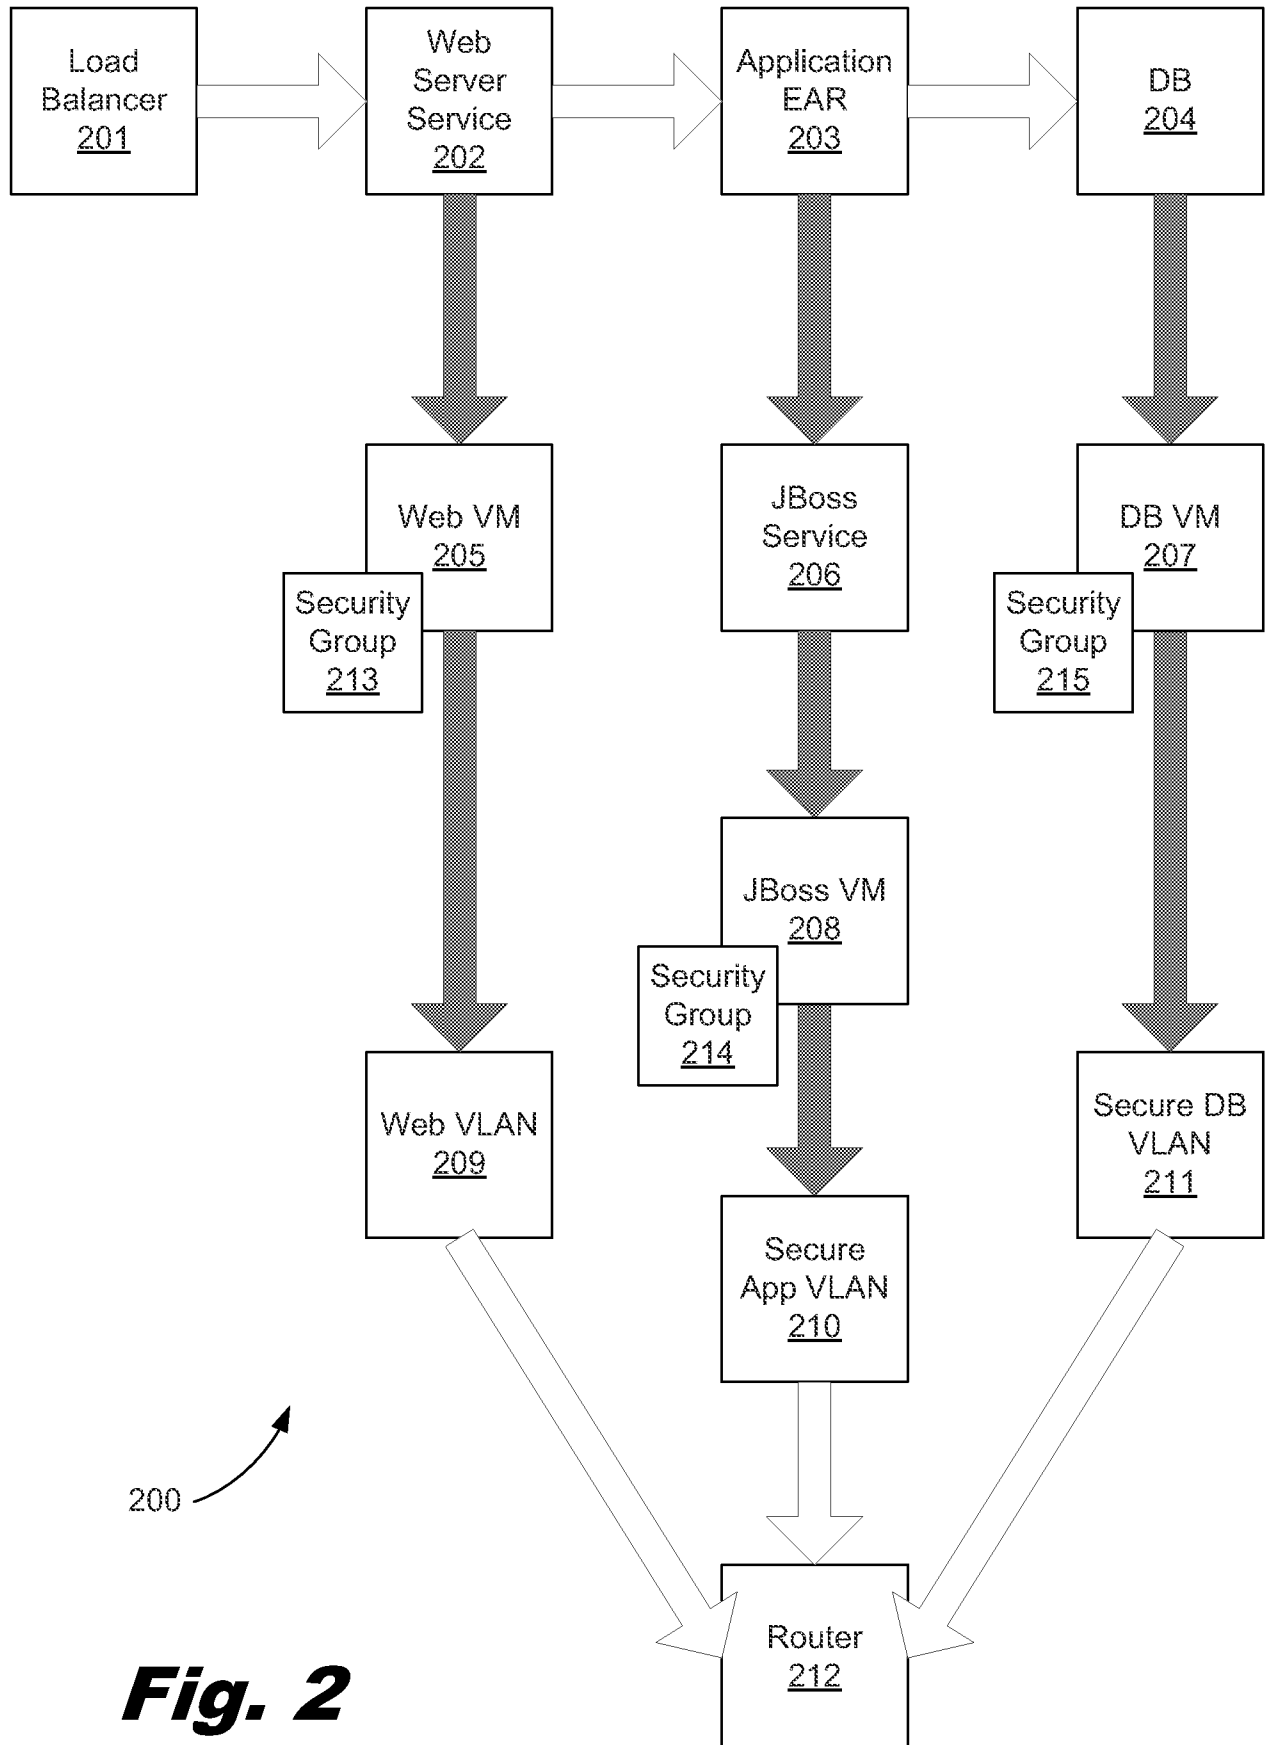


Fig. 2

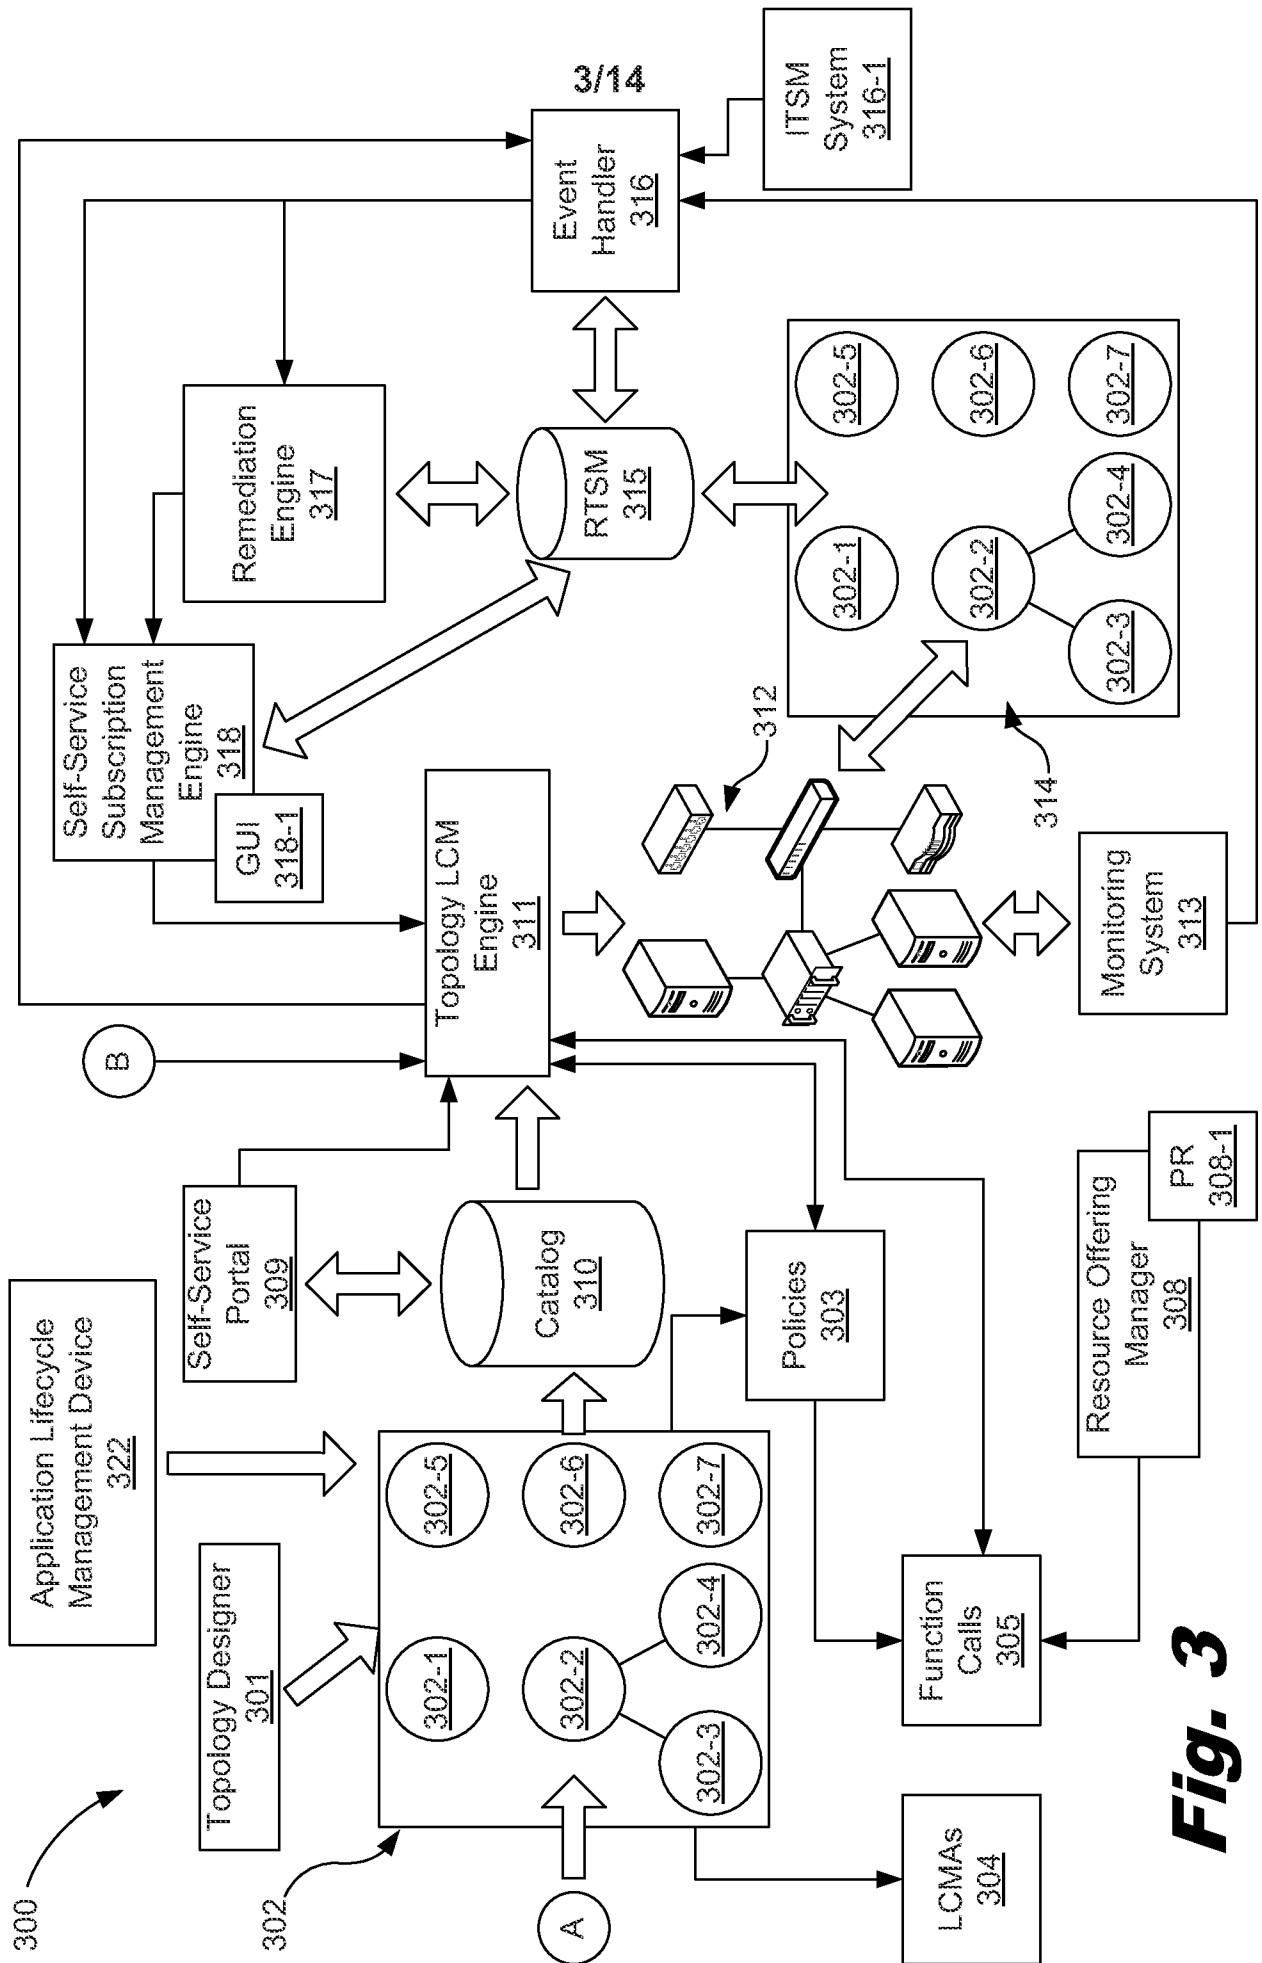


Fig. 3

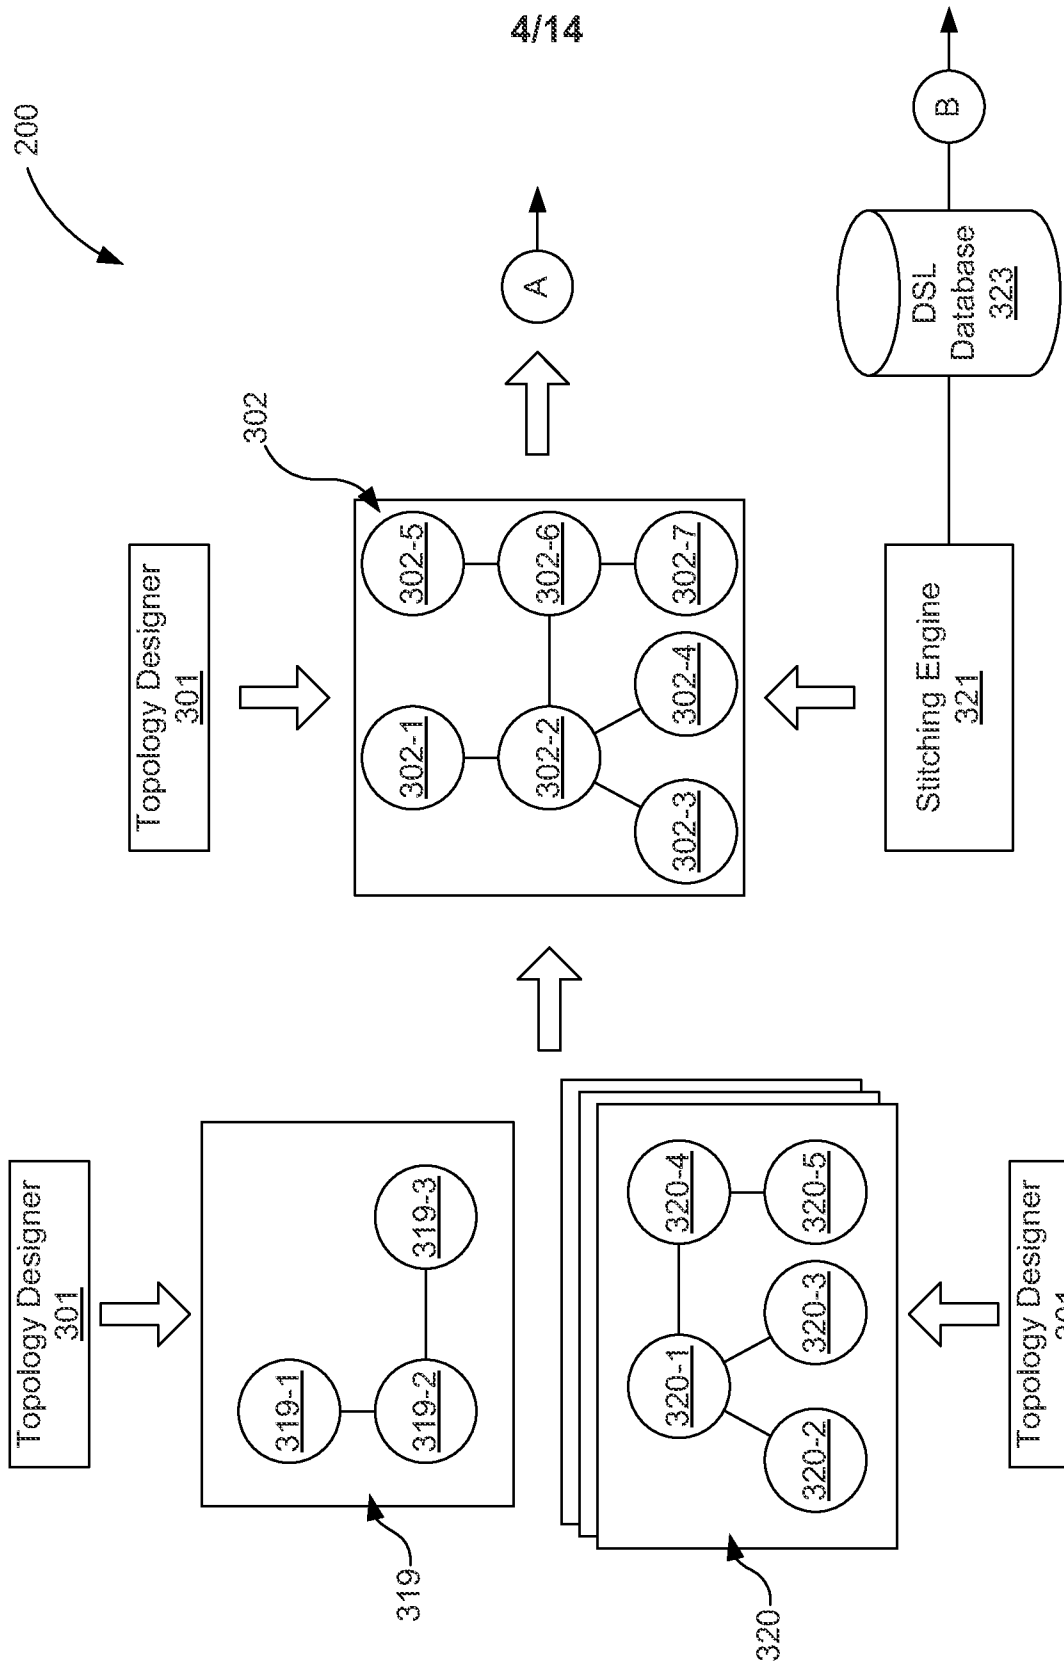


Fig. 4

5/14

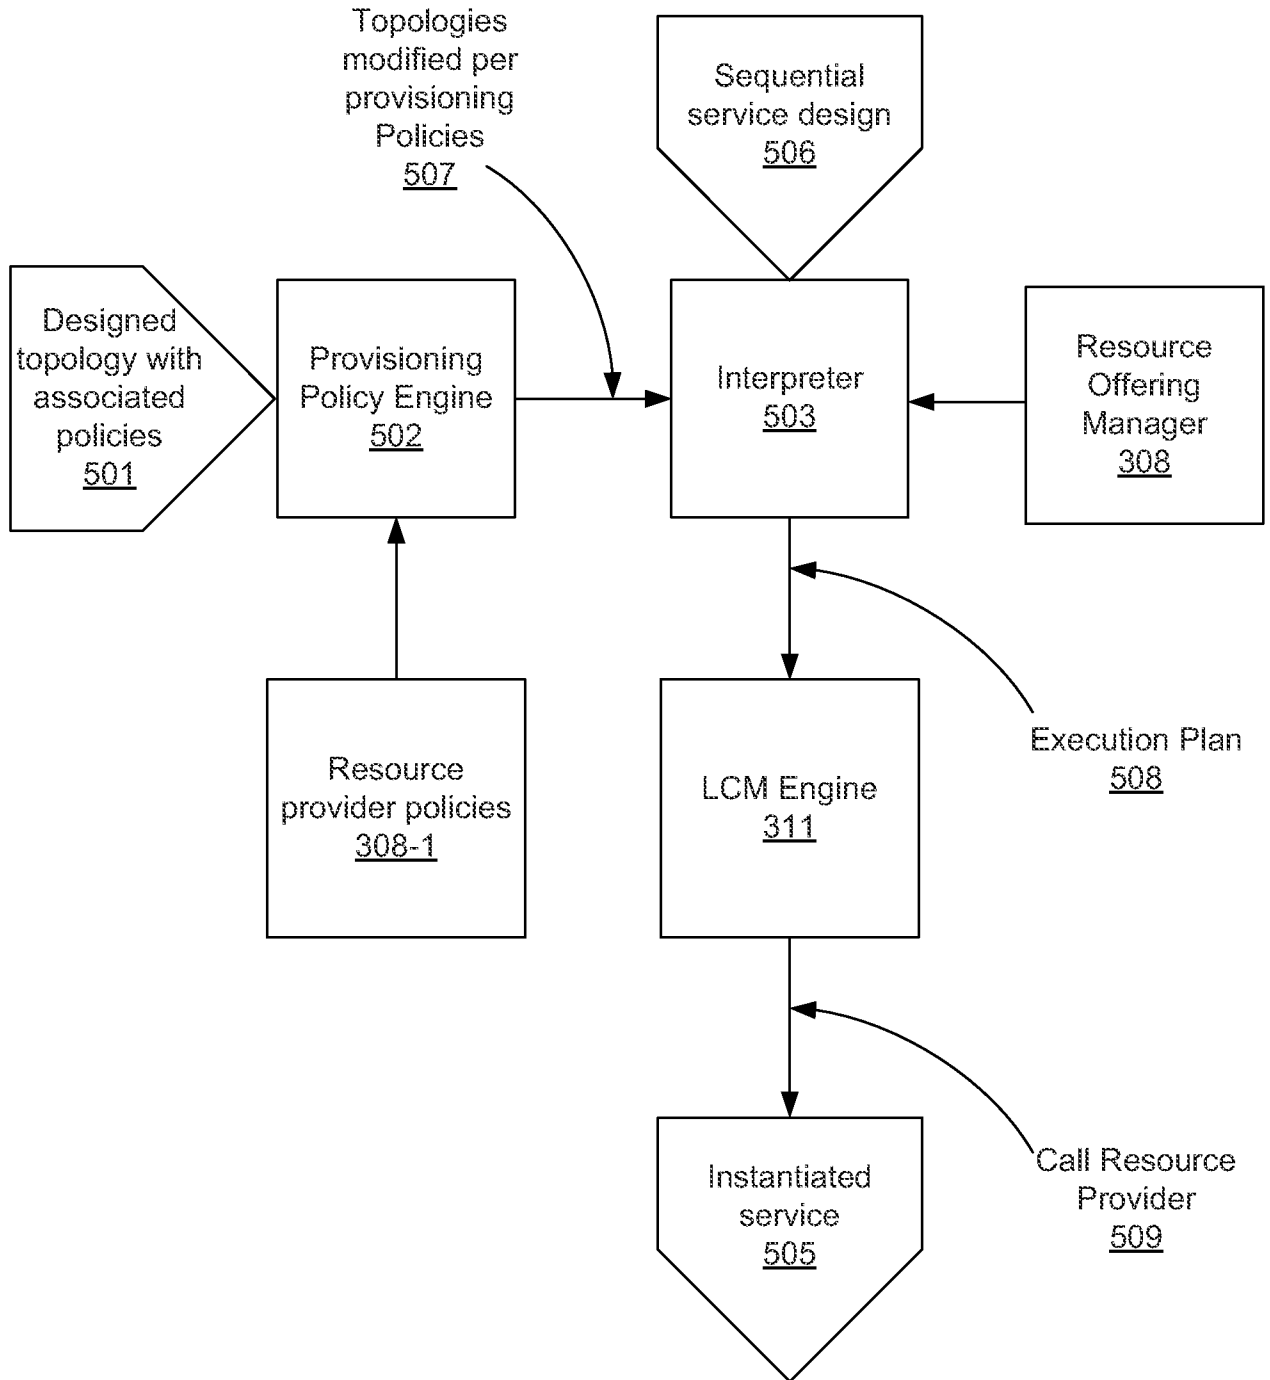


Fig. 5

6/14

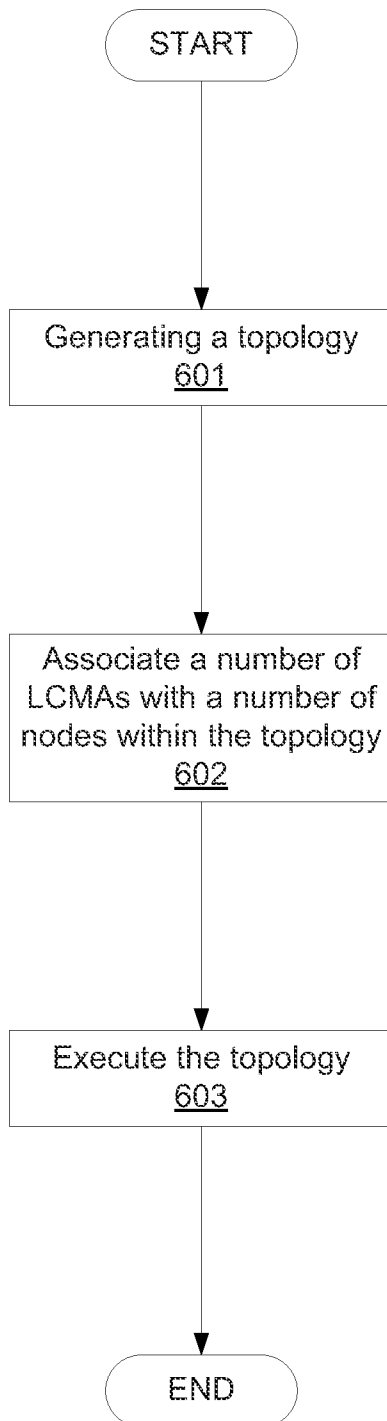


Fig. 6

7/14

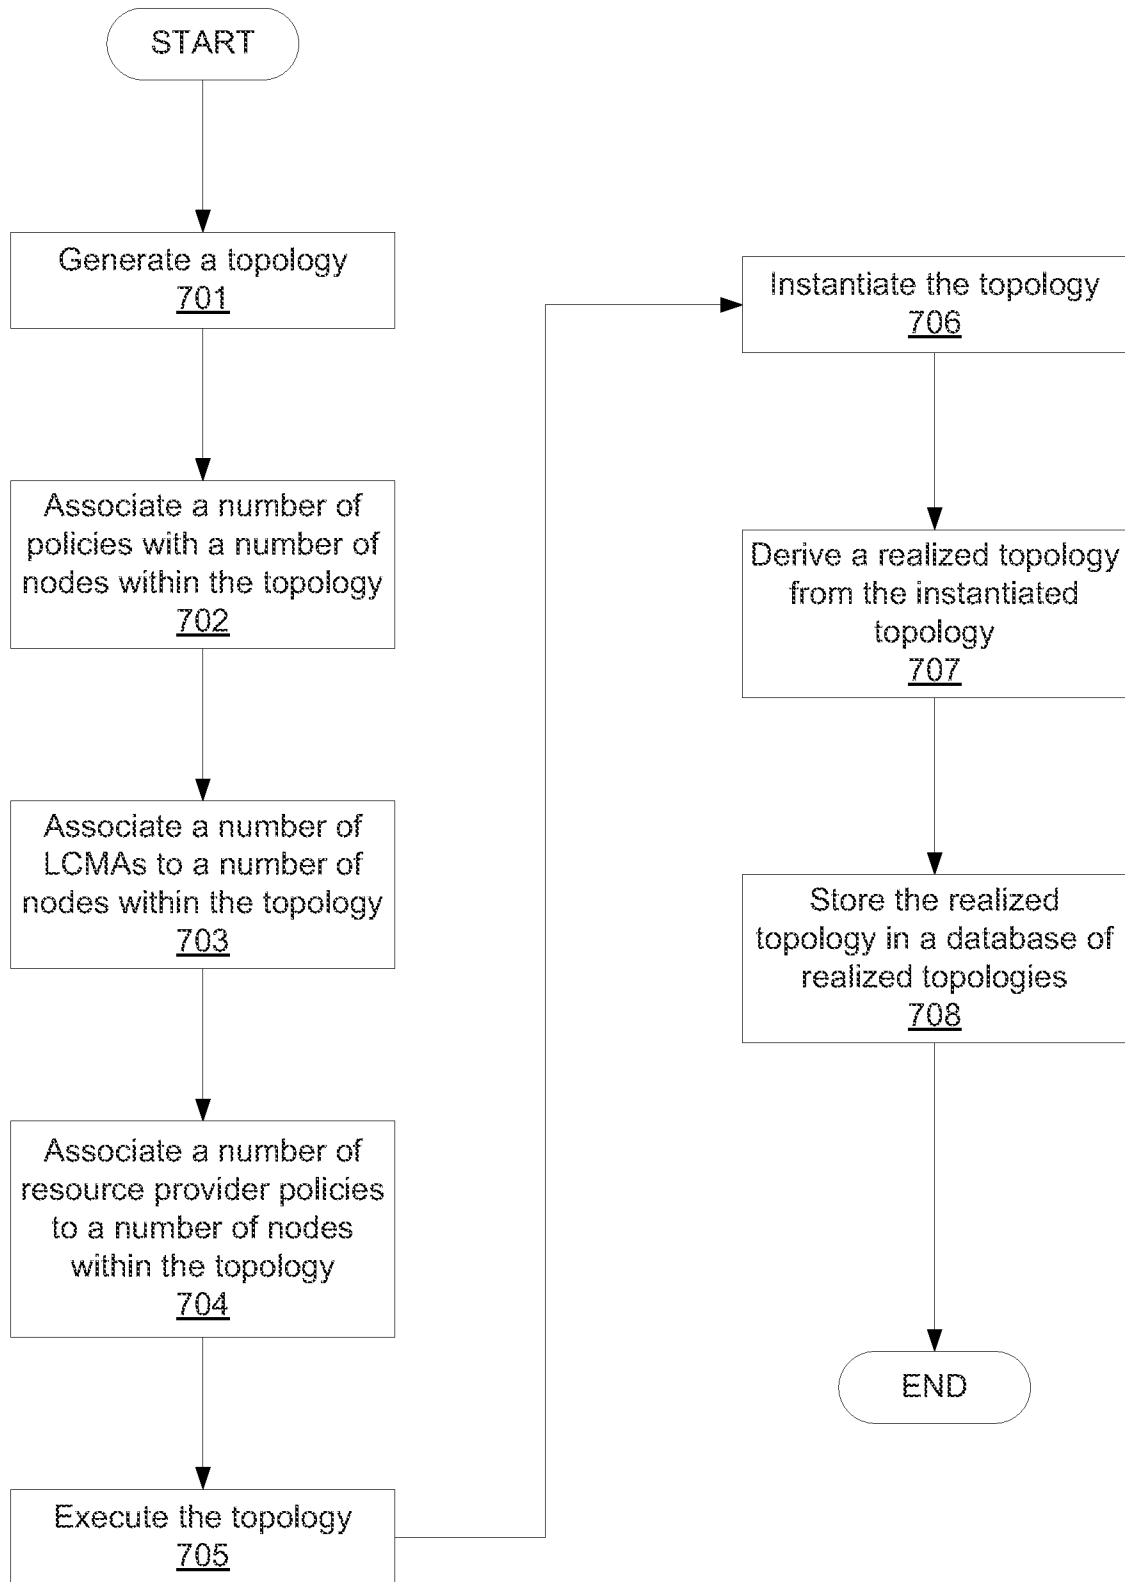


Fig. 7

8/14

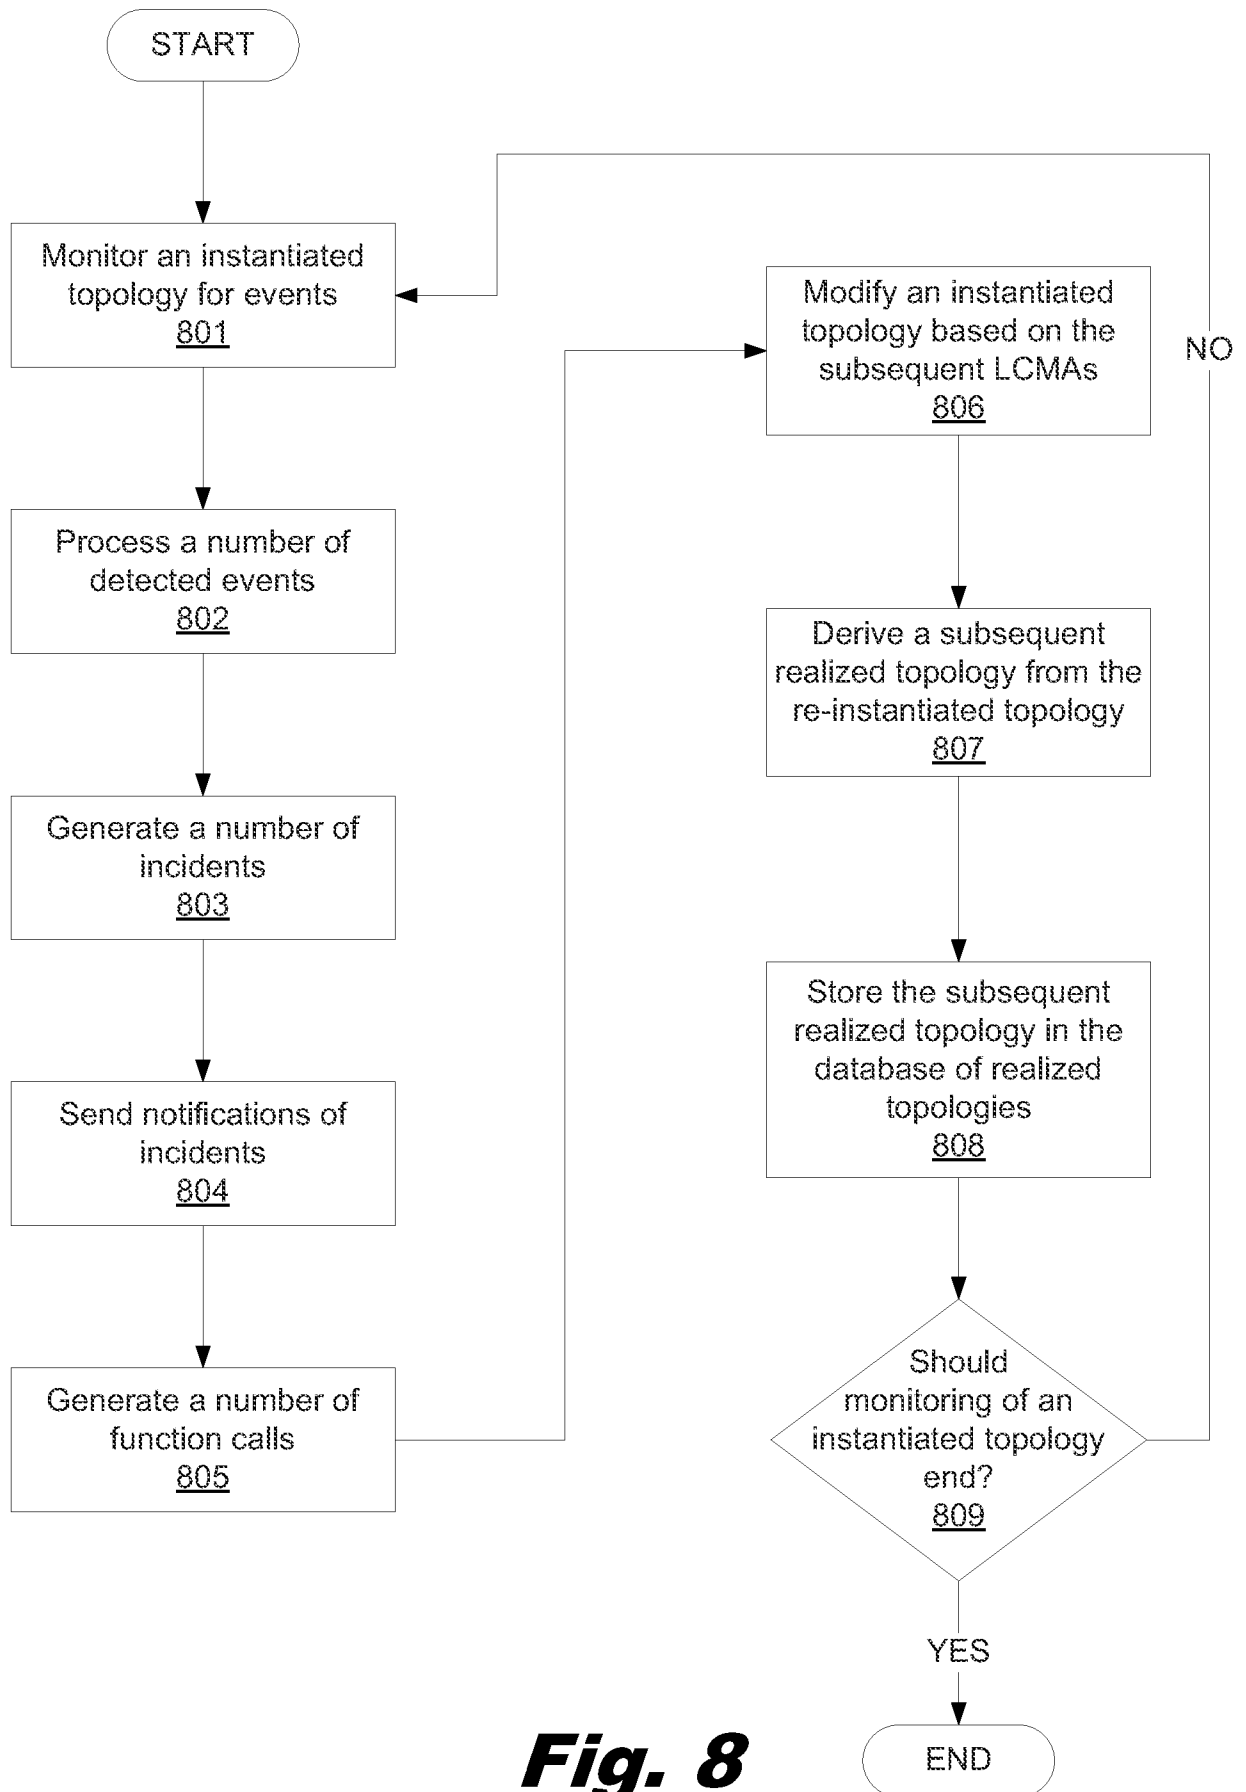



Fig. 8

9/14

**Fig. 9**

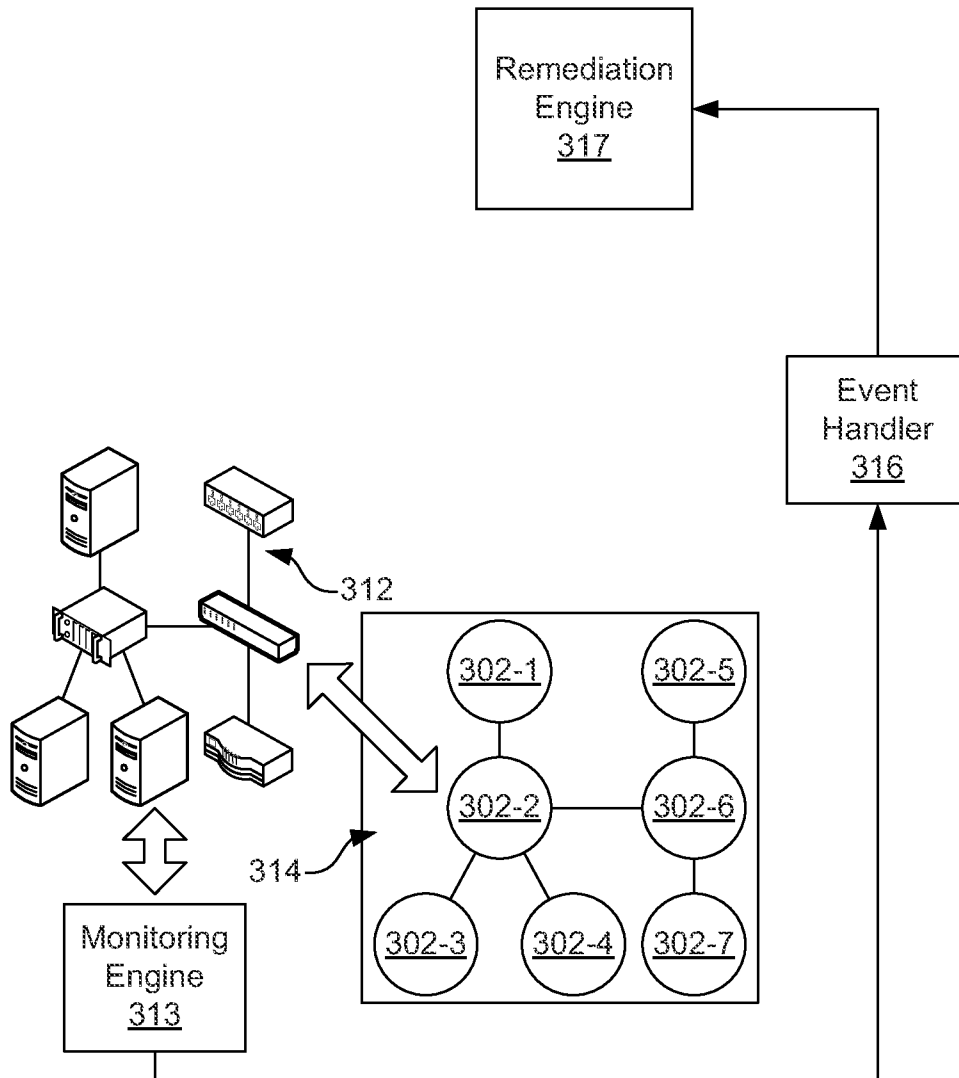


Fig. 10

11/14

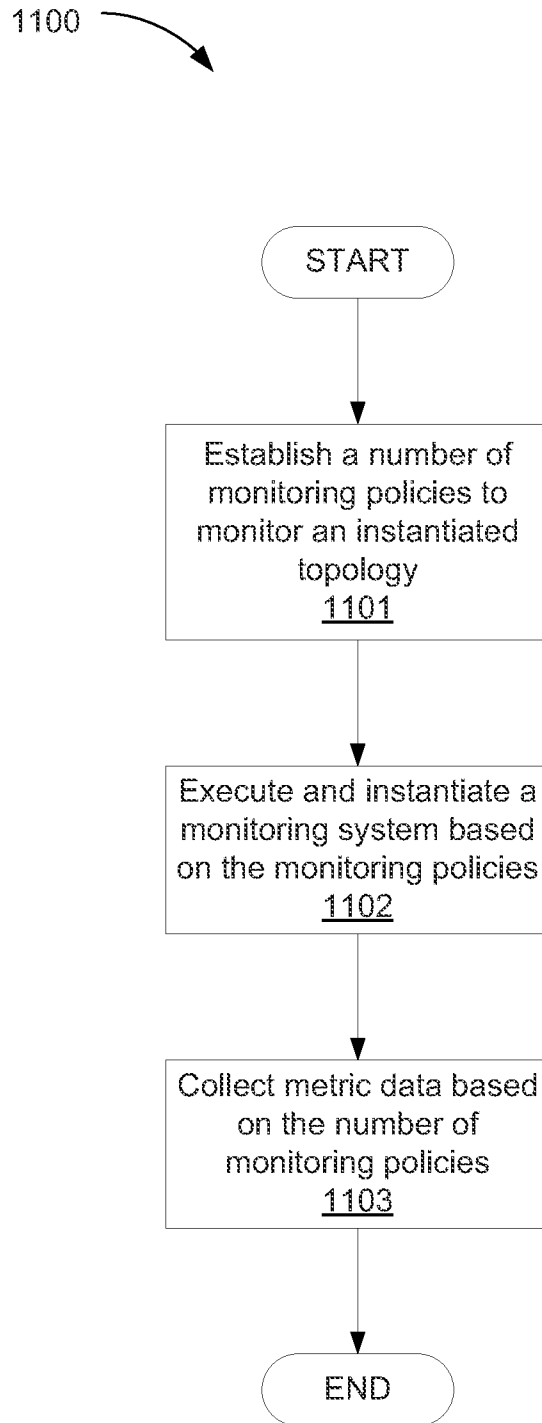


Fig. 11

12/14

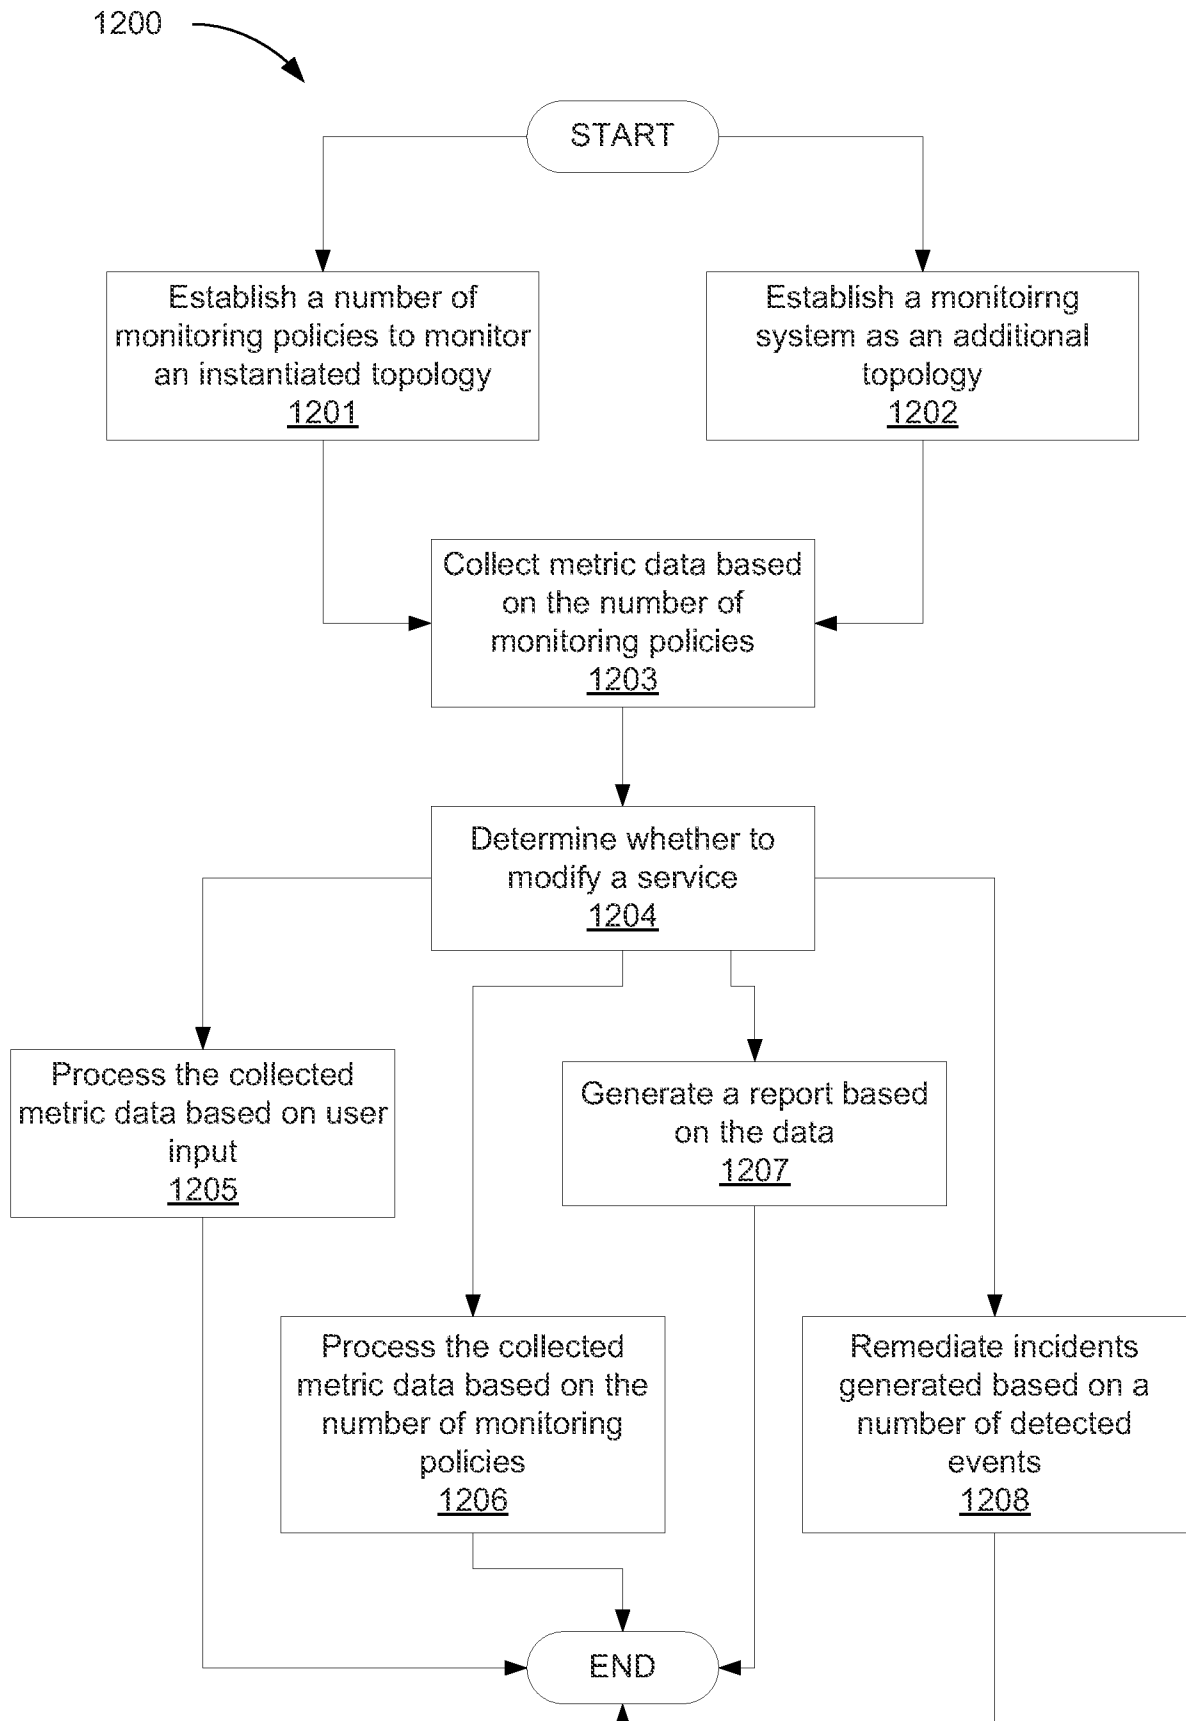


Fig. 12




Fig. 13




Fig. 14

A. CLASSIFICATION OF SUBJECT MATTER**G06F 17/30(2006.01)i, G06F 15/16(2006.01)i**

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F 17/30; G06F 15/16; G06F 15/173; G06F 15/177; G06F 3/01; H04L 12/56

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords:topology, cloud, monitor, policy, and similar terms.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2013-0086249 A1 (CHRISTOPHER J. WHITE et al.) 04 April 2013 See paragraphs 7-60; and figure 4a.	1-6,8-15
Y		7
Y	US 2011-0029882 A1 (DEVENDRA RAJKUMAR JAISINGHANI) 03 February 2011 See paragraphs 34, 56.	7
A	US 2011-0075674 A1 (LI ERRAN LI et al.) 31 March 2011 See paragraphs 24-52; and figures 1-2.	1-15
A	US 2011-0055707 A1 (TIM KIMMET) 03 March 2011 See paragraphs 15-50; and figures 1-6.	1-15
A	US 2011-0289119 A1 (JOSEPH WEIZHEN HU et al.) 24 November 2011 See paragraphs 24-83; and figures 1-6.	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

22 July 2014 (22.07.2014)

Date of mailing of the international search report

24 July 2014 (24.07.2014)

Name and mailing address of the ISA/KR

International Application Division
Korean Intellectual Property Office
189 Cheongsu-ro, Seo-gu, Daejeon Metropolitan City, 302-701,
Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

BYUN, Sung Cheal

Telephone No. +82-42-481-8262



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2013/067513

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2013-0086249 A1	04/04/2013	None	
US 2011-0029882 A1	03/02/2011	CN 102576354 A DE 112010003144 T5 GB 201203172 D0 GB 2485712 A US 2011-0029673 A1 US 2011-0029810 A1 US 2011-0029981 A1 US 2012-0102180 A1 US 2013-0080902 A1 US 8316305 B2 WO 2011-014827 A1 WO 2011-014830 A1 WO 2011-014835 A1	11/07/2012 05/07/2012 11/04/2012 23/05/2012 03/02/2011 03/02/2011 03/02/2011 26/04/2012 28/03/2013 20/11/2012 03/02/2011 03/02/2011 03/02/2011
US 2011-0075674 A1	31/03/2011	CN 102577270 A EP 2484061 A1 EP 2484061 B1 JP 05485403 B2 JP 2013-507044 A KR 10-1355721 B1 KR 10-2012-0063501 A US 8619779 B2 WO 2011-041159 A1	11/07/2012 08/08/2012 17/07/2013 07/05/2014 28/02/2013 24/01/2014 15/06/2012 31/12/2013 07/04/2011
US 2011-0055707 A1	03/03/2011	US 2011-0041066 A1 US 2013-0290853 A1 US 8479098 B2 US 8635534 B2	17/02/2011 31/10/2013 02/07/2013 21/01/2014
US 2011-0289119 A1	24/11/2011	None	