



- (51) International Patent Classification:
G06F 9/50 (2006.01) G06F 15/16 (2006.01)
- (71) International Application Number:
PCT/US2014/054998
- (22) International Filing Date:
10 September 2014 (10.09.2014)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
3766/CHE/2014 31 July 2014 (31.07.2014) IN
- (71) Applicant: HEWLETT PACKARD DEVELOPMENT COMPANY, L.P. [US/US]; 11445 Compaq Center Drive West, Houston, Texas 77070 (US).

- (72) Inventors: KUMAR, Munirathnam; Hewlett Packard, Sy. No 192 Whitefield Road, Mahadevapura Post, Bangalore 560048 (IN). SUPARNA, Adarsh; Hewlett Packard, Sy. No. 192 Whitefield Road, Mahadevapura Post, Bangalore 560048 (IN).
- (74) Agents: GARDINER, Austin, W. et al.; Hewlett-Packard Company, Intellectual Property Administration, 3404 East Harmony Road, Mail Stop 35, Mail stop 35, Fort Collins, Colorado 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, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM,

[Continued on next page]

(54) Title: CLOUD RESOURCE POOL




FIG. 3

(57) Abstract: In one implementation, a cloud resource system includes a resource engine and a configuration engine. The resource engine can identify a cloud resource based on a template and assign the cloud resource based on a subscription and the template. The configuration engine can maintain a pool of cloud resource object instances.

WO 2016/018438 A1

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.

SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE,

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))

Cloud Resource Pool

BACKGROUND

[0001] Services are commonly offered via the cloud to support client demand. A cloud can include a virtual shared pool of resources for providing the service offering. Resources of a cloud can include infrastructure, storage, network, and other middleware applications to support a service. A cloud can increase or decrease allocated resources based on a subscription for the service. For example, resources can be allocated to a customer upon initiation of a subscription and resources can be released when the subscription terminates.

BRIEF DESCRIPTION OF THE DRAWINGS

[0002] Figures 1 and 2 are block diagrams depicting example cloud resource systems.

[0003] Figure 3 depicts an example environment in which various cloud resource systems can be implemented

[0004] Figures 4A and 4B are block diagrams depicting example interactions of example cloud resource system components.

[0005] Figures 5-8 are flow diagrams depicting example methods of reusing a cloud resource.

DETAILED DESCRIPTION

[0006] In the following description and figures, some example implementations of cloud resource systems and/or methods of reusing a cloud resource are described. In a heterogeneous cloud environment, data center resources can be distributed across multiple providers. Created subscriptions can allocate resources objects from the cloud system based on a cloud broker application. For example, the cloud broker can interact

with a centralized repository of virtual machine images templates to provide a service from each template. Cancelled subscriptions can delete the realized resource objects from the cloud system and release the resources back to the resource provider. Creating a resource can be time consuming and/or processing intensive. For example, resource creation (or recreation) for a new subscription can take hours and can cause failure due to various availability factors.

[0007] Various examples described below relate to maintaining a pool of cloud resources based on a template. Maintaining resources in a pool consistent with a template of composite resources allows for resources to be added to a subscription using the template at the cloud broker layer instead of creating resources from scratch. The life cycle of the resource objects cached in the pool can be managed based on potential reusability, such as popularity or frequency of request. A template can be a model to promote popularity and reuse. Associating a cloud resource object instance with a template that is recognizable across the cloud broker layer allows for the cloud resource object instance to be reused and avoid recreation of popular resources requested by various subscriptions.

[0008] The terms "include," "have," and variations thereof, as used herein, mean the same as the term "comprise" or appropriate variation thereof. Furthermore, the term "based on," as used herein, means "based at least in part on." Thus, a feature that is described as based on some stimulus can be based only on the stimulus or a combination of stimuli including the stimulus. Furthermore, the term "maintain" (and variations thereof) as used herein means "to create, delete, add, remove, access, update, manage, and/or modify."

[0009] Figures 1 and 2 are block diagrams depicting example cloud resource systems. Referring to figure 1, the example cloud resource system 100 of figure 1 generally includes a data store 102, a resource engine 104 and a configuration engine 106. In general, the configuration engine 106 can maintain a pool of cloud resource objects for the resource engine 104 to assign to a subscription based on a template. The example cloud resource system 100 can also include an identification engine 108, a customization engine 110, and a designer engine 112 to facilitate cloud resource assignment.

[0010] The resource engine 104 represents circuitry or any combination of circuitry and executable instructions to manage cloud resources usable by a subscription based on a template. For example, the resource engine 104 can include circuitry or a combination of circuitry and executable instructions to identify a cloud resource based on a template and assign the cloud resource based on a relationship between a subscription and the template.

[0011] A cloud resource is any resource available from a cloud, such as a compute resource, a network resource, or a storage resource. Example cloud resources include a processor resource, a memory resource, an application, a database, a network flow, a firewall, a load balancer, etc. whether virtual or physical. A template represents a composition of cloud resources. A cloud resource object instance is a specific realization of a data structure class having a composition of cloud resources. The template acts as the blueprint (i.e. construction framework) for the service by providing the classes of resources and the connections between the resources to produce a cloud resource object instance having the composition designated by the template. Thus a cloud resource object instance can be particularly configured for a particular service (e.g. as designed by the designer of the template which is described in more detail with regards to the designer engine 112). A template can represent a resource stack associated with a service.

[0012] Subscriptions can be offered by a cloud broker to a customer. A subscription represents maintaining provisioning of a service or combination of services. A cloud service represents a functionality or combination of functionalities that supply the offerings of a subscription, such as offering a cloud storage resource. Services can be offered based on a distribution model, such as software as a service ("SaaS"), platform as a service ("PaaS"), and infrastructure as a service ("IaaS"). A relationship between a subscription and a template can be based on the service, such as the distribution model of the service. A service may utilize a template or combination of templates to realize the service (e.g. allocate resources, execute applications, and offer the service). The template can be related to the subscription based on the execution environment and operating conditions to perform or otherwise offer the service.

[0013] The resource engine 104 can automate the process of caching and reassigning a cloud resource for popular or frequently used offerings by managing the cloud resource object instances of each subscription offered by a cloud broker. The cloud broker layer is a beneficial logical layer to make resource reassignment decisions to catch resources at the time a subscription ends (or updates) and when a subscription begins (or updates) before the resource is released.

[0014] The resource engine 104 can verify the pool of cloud resource object instances contains a cloud resource object instance based on a template associated with the subscription. For example, identifying a cloud resource based on a template can include verifying a cloud resource object instance exists that has the composite of the cloud resources associated with the template of the subscription. If such a cloud resource object instance exists, a specific instance can be identified for attachment to a subscription and that particular cloud resource object instance can be assigned to the subscription. If no associated cloud resource object instance exists in the pool, the resource engine can cause a cloud resource object instance to be created based on the template and cause the cloud resource object to be added to the pool for attachment to a subscription. For another example, the number of cloud resource object instances associated with a template may meet a minimum threshold and the resource engine can cause an appropriate cloud resource object instance to be created and added to the pool.

[0015] The configuration engine 106 represents circuitry or any combination of circuitry and executable instructions to maintain a pool of cloud resource object instances. A pool of cloud resource object instances is a group of cloud resource object instances that are hibernating (e.g. in a paused state of operation) or otherwise not assigned to or operating for a subscription. The configuration engine 106 can maintain the pool based on a template. For example, the cloud resource object instances can have a composition of resources associated with a template. The configuration engine can maintain the pool of cloud resources based on reusability. For example, the pool can include cloud resource object instances from frequently requested templates.

[0016] The configuration engine 106 can manage pooled resource object instances based on limitations of the cloud system. In other words, the characteristics of

the cloud system can be used to configure the cache pool. Alternatively, or in addition to, the configuration engine 106 can manage the pool based on configurations provided by an administrator. Various factors may influence the service offering of a cloud system and resource objects offered from a pool managed by the configuration engine 106. For example, the pool can parameterize the lifetime of the resource object in the pool and a minimum and/or maximum number of resource objects of the pool based on the number of offerings of a service at the cloud broker layer.

[0017] The identification engine 108 represents circuitry or any combination of circuitry and executable instructions to identify a reusability potential of a cloud resource object instance. For example, the request for a template may exceed a minimum number of requests per day and the template can be labeled as frequently used. For another example, a template may be part of a number of services offered at the cloud broker layer and can be labeled as highly reusable. Reusability potential represents the likelihood that the cloud resource object instance would be useable in a future request. Reusability potential determinations can be based on administrator input, statistical analysis, and/or default parameters. For example, a set of templates may include a label that is statically set by an administrator. In another example, the system 100 can receive static ranges, such as a minimum or maximum number of requests for a service, and/or analysis methods to determine dynamic ranges that utilize historical information from the cloud broker layer. Reusability potential can be based on request frequency, popularity, service offerings, lifetime of subscription, service level agreement ("SLA") terms, and/or other factors associated with a cloud system managed by the cloud broker and services offered by the cloud system. A label can be represented as a value, a number, a string, a Boolean, a category, a tag, or any other structure capable of identifying a status.

[0018] A cloud resource object instance can be identified as having a reusability potential based on a label associated with the template and/or metadata associated with the cloud resource object instance. For example, the data structure or metadata associated with a cloud resource object instance can inherit a label from a template, and the object resource can be identified as reusable based on the existence of the label or on a match of the label with an identified reusability for the pool, such as a pool

of popular cloud resource objects can contain cloud resource object instance with a label of popular.

[0019] The customization engine 110 represents circuitry or any combination of circuitry and executable instructions to modify the cloud resource object instance. A cloud resource object instance can be modified by flexing the resources or otherwise configuring the resources of the cloud resource object instance based on the subscription. For example, the cloud resource object instance may include a base quantity of resources that are to be doubled for a new subscription request by another customer. For another example, the applications of the cloud resource object instance can be configured with an account associated with the user of the subscription. In this manner, the cloud resource object instance can be customized based on the request of the user, such as flexing the resources of the object instance (e.g. increasing or decreasing a resource) or adding additional resources.

[0020] The designer engine 112 represents circuitry or any combination of circuitry and executable instructions to maintain a template. For example, the designer engine 112 can be a combination of circuitry and executable instructions to model the template based on a composite resource context for a subscription request and publish the template for selection by users for future subscription requests. The designer engine 112 can attach a label to the template based on administrator input or when the template achieves the threshold of requests based on the identification engine 108.

[0021] The designer engine 112 can be used to design a composite resource template as a configuration framework of a combination of infrastructure and applications. For example, the combination of infrastructure and applications can be loadable on the cloud resource stack to provide a service offered. As used herein, the term "application" refers to a set of executable instructions that provide a service when executed by a processor resource. The designer engine 112 can provide a template of resource and connections between resources for selection and/or otherwise association with a subscription based on publishing the template in a form presentable by a display communicating with an application, such as a browser.

[0022] The data store 102 can contain information utilized by the engines 104, 106, 108, 110, and 112. For example, the data store 102 can store a cloud resource object instance, a label, a template, subscription parameters, and the like.

[0023] Figure 2 depicts the example cloud resource system 200 can be implemented on a memory resource 220 operatively coupled to a processor resource 222. The processor resource 222 can be operatively coupled to a data store 202. The data store 202 can be the same as the data store 102 of figure 1.

[0024] Referring to figure 2, the memory resource 220 can contain a set of instructions that are executable by the processor resource 222. The set of instructions can implement the system 200 when executed by the processor resource 222. The set of instructions stored on the memory resource 220 can be represented as a resource module 204, a configuration module 206, an identification module 208, a customization module 210, and a designer module 212. The processor resource 222 can carry out a set of instructions to execute the modules 204, 206, 208, 210, 212, and/or any other appropriate operations among and/or associated with the modules of the system 200. For example, the processor resource 222 can carry out a set of instructions to maintain a pool of cloud resource object instances, identify a cloud resource object instance based on a composite resource template and a reusability potential of the cloud resource object instance, configure a cloud resource object instance of the pool based on a subscription request, and assign a cloud resource object instance based on the subscription request. For another example, the set of instructions can cause the processor resource 222 to maintain a label associated with the cloud resource object instance, assign the cloud resource object instance based on the label, and flex a resource of the cloud resource object instance one of up and down based on the subscription request. The resource module 204, the configuration module 206, the identification module 208, the customization module 210, and the designer module 212 represent program instructions that when executed function as the resource engine 104, the configuration engine 106, the identification engine 108, the customization engine 110, and the designer engine 112 of figure 1, respectively.

[0025] The processor resource 222 can be any appropriate circuitry capable of processing (e.g. compute) instructions. For example, the processor resource 222 can

be a central processing unit ("CPU") that enables cloud resource management by fetching, decoding, and executing modules 204, 206, 208, 210, and 212. Example processor resources 222 include CPUs, semiconductor-based microprocessors, application specific integrated circuits ("ASIC"), a field-programmable gate array ("FPGA"). The processor resource 222 can be one or multiple processing elements capable of retrieving instructions from the memory resource 220 and executing those instructions. Such multiple processing elements can be integrated in a single device or distributed across devices. The processor resource 222 can process the instructions serially, concurrently, or in partial concurrence.

[0026] The memory resource 220 and the data store 202 represent a medium to store data utilized and/or produced by the system 200. The medium can be any non-transitory medium or combination of non-transitory mediums able to electronically store data, such as modules of the system 200 and/or data used by the system 200. For example, the medium can be a storage medium, which is distinct from a transitory transmission medium, such as a signal. The medium can be machine readable, such as computer readable. The medium can be an electronic, magnetic, optical, or other physical storage device that is capable of containing (i.e. storing) executable instructions. The memory resource 220 can be said to store program instructions that when executed by the processor resource 222 implements the system 200 of figure 2. The memory resource 220 can be integrated in the same device as the processor resource 222 or it can be separate but accessible to that device and the processor resource 222. The memory resource 220 can be distributed across devices. The memory resource 220 and the data store 202 can represent the same physical medium or separate physical mediums. The data of the data store 202 can include representations of data and/or information mentioned herein.

[0027] In the discussion herein, the engines 104, 106, 108, 110, and 112 of figure 1 and the modules 204, 206, 208, 210, and 212 of figure 2 have been described as circuitry or a combination of circuitry and executable instructions. Such components can be implemented in a number of fashions. Looking at figure 2, the executable instructions can be processor executable instructions, such as program instructions, stored on the memory resource 220, which is a tangible, non-transitory computer readable storage

medium, and the circuitry can be electronic circuitry, such as processor resource 222, for executing those instructions. The instructions residing on the memory resource 220 can comprise any set of instructions to be executed directly (such as machine code) or indirectly (such as a script) by the processor resource 222.

[0028] In one example, the executable instructions can be part of an installation package that when installed can be executed by the processor resource 222 to implement the system 200. In that example, the memory resource 220 can be a portable medium such as a compact disc, a digital video disc, a flash drive, or memory maintained by a computer device, such as a service device 334 of figure 3, from which the installation package can be downloaded and installed. In another example, the executable instructions can be part of an application or applications already installed. The memory resource 220 can be a non-volatile memory resource such as read only memory ("ROM"), a volatile memory resource such as random access memory ("RAM"), a storage device, or a combination thereof. Example forms of a memory resource 220 include static RAM ("SRAM"), dynamic RAM ("DRAM"), electrically erasable programmable ROM ("EEPROM"), flash memory, or the like. The memory resource 220 can include integrated memory such as a hard drive ("HD"), a solid state drive ("SSD"), or an optical drive.

[0029] Figure 3 depicts example environments in which various example cloud resource systems can be implemented. The example environment 390 is shown to include an example system 300 for managing cloud resources. The system 300 (described herein with respect to figures 1 and 2) can represent generally any circuitry or combination of circuitry and executable instructions to reuse a cloud resource using a template. The system 300 can include a resource engine 304, a configuration engine 306, an identification 308, a customization engine 310, and a designer engine 312 that are the same as the resource engine 104, the configuration engine 106, the identification engine 108, the customization engine 110, and the designer engine 112 of figure 1, respectively, and the associated descriptions are not repeated for brevity. As shown in figure 3, the system 300 can also include a monitor engine 314 and a user interface ("UI") engine 316 and the engines 304, 306, 308, 310, 312, 314, and 316 can be integrated into a compute device, such as a service device 334. The engine 304,

306, 308, 310, 312, 314, and 316 can be integrated via circuitry or as installed instructions into a memory resource of the compute device.

[0030] The monitor engine 314 represents circuitry or a combination of circuitry and executable instructions to monitor the subscription requests. The monitor engine 314 can operate in conjunction with the resource engine 304 and the configuration engine 306 to manage the cloud resources in a dynamic manner based on the information received by the monitor engine 314. The UI engine 316 represents circuitry or a combination of circuitry and executable instructions to provide a UI for a user to make a subscription request, such as create a subscription or cancel a subscription. The system 300 can manage a pool 328 of object instances 342 having a composite of cloud resources. The composite of cloud resources (i.e. cloud resource object instance) can be designed for services offered via the service devices 334 based on subscriptions managed by a cloud broker.

[0031] The example environment 390 can include compute devices, such as administrator devices 332, service devices 334, and user devices 336. An administrator device 332 can be used by an administrator to provide input to configure the system 300. For example, a service designer can create a service design template and publish the service design template via the administrator device 332 as shown in figure 4B. For another example, a business service manager can attach a label to a template identified as popular, as shown in figure 4B, and provide threshold or methods for determining popularity via an administrator device 332 and The service devices 334 represent generally any compute devices to respond to a network request received from a user device 336, whether virtual or real. For example, the service device 334 can operate a combination of circuitry and executable instructions to provide a network packet in response to a request for a page or functionality of an application. The user devices 336 represent generally any compute devices to communicate a network request and receive and/or process the corresponding responses. For example, a browser application may be installed on the user device 336 to receive the network packet from the service device 334 and utilize the payload of the packet to display an element of a page via the browser application.

[0032] The compute devices can be located on separate networks 330 or part of the same network 330. The example environment 390 can include any appropriate number of networks 330 and any number of the networks 330 can include a cloud compute environment. A cloud compute environment includes a virtual shared pool of compute resources. For example, networks 330 can be distributed networks comprising virtual computing resources. Any appropriate combination of the system 300 and compute devices can be a virtual instance of a resource of a shared pool of resources. The engines and/or modules of the system 300 herein can reside and/or execute "on the cloud" (e.g. reside and/or execute on a virtual shared pool of resources).

[0033] A link 338 generally represents one or a combination of a cable, wireless connection, fiber optic connection, or remote connections via a telecommunications link, an infrared link, a radio frequency link, or any other connectors of systems that provide electronic communication. The link 338 can include, at least in part, intranet, the Internet, or a combination of both. The link 338 can also include intermediate proxies, routers, switches, load balancers, and the like.

[0034] Referring to figures 1-3, the engines 104, 106, 108, 110, and 112 of figure 1 and/or the modules 204, 206, 208, 210, and 212 of figure 2 can be distributed across devices 332, 334, 336, or a combination thereof. The engine and/or modules can complete or assist completion of operations performed in describing another engine and/or module. For example, the resource engine 304 of figure 3 can request, complete, or perform the methods or operations described with the resource engine 104 of figure 1 as well as the configuration engine 106, the identification engine 108, and the customization engine 110 of figure 1. Thus, although the various engines and modules are shown as separate engines in Figures 1 and 2, in other implementations, the functionality of multiple engines and/or modules may be implemented as a single engine and/or module or divided in a variety of engines and/or modules. The engines of the system 300 can perform example methods described in connection with figures 4A, 4B, and 5-8.

[0035] Figures 4A and 4B are block diagrams depicting example interactions of example cloud resource system components. Referring to figure 4A, the example blocks of figure 4A generally include a cloud portal 448, a resource engine 404, and a

configuration engine 406. The cloud portal 448 represents a web portal for access to the cloud broker to manage subscriptions. The cloud portal 448 can provide both consumer access for users to select offerings and purchase subscriptions as well as administrative access for the service designers and business service managers. As shown in figure 4B, the service designer 456 can create a service design template 446 and publishes the service design template 456 via the cloud portal 448. The business service manager 444, as shown in figure 4B, can review the service design template created by the designer, attach a label if the template is popular, and adds the published template offering with tags 446 to the cloud portal 448.

[0036] Referring back to figure 4A, the portal 448 can provide an offering via a service template a tag 446 and a cloud user 450 can access the cloud portal 448 to create, cancel, or otherwise manage a subscription associable with the offering. Based on a request from the cloud user 450, the cloud portal 448 can interact with the resource engine 404 (and/or create a new resource at block 452) to assign an object instance associated with a template of a selected offering. The resource engine 404 can interact with the configuration engine 406 to access the pool of cloud resource objects. The default behavior 454, such as flexing or other customizations to the resources, can occur according the broker principles to provide the service to the cloud user based on a subscription.

[0037] Figures 5 and 6 are flow diagrams depicting example methods of reusing a cloud resource. Referring to figure 5, example methods managing a cloud resource can generally comprise maintaining a pool of cloud resource object instances, identifying a template, identifying a cloud resource object instance of the pool, and assigning the cloud resource object instance from the pool to a subscription based on the request.

[0038] At block 502, a pool of cloud resource object instances is maintained. Each cloud resource object instance includes a plurality of cloud resources. Thus, a cloud resource object instance is a composite of cloud resources associable with a template used to describe and offer a service. The pool of cloud resource object instances can be maintained based on a template. For example, the pool can contain

object instances associated with a set of templates designated as popular and be maintained based on the set of templates.

[0039] At block 504, a template is identified based on a request for a cloud service. A service of the request can be identified and the template of resources associated with the service can be identified. The template is a composite of cloud resources particularized for the service. In this manner, a template can be compared to a cloud resource object instance. A cloud resource object instance is identified based on the template at block 506. For example, a first cloud resource object instance of a plurality of cloud resource object instances in the pool can be identified for use with a service based on a match between the first cloud resource object instance and the template for the service. At block 508, the cloud resource object instance is assigned from the pool to a subscription based on the request. For example, the request may contain the selected offering and some initial parameters that describe the configuration of the cloud resource object instance, and a cloud resource object of the pool that fits the configuration can be assigned to the subscription associated with the request.

[0040] Figure 6 includes blocks similar to blocks of figure 5 and provides additional blocks and details. In particular, figure 6 depicts additional blocks and details generally regarding configuring the pool, maintaining a label of potential reusability, and configuring the cloud resource object. Blocks 606, 610, 612, and 616 are the same as blocks 502, 504, 506, and 508 of figure 5 and, for brevity, their respective descriptions are not been repeated.

[0041] At block 604, a pool is configured based on service limitations and/or administrator input. The configuration of the pool can be determined by parameters to limitations set on the cloud broker system or otherwise enable management of the pool.

[0042] At block 608, a label of reusability associated with a template is maintained. The label can be maintained within a data structure associated with a template or based on dynamic analysis, such as a function call to determine reusability using the template and/or service as the input. The label can be based on one of popularity of the template and frequency of request associated with the service of the template. The identification of the cloud resource object instance at block 612 can be based on potential reusability using the label associated with a template.

[0043] At block 614, the cloud resource object instance can be configured based on the request. For example, the cloud resource object can be customized based on the particulars of the subscription request (e.g. creating user data or accounts or flexing a resource). For another example, additional resource may be added to the subscription, such as when a template is used as a base for a service and additional functionality can utilize another tier of infrastructure or application components.

[0044] Figure 7 depicts a flow diagram of an example user request for a subscription. At block 702, a user requests a subscription and, at block 706, selects a service offering provided by the cloud broker via the cloud portal. The service offering of the subscription may start with a base template and other templates may be added to the subscription based on the selected service offerings. The request is sent to the resource engine 704 to manage the subscription request. A determination of popularity is made at block 708. If the offering is not popular, then a collection of resources may not be in the pool and the default behavior of creating a resource for the service can occur at block 710. If the offering is identified as popular, then the pool can be checked for availability of a cloud resource object instance. If a cloud resource object instance is available in the pool, the cloud resource object instance can be assigned to the subscription at block 718. If there is not an available instance to assign to the subscription (and the service offering is popular), a new resource object instance can be caused to be created at block 714. The newly created cloud resource object instance can be stored in a pool (or a pool can be created for the new resource object instance) at block 716. The updated pool can then be used to provide the newly created resource object to the subscription.

[0045] Figure 8 depicts a flow diagram of an example subscription cancellation request. At block 802, a user requests cancellation of their subscription via the cloud portal. In response to the cancellation request, the request is sent to the resource engine 804. At block 806, the subscription is checked for whether it is associated with a popular offering. For example, the tag associated with the template of the offering may indicate a popularity level (i.e. a level of potential reusability).

[0046] If the subscription is not identified as popular, the resources of the offering are deleted at block 808. If the subscription is identified as deriving from a popular

offering, the cloud resource object instance associated with the subscription is prepared for reuse by detaching, cleaning, and placing the cloud resource object instance into a pool for reuse by another subscription. As shown in figure 8, the cloud resource object instance is detached from the subscription at block 810, the resources of the cloud resource object instance are flexed down to a base level at block 812 (e.g. an average or minimum level of a resource that may be requested), the resources of the cloud resource object instance are placed into hibernation at block 814, and the resources of the cloud resource object instance are maintain in a cache for future reuse triggered by a future subscription request at block 816, whether by the same user or another user. The cloud resource object instance can be cleaned when the instance is released from the subscription to facilitate reuse with subscriptions not associated with the user.

[0047] Cleaning a cloud resource object instance can include removing custom configurations, reverting a snapshot to a base version, hibernating a resource (e.g. pausing a resource from execution) to avoid inefficiencies, sanitizing the resources from personal data, and flexing down the resource to a common level (such as decreasing the virtual CPU and memory). By cleaning the cloud resource object instance, the next request for the cloud resource object from pool can provide the cloud resource object instance without recreating the instance and the instance is ready for flexing or other customization. These custodial activities decrease the power foot print used by the resources during the time the resources reside in the pool awaiting assignment. In this manner, the pool can be maintained with cloud resource object instance available for configuration and reuse by a subscription.

[0048] Although the flow diagrams of figures 5-8 illustrate specific orders of execution, the order of execution may differ from that which is illustrated. For example, the order of execution of the blocks may be scrambled relative to the order shown. Also, the blocks shown in succession may be executed concurrently or with partial concurrence. All such variations are within the scope of the present description.

[0049] The present description has been shown and described with reference to the foregoing examples. However, other forms, details, and examples may be made without departing from the spirit and scope of the following claims.

CLAIMS

What is claimed is:

- 1 1. A cloud resource system comprising:
2 a resource engine to:
3 identify a cloud resource based on a template comprising a composite of
4 cloud resources; and
5 assign the cloud resource based on a subscription and the template; and
6 a configuration engine to maintain a pool of cloud resource object instances, a
7 first cloud resource object instance of the cloud resource object instances having a
8 composition of resources associated with the template.
- 1 2. The system of claim 1, comprising:
2 an identification engine to identify a reusability potential of the first cloud resource
3 object instance, wherein the configuration engine is to maintain the pool of cloud
4 resource object instances based on reusability.
- 1 3. The system of claim 1, comprising:
2 a customization engine to modify the cloud resource object instance.
- 1 4. The system of claim 1, wherein the resource engine is to :
2 verify the pool of cloud resource object instances contains the first cloud
3 resource object instance based on the template;
4 cause the first cloud resource object instance to be created based on the
5 template when the pool does not contain the first cloud resource object; and
6 cause the first cloud resource object instance to be added to the pool.
- 1 5. The system of claim 1, comprising:
2 a designer engine to:
3 model the template based on a composite resource context for a subscription;
4 and

5 publish the template.

1 6. A computer readable storage medium comprising a set of instructions executable by
2 a processor resource to:

3 maintain a pool of cloud resource object instances, a cloud resource object
4 based on a composite resource template;

5 identify a cloud resource object instance based on a composite resource
6 template and a reusability potential;

7 configure a cloud resource object instance of the pool based on a subscription
8 request; and

9 assign a cloud resource object instance based on the subscription request.

1 7. The medium of claim 6, wherein the set of instructions is executable by the
2 processor resource to:

3 maintain a label associated the cloud resource object instance, the label to
4 denote the reusability potential.

1 8. The medium of claim 6, wherein the set of instructions to configure the cloud
2 resource object instance, when executed, cause the processor resource to:

3 flex a resource of the cloud resource object instance one of up and down based
4 on a service of the subscription request.

1 9. The medium of claim 6, wherein the set of instructions is executable by the
2 processor resource to:

3 design the composite resource template as a configuration framework of a
4 combination of infrastructure and applications, the combination of infrastructure and
5 applications to be loadable on the cloud resource stack to provide a service.

1 10. A method of reusing a cloud resource comprising:

2 maintaining a pool of cloud resource object instances, a cloud resource object
3 instance to include a plurality of cloud resources;

4 identifying a template based on a request for a cloud service, the template
5 comprising a composite of the plurality of cloud resources,
6 identifying a cloud resource object instance based on the template; and
7 assigning the cloud resource object instance from the pool to a subscription
8 based on the request.

1 11. The method of claim 10, comprising:

2 maintaining a label of reusability associated with the template.

1 12. The method of claim 11, wherein the label is based on one of popularity of the

2 template and frequency of requests associated with the template and identifying the
3 cloud resource object instance is based on a potential reusability.

1 13. The method of claim 10, comprising:

2 configuring the cloud resource object instance based on the request; and

3 adding additional resources to the subscription.

1 14. The method of claim 10, comprising:

2 configuring the pool based on one of service limitations and administrator input.

1 15. The method of claim 10, comprising, in response to a subscription cancellation,:

2 identifying the cloud resource object instance is derived from a popular offering;

3 detaching the cloud resource object instance from the subscription;

4 cleaning the cloud resource object instance when the cloud resource object

5 instance is released from the subscription; and

6 placing the cloud resource object instance into the pool.




FIG. 1

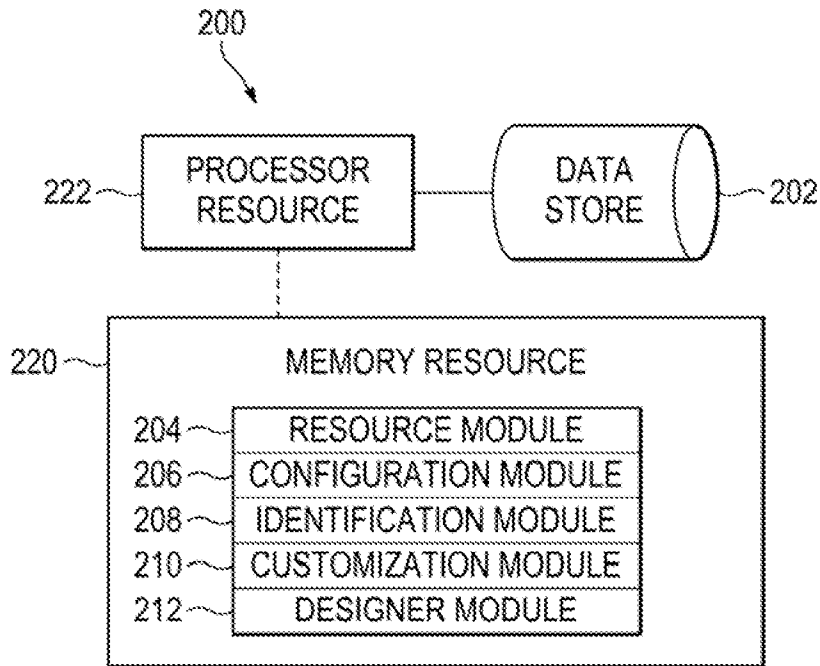


FIG. 2

2/6




FIG. 3

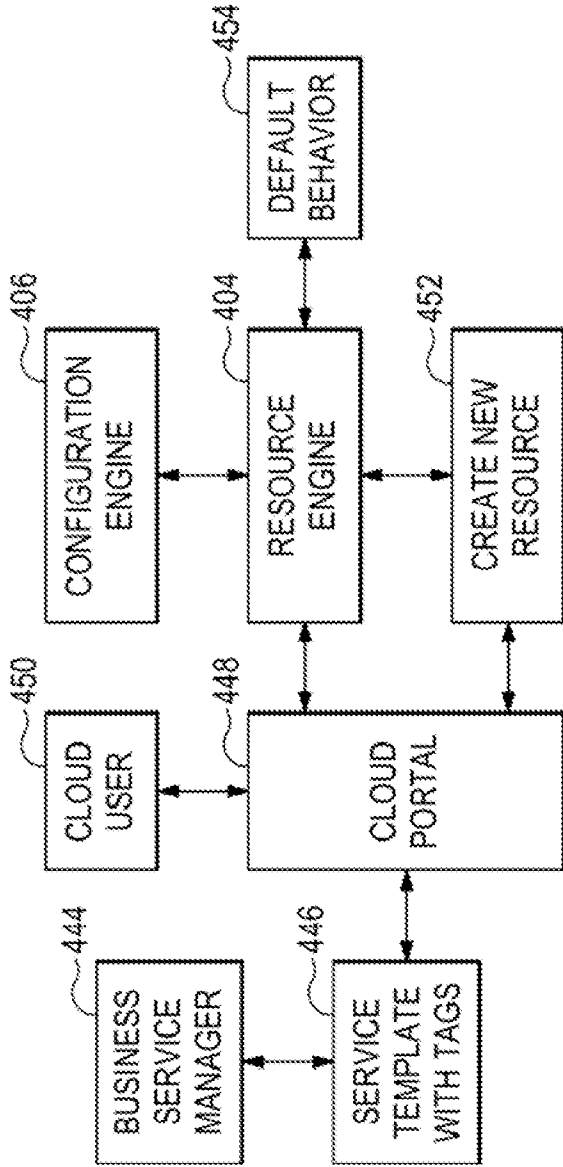


FIG. 4A




FIG. 4B

4/6

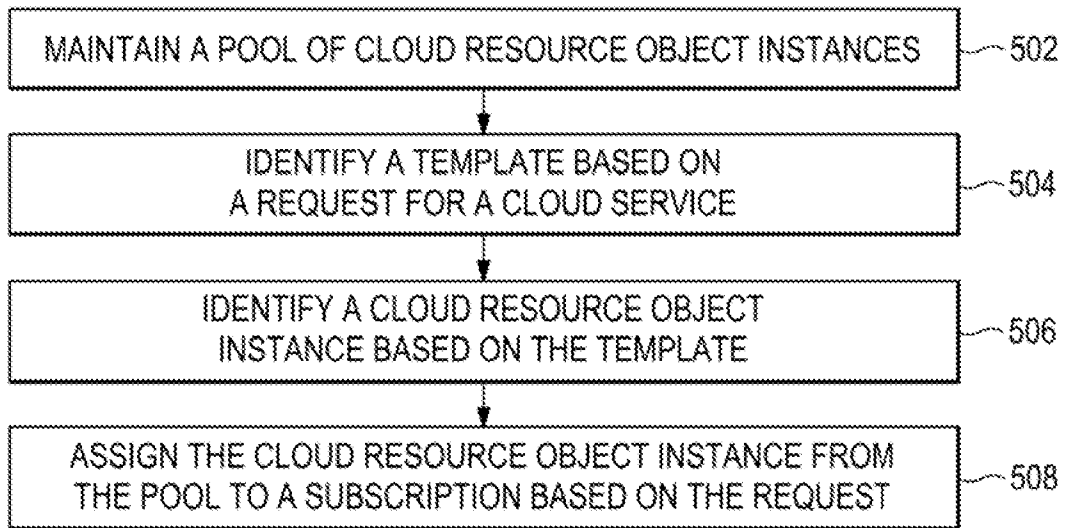


FIG. 5

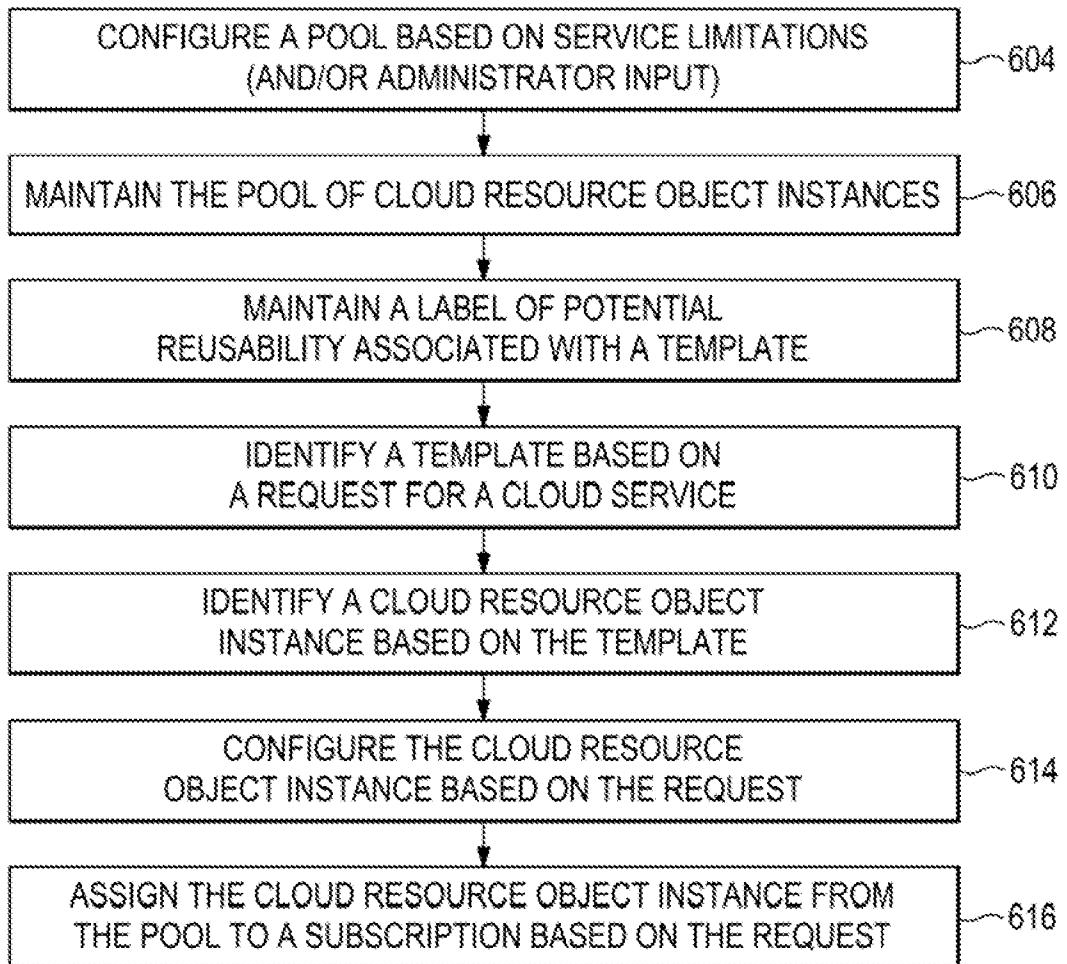


FIG. 6

5/6

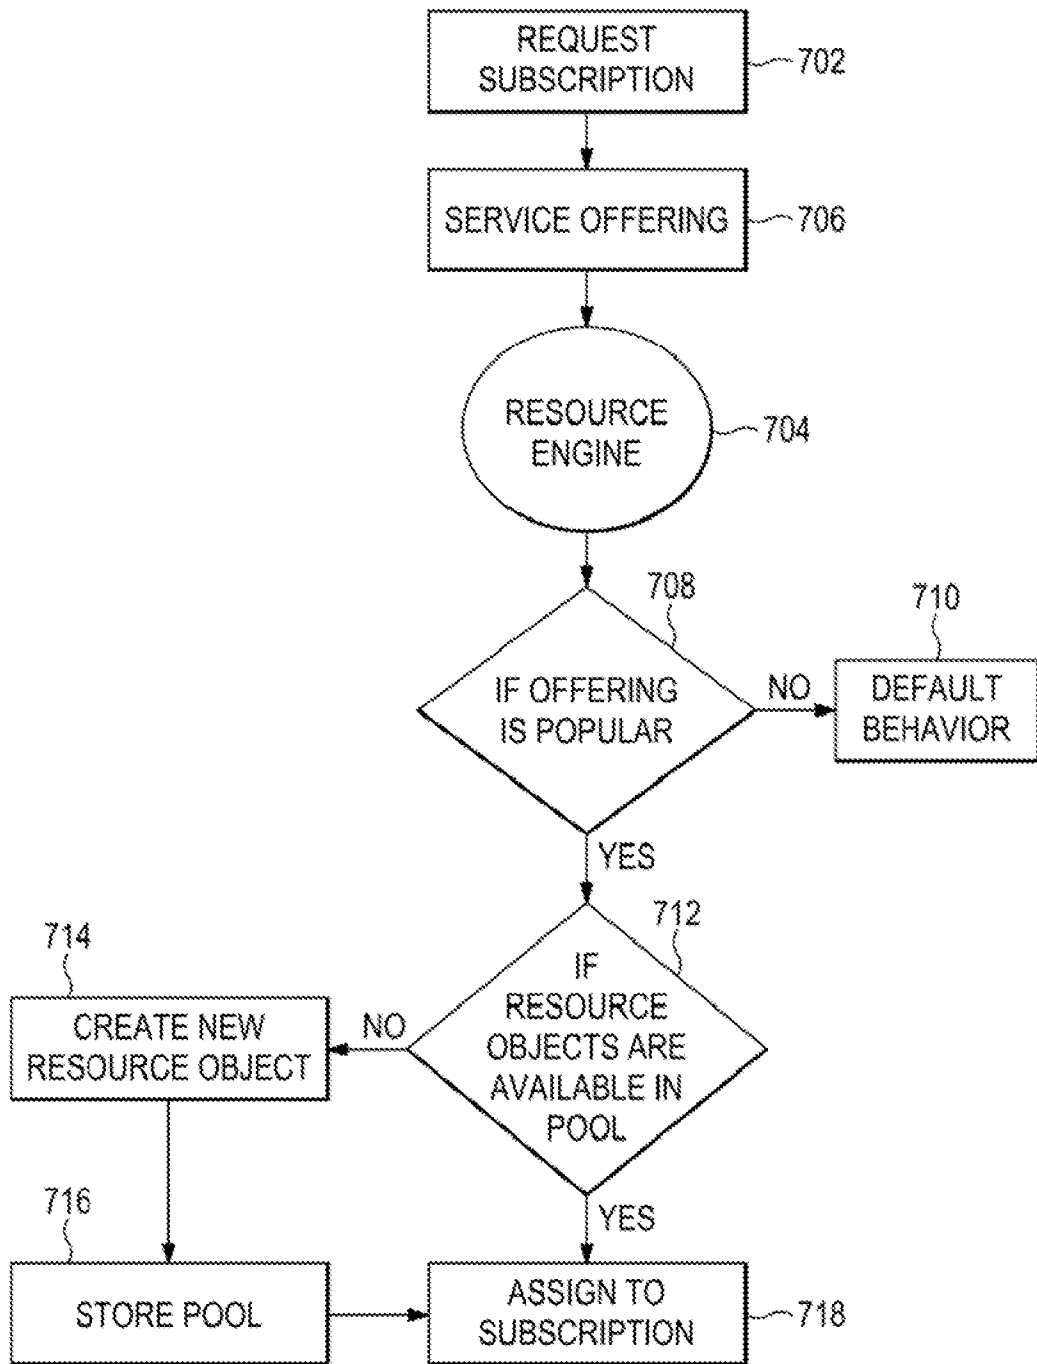


FIG. 7

6/6

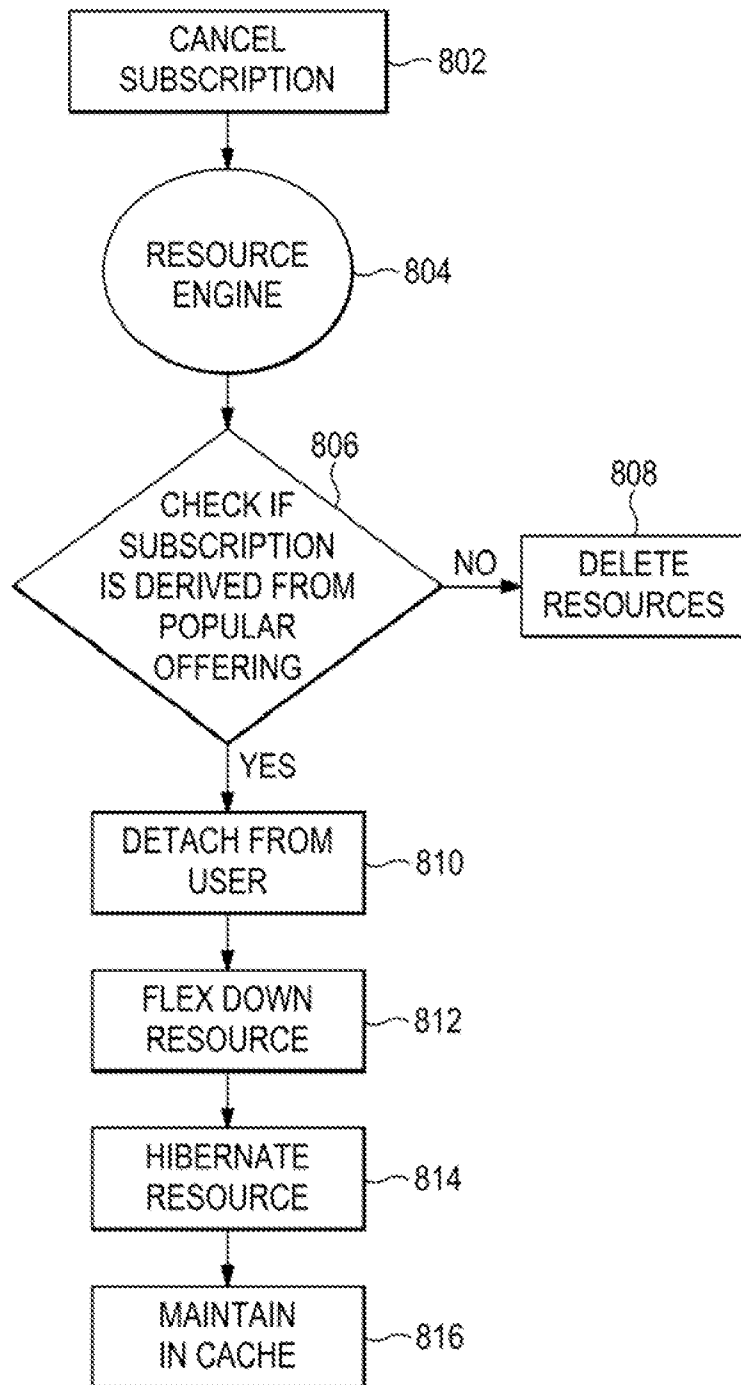


FIG. 8

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2014/054998**A. CLASSIFICATION OF SUBJECT MATTER****G06F 9/50(2006.01)i, G06F 15/16(2006.01)i**

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

B. FIELDS SEARCHEDMinimum documentation searched (classification system followed by classification symbols)
G06F 9/50; G06F 15/16; G06F 15/173; G06F 9/46; G06F 9/455Documentation 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 modelsElectronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & Keywords: cloud, resource, pool, allocation, reuse, template**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 8495197 B1 (NAGARGADDE, A. et al.) 23 July 2013 See abstract, col. 9, line 4 - line 10, line 47, claim 7, and figs. 2 and 3.	1-15
Y	US 2013-0326510 A1 (ADEKILE, O. et al.) 05 December 2013 See paragraphs [0054]-[0067], claim 1, and figs. 4 and 5.	1-15
A	US 7007075 B1 (COFFEY, J.) 28 February 2006 See abstract, col. 3, line 39 - col. 4, line 36, and fig. 1.	1-15
A	US 2013-0179574 A1 (CALDER, B. G. et al.) 11 July 2013 See abstract, paragraphs [0073]-[0078], and fig. 13.	1-15
A	US 2014-0059226 A1 (MESSERLI, A. J. et al.) 27 February 2014 See abstract, paragraphs [0031]-[0041], and fig. 1.	1-15

 Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

20 April 2015 (20.04.2015)

Date of mailing of the international search report

21 April 2015 (21.04.2015)

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

YU, Jintae

Telephone No. +82-42-481-8530



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2014/054998

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 8495197 B1	23/07/2013	None	
US 2013-0326510 A1	05/12/2013	None	
US 7007075 B1	28/02/2006	AU 1999-27903 A1 US 2002-0016768 A1 US 2004-0225617 A1 WO 99-44165 A1	15/09/1999 07/02/2002 11/11/2004 02/09/1999
US 2013-0179574 A1	11/07/2013	CN 104040526 A EP 2802998 A1 JP 2015-507281 A KR 10-2014-0111671 A US 8904008 B2 WO 2013-106257 A1	10/09/2014 19/11/2014 05/03/2015 19/09/2014 02/12/2014 18/07/2013
US 2014-0059226 A1	27/02/2014	WO 2014-031473 A2 WO 2014-031473 A3	27/02/2014 26/02/2015