

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

(19) World Intellectual Property Organization
International Bureau



(10) International Publication Number

WO 2016/049576 A1

(43) International Publication Date
31 March 2016 (31.03.2016)

(51) International Patent Classification:
G06F 9/50 (2006.01)

(21) International Application Number:
PCT/US2015/052459

(22) International Filing Date:
25 September 2015 (25.09.2015)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/055,575 25 September 2014 (25.09.2014) US
62/055,577 25 September 2014 (25.09.2014) US
62/055,856 26 September 2014 (26.09.2014) US
62/056,427 26 September 2014 (26.09.2014) US
14/864,723 24 September 2015 (24.09.2015) US

(71) Applicant: **ORACLE INTERNATIONAL CORPORATION** [US/US]; 500 Oracle Parkway, M/S 5op7, Redwood Shores, California 94065 (US).

(72) Inventors: **MORDANI, Rajiv**; 1038 Linz Terrace, Sunnyvale, California 94089 (US). **ISLAM, Nazrul**; 1555 Vista Club Circle, Apt. 202, Santa Clara, California 95054 (US).

(74) Agents: **MEYER, Sheldon, R.** et al.; Tucker Ellis LLP, One Market Plaza, Steuart Tower, Suite 700, San Francisco, California 94105 (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, 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, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

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

(54) Title: SYSTEM AND METHOD FOR USE OF A GLOBAL RUNTIME IN A MULTITENANT APPLICATION SERVER ENVIRONMENT




FIGURE 17

(57) Abstract: In accordance with an embodiment, described herein is a system and method for providing an end-to-end lifecycle in a multi-tenant application server environment. A lifecycle manager enables configuration integration across different components, for example application servers, traffic directors or load balancers, and database or other components, which may have their own partition-specific configurations. In accordance with an embodiment, the end-to-end lifecycle infrastructure includes a plugin-based mechanism which enables the different components within a multi-tenant application server environment to use a lifecycle plugin to provide their own component-specific lifecycle functionality. In accordance with an embodiment, the system includes support for a tenant onboarding functionality, which can be used to onboard a tenant for a service to a particular partition. In accordance with an embodiment, the system supports the use of a global runtime within the multi-tenant application server environment.

**SYSTEM AND METHOD FOR USE OF A GLOBAL RUNTIME IN A
MULTITENANT APPLICATION SERVER ENVIRONMENT**

5

COPYRIGHT NOTICE

10

A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

Claim of Priority and Cross-Reference to Related Applications:

[0001] This application claims the benefit of priority to U.S. Provisional Applications "SYSTEM AND METHOD FOR PROVIDING END-TO-END LIFECYCLE IN A MULTITENANT APPLICATION SERVER ENVIRONMENT" Application No. 62/055,575, filed September 25, 2014; "SYSTEM AND METHOD FOR TENANT ONBOARDING IN A MULTITENANT APPLICATION SERVER ENVIRONMENT", Application No. 62/055,577, filed September 25, 2014; "SYSTEM AND METHOD FOR DATABASE AND APPLICATION SERVER LIFECYCLE PLUGINS IN A MULTITENANT APPLICATION SERVER ENVIRONMENT", Application No. 62/055,856, filed September 26, 2014; and "SYSTEM AND METHOD FOR GLOBAL RUNTIME IN A MULTITENANT APPLICATION SERVER ENVIRONMENT", Application No. 62/056,427, filed September 26, 2014; and is related to U.S. Patent Application titled "SYSTEM AND METHOD FOR SUPPORTING MULTITENANCY IN AN APPLICATION SERVER, CLOUD, OR OTHER ENVIRONMENT", Application No. 14/601,883, filed January 21, 2015; each of which above applications are herein incorporated by reference.

Field of Invention:

[0002] Embodiments of the invention are generally related to application servers and cloud environments, and are particularly related to systems and methods for providing an end-to-end lifecycle, tenant onboarding, or global runtime, in a multitenant application server environment.

35

Background:

[0003] Application servers generally provide a managed computing environment within which software applications can be deployed and run. Cloud-based computing environments allow applications to run within and take advantage of the distributed resources provided by a cloud. Such environments can support large numbers of users or tenants, some of which may have specific requirements that are particular to that user or tenant.

When additional tenants are added to the system, the system must be able to address each tenant's needs, including associating other system components with the tenant as necessary. These are some examples of the type of environments in which embodiments of the present invention can be used.

5

Summary:

[0004] In accordance with an embodiment, described herein is a system and method for providing an end-to-end lifecycle in a multi-tenant application server environment. A lifecycle manager enables configuration integration across different components, for example application servers, traffic directors or load balancers, and database or other components, which may have their own partition-specific configurations. In accordance with an embodiment, the end-to-end lifecycle infrastructure includes a plugin-based mechanism which enables the different components within a multi-tenant application server environment to use a lifecycle plugin to provide their own component-specific lifecycle functionality. In accordance with an embodiment, the system includes support for a tenant onboarding functionality, which can be used to onboard a tenant for a service to a particular partition. In accordance with an embodiment, the system supports the use of a global runtime within the multi-tenant application server environment.

20 **Brief Description of the Figures:**

[0005] **Figure 1** illustrates a multitenant application server environment, in accordance with an embodiment.

[0006] **Figure 2** further illustrates a multitenant application server environment, in accordance with an embodiment.

25 [0007] **Figure 3** further illustrates a multitenant application server environment, in accordance with an embodiment.

[0008] **Figure 4** further illustrates a multitenant application server environment, in accordance with an embodiment.

30 [0009] **Figure 5** further illustrates a multitenant application server environment, in accordance with an embodiment.

[00010] **Figure 6** illustrates an end-to-end lifecycle in a multitenant application server environment, in accordance with an embodiment.

[00011] **Figure 7** illustrates an end-to-end lifecycle event diagram in a multitenant application server environment, in accordance with an embodiment.

35 [00012] **Figure 8** illustrates an end-to-end lifecycle topology in a multitenant application server environment, in accordance with an embodiment.

[00013] **Figure 9** illustrates another end-to-end lifecycle topology in a multitenant

application server environment, in accordance with an embodiment.

[00014] **Figure 10** illustrates a flowchart of method for providing end-to-end lifecycle in a multitenant application server environment, in accordance with an embodiment.

[00015] **Figure 11** illustrates tenant onboarding in a multitenant application server

5 environment, in accordance with an embodiment.

[00016] **Figure 12** further illustrates tenant onboarding in a multitenant application server environment, in accordance with an embodiment.

[00017] **Figure 13** illustrates a flowchart of method for providing tenant onboarding in a multitenant application server environment, in accordance with an embodiment.

10 [00018] **Figure 14** illustrates lifecycle plugins in a multitenant application server environment, in accordance with an embodiment.

[00019] **Figure 15** further illustrates lifecycle plugins in a multitenant application server environment, in accordance with an embodiment.

15 [00020] **Figure 16** illustrates a flowchart of method for supporting lifecycle plugins in a multitenant application server environment, in accordance with an embodiment.

[00021] **Figure 17** illustrates use of a global runtime in a multitenant application server environment, in accordance with an embodiment.

[00022] **Figure 18** further illustrates use of a global runtime in a multitenant application server environment, in accordance with an embodiment.

20 [00023] **Figure 19** illustrates a flowchart of method for supporting a global runtime in a multitenant application server environment, in accordance with an embodiment.

Detailed Description:

[00024] In accordance with an embodiment, described herein is a system and method

25 for providing an end-to-end lifecycle in a multi-tenant application server environment. A lifecycle manager enables configuration integration across different components, for example application servers, traffic directors or load balancers, and database or other components, which may have their own partition-specific configurations.

[00025] In accordance with an embodiment, the end-to-end lifecycle infrastructure

30 includes a plugin-based mechanism which enables the different components within a multi-tenant application server environment to use a lifecycle plugin to provide their own component-specific lifecycle functionality.

[00026] In accordance with an embodiment, the system includes support for a tenant

onboarding functionality, which can be used to onboard a tenant for a service to a particular

35 partition.

[00027] In accordance with an embodiment, the system supports the use of a global

runtime within the multi-tenant application server environment.

Application Server (e.g., Multi-Tenant, MT) Environment

[00028] **Figure 1** illustrates a system for supporting multi-tenancy in an application server, cloud, or other environment, in accordance with an embodiment.

[00029] As illustrated in Figure 1, in accordance with an embodiment, an application server (e.g., multi-tenant, MT) environment 100, or other computing environment which enables the deployment and execution of software applications, can be configured to include and operate according to a domain 102 configuration that is used at runtime to define an application server domain.

[00030] In accordance with an embodiment, the application server can include one or more partitions 104 that are defined for use at runtime. Each partition can be associated with a globally unique partition identifier (ID) and partition configuration, and can further include one or more resource groups 124, together with a reference to a resource group template 126 and/or partition-specific applications or resources 128. Domain-level resource groups, applications and/or resources 140 can also be defined at the domain level, optionally with a reference to a resource group template.

[00031] Each resource group template 160 can define one or more applications A 162, B 164, resources A 166, B 168, and/or other deployable applications or resources 170, and can be referenced by a resource group. For example, as illustrated in Figure 1, resource group 124 in partition 104 can reference 190 resource group template 160.

[00032] Generally, a system administrator can define partitions, domain-level resource groups and resource group templates, and security realms; while a partition administrator can define aspects of their own partition, for example, by creating partition-level resource groups, deploying applications to the partition, or referencing specific realms for the partition.

[00033] **Figure 2** further illustrates a system for supporting multi-tenancy in an application server, cloud, or other environment, in accordance with an embodiment.

[00034] As illustrated in Figure 2, in accordance with an embodiment, a partition 202 can include, for example, a resource group 205 which includes a reference 206 to a resource group template 210, a virtual target (e.g., virtual host) information 207, and a pluggable database (PDB) information 208. A resource group template (e.g., 210) can define, for example, a plurality of applications A 211 and B 212, together with resources such as a Java Message Server (JMS) server 213, store-and-forward (SAF) agent 215, mail session component 216, or Java Database Connectivity (JDBC) resource 217.

[00035] The resource group template illustrated in Figure 2 is provided by way of example; in accordance with other embodiments, different types of resource group templates and elements can be provided.

[00036] In accordance with an embodiment, when a resource group within a partition (e.g., 202) references 220 a particular resource group template (e.g., 210), information

associated with a particular partition can be used in combination with the referenced resource group template, to indicate a partition-specific information 230, for example a partition-specific PDB information. The partition-specific information can then be used by the application server to configure resources, for example a PDB resource, for use by the partition. For example, partition-specific PDB information associated with partition 202 can be used, by the application server, to configure 232 a container database (CDB) 236 with an appropriate PDB 238, for use by that partition.

[00037] Similarly, in accordance with an embodiment, a virtual target information associated with a particular partition can be used to define 239 a partition-specific virtual target 240, for use by the partition, e.g., baylandurgentcare.com, which can then be made accessible via a uniform resource locator (URL), e.g., <http://baylandurgentcare.com>.

[00038] **Figure 3** further illustrates a system for supporting multi-tenancy in an application server, cloud, or other environment, in accordance with an embodiment.

[00039] In accordance with an embodiment, a system configuration such as a config.xml configuration file, is used to define a partition, including configuration elements for resource groups associated with that partition, and/or other partition properties. Values can be specified per-partition using property name/value pairs.

[00040] In accordance with an embodiment, a plurality of partitions can be executed within a managed server / cluster 242, or a similar environment which can provide access to a CDB 243, and which are accessible via a web tier 244. This allows, for example, a domain or partition to be associated with one or more of the PDBs (of the CDB).

[00041] In accordance with an embodiment, each of the plurality of partitions, in this example partition A 250 and partition B 260, can be configured to include a plurality of resources associated with that partition. For example, partition A can be configured to include a resource group 251 which contains an application A1 252, application A2 254, and JMS A 256, together with a datasource A 257 associated with PDB A 259, wherein the partition is accessible via a virtual target A 258. Similarly, partition B 260 can be configured to include a resource group 261 which contains an application B1 262, application B2 264, and JMS B 266, together with a datasource B 267 associated with PDB B 269, wherein the partition is accessible via a virtual target B 268.

[00042] While several of the above examples illustrate use of CDB and PDBs, in accordance with other embodiments, other types of multi-tenant or non-multi-tenant databases can be supported, wherein a particular configuration can be provided for each partition, for example through the use of schemas, or the use of different databases.

35

Resources

[00043] In accordance with an embodiment, a resource is a system resource,

application, or other resource or object that can be deployed to a domain of the environment. For example, in accordance with an embodiment, a resource can be an application, JMS, JDBC, JavaMail, WLDF, data source, or other system resource or other type of object that can be deployed to a server, cluster, or other application server target.

5

Partitions

[00044] In accordance with an embodiment, a partition is a runtime and administrative subdivision or slice of a domain that can be associated with a partition identifier (ID) and configuration, and can contain applications and/or refer to domain-wide resources through the use of resource groups and resource group templates.

10

[00045] Generally, a partition can contain its own applications, refer to domain wide applications via resource group templates, and have its own configuration. Partitionable entities can include resources, for example JMS, JDBC, JavaMail, WLDF resources, and other components, such as JNDI namespace, network traffic, work managers, and security policies and realms. In the context of a multi-tenant environment, the system can be configured to provide tenant access to the administrative and runtime aspects of partitions associated with a tenant.

15

[00046] In accordance with an embodiment, each resource group within a partition can optionally reference a resource group template. A partition can have multiple resource groups, and each of them can reference a resource group template. Each partition can define properties for configuration data that is not specified in the resource group templates to which the partition's resource groups refer. This enables the partition to act as a binding of deployable resources defined in a resource group template, to specific values for use with that partition. In some cases, a partition can override configuration information specified by a resource group template.

20

[00047] In accordance with an embodiment, a partition configuration, as defined for example by a config.xml configuration file, can include a plurality of configuration elements, for example: "partition", which contains the attributes and child elements that define the partition; "resource-group", which contains the applications and resources deployed to the partition; "resource-group-template", which contains applications and resources defined by that template; "jdbc-system-resource-override", which contains a database-specific service name, user name, and password; and "partition-properties", which contains property key values that can be used for macro replacement in resource group templates.

30

[00048] Upon startup, the system can use the information provided by the configuration file to generate partition-specific configuration elements for each resource, from the resource group template.

Resource Groups

[00049] In accordance with an embodiment, a resource group is a named, fully-qualified collection of deployable resources that can be defined either at a domain or partition level, and can reference a resource group template. The resources in a resource group are considered fully-qualified in that the administrator has provided all of the information needed to start or connect to those resources, for example the credentials for connecting to a data source, or the targeting information for an application.

[00050] A system administrator can declare resource groups at the domain level, or at the partition level. At the domain level, a resource group provides a convenient way to group related resources. The system can manage resources declared in a domain-level resource group the same as ungrouped resources, so that the resources can be started during system start-up, and stopped during system shut-down. An administrator can also stop, start, or remove a resource in a group individually, and can act on all the resources in a group implicitly by operating on the group. For example, stopping a resource group stops all of the resources in the group that are not already stopped; starting the resource group starts any resources in the group that are not already started; and removing the resource group removes all of the resources contained in the group.

[00051] At the partition level, a system or partition administrator can configure zero or more resource groups in a partition, subject to any security restrictions. For example, in a SaaS use case, various partition-level resource groups can refer to domain-level resource group templates; while in a PaaS use case, partition-level resource groups can be created that do not refer to resource group templates, but instead represent applications and their related resources that are to be made available within that partition only.

[00052] In accordance with an embodiment, resource grouping can be used to group together applications and the resources they use as a distinct administrative unit within the domain. For example, in the medical records (MedRec) application described below, a resource grouping defines the MedRec application and its resources. Multiple partitions can run the same MedRec resource group, each using a partition-specific configuration information, such that the applications that are part of each MedRec instance are made specific to each partition.

Resource Group Templates

[00053] In accordance with an embodiment, a resource group template is a collection of deployable resources that are defined at a domain level, that can be referenced from a resource group, and some of the information required to activate its resources may not be stored as part of the template itself, such that it supports the specification of partition level configuration. A domain can contain any number of resource group templates, each of which

can include, for example, one or more related Java applications and the resources on which those applications depend. Some of the information about such resources may be the same across all partitions, while other information may vary from one partition to the next. Not all configuration need be specified at the domain level – partition level configuration can instead be specified in the resource group template through the use of macros, or property name/value pairs.

[00054] In accordance with an embodiment, a particular resource group template can be referenced by one or more resource groups. Generally, within any given partition, a resource group template can be referenced by one resource group at a time, i.e., not simultaneously by multiple resource groups within the same partition; however, it can be referenced at the same time by another resource group in a different partition. The object containing the resource group, e.g., the domain or partition, can use property name/value assignments to set the value of any tokens in the resource group template. When the system activates a resource group template using a referencing resource group, it can replace those tokens with values set in the resource group's containing object. In some cases, the system can also use statically-configured resource group templates and partitions to generate runtime configuration for each partition/template combination.

[00055] For example, in a SaaS use case, the system can activate the same applications and resources multiple times, including once for each partition that will use them. When an administrator defines a resource group template they can use tokens to represent the information that will be supplied elsewhere. For example, the username to use in connecting to a CRM-related data resource can be indicated in the resource group template as \\${CRMDataUsername}.

25 **Tenants**

[00056] In accordance with an embodiment, in a multi-tenant environment, such as a multi-tenant (MT) application server environment, a tenant is an entity that can be represented by, or otherwise associated with, one or more partitions and/or one or more tenant-aware applications.

[00057] For example, tenants can represent distinct user organizations, such as different external companies, or different departments within a particular enterprise (e.g., HR and Finance departments), each of which can be associated with a different partition. A tenant globally unique identity (tenant ID) is the association of a particular user, at a particular moment in time, with a particular tenant. The system can derive which tenant a particular user belongs to from the user identity, for example by referring to a user identity store. The user identity enables the system to enforce those actions that a user is authorized to perform, including, but not limited to, which tenant the user may belong.

[00058] In accordance with an embodiment, the system enables isolation of the administration and runtime of different tenants from each other. For example, tenants can configure some behaviors of their applications, and resources to which they have access. The system can ensure that a particular tenant cannot administer artifacts belonging to another tenant; and, at runtime, that the applications working on behalf of a particular tenant refer only to resources associated with that tenant, and not to resources associated with other tenants.

[00059] In accordance with an embodiment, a tenant-unaware application is one that contains no logic dealing with tenants explicitly, such that any resources that the application uses may be accessible regardless of what user submitted a request to which the application is responding. In contrast, a tenant-aware application includes logic that explicitly deals with tenants. For example, based on a user's identity the application can derive the tenant to which the user belongs and use that information to access tenant-specific resources.

[00060] In accordance with an embodiment, the system enables users to deploy applications that are explicitly written to be tenant-aware, so that application developers can obtain the tenant ID of a current tenant. The tenant-aware application can then use the tenant ID to handle multiple tenants that are using a single instance of the application.

[00061] For example, the MedRec application, which supports a single doctor's office or hospital, can be exposed to two different partitions or tenants, e.g., a Bayland Urgent Care tenant, and a Valley Health tenant, each of which is able to access separate tenant-specific resources, such as separate PDBs, without changing the underlying application code.

Exemplary Domain Configuration and Multi-Tenant Environment

[00062] In accordance with an embodiment, applications can be deployed to a resource group template at the domain level, or to a resource group that is scoped to a partition or scoped to the domain. Application configuration can be overridden using deployment plans specified per-application, or per-partition. Deployment plans can also be specified as part of a resource group.

[00063] **Figure 4** illustrates a domain configuration for use with an exemplary multi-tenant environment, in accordance with an embodiment.

[00064] In accordance with an embodiment, when the system starts a partition, it creates virtual targets (e.g., virtual hosts) and connection pools, including one for each partition, to respective database instances, according to the provided configuration.

[00065] Typically, each resource group template can include one or more related applications and the resources on which those applications depend. Each partition can provide configuration data that is not specified in the resource group templates to which it refers, by providing a binding of the deployable resources in resource group templates to

specific values associated with the partition; including, in some cases, overriding certain configuration information specified by the resource group template. This enables the system to activate an application represented by a resource group template differently for each partition, using the property values each partition has defined.

5 [00066] In some instances, a partition may contain resource groups that do not refer to resource group templates, or that directly define their own partition-scoped deployable resources. Applications and data sources that are defined within a partition are generally available only to that partition. Resources can be deployed so that they can be accessed from across partitions using partition:<partitionName>/<resource JNDI name>, or
10 domain:<resource JNDI name>.

[00067] For example, a MedRec application can include a plurality of Java applications, a data source, a JMS server, and a mail session. To run the MedRec application for multiple tenants, the system administrator can define a single MedRec resource group template 286, declaring those deployable resources in the template.

15 [00068] In contrast to domain-level deployable resources, the deployable resources declared in a resource group template may not be fully configured in the template, or cannot be activated as-is, since they lack some configuration information.

[00069] For example, the MedRec resource group template may declare a data source used by the applications, but it may not specify a URL for connecting to the database.

20 Partitions associated with different tenants, for example, partition BUC-A 290 (Bayland Urgent Care, BUC) and partition VH-A 292 (Valley Health, VH) can reference one or more resource group templates, by each including a MedRec resource group 293, 294 that references 296, 297 the MedRec resource group template. The reference can then be used to create 302, 306, the virtual targets / virtual hosts for each tenant, including a virtual host
25 baylandurgentcare.com 304 associated with the BUC-A partition, for use by the Bayland Urgent Care tenant; and a virtual host valleyhealth.com 308 associated with the VH-A partition, for use by the Valley Health tenant.

[00070] **Figure 5** further illustrates an exemplary multi-tenant environment, in accordance with an embodiment. As illustrated in Figure 5, and continuing the example from 30 above, in which two partitions reference the MedRec resource group template, in accordance with an embodiment, a servlet engine 310 can be used to support a plurality of tenant environments, in this example a Bayland Urgent Care Physician tenant environment 320, and a Valley Health Physician tenant environment 330.

[00071] In accordance with an embodiment, each partition 321, 331 can define a 35 different virtual target on which to accept incoming traffic for that tenant environment, and a different URL 322, 332 for connecting to the partition and to its resources 324, 334, including in this example either a bayland urgent care database, or a valley health database

respectively. The database instances can use compatible schemas, since the same application code will execute against both databases. When the system starts the partitions, it can create the virtual targets and connection pools to the respective database instances.

5

* * * * *

End-to-End Lifecycle

[00072] In accordance with an embodiment, described herein is a system and method for providing an end-to-end lifecycle in a multi-tenant application server environment. A 10 lifecycle manager enables configuration integration across different components, for example application servers, traffic directors or load balancers, and database or other components, which may have their own partition-specific configurations.

[00073] In accordance with various embodiments, some of the descriptions of the end-to-end lifecycle process that are provided herein use the following terms:

15 **[00074]** Tenancy: a right to use a component of the system, typically obtained for a limited duration by a tenant, to allow that tenant to customize the behavior of the system to a certain extent to suit their needs.

20 **[00075]** Tenancy context: contextual information, for example security, naming, or class loading information, which is used to determine the tenancy of a system, so that the system can then use that tenancy context to provide tenant-specific behavior, such as generating bills.

[00076] Partition context: a partition identifier (ID) associated with a current thread of execution, which the system can use to determine platform tenancy.

25 **[00077]** Resource tenancy: a similar property as "tenancy", used primarily when describing Oracle Fusion Middleware (FMW) / Fusion Application (FA) embodiments.

[00078] MT admin code: code that is deployed and activated before any tenant is onboarded for a multi-tenant system, and which is responsible for lifecycle management of the system as well as tenant runtimes.

[00079] MT admin runtime: a runtime representation of the MT admin code.

30 **[00080]** Tenant code: code that is specifically deployed for a tenant.

[00081] Tenant runtime: a runtime representation of the tenant code.

[00082] Layered products: software products which can interoperate with or otherwise utilize a multi-tenant application server environment, examples of which can include the FMW and FA families of products.

35 **[00083]** **Figure 6** illustrates an end-to-end lifecycle in a multitenant application server environment, in accordance with an embodiment.

[00084] As illustrated in Figure 6, in accordance with an embodiment, the system can

include a lifecycle coordination framework 400, including a lifecycle manager (LCM) 410 and lifecycle interface 411, which enables lifecycle management 412 of partition and other components, for example a traffic director component 414, database component 415, or other types of components 416, for use within the multitenant application server environment.

5 [00085] For example, in addition to enabling an application server (e.g., a WebLogic, WLS) component to be multitenant-aware, and to support the use of partitions within the application server, the system must be able to ensure that other components and products can interoperate with the application server partitions, for example that database (e.g., Oracle Database, DB), and traffic director or load balancer (e.g., Oracle Traffic Director, 10 OTD) instances can be properly configured and can thereafter similarly support those partitions.

15 [00086] To address this, in accordance with an embodiment, as part of creating a partition for use with an application server, then, in addition to creating the application server partition, the system allows creation of a slice of runtime in other components that will ultimately be wired up with the application server partition. The result is a system in which an application can be entirely run within its own environment from, e.g., an OTD layer, down to a database layer.

20 [00087] In accordance with an embodiment, the end-to-end lifecycle infrastructure includes a plugin-based mechanism which enables a plurality of different components within a multi-tenant application server environment to use a lifecycle plugin to provide their own component-specific lifecycle functionality. Any component which will be made available to be wired up with a partition can provide its own plugin. Lifecycle operations can then call the plugin interface, which in turn performs the actual creation of the component – for example, the creation of a traffic director instance, or the creation of a database instance, as 25 appropriate.

[00088] In accordance with an embodiment, in addition to creating application server and/or component partitions, each tenant can be associated with an environment, as further described below, wherein the environment acts as a container or container object within which the tenant of a partition can run their own applications.

30 [00089] For example, in accordance with an embodiment, the lifecycle manager and its pluggable extensible mechanism can be used to support the creation of components within the middleware or database stack, and to also support the wiring of those components, such as wiring up an application server (e.g., WLS) partition to use a datasource. The collection of wired components can then be provided within a partition or 35 tenant environment.

[00090] As another example, in accordance with an embodiment, the system can enable an application server partition to be associated with a database, and created during a

lifecycle, which then creates a datasource in the application server partition and makes it usable within a partition or tenant environment.

[00091] Similarly, in accordance with an embodiment, the lifecycle manager can be used to set up OTD routing for a partition, including setting up a URI or host name for use with the partition, such that a WLS partition running in a cluster is front-ended with an OTD instance, while the lifecycle manager enables the creation of an OTD instance for the partition, and its wiring to the WLS partition, for use within a tenant environment.

[00092] Additional advantages provided by the end-to-end lifecycle management include the ability to detect changes out-of-band.

[00093] For example, an update to an application server partition can be sent to other components that are wired to that partition. As another example, when a new resource group is associated with a new target, that updated information can be sent to other layered products, such that the changes in the partition are propagated to any other wired components.

[00094] In accordance with an embodiment, orchestration of the lifecycle can be performed in a variety of ways to create the various components or slices of runtime. For example, the system can individually create a partition, PDB, and OTD instance as REST calls to the lifecycle manager, which directs the lifecycle manager to, e.g., create a WLS partition, and create an OTD instance, or other slices of components. During orchestration, if a failure occurs, the system can accommodate it by, for example, retrying the REST endpoint for association and orchestration of its artifacts.

[00095] Exemplary end-to-end lifecycle use cases include providing a lifecycle manager on an admin server within a cluster, and/or other lifecycle managers on managed servers within the cluster. For example, in a single domain, a single lifecycle manager can be provided on the admin server. For multiple domains, or high-availability requirement models, in which a lifecycle manager can manage, e.g., multiple OTD instances on multiple servers, the lifecycle manager can be run within a cluster of its own front-ended by a load balancer, and from that domain the lifecycle manager can manage multiple runtimes.

[00096] In accordance with an embodiment, the end-to-end lifecycle configuration information can be saved, for example to an XML file in a single domain mode, or more typically to a database in a high-availability mode.

[00097] In accordance with an embodiment, the lifecycle manager can recognize that it is operating on behalf of a particular partition by, in response to receiving a call, e.g., via REST, for use with a registered runtime, determining from the call details an intended partition, and, based on what that call is attempting to create, calling one or more appropriate plugins.

[00098] For example, if the call is directed to creating a PDB for a partition, then the

lifecycle manager can respond by calling the database lifecycle plugin; or if the call is directed to creating an OTD instance for a partition, then the lifecycle manager can respond by calling the OTD lifecycle plugin, and so on.

[00099] In a single domain multiple partition environment, when an admin server fails,

5 which had been providing lifecycle manager functionality, the system can continue to act as a traditional application server (e.g., WLS) server environment would, although its administrative and management functionality may be disabled for a period of time.

[000100] In accordance with an embodiment, the system can use HK2 as a dependency injection mechanism, for looking up contracts and implementations of those

10 contracts as services. For example, using HK2, services can be injectable into components, and, once injected, their HK2 objects can satisfy the dependencies of other classes. Lifecycle plugins can be implemented as HK2 services and implement their contract at runtime. The system can look up a contract to find appropriate plugins that have been plugged in as services.

15 **[000101]** For example, in accordance with an embodiment, the system can recognize a partition plugin contract, or a runtime plugin, or other contract. Each partition or other plugin can be configured with an understanding of how to operate to create a partition or slice of partition within that component, so for example, a database partition plugin can be configured with an understanding of how to operate to create a slice of a database.

20 **[000102]** In accordance with an embodiment, the system can receive an indication as to the type of service required during a call (e.g., “type = database”), and then determine an appropriate, in this case, database plugin to be used.

25 **[000103]** Similarly, in accordance with an embodiment, for application server (e.g., WLS) and traffic director (e.g., OTD) components, plugins can be provided for those components that can be implemented as services, and subsequently used during creation or configuration of a partition.

[000104] Mechanisms such as HK2 also provide support for configuration management which can be used to define a configuration model stored in XML by, for example, default store or persistence type.

30 **[000105]** In accordance with an embodiment, configuration objects can be used to define the attributes for a particular component for use with a partition. These objects are similarly injectable, which allows looking up configuration information of a given partition by, for example, injecting the configuration object and looking up the information for that partition.

35

Environment

[000106] In accordance with an embodiment, a set of partitions across different

components can be linked together as an environment, to form a cohesive unit that serves a tenant's needs. While an environment can have more than one partition of each type, typically they will only have one of each type of partition. For example, in an FA scenario there can be more than one WLS partition from each of the domains in the multidomain 5 topology; in such a scenario there can be multiple associations between WLS and OTD and Database that must be supported.

Component / Runtime

[000107] In accordance with an embodiment, a component or runtime is a software component, for example an OTD, WLS, Coherence, or Database component. In accordance with other embodiments, additional types of software components can also be supported.

Lifecycle Manager

[000108] In accordance with an embodiment, the lifecycle manager enables configuration integration across different components, for example application servers, traffic directors or load balancers, and database or other components, which may have their own partition-specific configurations. The lifecycle manager can be implemented as a library providing a local API and/or REST access to clients to create lifecycle configurations, update those configurations, view those configurations, and destroy those configurations across an 20 environment.

Lifecycle Plugin

[000109] In accordance with an embodiment, a lifecycle plugin is a software module that provides a bridge between the lifecycle manager and a component, for example an 25 OTD, WLS, Coherence, or Database component. Lifecycle plugins can include classes that implement the runtime plugin and partition plugin (where supported) interfaces as HK2 services, which enable the lifecycle plugin to participate in configuration change transaction and notification events and take whatever action is appropriate to reflect that activity in a component. The plugin API is generally agnostic of the underlying protocol used to perform 30 the actual lifecycle operations for a runtime or a partition. The implementation of the plugin can choose how to connect to the runtime to perform the necessary operations. For example, in accordance with an embodiment, when an operation is initiated on a runtime via the runtime manager, such as quiesce, scaleUp or scaleDown, an implementation of the runtime plugin corresponding to the runtime type can be looked up via HK2, and the 35 corresponding method of the runtime plugin (e.g, RuntimePlugin) implementation invoked for the runtime plugin to take appropriate action.

Client

[000110] In accordance with an embodiment, a client refers to any client of a lifecycle manager, which interacts with the lifecycle manager via, e.g., REST APIs or in an embedded case via Java APIs. Examples of clients include Oracle Fusion Middleware Control (FMWC)

5 and Enterprise Manager Cloud Control (EMCC), and curl-based scripts. In accordance with an embodiment, the lifecycle manager REST APIs can be used within products such as Enterprise Manager when provisioning and updating partition configurations across components.

10 Lifecycle Use Cases

[000111] In accordance with an embodiment, the following are examples of lifecycle use cases: 1. Partition CRUD operations for a component, including the creation, deletion and updates of a partition; 2. Runtime operations for a component, including registration of runtimes, updates to the registered runtimes and unregistration of runtimes; and 3.

15 Association of partitions across different components, so that associating two partitions wires the two slices of runtimes (for example, a Database with WLS, or an OTD with WLS), so that the two can be used together.

Lifecycle Interactions

20 **[000112]** **Figure 7** illustrates an end-to-end lifecycle event diagram in a multitenant application server environment, in accordance with an embodiment.

[000113] As illustrated in Figure 7, in accordance with an embodiment, the lifecycle manager can allow orchestration of components associated with one or a plurality of plugins, such as an application server (e.g., WLS) lifecycle plugin 420, a traffic director (e.g., OTD) 25 lifecycle plugin 422, a database (e.g., Oracle Database) lifecycle plugin 424, and/or an in-memory data grid (e.g., Coherence) lifecycle plugin 426.

[000114] In accordance with an embodiment, in response to receiving a request to create an environment 430 (together with any associated metadata), the lifecycle manager orchestrates the creation and association of the various component partitions, and returns 30 the environment 432, for the tenant to use thereafter.

[000115] **Figure 8** illustrates an end-to-end lifecycle topology in a multitenant application server environment, in accordance with an embodiment.

[000116] As illustrated in Figure 8, in accordance with an embodiment, a simple or single domain usage can include one or more clients 440, an admin server 450, a cluster of 35 managed servers 452 (here indicated as MS1, MS2 and MS3), a database 454, and a traffic director / load balancer 460, wherein the lifecycle manager is running in the domain it manages, along with the other components such as OTD, Database and Coherence.

[000117] In this scenario, the end-to-end lifecycle configuration can be stored in an XML file and is made available on the managed servers to establish the tenancy at runtime. The lifecycle manager can be enabled out-of-the-box in some application server environments, where it will be running on the admin server and will not require any additional 5 setup.

[000118] **Figure 9** illustrates another end-to-end lifecycle topology in a multitenant application server environment, in accordance with an embodiment.

[000119] As illustrated in Figure 9, in accordance with an embodiment and a more complex scenario, such as FA provisioning using Enterprise Manager (EM) in an enterprise 10 deployment, wherein the lifecycle manager may be managing a lot of components and may also have high-availability requirements, the lifecycle manager can run in a separate lifecycle manager cluster / domain 471, as a plurality of EM/FA provisioning instances 472, 474, 476, each having lifecycle manager functionality 473, 475, 477.

[000120] In accordance with an embodiment, the lifecycle manager domain can be 15 used to manage other domains or components, such as domain 490 having a cluster 492 of managed servers (here indicated as MS1, MS2 and MS3).

[000121] In this scenario, the lifecycle manager can run on the managed servers in the lifecycle manager cluster, and the end-to-end lifecycle configuration can be persisted in the database, so that any domains or runtimes that need access to the information can access 20 the database to obtain that information.

[000122] In both of the above topologies, the components that can be managed within the lifecycle will be determined by the lifecycle plugins that are registered with the lifecycle manager, which in addition to configuration and creation can also support an elasticity functionality for scaling of runtimes as appropriate.

25 **Lifecycle Runtime**

[000123] In accordance with an embodiment, a lifecycle runtime (e.g., LifecycleRuntime) represents a component or managed service such as WLS, OTD, or Database. A lifecycle runtime can be managed by the lifecycle manager if an instance of the 30 runtime plugin is implemented by the component. The runtime plugin provides implementations that allow operations such as creation, deletion, update, scaling up, scale down and quiescing operations that can be performed on the runtime that is being managed via lifecycle manager. A runtime can either be created via the lifecycle manager, or can be registered with the lifecycle manager if created out of band for use by the lifecycle manager 35 via a register runtime (e.g., registerRuntime) method in the runtime manager.

[000124] In order to perform any operations on a runtime, an instance of the runtime must be looked up. In accordance with an embodiment, an instance of a lifecycle runtime

can be looked up using a runtime manager for the given runtime type as described below. A lifecycle runtime that represents managed services that support the use of a partition (a slice of the runtime) allow CRUD operations for that partition on the lifecycle runtime.

[000125] In accordance with an embodiment, in order to create a partition after looking

5 up a lifecycle runtime, a create partition (e.g., createPartition) method of the lifecycle runtime can be used to create a partition on the selected runtime. The way the partition is created is that the partition plugin for the lifecycle runtime will be looked up as a HK2 service for the runtime type, and the corresponding create method on the partition plugin is invoked.

10 **Runtime Manager**

[000126] In accordance with an embodiment, a runtime manager (e.g., RuntimeManager) can be provided as a singleton HK2 service that is used to manage runtimes. A runtime manager can be provided per Java virtual machine (JVM) that will manage the lifecycle of all runtimes on that JVM. Using a runtime manager, a runtime can

15 be created or registered. The runtime manager also provides operations to create, delete, update, scaleUp, scaleDown and quiescing of a runtime. The runtime manager will look up the runtime plugin as a HK2 service for that runtime type and invoke the corresponding operation on the plugin.

[000127] In accordance with an embodiment, once a runtime manager is looked up for

20 a given runtime type, then scaleUp and scaleDown operations on that runtime manager can be called for the specified runtime. The runtime manager will look up the corresponding runtime plugin for the service and invoke the scaleUp or scaleDown operation. This runtime support is particularly useful in addressing elasticity use cases.

[000128] In accordance with an embodiment, the runtime manager provides the

25 following operations for use with runtimes:

| Operation | Description |
|-------------------|--|
| createRuntime | Runtime can be created if the runtime plugin for the component supports creating the runtime. |
| deleteRuntime | Runtime can be deleted if the runtime plugin for the component supports deleting the runtime. |
| updateRuntime | Runtime can be updated if the runtime plugin for the component supports updating the runtime. |
| registerRuntime | Registers a pre-created runtime with LCM. At a minimum a runtime MUST be registered with LCM in order to perform subsequent operations on the runtime. |
| unregisterRuntime | Unregisters a runtime. A runtime can only be unregistered |

| | |
|--------------------------|--|
| | when LCM isn't managing any artifacts that reference the runtime. All partitions on the runtimes MUST be deleted before a runtime can be unregistered. If an attempt is made to unregister a runtime when there are references to it from LCM artifacts, an <code>IllegalStateException</code> MUST be thrown. |
| quiesceRuntime | If a runtime supports quiescing, the <code>quiesceRuntime</code> method on runtime manager will invoke the corresponding method for quiescing a runtime on the runtime plugin. If a runtime supports more than one phase for quiescing, for example <code>start_quiesce</code> and <code>end_quiesce</code> , the <code>quiesce</code> method could be called more than once with a phase. If a runtime doesn't have phases for quiescing, a null value for phase MUST be passed. If a runtime does not support quiescing, an <code>UnsupportedOperationException</code> MUST be thrown. |
| scaleUp | If a runtime supports scaling operations, the <code>scaleUp</code> method on runtime manager will scale the runtime up by invoking the corresponding runtime plugin to perform the operation of scaling up the runtime. If a runtime does not support <code>scaleUp</code> , an <code>UnsupportedOperationException</code> MUST be thrown. |
| scaleDown | If a runtime supports scaling operations, the <code>scaleDown</code> method on runtime manager will scale the runtime down by invoking the corresponding runtime plugin to perform the operation of scaling down the runtime. If a runtime does not support <code>scaleDown</code> , an <code>UnsupportedOperationException</code> MUST be thrown. |
| getRuntime(runtimeName) | Returns the runtime object with the specified name. If there is no runtime with the specified name, null will be returned. |
| getRuntimes(runtimeType) | Returns a list of runtimes for the specified runtime type. If there are no runtimes registered or created for the runtime type, null will be returned. |
| getRuntimes() | Returns a list of all the runtimes that LCM is managing. If no runtimes have been registered or created using LCM, null will be returned. |

[000129] In accordance with an embodiment, the following pseudo code illustrates how a WLS runtime and a Database runtime, which have been pre-created, are registered with lifecycle manager and are then subsequently used by lifecycle manager to create partitions

5 on the runtimes:

```

@Inject
private RuntimeManager rtmMgr;
// Database is the type of the Runtime and DB1 is the actual instance
rtmMgr.registerRuntime("Database", "DB1", dbProps);
5 // WebLogic is the type of Runtime and WLSMTRuntime1 is the actual instance
rtmMgr.registerRuntime("WebLogic", "WLSMTRuntime1", wlsProps);

// You can lookup a specific Runtime or have RuntimeManager select
// one based on the type of the Runtime and some internal heuristics to
10 determine
// which one to pick.
LifecycleRuntime wlsRuntime = rtmMgr.getRuntime("WLSRuntime1");

// Create a WLS Partition and a Slice/Partition of a Database
15 LifecyclePartition wlsPartition =
wlsRuntime.createPartition("cokeCRMWLSProdPartition", wlsPartitionProps);
// Get the Database runtime
Runtime dbRuntime = rtmMgr.getRuntime("DB1");

20 // Create a Database Partition (PDB in 12c Database case)
LifecyclePartition pdb = dbRuntime.createPartition("cokeCRMPDBProd", pdbProps);

```

Quiescing of Runtimes

25 [000130] In accordance with an embodiment, components such as WLS, which can include support for quiescing / graceful shutdown of runtimes, can support the quiescing capabilities via the lifecycle manager by implementing the quiescing support in the runtime plugin. Quiescing a runtime may have different phases. In such a scenario, the quiesce method of a runtime plugin may be invoked more than once, indicating the quiescing phase 30 via the phase parameter in the quiesce method.

35 [000131] For example, there may be a "start_quiesce" phase, indicating that the server is beginning its quiescing phase thus notifying that no new requests be forwarded to the server that is being quiesced. After the quiescing is done, the subsequent call will have a phase "end_quiesce" phase indicating that the server is now ready to start accepting new requests again. If a component supports additional phases, those phases can be supported by the runtime plugin to ensure complete functionality.

40 [000132] The "start_quiesce", "end_quiesce" are examples of phases that a component might support and not a normative list of phases that every component must support. If a component does not support quiescing then an UnsupportedOperationException must be thrown. If a component does not support phases, a null value is a valid value that would be

passed to the quiesce method.

Runtime Scaling

[000133] One of the advantages of a multi-tenant environment is to share hardware,

5 resources and services across multiple applications and user groups, achieving economies of scale and allowing administrators to have the flexibility in assigning capacity as the needs of these users vary over time.

[000134] In accordance with an embodiment, lifecycle manager APIs provide the capability for other elasticity frameworks to be able to scale up and scale down runtimes.

10 The runtime manager and runtime plugin provide scaleUp and scaleDown methods for users to call to scale a runtime.

[000135] In accordance with an embodiment, when scaleUp or scaleDown is called in runtime manager, the runtime plugin corresponding to the runtime type is looked up and the scaling operation is invoked on the instance of the runtime plugin. If a component does not

15 support scaling operations, the implementation of the runtime plugin must throw an UnsupportedOperationException.

Runtime Plugin Implementation

[000136] In accordance with an embodiment, illustrated below is a pseudo code for a

20 runtime plugin (for readability, this example does not implement all the abstract methods):

```
package com.oracle.weblogic.lifecycle.plugin.wls;
import com.oracle.weblogic.lifecycle.core.LifecyclePartitionImpl;
import org.jvnet.hk2.annotations.Service;
25 import com.oracle.weblogic.lifecycle.*;

@Service(name = "Weblogic")
public class WLSRuntimePlugin extends RuntimePlugin {

30    public void create(String runtimeName, LifecycleContext ctx) throws
        LifecycleException
    // If a component wants to support creation of Runtimes via LCM,
    // then the implementation of create for a Runtime will need to be provided
    }

35    public void delete(String runtimeName, LifecycleContext ctx) throws
        LifecycleException
    // If a component wants to support deleting of Runtimes via LCM,
    // then the implementation of delete for a Runtime will need to be provided
```

```

}

public void update (String runtimeName, LifecycleContext ctx) throws
LifecycleException
5 // If a component wants to support updating of Runtimes via LCM,
// then the implementation of update for a Runtime will need to be provided
}

public void scaleUp(String runtimeName, int scaleFactor, LifecycleContext ctx)
10 throws LifecycleException
// If a component wants to support scaleUp of Runtimes via LCM,
// then the implementation of scaleUp for a Runtime will need to be provided
}
...
15

```

Partition Plugin

[000137] In accordance with an embodiment, runtimes that support partitions (for example, slices of runtimes such as WLS, 12c Database etc) must provide an 20 implementation of a partition plugin. When an operation on a runtime is invoked, such as createPartition, the implementation of partition plugin corresponding to the runtime type will be looked up as an HK2 service, and the corresponding create method will be invoked on the partition plugin implementation class. It is up to the implementation class to connect to the runtime and create the slice of the runtime. Once created, the partition is registered with 25 lifecycle manager and persisted in the lifecycle- config.xml. In accordance with an embodiment, the runtime provides the following functions for partitions:

| Operation | Description |
|-------------------|---|
| createPartition | Create a partition on the runtime that supports creation of a slice of a runtime. Throws UnsupportedOperationException if the runtime does not support slicing a runtime. |
| deletePartition | Delete a partition on the runtime that supports deletion of a slice of a runtime. Throws UnsupportedOperationException if the runtime does not support slicing a runtime. |
| updatePartition | Update a partition on the runtime that supports updating of a slice of a runtime. Throws UnsupportedOperationException if the runtime does not support slicing a runtime. |
| registerPartition | Register a partition if the partition was created not using LCM but some other mechanism. Throws UnsupportedOperationException if a runtime does not support |

| | |
|---------------|---|
| | partitions / slicing a runtime. |
| getPartition | Get a partition by the name specified on the runtime that supports creation of a slice of a runtime. If the runtime supports partitions but does not have a partition with the specified name, a null value is returned. If the runtime does not support partitions / slicing of a runtime, this method MUST throw <code>UnsupportedOperationException</code> . |
| getPartitions | Returns a list of all partitions that have been created on the runtime. If the runtime supports partitions but does not have any partitions created on the runtime, a null value is returned. If the runtime does not support partitions / slicing of a runtime, this method MUST throw <code>UnsupportedOperationException</code> . |

[000138] In accordance with an embodiment, illustrated below is a pseudo code for a partition plugin. All operations in the plugin make a JMX connection to the actual WLS Runtime to perform the operations on the WLS Runtime in this example:

5

```

package com.oracle.weblogic.lifecycle.plugin.wls;
import com.oracle.weblogic.lifecycle.core.LifecyclePartitionImpl;
import org.jvnet.hk2.annotations.Service;
import com.oracle.weblogic.lifecycle.*;
10 @Service(name = "Weblogic")

public class WLSPartitionPlugin extends PartitionPlugin {
    @Override
15    public LifecyclePartition create(String partitionName, LifecycleContext ctx,
        Lifecycle LifecycleException {
        // Implementation of creating a WLS partition.
    }
    @Override
20    public abstract LifecyclePartition migrate(String partitionName, String phase,
        LifecycleContext ctx, LifecycleRuntime runtime) throws LifecycleException
    {
        //Implementation of migration of a WLS Partition / Resource Group.
    }
25    @Override
        public abstract LifecyclePartition quiesce(String partitionName, String phase,
        LifecycleContext ctx, LifecycleRuntime runtime) throws LifecycleException
    {
        //Implementation of migration of a WLS Partition / Resource Group.
    }
}

```

```

}

@Override
public void delete(String partitionName, LifecycleContext ctx, LifecycleRuntime
runtime) throws LifecycleException
5 //Implementation of delete of a WLS Partition
}

@Override
public LifecyclePartition update(String partitionName, LifecycleContext ctx,
Lifecycle LifecycleException {
10 // Implementation of      update of a WLS Partition configuration
}
@Override
public void associate(LifecycleContext ctx, LifecyclePartition wlsPartition,
LifecyclePartition otherPartition, LifecycleRuntime runtime) throws
15 LifecycleException
// Associate two partitions (WLS partition with DB for example or WLS with
OTD).
}
@Override
20 public void dissociate(LifecycleContext ctx, LifecyclePartition wlsPartition,
LifecyclePartition otherPartition, LifecycleRuntime runtime) throws
LifecycleException
//Dissociate previously associated partitions
}
25 }

```

Plugin Registration

[000139] As described above, in accordance with an embodiment, runtime plugins and partition plugins can be defined as an HK2 contract, with an implementation of these plugins as HK2 services. The HK2 service name must match the name of the service for which the 30 plugin is registered. Plugin registration can be done in two ways: The plugins are packaged as a jar file and registered with lifecycle manager via REST commands; or drop the plugin implementation jars in a directory specified by lifecycle manager implementation which must be configurable by the, e.g., WLS System Administrator. The lifecycle manager will scan the directory at startup to load the plugins and register with lifecycle manager implementation.

35 The default directory for placing the plugins is oracle_home/wlsserver/common/lifecycle/plugins.

Environment

[000140] In a typical production system, there will generally be more than one artifact 40 like a WLS partition, a Database partition, and others that will be required to run an

application. In accordance with an embodiment, an environment provides a grouping construct for all the associated pieces of an application. The lifecycle manager allows the user to create an environment, which can then be used by the system administrator to group partitions for a given tenant. In accordance with an embodiment, an environment provides

5 the following functions for partitions:

| Operation | Description |
|---------------------|--|
| addPartition | Adds the given partition to the environment. |
| removePartition | Removes the given partition from the environment. |
| getPartitions(type) | Returns a list of all partitions of the given type in the environment. If the environment does not have any partitions of the specified type, null will be returned. |
| getPartitions | Returns a list of all partitions in the environment. If the environment does not have any partitions null will be returned. |
| associate | Associates (wires) two partitions in the environment. For more details see section below. |
| dissociate | Dissociates two partitions in the environment. For more details see section below. |
| migrate | Initiates migration of a partition. For more details see section below. |

Association of Partitions with Environments

[000141] In accordance with an embodiment, once a partition is added to an

10 environment, it can be associated with the other partitions that were added to the environment.

[000142] For example, if an environment has a WLS partition and a database,

associating the two will result in a DataSource being created in WLS referring to the database with which it is being associated. To associate two partitions, the partitions first

15 need to be added to an environment. Once added to an environment, then an associate (e.g., Associate) method can be called on the environment passing in the two partitions that

need to be associated. When associate is called on an environment, the corresponding plugin's associate method is called along with a lifecycle context (e.g., LifecycleContext) object. The lifecycle context should have the necessary information that is passed to the

20 plugins to be able to do the necessary wiring of the two partitions. The associate method of partition plugins are invoked in the order in which they are passed into the environment's associate method call.

Dissociation of Partitions from Environments

[000143] In accordance with an embodiment, if a partition is no longer being used, references to that partition from other partitions in the environment need to be removed before it can be deleted. Dissociation provides a mechanism to un-wire two partitions in an environment. When dissociation is called, the partition plugin's corresponding dissociate method is invoked and it is the responsibility of the plugin to remove the configuration that wire the two partitions together.

[000144] For example, if a WLS partition is dissociated from an OTD partition configuration, the routing information that was configured in OTD will be removed and no traffic to the partition will be routed via OTD. The dissociate method of partition plugins are invoked in the order in which they are passed into the method call.

Partition Migration

[000145] In accordance with an embodiment, runtimes that support migration of partitions, need to implement the migrate method of partition plugin. When migrating a partition, the other partitions that are associated with it may also need to take action. To address this, a migrate method acts on the environment. The lifecycle manager infrastructure must invoke the migrate method of the partition that is being migrated first, and then call the migrate method of the other partitions that are in the environment.

[000146] In accordance with an embodiment, the migrating of a partition may have different phases. In such a scenario, the migrate method of a partition plugin may be invoked more than once, indicating the migration phase via the phase parameter in the migrate method.

[000147] For example, there may be a "start_migration" phase, indicating that the partition migration phase is beginning. After the migration is done, the subsequent call will have a phase "end_migration" phase indicating that the partition has been migrated. If a partition supports additional phases, those phases can be supported by the partition plugin to ensure complete functionality.

[000148] The "start_migration", "end_migration" are examples of phases that a partition might support and not a normative list of phases that every partition must support. If a partition does not support migration then an `UnsupportedOperationException` must be thrown. If a partition does not support phases, a null value is a valid value that would be passed to the migrate method.

35 Quiescing of Partitions

[000149] In accordance with an embodiment, runtimes that support quiescing of partitions will need to implement the quiesce method of partition plugin. The lifecycle

manager will check the state of partition to ensure that the partition is in a compatible state before trying to quiesce the partition.

Configuration Module

5 [000150] In accordance with an embodiment, the lifecycle manager includes a configuration module that manages the various configuration elements needed for its operation. The configuration module provides the necessary configuration infrastructure for various scenarios described here. The configuration module works in two modes: local - in which the lifecycle runs on admin server, configuration is persisted in an XML file in the WLS 10 config directory; and HA - in which the lifecycle can run on a cluster, configuration is persisted in a database. In either case, lifecycle configuration uses the HK2 framework to manage configurations.

Lifecycle Configuration File Propagation

15 [000151] In accordance with an embodiment, when changes are made to the configuration file on the admin server, those changes must be propagated to all managed servers. The managed servers use the lifecycle configuration to establish the right tenant context for applications running on the managed server. HK2 config objects do not allow a reparsing of the xml configuration file. Instead, a new object must be instantiated and the 20 updated xml file parsed fresh. When a configuration change is committed, the XML file is saved into a special directory, namely, the pending directory for the domain.

File Transfer to Managed Servers

25 [000152] In accordance with an embodiment, a LifecycleConfigDocument save() method is called by HK2 whenever a change has been committed. This method also saves a copy of the file to the pending directory, and then calls the following API which triggers the uploading of the changed file to managed servers:

```
ConfigurationManagerMBean.startEdit(...)  
30 ConfigurationManagerMBean.activate(...)
```

Watching for changes

35 [000153] In accordance with an embodiment, a FileWatcher daemon, started on the first parse of the configuration file, will detect that the xml file has changed. It will then create a new config object and parse the xml file. Any callers will be blocked trying to access information during the parsing. As soon as the parsing is complete, the modified data will be available for callers.

Accessing the Configured Data

[000154] In accordance with an embodiment, the HK2 configuration can be exposed with with LifecycleConfigFactory class, for example:

5

```

@Service
public class LifecycleConfigFactory implements PostConstruct {
    public LifecycleConfig getLifecycleConfig();
    public synchronized void reloadLifecycleConfig();
10
}

```

Edit Sessions

[000155] In accordance with an embodiment, for REST requests, that require a change in WLS configuration, each REST call creates and manages edit sessions on behalf of a 15 user. This is similar to the 'implicit changes' in development mode in the console.

LifecycleTask

[000156] There are certain operations performed via lifecycle manager that are naturally long running asynchronous operations. Examples of such operations are: quiescing 20 Server (runtime manager.quiesce); or starting a Server (runtime manager.quiesce). In such scenarios the semantic of lifecycle manager API can be made asynchronous, to return a LifecycleTask object from these APIs.

LifecycleTaskManager

25 [000157] Generally, when a user performs an asynchronous operation and receives a LifecycleTask object back, they can do one of two things with that object: get a status of the Asynchronous operation represented by the lifecycle task; or cancel the asynchronous operation represented by the lifecycle task. In accordance with an embodiment, the above operations are performed by the interfaces exposed by LifecycleTaskManager:

30

| Operation | Description |
|---------------|--|
| getTaskStatus | Get the status of an async operation represented by LifecycleTask. |
| cancelTask | Cancel the async operation represented by LifecycleTask. |

Lifecycle APIs and Classes

[000158] In accordance with various embodiments, examples of lifecycle interfaces and classes are described below. In accordance with other embodiments, other types of

interfaces and classes can be provided.

Lifecycle Manager

[000159] In accordance with an embodiment, an exemplary Lifecycle Manager interface

5 is described below:

```
10 package com.oracle.weblogic.lifecycle;
  import java.util.List;
  import org.jvnet.hk2.annotations.Contract;
15
  @Contract
  public interface LifecycleManager {
    /**
     *      @param environmentName
     *      @return object Environment
     *      This interface for creating a named logical Entity called Environment
     */
    public Environment createEnvironment(String environmentName);
    /**
     *      @param environmentName
     *      This interface for deleting a named logical Entity called * Environment
     */
    public void deleteEnvironment(String environmentName);
    /**
     *      @param environmentName
     *      @return object Environment
     *      This interface for getting Environment object given its name
     */
    public Environment getEnvironment(String environmentName);
30
    /**
     *      @param partitionId
     *      @param runtimeName
     *      @return object Environment
     *      This interface for getting Environment object partition Id and
     */
    public Environment getEnvironment(String partitionId, String runtimeName);
35
    /**
     *      @return List of Environment
     *      This interface has responsibility for getting all Environments.
     */
40
```

```

public List<Environment> getEnvironments();
5 /**
 *      This interface has responsibility for syncing all Environments.
 *      This call force sync partitions in an environment for all Environments
 */
public void syncEnvironments();
10 /**
 *      This interface has responsibility for syncing a named Environment.
 *      This call force sync partitions in an environment for given named
Environment
 */
public void syncEnvironment(String environmentName);
}

```

Runtime Manager.Java

¹⁵ [000160] In accordance with an embodiment, an exemplary Runtime Manager.Java interface is described below:

```

package com.oracle.weblogic.lifecycle;
import java.util.List;
20 import java.util.Map;
import org.jvnet.hk2.annotations.Contract;

@Contract
public interface RuntimeManager {
25 /**
 *      @param runtimeType The type of runtime for which the runtime is being
created
 *      @param runtimeName The name of the runtime for the runtime being created
 *      @param properties Properties to be passed to the runtime plugin
 */
30 *      @return
 *      @throws LifecycleException is thrown if the operation fails.
 */
public LifecycleRuntime createRuntime(String runtimeType, String runtimeName,
Map<String, PropertyValue> properties) throws LifecycleException
35 /**
 *      @param runtimeType The type of runtime being deleted (not sure if we
need this if we can determine the type from the runtime)
 *      @param runtimeName The name of the Runtime being deleted
 *      @throws LifecycleException thrown if the operation fails.
 */
40 *      @return
public void deleteRuntime(String runtimeType, String runtimeName)

```

```
throws LifecycleException;  
/**  
 *      @param runtimeType The type of runtime being updated (not sure if we  
 *      need this . if we can determine the type from the runtime)  
5   *      @param name The name of the Runtime being updated  
 *      @param properties Properties to be passed to the runtime plugin  
 *      @return  
 *      @throws LifecycleException thrown if the operation fails.  
 */  
10  public LifecycleRuntime updateRuntime(String runtimeType, String name,  
Map<String, PropertyValue> properties) throws LifecycleException  
/**  
 *      @param runtimeType The type of runtime being operated on  
 *      @param name The name of the Runtime being operated on  
15  *      @param phase Phase for operation  
 *      @param properties Properties to be passed to the runtime plugin  
 *      @return  
 *      @throws LifecycleException thrown if the operation fails.  
 */  
20  public LifecycleTask quiesceRuntime(String runtTimeType, String name, String  
phase, M  
/**  
 *      @param runtimeType The type of runtime being operated on  
 *      @param name The name of the Runtime being operated on  
25  *      @param phase Phase for operation  
 *      @param properties Properties to be passed to the runtime plugin  
 *      @return  
 *      @throws LifecycleException thrown if the operation fails.  
 */  
30  public LifecycleTask startRuntime(String runtTimeType, String name, String  
phase, Map  
/**  
 *      @param runtimeType The type of runtime being registered  
 *      @param runtimeName The name of the runtime being registered  
35  *      @param properties Properties to be passed to the runtime plugin  
 *      @throws LifecycleException if the register operation fails  
 */  
public void registerRuntime(String runtimeType, String runtimeName,  
Map<String, PropertyValue> properties) throws LifecycleException  
40  /**
```

```
*      @param runtimeType The type of runtime being unregistered
*      @param runtimeName The name of the runtime being unregistered
*      @throws LifecycleException if the register operation fails
*/
5   public void unregisterRuntime(String runtimeType, String runtimeName)
throws LifecycleException;
/***
*      @param runtimeType The type of runtime being scaled up
*      (not sure if we need this . if we can determine
10  *      the type from the runtime)
*      @param runtimeName The name of the Runtime being scaled up
*      @param scaleFactor The factor by which it is being scaled up.
*      @param properties Properties to be passed to the runtime plugin
*      @throws LifecycleException      thrown if the operation fails.
15  *      @throws UnsupportedOperationException if the operation isn't supported
*      by the runtime type
*/
public LifecycleTask scaleUp(String runtimeType, String runtimeName, int
scaleFactor,
20  Map properties) throws LifecycleException;
/***
*      @param runtimeType The type of runtime being scaled down
*      (not sure if we need this . if we can determine
*      the type from the runtime)
25  *      @param runtimeName The name of the Runtime being scaled down
*      @param scaleFactor The factor by which it is being scaled down.
*      @param properties Properties to be passed to the runtime plugin
*      @throws LifecycleException      thrown if the operation fails.
*      @throws UnsupportedOperationException if the operation isn't supported
30  *      by the runtime type
*/
public LifecycleTask scaleDown(String runtimeType, String runtimeName, int
scaleFactor
Map properties) throws LifecycleException;
35  /**
*      Get the specified Runtime
*
*      @param runtimeName Get the specific runtime by name.
*      @return
40  */
```

```
public LifecycleRuntime getRuntime(String runtimeName);
/**
 *      Get a List of runtimes for the specified runtime type.
 *
5   *      @param runtimeType Get all the runtimes for the given runtime type
 *      @return List of LifecycleRuntimes
 */
public List<LifecycleRuntime> getRuntimes(String runtimeType);
/**
10  *      Get a list of all the runtimes.
 *
 *      @return List of Runtimes.
 */
public List<LifecycleRuntime> getRuntimes();
15 }
```

Lifecycle Runtime

[000161] In accordance with an embodiment, an example Lifecycle Runtime interface is
20 described below:

```
package com.oracle.weblogic.lifecycle;
import java.io.File; import java.util.List; import java.util.Map;
import java.util.Properties;
25
public interface LifecycleRuntime {
/**
 *      Get the name of the Runtime
 *      @return the name of the runtime
30 */
public String getRuntimeName();
/**
 *      The type of service of the runtime
 *      @return the type of the runtime
35 */
public String getRuntimeType();

public void applyPartitionTemplate(File partitonTemplate);
/**
40  *      The properties that represent the runtime
 *      @return the type of the runtime
}
```

```
/*
public Properties getRuntimeProperties();

/**
 *      Create a partition on this Runtime
5  *      @param partitionName The name of the partition
 *      @param properties Properties passed to the plugin.
 *      @return
 */

public LifecyclePartition createPartition(String partitionName, Map properties)
10 throws
/**
 *      Delete the specified Partition
 *      @param partitionName The specified named Partition to be deleted
 */
15 public void deletePartition(String partitionName, Map properties) throws
LifecycleException
/**
 *      Update the specified Partition
 *      @param partitionName The specified named Partition being updated
20 *      @param properties Properties passed to the plugin.
 *      @return
 */

public LifecyclePartition updatePartition(String partitionName, Map properties)
throws
25 /**
 *      @return List of Partitions on this Runtime
 */
public List<LifecyclePartition> getPartitions()
/**
30 *      Get the specific named Partition.
 *      @param partitionName The name of the Partition to get.
 *      @return
 */

public LifecyclePartition getPartition(String partitionName);
35 /**
 *      Register a partition on this Runtime with LCM
 *      @param partitionName The name of the partition
 *      @param partitionId The ID of the partition
 *      @return
 */
40 *      @throws com.oracle.weblogic.lifecycle.LifecycleException
*/
```

```

public void registerPartition(String partitionName, String partitionId);

/*
 *      UnRegister a partition on this Runtime with LCM
5   *      @param partitionName The name of the partition
 *      @return
 *      @throws com.oracle.weblogic.lifecycle.LifecycleException
 */
public void unregisterPartition(LifecyclePartition partition);
10
}

```

Lifecycle Partition

[000162] In accordance with an embodiment, an exemplary Lifecycle Partition package
15 is described below:

```

package com.oracle.weblogic.lifecycle;
public interface LifecyclePartition {
    public String getName(); public String getType(); public String getId();
20    public String getRuntimeName();
}

```

Runtime Plugin

[000163] In accordance with an embodiment, an exemplary Runtime Plugin class is
25 described below:

```

package com.oracle.weblogic.lifecycle;
import org.jvnet.hk2.annotations.Contract;
30
@Contract
public abstract class RuntimePlugin {
    /**
     *      @param runtimeName
35    *      @param ctx
     *      @throws com.oracle.weblogic.lifecycle.LifecycleException
     *      Implementation of method should facilitate creation of runtimes using
     *      runtime
     */
40    public abstract void create(String runtimeName, LifecycleContext ctx) throws
        LifecycleException

```

```
/*
 *      @param runtimeName
 *      @param ctx
 *      @throws com.oracle.weblogic.lifecycle.LifecycleException
5   *      Implementation of method should facilitate deletion of runtimes using
 *      runtime
 */
public abstract void      delete(String runtimeName, LifecycleContext ctx)
throws LifecycleException
10  /**
 *      @param runtimeName
 *      @param ctx
 *      @throws com.oracle.weblogic.lifecycle.LifecycleException
 *      Implementation of method should facilitate updates of runtimes using
15  runtimeN
 */
public abstract void update (String runtimeName, LifecycleContext ctx) throws
LifecycleException
/**
20  *      @param runtimeName
 *      @param scaleFactor
 *      @param ctx
 *      @throws com.oracle.weblogic.lifecycle.LifecycleException
 *      Implementation of this method should facilitate scaleUp of runtime
25  */
public abstract LifecycleTask scaleUp(String runtimeName, int scaleFactor,
Lifecycle
/**
 *      @param runtimeName
30  *      @param scaleFactor
 *      @param ctx
 *      @throws com.oracle.weblogic.lifecycle.LifecycleException
 *      Implementation of this method should facilitate scaleDown of runtime
 */
35  public abstract LifecycleTask scaleDown(String runtimeName, int scaleFactor,
LifecycleContext
/**
 *      @param runtimeName
 *      @param ctx
 *      @throws com.oracle.weblogic.lifecycle.LifecycleException
40  */
```

```

*      Implementation of this method should facilitate quiesce of runtime
*/
public abstract LifecycleTask quiesce(String runtimeName, String phase,
LifecycleContext
5  /**
*      @param runtimeName
*      @param ctx
*      @throws com.oracle.weblogic.lifecycle.LifecycleException
*      Implementation of this method should facilitate start of runtime
10 */
public abstract LifecycleTask start(String runtimeName, String phase,
LifecycleContext
}

```

15

Partition Plugin

[000164] In accordance with an embodiment, an exemplary Partition Plugin class is described below:

```

20 package com.oracle.weblogic.lifecycle;
import org.jvnet.hk2.annotations.Contract;

@Contract
public abstract class PartitionPlugin {
25 /**
*      @param partitionName
*      @param ctx
*      @param runtime
*      @return
*      @throws com.oracle.weblogic.lifecycle.LifecycleException
*      Implementation of method should facilitate creation of a Partition for a
30 specific
*      Component that the Plugin represents.
*/
35 public abstract LifecyclePartition create(String partitionName,
LifecycleContext ctx, LifecycleRuntime runtime)
/**
throws LifecycleException;
*      @param partitionName
40 *      @param phase

```

```
*      @param ctx
*      @param runtime
*      @return
*      @throws com.oracle.weblogic.lifecycle.LifecycleException
5   *      Implementation of method should facilitate migration of a Partition for a
*      specific
*      Component that the Plugin represents.
*/
public abstract LifecyclePartition migrate(String partitionName,
10  String phase, LifecycleContext ctx, LifecycleRuntime runtime)
/***
throws LifecycleException;
*      @param partitionName
*      @param ctx
15  *      @param runtime
*      @throws com.oracle.weblogic.lifecycle.LifecycleException
*      Implementation of method should facilitate deletion of a Partition for a
*      specific
*      Component that the Plugin represents.
20  */
public abstract void delete(String partitionName,
LifecycleContext ctx, LifecycleRuntime runtime)
throws LifecycleException;
/***
25  *      @param partitionName
*      @param ctx
*      @param runtime
*      @return
*      @throws com.oracle.weblogic.lifecycle.LifecycleException
*      Implementation of method should facilitate update of a Partition for a
*      specific
30  *      Component that the Plugin represents.
*/
public abstract LifecyclePartition update(String partitionName,
35  LifecycleContext ctx, LifecycleRuntime runtime)
/***
throws LifecycleException;
*      @param partitionName
*      @param ctx
40  *      @param runtime
```

```
*      @return List<LifecycleTask>
*      @throws com.oracle.weblogic.lifecycle.LifecycleException
*      Implementation of method should facilitate quiesce of a Partition.
*/
5   public abstract List<LifecycleTask> quiesce(String partitionName,
String phase, LifecycleContext ctx, LifecycleRuntime runtime)
/***
throws LifecycleException;
*      @param partitionName
10  *      @param ctx
*      @param runtime
*      @return List<LifecycleTask>
*      @throws com.oracle.weblogic.lifecycle.LifecycleException
*      Implementation of method should facilitate start of a Partition.
15  */
public abstract List<LifecycleTask> start(String partitionName,
String phase, LifecycleContext ctx, LifecycleRuntime runtime)
throws LifecycleException;
/***
20  *      @param ctx
*      @param partition1
*      @param partition2
*      @param runtime
*      @throws com.oracle.weblogic.lifecycle.LifecycleException
25  *      Implementation of this method should facilitate association of two
Partitions.
*/
public abstract void associate(LifecycleContext ctx,
LifecyclePartition partition1, LifecyclePartition partition2,
30  /**
throws LifecycleException;
LifecycleRuntime runtime)
*      @param ctx
*      @param partition1
35  *      @param partition2
*      @param runtime
*      @throws com.oracle.weblogic.lifecycle.LifecycleException
*      Implementation of method should facilitate disassociation of two
Partitions.
40  */
```

```
1    public abstract void dissociate(LifecycleContext ctx,
2        LifecyclePartition partition1, LifecyclePartition partition2, LifecycleRuntime
3        runtime)
4        throws LifecycleException;
5    }
```

Environment

[000165] In accordance with an embodiment, an exemplary Environment interface is

10 described below:

```
1    package com.oracle.weblogic.lifecycle;
2    import java.util.List;
3    import java.util.Map;
4
5    public interface Environment {
6        /**
7         *
8         * @return Name of environment
9         */
10        public String getName();
11
12        /**
13         *
14         * @param partition
15         */
16        public void addPartition(LifecyclePartition partition);
17
18        /**
19         *
20         * @param partitionType
21         * @param partitionName
22         */
23        public void removePartition(String partitionType,
24            String partitionName);
25
26        /**
27         *
28         * @return
29         */
30        public List<LifecyclePartition> getPartitions();
31
32        /**
33         *
34         * @param partitionType
35         * @return
36         */
37    }
```

```
/*
public List<LifecyclePartition> getPartitions(String partitionType);
/***
*      Migrate the specified Partition
5  *      @param partition The specified named Partition being updated
*      @param runtime . the runtime on which the partition exists
*      @param properties Properties passed to the plugin.
*/
public void migratePartition(LifecyclePartition lifecyclePartition,
10 LifecycleRuntime runtime, String phase, Map properties) throws
LifecycleException;
/***
*
*      @param partition1
15 *      @param partition2
*/
public void associate(LifecyclePartition partition1, LifecyclePartition
partition2, Map partitionProperties) throws LifecycleException;
/***
20 *
*      @param partition1
*      @param partition2
*/
public void dissociate(LifecyclePartition partition1, LifecyclePartition
partition2, Map partitionProperties) throws LifecycleException;
/***
*
*      @param lifecyclePartition
*      @param phase
30 *      @param properties
*/
public Map<String, List<LifecycleTask>> quiesce(LifecyclePartition
lifecyclePartition, String phase, Map properties) throws LifecycleException ;
/***
35 *
*      @param lifecyclePartition
*      @param phase
*      @param properties
*/
40 public Map<String, List<LifecycleTask>> start(LifecyclePartition
lifecyclePartition, String phase, Map properties)
```

```
throws LifecycleException ;
}
```

5 Lifecycle Context

[000166] In accordance with an embodiment, an exemplary Lifecycle Context interface is described below:

```
package com.oracle.weblogic.lifecycle;
10 import java.util.Map;
public interface LifecycleContext {

    /**
     *      Map of properties that are passed from the client to the plugin
15    *      @return
     */
    public Map getProperties();
}
```

20

Lifecycle Operation Type

[000167] In accordance with an embodiment, an example Lifecycle Operation Type enumerator is described below:

```
25 package com.oracle.weblogic.lifecycle;
public enum LifecycleOperationType { CREATE_RUNTIME, DELETE_RUNTIME,
    UPDATE_RUNTIME, SCALE_UP_RUNTIME, SCALE_DOWN_RUNTIME, CREATE_PARTITION,
    DELETE_PARTITION, UPDATE_PARTITION, MIGRATE_PARTITION
}
```

30

Lifecycle Task

[000168] In accordance with an embodiment, an example Lifecycle Task interface is described below:

```
35 import java.util.Map;
public interface LifecycleTask
{
    /**
40     *      Get the type of component the task represents
     *
```

```

*      @return the type of the task (ex:wls)
*      @includeapi forpublicapi
*/
public String getComponentType();
5  /**
*      The properties that represent the task
*
*      @return the Properties of the task
*      Example: used to store things like servername in case of server operation
10 */
public Map getProperties();
}

```

15 Lifecycle Task Manager

[000169] In accordance with an embodiment, an exemplary Lifecycle Task Manager interface is described below:

```

import org.jvnet.hk2.annotations.Contract;
20 import java.util.Map;

@Contract
public interface LifecycleTaskManager {
    public String getTaskStatus(String runtimeName, String taskType, Map properties
25 ) throws public void cancelTask(String runtimeName, String taskType, Map
    properties ) throws LifecycleException
}

```

30 [000170] **Figure 10** illustrates a flowchart of method for providing end-to-end lifecycle in a multitenant application server environment, in accordance with an embodiment.

[000171] As illustrated in Figure 10, at step 510, an application server environment is provided at one or more computers, including a domain for execution of software applications, which supports one or more partitions, wherein each partition provides an 35 administrative and runtime subdivision of the domain, and wherein a partition can optionally include one or more resource groups having a collection of deployable applications or resources and/or referencing a resource group template.

[000172] At step 512, a lifecycle manager is provided, which enables configuration integration across different components, for example application servers, traffic directors or 40 load balancers, and database or other components, which may have their own partition-

specific configurations.

[000173] At step 514, a plurality of components are provided, each of which includes a lifecycle plugin that provides a bridge between the lifecycle manager and that component.

[000174] At step 516, during creation or configuration of a partition, the lifecycle

5 manager performs lifecycle operations that can call a plugin interface, which in turn performs the actual creation or configuration of a component, for example, the creation of a traffic director instance, or the creation of a database instance, for use with the partition.

* * * * *

10

Tenant Onboarding

[000175] In accordance with an embodiment, the system includes support for a tenant onboarding functionality, which can be used to onboard a tenant for a service to a particular partition.

15 **[000176]** For example, in accordance with an embodiment that includes Fusion applications (FA), the system can provide a mechanism to onboard an FA tenant for a service to a given partition. In accordance with an embodiment, each environment can be associated with a service in the upper stack. When onboarding the tenant, information such as which partition is the FA service using, the tenant's database associated with the service, 20 as well as the top-level tenant directory, can be specified.

25 **[000177]** In accordance with an embodiment, the system provides an API to map a tenant to a given partition. For incoming requests, the system can determine the tenant from the configuration and use that information to establish a tenancy context at runtime. When queried using, e.g., a Java Required Files (JRF) API, the appropriate tenancy context can then be made available to the application.

30 **[000178]** For example, as described above, the system can be used to create an environment, so that an e.g., FA tenant can be onboarded, and so that upper stack components, such as Oracle Platform Security Services (OPSS), and Service Oriented Architecture (SOA) can then obtain access to provision that environment, so that the finally-configured environment will have all of the components that the tenant will need.

35 **[000179]** In accordance with an embodiment, a tenant can have one or more different services. For example, a tenant can have a Customer Relationship Management (CRM) service and/or a SOA service. Generally, each service in the upper stack can be associated with one or more environment in the application server, which means that a tenant signing up for two services in the upper stack will generally receive two environments, with the environments being linked at the tenant level.

[000180] For example, if a tenant signs up for SOA and Business Process Management

(BPM) services, then the tenant will receive a SOA environment and a BPM environment, and these environments will both be linked by that tenant. Each application within an environment can use a remote call, e.g., as an HTTP or a T3 call, to another environment associated with that tenant, if necessary.

5 [000181] As described above, in accordance with an embodiment, the system recognizes the concepts of tenancy context and partition context, with the former generally containing more information. A particular service may include several slices or partitions of various components, for example an application server (e.g., WebLogic, WLS), database (e.g., Database), and traffic director or other load balancer (e.g., Oracle Traffic Director, 10 OTD). A tenant context maps to a partition context for a particular tenant, in the manner of indicating a service and how that service maps to the partition context.

[000182] **Figure 11** illustrates tenant onboarding in a multitenant application server environment, in accordance with an embodiment.

15 [000183] As illustrated in Figure 11, in accordance with an embodiment, a tenant onboarding API 550 and configuration interface 560 allows an administrator 562 to provide a tenant configuration 564. The tenant onboarding process can then use that information to create a tenant 565 having a UUID and environment 566, and map that tenant's environment 570 to the appropriate partition or domain-level resources.

20 [000184] In accordance with an embodiment, each tenant can be associated with a globally unique ID (UUID), and each service can be similarly associated with a UUID. This allows the system to operate in the context of whichever service is being called to look up the tenant context. The system also provides the ability to switch between different tenant contexts to execute a particular task, which is generally referred to herein in various embodiments as ExecuteAs functionality.

25 [000185] For example, in an application server (e.g., WLS) environment, a global or domain level runtime can be provided, as further described below. A product or component in the upper stack (e.g., FA, SOA, or BPM) may want to switch to a given tenant context, for example for onboarding a particular tenant to a particular service. When that tenant is onboarded, the upper stack software running at the global or domain level can temporarily 30 switch to that tenant, and perform any necessary tasks e.g., with the traffic director, or the database.

[000186] Generally, the system can switch from a global level context to a tenant context only, and not the other way around, such that an executeAs operation can only be executed from a global to tenant contexts, but not from a tenant to tenant context.

35

Switching of Tenancy Context

[000187] In accordance with an embodiment, the system can determine a partition

context automatically based on request attributes such as a target component that may be currently handling the request. The partition context of the thread generally changes whenever the request enters a component that is deployed to a different partition.

[000188] For example, when a servlet running in one partition calls a remote component such as an Enterprise Java Bean (EJB) or a servlet or a web service deployed in another partition, the remote call is executed in the partition hosting the target component.

[000189] In this case, the underlying platform takes care of setting the appropriate partition context before calling the business method of the target component. An application developer can make use of this facility when they want to execute code under a different partition context.

[000190] In accordance with an embodiment, tenancy context switching also supports simple use cases, for example where an application sets the context before making a call to a shared service which may run in the same thread sharing the component context of the caller. In such cases partition context switching is neither required nor desired.

Tenancy Context Switching as a Privileged Operation

[000191] In accordance with an embodiment, only privileged application code, such as the global runtime, can switch tenancy context. Therefore, the API must enforce appropriate security permission checks before allowing access. Permissions can be granted that include the target tenancy being set.

Protection against Illegal Tenancy Switching

[000192] In accordance with an embodiment, the table below describes which context switches are allowed, subject to security privileges of the code:

| Current Application Tenancy | New Application Tenancy | Possible or Not |
|------------------------------------|--------------------------------|------------------------|
| GLOBAL | Any valid tenancy context | Possible |
| GLOBAL | GLOBAL | Possible |
| Tenant | GLOBAL | Impossible |
| Tenant-i | Tenant-i | Impossible |
| Tenant-i | Tenant-j | Impossible |
| null | Any tenancy context | Impossible |

[000193] In accordance with an embodiment, in, for example, an FA scenario, the process of tenant onboarding can be initiated by communicating with the provisioning system. In some situations, such as the case of FA running on a multi-tenant application server such as a WLS server, the process of tenant onboarding requires creation of partitions

on runtimes which are setup by the provisioning sub-system.

[000194] In accordance with an embodiment, prior to onboarding a tenant, an environment representing the service for which the tenant is signing up for has to be created. The FA provisioning sub-system can call into the lifecycle manager API to first create an environment. Once an environment has been created, a tenant can be onboarded via a tenant admin (e.g., TenantAdmin) class, an example of which is illustrated below. The tenant admin can then expose an API (e.g., registerTenant) to register a tenant.

[000195] Once a tenant is registered, a database created by the provisioning infrastructure can be registered with the tenant admin for a given tenant. At this point, the system can onboard a tenant to its environment. Onboarding a tenant requires the service name and service type (available from FA) to be passed, in order to maintain a mapping for the tenant to an environment. The tenant mapping can be stored in lifecycle-config.xml.

[000196] In accordance with an embodiment, below is an exemplary code for tenant mapping stored in the lifecycle- config.xml:

```

15 <tenants>
  <tenant name="coke" id="123" top-level-dir="coke/top/level/dir">
    <service type="CRMSERVICE" name="CRMProd"
20   environment-ref="coke-crmenv-prod"/>
    <service type="CRMSERVICE" name="CRMTest"
      environment-ref="coke-crmenv-test"/>
      name="HCMProd" <service type="CRMSERVICE
      environment-ref="coke-hcmenv-prod"/>
    </tenant>
25 <tenant name="pepsi" id="456" top-level-dir="pepsi/top/level/dir">
    <service type="HCMService" name="HCMProd"
      environment-ref="pepsi-hcmenv-prod"/>
    </tenant>
  </tenants>
30

```

[000197] In accordance with an embodiment, at runtime, the tenant context (e.g., TenantContext) is available via a TenantContext / TenantContextFactory. An API to get a TenantContext for a given partition or tenant database is defined in the TenantManager class defined below. All of the tenant admin APIs for registering a tenant and the corresponding tenant database and environment can be exposed via a REST API when running in WLS. When running on Java SE the Java APIs can be used directly. An exemplary pseudo code is shown below describing the process above to onboard a tenant on WLS environment:

```
@Inject
```

```

TenantAdmin tenantAdmin

@Inject
RuntimeManager rtmMgr;
5
@Inject
LifecycleManager lcmMgr;

// Database is the type of the Runtime and DB1 is the actual instance
10 rtmMgr.registerRuntime("Database", "DB1", props);
rtmMgr.registerRuntime("WebLogic", "WLSMTRuntime1", props);

// You can lookup a specific Runtime or have RuntimeManager select one based on
type
15 // and some internal heuristics to determine which one to pick.
Runtime wlsRuntime = rtmMgr.getRuntime("WLSRuntime1");
or
Runtime wlsRuntime = rtmMgr.selectRuntime("WebLogic");

20 //Create a WLS Partition and a PDB
LifecyclePartition wlsPartition =
wlsRuntime.createPartition("cokeCRMWLSProdPartition", Runtime dbRuntime =
rtmMgr.getRuntime("DB1"));
LifecyclePartition pdb = dbRuntime.createPartition("cokeCRMPDBProd", props);
25

Environment cokeEnv = lcmMgr.createEnvironment("coke-crmenv-prod");
cokeEnv.addPartition(wlsPartition);
cokeEnv.addPartition(pdb);

30 //Tenant ID and name are provided by Provisioning to LCM Tenant tenant =
tenantAdmin.registerTenant("111", "coke"); tenantAdmin.onBoard(tenant,
"CRMService", "CRMPProd", cokeEnv);

```

35 **[000198]** In the code above, if the DB partition is not created using a lifecycle manager, but instead performed by provisioning, then provisioning will call registerDatabase as opposed to createPartition on the DB1 runtime object.

```

LifecycleRuntime runtime = rtmMgr.getRuntime("DB1");
40 runtime.registerDatabase("111", "cokeCRMPDBProd");

```

Tenant Context

[000199] In accordance with an embodiment, in FA / SaaS and similar scenarios where the applications themselves may be expected to be multi-tenant aware, then in addition to the partition ID and partition name, they need to know the tenant on behalf of which a certain operation is being performed. A tenant context (e.g., TenantContext) API can be made 5 available in the JRF layer for upper stack FMW / FA to use to get the appropriate tenant context.

```
10 TenantContextFactory tcFactory =
  TenantContextFactory.getTenantContextFactory(); TenantContext tenantContext =
  tcFactory.getTenantContext();
```

[000200] Once a tenant context is determined, a component can lookup the partition(s) associated with the environment for a given service.

```
15 Environment env = tenantContext.getEnvironment();
List<Partitions> partitions = env.getPartitions(); Iterator iterator =
partitions.iterator();
while (iterator.hasNext()) {
  LifecyclePartition partition = iterator.next();
20 if (partition.getType().equals("WebLogic")) {
  String partitionId = partition.getId(); String partitionName =
partition.getName();
}
}
```

25

Support for Java SE for tenant mapping.

[000201] In accordance with an embodiment, the following is an exemplary code which can be used in SE to retrieve a tenant context:

```
30 TenantContextFactory tcFactory = TenantContextFactory.getTenantContextFactory(); TenantContext tenantContext =
tcFactory.getTenantContext();
```

[000202] For Java SE, a method can be exposed in TenantManager which will fetch a 35 tenant context based on a tenant ID, service name and service Type. The setting up of tenant context can be done implicitly by the TenantContextFactory when getTenantContext() is called for the first time.

Config File Bootstrap in SE

[000203] In accordance with an embodiment, the tenant mapping can be stored in lifecycle-config.xml as mentioned in previous section. In SE, the system can look for this file in (a) a location pointed to by environment variable lifecycle.configDir; or (b) if a lifecycle.configDir variable is not set, then the user working directory, which can be read

5 using:

```
System.getProperty("user.dir");
```

Specifying Tenant Information in SE

[000204] In accordance with an embodiment, when running in JAVA SE mode, the information identifying a tenant can be read using the following system properties. lifecycle.tenantId lifecycle.serviceName lifecycle.serviceType; or by specifying the service UUID via system property lifecycle.serviceld

[000205] The above information uniquely identifies a tenant, and can be used to identify a unique tenant plus service combination present in the lifecycle configuration. The SE application assumes the identity and that is represented when TenantContextFactory.getTenantContext is called.

Tenant Context in SE

[000206] In accordance with an embodiment, the tenant context in SE is stored on InheritableThreadLocal. When getTenantContext() method is called for the first time, it uses the system properties to compute the tenant context and stores it also on InheritableThreadLocal.

[000207] Upon further calls to same method, the system checks to see if at tenant context is stored on Inheritable ThreadLocal and will return what has been saved. If in the middle of the program execution, the system properties are changed, these changes may or may not be detected and used to refresh the tenant context on the Inheritable Thread Local.

Switching with switchTenantContext (aka executeAs)

[000208] Figure 12 further illustrates tenant onboarding in a multitenant application server environment, in accordance with an embodiment.

[000209] As illustrated in Figure 12, in accordance with an embodiment, the runtime view of FA-MT includes an admin runtime 580, which in turn provides access to one or more global data sources 582 or proxy data sources 584, a plurality of tenant runtimes 590, 600, each with their own global data source 592, 602, and tenant data source 594, 604, a global PDB 610 and several tenant PDBs 612, 614. Each of these component can be made accessible to upper stack components via a plurality of URLs, for example <http://tenant1.opc>,

http://tenant2.opc, and http://internal.opc.

[000210] In accordance with various embodiments, there are several different usages of tenancy in FAMT runtime, a few examples of which are described below: used by MetaDataService to select the tenant database partition to use to pick tenant-specific 5 customization of some metadata.

[000211] For example, it can be used by OPSS to select the tenant database partition to pick security policies that need to be enforced. It can also be used by Audit Service to determine the audit policy for a tenant and select the tenant database partition to store the audit log. In all such cases, there is a need for the correct tenancy to be associated with the 10 thread of execution before the shared service is invoked.

[000212] When tenancy context is not explicitly set, it is derived from the current partition context of the thread as described below: When a request enters WLS, an appropriate container determines various details of the request's target component, and creates a ComponentInvocationContext object to represent the invocation. Similarly, when 15 WLS invokes a callback object as part of some asynchronous operation like timer expiration, it sets the correct invocation context in the thread. As part of ComponentInvocationContext, partition information is also maintained. The onboarding process described above provides the necessary mapping for the current partition to the current tenant value which is available via the tenancy context.

20

Global Tenancy

[000213] Typically, the system's administration code usually has a need to run in the context of an MT Administrator who represents all tenants. In accordance with an embedment, this special state is represented by a special tenant ID called "GLOBAL", which 25 is different from the null tenant ID that indicates there is no tenant context associated with the current request.

[000214] Any software code that wants to use GLOBAL Database must maintain a fixed reference to a DataSource that always connects to GLOBAL Database. Code running in the GLOBAL runtime does not have a fixed binding to a particular tenant DB, but instead 30 accesses the data from a tenant DB that is dependent on the operational context. Since tenant DBs are not pre-provisioned in the system, a proxy data source (e.g., ProxyDataSource) can be used, in which the code can continue to have a fixed JNDI reference to the proxy data source, which selects a particular underlying DB based on the current tenant context.

35

[000215] There are two distinct use cases requiring FMW/FA code to run outside WLS as described below: Certain life cycle operations and administrative operations happen in Java SE environment. As part of business process management, the server code

sometimes fork a long running Java process that's working on a specific tenant's context.

[000216] Having a single reference to a proxy data source which is made available via JNDI and the proxy data source takes care of delegating the calls to appropriate underlying database using tenancy context. For the above to work, the tenancy context must be set correctly in the thread in Java SE process. The Java SE process must be launched by passing in the GLOBAL tenant context using the system properties for tenant name and tenant id. The forked Java process uses this information and calls switchTenantContext to establish the tenancy context in the thread of execution.

10 **Global Tenancy check in Java SE**

[000217] In accordance with an embodiment, in order to be able to executeAs, the system must ensure that the current tenant context is Global. This can be done by checking that for a current tenant context, that the tenant ID is equal to 0.

15 **[000218]** **Figure 13** illustrates a flowchart of method for providing tenant onboarding in a multitenant application server environment, in accordance with an embodiment.

[000219] As illustrated in Figure 13, at step 630, an application server environment is provided at one or more computers, including a domain for execution of software applications, which supports one or more partitions, wherein each partition provides an 20 administrative and runtime subdivision of the domain, and wherein a partition can optionally include one or more resource groups having a collection of deployable applications or resources and/or referencing a resource group template.

[000220] At step 632, a lifecycle manager is provided, which enables configuration integration across different components, for example application servers, traffic directors or 25 load balancers, and database or other components, which may have their own partition-specific configurations.

[000221] At step 634, a tenant onboarding and/or configuration interface are provided, which allows for definition of a tenant configuration for a tenant to be onboarded to the system.

30 **[000222]** At step 636, during tenant onboarding, the system uses the tenant configuration information to create a tenant having a unique identifier, an environment, and a tenant context that maps a partition context for a particular tenant.

[000223] At step 638, the system can optionally switch to the context of the tenant via a tenant switching interface, and execute functions as or on behalf of that tenant, during the 35 tenant onboarding.

* * * * *

Component Lifecycle Plugins

[000224] In accordance with an embodiment, the end-to-end lifecycle infrastructure includes a plugin-based mechanism which enables the different components within a multi-tenant application server environment to use a lifecycle plugin to provide their own component-specific lifecycle functionality.

[000225] **Figure 14** illustrates lifecycle plugins in a multitenant application server environment, in accordance with an embodiment.

[000226] As described above, in accordance with an embodiment, the system can include a lifecycle coordination framework, including a lifecycle manager (LCM) and lifecycle interface, which enables lifecycle management of the partition and other components, such as a traffic director component, database component, or other components.

[000227] For example, as illustrated in Figure 14, the lifecycle interface can support the use of a partition lifecycle plugin 650, an application server lifecycle plugin 652, a traffic director lifecycle plugin 654, and a database lifecycle plugin 656, or other types of plugin.

[000228] Each lifecycle plugin can encapsulate the lifecycle needs for, for example, an application server, e.g., WebLogic, WLS; or traffic director or other load balancer, e.g., Oracle Traffic Director, OTD instance lifecycle. The lifecycle interface provides a common or generic process for each plugin to manage the lifecycle of its given component.

[000229] **Figure 15** further illustrates lifecycle plugins in a multitenant application server environment, in accordance with an embodiment.

[000230] As illustrated in Figure 15, the lifecycle plugins can contribute to a partition lifecycle management 660, or a database lifecycle management, 662.

[000231] For example, an OTD instance can request the lifecycle manager to create a WLS runtime, for use in creating a partition on a particular domain. The lifecycle manager delegates to each plugins create method, in effect providing an indication of the domain, and requesting the plugin to create a partition on that domain.

[000232] In accordance with various embodiments, similar plugin functionality can be provided to support, e.g., other database, traffic director, or other components, and to create, delete, stop, start, quiesce, migrate servers where available. Each plugin can be implemented for each component supported through the lifecycle manager.

A. Application Server Plugin

[000233] In accordance with an embodiment, the system can include an application server (e.g., WebLogic, WLS) lifecycle plugin implements the lifecycle APIs and manages the creation or configuration of application server (e.g., WLS) partitions.

Interaction with Lifecycle Manager

[000234] In accordance with an embodiment, the application server plugin can be registered with the lifecycle manager using a register-plugin operation. Subsequently, upon registration of an application server (e.g., WLS) runtime with the lifecycle manager, the client interacting with lifecycle manager can make a REST call to the lifecycle manager to create/update/delete partition on the specified runtime registered with lifecycle manager. The application server (e.g., WLS) plugin can then be invoked to perform lifecycle operations on that particular application server (e.g., WLS) runtime.

10 Communication with the Application Server

[000235] In accordance with an embodiment, the application server plugin can use JMX/REST APIs exposed by the application server (e.g., WLS) to interact with the server. The plugin can implement the partition plugin contract from lifecycle manager to allow lifecycle operations on targeted application server runtimes to do CRUD operation on application server partitions. Following are some exemplary operations which can be done on a, e.g., WLS runtime for partitions:

| Lifecycle Operation | Description | Contract implemented by App Server Plugin |
|---------------------|---|--|
| createPartition | Creates WLS partition based on information provided in lifecycle context | public abstract LifecyclePartition create(String partitionName, LifecycleContext ctx, LifecycleRuntime runtime) throws LifecycleException; |
| updatePartition | Updates WLS partition based on information provided in lifecycle context | public abstract LifecyclePartition update (String partitionName, LifecycleContext ctx, LifecycleRuntime runtime) throws LifecycleException; |
| deletePartition | Deletes WLS partition based on information provided in lifecycle context | public abstract void delete(String partitionName, LifecycleContext ctx, LifecycleRuntime runtime) throws LifecycleException; |
| associatePartition | Associates WLS partition with a partition of other Type. [Example of other Type being OTD, database, etc] | public abstract void associate(LifecycleContext partition1Ctx, LifecycleContext partition2Ctx, LifecyclePartition partition1, LifecyclePartition partition2, LifecycleRuntime runtime) |

| | | |
|---------------------|---|--|
| | | throws LifecycleException; |
| dissociatePartition | Dissociate WLS partition from a partition of otherType. | public abstract void dissociate(LifecycleContext partition1Ctx, LifecycleContext partition2Ctx, LifecyclePartition partition1, LifecyclePartition partition2, LifecycleRuntime runtime) throws LifecycleException; |
| migratePartition | Migrates WLS partition to the newly available targets | public abstract LifecyclePartition migrate (String partitionName, LifecycleContext ctx, LifecycleRuntime runtime) throws LifecycleException; |

[000236] In accordance with an embodiment, the application server plugin can implement the runtime plugin contract from the lifecycle manager to allow lifecycle operations on targeted application server (e.g., WLS) runtimes, for example:

5

| Runtime Operation | Description | Contract implemented by App Server Plugin |
|-------------------|--|--|
| scaleup | Scales up weblogic server instance present in the pre-configured dynamic cluster in this weblogic domain | public abstract void scaleUp(String runtimeName, int scaleFactor, LifecycleContext ctx) throws LifecycleException; |
| scaledown | Scales down weblogic server instance present in the pre-configured dynamic cluster in this weblogic domain | public abstract void scaleDown(String runtimeName, int scaleFactor, LifecycleContext ctx) throws LifecycleException; |
| quiesce | Quiesces a weblogic server instance | public abstract void quiesce(String runtimeName, LifecycleContext ctx) throws LifecycleException; |

Runtime properties for Lifecycle Operations on Partitions

[000237] In accordance with an embodiment, the application server plugin uses the following properties to connect to an application server, e.g., WLS runtime. These properties 10 are passed on to the lifecycle operation via lifecycle runtime object as specified in the contract above:

```
Properties runtimeProperties = runtime.getRuntimeProperties();
```

| Property | Description | Property key |
|----------|----------------|---|
| Host | Weblogic host | runtimeProperties.getProperty("hostname") |
| Port | Weblogic port | runtimeProperties.getProperty("port") |
| username | Admin username | runtimeProperties.getProperty("username") |
| password | Admin password | runtimeProperties.getProperty("password") |

Lifecycle Context for WLS Plugin

[000238] In accordance with an embodiment, any CRUD operation done using an application server plugin can be customized based on the information/data provided to the plugin in the lifecycle context. The client can specify the properties as part of REST call. The lifecycle manager populates the data into the lifecycle context and passes the property back to the plugin. For example, when a client makes a REST call for createPartition for WLS partition, they can specify the properties to specify the attributes of partition.

10 **Lifecycle Context for create/update partition**

[000239] In accordance with an embodiment, following are some exemplary properties which can be specified by the client while create/update partition operation. These configurations are created as part of partition creation/update and are specific to the partition:

| Property | Description |
|------------------------|--|
| resourcegroup | Resource group Name. |
| resourcegrouptemplate | Resource group template. Assumption is the resource group template specified should be existing in the domain. |
| partitionPDBinfo | Pluggable database name for this partition. |
| datasource | datasource name for this specific PDB. |
| url | url for this specific PDB. |
| username | username for this specific PDB. |
| password | password for this specific PDB. |
| partitionproperty | partition property key/name. |
| partitionpropertyvalue | partition property value. |

15

[000240] In accordance with an embodiment, following are some exemplary properties which can be specified by the client while create/update partition. Failure to find these entities in the application server (e.g., WLS) Server would result in exception being thrown from the application server plugin:

20

| Property | Description |
|----------|-------------|
| | |

| | |
|---------------------------|--|
| resourcegroup:target | Target to be associated with this resourcegroup. [Comma separated list of targets] |
| securityrealm | Security realm to be associated with this partition |
| partition:defaulttarget | Comma separated list of targets |
| partition:availabletarget | Available targets [Comma separated list of targets] |

Lifecycle Context for Delete Partition

[000241] In accordance with an embodiment, no additional property other than partition name is needed to delete an application server (e.g., WLS) partition.

5

Lifecycle Context for Associate Partition

[000242] In accordance with an embodiment, association of an application server instance with another component may involve updating the application server (e.g., WLS) partition with some new properties, or may involve providing the necessary properties to the other component for association.

10

[000243] Any component which wants to participate in association needs to make sure to pass in the correct set of properties in lifecycle context in the format specified above.

15

[000244] In the case of WLS-OTD association, in this case, the OTD partition expects the cluster-name, cluster-members and target information to do the association. The application server plugin makes a JMX/REST call to the application server runtime to fetch the correct targeting information based on the partition name. These properties are added to the application server's lifecycleContext as part of associatePartition implementation. As part of an association call to the plugin, a complete snapshot of target configuration will be added to the lifecycle context. Then, the lifecycle manager calls the associatePartition implementation of OTD plugin, which uses the lifecycle context (which has the properties of WLS targets) to do the wiring of origin server pool and the application server partition.

20

[000245] In accordance with an embodiment, when an application server (e.g., WLS) partition is associated with an OTD partition, and if the partition is targeted to a configured cluster, if server is not running then the lifecycle manager will get configuration information available for Server's ListenAddress and Port and send it to OTD. If servers are running, then we will get the runtime information for Server available from RuntimeMBean and pass that to OTD via context information.

25

[000246] In a dynamic cluster, there is no configuration information available for ListenAddress, although the configuration time information is available for ports for dynamic servers. If dynamic servers are not running when an associate call happens then the lifecycle manager will get the configured port and the node manager's ListenAddress (if dynamic servers are associated with a machine) and send this to other plugins. If dynamic

servers are running when associate call happens then the lifecycle manager will get the server listen address and port from runtime Mbean and send it to the OTD plugins.

[000247] In the case of WLS-DB Association, in this case, the application server (e.g., WLS) plugin's associatePartition method uses the DB Partition's lifecycleContext to populate

5 PDB Info and update the WLS partition with the PDB info.

Lifecycle Context for Dissociate Partition

[000248] In accordance with an embodiment, dissociation of the application server (e.g., WLS) with another component may involve removing/updating certain WLS partition 10 configuration. Any component which wants to participate in dissociation needs to make sure to pass in the correct set of properties in lifecycle context in the format specified above. Following is an example which explains dissociation between WLS and DB partitions.

[000249] WLS-DB Dissociation: In this case, the WLS plugin's dissociatePartition method uses the DB partition's lifecycleContext to populate PDB Info name and update the 15 WLS partition by removing the PDB info. The WLS plugin expects the PDB Info name in the lifecycleContext.

[000250] WLS-OTD Dissociation: In this case, the OTD partition expects the cluster-name, cluster-members and target information to do the dissociation. WLS plugin makes JMX/REST call to the WLS runtime to fetch the correct targeting information based on 20 the partition name. These properties are added to the WLS's lifecycleContext as part of dissociatePartition implementation of WLS plugin.

WLS Partition Configuration Discovery and Association in an LCM Environment

[000251] In accordance with an embodiment, the partition configuration may be 25 created/updated/deleted outside of the lifecycle manager environment/scope. For example, an Administrator might invoke WLST command to create a partition, or use a WLS admin console to update a partition. In any circumstance, it needs to be made sure that the partition information present with the lifecycle manager is recent and the most updated one.

30 Effect of Lifecycle Manager Runtime Scaling on Application Server Configurations

[000252] In accordance with an embodiment, a ScaleUp and scaleDown operation in the application server (e.g., WLS) plugin can be used to provide the functionality of starting and stopping the server instances which are part of the dynamic cluster. The WLS plugin might use node manager to interact with the machines/servers to start/stop the servers.

[000253] Elasticity frameworks can make a scale up or scale down call to lifecycle 35 manager with the cluster information, scale factor, runtime type (i.e. WLS) Interceptors which are registered with the scaleUp/scaleDown are invoked before the scaleup/scaledown

operation to validate and perform the necessary infrastructure needed for the cluster to scale up.

[000254] For example, to make the necessary calls to WLS runtime to determine/create the number of datasources needed for the new servers to be brought up in the cluster. Once the pre-requisites to scaleup/scaledown the cluster is satisfied (i.e., any pre-invoke of the interceptors is satisfied), the real scaleup or scale down happens. The interceptors are called again after the operation has completed in the plugin to satisfy the post-invoke requirements.

[000255] For example, assuming that a number of interceptors will be used in order to scale a WLS cluster, described below is one such sequence.

[000256] Assume that the following interceptors have been added: WLS Interceptor, OTD Interceptor, JMS Interceptor, DataSource Interceptor when WLSScaleDownAction from Elasticity module calls lcm.scaleDown(clusterName, factor, payload): 1. WLS Interceptor checks if scaling will be within cluster limits. It then calls ctx.proceed(); 2. OTD Interceptor will make the necessary OTD config changes to mark the selected targets to be drained. OTD Interceptor will wait till OTD metrics indicate that the requests have been drained. It then calls ctx.proceed(); 3. JMS Interceptor will wait for the JMS backlogs to drain. It then calls ctx.proceed(); 4. DataSource Interceptor will ensure that all JDBC Connections have been returned to the pool. It then calls ctx.proceed(); 5. The service is scaled down (called by ElasticServiceManager after the last interceptor).

[000257] On the return path of Interceptor call chain: 1. DataSourceInterceptor will increase the max-connection-pool-size appropriately for the remaining alive managed servers; 2. JMS Interceptor does nothing; 3. OTD Interceptor does nothing; 4. WLS Interceptor does nothing.

Invocation of Application Server Plugin using REST Requests

[000258] In accordance with an embodiment, lifecycle manager clients can use REST requests to do operations on the WLS, including for example: 1. Register a WLS runtime and the properties to be passed in the REST request; 2. Create empty WLS partition; 3. Create WLS partition with resource group and targeting; 4. Add WLS Partition to environment (assuming environment test1 is already created); 5. Associate WLS Partition with another partition; 6. Dissociate WLS Partition with another partition; 6. Delete WLS Partition; 7. Update WLS Partition.

B. Database Plugin

[000259] In accordance with an embodiment, the system can include a database (e.g., an Oracle Database) lifecycle plugin which implements the lifecycle APIs and manages the

creation or configuration of database (e.g., Database) partitions.

[000260] In accordance with an embodiment, the multitenant-aware architecture of a database component such as the Oracle Database enables it to function as a multitenant container database (CDB) that includes customer-created pluggable databases (PDBs).

5 Each PDB acts as a portable collection of schemas, schema objects, and nonschema objects. From the perspective of the lifecycle manager module, the database component manages its own partition configuration using the PDB features built into the database machinery. The database lifecycle plugin can leverage this functionality to perform PDB management operations when called upon by the lifecycle manager.

10

Interaction with Lifecycle Manager

[000261] In accordance with an embodiment, the database plugin is an implementation of the lifecycle manager partition plugin interface, and is registered with lifecycle manager before it can perform any operations. Subsequently, upon registration of a database runtime

15 with the lifecycle manager, the client interacting with lifecycle manager would make a REST call to the lifecycle manager to perform PDB operations on the specified runtime registered with lifecycle manager.

Runtime Registration

20 **[000262]** In accordance with an embodiment, before an operation can be performed on a database, it has to be registered with the lifecycle manager. The registration process requires information which is used to connect to the database. For example, a database runtime can be registered with the following properties:

| Name | Description |
|---------------|---|
| connectionURL | URL to connect to the database. |
| username | User that has administrative privileges to perform PDB management operations on the database. |
| password | Password for the database user. |

25

Lifecycle operations on the database

[000263] In accordance with an embodiment, the database lifecycle plugin communicates with the database using a remote SQL connection via JDBC. Properties

30 specified during the runtime registration are used to connect to the database. A number of lifecycle operations map to corresponding SQL commands, various clauses of which are determined from the properties passed on in the lifecycle context object, and passed to the

plugin, for example:

| Lifecycle partition Operation | Plugin Operation |
|-------------------------------|---|
| create | Execute the CREATE PLUGGABLE DATABASE SQL Command |
| delete | Execute the DROP PLUGGABLE DATABASE SQL Command |
| update | Execute the ALTER PLUGGABLE DATABASE SQL Command |
| associate | Associate the database partition with a WLS partition |
| quiesce | Execute the ALTER PLUGGABLE DATABASE CLOSE command |
| start | Execute the ALTER PLUGGABLE DATABASE OPEN command |

[000264] In accordance with an embodiment, the following are details of the above

5 operations, with examples:

Create

[000265] In accordance with an embodiment, the create operation performs the following tasks: (1) Create a pluggable database on the specified database runtime. The partition ID of the newly created PDB is obtained from the V\$PDBS table. The partition name is the same as the PDB name; (2) Optionally, create a database user; (3) Optionally, create a CDB plan directive for the PDB using an existing CDB plan. A number of options are available for the above tasks. These options are passed as properties in the lifecycle context object to the plugin.

15

Delete

[000266] In accordance with an embodiment, the delete operation can perform the following tasks: (1) drop the pluggable database, using the PDB name. PDB users will also be dropped automatically; (2) Optionally, drop CDB plan directive for the PDB Sample SQL executed by the plugin for above tasks:

Update

[000267] In accordance with an embodiment, the update operation modifies the attributes of a pluggable database, by executing an ALTER PLUGGABLE DATABASE command to perform the updates to the PDB.

Associate

[000268] In accordance with an embodiment, the associate operation is relevant when

a WLS partition is associated with a PDB. The database plugin passes on the database connection URL to the WLS plugin, during association. If a user was created while creating the PDB by lifecycle manager, then the username and password of that user is also passed to the WLS partition plugin.

5

Quiesce

[000269] In accordance with an embodiment, the quiesce operation closes the specified PDB using the ALTER PLUGGABLE DATABASE command.

10

Start

[000270] In accordance with an embodiment, the start operation starts the specified PDB using the ALTER PLUGGABLE DATABASE command.

Lifecycle Context Schema

15

[000271] In accordance with an embodiment, various SQL clauses can be formed from properties passed to the plugin via a lifecycle context object. Some of the common properties include the following:

| Property | Description |
|--------------------|--|
| name | Name of the pluggable database. |
| admin user | User name for the PDB. |
| identified by | Password for the PDB admin user. |
| storage | STORAGE clause for the SQL statement creating the PDB. |
| default tablespace | DEFAULT TABLESPACE clause for the SQL statement creating the PDB. |
| file_name_convert | FILE_NAME_CONVERT clause for the SQL statement creating the PDB. |
| tempfile_reuse | TEMPFILE_REUSE clause for the SQL statement creating the PDB. |
| roles | ROLES clause for the SQL statement creating the PDB. |
| from | FROM clause for the SQL statement creating the PDB. This is used when cloning a PDB. |

20

[000272] In accordance with an embodiment, the SQL commands and various clauses can be specified as lifecycle context properties. Depending on support in the database, most of these can be used when altering the PDB as well.

Resource Manager Properties

[000273] In accordance with an embodiment, when creating a PDB, a CDB plan directive can also be created. Properties related to this should be specified under the key "resourceManager". Following are exemplary properties for the "resourceManager"

| Property | Description |
|-----------------------|---|
| plan | Name of the existing CDB plan. This is required. |
| shares | Specifies the share of resource allocation for the pluggable database. |
| utilization_limit | Specifies the maximum percentage of CPU that the pluggable database can utilize. |
| parallel_server_limit | Specifies the maximum percentage of parallel servers that the pluggable database can use. |

5

User Properties

[000274] In accordance with an embodiment, when creating a PDB, a user can also be created. The properties related to the user should be specified under the key "databaseUser".

10

[000275] **Figure 16** illustrates a flowchart of method for supporting lifecycle plugins in a multitenant application server environment, in accordance with an embodiment.

[000276] As illustrated in Figure 16, at step 680, an application server environment is provided at one or more computers, including a domain for execution of software applications, which supports one or more partitions, wherein each partition provides an administrative and runtime subdivision of the domain, and wherein a partition can optionally include one or more resource groups having a collection of deployable applications or resources and/or referencing a resource group template.

15

[000277] At step 682, a lifecycle manager is provided, which enables configuration integration across different components, for example application servers, traffic directors or load balancers, and database or other components, which may have their own partition-specific configurations.

20

[000278] At step 684, a plurality of components are provided, each of which includes a lifecycle plugin that provides a bridge between the lifecycle manager and that component.

25

[000279] At step 686, during creation or configuration of a partition, the plurality of lifecycle plugins implement an interface that allows the lifecycle manager to perform lifecycle operations that can call a plugin, so that the plugin can operate with its associated component to perform an actual creation or configuration of the component.

* * * * *

5 Global Runtime

[000280] In accordance with an embodiment, the system supports the use of a global runtime within the multi-tenant application server environment.

[000281] As described above, in accordance with various embodiments, such as those that use Fusion Apps or Fusion Middleware, there may be requirements for running things in 10 a global context for all partitions.

[000282] For example in a Metadata Services (MDS) component which is provided as part of Fusion Middleware (FMW), there is a requirement that there be one global cache for all of the shared documents and then partition specific caches for the customizations that are specific to the tenants. In addition to having such a global cache, there is also the 15 requirement for notifying all the partitions when the cache has been purged or updated, so that the MDS instance running in the partitions can apply the customizations on the updated cache.

[000283] Additionally, at a high level in a multi-tenant environment there may be artifacts or other components that need to be shared across the entire domain, including any 20 partitions, referred to herein as global artifacts. These artifacts reside or run at the global or domain level in a multitenant container which is referred as the global runtime. The partition specific artifacts or components are the ones that are reside or run in the context of a partition.

[000284] In accordance with an embodiment, the global runtime does not run in any 25 specific partition, but is instead scoped to a domain. When running in, for example, a Java Required Files (JRF) domain, with the tenant mapping for Fusion Applications (FA) and multi-tenancy flag enabled, code running in the global runtime can switch to the context of a tenant and execute something on behalf of a tenant via a tenant switching (e.g., switchTenant) API.

[000285] In accordance with an embodiment, one or more resource groups can also be 30 specified at the domain level for use in the global runtime. Targeting of domain-level resource groups is set at the resource group, not at the individual resources within the resource group. All resources in the group will be targeted to all targets specified for the resource group. When targeted to a cluster, the applications run on all the managed servers in the cluster.

[000286] In accordance with an embodiment, at the domain level the resource group is a convenient way to group related resources. The system manages the resources declared

in a domain-level resource group the same as ungrouped resources: for example, they are started during system start-up and stopped during system shut-down. Further, an administrator can stop, start, or remove any single resource in a group individually and can also act on all the resources in a group implicitly by operating on the group.

5 [000287] For example, stopping the group stops all the resources in the group that are not already stopped. Starting the group starts any resources in the group that are not already started. Removing the group removes all the resources contained in the group.

10 [000288] In accordance with an embodiment, the global runtime resource group(s) must be started before all the other resource groups defined in partition(s). This allows the global runtime to be initialized before any of the partitions and, if necessary to load shared classes using the shared class loader to share among partitions or other applications / resource groups it can do so before initializing the partitions.

15 [000289] In accordance with an embodiment, when running in the global runtime, the partition name will return GLOBAL, and the id will be 0.

20 [000290] When running in a JRF domain with the multi-tenancy flag enabled, the tenant context will have tenancy set as GLOBAL, and the tenant id will also be 0.

25 [000291] In accordance with an embodiment, resource groups defined at the global runtime must be targeted to the same managed servers / clusters as the partitions that applications running in the global runtime may want to operate upon or share classes with.

30 When migrating a partition (resource groups) from one cluster to another, the global runtime must either be running on the global runtime where the partition is being migrated, or, the migration feature must ensure that the resource groups defined at the global runtime are migrated prior to migrating the resource group of a partition.

35 [000292] **Figure 17** illustrates use of a global runtime in a multitenant application server environment, in accordance with an embodiment.

40 [000293] In accordance with an embodiment, the global runtime can be used for those applications which are intended to run at the domain level. If a particular component or application wants to use executeAs, it must run at the domain level in the global runtime, wherein the global runtime is the same as or acts as the domain runtime.

45 [000294] In accordance with an embodiment, the ExecuteAs and tenant onboarding are provided as a JRF type functionality, for upper stack usage. An application running in the global runtime can be seen by the various partitions depending on the security settings afforded those partitions. Similarly, an application running at the domain (global) level with its own UUID can be made available to the partitions.

50 [000295] For example, a global PDB can be shared across tenants, wherein the global PDB is set up by the upper stack, which is a common configuration setup used with SaaS usecases, and not as commonly with consolidation or PaaS use cases.

[000296] As illustrated in Figure 17, in accordance with an embodiment, the global runtime 700 can include one or more applications, and can use a tenant switching API 702 to perform an executeAs operation 704, 705 on behalf of a tenant, to access or operate on behalf of, one or more partitions 710, 712, 714.

5 **[000297]** In Fusion Apps and Fusion Middleware scenarios there are requirements for running things in the global context for all Partitions. For example, as described above in the context of Metadata Services (MDS), there is a requirement that there be one global cache for all the shared documents and then partition specific caches for the customizations that are specific to the tenants; and also that all of the partitions be notified when the cache has
10 been purged or updated so that the MDS instance running in the partitions can apply the customizations on the updated cache.

Partition and Tenant Name and ID

15 **[000298]** As described above, in accordance with an embodiment, when running in the global runtime, the partition name will return GLOBAL, and the ID will be 0. When running in a JRF domain with the multitenancy flag enabled, the tenant context will be have tenancy set as GLOBAL, and the tenant ID will also be 0.

Migration and Global Runtime

20 **[000299]** As described above, in accordance with an embodiment, resource groups defined at the global runtime must be targeted to the same managed servers / clusters, since the partitions that applications running in the global runtime may want to operate upon or share classes with.

25 **[000300]** In accordance with an embodiment, when migrating a partition (or its resource groups) from one cluster to another, the global runtime must either be running on the cluster where the partition is being migrated, or, the migration feature must ensure that the resource groups defined at the global runtime are migrated prior to migrating the resource group of a partition.

SaaS scenario Example

30 **[000301]** The example below shows use of resource group at the domain level in the SaaS scenario where a Customer Relationship Management (CRM) application is deployed to the global runtime. The resource group template defined at the domain level can be used by both the global runtime as well as partitions. For readability, the example below shows just the resource group at the domain level:

35

```
<domain>
<resource-group-template>
```

```
<name>CRMResourceGroupTemplate</name>
<app-deployment>
<name>CRM</name>
...
5  </app-deployment>
<app-deployment>
<name>physician</name>
...
</app-deployment>
10 <app-deployment>
<name>browser-starter</name>
...
</app-deployment>
<app-deployment>
15 <name>SamplesSearchWebApp</name>
...
</app-deployment>
<jms-server>
<name>CRMJMSServer</name>
20 </jms-server>
<mail-session>
<name>mail/CRMMailSession</name>
<jndi-name>mail/CRMMailSession</jndi-name>
<properties>mail.user=joe;mail.host=mail.mycompany.com</properties>
25 </mail-session>
<jms-system-resource>
<name>CRM-jms</name>
<sub-deployment>
<name>CRMJMSServer</name>
30 </sub-deployment>
<descriptor-file-name>jms/CRM-jms.xml</descriptor-file-name>
</jms-system-resource>
<jdbc-system-resource>
<name>CRMGlobalDataSourceXA</name>
35 <descriptor-file-name>jdbc/CRMGlobalDataSourceXA-9669-jdbc.xml</descriptor-file-n
ame
</jdbc-system-resource>
<saf-agent>
```

```

<name>WsrmAgent</name>
<service-type>Both</service-type>
</saf-agent>
</resource-group-template>
5  <resource-group>
<name>CRM</name>
<resource-group-template>CRMResourceGroupTemplate</resource-group-template>
</resource-group>
</domain>
10

```

10 [000302] In accordance with an embodiment, applications and shared libraries can be deployed to a resource group at the domain level as well.

Global PDB and Tenant-Specific PDB

15 [000303] In accordance with an embodiment, when running, e.g., an Oracle 12c Database, many PDBs can run in one container database (CDB). With the creation of a partition for a tenant, a PDB will be created and the data source configured for the partition to use.

20 [000304] In accordance with an embodiment, in addition to the tenant-specific PDB there can optionally be a global PDB that is shared across all tenants. For example, the metadata (documents) in MDS can be stored in a global PDB and the tenant-specific customizations in the tenant-specific PDBs.

25 [000305] In accordance with an embodiment, in order to be able to use the shared data from the global PDB, the partitions need access to both the PDBs (global and tenant-specific). Each partition can be configured with an additional, read only data source to the global PDB to be able to access the shared data, and also a data source to the tenant-specific PDB that will be used to store tenant-specific data. The shared application code that is running in the global runtime will be able to write to the global PDB and will be configured with a read / write data source to the global PDB.

30 [000306] **Figure 18** further illustrates use of a global runtime in a multitenant application server environment, in accordance with an embodiment.

[000307] As illustrated in Figure 18, in accordance with an embodiment, each partition can be configured with one or more data sources 701, 711, 721, including, for example, that the global runtime has access to the global PDB 702 via data source DS2.

35 [000308] As further illustrated in Figure 18, the tenants that run in Partition 1 (710) and Partition 2 (720) have read / write access via DS1 to their specific PDBs (712, 722 respectively) and a read only DS2 (same data source shared across all partitions) to access

the read only data in the global PDB.

DataSources

[000309] In accordance with an embodiment, each partition will have a data source that points to its PDB instance for writing tenant-specific data. When running in a partition in an application server, e.g., WLS, the name of the data source does not have to change for each partition. Instead, the application server isolates the data source in the JNDI tree on a per partition basis. These datasources can be defined in the resource group template, which will be referenced from the resource group for each tenant's partition.

10

Executing at Global Runtime and Partition

[000310] In accordance with an embodiment, applications or shared libraries that are deployed to the global runtime can be shared with all tenant-specific applications running in a partition. In addition to being shared across the partition, applications running in the global runtime can also execute in the context of a partition running in a specified tenant's context. However the reverse is not true; a tenant-specific application running in a partition cannot execute in the context of the global runtime. The table below shows which context switching is allowed:

| Source Scope of application | Target scope of application | Permit / Denied |
|-----------------------------|-----------------------------|-----------------|
| Global runtime | Global runtime | Allowed |
| Global runtime | Partition | Allowed |
| Partition | Global runtime | Denied |
| Partition 1 | Partition 1 | Allowed |
| Partition 1 | Partition 2 | Denied |

20

[000311] **Figure 19** illustrates a flowchart of method for supporting a global runtime in a multitenant application server environment, in accordance with an embodiment.

[000312] As illustrated in Figure 19, at step 740, an application server environment is provided at one or more computers, including a domain for execution of software applications, which supports one or more partitions, wherein each partition provides an administrative and runtime subdivision of the domain, and wherein a partition can optionally include one or more resource groups having a collection of deployable applications or resources and/or referencing a resource group template.

30

[000313] At step 742, a global runtime is provided, which does not run in any specific partition, but is instead scoped to a domain, wherein one or more resource groups can

optionally be specified at the domain level for use by the global runtime.

[000314] At step 744, upon receiving a request to execute a function by the global runtime on behalf of a tenant, the global runtime can switch to the context of the tenant via a tenant switching interface, and execute the function as or on behalf of that tenant.

5 **[000315]** The present invention may be conveniently implemented using one or more conventional general purpose or specialized digital computer, computing device, machine, or microprocessor, including one or more processors, memory and/or computer readable storage media programmed according to the teachings of the present disclosure. Appropriate software coding can readily be prepared by skilled programmers based on the 10 teachings of the present disclosure, as will be apparent to those skilled in the software art.

15 **[000316]** In some embodiments, the present invention includes a computer program product which is a non-transitory storage medium or computer readable medium (media) having instructions stored thereon/in which can be used to program a computer to perform any of the processes of the present invention. The storage medium can include, but is not limited to, any type of disk including floppy disks, optical discs, DVD, CD-ROMs, microdrive, and magneto-optical disks, ROMs, RAMs, EPROMs, EEPROMs, DRAMs, VRAMs, flash memory devices, magnetic or optical cards, nanosystems (including molecular memory ICs), or any type of media or device suitable for storing instructions and/or data.

20 **[000317]** The foregoing description of the present invention has been provided for the purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise forms disclosed. Many modifications and variations will be apparent to the practitioner skilled in the art. The modifications and variations include any relevant combination of the disclosed features. The embodiments were chosen and described in order to best explain the principles of the invention and its practical application, thereby 25 enabling others skilled in the art to understand the invention for various embodiments and with various modifications that are suited to the particular use contemplated. It is intended that the scope of the invention be defined by the following claims and their equivalents.

Claims:

What is claimed is:

- 5 1. A system for supporting a global runtime in a multitenant application server environment, comprising one or more computers, comprising an application server environment executing thereon, together with
 - a plurality of resources which can be used within the application server environment;
 - one or more partitions, wherein each partition provides a subdivision of a domain, that
- 10 can be associated with a tenant; and
 - a global runtime which is scoped to the domain so that one or more resources or resource groups can be specified at the domain level for use in the global runtime.
2. The system of claim 1, wherein one or more resource group templates are used to define groupings of deployable resources within the domain.
- 15 3. The system of any of claims 1 to 2, wherein targeting of domain-level resource groups can be set at a resource group, so that all resources within that resource group will be targeted to all targets specified for that resource group.
- 20 4. The system of any of claims 1 to 3, wherein, when a resource group is targeted to a cluster, applications within that resource group are run on all managed servers in that cluster.
- 25 5. A method for supporting a global runtime in a multitenant application server environment, comprising:
 - providing, at one or more computers, comprising an application server environment executing thereon,
 - a plurality of resources which can be used within the application server environment,
 - one or more partitions, wherein each partition provides a subdivision of a domain, that
- 30 can be associated with a tenant; and
 - providing a global runtime which is scoped to the domain so that one or more resources or resource groups can be specified at the domain level for use in the global runtime.
- 35 6. The method of claim 5, wherein one or more resource group templates are used to define groupings of deployable resources within the domain.

7. The method of any of claims 5 to 6, wherein targeting of domain-level resource groups can be set at a resource group, so that all resources within that resource group will be targeted to all targets specified for that resource group.

5 8. The method of any of claims 5 to 7, wherein, when a resource group is targeted to a cluster, applications within that resource group are run on all managed servers in that cluster.

9. A non-transitory computer readable storage medium, including instructions stored thereon which when read and executed by one or more computers cause the one or more 10 computers to perform the method comprising:

providing an application server environment, together with a plurality of resources which can be used within the application server environment, one or more partitions, wherein each partition provides a subdivision of a domain, that can be associated with a tenant; and

15 providing a global runtime which is scoped to the domain so that one or more resources or resource groups can be specified at the domain level for use in the global runtime.

10. The non-transitory computer readable storage medium of claim 9, wherein one or 20 more resource group templates are used to define groupings of deployable resources within the domain.

11. The non-transitory computer readable storage medium of any of claims 9 to 10, wherein targeting of domain-level resource groups can be set at a resource group, so that all 25 resources within that resource group will be targeted to all targets specified for that resource group.

12. The non-transitory computer readable storage medium of any of claims 9 to 11, wherein, when a resource group is targeted to a cluster, applications within that resource 30 group are run on all managed servers in that cluster.

13. A computer program comprising program instructions for execution on one or more computer systems, wherein the program instructions when executed causes the one or more computer systems to perform the method of any of claims 5 to 8.

35

14. A computer program product comprising the computer program of claim 13 stored in a non-transitory machine readable data storage medium.








FIGURE 1

**FIGURE 2**

**FIGURE 3**

**FIGURE 4**

**FIGURE 5**









FIGURE 6

**FIGURE 7**

**FIGURE 8**

**FIGURE 9**

**FIGURE 10**






FIGURE 12




FIGURE 13




FIGURE 14





FIGURE 15

**FIGURE 16**




FIGURE 17





FIGURE 18

**FIGURE 19**

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2015/052459

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F9/50
ADD.

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

B. FIELDS SEARCHED

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

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

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

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
|-----------|--|-----------------------|
| X | WO 2014/039918 A1 (ORACLE INT CORP [US]) 13 March 2014 (2014-03-13) paragraphs [0008], [0044], [0056], [0057], [0005], [0040], [0043], [0039], [0137] ----- | 1-14 |
| X | WO 2014/039892 A1 (ORACLE INT CORP [US]) 13 March 2014 (2014-03-13) paragraphs [0033], [0029], [0026], [0007], [0040], [0034], [0032], [0052] ----- | 1-14 |
| X | WO 2014/039919 A1 (ORACLE INT CORP [US]) 13 March 2014 (2014-03-13) paragraphs [0035], [0040], [0053], [0073], [0139], [0133], [0134], [0137] ----- -/- | 1-14 |

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier application or patent but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

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

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

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

"&" document member of the same patent family

| | |
|--|--|
| Date of the actual completion of the international search | Date of mailing of the international search report |
| 16 December 2015 | 04/01/2016 |
| Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016 | Authorized officer Simion, C |

INTERNATIONAL SEARCH REPORT

| |
|------------------------------|
| International application No |
| PCT/US2015/052459 |

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
|-----------|--|-----------------------|
| X | US 2014/280975 A1 (MORDANI RAJIV [US] ET AL) 18 September 2014 (2014-09-18) paragraphs [0022], [0023], [0025], [0030], [0032], [0041], [0059] ----- | 1-14 |

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2015/052459

| Patent document cited in search report | Publication date | Patent family member(s) | | | Publication date |
|--|------------------|-------------------------|---------------|-------|------------------|
| WO 2014039918 | A1 13-03-2014 | CN | 104737517 A | | 24-06-2015 |
| | | EP | 2893683 A1 | | 15-07-2015 |
| | | WO | 2014039918 A1 | | 13-03-2014 |
| ----- | ----- | ----- | ----- | ----- | ----- |
| WO 2014039892 | A1 13-03-2014 | US | 2014075412 A1 | | 13-03-2014 |
| | | WO | 2014039892 A1 | | 13-03-2014 |
| ----- | ----- | ----- | ----- | ----- | ----- |
| WO 2014039919 | A1 13-03-2014 | CN | 104756460 A | | 01-07-2015 |
| | | EP | 2893684 A1 | | 15-07-2015 |
| | | JP | 2015529366 A | | 05-10-2015 |
| | | US | 2014075032 A1 | | 13-03-2014 |
| | | WO | 2014039919 A1 | | 13-03-2014 |
| ----- | ----- | ----- | ----- | ----- | ----- |
| US 2014280975 | A1 18-09-2014 | | NONE | | |
| ----- | ----- | ----- | ----- | ----- | ----- |