

(51) International Patent Classification:
G06F 9/50 (2006.01)(21) International Application Number:
PCT/US2018/059913(22) International Filing Date:
09 November 2018 (09.11.2018)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/587,194 16 November 2017 (16.11.2017) US
16/115,414 28 August 2018 (28.08.2018) US(71) Applicant: **MICROSOFT TECHNOLOGY LICENSING, LLC** [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).(72) Inventors: **FONTOURA, Marcus Felipe**; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **BIANCHINI, Ricardo Gouvea**; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **BABLANI, Girish**; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US).(74) Agent: **MINHAS, Sandip S.** et al.; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US).

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

(54) Title: PREDICTIVE RIGHTSIZING FOR VIRTUAL MACHINES IN CLOUD COMPUTING SYSTEMS




FIG. 2D

(57) Abstract: Various embodiments, methods, and systems for implementing a predictive rightsizing system are provided. Predicted rightsized deployment configurations are generated for virtual machine "VM" deployments having deployment configurations that are modified to predicted rightsized deployment configurations based on a prediction engine. In operation, a VM deployment, associated with a request to deploy one or more VMs on a node, is accessed at a predictive rightsizing controller. A predicted resource utilization for the VM deployment is generated at the prediction engine and accessed at the predictive rightsizing controller. The predicted resource utilization is generated based on a prediction engine that uses past behaviors and features associated with previous VM deployments. Based on the predicted resource utilization, a predicted rightsized deployment configuration is generated for the VM deployment. The predicted rightsized deployment configuration is communicated such that the request to deploy the VMs is executed based on the predicted rightsized deployment configuration.

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

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

PREDICTIVE RIGHTSIZING FOR VIRTUAL MACHINES IN CLOUD COMPUTING SYSTEMS

FIELD

5 **[0001]** The present disclosure relates generally to cloud computing systems and more particularly to resource management for virtual machines in cloud computing systems.

BACKGROUND

[0002] Users rely on applications and services to perform various computing tasks. Cloud computing systems may offer building, deployment and management functionality for different types of applications and services. Cloud computing system providers may
10 implement virtual systems (e.g., virtual machines “VMs”) to maximize resource usage. Virtualization can support sharing and managing hardware, allowing for multiple environments which are isolated from one another, yet exist on the same physical machine. Effective resource management of cloud computing systems, including virtual
15 systems, can be challenging. For example, individual VMs are supposed to have the right levels of central processing, memory and disk space. However, resources allocated to VMs may at times be under-utilized, which results in inefficiency.

SUMMARY

[0003] Embodiments of the present invention relate to methods, systems and computer
20 storage media for providing predicted rightsized deployment configurations using a predictive rightsizing system. In particular, predicted rightsized deployment configurations are generated for virtual machine “VM” subscriptions having deployment configurations that are modified to predicted rightsized deployment configurations based on a prediction engine of a resource manager. The prediction engine may use a machine learning model
25 based on past behaviors and features associated with previous VM deployments to make resource utilization predictions (e.g., maximum resource utilization predictions) for VM deployments.

[0004] In operation, at a high level, a predictive rightsizing controller can communicate to a resource manager VM deployment. The VM deployment is associated with a request
30 to deploy one or more VMs on one or more nodes of a cloud computing system. The resource manager has a prediction engine that operates to predict resource utilizations for VM deployments based on past behaviors and features of previous VM deployments. The prediction engine generates a predicted resource utilization for the VM deployment. The predicted resource utilization can specifically be a predicted maximum resource utilization

of the VM deployment. The predictive rightsizing controller receives the predicted resource utilization. Based on the predicted resource utilization, the predictive rightsizing controller generates a predicted rightsized deployment configuration for the VM deployment. The predicted rightsized deployment configuration is communicated such
5 that the request to deploy one or more VMs can be executed based on the predicted rightsized deployment configuration.

[0005] In embodiments, the predicted resource utilization is associated with a confidence score for the predictions (“confidence score”). The confidence score can be associated with a confidence score threshold. The confidence score and the confidence score
10 threshold can be used to determine whether to generate or recommend deploying the predicted rightsized deployment configuration of the VM deployment. Further, a customer may choose to opt-in to automatically execute the request to deploy the VM deployment based on the predicted rightsized deployment configuration, while another customer may simply have the predicted rightsized deployment configuration communicated to them via
15 an interface.

[0006] Further areas of applicability of the present disclosure will become apparent from the detailed description, the claims and the drawings. The detailed description and specific examples are intended for purposes of illustration only and are not intended to limit the scope of the disclosure.

20 BRIEF DESCRIPTION OF DRAWINGS

[0007] FIG. 1 is a simplified schematic of an example of a resource management system including a prediction engine and a virtual machine allocator for a cloud computing system.

[0008] FIG. 2A is a simplified schematic of an example of the prediction engine of the
25 resource management system of FIG. 1.

[0009] FIG. 2B is a simplified schematic of an example bucketizing chart for the resource management system of FIG. 1.

[0010] FIG. 2C is a simplified schematic of an example of feature score chart of the resource management system of FIG. 1.

[0011] FIG. 2D is a flow diagram showing an exemplary method for implementing a predictive rightsizing system, in accordance with embodiments described herein.

[0012] FIG. 3 is a functional block diagram of a simplified example of a cloud computing system that utilizes the resource management system of FIG. 1 to deploy virtual machines.

[0013] FIG. 4 is a functional block diagram of a simplified example of a data center shown in FIG. 3 that utilizes the resource management system of FIG. 1 to deploy virtual machines.

5 [0014] FIG. 5 is a functional block diagram of a simplified example of a cluster shown in FIG. 3.

[0015] FIG. 6 shows a simplified example of a healing buffer and a scale-out/growth buffer in a cluster.

[0016] FIG. 7 shows examples of different types of virtual machines.

10 [0017] FIG. 8 is a functional block diagram of a simplified example of a distributed network system.

[0018] FIG. 9 is a functional block diagram of a simplified example of a client device used in the distributed network system of FIG. 8.

[0019] FIG. 10 is a functional block diagram of a simplified example of a server used in the distributed network system of FIG. 8.

15 DESCRIPTION

[0020] Cloud computing system providers often rely on resource management tools to support efficient utilization of cloud computing system resources. The efficiency of VM allocation and resource management in the cloud computing system can be significantly improved if accurate predictions are available about future VM behaviors. Historical data about VMs can be leveraged to produce such predictions. For example, the historical data may indicate that many VM roles exhibit consistent peak CPU utilizations over time; VM roles that execute user-facing workloads consistently do so over time; tenant deployments, deployment sizes or deployment configurations are unlikely to vary widely over time; and so on. In these cases, past behaviors are reasonable predictors of future behaviors.

25 [0021] The resource management system (of U.S. Patent Application Serial No. 15/424,305, filed February 03, 2017, entitled “RESOURCE MANAGEMENT FOR VIRTUAL MACHINES IN CLOUD COMPUTING SYSTEMS”) discloses a system that learns from past behaviors and, when requested, produces predictions of future behaviors. For example, before selecting servers to run a particular VM deployment, the VM allocator can obtain from the resource management system any predictions on the potential behavior of the VMs that the VM allocator is about to allocate. With this information, the VM allocator can make a more informed set of selections.

30 [0022] The resource management system can further address another specific issue in cloud computing systems, VM rightsizing. VM rightsizing can generally refer to assessing

and optimizing allocated VM resources. For example, CPU, memory, disk space, networking and other identified resources that are part of a VM can be monitored and changed as part of resource management to efficiently utilize resources in the cloud computing system. Rightsizing can ensure maximum performance of workloads and efficient use of underlying hardware.

[0023] Conventional rightsizing recommendations usually include rightsizing systems that monitor each virtual machine for a defined period of time (e.g., number of days) before issuing a recommended rightsize for the VM. It is common to monitor a VM for a period of time that would usually provide a comprehensive view of the lifecycle of the VM. For example, a VM can be monitored for 7, 8 or even 14 days prior to recommending a rightsize for the VM and making changes to the allocated resources of the VM. In this regard, if a customer requests a virtual machine deployment for peak performance, the rightsizing system can evaluate the VM for a period of time to understand what the peak performance behaviors are, and resource utilizations for the VM.

[0024] It may be possible that only a small percentage of virtual machines remain alive for the duration of the defined period for monitoring the VM to determine the recommended rightsized VM. It is common, for an exemplary cloud computing system, to have a significant number (e.g., 90%) of VMs last just 1 day, while a small percent (e.g., 4%) last 8 days and even smaller percent (e.g. 3%) last 14 days or more. Accordingly, the conventional model of monitoring a VM to make rightsizing recommendations for the particular VM seems flawed, in that, most VMs do not last long enough to enable a recommendation.. As such, a comprehensive rightsizing system that can determine and implement recommended rightsized deployment for virtual machines while circumventing the conventional monitoring period would improve cloud computing system resource management with regard to rightsizing VMs.

[0025] Embodiments of the present invention are directed to simple and efficient methods for providing predicted rightsized deployment configurations using a predictive rightsizing system. In particular, predicted rightsized deployment configurations are generated for virtual machine “VM” deployments having deployment configurations that are modified to predicted rightsized deployment configurations based on a prediction engine of a resource manager. Based on past behaviors and features associated with previous VM deployments, the prediction engine uses a machine learning model to make resource utilization predictions (e.g., maximum resource utilization predictions) for VM deployments.

[0026] At a high level, a customer may request a VM deployment (i.e., tenant deployment). The VM deployment can include 10 A4s, where an A4 can be a type of VM having specific resource types (e.g., CPU, memory, disk, network card, etc.). The 10 A4s have a deployment configuration (i.e., identified or allocated resources for resource types of the 10 A4s). Based on receiving the request to deploy the VM deployment, a rightsizing operation, at a predictive rightsizing controller, is triggered for the VM deployment. The rightsizing operation can include a request to a resource manager for a predicted resource utilization for the VM deployment. The predicted resource utilization can specifically be a predicted maximum resource utilization of the VM deployment. The predicted maximum resource utilization can be, for example, 10% CPU, 50% memory, 20% disk, and 5% network. Based on the predicted resource utilization, a predicted rightsized deployment configuration (e.g., 5 A2s) for the VM deployment is determined. In embodiments, the predicted resource utilization is associated with a confidence score. The confidence score is associated with a confidence score threshold. The confidence score and the confidence score threshold can be used to determine whether to generate or recommend deploying the predicted rightsized deployment configuration of the VM deployment. Also, a customer may choose to opt-in to having the VM deployment automatically modified to a predicted rightsized deployment configuration, while another tenant may simply have the predicted rightsized deployment configuration communicated to them via an interface, such as a recommended predicted rightsized deployment configuration.

[0027] Embodiments of the present invention can be described by way of an exemplary customer, components, and systems of a cloud computing system. The customer can initially generate a request to deploy a virtual machine (e.g., via a virtual machine subscription) on one of several nodes of the cloud computing system running several virtual machines. The request can be received through a portal or through any other interface with the cloud computing system that provides for requesting virtual machine deployments. The VM deployment can refer to the one or more virtual machines and the corresponding resources of the VM deployment. The VM deployment can also refer to an electronic representation (e.g., a deployment configuration or deployment configuration file or template) of one or more virtual machines and corresponding resources of the VM deployment.

[0028] A virtual machine deployment may be associated with a subscription. A subscription can generally refer to a set of identified resources or selected offerings (e.g., virtual machine deployment instance) from a cloud computing system. A subscription can

be associated with a subscription identifier (ID) that identifies the subscription to be implemented. A virtual machine deployment can include resource types that correspond to the subscription. Resource types (e.g., CPU, GPU, RAM, Total Storage MB, NIC_Ghz, VMs, and HDD IOPS) can be predefined or selectively configured by the customer. For example, a virtual machine deployment having 10 A4s, where an A4 is a type of virtual machine, would have a deployment configuration having the specific resources corresponding to implementing the 10 A4s.

[0029] A cloud computing system provider may support VMs (e.g., a VM deployment) that are deployed on a node or host. For example, a host can support a first VM deployment for five Small VMs or another VM deployment for three Small VMs and two medium VMs, or yet another VM deployment for one GPU Large VM and two Small VMs, etc. each with additional resource types from the exemplary list of resource types identified above.

[0030] Customers may make VM deployment size selections that do not efficiently correspond to the usage. In particular, a customer may select a VM deployment size that under-utilizes the resources associated with the VM deployment. Ideally, a selected VM deployment size should correspond to an expected maximum resource utilization. Cloud computing systems may implement a resource management system (e.g., a resource manager) to rightsize VM deployment configurations.

[0031] Embodiments described herein operate based on a prediction engine, of a resource manager, which can make predictions of resource utilization. The prediction engine can make average and maximum resource utilization predictions for selected resources. For example, a deployment configuration can include a CPU utilization metric and the prediction engine can generate a predicted maximum CPU utilization metric, using techniques described herein. In this regard, the prediction engine can generate a predicted resource utilization for a plurality of resources of a subscription.

[0032] VM rightsizing in the predictive rightsizing system can be based on the capacity for the prediction engine to make predictions based on VM deployments. The resource engine uses a subscription ID and resource types associated with a VM deployment to make predictions for resource utilization. The subscription ID and selected features are used as input into a predictive model. Resource utilization predictions can be based on a predictive model that uses historical operations or past behaviors of previous VM deployments. The resource utilization prediction model can be a machine learning model that is based on features and learning techniques to generate resource utilization

predictions. The prediction engine can use data mining and machine learning techniques to identify correlations that can be leveraged in predictive VM rightsizing.

[0033] The prediction engine can leverage multiple techniques and algorithms to produce a prediction, in particular, a maximum resource utilization prediction. For example, the prediction engine can use regression trees to produce resource utilization predictions, the machine can also use Fast Fourier Transforms on the CPU utilization times series to infer whether a VM role is likely to be user-facing. Other variations and combinations of techniques and algorithms for making resource utilization predictions and other types of predictions are contemplated with embodiments of the present disclosure.

[0034] The prediction engine can receive a large set of customer, tenant, allocation and usage features as input for learning purposes. In particular, previous VM deployments features and past behaviors can be used to predict a maximum utilization of resource for a new VM deployment. Features can include, but not limited to, a lifetime of the VMs (i.e., length of time the VMs are active in the system), service type, operating system, VM stock keeping unit (SKU), number of cores, virtual machine type, last deployed, and VM memory. After an online training period of the predictive model in the predictive engine, a predictive rightsizing controller can query the predictive engine based on a VM deployment having a set of features to receive a predicted resource utilization.

[0035] The predictive rightsizing system can include other aspects (e.g., bucketizing, feature scores and confidence scores) as part of the framework for providing predicted rightsized deployment configurations. With reference to bucketizing, the predictive model can implement bucketizing as part of the model. By way of example, a categorical variable or feature - CPU utilization metric (e.g., 53%) – indicating a VM's past behavior, can be identified. However, the CPU utilization metric at 53% is a narrow indicator of CPU utilization. As such, the CPU utilization metric can be subdivided, or bucketized, into different groups or buckets (e.g., 0-25%, 25-50%, 50%-75%, 75%-100%) collapsing sparsely populated individual CPU utilization metric indicators.

[0036] As shown in FIG. 2B, FIG. 2B plots recall for each bucket independently. The y-axis percentage represents the percentage of VM CPU utilizations in each range - x-axis (i.e., 0-25%, 25-50%, 50%-75%, 75%-100%) – that has been predicted correctly. The VMs have been grouped into 4 VM buckets, each having at or around the indicated percentage of recall for VMs (i.e., VM0 (89%), VM1 (73%), VM2 (71%), and VM3 (86%)) for their corresponding CPU usage buckets (e.g., 0-25%, 25-50%, 50%-75%, 75%-100%). For example, out of all VMs that show CPU utilization between 0-25%, the system

predicted 89% correctly as being in the correct bucket, while 11% were predicted as being in other buckets. As mentioned, for a selected feature (e.g., CPU utilization) it can be difficult to predict an exact percentage of utilization of the resource (e.g., 53%) compared to a predefined range (e.g., 50% - 75%). In this regard, the predictive model only has to
5 make a prediction for a selected resource type within a particular range. As such, resource utilization prediction is based on bucketizing a VM's resource behaviors into selected ranges.

[0037] Turning to feature scores, a feature score is an indicator of the importance of a particular feature in generating a predicted resource utilization. In other words, a feature
10 refers to a processing element determined by the machine learning model to be relevant in making an accurate prediction. For example, as shown in FIG. 2C, in this scenario, the percentage of VMs of the same subscription classified into bucket 1 to date, is the most important feature for making a resource utilization prediction, while VM memory feature is the least important feature in making the resource utilization prediction.

[0038] With reference to the confidence score for each prediction, like with any other prediction system, prediction errors are possible. The confidence score can be generated based on a precision score and a recall score. The precision score is based on the fraction of correct bucket predictions made in the full set of predictions, and the recall is the fraction of correct bucket predictions made in the space of all elements of the correct
15 bucket. With reference to FIG. 2B, the precision score is 0.94 and recall score 0.73, assuming that any predictions with a confidence score lower than 0.6 are discarded. When no predictions are discarded, overall prediction accuracy is 0.83 or 83%.

[0039] As such, a confidence score threshold (e.g., 75%) can be defined such that a predicted rightsized deployment configuration is generated or not generated based on
25 whether the desired confidence is met or not met. By way of example, for low confidence scores (e.g., scores < 75%), the predictive rightsizing controller can be configured not to generate a predictive rightsized deployment configuration and thus not make a recommendation to the customer for their VM deployment. While for a high precision score (e.g., score > 75), the predictive rightsizing controller can be configured to generate
30 a predictive rightsized deployment configuration and then make a recommendation to the customer for their VM deployment. It is contemplated that the confidence score threshold is selectively configurable based on a plurality of factors associated with the cloud computing system. For example, a customer may have a service level agreement where they opt-in to automatic predicted rightsized deployments for confidence scores above a

threshold, or the cloud computing system provider may select or modify the confidence score threshold based on observed outcomes of the prediction engine. Other variations and combinations of confidence scores and threshold and action taken based on them are contemplated with embodiments described herein.

5 **[0040]** Advantageously, the prediction engine uses specifically identified features of VM deployments and corresponding past behaviors for previous VM deployments to make instantaneous rightsized VM recommendations for a new VM deployment without having to evaluate the new VM deployment for a period of time. The rightsized VM deployment configuration recommendation can also be selectively deployed depending on customer
10 elections. The prediction engine also operates to generate a predicted resource utilization for one or a combination of resources, where the predicted resource utilization can be a predicted maximum resource utilization for the VM deployment. Several techniques (e.g., machine learning, bucketing, feature scores, and confidence scores) can be part of the framework for generating and implementing the predicted resource utilization for
15 rightsizing VMs. Moreover, the predicted rightsized deployment for a VM deployment ultimately improves the operation of several hosts or nodes running the VMs, in that, underlying hardware resources are efficiently used and the overall technology field of rightsizing systems is improved as customers, and cloud computing system providers, benefit from rightsized VMs.

20 **[0041]** FIG. 1 shows a resource management system 1 including a predictive rightsizing controller 9. FIG. 2A illustrates the prediction engine 2 of the resource management system 1 with its input features, learning techniques, and output predictions. The prediction engine 2 can use data mining and machine learning techniques to identify correlations that can be leveraged in resource management. For example, the prediction
25 engine 2 can receive a large set of customer, tenant, allocation, and usage features as inputs for learning purposes. Examples of these features include customer name, customer type (e.g., internal, external), deployment type (PaaS vs IaaS), time of deployment, number of VMs, VM roles, VM IDs, VM resource utilizations, and times of shutdown.

[0042] With reference to FIG. 2D, a flow diagram is provided illustrating a method for
30 implementing a predictive rightsizing system. The methods can be performed using the predictive rightsizing system described herein. In embodiments, one or more computer storage media having computer-executable instructions embodied thereon that, when executed, by one or more processors, can cause the one or more processors to perform the methods in the predictive rightsizing system.

[0043] In particular, FIG. 2D is a flow diagram that illustrates a method 200 for implementing a predictive rightsizing system. At 210, the resource manager collects data for VM deployments in a cloud computing system. At 212, the predictive rightsizing controller receives a virtual machine deployment associated with a request to deploy a virtual machine on a node from a plurality of nodes running a plurality of virtual machines in the cloud computing system. At 214, based on collected data, the prediction engine generates a predicted resource utilization for the virtual machine deployment. The predicted resource utilization is generated using the prediction engine that uses past behaviors and a plurality of features associated with previous VM deployments to make resource utilization predictions. Predicted resource utilizations are associated with corresponding confidence scores such that predicted rightsized deployments are generated when a confidence score threshold is met;

[0044] At 216, determine whether a confidence score for the predicted resource utilization meets a confidence score threshold. At 218, end, to not generate or recommend a predicted rightsized deployment configuration when the confidence score does not meet the confidence score threshold. At 220, based on the predicted resource utilization and the confidence score meeting the confidence score threshold, the predictive rightsizing controller generates a predicted rightsized deployment configuration for the VM deployment. At 222, determine whether a customer has opted-in to automatically executing the request to deploy the VM deployment based on the predicted rightsized deployment configuration. At 224, communicate the predicted rightsized deployment configuration to an interface for the customer, when the customer has not opted-in to automatically executing the request to deploy the VM deployment based on the predicted rightsized deployment configuration. At 426, execute the request based on the predicted rightsized deployment configuration when the customer has opted-in to automatically executing the request to deploy the VM deployment based on the predicted rightsized deployment configuration.

[0045] Examples of a cloud computing system, a cluster, healing and growth buffers, and types of VMs are described with reference to FIGS. 3-7.

[0046] FIG. 3 shows a simplistic example of a cloud computing system (CCS) according to the present disclosure. The cloud computing system 10 includes a cloud controller 12 and at least one data center 14. While only one data center 14 is shown for simplicity, the cloud controller 12 can interface with a plurality of data centers. Further, while the data center 14 is shown as being local to the cloud controller 12, one or more

data centers may be geographically remote from the cloud controller 12, may be located in different geographic locations (e.g., in different time zones, different countries or continents, and so on), and may communicate with the cloud controller 12 via various networks.

5 **[0047]** Each data center 14 includes a plurality of fabric controllers 32-1, 32-2, . . . , 32-n (collectively fabric controllers 32) and corresponding clusters 34-1, 34-2, . . . , 34-n (collectively clusters 34). Each fabric controller 32 controls a respective cluster 34. Each cluster 34 includes a plurality of racks (shown in FIGS. 4-5), and each rack includes a plurality of nodes (shown in FIG. 5), which are also called servers, hosts, or machines
10 throughout the present disclosure. Each fabric controller 32 is associated with an allocator 36 that allocates resources within the cluster 34 for instances of customer services hosted on the cluster 34.

[0048] The cloud controller 12 includes a portal 20 and a software development kit (SDK) 22 that the customers can use to select resources and request service deployment.

15 The cloud controller 12 further includes a cloud resource manager 24, a compute resource provider 26, and a front-end 28. The front-end 28 interfaces with the fabric controllers 32. The cloud resource manager 24 receives the customer selections and forwards the customer selections to the compute resource provider 26. The compute resource provider 26 generates a tenant model based on the customer selections. The compute resource
20 provider 26 provisions resources to the customer services according to the tenant model generated based on the customer selections. The compute resource provider 26 provisions storage, networking, and computing resources by interfacing with a cloud storage (Xstore) 30, a network resource provider 31, and the fabric controllers 32. The compute resource provider 26 and the fabric controllers 32 can at least partially implement the resource
25 management system 1 shown in FIGs. 1-2.

[0049] FIG. 4 shows a simplistic example of a data center (e.g., the data center 14 shown in FIG. 3). The data center 14 includes a VM allocator 50 and the clusters 34. The VM allocator 50 can implement the VM allocator 4 of FIG. 1. The VM allocator 50 includes a cluster selector 52 and a plurality of admission controllers 54-1, 54-2, . . . , 54-n
30 (collectively admission controllers 54). Each admission controller 54 is associated with a corresponding cluster 34. Each cluster 34 includes an allocation and healing controller 60 (shown as allocation and healing controllers 60-1, 60-2, . . . , and 60-n; one allocation and healing controller per cluster) and one or more racks 62 of nodes (also called servers, hosts, or machines; and shown as racks 62-1, 62-2, . . . , 62-n). The allocation and healing

controller 60 can implement the VM allocator 36 of FIG. 3.

[0050] Allocating a VM can be a multilevel allocation operation. The VM allocator 50 first selects one of the clusters 34 in which to allocate a VM in association with the corresponding admission controller 54. After the VM allocator 50 selects one of the clusters 34 to allocate the VM, the allocation and healing controller 60 of the selected cluster 34 places the VM on one or more of the nodes in one or more of the racks 62 in the selected cluster 34 depending on the number of update and fault domains and other resources specified by the customer and based on the predictions provided by the prediction engine 2.

[0051] Based on VM activity in the clusters 34, a cluster and service update controller 56 provides updates to the cluster selector 52 and the resource management system 1, which can be partly implemented by the VM allocator 50 and the allocation and healing controllers 60. For example, the VM activity may include activation and termination of one or more VM's in the clusters 34. The cluster and service update controller 56 may also provide updates to the cluster selector 52 and the resource management system 1 regarding utilization of growth buffers due to service scale out and utilization of healing buffers due to node/rack failures, which are explained below with reference to FIG. 6.

[0052] FIG. 5 shows an example of a cluster (e.g., the cluster 34 shown in FIGS. 3-4). Each cluster 34 includes the plurality of racks 62. Each rack 62 includes a rack controller 70 and a plurality of nodes 72. Each rack controller 70 includes a power controller (not shown) that controls power allocation to the plurality of nodes 72 in the rack 62.

[0053] FIG. 6 shows an example of a scale out/growth buffer and a healing buffer in a cluster (e.g., the cluster 34 shown in FIGS. 3-5). Each of the scale out/growth buffer and the healing buffer can comprise a plurality of the nodes 72 within the cluster 34. If a node 72 or a rack 62 fails in the cluster 34, the VM's running on the failed component are healed using the plurality of nodes in the healing buffer. If a service needs to scale out (i.e., expand beyond the allocated resources, which requires additional resources than the allocated resources), extra VM's are placed within the cluster 34. To address one or both of these situations (healing and/or scale-out), sufficient nodes are allocated to the respective buffers.

[0054] When the new deployment threshold is reached in the cluster 34, no new VM deployments can occur into the cluster 34. When the scale out threshold is reached in the cluster 34, the existing tenants cannot grow anymore in the cluster 34. Any attempt to grow results in a scale out failure. If the healing buffer is exhausted in the cluster 34, any

node/rack failures cannot be healed, and a healing failure occurs. Setting the thresholds determines the size of the buffers (e.g., too small, too large, or balanced). The thresholds can be optimized and failures can be minimized based on the predictions from the resource management system 1.

5 **[0055]** FIG. 7 shows examples of different types of VMs depending on their resource requirements. These are only examples, and additional types of VM categories are contemplated. Due to the different resource requirements, the type or category of a VM is an important consideration in the allocation decision. For example, the resource management system 1 can collect and analyze data for many types of VMs. Based on the
10 data, the resource management system 1 can use information regarding the type of a VM that is to be deployed and predict whether the VM will be user-facing or non-user-facing, the lifetime of the VM, and so on.

[0056] Accordingly, the resource management system 1, when implemented by the cloud computing system 10, can provide predicted rightsized deployment configurations
15 using a predictive rightsizing system. In particular, predicted rightsized deployment configurations are generated for virtual machine “VM” subscriptions having deployment configurations that are modified to predicted rightsized deployment configurations based on a prediction engine of a resource manager. The prediction engine may use a machine learning model based on past behaviors and features associated with previous VM
20 deployments to make resource utilization predictions (e.g., maximum resource utilization predictions) for VM deployments.

[0057] Embodiments described in the paragraphs above may be combined with one or more of the specifically described alternatives. In particular, an embodiment that is claimed may contain a reference, in the alternative, to more than one other embodiment.
25 The embodiment that is claimed may specify a further limitation of the subject matter claimed.

[0058] In a first independent embodiment, a predictive rightsizing system for rightsizing virtual machines is provided, the system comprising: a processor and memory; and machine readable instructions stored in the memory, when executed by the processor,
30 configure the processor to: access a virtual machine deployment associated with a request to deploy a virtual machine on a node from a plurality of nodes running a plurality of virtual machines in a cloud computing system; access a predicted resource utilization for the virtual machine deployment, wherein the predicted resource utilization is generated based on a prediction engine that uses past behaviors and a plurality of features associated

with previous virtual machine deployments to make resource utilization predictions; based on the predicted resource utilization, generate a predicted rightsized deployment configuration for the virtual machine deployment; and communicate the predicted rightsized deployment configuration such that the request to deploy the virtual machine is executed based on the predicted rightsized deployment configuration.

[0059] In another embodiment, the predicted resource utilization is a predicted maximum resource utilization for the virtual machine deployment, where the predicted resource utilization is associated with one or more resource types. The virtual machine deployment comprises a deployment configuration comprising the virtual machine and one or more resource types.

[0060] In another embodiment, the plurality of features are associated with previous virtual machine deployments in the cloud computing system, where based on the plurality features, the past behaviors for the previous virtual machine deployments are used generate the predicted resource utilization.

[0061] In another embodiment, generating the predicted rightsized deployment configuration for the virtual machine deployment is based on the predicted resource utilization to circumvent a monitoring period for monitoring virtual machine deployments prior to generating a recommended rightsized deployment.

[0062] In another embodiment, generating the predicted rightsized deployment configuration further comprises: based on the predicted resource utilization, modifying one or more resources in a deployment configuration of the virtual machine deployment.

[0063] In another embodiment, the request is associated with a customer of the cloud computing system, wherein the customer has opted-in to automatically executing the request based on the predicted rightsized deployment configuration.

[0064] In another embodiment, the predicted resource utilizations are associated with corresponding confidence scores such that predicted rightsized deployment configurations are generated when a confidence score threshold is met.

[0065] In a second independent embodiment, one or more computer storage media having computer-executable instructions embodied thereon that, when executed, by one or more processors, cause the one or more processors to perform a method for rightsizing virtual machines, the method comprises accessing a virtual machine deployment associated with a request to deploy a virtual machine on a node from a plurality of nodes running a plurality of virtual machines in a cloud computing system; accessing a predicted resource utilization for the virtual machine deployment, wherein the predicted resource

utilization is generated based on a prediction engine that uses past behaviors and a plurality of features associated with previous virtual machine deployments to make resource utilization predictions, wherein predicted resource utilizations are associated with corresponding confidence scores such that predicted rightsized deployment configurations are generated when a confidence score threshold is met; based on the predicted resource utilization and a confidence score meeting the confidence score threshold, generating a predicted rightsized deployment configuration for the virtual machine deployment; and communicating the predicted rightsized deployment configuration such that the request to deploy the virtual machine is executed based on the predicted rightsized deployment configuration.

[0066] In another embodiment, the predicted resource utilization is a predicted maximum resource utilization for the virtual machine deployment, the predicted resource utilization is associated with one or more resource types, the virtual machine deployment comprises a deployment configuration comprising the virtual machine and one or more resource types.

[0067] In another embodiment, the plurality of features are associated with previous virtual machine deployments in the cloud computing system, where, based on the plurality features, the past behaviors for the previous virtual machine deployments are used generate the predicted resource utilization.

[0068] In another embodiment, generating the predicted rightsized deployment configuration for the virtual machine deployment is based on the predicted resource utilization to circumvent a monitoring period for monitoring virtual machine deployments prior to generating a recommended rightsized deployment.

[0069] In another embodiment, generating the predicted rightsized deployment configuration further comprises: based on the predicted resource utilization modifying one or more resources in a deployment configuration of the virtual machine deployment.

[0070] In another embodiment, the request is associated with a customer of the cloud computing system, where the customer has opted-in to automatically executing the request based on the predicted rightsized deployment configuration.

[0071] In another embodiment, the confidence score threshold is selectively configured based a plurality of factors associated with the cloud computing system.

[0072] In a third independent embodiment, a method for rightsizing virtual machines is provided, the method comprising: accessing a virtual machine deployment associated with a request to deploy a virtual machine on a node; accessing a predicted resource utilization

for the virtual machine deployment, where the predicted resource utilization is generated based on a prediction engine that uses past behaviors and a plurality of features associated with previous virtual machine deployments make resource utilization predictions; based on the predicted resource utilization, generating a predicted rightsized deployment configuration for the virtual machine deployment; and processing the predicted rightsized deployment configuration.

[0073] In another embodiment, the predicted resource utilization is a predicted maximum resource utilization for the virtual machine deployment, wherein the predicted resource utilization is associated with one or more resource types, wherein the virtual machine deployment comprises a deployment configuration comprising the virtual machine and one or more resource types.

[0074] In another embodiment, the plurality of features are associated with previous virtual machine deployments in the cloud computing system; wherein, based on the plurality features, the past behaviors for the previous virtual machine deployments are used generate the predicted resource utilization.

[0075] In another embodiment, generating the predicted rightsized deployment configuration for the virtual machine deployment is based on the predicted resource utilization to circumvent a monitoring period for monitoring virtual machine deployments prior to generating a recommended rightsized deployment.

[0076] In another embodiment, the predicted resource utilizations are associated with corresponding confidence scores such that predicted rightsized deployment configurations are generated when a confidence score threshold is met.

[0077] In another embodiment, processing the predicted rightsized deployment configuration comprises performing one of the following: causing the request to deploy the virtual machine to be executed based on the predicted rightsized deployment configuration; and communicating the predicted rightsized deployment configuration to an interface.

[0078] FIG. 8 shows a simplified example of a distributed network system 100. The distributed network system 300 includes a network 310, one or more client devices 320-1, 320-2, . . . , and 320-N (collectively client devices 320) (where N is an integer greater than or equal to one), and a server 330. The network 310 may include a local area network (LAN), a wide area network (WAN) such as the Internet, or other type of network (collectively shown as the network 310). While only one server is shown, the distributed network system 300 may include multiple servers. The client devices 320 communicate

with the server 330 via the network 310. The client devices 320 and the server 330 may connect to the network 310 using wireless and/or wired connections to the network 310.

[0079] One or more servers 330 and the client devices 320 may implement one or more components of the cloud computing system 10 shown in FIG. 3. For example, one server 330 may implement the cloud controller 12 or the compute resource provider 26 of the cloud controller 12 while one or more client devices 320 may implement the fabric controllers 32. Alternatively, one or more servers 330 may implement one or more components of the cloud controller 12. Many different configurations of implementations are contemplated.

[0080] The server 330 may provide multiple services to the client devices 320. For example, the server 330 may execute a plurality of software applications. The server 330 may host multiple databases that are utilized by the plurality of software applications and that are used by the client devices 320. In addition, the server 330 and the client devices 320 may execute applications that implement one or more components of the resource management system 1 and one or more methods for the use cases described above.

[0081] FIG. 9 shows a simplified example of the client device 320. The client device 320 may typically include a central processing unit (CPU) or processor 350, one or more input devices 352 (e.g., a keypad, touchpad, mouse, and so on), a display subsystem 354 including a display 356, a network interface 358, a memory 360, and a bulk storage 362.

[0082] The network interface 358 connects the client device 320 to the distributed network system 300 via the network 310. For example, the network interface 358 may include a wired interface (e.g., an Ethernet interface) and/or a wireless interface (e.g., a Wi-Fi, Bluetooth, near field communication (NFC), or other wireless interface). The memory 360 may include volatile or nonvolatile memory, cache, or other type of memory. The bulk storage 362 may include flash memory, a hard disk drive (HDD), or other bulk storage device.

[0083] The processor 350 of the client device 320 executes an operating system (OS) 364 and one or more client applications 366. The client applications 366 include an application to connect the client device 320 to the server 330 via the network 310. The client device 320 accesses one or more applications executed by the server 330 via the network 310. The client applications 366 may also include an application that implements all or some aspects of one or more components of the resource management system 1 and one or more methods for the use cases described above.

[0084] FIG. 10 shows a simplified example of the server 330. The server 330 typically

includes one or more CPUs or processors 370, one or more input devices 372 (e.g., a keypad, touchpad, mouse, and so on), a display subsystem 374 including a display 376, a network interface 378, a memory 380, and a bulk storage 382.

[0085] The network interface 378 connects the server 330 to the distributed network system 300 via the network 310. For example, the network interface 378 may include a wired interface (e.g., an Ethernet interface) and/or a wireless interface (e.g., a Wi-Fi, Bluetooth, near field communication (NFC), or other wireless interface). The memory 380 may include volatile or nonvolatile memory, cache, or other type of memory. The bulk storage 382 may include flash memory, one or more hard disk drives (HDDs), or other bulk storage device.

[0086] The processor 370 of the server 130 executes an operating system (OS) 384 and one or more server applications 386. The server applications 386 may include an application that implements all or some aspects of one or more components of the resource management system 1 and one or more methods for the use cases described above. The bulk storage 382 may store one or more databases 388 that store data structures used by the server applications 386 to perform respective functions.

[0087] The foregoing description is merely illustrative in nature and is in no way intended to limit the disclosure, its application, or uses. The broad teachings of the disclosure can be implemented in a variety of forms. Therefore, while this disclosure includes particular examples, the true scope of the disclosure should not be so limited since other modifications will become apparent upon a study of the drawings, the specification, and the following claims. It should be understood that one or more steps within a method may be executed in different order (or concurrently) without altering the principles of the present disclosure. Further, although each of the embodiments is described above as having certain features, any one or more of those features described with respect to any embodiment of the disclosure can be implemented in and/or combined with features of any of the other embodiments, even if that combination is not explicitly described. In other words, the described embodiments are not mutually exclusive, and permutations of one or more embodiments with one another remain within the scope of this disclosure.

[0088] Spatial and functional relationships between elements (for example, between modules, circuit elements, semiconductor layers, etc.) are described using various terms, including “connected,” “engaged,” “coupled,” “adjacent,” “next to,” “on top of,” “above,” “below,” and “disposed.” Unless explicitly described as being “direct,” when a

relationship between first and second elements is described in the above disclosure, that relationship can be a direct relationship where no other intervening elements are present between the first and second elements, but can also be an indirect relationship where one or more intervening elements are present (either spatially or functionally) between the first and second elements. As used herein, the phrase at least one of A, B, and C should be construed to mean a logical (A OR B OR C), using a non-exclusive logical OR, and should not be construed to mean “at least one of A, at least one of B, and at least one of C.”

[0089] In the figures, the direction of an arrow, as indicated by the arrowhead, generally demonstrates the flow of information (such as data or instructions) that is of interest to the illustration. For example, when element A and element B exchange a variety of information but information transmitted from element A to element B is relevant to the illustration, the arrow may point from element A to element B. This unidirectional arrow does not imply that no other information is transmitted from element B to element A. Further, for information sent from element A to element B, element B may send requests for, or receipt acknowledgements of, the information to element A.

[0090] The term memory is a subset of the term computer-readable medium or machine-readable medium. The term computer-readable medium or machine-readable medium, as used herein, does not encompass transitory electrical or electromagnetic signals propagating through a medium (such as on a carrier wave); the term computer-readable medium or machine-readable medium may therefore be considered tangible and non-transitory. Non-limiting examples of a non-transitory, tangible computer-readable medium or machine-readable medium are nonvolatile memory circuits (such as a flash memory circuit, an erasable programmable read-only memory circuit, or a mask read-only memory circuit), volatile memory circuits (such as a static random access memory circuit or a dynamic random access memory circuit), magnetic storage media (such as an analog or digital magnetic tape or a hard disk drive), and optical storage media (such as a CD, a DVD, or a Blu-ray Disc).

[0091] In this application, apparatus elements described as having particular attributes or performing particular operations are specifically configured to have those particular attributes and perform those particular operations. Specifically, a description of an element to perform an action means that the element is configured to perform the action. The configuration of an element may include programming of the element, such as by encoding instructions on a non-transitory, tangible computer-readable medium associated

with the element.

[0092] The apparatuses and methods described in this application may be partially or fully implemented by a special purpose computer created by configuring a general purpose computer to execute one or more particular functions embodied in computer programs.

5 The functional blocks, flowchart components, and other elements described above serve as software specifications, which can be translated into the computer programs by the routine work of a skilled technician or programmer.

[0093] The computer programs include processor-executable instructions that are stored on at least one non-transitory, tangible computer-readable medium. The computer
10 programs may also include or rely on stored data. The computer programs may encompass a basic input/output system (BIOS) that interacts with hardware of the special purpose computer, device drivers that interact with particular devices of the special purpose computer, one or more operating systems, user applications, background services, background applications, etc.

15 **[0094]** The computer programs may include: (i) descriptive text to be parsed, such as HTML (hypertext markup language), XML (extensible markup language), or JSON (JavaScript Object Notation) (ii) assembly code, (iii) object code generated from source code by a compiler, (iv) source code for execution by an interpreter, (v) source code for compilation and execution by a just-in-time compiler, etc. As examples only, source code
20 may be written using syntax from languages including C, C++, C#, Objective-C, Swift, Haskell, Go, SQL, R, Lisp, Java®, Fortran, Perl, Pascal, Curl, OCaml, Javascript®, HTML5 (Hypertext Markup Language 5th revision), Ada, ASP (Active Server Pages), PHP (PHP: Hypertext Preprocessor), Scala, Eiffel, Smalltalk, Erlang, Ruby, Flash®, Visual Basic®, Lua, MATLAB, SIMULINK, and Python®.

25 **[0095]** None of the elements recited in the claims are intended to be a means-plus-function element within the meaning of 35 U.S.C. §112(f) unless an element is expressly recited using the phrase “means for,” or in the case of a method claim using the phrases “operation for” or “step for.”

CLAIMS

1. A predictive rightsizing system for rightsizing virtual machines, the system comprising:
a processor and memory; and
machine readable instructions stored in the memory, when executed by the processor, configure the processor to:
 - access a virtual machine deployment associated with a request to deploy a virtual machine on a node from a plurality of nodes running a plurality of virtual machines in a cloud computing system;
 - access a predicted resource utilization for the virtual machine deployment, wherein the predicted resource utilization is generated based on a prediction engine that uses past behaviors and a plurality of features associated with previous virtual machine deployments to make resource utilization predictions;
 - based on the predicted resource utilization, generate a predicted rightsized deployment configuration for the virtual machine deployment; and
 - communicate the predicted rightsized deployment configuration such that the request to deploy the virtual machine is executed based on the predicted rightsized deployment configuration.
2. The system of claim 1, wherein the predicted resource utilization is a predicted maximum resource utilization for the virtual machine deployment, wherein the predicted resource utilization is associated with one or more resource types, wherein the virtual machine deployment comprises a deployment configuration comprising the virtual machine and one or more resource types.
3. The system of claim 1, wherein the plurality of features are associated with previous virtual machine deployments in the cloud computing system; wherein, based on the plurality features, the past behaviors for the previous virtual machine deployments are used generate the predicted resource utilization.
4. The system of claim 1, wherein generating the predicted rightsized deployment configuration for the virtual machine deployment is based on the predicted resource utilization to circumvent a monitoring period for monitoring virtual machine deployments prior to generating a recommended rightsized deployment.
5. The system of claim 1, wherein generating the predicted rightsized deployment configuration further comprises: based on the predicted resource utilization, modifying one or more resources in a deployment configuration of the virtual machine deployment.

6. The system of claim 1, wherein predicted resource utilizations are associated with corresponding confidence scores such that predicted rightsized deployment configurations are generated when a confidence score threshold is met.

7. One or more computer storage media having computer-executable instructions embodied thereon that, when executed, by one or more processors, cause the one or more processors to perform a method for rightsizing virtual machines, the method comprising:

accessing a virtual machine deployment associated with a request to deploy a virtual machine on a node from a plurality of nodes running a plurality of virtual machines in a cloud computing system;

accessing a predicted resource utilization for the virtual machine deployment, wherein the predicted resource utilization is generated based on a prediction engine that uses past behaviors and a plurality of features associated with previous virtual machine deployments to make resource utilization predictions, wherein predicted resource utilizations are associated with corresponding confidence scores such that predicted rightsized deployment configurations are generated when a confidence score threshold is met;

based on the predicted resource utilization and a confidence score meeting the confidence score threshold, generating a predicted rightsized deployment configuration for the virtual machine deployment; and

communicating the predicted rightsized deployment configuration such that the request to deploy the virtual machine is executed based on the predicted rightsized deployment configuration.

8. The media of claim 7, wherein the confidence score threshold is selectively configured based a plurality of factors associated with the cloud computing system.

9. A method for rightsizing virtual machines, the method comprising:

accessing a virtual machine deployment associated with a request to deploy a virtual machine on a node;

accessing a predicted resource utilization for the virtual machine deployment, wherein the predicted resource utilization is generated based on a prediction engine that uses past behaviors and a plurality of features associated with previous virtual machine deployments make resource utilization predictions;

based on the predicted resource utilization, generating a predicted rightsized deployment configuration for the virtual machine deployment; and

processing the predicted rightsized deployment configuration.

10 The method of claim 9, wherein processing the predicted rightsized deployment configuration comprises performing one of the following:

causing the request to deploy the virtual machine to be executed based on the predicted rightsized deployment configuration; and

communicating the predicted rightsized deployment configuration to an interface.




FIG. 1




FIG. 2A




FIG. 2B





FIG. 2C

4/9






FIG. 3




FIG. 5




FIG. 6




FIG. 4

Type A: High-value			
Type	Cores	Ram	
A0	1	.768	
A1	1	1.75	
A2	2	3.5	
A3	4	7	
A4	8	14	
A5	2	14	
A6	4	28	
A7	8	56	
A8	8	56	
A9	16	112	
A10	8	56	
A11	16	112	

Type D: Low-latency, SSD			
Type	Cores	Ram	
D1	1	3.5	
D2	4	7	
D3	4	14	
D4	8	28	
D11	2	14	
D12	4	28	
D13	8	56	
D14	16	112	

Type G: Extreme Perf., SSD			
Type	Cores	Ram	
G1	2	28	
G2	4	56	
G3	8	112	
G4	16	224	
G5	32	448	

FIG. 7




FIG. 9




FIG. 8




FIG. 10

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2018/059913

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	US 2015/317081 A1 (SINGH ASHISH KUMAR [IN] ET AL) 5 November 2015 (2015-11-05) paragraphs [0002], [0021] - [0036], [0043], [0044], [0049], [0054], [0055] -----	1-10
A	US 2010/058342 A1 (MACHIDA FUMIO [JP]) 4 March 2010 (2010-03-04) paragraphs [0019], [0023], [0025], [0037] -----	1-10
A	WO 2015/027865 A1 (IBM [US]; IBM CHINA CO LTD [CN]) 5 March 2015 (2015-03-05) paragraphs [0004] - [0007], [0049] - [0051] -----	1-10



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

18 February 2019

Date of mailing of the international search report

25/02/2019

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

Buzgan, C

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2018/059913

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 2015317081	A1	05-11-2015	NONE	

US 2010058342	A1	04-03-2010	JP 5262724 B2	14-08-2013
			JP WO2008084826 A1	06-05-2010
			US 2010058342 A1	04-03-2010
			WO 2008084826 A1	17-07-2008

WO 2015027865	A1	05-03-2015	CN 104424013 A	18-03-2015
			US 2016162308 A1	09-06-2016
			US 2017308393 A1	26-10-2017
			WO 2015027865 A1	05-03-2015
