



(19) **United States**

(12) **Patent Application Publication** (10) **Pub. No.: US 2018/0331927 A1**

Eberlein et al.

(43) **Pub. Date: Nov. 15, 2018**

(54) **RESOURCE COORDINATE SYSTEM FOR DATA CENTERS**

(71) Applicant: **SAP SE**, Walldorf (DE)

(72) Inventors: **Peter Eberlein**, Malsch (DE); **Volker Driesen**, Heidelberg (DE)

(73) Assignee: **SAP SE**, Walldorf (DE)

(21) Appl. No.: **15/591,683**

(22) Filed: **May 10, 2017**

Publication Classification

(51) **Int. Cl.**
H04L 12/26 (2006.01)

(52) **U.S. Cl.**
CPC **H04L 43/0876** (2013.01); **H04L 43/06** (2013.01)

(57) **ABSTRACT**

Embodiments allow identifying unique locations, proximity, and constraints of various individualization resource units in a data center. A plurality of processes may compete for available resources (e.g., CPU, network bandwidth, memory) in a shared environment, which may be virtualized (e.g., comprising host, VM, container, application server instance resource units). A scheduling infrastructure determines individualization unit locations according to an identifier reflecting resource hierarchies. For each process, the data sending infrastructure sends location data as an ID stack to a central monitoring instance for analysis. That central monitoring instance can reference the resource location information to identify process proximity and competition for resources constraining process performance. Resource location information in the form of coordinates provided by the ID stack, offer valuable information regarding resource consumption and other information relevant for process monitoring. Embodiments may be particularly useful in diagnosing bottlenecks arising from processes sharing resource units in virtualized environments.





FIG. 1



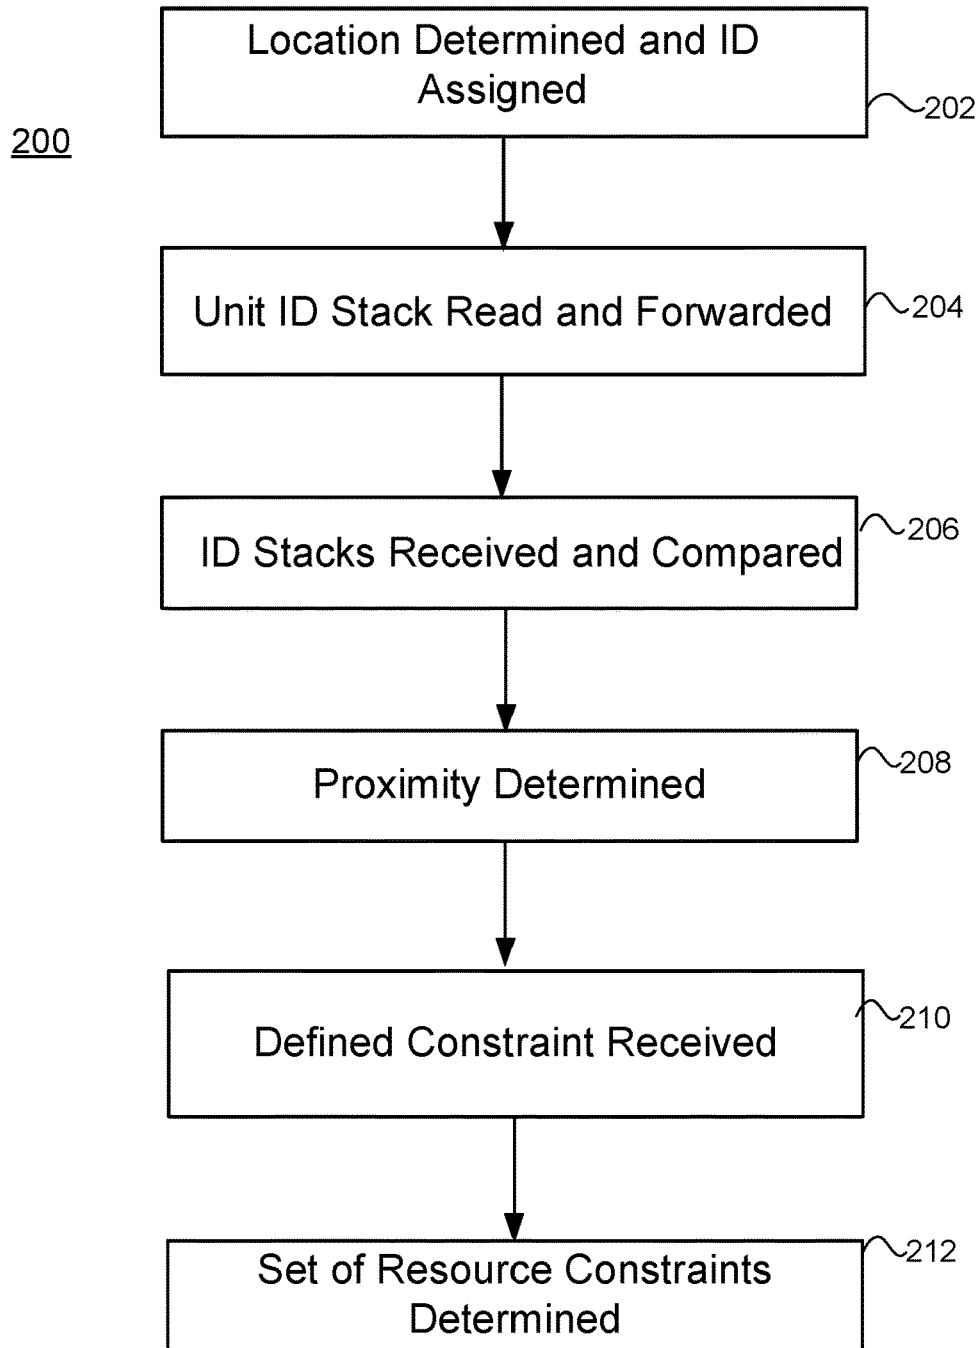
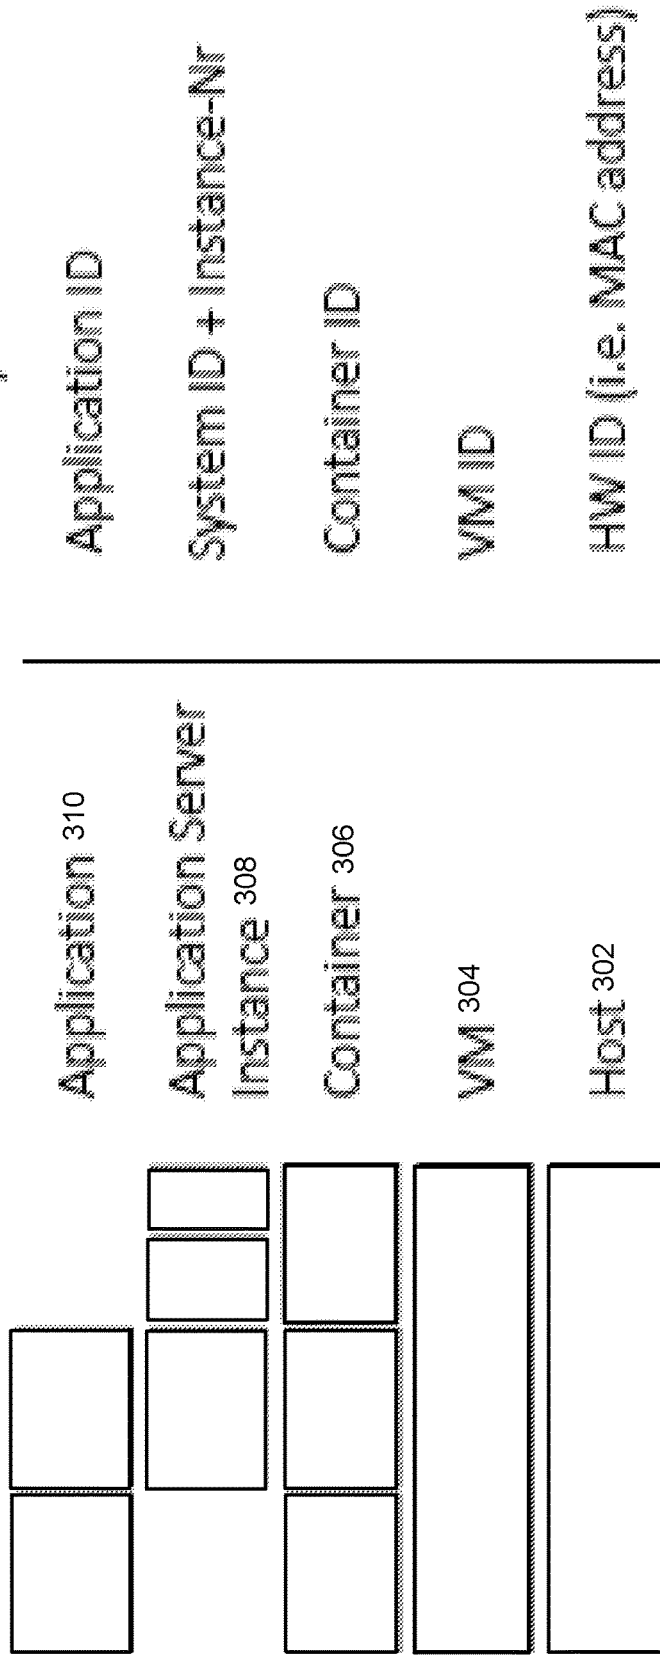


FIG. 2

FIG. 3

Define unique identifiers
for all layers:






FIG. 4

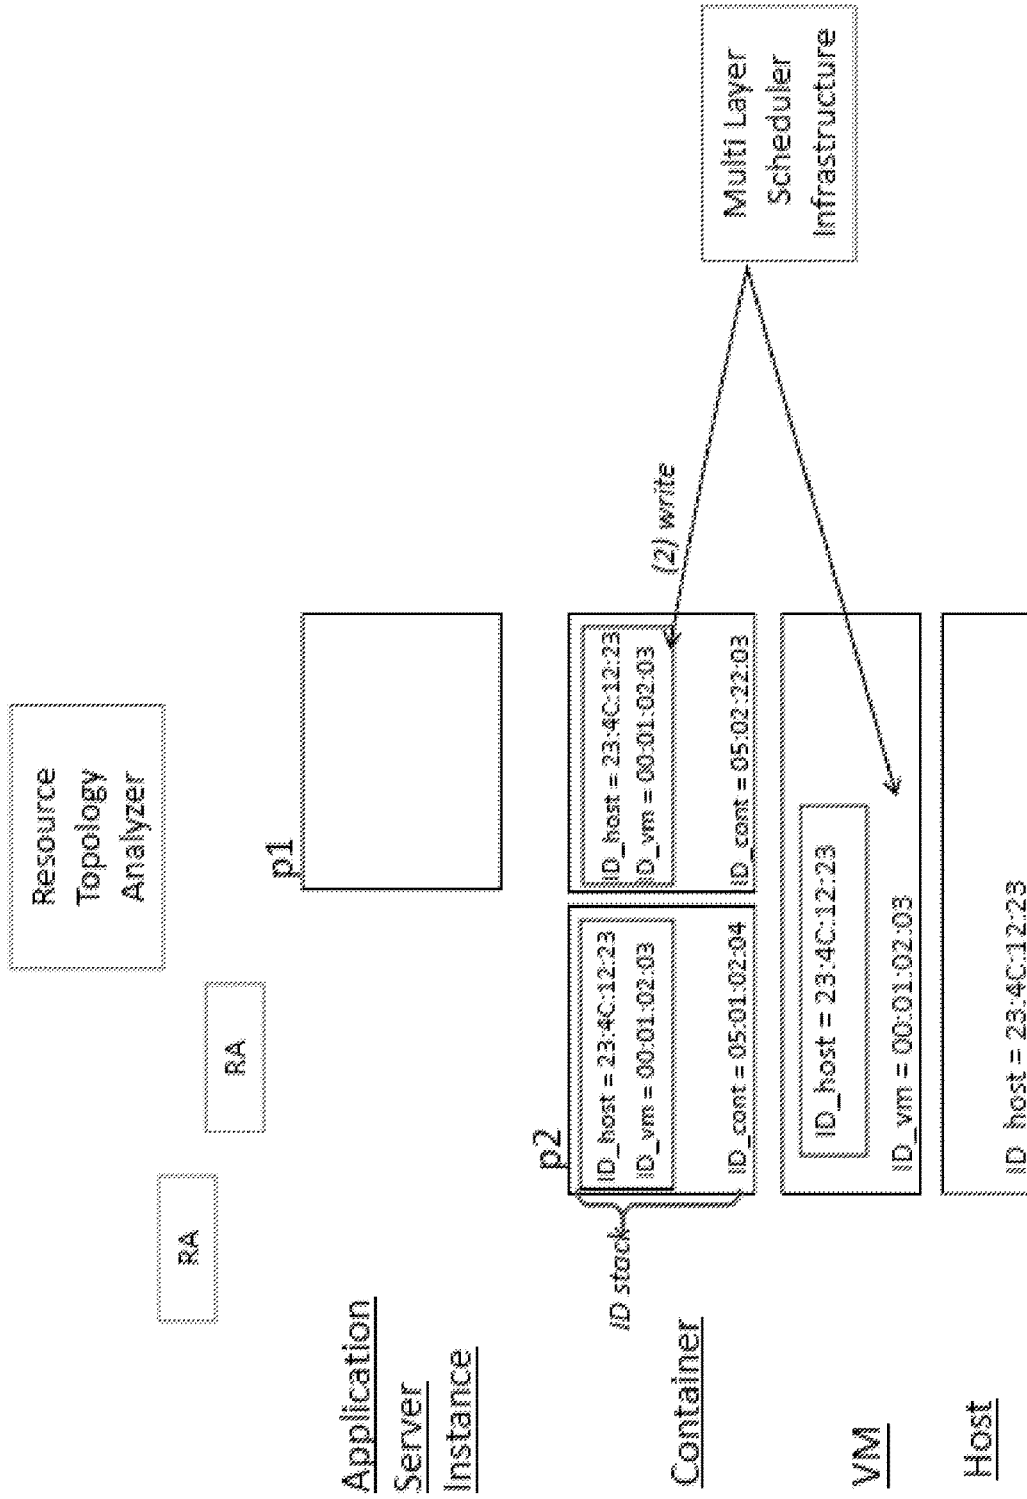


FIG. 5

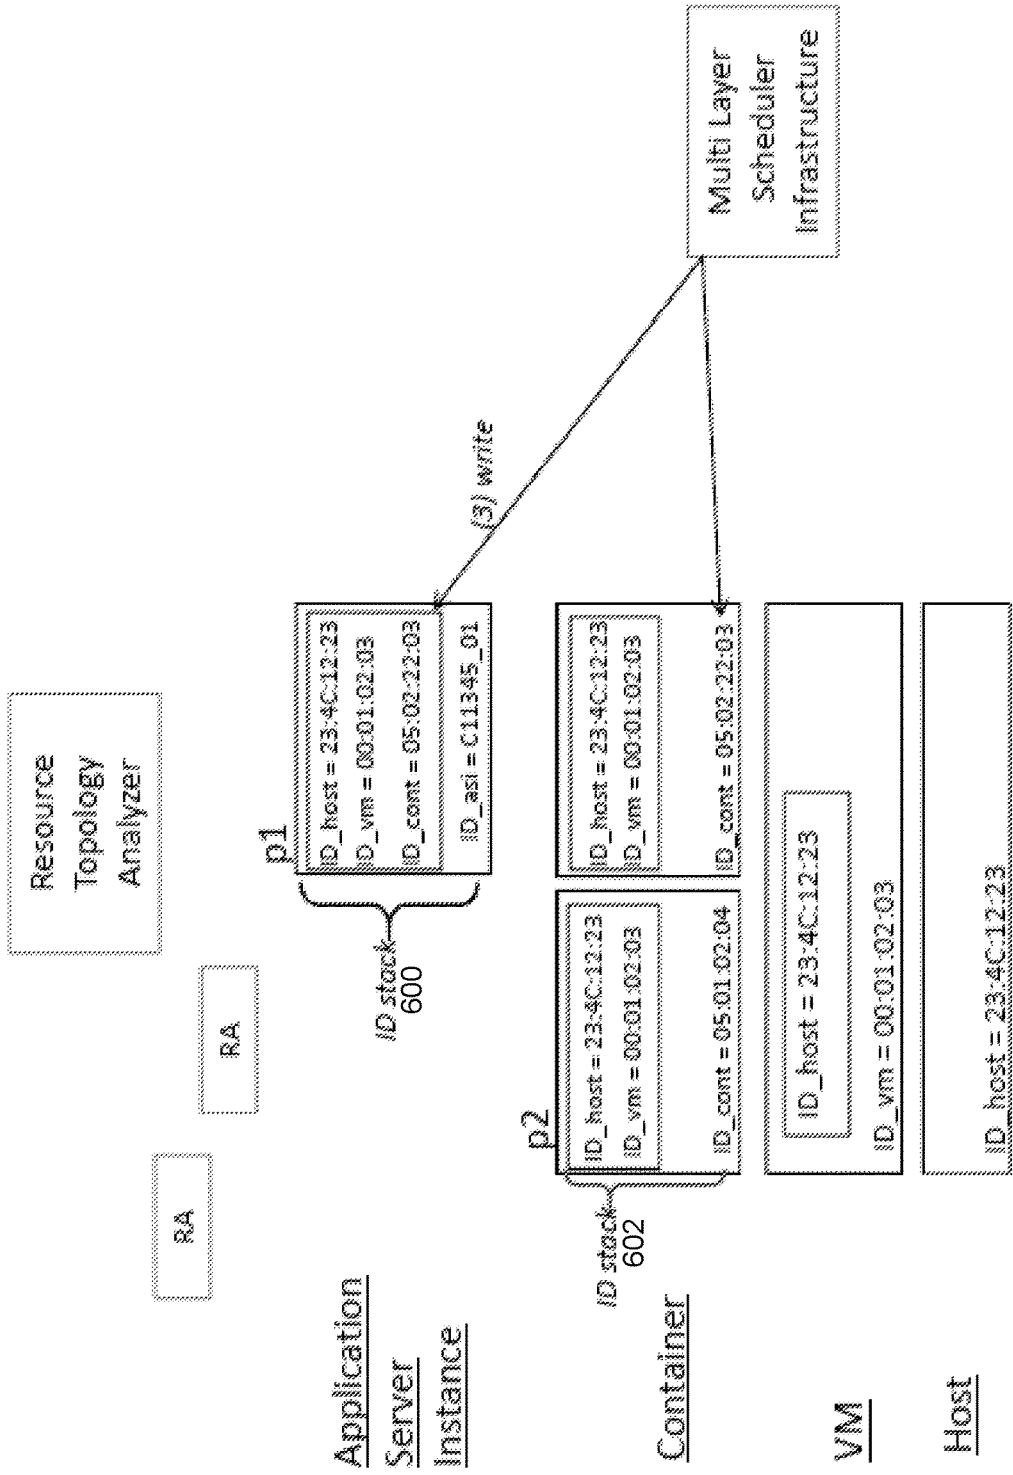


FIG. 6




FIG. 7

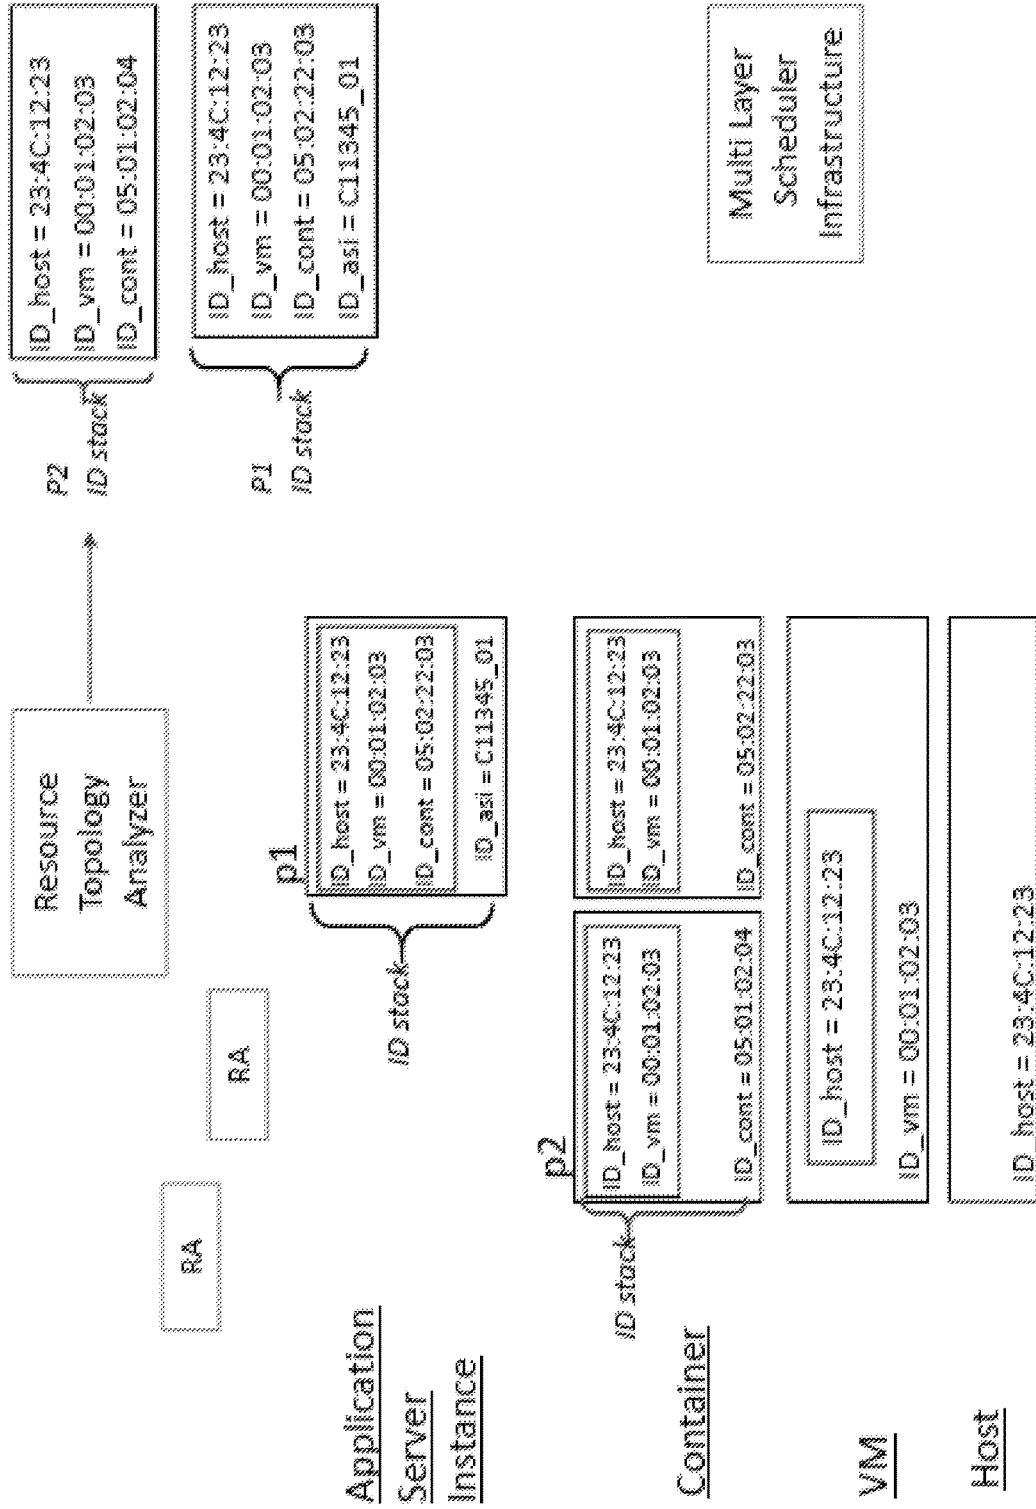


FIG. 8




FIG. 9




FIG. 10




FIG. 11

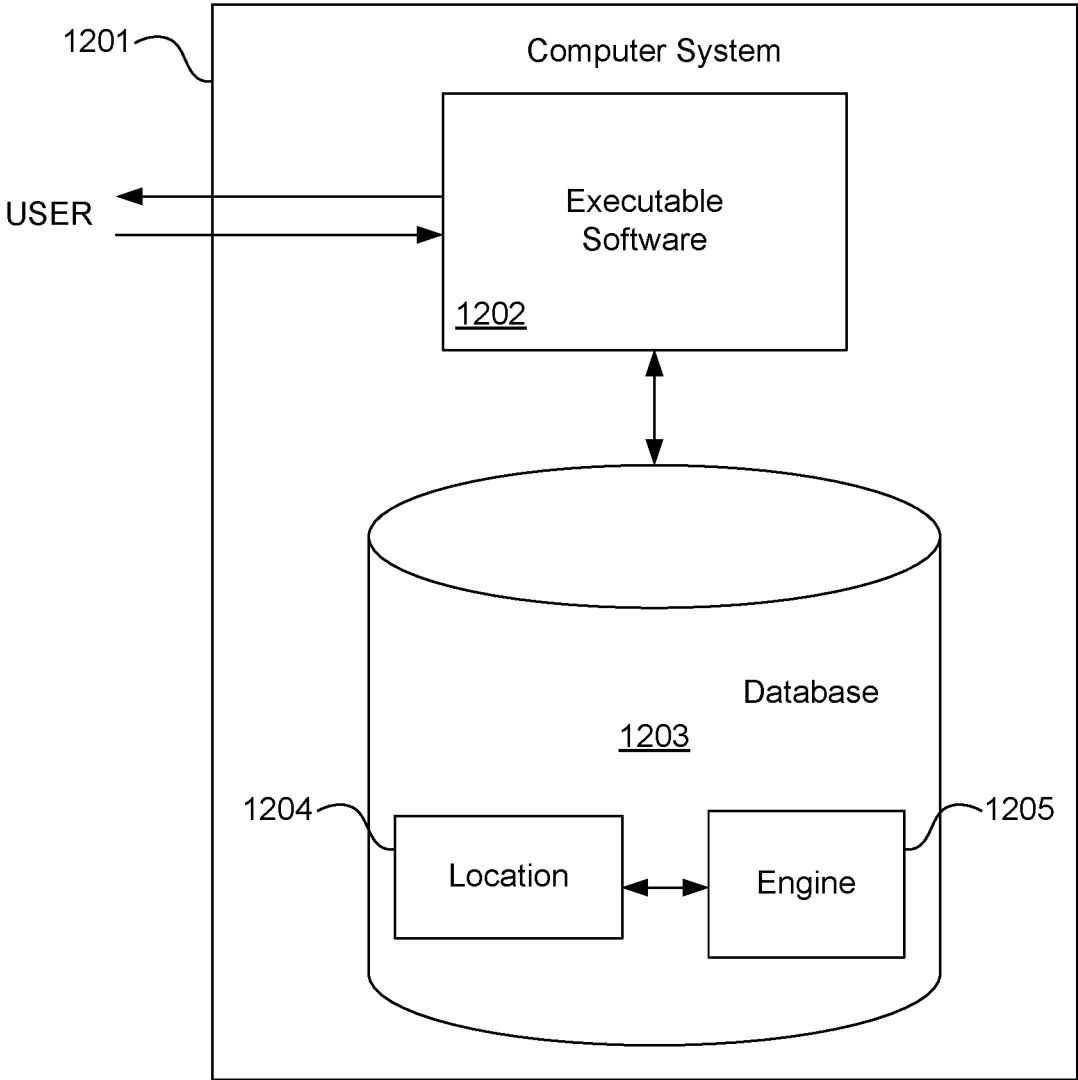


FIG. 12




FIG. 13

RESOURCE COORDINATE SYSTEM FOR DATA CENTERS

BACKGROUND

[0001] Unless otherwise indicated herein, the approaches described in this section are not prior art to the claims in this application and are not admitted to be prior art by inclusion in this section.

[0002] The push for cost reduction in modern data centers leads to increasing hardware utilization. Idle hardware is considered a cost problem, and mechanisms are created to minimize idle times.

[0003] However, running more processes to minimize idle hardware also implies that these processes compete for resources. This can lead to slower execution, longer run-times, and even the abortion of processes (e.g., if no more memory can be allocated).

[0004] Virtualization layers, virtual machines (VMs), containers, and application servers are various mechanisms that are used to optimize hardware utilization. These allow “overprovisioning”: VMs can be defined for certain hardware, which in sum offer more resources than the hardware that is available.

[0005] Overprovisioning succeeds most of the time, where all VMs are not encountering peak demand conditions. However where hardware limits are reached and processes are slowed down or aborted, it becomes increasingly complicated to identify the root cause of the delays. This task is made more difficult by the presence of abstraction layers of VMs, containers, and application server instances (ASIs).

SUMMARY

[0006] Embodiments offer a coordinate system allowing the identification of resources in a data center. Particular embodiments permit the determination of unique locations, proximity, and constraints of various individualization resource units, including virtualized resources. Embodiments may thus be particularly useful in diagnosing resource bottlenecks arising from processes competing for resource units in virtualized environments.

[0007] Specifically, plurality of processes may compete for available resources (e.g., CPU, network bandwidth, memory) in a shared environment, which may be virtualized (e.g., comprising host, VM, container, application server instance resource units). A scheduling infrastructure initially defines individualization resource unit locations as coordinates in the form of an identifier (which may reflect resource hierarchies). For each process, a data sending infrastructure sends location data in the form of an ID stack to a central monitoring instance for analysis of the resource situation. That central monitoring instance can reference the resource location information to identify process proximity and competition for resources constraining process performance. Resource location information in the form of coordinates provided by the ID stack, offer valuable information regarding resource consumption and other information relevant for process monitoring.

[0008] An embodiment of a computer-implemented method comprises assigning unique identifiers to each of a plurality of resources within a hierarchy of a computing environment to define location information, storing the location information, and reporting the location information to a central monitoring instance. The central monitoring

instance determines from the location information, a first set of resources consumed by a first process based upon a first identifier stack comprising one or more of the unique identifiers. The central monitoring instance displays the first set of resources.

[0009] A non-transitory computer readable storage medium embodies a computer program for performing a method comprising assigning unique identifiers to each of a plurality of resources within a hierarchy of a computing environment to define location information. The method further comprises storing the location information, and reporting the location information to a central monitoring instance. The central monitoring instance determines from the location information, a first set of resources consumed by a first process based upon a first identifier stack comprising one or more of the unique identifiers. The central monitoring instance receives from a modeling infrastructure, a constraint of one resource of the first set of resources, the central monitoring instance displays the first set of resources and the constraint.

[0010] An embodiment of a computer system comprises one or more processors and a software program, executable on said computer system. The software program is configured to cause an in-memory database engine to assign unique identifiers to each of a plurality of resources within a hierarchy of a computing environment to define location information, to store the location information in an in-memory database, and to determine from the location information, a first set of resources consumed by a first process based upon a first identifier stack comprising one or more of the unique identifiers. The software program is further configured to cause an in-memory database engine to determine from the location information, a second set of resources consumed by a second process based upon a second identifier stack comprising one or more of the unique identifiers, to compare the first identifier stack with the second identifier stack to determine a proximity between the first process and the second process, and to communicate the proximity for display.

[0011] Certain embodiments further comprise the central monitoring instance receiving from a modeling infrastructure a constraint of one resource of the first set of resources, and the central monitoring instance revealing the one resource as a performance bottleneck of the first process based upon the constraint.

[0012] Some embodiments further comprise the central monitoring instance determining from the location information, a second set of resources consumed by a second process based upon a second identifier stack comprising one or more of the unique identifiers, and the central monitoring instance comparing the first identifier stack with the second identifier stack to determine a proximity between the first process and the second process.

[0013] Particular embodiments may further comprise the central monitoring instance receiving from a modeling infrastructure, a constraint of one resource identified as proximate between the first process and the second process.

[0014] Various embodiments may further comprise the central monitoring instance revealing the one resource as a performance bottleneck of the first process based upon the constraint.

[0015] According to some embodiments the computing environment comprises a virtualization computing environ-

ment, and the constraint comprises a number of virtual machines supported by a hardware central processing unit.

[0016] In particular embodiments the computing environment comprises a network topology including a switch, and the constraint comprises a bandwidth of the switch.

[0017] In various embodiments the computing environment comprises a common resource shared between the first process and the second process, and the constraint comprises a number of operations per second of the common resource.

[0018] According to certain embodiments the common resource comprises an input-output device.

[0019] In particular embodiments the common resource comprises a memory.

[0020] In some embodiments the computing environment comprises a virtualization computing environment, and the hierarchy comprises a virtual machine.

[0021] The following detailed description and accompanying drawings provide a better understanding of the nature and advantages of embodiments.

BRIEF DESCRIPTION OF THE DRAWINGS

[0022] FIG. 1 shows a simplified diagram of a system according to an embodiment.

[0023] FIG. 2 shows a simplified flow diagram of a method according to an embodiment.

[0024] FIG. 3 shows a resource hierarchy in a virtualization environment.

[0025] FIGS. 4-9 shows details of a process flow for resource analysis according to a specific example involving a virtualization environment.

[0026] FIGS. 10-11 show details of a process flow for resource analysis according to a specific example involving a network topology environment.

[0027] FIG. 12 illustrates hardware of a special purpose computing machine according to an embodiment that is configured to perform resource analysis.

[0028] FIG. 13 illustrates an example computer system.

DETAILED DESCRIPTION

[0029] Described herein are methods and apparatuses implementing resource coordinate systems for data centers. In the following description, for purposes of explanation, numerous examples and specific details are set forth in order to provide a thorough understanding of embodiments according to the present invention. It will be evident, however, to one skilled in the art that embodiments as defined by the claims may include some or all of the features in these examples alone or in combination with other features described below, and may further include modifications and equivalents of the features and concepts described herein.

[0030] FIG. 1 shows a simplified view of a system configured to implement resource analysis according to an embodiment. Specifically, system 100 comprises engine 102 comprising a scheduler 104, reporting agent(s) 106, and central monitoring instance 108.

[0031] Engine 102 is configured to communicate with a data center 110 comprising a plurality of individualization resource units 112 that are arranged according to a resource hierarchy 114. One example of such a resource hierarchy is a virtualization environment comprising host, VMs, containers, and application server instances. Another possible

example of such a resource hierarchy is a network topology comprising cards connected by switches having different bandwidths.

[0032] A first process (p1) 120 operates within the data center, consuming a set of various individualization resource units thereof, from parent (lower) to child (middle) and (grandchild) upper. A second process (p2) 122 also operates within the same data center to consume various (parent/child) individualization resource units, some of which may be shared with the first process.

[0033] In order to provide an insightful analysis of resource consumption by data center processes, FIG. 2 shows a process flow of various actions that are taken by the engine. Specifically, at 202 the scheduler determines the location of each individual resource unit within the data center, assigning each unit an identifier. In this manner a resource coordinate system is created, which can be used to map processes to individualization resource units.

[0034] At 204, the reporting agent reads a respective stack of unit IDs 130, 132 from each resource unit, the stack corresponding to each process running in the data center. These unit ID stacks are then forwarded by the reporting agent to the central monitoring instance for storage in non-transitory computer-readable storage medium 134.

[0035] At 206, the central monitoring instance receives the data sets from the reporting agent(s). The monitoring instance extracts the IDs and compares them for different data sets of different processes.

[0036] At 208, where comparison of the reported stacks reveals the presence of the same identifiers, proximity 136 between processes is indicated. For example, in the simplified embodiment of FIG. 1, the ID stacks of the processes p1 and p2 reveal that they both include resource individualization units •I and •A (e.g., which in a virtualization environment may correspond to the same VM running on the same host).

[0037] At 210 a resource modeling infrastructure (e.g., 150 of FIG. 1) may define constraints 152, 154 of individual resource units. For example, the resource modeling infrastructure may define a constraint in the form of a maximum number of VM units that are able to be supported by a specific single physical CPU of the hardware. In another example, the resource modeling infrastructure may define a maximum bandwidth available to a particular switch in a network topology. This constraint information is then sent to the central monitoring instance.

[0038] At 212, the central monitoring instance receives the individual unit constraint information from the modeling infrastructure. The monitoring instance finds matching IDs of units of different processes, and reads the constraints (e.g., from child-to-parent), thus providing to a user 160 a set of constraint resources 156 for the competing processes.

[0039] Various details of implementing a coordinate system for data center resource analysis according to particular embodiments, are now discussed in connection with the following example of FIGS. 3-10.

Example

[0040] FIG. 3 shows an example of a hardware environment 300 comprising a plurality of CPU resource units configured in a hierarchical manner, e.g.:

[0041] host 302;

[0042] virtual machine (VM) 304;

[0043] containers 306;

[0044] application server instances 308; and

[0045] application 310.

[0046] Monitoring data and runtime evaluation of processes deployed in that hardware environment may reveal that one or more of these processes is running slower than usual or expected. Then, a root-cause-analysis process commences in order to identify the origin of the problem.

[0047] The next level of analysis can indicate that one of the resources was constrained, and that the slowed process did not receive the expected set of resources. Application of a resource coordinate system according to embodiments may aid in identifying such a constrained resource.

[0048] As discussed previously, initially a location of each of the individualization resource units of the data center environment is determined. Here, FIG. 3 generally shows the assignment of unique IDs to the individualization units in each layer of the resource hierarchy of the virtualization environment. That unique ID provides the coordinate system serving to provide a location of the individualization resource unit within the data center environment.

[0049] Specifically, for each individualization unit, a central scheduling infrastructure (e.g. the one creating, starting and moving the unit) injects the unique ID of the underlying individualization unit. This ID is updated by the scheduling infrastructure, if the underlying individualization unit changes (e.g. a move to another unit).

[0050] Accordingly, FIG. 4 shows a multi-layer scheduling infrastructure (MLSI) 400 (which can be a group of scheduling infrastructures). The MLSI first assigns to each parent individualization unit (e.g., host), a unique ID 402. Here, the infrastructure comprises only a single parent (host), which is assigned the media access control (MAC) address as the unique ID.

[0051] As shown in FIG. 5, the MLSI then passes the unique ID of the parent unit to the child unit contained within the parent unit, with each parent unit potentially having several child units, and each child potentially having several parent units. If the parent changes during process runtime, the MLSI passes the new parent unit ID to the child, thereby allowing the child unit to know its parent unit.

[0052] As shown in FIG. 6, the MLSI also manages parents of parents (i.e. grandparents) and children of children (i.e. grand children). The MLSI thus passes the grand parent unit IDs to a child unit, reading them from the parent unit ID store (and grand-grand parent unit IDs and so forth).

[0053] At the conclusion of the location procedure shown in FIG. 6, the specific set of resources of particular processes p1 and p2 can thus be represented by an ID stack (e.g., 600, 602). In particular, the unique ID of the host is injected into the VM, the unique ID of the VM into the container, the unique ID of the container into the app server instance. The location of the resource is then a set of IDs (ID stack) for the direct environment, the underlying environment and so on.

[0054] Next, FIG. 7 shows coordinate and resource consumption reporting agents (RA) 700 and 702 for the respective processes p1 and p2, reading the stack of unit IDs from the unit ID store of the current environment. The RA sends the stack of unit IDs to the resource topology analyzer RTA 704 for further processing. For each send event, the RA reads the unit IDs again in order to capture updated IDs (due, e.g., to a shifting of resources allocated to the process).

[0055] FIG. 8 shows the activity of the RTA upon receiving data sets from all connected RAs, to derive location and proximity. In particular, the RTA extracts the IDs from the data sets.

[0056] The RTA then compares the IDs from different data sets, sent by different RAs for different processes. Matching IDs resulting from this comparison identify proximity between the resources.

[0057] Proximity may be derived from resource location as follows. Here, the two processes p1 and p2 have two ID stacks:

[0058] Stack_p1 (ID1_host, ID1_vm, ID1_cont, ID1_asi)

[0059] Stack_p2 (ID2_host, ID2_vm, ID2_cont).

It is noted that not all resource layers are mandatory (e.g., the process p2 does not utilize an application service instance resource).

[0060] The two processes p1 and p2 share resource constraints, if one of the IDs in the stack match, e.g.:

[0061] ID1_host=ID2_host, indicating both processes run on the same host;

[0062] ID1_vm=ID2_vm, indicating both processes run in the same VM; and

[0063] ID1_cont≠ID2_cont, indicating the processes run in different containers.

[0064] This location-based procedure also works if the two stacks are not using the same layers. For example:

[0065] Stack_p1 (ID1_host, ID1_vm, ID1_asi)—a process running in an application server running in a VM; and

[0066] Stack_p2 (ID2_host, ID2_cont)—a process running in a container.

[0067] Thus FIG. 8 shows that processes p1 and p2 share resource constraints because one of the IDs in the stack match:

[0068] ID1_host=ID2_host, indicating both processes run on the same host.

[0069] FIG. 9 shows the activity of a resource modeling infrastructure (RMI) 900 to analyze constraints. In particular, the RMI defines for individual units, constraints of certain resource types or reading from physical machines. For example, here the RMI may indicate the limit of processing power (e.g., in operations per second) of a particular physical CPU, which is comprised of virtualized components (e.g., VMs, containers, etc.).

[0070] The constraint is read from the host with the ID: ID_host (=ID1_host=ID2_host). These constraints are the maximum value for the two processes with matching IDs, but there may be more processes using the same host and competing on the resource. These can be found, for example by searching the monitoring system for more processes reporting a host with the ID: ID_host.

[0071] In this manner, embodiments allow identifying if two processes run in the same VM, and the VM is constrained to have 2 CPU only, even if the processes would run in an application server instance and the application server in a container.

[0072] The RMI then sends these constraints to the RTA, together with the IDs. In this example these IDs indicate host, VM, container, application server instance, etc.

[0073] The RTA then proceeds to store the resource constraints together with IDs. The RTA finds matching IDs of units of processes, and reads the provided constraints. The RTA then finds the parent units of the units with the

matching IDs, and reads the constraints from those parent units, thus providing a set of constraint of resource types for two processes.

[0074] According to certain embodiments, a procedure for constraint analysis utilizing coordinate location data may move from child-to-parent in order to identify bottlenecks. In alternative embodiments the procedure may move from parent-to-child through a resource hierarchy.

[0075] The specific embodiments described above are provided for purposes of illustration only, and embodiments are not limited to them. For example, while the above example has described constraints defined on different levels of CPU virtualization, embodiments are not limited to such environments.

[0076] For example, alternative embodiments may allow analysis of resource constraints defined on network topology. Consider a network topology shown in FIG. 10, with one switch serving three hosts, having a 10 GB/s line to the next switch serving a larger number of other switches and hosts, itself having a 100 GB/s line.

[0077] Considering a first level of the resource hierarchy of the network topology, hosts may each have nominal network bandwidth of 10 GB/s. Since the switch also has only 10 GB/s, the sum of the hosts can in total only be 10 GB/s.

[0078] However, a process on a host can experience lower bandwidth than the network card specification of 10 GB/s, as other processes on the same switch could consume bandwidth.

[0079] Considering a second level of the resource hierarchy of the network topology, even if the hosts attached to the same switch consume less than 10 GB/s bandwidth, a constraint on another switch can constrain the bandwidth available for a process. Thus even if the other hosts at the same switch do only consume 1 GB/s, a host may not get the remaining 9 GB/s, as the next switch may be loaded already with 99 GB/s and only has 100 GB/s capacity.

[0080] In this example, the topology of the network definition may be initially read by the scheduler component in order to define the resource coordinate system. Then, as shown in FIG. 11, an ID is defined for each unit in the network (network card, switch, switch . . .):

[0081] Stack_process<x>(ID<x>_card, ID<y>_switch, ID<z>_switch).

[0082] With such a resource coordinate system in place determining locations of individualization resource units, the monitoring instance is then poised to provide insight into resource consumption. The centralized monitoring instance may work with a modeling infrastructure to determine proximity and constraints as has been described above.

[0083] It is further noted that embodiments are not limited to analysis of resource constraints that are due to virtualization and/or network topology. Other examples of constrained resources can include input-output (IO) devices (e.g., memory), which may perform at a given number of operations per second. For example, processes operating in different containers can use network mounted disk drives and thus share potentially the same storage capacity and IO.

[0084] In certain embodiments, the resource analysis may be implemented by a database engine, for example as present in an in-memory database. One example of such an in-memory database engine is that of the HANA in-memory database available from SAP SE of Walldorf, Germany.

[0085] According to particular embodiments, the processing power available to an in-memory database engine may be leveraged to perform one or more aspects as have been described herein. For example, the in-memory database engine may be employed to perform one or more of the functions described above: scheduling, reporting, topology analysis, and/or resource modeling to analyze constraints.

[0086] In summary, resource analysis according to embodiments may find use in analyzing a resource problem associated with one particular process. Such a single-process analysis could be conducted as follows.

[0087] A process p1 has problems with one resource (e.g. CPU).

[0088] Embodiments use the resource coordinate system to identify various individualization units of the CPU, and read the constraints for every layer.

[0089] The process loops over the resource layers, looking in a current layer for processes with competing resources.

[0090] If this is a layer with higher layers above (e.g., children), then sum iteratively the sum of all processes with the same ID on that layer. That is, this is a query on matching IDs of a certain type for all monitored processes.

[0091] If the sum of the resources matches the constraints, then a bottleneck has been found. If the sum of the resources does not match the constraints, then the analysis moves to the next layer (the parent) in order to identify the bottleneck.

[0092] Resource analysis according to embodiments may also find use in analyzing a resource problem that may be associated with multiple processes. That is, a multi-process resource analysis could be conducted to identify particular resource units that are the subject of competition between suspected processes. Such a multi-process analysis could be conducted as follows.

[0093] Identify if the multiple processes have matching IDs in the coordinate system of resource units (e.g. CPU or network bandwidth).

[0094] If the processes do not have matching coordinates (e.g., their ID stacks do not reflect common individualization resource units), then the processes do not compete and the root of the resource problem does not lie in competition between those processes.

[0095] If the suspect processes do have matching coordinates (e.g., share particular individualization resource unit IDs in their stacks), they are likely in competition for at least one individualization resource unit. A user may then carefully inspect the constraints for the individualization resource unit having that shared ID, and potentially analyze more processes to hone in upon the likely source of the resource problem.

[0096] FIG. 12 illustrates hardware of a special purpose computing machine configured to implement resource analysis according to an embodiment. In particular, computer system 1201 comprises a processor 1202 that is in electronic communication with a non-transitory computer-readable storage medium comprising a database 1203. This computer-readable storage medium has stored thereon code 1205 corresponding to an engine. Code 1204 corresponds to location information. Code may be configured to reference data stored in a database of a non-transitory computer-readable storage medium, for example as may be present locally or in a remote database server. Software servers

together may form a cluster or logical network of computer systems programmed with software programs that communicate with each other and work together in order to process requests.

[0097] An example computer system 1300 is illustrated in FIG. 13. Computer system 1310 includes a bus 1305 or other communication mechanism for communicating information, and a processor 1301 coupled with bus 1305 for processing information. Computer system 1310 also includes a memory 1302 coupled to bus 1305 for storing information and instructions to be executed by processor 1301, including information and instructions for performing the techniques described above, for example. This memory may also be used for storing variables or other intermediate information during execution of instructions to be executed by processor 1301. Possible implementations of this memory may be, but are not limited to, random access memory (RAM), read only memory (ROM), or both. A storage device 1303 is also provided for storing information and instructions. Common forms of storage devices include, for example, a hard drive, a magnetic disk, an optical disk, a CD-ROM, a DVD, a flash memory, a USB memory card, or any other medium from which a computer can read. Storage device 1303 may include source code, binary code, or software files for performing the techniques above, for example. Storage device and memory are both examples of computer readable mediums.

[0098] Computer system 1310 may be coupled via bus 1305 to a display 1312, such as a cathode ray tube (CRT) or liquid crystal display (LCD), for displaying information to a computer user. An input device 1311 such as a keyboard and/or mouse is coupled to bus 1305 for communicating information and command selections from the user to processor 1301. The combination of these components allows the user to communicate with the system. In some systems, bus 1305 may be divided into multiple specialized buses.

[0099] Computer system 1310 also includes a network interface 1304 coupled with bus 1305. Network interface 1304 may provide two-way data communication between computer system 1310 and the local network 1320. The network interface 1304 may be a digital subscriber line (DSL) or a modem to provide data communication connection over a telephone line, for example. Another example of the network interface is a local area network (LAN) card to provide a data communication connection to a compatible LAN. Wireless links are another example. In any such implementation, network interface 1304 sends and receives electrical, electromagnetic, or optical signals that carry digital data streams representing various types of information.

[0100] Computer system 1310 can send and receive information, including messages or other interface actions, through the network interface 1304 across a local network 1320, an Intranet, or the Internet 1330. For a local network, computer system 1310 may communicate with a plurality of other computer machines, such as server 1315. Accordingly, computer system 1310 and server computer systems represented by server 1315 may form a cloud computing network, which may be programmed with processes described herein. In the Internet example, software components or services may reside on multiple different computer systems 1310 or servers 1331-1335 across the network. The processes described above may be implemented on one or more servers, for example. A server 1331 may transmit actions or

messages from one component, through Internet 1330, local network 1320, and network interface 1304 to a component on computer system 1310. The software components and processes described above may be implemented on any computer system and send and/or receive information across a network, for example.

[0101] The above description illustrates various embodiments of the present invention along with examples of how aspects of the present invention may be implemented. The above examples and embodiments should not be deemed to be the only embodiments, and are presented to illustrate the flexibility and advantages of the present invention as defined by the following claims. Based on the above disclosure and the following claims, other arrangements, embodiments, implementations and equivalents will be evident to those skilled in the art and may be employed without departing from the spirit and scope of the invention as defined by the claims.

What is claimed is:

1. A computer-implemented method comprising:
 - assigning unique identifiers to each of a plurality of resources within a hierarchy of a computing environment to define location information;
 - storing the location information;
 - reporting the location information to a central monitoring instance;
 - the central monitoring instance determining from the location information, a first set of resources consumed by a first process based upon a first identifier stack comprising one or more of the unique identifiers; and
 - the central monitoring instance displaying the first set of resources.
2. A method as in claim 1 further comprising:
 - the central monitoring instance receiving from a modeling infrastructure, a constraint of one resource of the first set of resources; and
 - the central monitoring instance revealing the one resource as a performance bottleneck of the first process based upon the constraint.
3. A method as in claim 1 further comprising:
 - the central monitoring instance determining from the location information, a second set of resources consumed by a second process based upon a second identifier stack comprising one or more of the unique identifiers; and
 - the central monitoring instance comparing the first identifier stack with the second identifier stack to determine a proximity between the first process and the second process.
4. A method as in claim 3 further comprising:
 - the central monitoring instance receiving from a modeling infrastructure, a constraint of one resource identified as proximate between the first process and the second process.
5. A method as in claim 4 further comprising:
 - the central monitoring instance revealing the one resource as a performance bottleneck of the first process based upon the constraint.
6. A method as in claim 4 wherein:
 - the computing environment comprises a virtualization computing environment; and
 - the constraint comprises a number of virtual machines supported by a hardware central processing unit.

7. A method as in claim 4 wherein:
the computing environment comprises a network topology including a switch; and
the constraint comprises a bandwidth of the switch.
8. A method as in claim 4 wherein:
the computing environment comprises a common resource shared between the first process and the second process; and
the constraint comprises a number of operations per second of the common resource.
9. A method as in claim 8 wherein the common resource comprises an input-output device.
10. A method as in claim 9 wherein the common resource comprises a memory.
11. A method as in claim 1 wherein:
the computing environment comprises a virtualization computing environment; and
the hierarchy comprises a virtual machine.
12. A non-transitory computer readable storage medium embodying a computer program for performing a method, said method comprising:
assigning unique identifiers to each of a plurality of resources within a hierarchy of a computing environment to define location information;
storing the location information;
reporting the location information to a central monitoring instance;
the central monitoring instance determining from the location information, a first set of resources consumed by a first process based upon a first identifier stack comprising one or more of the unique identifiers;
the central monitoring instance receiving from a modeling infrastructure, a constraint of one resource of the first set of resources; and
the central monitoring instance displaying the first set of resources and the constraint.
13. A non-transitory computer readable storage medium as in claim 12 wherein the method further comprises:
the central monitoring instance determining from the location information, a second set of resources consumed by a second process based upon a second identifier stack comprising one or more of the unique identifiers; and
the central monitoring instance comparing the first identifier stack with the second identifier stack to determine a proximity between the first process and the second process.
14. A non-transitory computer readable storage medium as in claim 13 wherein the method further comprises:
the central monitoring instance revealing the one resource as a performance bottleneck of the first process based upon the constraint.
15. A non-transitory computer readable storage medium as in claim 12 wherein:
the computing environment comprises a virtualization computing environment; and
the hierarchy comprises a virtual machine.
16. A computer system comprising:
one or more processors;
a software program, executable on said computer system, the software program configured to cause an in-memory database engine to:
assigning unique identifiers to each of a plurality of resources within a hierarchy of a computing environment to define location information;
store the location information in an in-memory database;
determine from the location information, a first set of resources consumed by a first process based upon a first identifier stack comprising one or more of the unique identifiers;
determining from the location information, a second set of resources consumed by a second process based upon a second identifier stack comprising one or more of the unique identifiers;
compare the first identifier stack with the second identifier stack to determine a proximity between the first process and the second process; and
communicate the proximity for display.
17. A computer system as in claim 16 wherein the software program is further configured to cause the in-memory database engine to:
receive a constraint of one resource of the first set of resources;
store the constraint in the in-memory database; and
reveal the one resource as a performance bottleneck based upon the constraint.
18. A computer system as in claim 17 wherein:
the computing environment comprises a virtualization computing environment; and
the hierarchy comprises a virtual machine.
19. A computer system as in claim 17 wherein:
the computing environment comprises a network topology including a switch; and
the constraint comprises a bandwidth of the switch.
20. A computer system as in claim 17 wherein the constraint comprises a number of operations per second of the one resource.

* * * * *