

**(12) STANDARD PATENT**  
**(19) AUSTRALIAN PATENT OFFICE**

(11) Application No. **AU 2012259086 B2**

(54) Title  
**Cross-cloud management and troubleshooting**

(51) International Patent Classification(s)  
**G06F 15/16** (2006.01) **G06F 9/44** (2006.01)  
**G06F 9/06** (2006.01)

(21) Application No: **2012259086** (22) Date of Filing: **2012.05.18**

(87) WIPO No: **WO12/162171**

(30) Priority Data

(31)	Number	(32)	Date	(33)	Country
	<b>13/111,956</b>		<b>2011.05.20</b>		<b>US</b>

(43) Publication Date: **2012.11.29**

(44) Accepted Journal Date: **2016.09.22**

(71) Applicant(s)  
**Microsoft Technology Licensing, LLC**

(72) Inventor(s)  
**Iyer, Kannan C.;Watson, Eric B.**

(74) Agent / Attorney  
**Davies Collison Cave Pty Ltd, Level 15 1 Nicholson Street, MELBOURNE, VIC, 3000**

(56) Related Art  
**US 2011/0055377 A1**  
**US 2011/0078303 A1**



## (51) International Patent Classification:

G06F 15/16 (2006.01) G06F 9/44 (2006.01)  
G06F 9/06 (2006.01)

## (21) International Application Number:

PCT/US2012/038647

## (22) International Filing Date:

18 May 2012 (18.05.2012)

## (25) Filing Language:

English

## (26) Publication Language:

English

## (30) Priority Data:

13/111,956 20 May 2011 (20.05.2011) US

(71) Applicant (for all designated States except US): **MICROSOFT CORPORATION** [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).(72) Inventors: **IYER, Kannan C.**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **WATSON, Eric B.**; c/o Microsoft Corporation, LCA - International Patents, 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, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.


(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

## Declarations under Rule 4.17:

— as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))

[Continued on next page]

## (54) Title: CROSS-CLOUD MANAGEMENT AND TROUBLESHOOTING



**FIG. 1**

(57) Abstract: A cloud management system is described herein that provides the ability for an application to span two or more clouds while allowing operation, management, and troubleshooting of the distributed application as a single application. The system provides infrastructure that communicates across datacenters for execution and for centralizing knowledge of instances of an application that are running at different locations. The infrastructure provided by the system monitors both the application and connections between the clouds, with intelligence to know if issues are within the application or because of the connection between the clouds. The system coordinates management functions across multiple cloud platforms/locations. Thus, the cloud management system creates a single monitoring and troubleshooting interface and knowledge and execution fabric across multiple clouds so that applications spread across multiple clouds can be monitored, managed, and debugged more easily.



— as to the applicant's entitlement to claim the priority of  
the earlier application (Rule 4.17(iii))

— before the expiration of the time limit for amending the  
claims and to be republished in the event of receipt of  
amendments (Rule 48.2(h))

**Published:**

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

**(88) Date of publication of the international search report:**

21 February 2013

**BACKGROUND**

[0001] Datacenters provide servers for running large applications. Enterprises often use datacenters to run core business functions such as sales, marketing, human resources, billing, product catalogs, and so forth. Datacenters may also run customer-facing applications, such as web sites, web services, email hosts, databases, and many other applications. Datacenters are typically built by determining an expected peak load and providing servers, network infrastructure, cooling, and other resources to handle the peak load level. Datacenters are known for being very expensive and for being underutilized at non-peak times. They also involve a relatively high management expense in terms of both equipment and personnel for monitoring and performing maintenance on the datacenter. Because almost every enterprise uses a datacenter of some sort, there are many redundant functions performed by organizations across the world.

[0002] Cloud computing has emerged as one optimization of the traditional datacenter. A cloud is defined as a set of resources (e.g., processing, storage, or other resources) available through a network that can serve at least some traditional datacenter functions for an enterprise. A cloud often involves a layer of abstraction such that the applications and users of the cloud may not know the specific hardware that the applications are running on, where the hardware is located, and so forth. This allows the cloud operator some additional freedom in terms of rotating resources into and out of service, maintenance, and so on. Clouds may include public clouds, such as MICROSOFT™ Azure, Amazon Web Services, and others, as well as private clouds, such as those provided by Eucalyptus Systems, MICROSOFT™, and others. Companies have begun offering appliances (e.g., the MICROSOFT™ Azure Appliance) that enterprises can place in their own datacenters to connect the datacenter with varying levels of cloud functionality.

[0003] Enterprises with datacenters incur substantial costs building out large datacenters, even when cloud-based resources are leveraged. Enterprises often still planned for “worst-case” peak scenarios and thus include an amount of hardware at least some of which is rarely used or underutilized in terms of extra processing capacity, extra storage space, and so forth. This extra amount of resources incurs a high cost for little return. Customers using cloud based computing on premise expect to be able to use capacity in another compatible cloud (e.g., a second instance of their own in another location, Microsoft’s public cloud, and so forth) for peak capacity times, for disaster recover scenarios, or just for capacity management. Doing so is much less expensive than

building out for the worst-case scenario and then doubling for redundancy. In addition, they expect to be able to manage (e.g., troubleshoot, operate) applications split across multiple clouds. Today, applications, cloud management, and troubleshooting do not operate across clouds or other datacenters.

- 5 [0003a] It is desired to address or ameliorate one or more disadvantages or limitations associated with the prior art, or to at least provide a useful alternative.

#### SUMMARY

- 10 [0003b] In one embodiment, the present invention provides a computer-implemented method to handle a request from a management tool to access application management data from distributed application instances, the method comprising: receiving from an application management tool a management tool request to access management data related to an application running instances in one or more datacenters, wherein receiving the management tool request comprises receiving a request from a performance monitoring tool to access status information describing operation of one or more application instances; identifying one or more types of management data that satisfy the received request; determining a distribution of the application that includes two or more instances of the application; gathering management data to satisfy the request from each distributed application instance; unifying gathered management data to provide a uniform response to the received management tool request; and reporting the gathered and unified management data in response to the received management tool request, wherein the preceding steps are performed by at least one processor.

- 20 [0003c] In another embodiment, the present invention provides a computer system for cross-cloud computing for cross-cloud management and troubleshooting, the computer system comprising: a processor and memory configured to execute software instructions embodied within the following components; a location management component that manages information about multiple datacenter locations at which instances of an application are running; a location data store that stores information that describes locations at which instances of the application are running; a tool interface component that provides an interface to the computer system through which one or more tools can access management and troubleshooting information for the application; one or more management tools that connect to the tool interface component to access management information or perform application troubleshooting; a data migration component that migrates management information at one or more remote locations where the application is running back to a home location of the application; and a troubleshooting component

that performs one or more troubleshooting tasks on the application at one or more locations.

[0003d] In a further embodiment, the present invention provides a computer-readable storage device comprising instructions for controlling a computer system to report data back from and handle troubleshooting requests at a location of a remote application instance, wherein the instructions, upon execution, cause a processor to perform actions comprising: receiving management data at a remote application instance handling a portion of load generated by requests from clients of the application; determining a home location of the application where an administrator can access management data reported by multiple instances of the application running at one or more distributed remote locations, the home location being a location where the application operates under steady conditions; sending the received management data from the remote application instance to the determined home location of the application; receiving a troubleshooting request from a management tool run at the home location to troubleshoot the remote application instance; performing one or more troubleshooting actions in response to the received troubleshooting request; and sending a troubleshooting result to the home location in response to the received troubleshooting request.

[PARAGRAPH 4 DELETED]

[0005] This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used to limit the scope of the claimed subject matter.

#### BRIEF DESCRIPTION OF THE DRAWINGS

[0005a] Preferred embodiments of the present invention are hereinafter described, by way of example only, with reference to the accompanying drawings, in which:

[0006] Figure 1 illustrates an application running in two clouds with associated management infrastructure, in one embodiment.

[0007] Figure 2 is a block diagram that illustrates components of the cloud management system, in one embodiment.

[0008] Figure 3 is a flow diagram that illustrates processing of the cloud management system to handle a request from a management tool to access data from distributed application instances, in one embodiment.

[0009] Figure 4 is a flow diagram that illustrates processing of the cloud management system to report data back from and handle troubleshooting requests at a location of a remote application instance, in one embodiment.

#### DETAILED DESCRIPTION

5 [0010] A cloud management system is described herein that provides the ability for an application to span two or more clouds (which may be across great distances), while allowing operation, management, and troubleshooting of the distributed application as a single application. The system provides infrastructure that communicates across datacenters for execution and for centralizing knowledge of instances of an application  
10 that are running at different locations. For example, the system may centralize logging, performance tracking, and other management functions, regardless of where the application is running. In some cases, the system provides a computing appliance that an enterprise can place in its own private datacenter that allows an administrator to distribute at least some application loads to a public cloud or other separate locations, while  
15 providing unified management via the computing appliance.

[0011] The infrastructure provided by the cloud management system monitors both the application and connections between the clouds, with intelligence to know if issues are within the application or because of the connection between the clouds. The system coordinates management functions across multiple cloud platforms/locations (from one  
20 cloud's infrastructure, tasks are coordinated to execute across two or more clouds). If an administrator wants to debug the application, the system allows live debugging at the correct location through a seamless, unified interface. Thus, the cloud management system creates a single monitoring and troubleshooting interface and knowledge and execution "fabric" across multiple clouds so that applications spread across multiple  
25 clouds can be monitored, managed, and debugged more easily.

[0012] Figure 1 illustrates an application running in two clouds with associated management infrastructure, in one embodiment. In some embodiments, the cloud management system involves the application (and/or the administrator) using infrastructure in one cloud that has the data/access to the data at all locations to be able to  
30 fully monitor/troubleshoot the application. As an example, consider an application with instances running in two clouds, cloud 110 and cloud 150 as shown in Figure 1. Cloud

110 includes a MICROSOFT™ Azure appliance instance 120 that includes infrastructure 130. The appliance instance 120 includes application instance 125 that is running role 140 and role 145. A second cloud 150 includes application instance 155 that is running role 160 and role 170. The second cloud 150 also includes infrastructure 180.

- 5 The appliance instance 120 knows about each of the roles and that they are part of the same application. The infrastructure plumbing at each location allows the appliance instance 120 to retrieve information about role 160 and role 170 executing at the second cloud 150. The system can distribute either individual roles, whole applications, or both. With all of the management data (e.g., logs from the applications, machines, and
- 10 infrastructure), the system can assess the health of the application just as if all of the roles were local by applying pre-defined health rules. The system can also see the infrastructure health across both locations as well as the connection 190 in between to assess if a problem is occurring with the application or the infrastructure/network.

- [0013] Similarly, when automated or manual troubleshooting or remediation steps are
- 15 needed, the infrastructure 130 in cloud 110 can coordinate with the infrastructure 180 in cloud 150 to provide troubleshooting and debugging support. For example, the system fabric can reach across locations to execute an application wide update, shutdown, and so forth. Those of ordinary skill in the art will recognize numerous ways to perform cross-location control. For example, infrastructure 130 may directly control infrastructure 180,
- 20 infrastructure 130 may request infrastructure 180 to execute on infrastructure 130's behalf, and so forth. Likewise, with operator/administrator troubleshooting tools (e.g., monitoring visualization, alerting, log and configuration data viewing, and so on), the location of the applications and infrastructure is available and logically displayed, but does not involve separate tools and mental gymnastics from the administrator to put together. For instance,
- 25 when troubleshooting and viewing data on all roles, if the administrator 105's next step is using one or more tools 195 to view the application's logs or start a remote session to the role instance, the system connects the administrator 105 directly, regardless of at which location the role is residing.

- [0014] The design of the cloud management system provides simplified and consistent
- 30 running of a service across multiple clouds/location. The system moves the definition of "a computing resource" from a server, beyond a datacenter to a portion of the internet (the datacenters and the connection between them). This allows service level agreements (SLAs) to be defined, monitored, and managed at the service level – which is what service owners often care most about.



[0015] In some embodiments, the cloud management system works in cooperation with a cloud migration system that seamlessly migrates applications from one location to another as needed, called bursting. The cloud migration system provides capacity management and disaster recovery by detecting peak load conditions and automatically moving computing to another source (and back) and by providing computing across two or more clouds and moving completely to one in the case of a disaster at one site. This allows enterprises to plan for local resources for a sustained level of load and to leverage cloud-based resources for peak or other unusual loads. In many cases, an enterprise's business is such that a particular time of year is busier, and extra resources may only be needed during those times. For example, tax-planning enterprises are particular busy in mid-April, e-commerce sites experience holiday rushes around Thanksgiving and Christmas, and so on. The cloud migration system monitors loads within a datacenter and detects a threshold that indicates that the current load is nearing the datacenter's capacity. For example, the system may monitor central processing unit (CPU) usage, memory usage, storage usage, network bandwidth, and other metrics to determine how well the datacenter is handling the current load. The system may also observe trends (e.g., a rate of acceleration of resource usage) to determine whether the threshold has been or will soon be reached.

[0016] Upon detecting that the threshold will be reached, the cloud migration system facilitates an orderly move of at least some datacenter load to another datacenter or cloud-based resources. For example, the system may migrate some peak load to a public cloud. Because cloud pricing models may vary, the system may factor cost into the decision. For example, the system may prefer to host as much load as possible at the enterprise datacenter to reduce cost, while leveraging cloud resources only to the extent needed to satisfy client requests. The system may also provide management and monitoring tools that provide a consistent experience for information technology (IT) personnel regardless of where particular loads are run (e.g., locally within the enterprise or publicly using a cloud). The system may also provide planning tools to help decide appropriate workloads or applications for moving to other resources during high loads. For example, applications may have various compliance/regulatory or networking/design limitations that make them more or less suitable for migration. The system can also be used as a disaster recovery architecture at a datacenter/network level to manage fast workload transition in case of disaster. If a datacenter resource permanently fails, the system can quickly and efficiently migrate additional load to the cloud or other resources so that clients of the datacenter are

unaffected or are less affected by the failure. Thus, the cloud migration system allows enterprises to build smaller and more efficient datacenters that leverage other resources for rare extra loads.

5 [0017] The cloud management system works with the cloud migration system to provide seamless management and troubleshooting as applications are migrated from one location to another. As described above, the cloud migration system may move resources between a datacenter and the cloud on a temporary (i.e., bursting) or permanent (i.e., disaster recovery) basis. Temporary movements include bursting an application or other load for a short time period to handle a peak or other high load that exceeds the datacenter's  
10 capacity. A temporary movement may include bursting an entire application or splitting the application's load across two or more locations. Permanent movements include longer-term migration of loads due to a failure of hardware in the datacenter, a more sustained increase in capacity needs, a desire to globally distribute an application with dynamic load balancing, and so forth. Following are several example scenarios in which  
15 the system may be used by an enterprise.

[0018] In the first example, an enterprise bursts application load to a public cloud to manage capacity. The business decision maker (i.e., CEO, CFO, or VP Marketing/sales) and datacenter systems administrator decide it would be more cost effective and provide a better customer experience to burst some work to the public cloud at their top three peak  
20 day level of use/traffic per year, and maintain their own datacenter (potentially with a cloud appliance) at their peak monthly usage level. They sign business agreements with the cloud provider to burst work into the cloud and project estimates of when and how much work that would be. Their account is set up and the information entered into the cloud appliance. During a planning phase, the administrator runs a test with a test  
25 application from the cloud provider that ensures the connection is working properly. The administrator then sets the capacity values (e.g., threshold) for initiating bursting of applications that keep capacity at the specified level in a capacity management tool. The administrator goes into the tool to further specify the applications that are eligible to move in this situation (e.g., no regulatory issues with temporary movement, good technical fit).

30 [0019] The day comes when use exceeds the limits and the system automatically moves applications to the public cloud. Alerts are thrown in the monitoring/usage systems when capacity is within 5% of the bursting being initiated, when the system does burst, what the system bursts, and when the system brings the applications back. An explicit log is kept of all compute resources and/or storage moved and the administrator is alerted to go to

their public cloud account for billing. A review of the bursting parameters and applications tagged as movable is reviewed in the regular capacity planning meetings with in the enterprise datacenter group and management.

**[0020]** In a second example, an enterprise splits applications across clouds to manage capacity. This scenario is similar to the scenario above except for the type of application moved is more complex, so is split out to prioritize differently. The company decides to have a relationship with the cloud provider for splitting applications into the cloud (a form of bursting). In this case, a large application was pre-identified as a bursting candidate. When the capacity reaches the threshold, 50 of the 100 worker instances are automatically moved to the public cloud. The application is now split across two appliance instances or cloud instances, with all monitoring and billing data being sent to the starting instance so it can be centrally managed. A cloud appliance in the enterprise's own datacenter has troubleshooting tools to help debug possible issues of split applications (e.g., networking issues, network bandwidth/latency issues, fabric communication, and so on). When the capacity situation has subsided on the appliance, the 50 worker instances are moved back to the appliance and it is a normally functioning application again.

**[0021]** In another example, a cloud provider decides to burst from one cluster to another. The public cloud capacity planning team decides that a cluster in the Chicago datacenter is critically full but wants to maintain high utilization. They set up bursting to an underutilized cluster in a West coast datacenter when utilization gets to 90%. The administrator goes into the capacity management tool and chooses appropriate customers/applications (e.g., with low data usage) to be move candidates. The day comes when the usage of the Chicago cluster reaches the threshold and the system automatically moves the selected applications (e.g., 10% of the cluster's apps) to the West coast datacenter for one day. As usage returns below the threshold, the system moves the applications back to Chicago. The system notifies a designated monitoring team proactively of the burst to be able to answer customer questions.

**[0022]** In another example, the system is used for cross-cloud portfolio management. An enterprise decides that to manage capacity efficiently on their cloud appliance, they want to put all variable demand apps in a public cloud and their constant demand applications in the appliance or local datacenter resources (and thus be able to run the appliance at higher utilization). While they want their computing resources split, they still want a global view across all of their application's health, to have their application developers manage applications the same way, and to maintain a single view of

departmental billing across both (e.g., what costs to allocate to the consumer sales groups, internal IT, B2B sales, and so forth). The enterprise is able to set up the aggregation accounts with the public cloud with the same groups as the appliance and get billing data to integrate on their side. Similarly, they are able to get application-programming interface (API) access to the public cloud monitoring data for the platform their apps are running on as well as application level monitoring, so their network operation center (NOC) has a complete and consistent view of the state of the computing activity of the enterprise.

5  
10  
15  
20  
[0023] In another example, an enterprise sets up a globally distributed application with dynamic load balancing. An enterprise customer wants to manage capacity across two or more cloud instances and has a significant amount of their load in independent but geographically distributed instances (e.g., Bing search with a US and UK datacenter that both serve German queries). Under normal circumstances, a global traffic manager sends 50% traffic to each location. When load gets high at the primary location, the system instructs the load balancer to send 75% of the traffic to the UK system, thus freeing up capacity from the US cloud instance, bringing it to acceptable levels. When capacity returns to normal, the system tells the load balancer to return to the 50/50 split. A variation of this is for the public cloud to be used as a secondary datacenter (with say 1% of load, the customer's site with the appliance to be the other 99%). In the case of a disaster or other reason to move load from the customer's site, 100% of traffic is shifted to the public cloud.

25  
[0024] In another example, an enterprise has reached its datacenter's capacity and needs extra computing resources, but does not yet have the available capital to expend to expand the datacenter. In this case, the company can use a public cloud for spillover until they can get the hardware purchase completed.

30  
[0025] Figure 2 is a block diagram that illustrates components of the cloud management system, in one embodiment. The system 200 includes a location management component 210, location data store 220, tool interface component 230, one or more management tools 240, a data migration component 250, a troubleshooting component 260, and a billing component 270. Each of these components is described in further detail herein.

[0026] The location management component 210 manages information about multiple datacenter locations at which instances of an application are running. The component 210 includes information describing how to reach each location, connections available for retrieving management information, user accounts to use for each location with associated

security credentials, application and datacenter components from which to gather troubleshooting information and send troubleshooting commands, and so forth. The location management component 210 receives information describing any migration of application loads or bursting from one datacenter/cloud to another and updates the managed information so that the component 210 has a complete picture of all of the locations where the application is running. This allows the system 200 to present the complete picture and to make management of applications uniform, no matter where or at how many locations the applications are running. As conditions change and applications are distributed, the location management component 210 can present management tools with a comprehensive set of management data.

[0027] The location data store 220 stores information that describes locations at which instances of the application are running. The data store 220 may include one or more files, file system, hard drives, databases, cloud-based storage services, or other facilities for persisting information between sessions with the system 200. The stored information may include connection information, user roles, sources of management data, available log files, and any other information related to management or troubleshooting of applications distributed to multiple locations.

[0028] The tool interface component 230 provides an interface to the system 200 through which one or more tools can access management and troubleshooting information for the application. The interface may include one or more web pages, web services, application-programming interfaces (APIs), or other interfaces through which an administrator or tools can directly or programmatically access management and troubleshooting information of the system 200. In some embodiments, the tool interface component 230 provides an initial connection point for tools to access information related to the application at a cloud-computing appliance located within an enterprise's private datacenter. The appliance may manage migration and distribution of application loads to a public cloud or other datacenter, and provides a central point of contact for tools that gather management information or provide application troubleshooting.

[0029] One or more management tools 240 connect to the tool interface component 230 to access management information or perform application troubleshooting. The tools may include log viewers, reporting tools, debugging tools, or other tools that display information about or assist in resolving problems with a running application. The management tools 240 may include tools designed to work with a local application, and the system 200 provides the tools with information describing a distributed application

running at multiple locations without the tool's knowledge. This allows existing tools that administrators rely upon to be used even as automatic application load migration is introduced into a datacenter or cloud. In other cases, tools may be specifically written to understand distributed applications and to provide specific management information or troubleshooting related to multiple locations. The tool interface component 230 may provide multiple interfaces through which management tools 240 connect to the system 200 using paradigms that are understood by each tool.

5 [0030] The data migration component 250 migrates management information at one or more remote locations where the application is running back to a home location of the application. The home location may include a private datacenter, location or a cloud-computing appliance, or other location where the application normally runs under steady conditions. Upon hitting a certain level of load (e.g., peak or periodic bursts), the application may migrate some load to one or more other datacenters or public clouds to help satisfy client requests. These other locations generate management data, such as log  
10 files, transaction data, and so on, just like the home location, and the data migration component 250 migrates this data back to the home location or provides access to the data from the home location, so that the management tools 240 can provide a comprehensive picture of the application's activity to administrators.

[0031] The troubleshooting component 260 performs troubleshooting tasks on the application at one or more locations. Troubleshooting may include debugging, processing test data, or other forms of determining whether an application is operating correctly. Troubleshooting is generally well understood at the home location, but becomes more complex as an application begins to span multiple datacenters or clouds. The cloud management system 200 isolates management tools 240 and administrators from this  
25 complexity by providing a uniform interface through which tools and administrators access management information and perform troubleshooting at multiple locations. Thus, if a management tool allows an administrator to place a breakpoint at or receive trace information from a particular piece of application code at the home location, then the troubleshooting component 260 makes it just as easy to do so at a remote cloud-based  
30 instance of the application. The tools and administrator may even be unaware of all of the locations where the application is running, but can still perform management tasks as if the application were only executing at the home location.

[0032] The billing component 270 reports billing information related to one or more locations where an application is running. One common management task is to manage

computing costs, and public clouds often charge based on metrics related to workload (e.g., computing time, storage space used, and so forth). It may be useful for an administrator to gather a picture of the costs that application instances are incurring at various locations, and the cloud management system 200 can optionally provide the billing component 270 to gather this type of information so that the information can be reported through management tools and reports.

5 [0033] The computing device on which the cloud management system is implemented may include a central processing unit, memory, input devices (e.g., keyboard and pointing devices), output devices (e.g., display devices), and storage devices (e.g., disk drives or other non-volatile storage media). The memory and storage devices are computer-readable storage media that may be encoded with computer-executable instructions (e.g., software) that implement or enable the system. In addition, the data structures and message structures may be stored or transmitted via a data transmission medium, such as a signal on a communication link. Various communication links may be used, such as the Internet, a local area network, a wide area network, a point-to-point dial-up connection, a cell phone network, and so on.

15 [0034] Embodiments of the system may be implemented in various operating environments that include personal computers, server computers, handheld or laptop devices, multiprocessor systems, microprocessor-based systems, programmable consumer electronics, digital cameras, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, set top boxes, systems on a chip (SOCs), and so on. The computer systems may be cell phones, personal digital assistants, smart phones, personal computers, programmable consumer electronics, digital cameras, and so on.

25 [0035] The system may be described in the general context of computer-executable instructions, such as program modules, executed by one or more computers or other devices. Generally, program modules include routines, programs, objects, components, data structures, and so on that perform particular tasks or implement particular abstract data types. Typically, the functionality of the program modules may be combined or distributed as desired in various embodiments.

30 [0036] Figure 3 is a flow diagram that illustrates processing of the cloud management system to handle a request from a management tool to access data from distributed application instances, in one embodiment. Beginning in block 310, the system receives from a management tool a request to access management data related to an application

running instances in one or more datacenters. For example, a performance-monitoring tool may request status information describing how many client requests the application is handling, resource usage of the application, or other information from the application.

The system may receive the tool request through an API the system exposes to tools for  
5 requesting management data. The API may comprise a uniform interface for accessing management data irrespective of where or at how many locations application instances are running.

[0037] Continuing in block 320, the system identifies one or more types of management data that satisfy the received request. For example, the system may determine that the  
10 request asks for log information that is produced by each instance of the application. Identifying the requested data allows the system to determine which information to gather from each application instance or whether the data is already gathered locally from data pushed to a central location by each application instance.

[0038] Continuing in block 330, the system determines a distribution of the application that includes two or more instances of the application. The distribution determines where  
15 the application is running and where the system will find management data to satisfy the request. The system may include a data store that tracks information describing each burst or other migration of application load to and from other datacenters, so that the system is aware of each location where application instances are running. Upon receiving the  
20 management tool request, this information allows the system to determine from where to gather management data.

[0039] Continuing in block 340, the system gathers management data to satisfy the request from each distributed application instance. The instances may include an instance in a local private datacenter, a remote private datacenter, a private cloud computing  
25 facility, a public cloud computing facility, spare resources offered by other private datacenters, and so on. The system contacts each instance of the application or accesses previously sent information from each instance that contains information (such as performance data, failures, and so forth) to satisfy the received management tool request.

[0040] Continuing in block 350, the system optionally sends one or more  
30 troubleshooting commands to one or more remote application instances. For example, if one location is experiencing failures, the administrator may use a management tool to request additional trace information, to send one or more test requests, or to perform other types of debugging. The remote application instances carry out the troubleshooting



commands and report requested data back to a central location where the management tool can access the information.

- 5 [0041] Continuing in block 360, the system unifies gathered data to provide a uniform response to the received management tool request. In this way, management tools need not be written to include an understanding of the various potential distributions of applications managed by the tools. The system can thus freely migrate the application from location to location or to multiple locations as needed to handle application loads, while still providing administrators with a straightforward management and troubleshooting experience.
- 10 [0042] Continuing in block 370, the system reports the gathered and unified management data in response to the received management tool request. The system may send the data through the interface on which the request was received or through a notification interface or other facility for providing data to the tool. After block 370, these steps conclude.
- 15 [0043] Figure 4 is a flow diagram that illustrates processing of the cloud management system to report data back from and handle troubleshooting requests at a location of a remote application instance, in one embodiment. Beginning in block 410, the system receives management data at a remote application instance handling a portion of load generated by requests from clients of the application. The management data may include
- 20 performance data, log information, error details, statistical information, sales history, or other indications of application operation useful for management the application.
- [0044] Continuing in block 420, the system determines a home location of the application where an administrator can access management data reported by multiple instances of the application running at distributed remote locations. The application
- 25 instance may receive configuration information from the home location upon creation of the instance that specifies where the home location can be contacted and that the application instance is a remote instance of the application. The system may migrate applications to multiple locations to handle peak loads, perform low priority tasks at locations where processing is off-peak and thus cheaper, or for other reasons determined
- 30 by an administrator. The application may have a home location that is where the application normally runs and may handle peak or other loads at one or more distributed remote locations.
- [0045] Continuing in block 430, the system sends the received management data from the remote application instance to the determined home location of the application. The

system may periodically migrate data generated at distributed instances back to the home location so that management data is available in one place at the home location for the convenience of administrators and management tools. The system may also migrate data on demand or as requested by various tools (see, e.g., Figure 3). In some cases, the system  
5 may burst application loads to remote locations for short durations and then collect information related to the application's execution when the loads are migrated back to the home location and the remote instances are finished.

[0046] Continuing in block 440, the system optionally receives a troubleshooting request from a management tool run at the home location to troubleshoot the remote  
10 application instance. The troubleshooting requests may include debug breakpoints, a request for detailed trace information, or other commands or requests to perform troubleshooting actions.

[0047] Continuing in block 450, the system performs one or more troubleshooting actions in response to the received troubleshooting request. The action may include  
15 setting a debug breakpoint, turning up a logging level, sending test data to the application, or performing any other action specified by the request for determining whether the application is operating properly.

[0048] Continuing in block 460, the system sends a troubleshooting result to the home location in response to the received troubleshooting request. By providing a facility for  
20 executing troubleshooting commands remotely, the system allows a troubleshooting tool operating at the home location to troubleshoot application instances no matter where the instances are running, and allows the system to seamlessly migrate instances of the application to various locations without interrupting the ability of an administrator to manage and troubleshoot the application. After block 460, these steps conclude.

[0049] In some embodiments, the cloud management system migrates application load by modifying domain name service (DNS) records. The system may modify a DNS server to point incoming client requests to one or more new destination Internet Protocol (IP) addresses to direct loads away from a source datacenter over to a target datacenter/cloud. A global traffic manager (GTM) often points clients to the nearest server for handling their  
25 requests, and these solutions can be modified to redirect traffic based on load or other conditions. Thus, when one datacenter becomes overloaded or near capacity, the system may inform the GTM to direct at least some client requests to a new location that can  
30 handle the excess load. Similarly, the system may provide a DNS or other address to

which management tools can address management requests and be connected to application instances no matter where they reside.

[0050] In some embodiments, the cloud management system migrates log and other data back from the target computing resources after migration conditions have eased. For example, following a period of peak load, the system may migrate all application loads back to the original datacenter, and may pull information generated at the target datacenter, such as application logs, back to the original datacenter for later analysis. For some applications, tracking client requests may be a matter of regulatory compliance or simply useful for debugging and reporting. In either case, consolidating the logs at the source location may be part of a successful migration back to the source location.

[0051] In some embodiments, the cloud management system allocates a dynamically variable amount of application load between a source computing resource and one or more target computing resources. For example, the system may dynamically route requests to keep the source computing resource at or near full capacity and only send requests out to external computing resources that the source computing resource could not successfully handle. Such decisions may be a matter of cost, data safety, or other considerations to migrate out as little application load as needed or to place application loads where they can be performed cheapest or most efficiently. In some cases, the decisions may be based on regulatory requirements of applications. For example, applications subject to healthcare or other recordkeeping laws may have restrictions about the datacenters/clouds in which they can operate.

[0052] In some embodiments, the cloud management system provides various options for disaster recovery. In some cases, the system may enlist resources at an external datacenter to monitor a main datacenter for outages. If the external datacenter becomes unable to reach the main datacenter, then the external datacenter may determine that a disaster has occurred and move application loads to the external datacenter. In past systems, it was typical for an organization to maintain 200% of needed capacity (at substantial expense) in order to successfully handle disasters. With the cloud management system, the organization can maintain a lower amount of available capacity at a second location (e.g., 10%) and can rapidly request more as needed in the event of a failure. Much like insurance, the likelihood of all clients of a cloud provider failing at the same time and requesting a high capacity spare is low, such that multiple clients can share a set of redundant secondary resources to be used in the event of failure of primary resources.

The system may also re-home management tools and troubleshooting resources to point to the new location following disaster recovery so that management continues uninterrupted.

[0053] From the foregoing, it will be appreciated that specific embodiments of the cloud management system have been described herein for purposes of illustration, but that various modifications may be made without deviating from the scope of the invention. Accordingly, the invention is not limited except as by the appended claims.

[0054] Throughout this specification and the claims which follow, unless the context requires otherwise, the word "comprise", and variations such as "comprises" and "comprising", will be understood to imply the inclusion of a stated integer or step or group of integers or steps but not the exclusion of any other integer or step or group of integers or steps.

[0055] The reference in this specification to any prior publication (or information derived from it), or to any matter which is known, is not, and should not be taken as an acknowledgment or admission or any form of suggestion that that prior publication (or information derived from it) or known matter forms part of the common general knowledge in the field of endeavour to which this specification relates.

THE CLAIMS DEFINING THE INVENTION ARE AS FOLLOWS:

1. A computer-implemented method to handle a request from a management tool to access application management data from distributed application instances, the method comprising:

5 receiving from an application management tool a management tool request to access management data related to an application running instances in one or more datacenters, wherein receiving the management tool request comprises receiving a request from a performance monitoring tool to access status information describing operation of one or more application instances;

10 identifying one or more types of management data that satisfy the received request; determining a distribution of the application that includes two or more instances of the application;

gathering management data to satisfy the request from each distributed application instance;

15 unifying gathered management data to provide a uniform response to the received management tool request; and

reporting the gathered and unified management data in response to the received management tool request,

wherein the preceding steps are performed by at least one processor.

20 2. The method of claim 1 wherein receiving the management request comprises receiving the management tool request through a programmatic application programming interface (API) exposed to tools for requesting management data.

3. The method of claim 2 wherein the API comprises a uniform interface for accessing management data without requiring tools to understand where or at how many  
25 locations application instances are running.

4. The method of claim 1 wherein identifying one or more types of management data comprises determining that the management tool request asks for information that is produced by each instance of the application.

5. The method of claim 1 wherein identifying one or more types of management data  
30 comprises determining which information to gather from each application instance and whether the one or more types of management data is already gathered locally from data pushed to a central location by each application instance.

6. The method of claim 1 wherein determining the distribution of the application comprises determining where the application is running and where the system will find management data to satisfy the request.

7. The method of claim 1 wherein gathering management data comprises accessing at least one instance in a private datacenter and at least one instance in a cloud computing facility.

8. The method of claim 1 wherein gathering management data comprises contacting each instance of the application to satisfy the received management tool request.

9. The method of claim 1 further comprising sending one or more troubleshooting commands to one or more remote application instances, wherein the remote application instances carry out the troubleshooting commands and report requested data back to a central location where the management tool can access information associated with the reported requested data.

10. The method of claim 1 wherein unifying gathered data comprises formatting data so that management tools need not be written to include an understanding of the various potential distributions of applications managed by the management tools.

11. The method of claim 1 wherein unifying gathered data comprises formatting data so that the computer system can freely migrate the application from location to location or to multiple locations as needed to handle application loads, while still providing administrators with a consistent management and troubleshooting experience.

12. The method of claim 1 wherein reporting the data comprises sending the data to the management tool through an interface on which the management tool request was received.

13. A computer system for cross-cloud computing for cross-cloud management and troubleshooting, the computer system comprising:

a processor and memory configured to execute software instructions embodied within the following components;

a location management component that manages information about multiple datacenter locations at which instances of an application are running;

a location data store that stores information that describes locations at which instances of the application are running;

a tool interface component that provides an interface to the computer system through which one or more tools can access management and troubleshooting information for the application;

one or more management tools that connect to the tool interface component to access management information or perform application troubleshooting;

a data migration component that migrates management information at one or more remote locations where the application is running back to a home location of the

application; and

a troubleshooting component that performs one or more troubleshooting tasks on the application at one or more locations.

14. The system of claim 13 wherein the location management component includes information describing how to reach each location, connections available for retrieving management information, credentials for user accounts to use on each connection, and one or more application and datacenter components from which to gather troubleshooting information and send troubleshooting commands.

15. The system of claim 13 wherein the location management component receives information describing dynamic migration of application loads from one datacenter or cloud to another and updates the management information so that the location management component has a complete picture of all of the datacenter locations where the application is running.

16. The system of claim 13 wherein the tool interface component provides an initial connection point for tools to access information related to the application at a cloud-computing appliance located within an enterprise's private datacenter.

17. The system of claim 13 wherein at least one management tool is designed to work with a local application, and the system provides the tool with information describing a distributed application running at multiple locations without the at least one management tool's knowledge.

18. The system of claim 13 further comprising a billing component that reports billing information related to one or more locations where the application is running.

19. A computer-readable storage device comprising instructions for controlling a computer system to report data back from and handle troubleshooting requests at a location of a remote application instance, wherein the instructions, upon execution, cause a processor to perform actions comprising:

receiving management data at a remote application instance handling a portion of load generated by requests from clients of the application;

determining a home location of the application where an administrator can access management data reported by multiple instances of the application running at one or more

distributed remote locations, the home location being a location where the application operates under steady conditions;

    sending the received management data from the remote application instance to the determined home location of the application;


5       receiving a troubleshooting request from a management tool run at the home location to troubleshoot the remote application instance;

    performing one or more troubleshooting actions in response to the received troubleshooting request; and

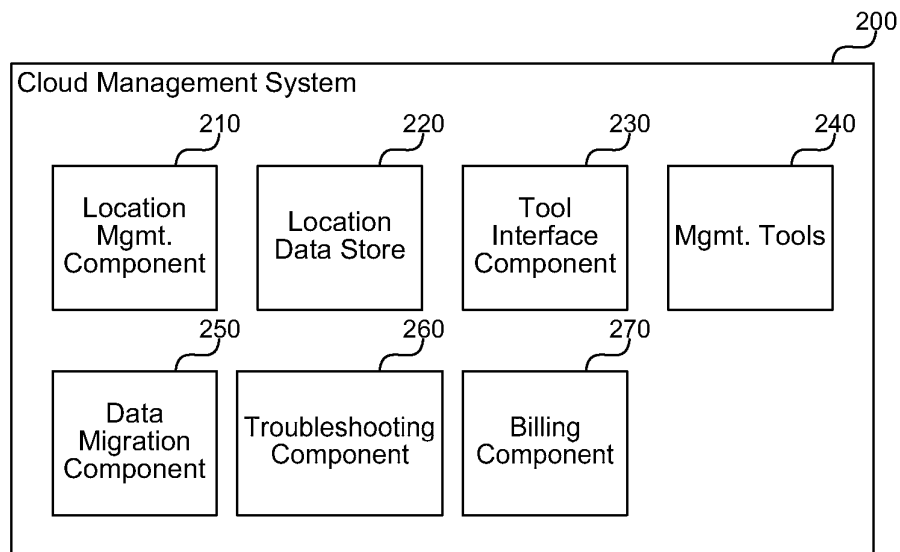
10       sending a troubleshooting result to the home location in response to the received troubleshooting request.




1/4

**FIG. 1**

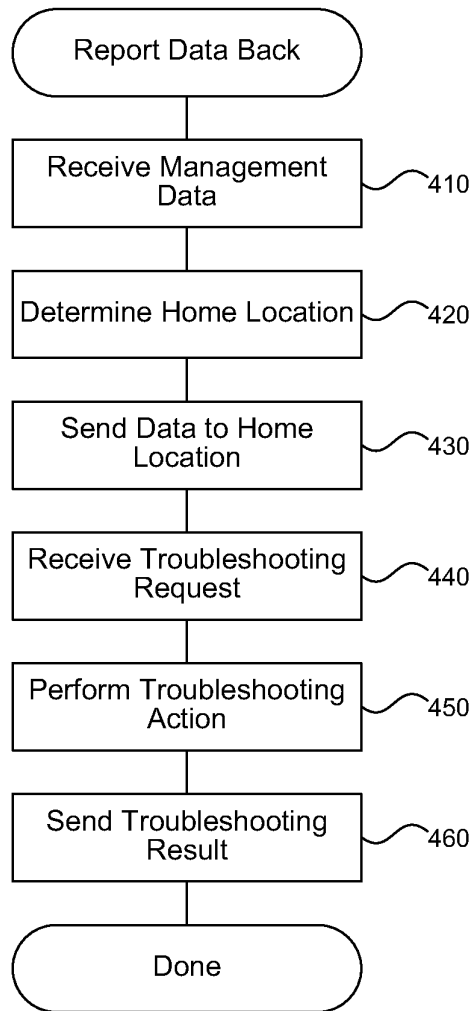
2/4

**FIG. 2**

3/4

**FIG. 3**

4/4

**FIG. 4**