



(19) **United States**

(12) **Patent Application Publication**
Jackson

(10) **Pub. No.: US 2015/0229586 A1**

(43) **Pub. Date: Aug. 13, 2015**

(54) **SYSTEM AND METHOD OF BROKERING CLOUD COMPUTING RESOURCES**

G06Q 20/14 (2006.01)
G06Q 30/02 (2006.01)
G06Q 30/04 (2006.01)

(71) Applicant: **Adaptive Computing Enterprises, Inc.**,
Provo, UT (US)

(52) **U.S. Cl.**
CPC *H04L 47/827* (2013.01); *G06Q 30/0269*
(2013.01); *G06Q 30/04* (2013.01); *G06Q 20/14*
(2013.01); *H04L 41/5006* (2013.01)

(72) Inventor: **David B. Jackson**, Spanish Fork, UT
(US)

(21) Appl. No.: **14/691,120**

(22) Filed: **Apr. 20, 2015**

Related U.S. Application Data

(63) Continuation of application No. 13/418,777, filed on Mar. 13, 2012, now Pat. No. 9,015,324, which is a continuation-in-part of application No. 12/752,622, filed on Apr. 1, 2010, now Pat. No. 8,370,495, which is a continuation of application No. 11/276,856, filed on Mar. 16, 2006, now Pat. No. 7,698,430.

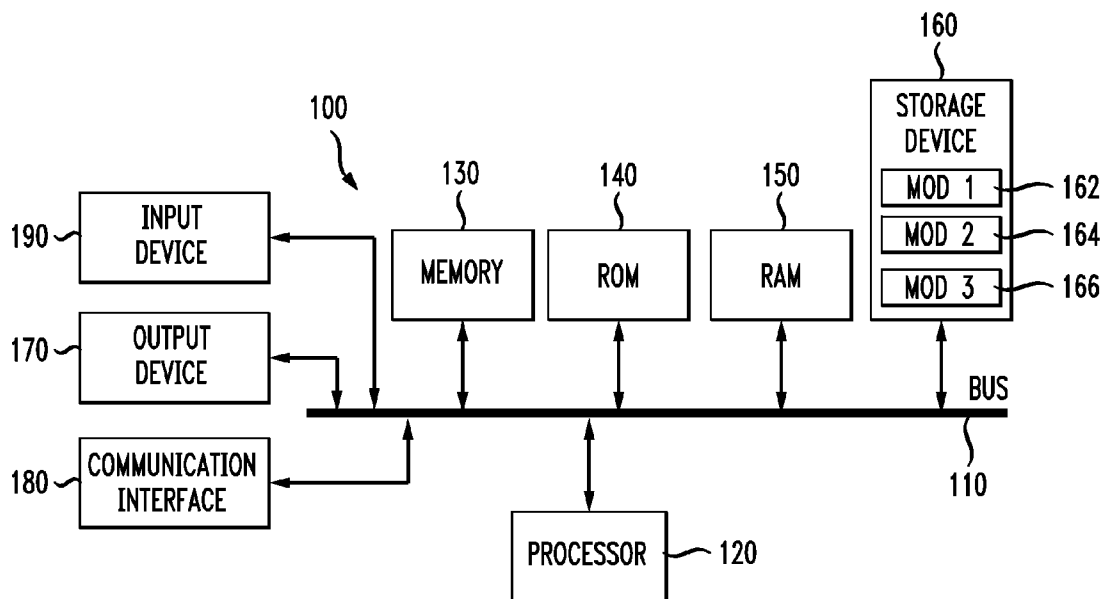
(60) Provisional application No. 60/662,240, filed on Mar. 16, 2005.

Publication Classification

(51) **Int. Cl.**
H04L 12/911 (2006.01)
H04L 12/24 (2006.01)

(57) **ABSTRACT**

System, method, and tangible computer-readable storage media are disclosed for providing a brokering service for compute resources. The method includes, at a brokering service, polling a group of separately administered compute environments to identify resource capabilities and information, each compute resource environment including the group of managed nodes for processing workload, receiving a request for compute resources at the brokering service system, the request for compute resources being associated with a service level agreement (SLA) and based on the resource capabilities across the group of compute resource environments, selecting compute resources in one or more of the group of compute resource environments. The brokering service system receives workload associated with the request and communicates the workload to the selected resources for processing. The brokering services system can aggregate resources for multiple cloud service providers and act as an advocate for or a guarantor of the SLA associated with the workload.



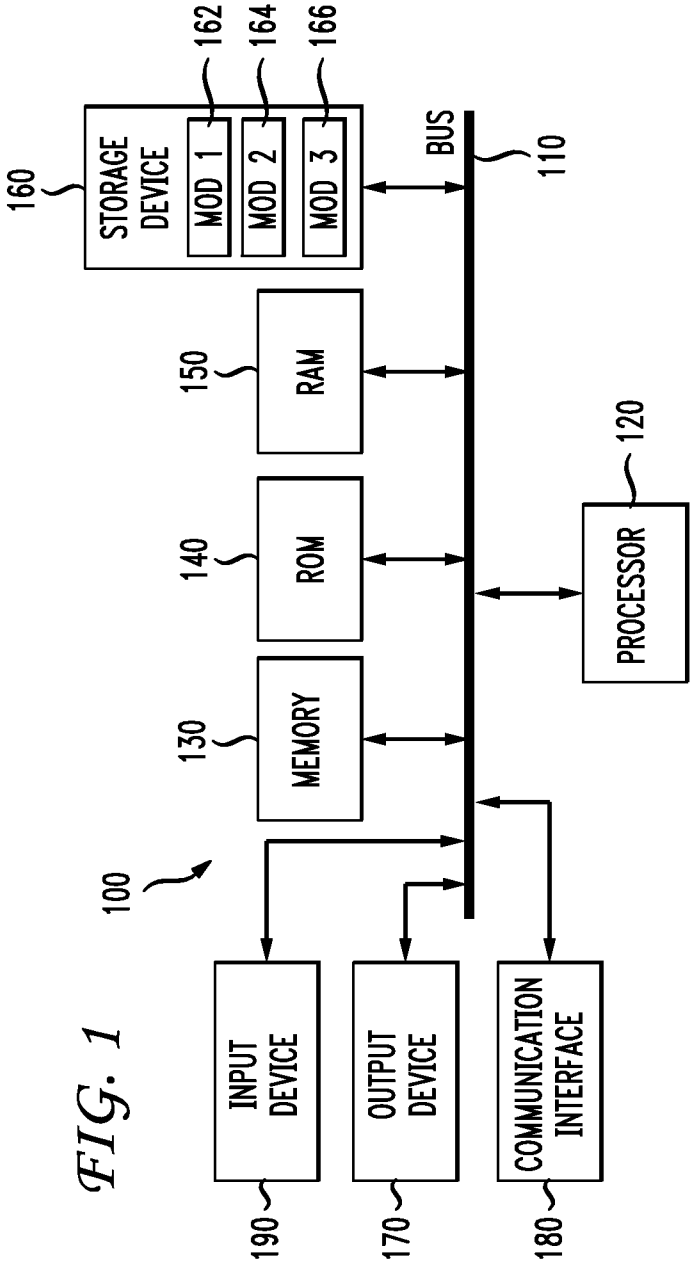
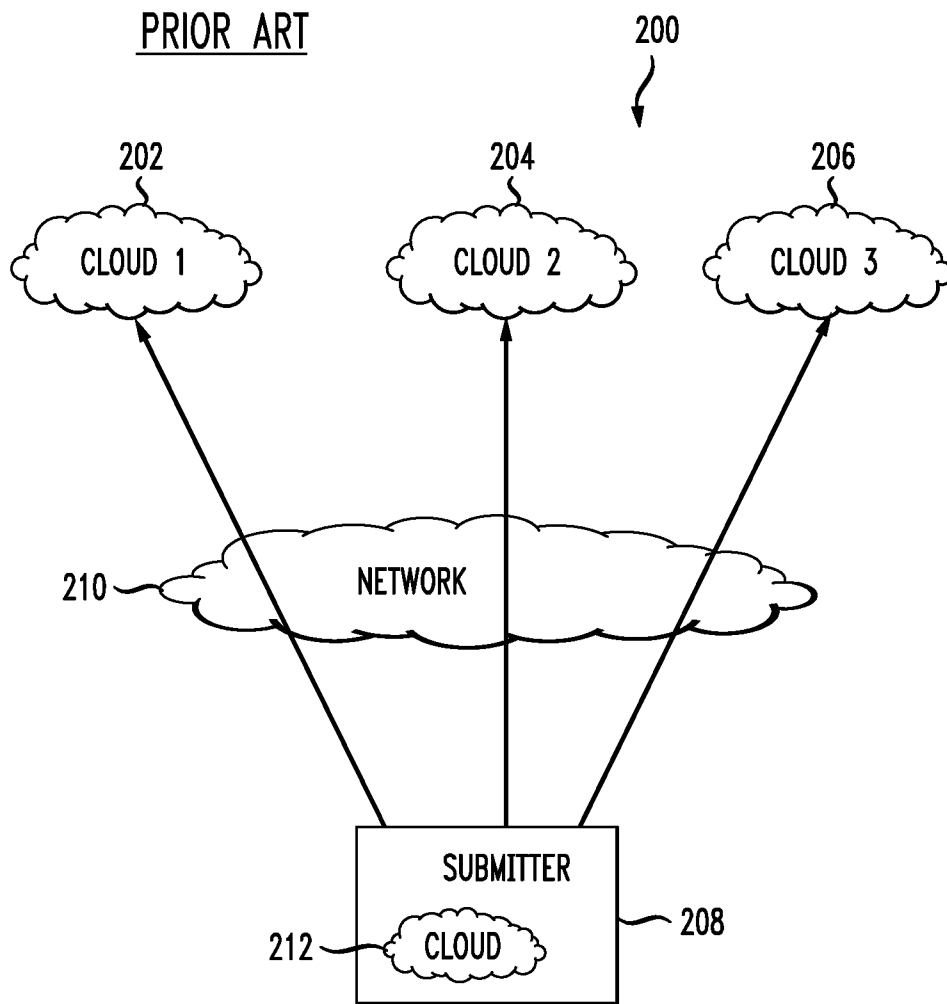


FIG. 1

FIG. 2

PRIOR ART



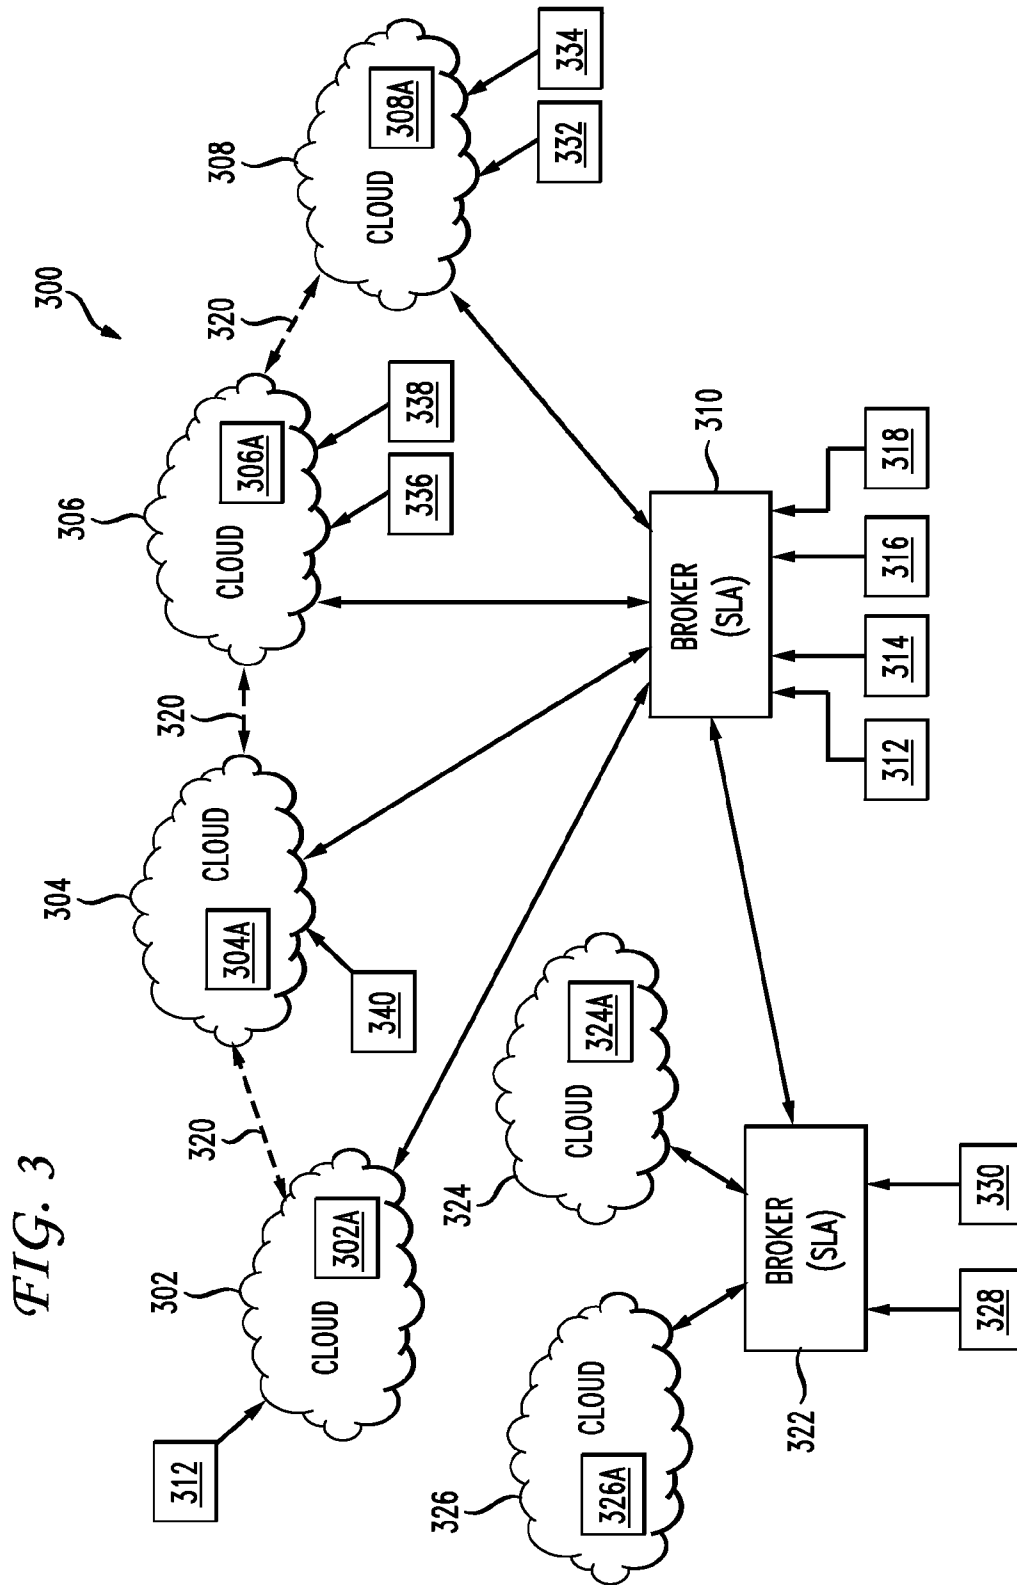


FIG. 3

FIG. 4

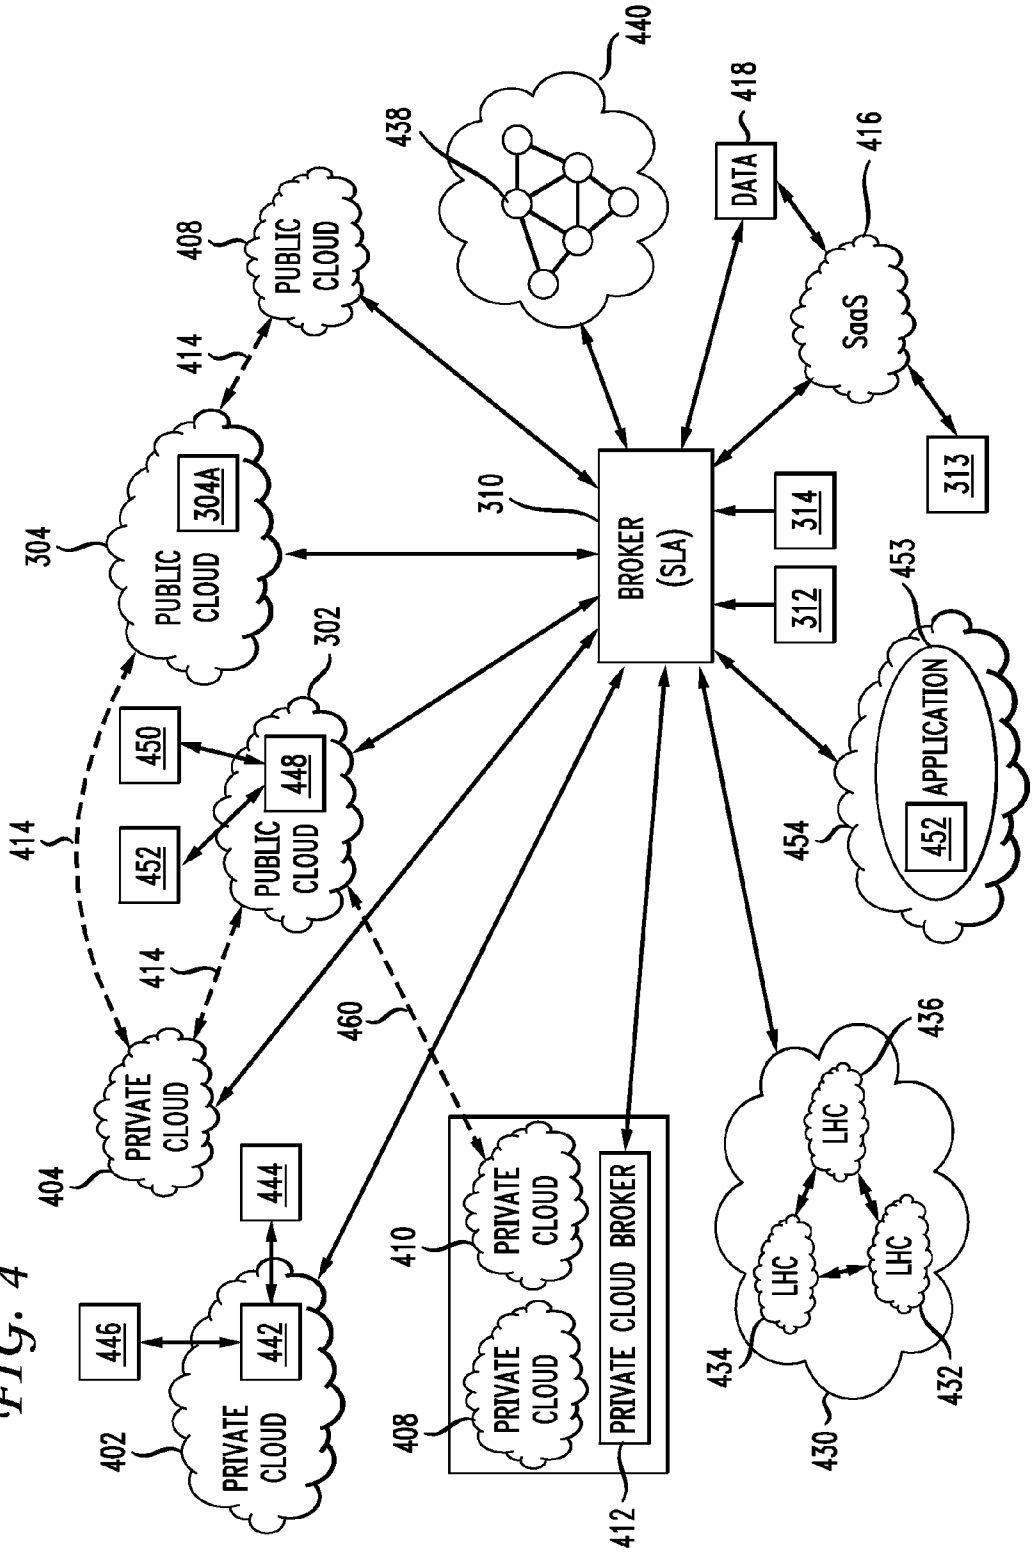


FIG. 5

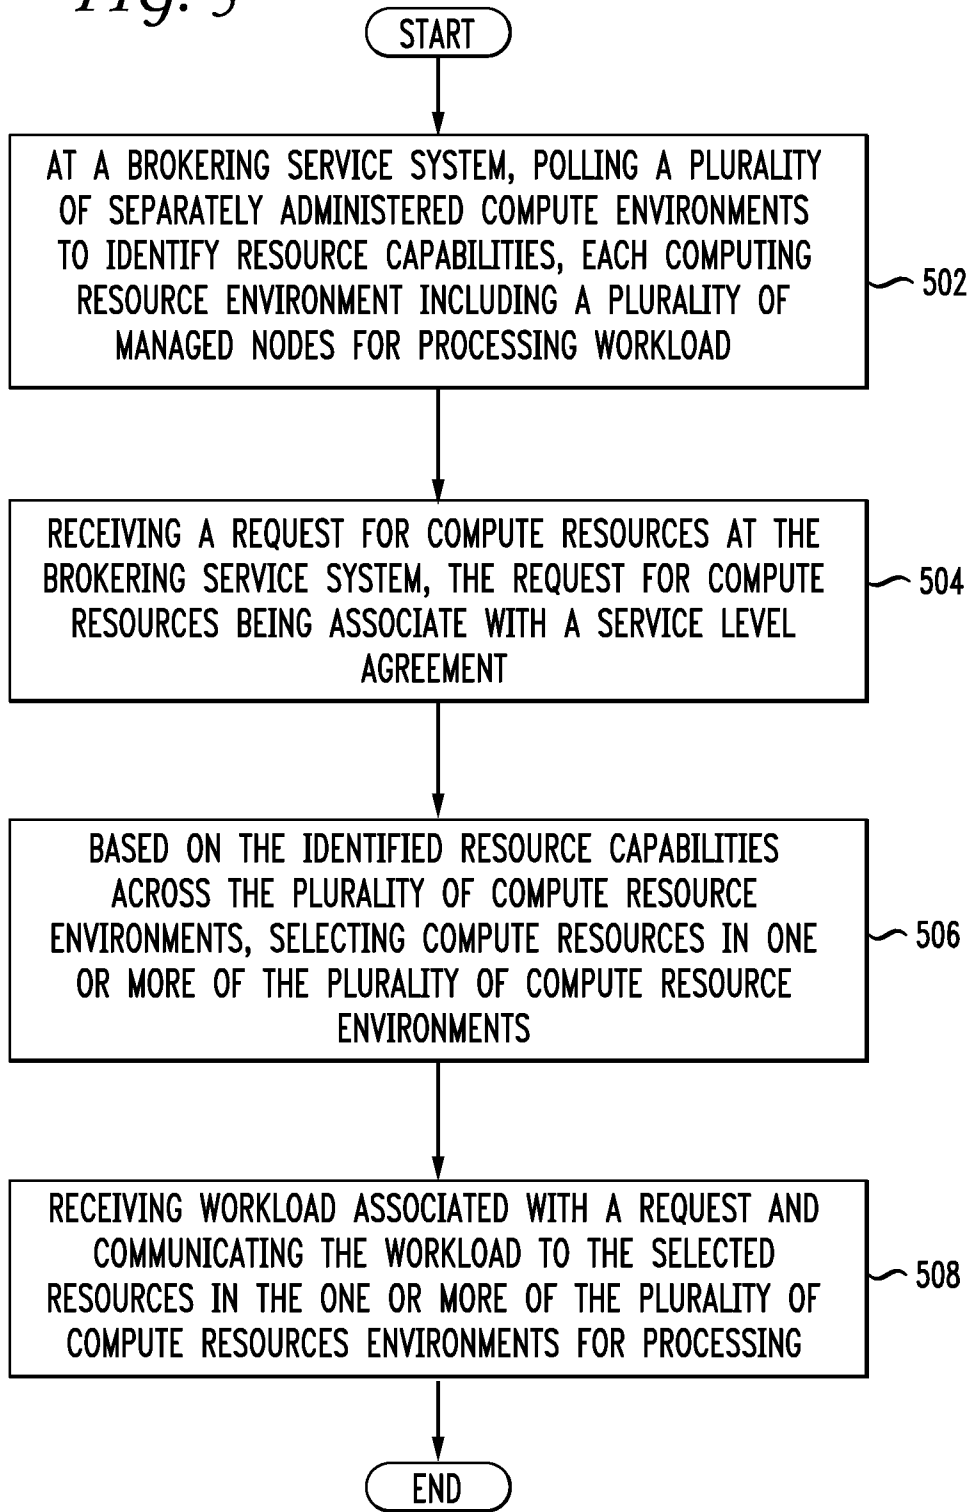


FIG. 6

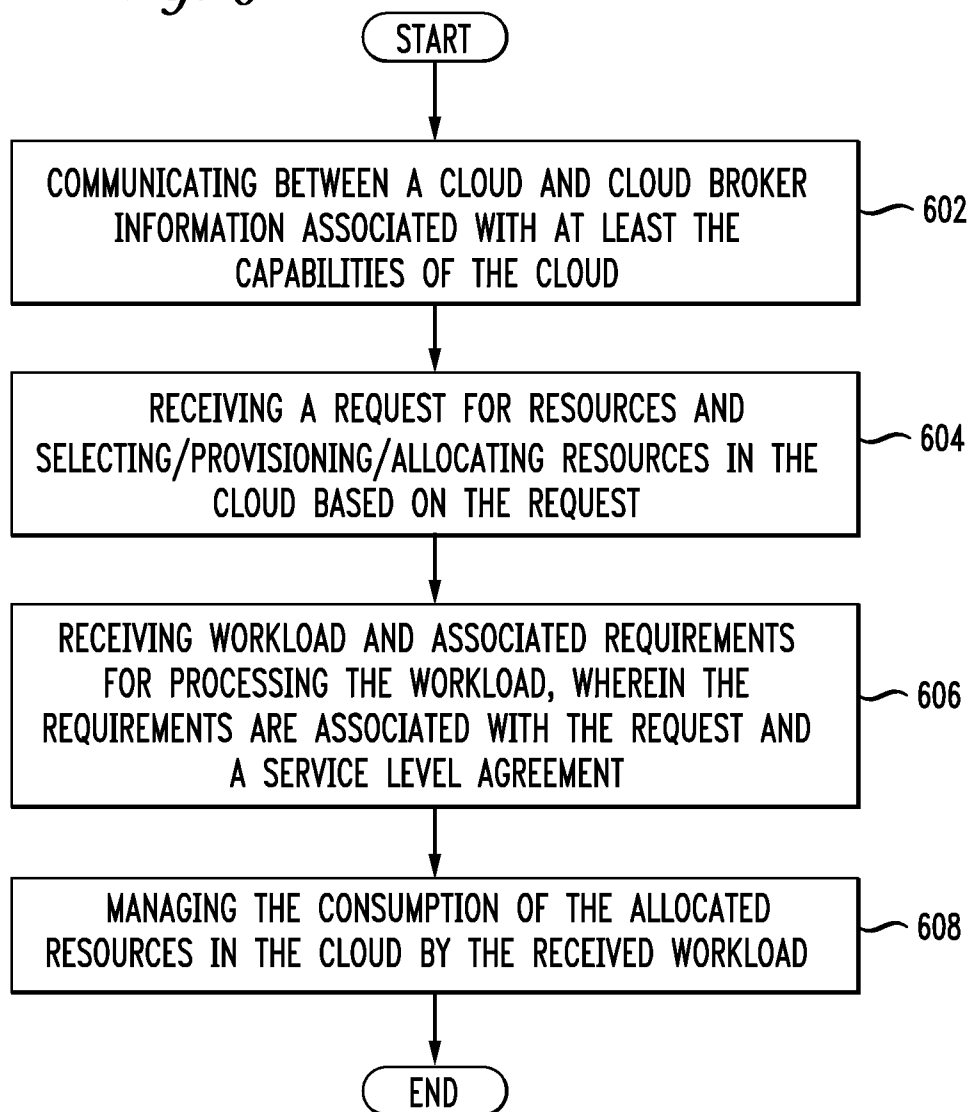


FIG. 7

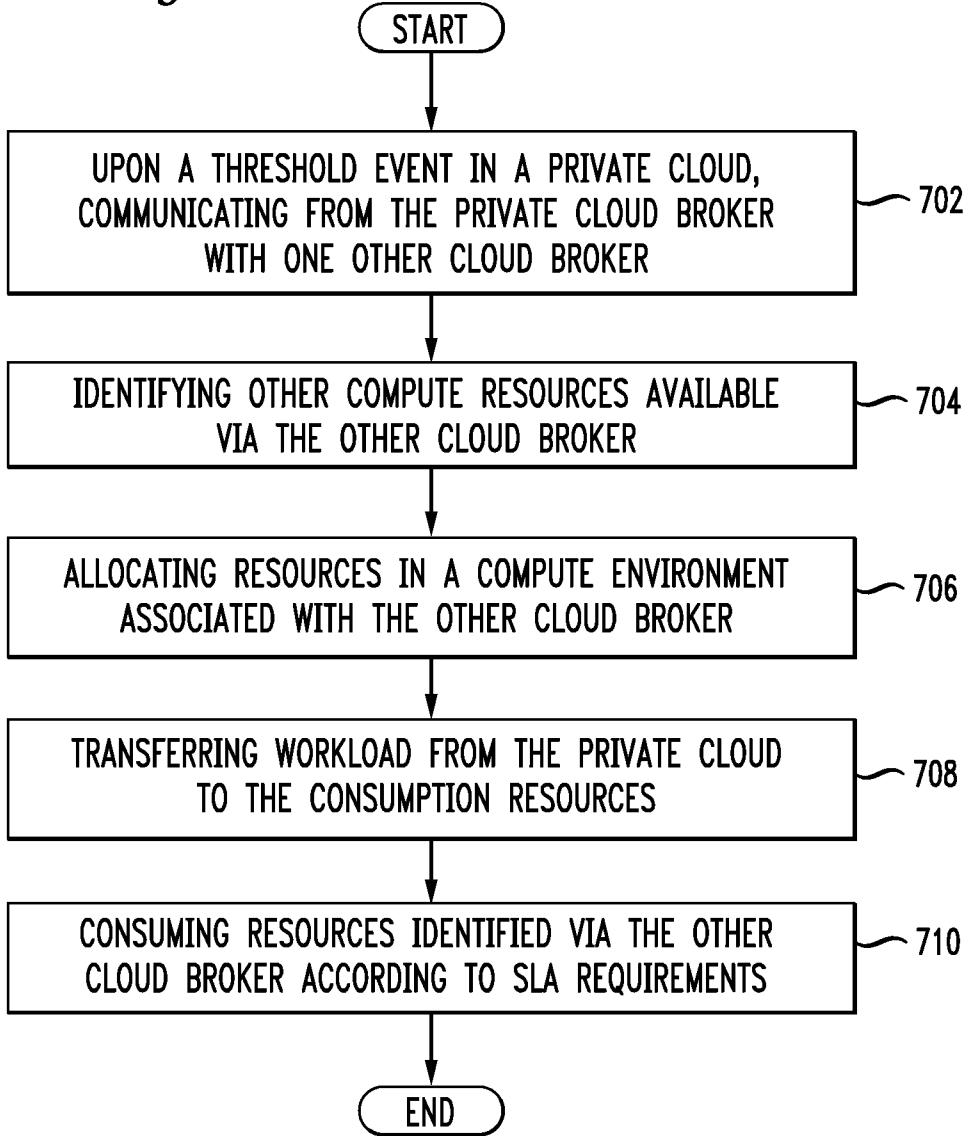


FIG. 8

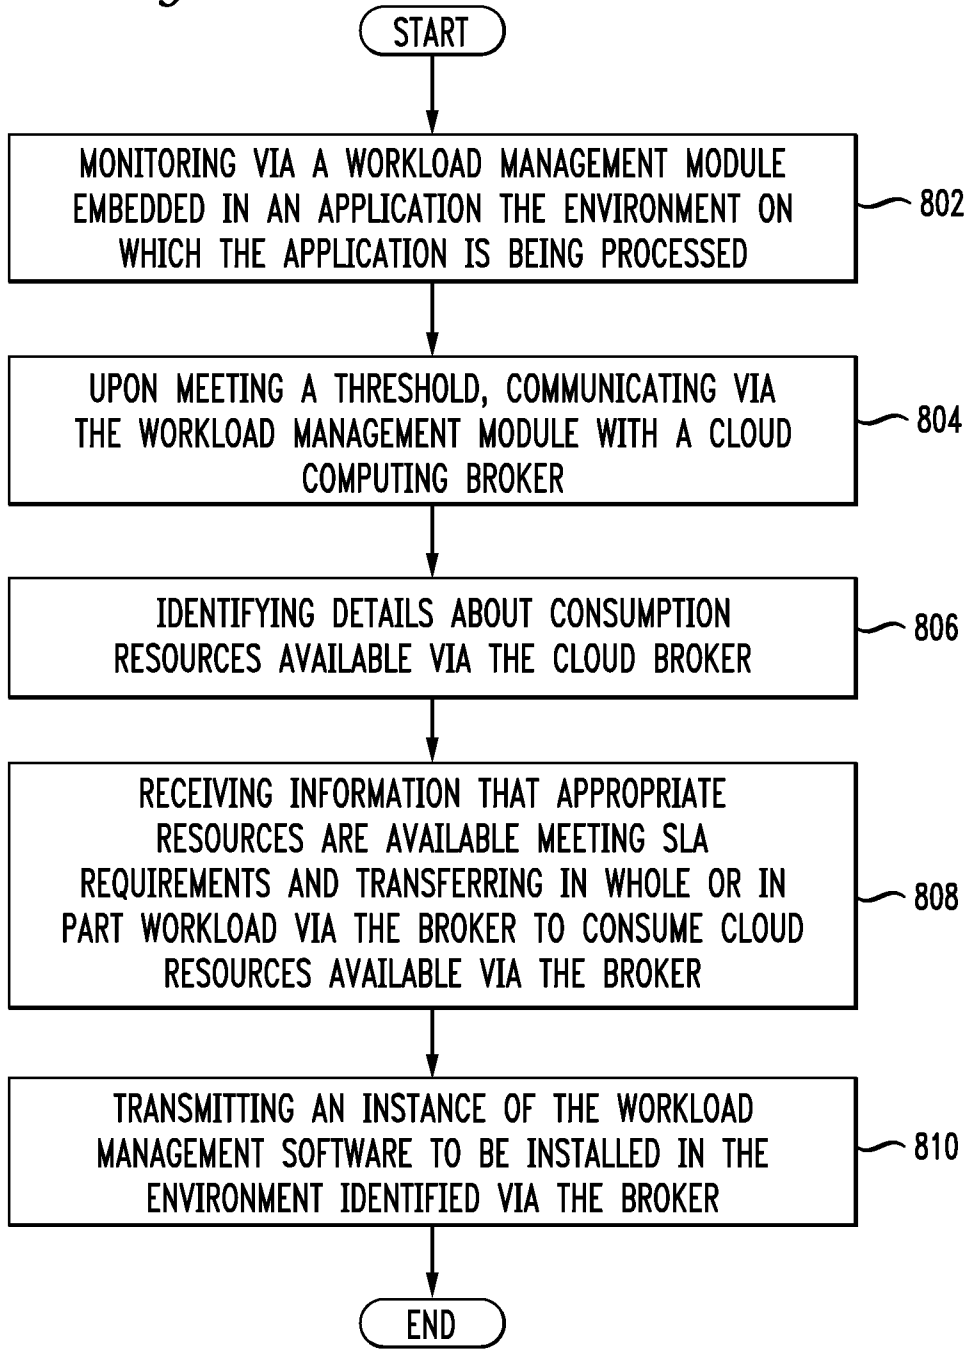
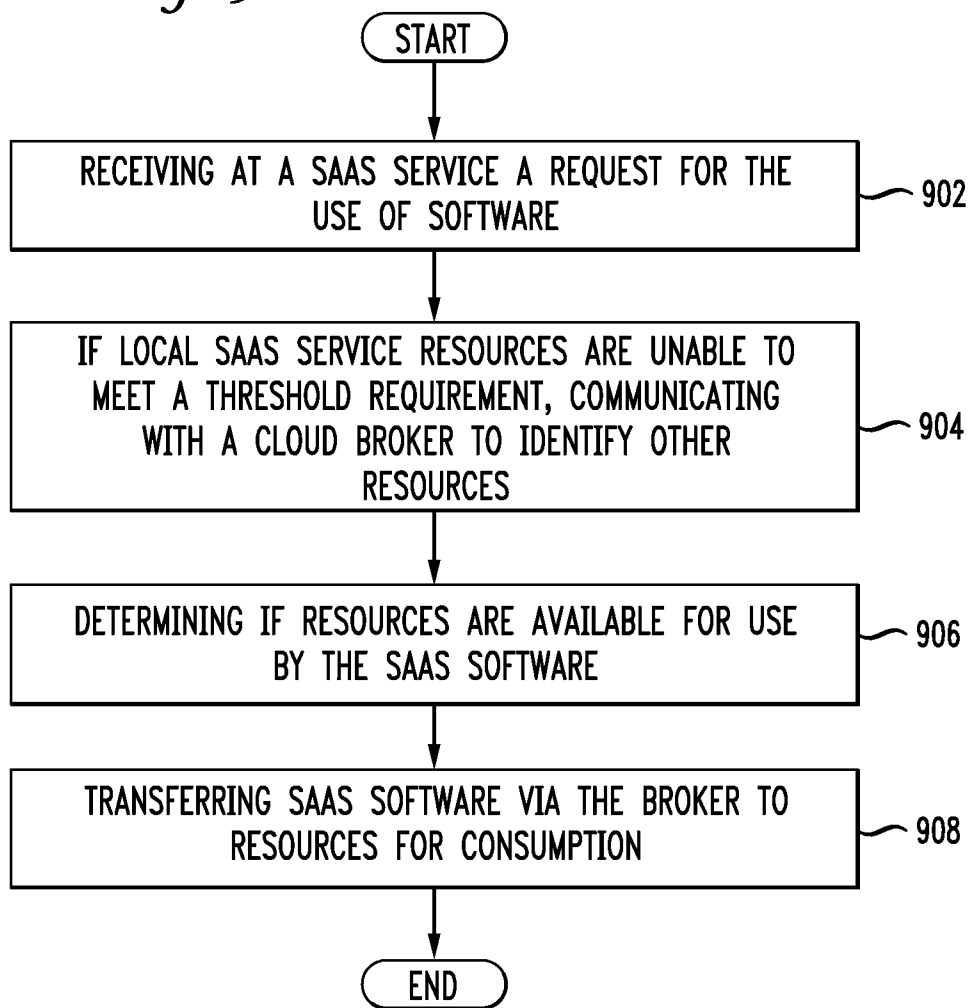


FIG. 9



SYSTEM AND METHOD OF BROKERING CLOUD COMPUTING RESOURCES

PRIORITY INFORMATION

[0001] The present application is a continuation of U.S. patent application Ser. No. 13/418,777, filed Mar. 13, 2012, which is a continuation-in-part of U.S. patent application Ser. No. 12/752,622, filed Apr. 1, 2010, which is a continuation of U.S. patent application Ser. No. 11/276,856, filed Mar. 16, 2006, now U.S. Pat. No. 7,698,430, which claims the benefit of U.S. Provisional Patent Application No. 60/662,240, filed Mar. 16, 2005 and the present application is also a continuation of U.S. patent application Ser. No. 12/503,424, filed Jul. 15, 2009. The contents of each of the above-identified applications are incorporated herein by reference in their entirety.

BACKGROUND

[0002] 1. Technical Field

[0003] The present disclosure relates to managing computer resources and more specifically to brokering access to resources in multiple compute resource environments, such as multiple private and/or public clouds.

[0004] 2. Introduction

[0005] Cloud computing is described as a mechanism of providing scalable and virtualized compute resources over the Internet to the public or internally for a company, university, or other organization. Typically a computing cloud is made up of a group of individual computers connected by high-speed communications and some form of management software guiding the collective actions and usage of the individual computers. However, other cloud configurations exist. For example, FIG. 2 illustrates an example environment **200** in which a submitter **208** can choose from a number of clouds **202**, **204**, **206** to submit workload for processing. Currently, a number of vendors provide cloud computing services, including Google™, Amazon®, and Yahoo®. FIG. 2 shows cloud **1202**, cloud **2204** and cloud **3206**. Typically these clouds are accessed via the Internet **210**.

[0006] The infrastructure in a cloud is a data center with numerous servers typically with different levels of virtualization technologies. Access to each cloud is presented as a single point for the computing needs of customers. Often service level agreements (SLAs) with each cloud will exist which promise or guarantee a particular quality of service which can meet the requirements of a submitter **208**.

[0007] Many advantages exist for cloud computing, including the ability of companies to avoid capital expenditures by only paying for consumed resources that they use. Some challenges also exist with the current state of cloud computing. One issue is with individual cloud vendors. Users may be limited to the service levels and applications or service providers that the cloud vendor **202**, **204**, **206** is willing to offer. Users may have limited ability to install applications and to perform certain tasks. Only particular operating environments may be available which do not match or have an affinity to a user's workload. In some cases, experts have argued that cloud computing is regression to the mechanism of mainframe computing prior to the 1970's.

[0008] Another feature of cloud computing is also shown in FIG. 2. In this example, assume that each submitter has its own cloud **212**. This could be termed a private cloud which is an internal on-demand center for an entity such as a large company. Clouds **202**, **204**, **206** may be public clouds and

companies such as IBM provide integration services to enable submitters **208** with private clouds **212** utilize a hybrid environment in which overflow workload transfers from a private cloud **212** to one of the public clouds **202**, **204**, **206**. In this manner, a hybrid cloud can span both the corporate data centers and public services. However, a submitter **208** must identify and submit work to each cloud individually. This generally requires some kind of business relationship with one or more clouds, which can be cumbersome, expensive and administratively difficult for the owner of private cloud **210**. The submitter **210** may use different cloud computing environments. The capability for this hybrid environment can be enabled by a combination of servers such as IBM's Blade Center HS22 Intel-based Blade Servers, the IBM Tivoli Service Management, monitoring and provisioning tools as well as VMware Virtualized Images and other software that can support specific workload types. Such hybrid environments, while enabling companies to run normal workloads in a private data center **212** and transmit overflow data and applications to public clouds, still suffer from some of the same issues set forth above.

SUMMARY

[0009] Additional features and advantages of the disclosure will be set forth in the description which follows, and in part will be obvious from the description, or may be learned by practice of the principles disclosed herein. The features and advantages of the disclosure may be realized and obtained by means of the instruments and combinations particularly pointed out in the appended claims. These and other features of the present disclosure will become more fully apparent from the following description and appended claims, or may be learned by the practice of the principles set forth herein.

[0010] Disclosed are systems, methods and tangible computer-readable storage media which provide improved and simplified access to cloud computing services. The method is computer-implemented and is preferably performed via one or more processors in a computing device or system. The method includes, at a brokering service system, periodically polling a group of separately administered compute environments to identify resource capabilities and/or other data associated with the environment such as availability, cost, reliability, etc., each compute resource environment including a group of managed nodes for processing workload. Preferably, each compute resource environment registers with the brokering service system. The brokering service aggregates resources for multiple cloud service environments and presents a single united interface to users. The brokering service system can preferably be a server or servers with the necessary communication mechanisms known in the art to communicate between outside users and more than one separately administered compute environments.

[0011] An example of separately administered compute environments can be individual public clouds provided by a company such as IBM®, Amazon®, Google™, Microsoft®, Yahoo® and so forth. As noted above, an example public cloud can be a combination of IBM's Blade Center HS22 Intel-based Blade Servers, IBM Tivoli Service Management, monitoring and provisioning tools, VMware Virtualized Imaging Software and other software designed to support specific workload types. Another example workload management software is the Moab workload manager and other management products from Cluster Resources, Inc. In one respect, such clouds have been termed a "public cloud" or

“external cloud”. Other clouds may include a “private cloud” or an “internal cloud” which may be managed and only available internally in the government, a corporation, or other organization.

[0012] The method further includes receiving a request for compute services at the brokering system, the request for compute resources being associated with a service level agreement and, based on the identified resource information across the group of compute resource environments, selecting compute resources in one or more of compute resource environments. The method next can include receiving workload associated with the request and communicating the workload to the selected resources in the group of compute resource environments for processing. The selection of compute resources ensures that the processing complies with the service level agreement.

[0013] This disclosure provides a mechanism through the use of the brokering service system of providing a standardized portal or environment for submitting workload to one of a number of compute resource environments or clouds. This single interface to various clouds can improve the marketability and availability of cloud computing and greatly simplify the process for submitters whether they be individuals or companies. The brokering service system not only is an advocate for ensuring that service level agreements are complied with by the selected resources of the group of clouds, but can also provide the necessary translations in order to accommodate and match the workload with the selected resources in one or more of the group of clouds. In other words, the brokering service system can translate or modify the workload because it was not programmed or configured for the selected environment. In one example, assume that the workload has an affinity for being processed on a Microsoft operating system. The affinity can be a soft or a hard requirement. An example of a soft requirement is an application that processes data 30% faster in a Microsoft environment. The affinity is desirable, but not necessary. An example of a hard requirement is an application that relies on core libraries that are only available in a Microsoft environment. This affinity is a necessary requirement. However, the brokering service system has selected resources within a cloud that are running at the Linux operating system. The brokering service system can either modify the workload so that its affinity is more suitable to a Linux operating system environment, if possible, or the brokering service system can reprovision the Linux nodes with a Microsoft operating system and then submit the workload into those resources for consumption in an environment that has been modified to better suit the workload itself. The brokering system can also monitor the environments and if nodes with the Microsoft operating system become available it can migrate a reservation of nodes to new nodes which better fit the workload which is yet to consume compute resources. In this manner, the brokering services system provides easy and efficient accessibility to resources within multiple clouds.

BRIEF DESCRIPTION OF THE DRAWINGS

[0014] In order to describe the manner in which the above-recited and other advantages and features of the disclosure can be obtained, a more particular description of the disclosure briefly described above will be rendered by reference to specific embodiments thereof which are illustrated in the appended drawings. Understanding that these drawings depict only exemplary embodiments of the disclosure and are

not therefore to be considered to be limiting of its scope, the principles will be described and explained with additional specificity and detail through the use of the accompanying drawings in which:

- [0015]** FIG. 1 illustrates the standard system;
- [0016]** FIG. 2 illustrates a prior art approach to accessing cloud computing;
- [0017]** FIG. 3 illustrates an example brokering service;
- [0018]** FIG. 4 illustrates another example brokering service;
- [0019]** FIG. 5 illustrates a method of the present disclosure;
- [0020]** FIG. 6 illustrates another method embodiment;
- [0021]** FIG. 7 illustrates another method embodiment;
- [0022]** FIG. 8 illustrates another method embodiment; and
- [0023]** FIG. 9 illustrates yet another method embodiment.

DETAILED DESCRIPTION

[0024] Various embodiments of the disclosure are discussed in detail below. While specific implementations are discussed, it should be understood that this is done for illustration purposes only. A person skilled in the relevant art will recognize that other components and configurations may be used without parting from the spirit and scope of the disclosure.

[0025] With reference to FIG. 1, an exemplary system **100** includes a general-purpose computing device **100**, including a processing unit (CPU or processor) **120** and a system bus **110** that couples various system components including the system memory **130** such as read only memory (ROM) **140** and random access memory (RAM) **150** to the processor **120**. These and other modules can be configured to control the processor **120** to perform various actions. Other system memory **130** may be available for use as well. It can be appreciated that the disclosure may operate on a computing device **100** with more than one processor **120** or on a group or cluster of computing devices networked together to provide greater processing capability. The processor **120** can include any general purpose processor and a hardware module or software module, such as module **162**, module **2 164**, and module **3 166** stored in storage device **160**, configured to control the processor **120** as well as a special-purpose processor where software instructions are incorporated into the actual processor design. The processor **120** may essentially be a completely self-contained computing system, containing multiple cores or processors, a bus, memory controller, cache, etc. A multi-core processor may be symmetric or asymmetric.

[0026] The system bus **110** may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. A basic input/output (BIOS) stored in ROM **140** or the like, may provide the basic routine that helps to transfer information between elements within the computing device **100**, such as during start-up. The computing device **100** further includes storage devices **160** such as a hard disk drive, a magnetic disk drive, an optical disk drive, tape drive or the like. The storage device **160** can include software modules **162**, **164**, **166** for controlling the processor **120**. Other hardware or software modules are contemplated. The storage device **160** is connected to the system bus **110** by a drive interface. The drives and the associated computer readable storage media provide nonvolatile storage of computer readable instructions, data structures, program modules and other data for the computing device **100**. In one aspect, a hardware module that performs a particular function includes

the software component stored in a tangible and/or intangible computer-readable medium in connection with the necessary hardware components, such as the processor **120**, bus **110**, display **170**, and so forth, to carry out the function. The basic components are known to those of skill in the art and appropriate variations are contemplated depending on the type of device, such as whether the device **100** is a small, handheld computing device, a desktop computer, or a computer server.

[0027] Although the exemplary embodiment described herein employs the hard disk **160**, it should be appreciated by those skilled in the art that other types of computer readable media which can store data that are accessible by a computer, such as magnetic cassettes, flash memory cards, digital versatile disks, cartridges, random access memories (RAMs) **150**, read only memory (ROM) **140**, a cable or wireless signal containing a bit stream and the like, may also be used in the exemplary operating environment. Tangible computer-readable storage media expressly exclude media such as energy, carrier signals, electromagnetic waves, and signals per se.

[0028] To enable user interaction with the computing device **100**, an input device **190** represents any number of input mechanisms, such as a microphone for speech, a touch-sensitive screen for gesture or graphical input, keyboard, mouse, motion input, speech and so forth. The input device **190** may be used by the presenter to indicate the beginning of a speech search query. An output device **170** can also be one or more of a number of output mechanisms known to those of skill in the art. In some instances, multimodal systems enable a user to provide multiple types of input to communicate with the computing device **100**. The communications interface **180** generally governs and manages the user input and system output. There is no restriction on operating on any particular hardware arrangement and therefore the basic features here may easily be substituted for improved hardware or firmware arrangements as they are developed.

[0029] For clarity of explanation, the illustrative system embodiment is presented as including individual functional blocks including functional blocks labeled as a “processor” or processor **120**. The functions these blocks represent may be provided through the use of either shared or dedicated hardware, including, but not limited to, hardware capable of executing software and hardware, such as a processor **120**, that is purpose-built to operate as an equivalent to software executing on a general purpose processor. For example the functions of one or more processors presented in FIG. **1** may be provided by a single shared processor or multiple processors. (Use of the term “processor” should not be construed to refer exclusively to hardware capable of executing software.) Illustrative embodiments may include microprocessor and/or digital signal processor (DSP) hardware, read-only memory (ROM) **140** for storing software performing the operations discussed below, and random access memory (RAM) **150** for storing results. Very large scale integration (VLSI) hardware embodiments, as well as custom VLSI circuitry in combination with a general purpose DSP circuit, may also be provided.

[0030] The logical operations of the various embodiments are implemented as: (1) a sequence of computer implemented steps, operations, or procedures running on a programmable circuit within a general use computer, (2) a sequence of computer implemented steps, operations, or procedures running on a specific-use programmable circuit; and/or (3) interconnected machine modules or program engines within the programmable circuits. The system **100** shown in FIG. **1** can

practice all or part of the recited methods, can be a part of the recited systems, and/or can operate according to instructions in the recited tangible computer-readable storage media. Generally speaking, such logical operations can be implemented as modules configured to control the processor **120** to perform particular functions according to the programming of the module. For example, FIG. **1** illustrates three modules Mod1 **162**, Mod2 **164** and Mod3 **166** which are modules configured to control the processor **120**. These modules may be stored on the storage device **160** and loaded into RAM **150** or memory **130** at runtime or may be stored as would be known in the art in other computer-readable memory locations.

[0031] FIG. **2** illustrates an example prior art approach to cloud computing. Cloud **1 202**, cloud **2 204** and cloud **3 206** can represent a mixture of public clouds and/or private cloud. Again, a public cloud provides resources which may be accessible by the public or companies which are provisioned and available on a self-service basis over the Internet or other network **210** via a web application. The public cloud has the ability to enable fine-grained knowledge and control of the use of resources in the cloud by individual users and can track and bill users on that basis.

[0032] A private cloud can provide an entity managing the cloud enhanced data security, corporate governance and reliability. An illustration of a private cloud is shown in FIG. **2** as cloud **212** associated with the submitter and inaccessible to outside compute jobs. As noted above, one example of the use of cloud computing is the ability of a service to simply provide overflow computing services from a company’s private cloud **212** into one or more public clouds **202**, **204**, **206**. The configuration of each of these clouds is generally known to those in skill of the art. For example, a cloud may use IBM’s Blade Center HS22 Intel-Based Blade Servers, a workload management software for management, monitoring and provisioning. Virtualization software from VMware and other software may be implemented in order to manage the workload in any particular cloud. Such virtualization approaches can incur an additional cost that companies must buy, build, pay for and manage in public and/or private clouds. Typically private clouds **212** are used to be highly flexible compute resources for an enterprise which may be used in management for the compute needs of that enterprise. A private cloud **212** may also be made available to partners within a supply chain or other trusted third-parties to gain access to all or part of the private cloud of any individual company. For example, a company can allow third-party access to data stored in a private cloud, but not processing power, or vice versa.

[0033] FIG. **3** is an example environment **300** illustrating the principles of the present disclosure. Broker **310** represents a system or a device such as that shown in FIG. **1** or a group of servers or devices combined together to manage the requirements of and be able to perform the processes disclosed herein. Other hardware configuration hereinafter developed may also be used for the brokering system or device, which configurations are incorporated herein by reference and would be known to those of skill in the art.

[0034] Brokering system **310** receives requests from users **312**, **314**, **316** and **318**. An example of these users may be individual users with processing needs, companies which have jobs which need to consume resources, or any other entity that needs compute resources. Clouds **302**, **304**, **306** and **308** represent primarily public clouds which are available to process jobs according to the user needs but can also

represent in some cases private clouds as well. For the purposes of the discussion in FIG. 3, the clouds illustrated are public clouds unless stated otherwise. However, the clouds shown can be a combination of public and private clouds.

[0035] As noted above, each cloud 302, 304, 306, 308 may have differing characteristics. The broker 310 polls each of the respective clouds in order to identify the various capabilities and characteristics including any type of parameter associated with each respective cloud. An example of the different characteristics may include different resource types, a different cost per unit of a compute resource, the type of energy which is used to power the compute environment, SLA availability and parameters offered, and so forth. One compute environment may be powered by wind power, while another from a coal-fired power plant, or a nuclear power plant. Another cloud 308 may be powered by biofuel. Accordingly, one parameter which can be identified is the cost of the energy which may be required to power and cool the individual cloud.

[0036] Other parameters include the service level agreements which are warranted by the respective clouds, 302, 304, 306, 308. For example, a cloud 302 may provide a five cents per unit cost but only be able to provide a mid-level service level agreement and a low-level reliability for the processing of jobs submitted to that environment. Environment 304 may have a high-level service level agreement that is warranted but its cost may be eight cents per unit. Accordingly, each environment 302, 304, 306 and 308 will typically be managed by a workload manager or some other management software that will determine how resources within each environment are consumed. For example, module 302A, 304A, 306A and 308A generally represent this system which will manage the resources in these various environments. The examples of what are managed of course are compute jobs that are submitted into each cloud such as jobs 342, 340, 338, 332, and 334. Typically third party requesters submit jobs directly to each respective cloud. A job represents any computing or other workload that a consumer resources in the compute environment such as a weather analysis, an accounting task, or database search, etc.

[0037] The ability of the brokering service system to act as a broker to more than one cloud environment is enabled by its periodically polling the group of separately administered compute environments 302, 304, 306 and 308 to identify the resource capabilities and information. In one variation, each separately administered compute environments 302, 304, 306, 308 reports to the brokering service system 310 when its resource capabilities/information change instead of or in combination with the brokering service system 310 polling the environments. Resource capabilities also encompass any parameter associated with the environment such as cost per consumption hour, availability of resources, types of resources, a need for provisioning of particular software or operating systems, cost of associated human labor, and so forth. Another example of this type of provisioning can be illustrated by brokering services system 310 identifying that cloud 302 primarily offers a Linux operating system with an inexpensive rate, but only a mid-level service level agreement guarantee. Compute environment 304 may offer a hybrid Microsoft and Linux operating environment with a high service level agreement guarantee but with a more expensive cost per unit (such as per node per hour or any other unit used to identify cost). Other parameters may be identified as well, such as the processor speed for processors in each server in

the environment, a data through-put speed, hard drive latency or accessibility speed, and so forth. Each compute resource environment typically includes more than one managed node for processing the workload in each environment. In order to communicate with and function with broker 310, each compute resource environment will register with the brokering system. In one aspect, a cloud does not register but does make data available information to the broker 310 for determining whether to send workload.

[0038] The brokering system 310 in advance can develop a relationship with a number of clouds which can be any combination of public and private clouds. As shall be discussed herein, the ability of the broker to identify, aggregate, communicate, and manage compute resources across a number of different clouds can greatly simplify the ability of workload to be processed on compute resources that match SLA requirements for the requestors 312, 314, 316, 318, 322. In other words, the broker 310 provides the most easy and efficient supply chain management between a user who desires compute resources for workload and the consumption of selected resources by that workload. Clouds may also have restrictions on the types of resources that are available, or the types of users (such as those with a certain security level) that can use resources within the respective environment.

[0039] Users 312, 314, 316, 318 can submit workloads to the broker 310 which then identifies suitable public and/or private clouds and submits the job on behalf of the user. Users 312, 314, 316, 318 can query a broker to determine which public and/or private clouds are capable of servicing the workload within the SLA terms. Users submit jobs directly to the appropriate cloud(s) based on information from the broker 310. Therefore, rather than transmitting workload, the broker 310 just passes information about the clouds to the requestor. In the case of direct to cloud submissions, the workload can include a referral code or other identifying information which indicates that the broker 310 directed the user to the particular cloud(s). The broker can offer both of these services simultaneously.

[0040] The brokering services system 310 may utilize software that communicates seamlessly with management software in the various compute environments. For example, if the brokering service system 310 utilizes Moab, the Moab workload manager or any of its various components from Cluster Resources or from Adapted Computing Enterprises, and the workload management module 302A in compute environment 302 also utilizes the same workload management software, there can be a direct ability to identify with a high level of confidence the capabilities and resources available in cloud 302. However, if another cloud 304 utilizes a different type of workload management module 304A, then the necessary translations, estimations or intelligent predictions of the resource capabilities of a particular environment 304 can be performed by the brokering services system 310. Accordingly, there can be a confidence level associated with the knowledge that is received from polling the separately administered compute environments. The brokering system 310 can adapt workload distribution as the confidence level changes and as learning algorithms interact with and record metrics associated with various environments.

[0041] Other information which may be identified includes the quantity, capability, likelihood of failure and so forth of the particular environment. Again, without a match in the workload management software between the broker and any respective cloud, or if a cloud is not registered with the broker,

an effort will be made from the broker to query for the status, cost, statistics and so forth of any cloud it can communicate with. Although it is likely that the amount of information associated with that cloud will be reduced, as is discussed more fully below, learning algorithms are used to gather additional data associated with a respective cloud based on the consumption of workload within that cloud. Accordingly, the broker **310** will improve over time its ability to identify resources for clouds that it cannot gain full information about.

[0042] Individual clouds may or may not register with the broker. Registering with the broker involves providing a sufficient communication link and contractual understanding between the broker and the registered cloud for the brokering of cloud services through broker **310**. Such a registration improves the communication and ease with which end users will be able to access cloud computing services through the broker **310**. But again, the disclosure notes that even for unregistered clouds, some information may be obtained and thus workload can still flow to an unregistered cloud for consuming resources. When a user submits a workload, the user can indicate a preference for registered or unregistered clouds.

[0043] The polling of each of the separately administered compute environments or clouds can occur on a static periodic basis or dynamically by the brokering service system **310**. Any polling basis may be used. For example, it may occur every half hour, daily, or on a per job basis. A large batch job can be submitted every evening at midnight for processing. In preparation for most advantageously ensuring that the optimal compute resources are identified and matched for processing the batch job, this system can automatically schedule to receive an updated polling of all of the separately administered compute environments in order to have a current snapshot of the resource capabilities across the different environments. The service level agreement can require polling at a certain minimum interval.

[0044] Having the identified resource capabilities of the various clouds **302**, **304**, **306**, **308**, the brokering service system can receive a request for compute resources at a first time. The request is preferably associated with a service level agreement. In other words the request or the requestor can identify various parameters associated with the request. Such parameters can include a required cost, a required performance level, a required guarantee of the availability of resources, an amount of resources, and so on. For example, the requestor can set forth that they desire the cheapest cost with a minimum level of performance. Based on the identified resource capabilities across the compute resource environments, the brokering service system **310** selects compute resources in one or more of the compute resource environments **302**, **304**, **306**, **308**. The selection may involve identifying all of the resource being in one environment such as environment **302**. In some cases, a job with a workload associated request may be both computationally intensive as well as data intensive. In this case, the broker **310** can identify some resources from one environment **302** and perhaps other resources from another environment **306** and actually split the job up amongst more than one compute environment. The broker servicing system **310** can also instruct or communicate with workload managers in the respective environments to ensure that those resources are reserved or scheduled for processing workload at a second time which is later than the first time as disclosed herein and via the incorporated appli-

cations below. In this regard, the broker **310** becomes the enforcer of SLA requirements from the requestor associated with workload.

[0045] The process of managing the selection and reservation and actual consumption of resources may use many of the principles in the following patent applications: U.S. patent application Ser. No. 10/530,583, filed Apr. 7, 2005; U.S. patent application Ser. No. 11/751,899, filed May 22, 2007, both disclosing providing advanced reservations in a compute environment; U.S. patent application Ser. No. 10/530,582, filed Aug. 11, 2006 disclosing co-allocating a reservation spanning different compute resource types; U.S. patent application Ser. No. 10/530,581, filed Aug. 11, 2006 disclosing self-optimizing reservation in time of compute resources; U.S. patent application Ser. No. 10/530,577, filed Mar. 11, 2005 disclosing providing a self-optimizing reservation in space of compute resources; U.S. patent application Ser. No. 11/208,138, filed Aug. 19, 2005 disclosing providing dynamic roll-back reservations in time; U.S. patent application Ser. No. 11/629,940, filed Dec. 18, 2006 disclosing providing reservation masks within a compute environment; U.S. patent application Ser. No. 11/268,857, filed Nov. 8, 2005, now U.S. Pat. No. 7,356,770; U.S. patent application Ser. No. 12/033,386, filed Feb. 19, 2008 both disclosing graphically managing and monitoring a compute environment; U.S. patent application Ser. No. 11/155,090, filed Jun. 17, 2005 disclosing using transaction IDs for managing reservations of compute resources within a compute environment; U.S. patent application Ser. No. 11/155,347, filed Jun. 17, 2005 disclosing providing threshold-based access to compute resources; U.S. patent application Ser. No. 10/530,576, filed Mar. 11, 2005 disclosing providing multi-resource management support in a compute environment; U.S. patent application Ser. No. 11/718,867, filed May 8, 2007 disclosing providing system jobs within a compute environment; U.S. patent application Ser. No. 11/155,091, filed Jun. 17, 2005 disclosing providing dynamic provisioning within a compute environment; U.S. patent application Ser. No. 10/589,339, filed Aug. 11, 2006, now U.S. Pat. No. 7,490,325 disclosing providing intelligent pre-staging of data in a compute environment; U.S. patent application Ser. No. 11/276,852, filed Mar. 16, 2006 disclosing providing a virtual private cluster; U.S. patent application Ser. No. 10/530,578, filed Mar. 11, 2005 disclosing providing object triggers; U.S. patent application Ser. No. 10/530,580, filed Apr. 7, 2005 disclosing providing object messages in a compute environment; U.S. patent application Ser. No. 10/530,575, filed Feb. 4, 2008 disclosing enforcing future policies in a compute environment; U.S. patent application Ser. No. 11/207,438, filed Aug. 19, 2005 disclosing interfacing a workload manager and scheduler with an identity manager; U.S. patent application Ser. No. 11/276,013, filed Feb. 9, 2006 disclosing providing a fixed time offset based on a dedicated co-allocation of a common resource set; U.S. patent application Ser. No. 11/276,853, filed Mar. 16, 2006 disclosing automatic workload transfer to an on-demand center; U.S. patent application Ser. No. 11/276,854, filed Mar. 16, 2006 disclosing simple integration of an on-demand compute environment; U.S. patent application Ser. No. 11/276,855, filed Mar. 16, 2006 disclosing reserving resources in an on-demand compute environment; U.S. patent application Ser. No. 11/276,856, filed Mar. 16, 2006 disclosing an on-demand compute environment; U.S. patent application Ser. No. 11/279,007, filed Apr. 7, 2006 disclosing on-demand access to compute

resources; U.S. patent application Ser. No. 11/763,010, filed Jun. 14, 2007 disclosing optimized multi-component co-allocation scheduling with advanced reservations for data transfers and distributed jobs; U.S. patent application Ser. No. 11/616,156, filed Dec. 26, 2006 disclosing co-allocating a reservation spanning different compute resources types; U.S. patent application Ser. No. 12/023,722, filed Jan. 31, 2008 disclosing managing a hybrid compute environment; U.S. patent application Ser. No. 12/179,142, filed Jul. 24, 2008 disclosing managing energy consumption in a compute environment; U.S. patent application Ser. No. 12/245,276, filed Oct. 3, 2008 disclosing dynamically managing data-centric searches. Each of these patent applications is incorporated herein by reference.

[0046] The various principles incorporated in by reference above provide support to one of skill in the art in order to have necessary background information for brokering of cloud computing services. While any particular algorithm is not necessary for enabling the principles disclosed herein, the information will be helpful to one of skill in the art. Such principles include various methods for managing advanced reservations in a compute environment, collocating a reservation, spanning different compute resource types, self-optimizing reservations in time and or space, providing dynamic rollback reservations in time, providing reservation masks within a compute environment, providing transaction IDs for managing reservations in compute resources, providing threshold-based access to compute resources, providing multi-resource management support, providing system jobs, providing dynamic provisioning, providing intelligent pre-staging of data, providing a virtual private cluster, providing object triggers, providing object messages, enforcing future policies, interfacing a workload manager and scheduler with an identity manager, providing fixed-time offset-based dedicated co-allocation of common resource sets, workload transfer to an on-demand center, simple integration of an on-demand compute environment, reserving resources in an on-demand compute environment, on-demand access to compute resources, optimizing multi-component co-allocation scheduling with advanced reservations for data transfers and distributed jobs, co-allocating a reservation spanning different compute resource types, managing a hybrid compute environment such as having multiple operating systems that may be provisioned and repositioned according to workload need and managing energy consumption in the compute environment.

[0047] After the broker **310** utilizes its relationships with the various clouds **302**, **304**, **306**, **308** and gathers the necessary data from each of the clouds, as it receives requests for processing workload, it can analyze all of the clouds using the various principles set forth above and disclosed herein to select the appropriate resources in one or more clouds for processing the workload. In one example, requestor **312** may include in their SLA requirement that is enforced by the broker **310**, a requirement that if its workload is being processed in cloud **308**, that if their performance level drops below a threshold or there is a failure of services from cloud **308**, that the broker **310** can intelligently apply those policies that can then migrate **320** workload to another cloud **306** with the beneficial result of providing business continuity for its computing requirements. In this respect, the present disclosure provides control of the SLA to be with the consumer. The broker **310** may only need to engage in a consumer agreement with various cloud providers and utilize that consumer agree-

ment amongst multiple clouds in order to successfully act as a broker for third party requestors. The more clouds registered with or in communication with the broker, the better service that the broker can provide to third party requestors and the wider the variety of possible SLA requirements the broker can enforce.

[0048] The broker **310** can provide brokering services in several ways. First, the broker may provide information only. In this case, the requestor **312**, **314**, **316** or **318** may simply receive information back about available resources. For example, the information may be that cloud **304** is offering a discount for processing a workload assuming that a user can wait 24 hours. The broker **310** can engage in interaction with the submitter and manage the relationship between the respective cloud and the user to receive a commitment that the user will accept and then manage receiving the workload and transmitting it to the selected resources in the respective cloud or clouds.

[0049] In another aspect, the broker **310** selects the appropriate resources (based on factors disclosed herein) and then route the received workload to those selected resources. The selected resources can be reserved and managed by the workload management software **302A**, **304A**, **306A**, **308A**. Requesters **312**, **314**, **316** and **318** provide workload to the broker **310** which then communicates the workload to the selected resources in one or more of the compute environments. Brokering service system **310** can be an advocate for or a guarantor of the service level agreement (SLA) associated with the request. The individual environments **302**, **304**, **306** and **308** provide respective SLAs, but the analysis done by the broker **310** can enforce a SLA requirement associated with the workload or requester. One of the concepts disclosed herein is that the broker **310** is the guarantor of an SLA although it does not provide compute resources per se to the requestor. By being a brokering system, broker **310** can identify the resources that will ensure compliance with the requestor SLA. This can be accomplished by reserving resources based on SLAs, migrating reservations to maintain SLA requirements, or other means.

[0050] In some instances, the workload itself may need to be modified or prepared to match the selected compute resources. For example, the selected resources may be running the Microsoft operating system and the workload was initially prepared or primarily designed to be processed on a Linux operating environment. The brokering service system **310** can evaluate and process, if necessary, workload associated with the request for compute services resources in order to match the workload with the selected compute resources. The appropriate modifications can be made and thus ensuring that the workload when it ultimately consumes resources within the environment will do so efficiently and quickly. In some cases, different stages of a compute job are optimally executed in different compute environments. In such cases, the broker can establish a compute job workflow between multiple nodes or clouds which provide the optimal compute environment for that particular stage. For instance, the compute job workflow can prepare initial data in a Linux environment, process the data in a Windows environment, and prepare reports in a BSD environment. The broker can appropriately establish a pipeline between multiple nodes or clouds so the compute job is processed in an optimal environment in an optimal location at each stage.

[0051] In another example, the broker **310** can determine that rather than translating or modifying the workload, the

selected resources within the one or more clouds can be modified or reprovisioned in order to have more affinity for the workload. Taking the above example, if the workload was initially prepared or primarily designed to be processed in a Linux operating environment, the broker **310** can reprovision a Microsoft environment to a Linux operating environment and in essence create a virtual private cluster or a virtual private cloud within a larger public cloud and then transfer the workload to that environment for consumption. In this manner, the present disclosure provides a standardized brokering environment for communicating workload into the cloud. Providing a single interface to the various clouds greatly simplifies the ability of any individual and user to obtain cloud computing services that match or meet their SLA requirements without the need to individually poll or communicate with various cloud service providers.

[0052] Once the workload is communicated to the selected compute environment to the selected resources which can span one or more of the compute environments, the brokering system server **310** can further analyze parameters associated with how the selected compute environment is processing the communicated workload and if a particular threshold is met, identify and select new compute resources from the compute resource environments and migrate all or part of the communicated workload to the selected new compute resources. This migration can be represented by line **320**. In this case, the broker can communicate with the various workload management modules **302A**, **304A**, **306A**, **308A** within each environment and inasmuch as the broker **310** has a more global view of workload and resources amongst the multiple environments, it can manage the migration of workload to new compute resources.

[0053] One example of where this can be helpful can include the ability of the broker **310** to keep prices low by moving workload to new environments that are always processing at night when the resources are cheaper (i.e. when electricity prices are off-peak) and the environment is operating at a cooler temperature. For example, if a particularly large job were to take 24 hours to process, broker **310** can continue to migrate the processing of the workload to new compute environments throughout the world during the 24 hour period to ensure that it is always processed at night. Any type of threshold associated with the processing of the workload can apply. For example, if a particular workload requires a high-level of reliability of the resources, and if the workload is currently being processed in cloud **304** but a threshold becomes met which involves the reliability of the availability or the continued availability of resources in cloud **304** drops below a particular threshold, then the broker **310** can manage the migration of those resources to cloud **306** which maintains a high-reliability report and expectation.

[0054] In another aspect, the broker **310** can manage the ability to provide “hot-hot” processing for the workload. The broker can identify two similar but separate sets of compute resources which may be in the same environment or may span different environments, which can in parallel or alternately process the workload. Assume that workload is processed partially in cloud **302** and partially in cloud **304**. The broker **310** can manage data or processing such that it alternates between the resources in cloud **302** and other resources in cloud **304** such that the reliability of the job not failing can dramatically increase inasmuch that as both separate sets of resources are “hot” and currently processing portions of the workload. If one of the sets of resources goes offline or goes

down, then the other set can immediately continue to process all of the remainder of the job.

[0055] In another aspect, the broker **310** can select resources including one set in one compute environment that process the workload and the back-up set of resources in the same or different compute environment that does not process the workload unless the one set of compute resources is unable to process the workload according to the service level agreement. A “hot-cold” environment creates a back-up set of resources reserved and made available in case the primary processing resources become unavailable or unable to process the workload. In this manner, a requestor **312**, **314**, **316**, **318** can request any variety of environments for processing their workload beyond just identifying x number of nodes, and have that environment identified, selected, reserved and made available throughout multiple clouds. Therefore, in addition to SLA requirements, the broker **310** can configure increasingly complex environments for workload according to the requestor’s desires and needs.

[0056] A “hot-cold” environment can have a size 5 primary operating environment which is reserved and allocated for processing workload, while simultaneously reserving a size 3 back-up environment which, if necessary, can stand by and process the workload if the size 5 primary environment experiences a power failure or some other event which eliminates its ability to process the workload. The size 3 back-up environment may take longer to process the workload but less than the full amount of resources are limited as back-up resources in the size 3 environment when compared to the size 5 primary environment. In this respect, the broker **310** can manage complicated arrangements for users in order to seek to provide the best services and back-up capabilities as is desirable or affordable.

[0057] In another example, two workloads are separate but related and share at least a portion of a same set of input data. The broker **310** can place the two workloads in close proximity to each other, possibly even within the same physical facility, so that each workload can easily share the same dataset with reduced latency rather than duplicating the input data to two locations. This approach can be useful to minimize reserved storage space and possibly reduce the cost to the customer for performing the workload. The broker **310** considers the SLA requirements for both workloads when selecting where the workloads should be processed.

[0058] FIG. 3 further illustrates another aspect of the present disclosure which involves workload being communicated between brokers. A secondary broker **322** is illustrated which communicates with compute environments **324** and **326**. Each of these has respective workload management modules **324A** and **326A**. Users **328** and **330** can submit requests and workload to broker **322**. One broker **322** can communicate with another broker **310**. All of this can be done transparent to the end users. Broker **322** may not be able to find ideal or satisfactory compute resources within environments **324**, **326** but may need to communicate the request to broker **310**. Inasmuch as broker **310** has access to more environments and more available resources, part or all of the workload submitted by a user **328** may need to be communicated from broker **322** to broker **310** in order to satisfy the service level agreement associated with the request from the user **328**.

[0059] Broker **310** can also control the access to the resources within the compute environment based on an authorization associated with the submitter of the request. For

example, the requestor **312** may be a government employee or unit with a high security requirement. In this case, with that information communicated to broker **310**, the broker may only poll or only evaluate resources within the environment **306** and **308**, which may be behind a government firewall, internal to the government or have a sufficiently secure system in order to process the workload. The particular available workload compute resources from the resource environments is narrowed or only available for selection based on the authorization associated with the submitter **312** of that request.

[0060] Having the broker **310** communicating with submitters **312**, **314**, **316**, **318** and with the various compute environments enables the broker **310** to obtain knowledge about the desires and needs of individual requestors. Parameters associated with the known needs of the submitters, their type of workload, how previously selected resources processed that workload, and so forth can be used to select advertisements that are presented to a requestor from the brokering service system **310**. Fees may be charged to the requestor for the brokering of compute services which may be flat fees, a percentage of the fee which is charged to the individual environments **302**, **304**, **306**, **308** for the consumption of resources within those environments, fees for advertisements, click through fees, and so forth. It is contemplated that any number of mechanisms may be used to charge users for the use of the resources and the opportunity to use the brokering service.

[0061] The broker server **310** can provide a customer with a graphical display of progress and location of submitted compute jobs within various public and private clouds. The customer can manipulate or suggest different courses of action to the broker server **310** through the graphical display and see in advance the potential outcomes of such suggestions. For example, the user can suggest splitting the workload over three clouds instead of over two clouds. Feedback on cost and availability of such as approach can then be presented to the user. The system can demonstrate real-time performance metrics of the current configuration and project performance of the suggested change. In this way, the broker server **310** can allow users to select less than optimal compute environment(s) for security, personal preference, benchmarking, or other reasons. The customer can also provide the broker server **310** with performance information and characteristics of the customer's private cloud. In this way, the broker server **310** can calculate differential metrics with respect to expected workload performance of the customer's private cloud and the broker's allocation of the workload. For instance, the broker server **310** can show that the job is being performed at an overall time savings, money savings, or some other absolute, relative, or time-based metric.

[0062] Additionally, the knowledge that is obtained both about individual requestors, their type of workload and how previously selected resource efficiently and appropriately process that workload can be used to make improved selection choices. As workload is processed by either registered or non-registered clouds, information is gathered and stored and analyzed associated with how that cloud performs. As further intelligence is gathered, improved choices are made in terms of selected resources. For example, additional selections by the broker **310** may avoid the non-registered cloud based on its performance. Individual servers within clouds may be utilized to make improved decisions in the future. As noted above, registered clouds or clouds that use complementary workload management software to that used by the broker **310** will enable more confidence in the data obtained for each

cloud and thus the learning algorithms may improve even quicker the identification and selection of resources for managing workload. As an example of this in operation, assume that every Friday night a large accounting company has a large batch job that is processed over the weekend for providing updated data Monday morning. As the batch job is processed week after week, the broker **310** in running its learning algorithms may initially select resources from cloud **302** but over time may consistently identify the best resources are combined between clouds **304** and **306**. In this case, later the workload may either be migrated during processing or the initial selection of resources will be selected for the similar type of batch job that is periodically received for processing by the broker **310**. As noted above, techniques such as co-allocation, timing, weighting of the data, and using all of the information that is achieved when identifying and analyzing the various resources, are utilized for finding and reserving resources within one or more of the group of clouds to ultimately provide the resources.

[0063] Decisions on which resources to select from the group of clouds can be made based at least in part on a co-allocation analysis of different types of resources. For example, CPUs and data accessing speeds for hard drives illustrate the co-allocation process. Inasmuch as co-allocation involves different types of resources, the most suitable resources for any particular workload may span different cloud computing sites. Inasmuch as the broker **310** has more of a global view of all of the various clouds, a co-allocation analysis may provide the best resources for any given workload whether those resources spanned multiple clouds.

[0064] One mechanism by which workload can be communicated through the broker **310** to an individual cloud can be to provide an instance of workload management software actually within the selected cloud. For example, assume that requestor **312** requests compute resources and broker **310** selects cloud **304** to provide those resources. Rather than simply communicating workload to the cloud for consumption, the broker **310** can manage the creation of an instance of workload management software in cloud **304**. This can be represented in this particular example by module **304A**. The Moab branded software from Cluster Resources may be an example of the type of software an instance of which can be installed or provisioned to manage the particular workload of a requestor.

[0065] In this scenario, several benefits are realized. First, the workload management software **304A** can perform necessary provisioning or modification of the reserved nodes in order to duplicate or create a particular environment which is most suitable for the workload from the requestor **312**. Furthermore, having an instance of the workload management software **304A** on cloud **304** can enable more efficient communication between the cloud **304** and the broker **310**. Therefore, one aspect of this disclosure involves receiving a request from a requestor **312**, at a brokering system **310**, identifying resources within a cloud **304** for consumption, installing an instance of workload management software **304A** within that selected cloud **304**, providing a modification, if necessary, of the selected resources via the installed workload management software **304A**, providing necessary communication of SLA and other requirements between the broker **310** and the workload management software **304A** on the cloud **304**, and consuming the resources associated with the request according to transmitted workload to the cloud **304**. This creates a package of consumed resources that can grow and shrink according to

requirements and is essentially a virtual private cloud. These resources then are configured to match middleware, operating system, particular libraries and applications or any other requirements of the workload so that these are not raw resources but dynamically created. The environment that is created in cloud **304** can be provisioned to match a private cloud such as cloud **212**. In other words, if a company has a private cloud **212** and utilizing a broker **310** because they have overflow workload and need additional resources from a public cloud, this mechanism of providing an instance of workload management software can receive from the requestor **312** information associated with the environment of the private cloud, create the instance **304A** of workload management software on cloud **304** and duplicate the environment such that the overflow workload will most efficiently and easily processed in the other cloud environment. This may be accomplished through virtualization through physical provisioning or any other mechanisms of modifying and managing the environment in this concept of a virtual private cloud. Principles applicable to this concept can be found in U.S. application Ser. No. 11/276,852, referenced above and incorporated herein by reference.

[0066] FIG. 4 illustrates a system with various components that is similar in some respects to FIG. 3 with additional features. Feature **430** illustrates a local cloud network with various local hosting clouds (LHC) **432**, **434**, **436**. These represent unifying various disparate resources typically within a company or a group of companies. This local cloud network **430** may sell its resources to other third party users. Software such as that offered by Cluster Resources can monitor, identify, and aggregate resources for the local cloud network **430** for private or public consumption. As is shown, if the resources available in the local cloud network **430** are not sufficient for handling all of its workload, it can communicate with the broker **310** as noted herein and tap into other cloud networks such that workload can be communicated as instructed by the broker to consumption resources for processing.

[0067] Application **453** represents a software application that may be used by a third party that has a need to tap into additional consumption resources. This can represent any conceivable software application run by a third party. As shown, an instance of workload management software **452** is shown as embedded within or linked to application **453**. Environment **454** represents a private cloud, cluster, grid or any other compute environment that is running the application **453**. An example of such workload management software is the Moab branded software from Cluster Resources. In this case, as the application **453** operates in a private cloud or a local environment **454** if the software module **452** can monitor that environment and the consumption of resources within that environment. Software **452** can intelligently determine if more resources are needed. If so, software **452** can communicate with the broker **310** and determine whether additional resources are available and at what SLA requirements and costs. For example, resources in public cloud **304** can be utilized by application **453**. Then, as the use of those resources is complete, application **453** can then withdraw from the public cloud **304** and continue running only in its local environment. Using the broker **310**, application **453** can obtain the best service without needing to establish and maintain a relationship with all of the available clouds. As application **453** runs in the cloud environment **304**, the instance of

management software **452** can be already embedded and used to customize the consumed resources in the cloud **304**.

[0068] Also shown in FIG. 4 is communication between broker **310** and software as a service (SaaS) **416**. As is known in the art, SaaS is a method of software deployment in which a provider **416** licenses an application to customers for use of its service on-demand. The end user does not purchase or install the software on a local computer but the software vendors host the application in their own environment **416** for use by an end user. In some cases, the software may be downloaded to a client device but disabled after the on-demand contract expires. There are many benefits to the SaaS approach in which the end user does not have the burden of purchasing a device that includes every necessary software application. Licensing requirements are also simplified as well as the need to continually update the software with corrections such as patches and upgrades and so forth. SaaS can also take advantage of the Service-Oriented Architecture (SOA) which enables software applications to communicate with each other. One aspect of the present disclosure is the ability of end users **312**, **314** to access SaaS software from a vendor **416** via the broker **310**, including using cloud resources.

[0069] The ability of the broker **310** to communicate with multiple clouds can also improve the ability of an SaaS service vendor **416** to avoid the cost of managing their own servers in a local environment. The broker **310** can communicate with the SaaS vendor **416** and identify and select the consumption resources within the various cloud providers for the consumption of resources utilizing the SaaS software. Therefore, this aspect of the disclosure involves receiving a request at a broker **310** for using software as a service on-demand (**602**), coordinating with a software on-demand vendor to provide a software application (**604**) identifying and selecting consuming resources from a cloud environment (**606**), and running the selected software application in the selected resources in the cloud environment (**608**). Once the SaaS software has completed its job, and the on-demand contract expires, the resources within the cloud are released for other users and any necessary data and reporting based on the SaaS software can be transmitted back to the requestor **312**, **314**. Licenses for the on-demand use of the software and charging for the use of cloud resources can be also managed by the broker **310** between the selected cloud and the SaaS vendor **416** and invoiced to the requestor **312** **310** in an efficient manner.

[0070] Also as shown in FIG. 4, a database **418** may communicate with the broker and/or the SaaS vendor **416** as necessary. Data service **418** represents a source of data in this respect, but broker **310** can also be a data broker such as scientific data, business intelligent, temperature data, and as such may be a stream for data on demand. In this respect, FIG. 4 also illustrates a requestor **312**, **314** which may require access to data **418** (independent of the SaaS vendor) which may need to utilize consumption resources within the cloud. A requestor can access the data such as might be available from a RSS feed, satellite data, or any other data source, which may need to be combined or utilized with other software services but that need consumption resources within the cloud. Accordingly, data is available to be accessed by the broker **310**. As usage of a particular SaaS application increases, the broker **310** can automatically and transparently expand the resources allocated to that SaaS application within a cloud.

[0071] The above approach enables the SaaS software to be plugged in, made available and hosted or consumed on any particular resource. There may be more than one instance of the software, but the broker 310 will manage the rights, licensing, and metering to charge consumers according to SLA or other contractual relationships. An entity can provide a data mining service 418 that can provide intelligence or simply data to consumers. In this respect, if user 312, 314 through broker 310 is accessing data, it may simply be accessing a service that simply receives data such as business intelligence. The user may simply get the results of the data. This can include such things as in Google intelligence on click-through data, news sources, weather information, and so forth. Accordingly, accessing data sets such that it becomes available through the use of broker 310. In this regard, this embodiment includes receiving a request from a user 312 at a cloud broker 310 that involves a request for data and/or a need for consumption resources. Broker 310 communicates with data source 418 for providing that data and, if necessary identifies and selects consumption resources within a number of cloud environments and identifies and utilizes the consumption resources for processes that are associated with the data. Finally, the broker 310 can manage the reporting of the data or processing associated with the data to the user 312. In this way, the broker 310 can interface with the appropriate computational resources with needed specific sets of software to get the results desired by the user 312.

[0072] An example of the above approach is in human resources where applications involve a number of people simply consuming a service. These people may need to access a certain set of data, which may be internal to a company, and the value of that data may be part of what is provided back to the end user as a service. For example, companies like Gartner may access data 418 available from various data providers with regards to purchasing choices by large companies and this data may be accessible via broker 310.

[0073] In another aspect, broker 310 can provide a particular service to end user requestors 312, 314 or other environments 430, 440. For example, broker 310 can provide entities with the ability to have a continuity of compute resources. In this respect, with an appropriate SLA in place, the broker 310 can provide a local environment 430 or private cloud 402 with the ability to manage and reduce down time of their own environments. In this case, assume private cloud 402 within an entity or a company goes down because of a power outage. Broker 310 can provide the necessary business continuity in case they have such a failure and immediately find resources that can provide that continuity and shift the workload from the down cloud 402 to, for example, a public cloud 408 that has available resources. This can be for internal consumption requirements as well as providing continuity for websites. As shown in private cloud 402, web server or web servers 442 provide communication to clients 444 and 446 over the Internet (not shown). This represents end users accessing data via one or more web servers. If private cloud 402 goes down, rather than irritating customers and losing access to those web services, broker 310 can identify resources within a cloud such as cloud 302 and duplicate the necessary data for those web servers on resources 448. In this instance, public cloud also includes the necessary communicating means with the Internet such that clients 450 and 456 can continue to go to that particular uniform resource locator (URL) and access that website. In this mechanism, the continuity of the availability of the company's website is preserved through the

ability of the broker 310 to immediately respond to some kind of failure and maintain that continuity. The broker 310 can act as a redirection server or can designate one or more nodes to act as redirection servers to redirect HTTP requests for the down cloud 402 to other suitable clouds. The redirection techniques can be static or dynamic and are preferably done in real time to reflect the actual status of the cloud(s) operating on the website. In one example, companies can plan for spikes in website usage by preestablishing additional resources with a cloud to handle expected increased traffic.

[0074] In yet another aspect, the broker 310 enables the ability of energy to be provided as a utility. For example, broker 310 can monitor the number of data centers that will send workload to since they have the ability to manage the energy consumption in those environments. For example, utility companies can pay a cloud in a location with high electricity demand more money to reduce their energy consumption and to get the same workload done by transferring the workload to another cloud. Public cloud 408 can be in an electrical grid that is operating on a very hot day in which there is a surge of demand for electricity. In this case, the energy company, rather than entering into a black out or a brown out, can request via broker 310, public cloud 408 to transfer workload to another public cloud in a cool environment which does not have a demand for energy consumption. Broker 310 can manage that transfer to cloud 304 and even manage the payment and negotiation of funds to enable that transfer. In this respect, utility companies can more efficiently and effectively manage the actual consumption of resources by the public clouds which are already large consumers of energy.

[0075] Broker 310 in this respect can also establish itself as an integral part of a national or international energy response service. Broker 310 can balance and manage local, national, international, or worldwide workload via its intelligent workload management algorithms.

[0076] Broker 310 in this respect can also manage carbon credits. Where various public clouds can be more energy efficient than others, broker 310 can also manage the differentials in efficient energy consumption and can manage the buying and selling of carbon credits between the various clouds with which it has a relationship. Therefore, the one aspect of the present disclosure is the broker 310 having the capability of managing carbon credits between data centers. In one example, assume public cloud 408 has an efficient energy consumption in which its energy is received from a nuclear power plant and its efficiency is such that it has carbon credits which it can sell. Assume that cloud 304 is inefficient and is in need of purchasing carbon credits because of its energy consumption approach. Broker 310 can manage the purchase of carbon credits by public cloud 304 from public cloud 408. In addition, broker 310 can utilize the knowledge gained by virtue of its green computing algorithms incorporated herein by reference above from application Ser. No. 12/179,142 and also have the ability to manage the distribution of consumption resources such that independent clouds may not reach a threshold and have to purchase carbon credits from other clouds. In this respect, assuming that individual cloud environments utilize software algorithms that allow the energy consumption management of their environment to communicate with broker 310 those energy consumption parameters and capabilities, the broker 310 can most efficiently distribute workload amongst the various clouds in order to achieve in a more globalized view, the most efficient

use of energy. The broker **310** can track information on what has been consumed in various environments and identify where the cheaper environments are with regards to energy consumption. As various clouds may hit targets, and as the cost of a carbon credit may rise too high, broker **310** can, dynamically or under the instruction of a user or a public cloud, shift the consumption of resource to another type of energy or to a type of cloud. In one aspect, the various clouds may be in various countries with different climate change policies. For example, if Europe has a more advanced climate change policy than China or the United States, the SLA requirements enforced by broker **310** may include details for how the markets for carbon trading are to be managed and triggers could be in place which would cause the broker **310** to transfer workload amongst the clouds based at least in part on these factors.

[**0077**] Feature **440** illustrates a network with sub-administrative elements **438**. This may represent a compute domain, a data domain, a cluster or grid, or an on-demand service center in which the network operates like a utility that can store or manipulate data or do computational work. This is illustrated in FIG. **4** as another aspect of this disclosure in which the broker **310** can communicate with such a network **440** and provide the ability of brokering consumption or other types of resources within the various clouds such as network resources as needed by network **440**.

[**0078**] FIG. **5** illustrates an example of a method embodiment of the present disclosure. The method involves brokering compute resources and includes, at a brokering service system, polling a group of separately administered compute environments to identify resource capabilities, each computing resource environment including a group of managed nodes for processing workload (**502**) (each compute resource environment may be registered or not with the broker), receiving a request for compute resources at the brokering service system, the request for compute resources being associated with a service level agreement (**504**), based on the identified resource capabilities across the group of compute resource environments, selecting compute resources in one or more of the group of compute resource environments (**506**), and receiving workload associated with a request and communicating the workload to the selected resources in the one or more of the group of compute resources environments for processing (**508**). The selection of compute resources ensures that the processing complies with the service level agreement. Additional steps of the method include evaluating the processing and the workloads affinity to the selected compute resources. If necessary, the workload associated with the request for compute resources to match the workload with the selected compute resources. The polling can further identify a cost for available compute resources and each of the compute resource environments.

[**0079**] After communicating the workload to the selected compute environment, the method can further include analyzing parameters associated with how the selected compute environment is processing the communicated workload and, if a threshold is met, identifying and selecting new compute resources from the group of compute resource environments and migrating the communicated workload to the selected new compute resources. When the selected compute resources span more than one compute resource environment in time, the method can further include migrating the workload to compute resources within the more than one compute environment. As noted above, this can involve migrating to

cheap resources, resources processing at night in cooler environments, and so forth. In the case of encryption-related export restrictions or other applicable international laws, the brokering system can further limit encryption-related workloads, for example, to domestic clouds or international clouds which are not affected by export restrictions.

[**0080**] In another aspect, the brokering system **310** can include learning algorithms which, in time, improve their ability to understand the needs of workload which are submitted into the various environments as well as the environments themselves. For example, if a particular type of workload is received from an accounting company, and the original selection of resources is in cloud **302** for processing the accounting workload, the brokering service **310** can predict how that workload will be processed in environment **302**. However, after the environment actually processes the workload, improvements and intelligence can be gained disclosing how well and efficiently the environment process the workload as well as how well the workload needs matched with the capabilities of the environment. Over time, the brokering service **310** may actually select different resources which will more suitably match the workload. The broker **310** can record expected and actual performance of various clouds to extrapolate future performance and assign workloads based on expected performance.

[**0081**] FIG. **6** is a method presented within the frame of reference of how the interaction with the broker is viewed from the point of view of the public or private cloud. As shown in FIG. **6**, the method includes receiving from a cloud broker **310** a polling to request information about the cloud (**602**). In this regard, the communicated information about the capabilities and information associated with a particular cloud may be exchanged in any number of ways. As is noted herein, the broker may periodically or on some other basis request information about the cloud. In another aspect, the cloud may periodically transmit information to the broker about its capabilities or a change in capabilities. For example, if the cloud has a large job that is scheduled to consume resources but that job gets cancelled and those resources become available, the cloud can transmit information to the broker **310** which notifies the broker **310** of the freed-up resources that have become available. Alternate pricing and other information for those resources may also be offered which can be used by the broker **310** to shift workload to those resources according to the individual SLA requirements for the various requestors. Requestors can then take advantage of a reduced pace and more immediate availability and those cycles are not wasted. The cloud then can receive a request for resources and select/provision/allocate resources based on the request (**604**). The cloud can include workload management software which can select, reserve, and/or provision necessary functionality or software within selected resources. In general with knowledge of the requirements from the broker, the cloud will prepare the appropriate resources for consumption by the workload. Next, the cloud receives workload and associated requirements for processing the workload, wherein the requirements are associated with the request and an SLA (**606**). Finally, the cloud processes the workload for the workload consumes resources in the cloud with the selected resources (**608**).

[**0082**] Other embodiments include additional views of the operation of the broker from different entities within the system. For example, the broker **322** shown in FIG. **3** illustrates and embodiment of the present disclosure in which

requests are made for cloud computing services from requestors 328 and 330. The broker communicates with clouds 324, 326 within its group of resources. If a threshold is not met or the available clouds to broker 322 cannot satisfy the workload of requestors 328, 330, broker 322 then polls or transmits a request to broker 310 who can return information to broker 322 with regards to the availability of resources in its group of clouds 302, 304, 306, 308 that can meet the SLA requirements of requestors 328 or 330. Broker 322 can then communicate with broker 310 to receive workload from requestors 328, 330 and enable the communication of that workload to a cloud having a relationship with broker 310 such that the requestors workload can be consumed according to their requirements.

[0083] Another embodiment can be described in view of another entity such as private cloud 410 shown in FIG. 4. This embodiment is shown in FIG. 7. The private cloud 410 operates to receive workload from a private cloud broker 412. However, if workload in the private cloud 410 becomes too demanding or a threshold event occurs that prevents the private cloud 410 from being able to offer compute resources according to private SLA requirements, then the private cloud 410 communicates with the private cloud broker 412 (702). The private cloud broker communicates with cloud broker 310 to identify what other compute resources are available in other private clouds 402, 404 or public clouds 304, 308 (704). The cloud broker 310 communicates with its various computing environments to ensure that resources are reserved/provisioned/allocated in preparation for their consumption by workload from the private cloud 410 (706). Workload then flows from private cloud 410 through communication with the private broker 412 and broker 310 to another consumption environment such as public cloud 302 (708). The communication of the workload is shown via line 460 which can be through any network or communication means which is available. In other words, the transfer of workload does not necessarily travel through private broker 412 and broker 310 to a public cloud 302 from the private cloud 420 but the brokers can manage the transition of workload from a private cloud 410 to another compute environment like public cloud 302.

[0084] The private cloud broker 412 therefore does not need to “poll” clouds 302, 304, 408 or other environments 440. The broker 310 has all of that information aggregated. Broker 412 may only be polling broker 310 and optionally other brokers (not shown). In this manner, broker 412 can easily obtain a view of other resources available to it for additional processing capabilities. All necessary SLA requirements, security, compute environment provisioning or allocating, can be managed between broker 412, broker 310 and the resulting cloud environment 302 that has the resources to be consumed (310).

[0085] Note that other various embodiments can be developed based on the disclosure set forth herein but while viewing the processing from another entity including SaaS 416, data source 418, application 453 with an embedded workload management module 452, environments 430, 440, and so forth. For example, FIG. 8 illustrates a method embodiment of an application 453 having an embedded copy of a workload management software module 452. The module 452 monitors the environment 453 on which the application is being processed (802). This can include how closely the environment 454 is able to meet SLA requirements or other internal thresholds. If a threshold is met which indicates that the current or perhaps projected capability of the environment 454 is unable

to meet the appropriate standards, then the module 452 communicates with the broker 310 (804). The module 452 can include in the communication details about the preferred environment 454 and receive information back regarding the aggregated resources in the various clouds 302, 304, 408, etc. communicating with the broker 310 (806). If the appropriate resources are available at acceptable rates and capabilities (for example in cloud 408), then the broker 310 insures that the resources are reserved, provisioned, allocated, etc. such that the workflow for application 453 can be transferred in whole or in part to the cloud 408 for consuming additional resources (808). Part of the workload may continue to be processed in environment 454. For example, the data-centric portion may be shifted to the cloud 408. An instance of workload management software 452 can be optionally installed to manage the virtual private cloud (resources) created in cloud 408 for the job (810). As the job completes the resources are made available to other users, the system processes appropriate charges for the use of the resources and finalizes data reporting.

[0086] FIG. 9 illustrates an embodiment from the viewpoint of an entity offering an SaaS service 416. In this embodiment, the SaaS service 416 receives a request for use of software from a user 313 (902). The user request may include SLA requirements for the use of the SaaS software which a local compute environment (not shown) for the SaaS service cannot meet. The SaaS service 416 communicates with broker 310 to identify whether consumption resources are available for the SaaS service software (904). Certain SLA requirements are preferably included in the communication. If resources are available, then between the broker 310 and clouds 302, 304, 408, etc., resources are reserved, allocated, provisioned if necessary with operating systems or other requirements (906), and the SaaS software is then loaded onto the allocated resources for consumption (908). As the SaaS job completes, reporting, charging, and other housecleaning matters are finalized and the resources are made available to other users. User 313 may only be charged a single fee, some of which will be allocated to the SaaS service and some of which the SaaS service may forward to the cloud 408 (or whatever cloud had the resources), for the consumption of the resources and software, licensing etc.

[0087] Embodiments within the scope of the present disclosure may also include tangible computer-readable storage media for carrying or having computer-executable instructions or data structures stored thereon. Such computer-readable storage media can be any available media that can be accessed by a general purpose or special purpose computer, including the functional design of any special purpose processor as discussed above. By way of example, and not limitation, such computer-readable media can include RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to carry or store desired program code means in the form of computer-executable instructions, data structures, or processor chip design. When information is transferred or provided over a network or another communications connection (either hardwired, wireless, or combination thereof) to a computer, the computer properly views the connection as a computer-readable medium. Thus, any such connection is properly termed a computer-readable medium. Combinations of the above should also be included within the scope of the computer-readable media.

[0088] Computer-executable instructions include, for example, instructions and data which cause a general purpose computer, special purpose computer, or special purpose processing device to perform a certain function or group of functions. Computer-executable instructions also include program modules that are executed by computers in stand-alone or network environments. Generally, program modules include routines, programs, components, data structures, objects, and the functions inherent in the design of special-purpose processors, etc. that perform particular tasks or implement particular abstract data types. Computer-executable instructions, associated data structures, and program modules represent examples of the program code means for executing steps of the methods disclosed herein. The particular sequence of such executable instructions or associated data structures represents examples of corresponding acts for implementing the functions described in such steps.

[0089] Those of skill in the art will appreciate that other embodiments of the disclosure may be practiced in network computing environments with many types of computer system configurations, including personal computers, hand-held devices, multi-processor systems, microprocessor-based or programmable consumer electronics, network PCs, mini-computers, mainframe computers, and the like. Embodiments may also be practiced in distributed computing environments where tasks are performed by local and remote processing devices that are linked (either by hardwired links, wireless links, or by a combination thereof) through a communications network. In a distributed computing environment, program modules may be located in both local and remote memory storage devices.

[0090] Although the above description may contain specific details, they should not be construed as limiting the claims in any way. Other configurations of the described embodiments are part of the scope of this disclosure. For instance, while the principles disclosed herein are generally discussed in terms of a public cloud, a private cloud 412 can also receive workloads from a private cloud broker. The principles herein are applicable to all cloud compute environments. Those skilled in the art will readily recognize various modifications and changes that may be made to the principles described herein without following the example embodiments and applications illustrated and described herein, and without departing from the spirit and scope of the disclosure.

I claim:

1. A method of brokering compute resources, the method comprising:

at a brokering service system, polling a plurality of compute resource environments to identify a type of energy that powers each of the plurality of compute resource environments, each of the plurality of compute resource environments comprising a plurality of aggregated nodes;

based on a request for compute resources at the brokering service system and the type of energy that powers each of the plurality of compute resource environments, reserving compute resources from the plurality of compute resource environments to yield reserved compute resources for consumption;

receiving a workload associated with the request and communicating the workload associated with the request to the reserved compute resources for processing, wherein

reserving the compute resources ensures that the processing at the second time complies with a service level agreement;

after communicating the workload associated with the request to the reserved compute resources, analyzing parameters associated with how the reserved compute resources are processing the workload associated with the request; and

if a threshold is met, identifying and reserving new compute resources from the plurality of compute resource environments and migrating the workload to the new compute resources.

2. The method of claim 1, further comprising: evaluating and processing, if necessary, the workload associated with the request to match the workload associated with the request with the reserved compute resources.

3. The method of claim 1, wherein the polling further identifies a cost for available compute resources in each of the plurality of compute resource environments.

4. The method of claim 1, wherein the workload associated with the request is batch workload.

5. The method of claim 1, wherein the reserved compute resources span more than one compute resource environment of the plurality of compute resource environments.

6. The method of claim 1, wherein the reserved compute resources span more than one compute resource environment over time, and wherein the method further comprises:

migrating the workload associated with the request over time to compute resources within the more than one compute resource environment.

7. The method of claim 1, wherein the reserved compute resources comprise two similar but separated sets of compute resources, each set being in a differently administered compute resource environment, but that both alternately process the workload associated with the request.

8. The method of claim 1, wherein the reserved compute resources comprise one set of compute resources in one compute resource environment that process the workload associated with the request and a backup set of resources in a different compute resources environment that does not process the workload associated with the request unless the one set of compute resources is unable to process the workload associated with the request according to the service level agreement.

9. The method of claim 5, wherein the a first portion of the reserved compute resources comprises data centric resources in a first compute resource environment and a second portion of the reserved compute resources comprises compute intensive resources in a second compute resource environment, wherein the data centric resources are resources associated with storing data and compute intensive resources comprise processing units in the second compute resource environment.

10. The method of claim 1, wherein which compute resources of the plurality of compute resource environment are available for reservation is based on an authorization associated with a submitter of the request.

11. The method of claim 1, further comprising presenting an advertisement to a requestor.

12. The method of claim 1, further comprising charging a fee to a requestor for brokering compute resources.

13. The method of claim 1, further comprising receiving a fee from the compute resource environments having selected compute resources that process the workload associated with

the request based on fees charged by the compute resource environments having reserved compute resources that process the workload.

14. The method of claim **1**, wherein the request is received from a requesting brokering service system that polls a second plurality of compute environments to identify resource capabilities, each compute resource environment associated with the requesting brokering service system comprising a plurality managed nodes for processing workload and each compute resource environment associated with the requesting brokering service system having registered with the requesting brokering service system.

15. The method of claim **1**, wherein each of the plurality of compute resource environment is registered with the brokering service system.

16. The method of claim **1**, wherein the type of energy comprises one of nuclear energy, coal-based energy and wind energy.

17. A system for brokering compute resources, the system comprising:

a processor; and

a non-transitory computer-readable storage medium operatively coupled to the processor and storing instructions which, when executed by the processor, cause the processor to perform operations comprising:

polling a plurality of compute resource environments to identify a type of energy that powers each of the plurality of compute resource environments, each of the plurality of compute resource environments comprising a plurality of aggregated nodes;

reserving compute resources from the plurality of compute resource environments based on a request for compute resources at the brokering service system and the type of energy that powers each of the plurality of compute resource environments to yield reserved compute resources such that the reserved compute resources can be consumed;

receiving a workload associated with the request and communicating the workload to the reserved compute resources for processing, wherein reserving compute resources ensures that the processing at the second time complies with a service level agreement;

after communicating the workload associated with the request to the reserved compute resources, analyzing parameters associated with how the reserved compute resources are processing the workload associated with the request; and

if a threshold is met, identifying and reserving new compute resources from the plurality of compute resource environments and migrating the workload to the new compute resources.

18. The system of claim **17**, the non-transitory computer-readable storage medium further comprising instructions which, when executed by the processor, cause the processor to perform further operations comprising:

evaluating and processing, if necessary, the workload associated with the request to match the workload associated with the request with the reserved compute resources.

19. They system of claim **17**, wherein the type of energy comprises one of nuclear energy, coal-based energy and wind energy.

20. A non-transitory computer-readable storage medium storing a computer program having instructions which, when executed by a processor, cause the processor to perform operations comprising:

at a brokering service system, polling a plurality of separately administered compute resource environments to identify a type of energy that powers each of the plurality of separately administered compute resource environments, each of the plurality of separately administered compute resource environments comprising a plurality of aggregated nodes for processing workload;

based on a request for compute resources at the brokering service system and the type of energy that powers each of the plurality of separately administered compute resource environments, reserving compute resources in one or more of the plurality of separately administered compute resource environments to yield reserved compute resources;

receiving a workload associated with the request and communicating the workload associated with the request to the reserved resources in the plurality of separately administered compute resource environments for processing at a second time which is later than the first time, wherein the reservation of compute resources ensures that the processing complies with a service level agreement;

after communicating the workload associated with the request to the reserved compute resources, analyzing parameters associated with how the reserved compute resources are processing the workload associated with the request; and

if a threshold is met, identifying and reserving new compute resources from the plurality of separately administered compute resource environments and migrating the workload to the new compute resources.

21. The non-transitory computer-readable storage medium of claim **19**, the instructions further comprising:

evaluating and processing, if necessary, the workload associated with the request to match the workload associated with the request with the reserved compute resources.

22. The non-transitory computer-readable storage medium of claim **19**, wherein the type of energy comprises one of nuclear energy, coal-based energy and wind energy.

23. A system for processing workload from a brokering service, the system comprising:

a processor; and

a memory operatively coupled to the processor and storing instructions which, when executed by the processor, cause the processor to perform operations comprising:

at a brokering service system, polling a plurality of compute resource environments to identify a type of energy that powers each of the plurality of compute resource environments, each of the plurality of compute resource environments comprising a plurality of aggregated nodes for processing a workload;

based on a request for compute resources at the brokering service system and type of energy that powers each of the plurality of compute resource environments, reserving compute resources from the plurality of compute resource environments to yield reserved compute resources such that the workload can consume the reserved compute resources;

receiving a workload associated with the request and communicating the workload to the reserved compute

resources for processing, wherein reserving the compute resources ensures that the processing at the second time complies with the service level agreement; after communicating the workload associated with the request to the reserved compute resources, analyzing parameters associated with how the reserved compute resources are processing the workload associated with the request; and
if a threshold is met, identifying and reserving new compute resources from the plurality of compute resource environments and migrating the workload to the new compute resources.

24. The system of claim **23**, wherein the type of energy comprises one of nuclear energy, coal-based energy and wind energy.

* * * * *