



- (51) **International Patent Classification:**
G06F 15/16 (2006.01) G06F 9/06 (2006.01)
- (21) **International Application Number:**
PCT/US2013/026907
- (22) **International Filing Date:**
20 February 2013 (20.02.2013)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
61/600,888 20 February 2012 (20.02.2012) US
- (71) **Applicant:** VIRTUSTREAM CANADA HOLDINGS, INC. [US/US]; 4800 Montgomery Lane, Suite 1100, Bethesda, MD 20814 (US).
- (72) **Inventor:** ANDERSON, Derek; 4401 Blauson Boulevard, #58, Abbotsford, BC V3G 0A4 (CA).
- (74) **Agents:** LAZAR, Dale, S. et al.; DLA Piper LLP US, P.O. Box 2758, Reston, VA 20195 (US).
- (81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY,


BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

Published:

- with international search report (Art. 21(3))
- with amended claims (Art. 19(1))

(54) **Title:** SYSTEMS AND METHODS INVOLVING VIRTUAL MACHINE HOST ISOLATION OVER A NETWORK



(57) **Abstract:** Systems and methods for virtual machine host isolation are disclosed. According to one implementation, an illustrative system may include a first compute node configured to be operatively coupled to a second compute node via a first application server, and to a third compute node via a second application server. In operation, the first compute node may be configured to receive an instruction from the second compute node via the first server to define a virtual machine, and send an instruction to the third compute node via the second server to define the virtual machine.

WO 2013/126430 A1

**SYSTEMS AND METHODS INVOLVING VIRTUAL MACHINE
HOST ISOLATION OVER A NETWORK**

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application is based on and derives the benefit of the filing date of United States Provisional Patent Application No. 61/600,888, filed February 20, 2012. The entire content of this application is herein incorporated by reference in its entirety.

FIELD

[0002] The present disclosure is directed generally to systems and methods involving virtual machine host isolation.

BRIEF DESCRIPTION OF THE FIGURES

[0003] The accompanying drawings, which constitute a part of this specification, illustrate various implementations and aspects of the innovations herein and, together with the description, help illustrate the principles of the present inventions. In the drawings:

[0004] Figure 1 is a block diagram of an illustrative system, according to a disclosed implementation.

[0005] Figure 2 is a flow chart depicting an illustrative method related to virtual machine host isolation, according to a disclosed implementation.

[0006] Figure 3 is a flow chart depicting another illustrative method related to virtual machine host isolation, according to a disclosed implementation.

[0007] Figure 4 is a flow chart depicting a further illustrative method related to virtual machine host isolation, according to a disclosed implementation.

[0008] Figure 5 is a flow chart depicting still another illustrative method related to virtual machine host isolation, according to a disclosed implementation.

[0009] Figure 6 is a flow diagram depicting an illustrative method of operation for isolating virtual machine host(s), according to an implementation.

DETAILED DESCRIPTION OF ILLUSTRATIVE IMPLEMENTATIONS

[00010] Reference will now be made in detail to the inventions herein, examples of which are illustrated in the accompanying drawings. The implementations set forth in the following description do not represent all implementations consistent with the claimed inventions. Instead, they are merely some examples consistent with certain aspects related to the present innovations. Wherever possible, the same reference numbers will be used throughout the drawings to refer to the same or like parts.

[00011] According to some aspects of the present innovations, implementations may relate to a system for isolating virtual machine hosts over a network using a federated downstream cluster. In one illustrative embodiment, a system may include a first compute node configured to be operatively coupled to (1) a second compute node via a first application server such as an extensible messaging and presence protocol (XMPP) server, and (2) a third compute node via a second application server such as another XMPP server. Here, for example, the first compute node may be configured to be included within a federated cluster that includes the third compute node. In operation, the first compute node may be configured to receive an instruction from the second compute node via the first

server to define a virtual machine. Further, the first compute node may be configured to send an instruction to the third compute node via the second server to define the virtual machine. In some embodiments, the first XMPP server can be the same as the second XMPP server.

[00012] With regard to the discussions below, a system using federated downstream clusters to isolate virtual machine hosts can be used to geographically and/or logically isolate virtual machine hosts. FIG. 1 depicts a block diagram of an illustrative system that may be involved with isolating virtual machine hosts ("system") 100, according to an embodiment. System 100 includes a federated cluster 110 including a compute node 102 and a compute node 104, an application server cluster 120 such as a jabber cluster, a compute node 106, and a network 140. In some implementations, compute node 102 of federated cluster 110 includes a downstream manager 112, compute node 104 of federated cluster 110 includes a downstream agent component 114, and compute node 106 includes an upstream manager component 130.

[00013] In some implementations detailed below, the application server cluster 120 is described as a jabber cluster that includes one or more extensible messaging and presence protocol (XMPP) servers logically connected and configured to support XMPP communication between upstream manager component 130 and downstream manager component 112, and between downstream manager component 112 and downstream agent component 114. However, various other processing components and messaging/communication protocols other than, or in conjunction with, XMPP may be utilized to process the command and control for the cluster(s), including but not limited to AMQP, ZeroMQ, and HTTP, among others. Here, for example, processing and

communication may also take place via hybrid protocols, such as combination(s) of HTTP and XMPP. Thus, while some discussions of FIG. 1 may describe system 100 as including a jabber cluster, in some embodiments system 100 may comprise other application server(s) and/or protocols. For instance, in certain embodiments, a variety of processing components and/or protocols may be utilized, provided that features such as topology, isolation and message content set forth in the detailed implementations shown and described herein are achieved. Here, for example, such application servers may utilize or include messaging or communication protocol(s) that require low processing and memory resources, are standardized, are customizable, are point-to-point, and/or are configured to send and receive state and/or provisioning messages.

[00014] Additionally, while FIG. 1 may describe system 100 as including a single cluster or jabber cluster, in other embodiments, system 100 can include more than one cluster or jabber cluster. In such plural jabber cluster embodiments, for example, a first jabber cluster can include one or more servers logically connected and configured to support communication between the upstream manager 130 and the downstream manager 112, and a second jabber cluster can include one or more servers logically connected and configured to support communication between the downstream manager 112 and the downstream agent 114. In this manner, communication between downstream manager 112 and downstream agent 114 may be secured separately from communication between upstream manager 130 and downstream manager 112, and may continue in the event of a failure of any portion of the first jabber cluster.

[00015] According to implementations herein, upstream manager component 130, downstream manager component 112, and downstream agent component

114 can be a software and/or hardware module(s) located at compute nodes, such as, for example, at compute nodes 102, 104, and/or 106. A compute node can be any type of device configured to send data to one or more devices within system 100 and to devices outside of system 100 via network 140, and/or receive data from devices included in network 140 and/or from devices outside of system 100 via network 140. In some embodiments, the compute node can be configured to function as, for example, virtual machine host, a server device (e.g., a web server device), a network management device, a data repository and/or the like. The compute node can be configured to define and send provision and/or action instructions, and/or add, remove, lock, revise and/or edit a virtual machine.

[00016] In some implementations, the compute nodes may include one or more memory 136 and/or processor 138 devices or components. The memory 136 can be, for example, a random access memory (RAM), a memory buffer, a hard drive, a database, an erasable programmable read-only memory (EPROM), an electrically erasable read-only memory (EEPROM), a read-only memory (ROM) and/or so forth. In some embodiments, the memory 136 of the compute node includes data used to define, send, and receive, instructions, messages, poll requests and results, etc. In some embodiments, the memory 136 stores instructions to cause the processor 138 to execute modules, processes and/or functions associated with such a system 100.

[00017] The processor(s) 138 of the compute nodes, such as, for example, compute nodes 102, 104, and/or 106, can be any suitable processing device configured to run and/or execute within system 100. In some embodiments, the processor can be a general purpose processor, a Field Programmable Gate

Array (FPGA), an Application Specific Integrated Circuit (ASIC), a Digital Signal Processor (DSP), and/or the like. The processor can be configured to execute modules, functions and/or processes to operate system 100.

[00018] According to some implementations, the upstream manager component 130 may be a software and/or hardware module(s) located at compute node 106, and may be configured to logically connect federated cluster 110 to network 140 via the application server cluster 130, such as a jabber cluster. Upstream manager component 130 may also be configured to receive hypertext transport protocol (HTTP) instructions from a device 145 of a user, for example, a customer or network administrator 150, via network 140. In some embodiments, the upstream manager component 130 may be configured to receive instructions via network 140 in other protocols, for example, representative state transfer (REST). Additionally, upstream manager component 130 may be configured to define an XMPP instruction(s) to downstream manager 112 via the application server cluster 120. In some embodiments, an XMPP instruction can include a provision instruction for defining a virtual machine.

[00019] In certain implementations, the upstream manager component 130 may be configured to receive provision results from downstream manager 112, may store those results, and may send status updates and/or results to the user etc. via network 140. Further, the upstream manager component 130 may be configured to logically interact with federated cluster 110 as a single compute node. Here, for example, the upstream manager component may be configured to treat a first federated cluster including a single downstream manager and a one or more downstream agents as a single compute node having a first amount of computing capacity, and may treat a second federated cluster including a

single downstream manager and ten downstream agents as a single compute node have a second amount of computing capacity distinct from the other compute node. In this manner, upstream manager component 130 need not store state and other information for each individual downstream agent, but instead may only store state and other information for each federated cluster. Upstream manager component 130 may also be configured to include customer data, such as available capacity, virtual machine quotas, and the like. In this manner, upstream manager component may accept or decline a request to provision a virtual machine, and/or can determine provisioning option based on that customer information. With regard to an associated compute node, compute node 106 may include one or more repositories 144A and/or one or more databases 146A, in addition to the memory 136A to store system and/or customer data, provisioning rules, etc.

[00020] As set forth herein, federated cluster 110 may be a cluster of compute nodes or multiple cooperating compute nodes that may not be centrally managed. In some implementations, federating the cluster of compute nodes can include designating one of the compute nodes as the location for a downstream manager and the remaining compute nodes as the location for downstream agents. The compute node designated as the location for the downstream manager can operate an agent emulation module to couple the federated cluster to an upstream manager via an jabber cluster, and can present the federated cluster to the upstream manager as a single downstream agent including the available capacity of the whole cluster of compute nodes.

[00021] In the illustrative system shown in FIG. 1, federated cluster 110 includes downstream manager component 112 and downstream agent component 114.

While shown in FIG. 1 as including a single downstream agent component 114, in some embodiments, federated cluster 110 can include any number of downstream agent components 114. Federated cluster 110 may be configured to communicate with the upstream manager 130 via application sever cluster and may be configured to logically appear to upstream manager 130 as a single compute node. Federated cluster 110 may further include or operate with local storage locations 136B, 136C and/or remote storage locations 136A, 136E for virtual machine, one or more local or remote databases 144A, 144D, 144E, and/or one or more local or remote repositories 146A, 146D, 146E, where such local databases and/or repositories may also or instead be located within specific compute nodes 102, 104. In embodiments, here, for example, such one or more local repositories 146D can be synced to an upstream repository 146A.

[00022] In some implementations, downstream manager component 112 of federated cluster 110 may be a software and/or hardware module(s) located at compute node 102 and may be configured to be logically connected to the compute node 106 and upstream manager component 130 via the application server cluster 120 and logically connected to downstream agent component 114 via the application server cluster 120. In this manner all traffic between downstream manager component 112 and upstream manager component 130, and all traffic between downstream manager component 112 and downstream agent component 114 may be sent and received via the application server cluster 120. Downstream manager component 112 can be configured to receive an instruction, such as an XMPP instruction(s), from upstream manager component 130. Downstream manager component 112 may also be configured to (1) define a provision request associated with downstream agent component

114, and (2) send that request to downstream agent component 114 via the application server cluster 120. In some embodiments, the provision request can include a request to instantiate a virtual machine. In such embodiments, the provision request refers to requests to both instantiate and provision a virtual machine. Downstream manager component 112 may be configured to receive an indication from downstream agent 114 indicating whether the provisioning was successful, and may also be configured to send a message indicating the provision result to the upstream manager 130 via jabber cluster 120.

[00023] Downstream agent component 114 of federated cluster 110 may be a software and/or hardware module(s) located at compute node 104 and may be configured to be logically connected to compute node 102 and downstream manager component 112 via the application server cluster 120. In this manner, all traffic between downstream manager component 112 and downstream agent component 114 may be sent and received via the application server cluster 120. Downstream agent component 112 may be configured to receive a provision request from downstream manager component 112 via the application server cluster 120. In some embodiments, the provision request may include a request to instantiate a virtual machine. In such embodiments, the provision request refers to requests to both instantiate and provision a virtual machine.

Downstream agent component 114 may be configured to send an indication to downstream manager component 112 indicating whether the provisioning was successful. Downstream agent component 114 may be configured to define a virtual machine in response to an instruction from downstream manager component 112. In some embodiment, defining a virtual machine may include unarchiving files, templates, etc.

[00024] According to various implementations herein, network 140 may be any type of network (e.g., a local area network (LAN), a wide area network (WAN), a virtual network, a telecommunications network) implemented as a wired network and/or wireless network. A user may communicate with system 100 via network 140. As such, in some embodiments, a user devices 145 may send data to and/or receive data from the compute node 106 using one or more communication modes (e.g., email, text messages, instant messages, optical pattern transmissions, using a mobile device application, via a website, using a personal computer (PC) application, an interactive advertisement, an ITV, TCP/IP transmissions, interactive voice response (e.g., via touch tones and/or voice recognition), etc.) that may be transmitted to the compute node 106 using a common network.

[00025] Consistent with present implementations, system 100 may be configured such that state information of federated cluster 110 is sent, using, for example, an XMPP message, to upstream manager component 130. In this manner, upstream manager component 130 may have access to much of the computing capacity of the federated cluster 110 while only updating state information when a state change within federated cluster 110 occurs. This allows federated cluster 110 to operate with less administrative overhead, making system 100 and federated cluster 110 more scalable. Furthermore, because the application server cluster 120 sees federated cluster 110 as a single downstream agent, any compute node and/or associated downstream agent of federated cluster 110 may act as a service endpoint, for example, a location to instantiate a virtual machine. In this manner, if any portion of federated cluster 110, for example, a portion of compute node 104 and/or downstream manager

114 becomes unavailable, upstream manager component 130 may send a provision request to another location, may reboot an offline virtual machine in another location, and/or may migrate a virtual machine to another location.

[00026] System 100 may be configured such that downstream manager component 112 receives commands and requests from upstream manager component 130, but need not be controlled by upstream manager component 130. In such embodiments, upstream manager component 130, and associated users, may have access to federated cluster 110's computing capacity without having complete control of federated cluster 110's compute capacity. In this manner, federated cluster 110, via downstream manager component 112, may be configured to limit, via quotas, permissions, etc., the computing capacity available to upstream manager component 130. Indeed, here, isolating control of the federated cluster 110 from upstream manager component 130 may prevent upstream manager component 130, or a user of upstream manager component 130, from monopolizing the compute capacity of federated cluster 110.

[00027] In implementations herein, system 100 may be configured such that compute node 106 may be operatively coupled to federated cluster 110 via an application server cluster 120 using lower quality links, for example, wireless area network (WAN) link, or internet links. Similarly, system 100 may be configured such that compute node 102 and compute node 104, within federated cluster 110, may be operatively coupled to federated cluster 110 via application server cluster 120 using lower quality links. System 100 may be configured such that if a link through the cluster 120 fails, either between upstream manager component 130 and downstream manager component 112, or between

downstream manager component 112 and downstream agent component 114, system 100 may failover to a new link. In some embodiments, a waiting period may pass before a new link is initiated, for example, to prevent premature failover. By way of example, upstream manager component 130 may detect that a communication link to downstream manager 112 has failed. Upstream manager component 130 may wait a predetermined amount of time, and may then establish a new link to downstream manager component 112 through jabber cluster 120. This allows system 100 to have improved reliability and fault recovery.

[00028] FIG. 2 depicts a flow chart of an illustrative method 200 of operating a system for isolating virtual machine hosts, according to an embodiment, and with reference to the system 100 depicted in FIG. 1. In FIG. 2, as with other methods herein, processing related to performance of the method steps may be performed by or among one or more of the various entities within the system. For example, the method steps or processing related thereto may be performed by one or more processing elements in the system, such as via the federated cluster 110, one or more compute nodes, including but not limited to those within the federated cluster, one or more network administrator components 150, and/or one or more management components 160. According to the exemplary implementation shown in FIG. 2, an illustrative method may comprise a first, optional step of federating a cluster of compute nodes including designating a compute nodes as the location for a downstream manager and remaining compute nodes as the location for downstream agents 202. However, such federation may occur prior to the processing described in more detail below, for example, by another entity. Method 200 may then proceed to and includes

receiving, at a first compute node, an instruction from the second compute node (e.g., via a first XMPP server) to define a virtual machine, at 204. Following the details set forth in the implementation of FIG. 1, downstream manager component 112 may receive an instruction from the upstream manager component 130 via an application server cluster 120 to define a virtual machine. An instruction is sent from the first compute node to a third compute node via the second XMPP server to define the virtual machine, at 206. Downstream manager component 112 may then send an instruction to downstream agent component 114 via the cluster 120 to define a virtual machine.

[00029] Figure 3 is a flow chart depicting another method of processing information related to isolating virtual machine hosts, according to a disclosed implementation. The processing of FIG. 3 may be performed by or among one or more of the various entities in the system 100, as explained above.

[00030] The illustrative method 300 of FIG 3 may begin with the steps of FIG 2, namely receiving, at a first compute node, an instruction from the second compute node (e.g., via a first XMPP server) to define a virtual machine, at 302. Further, an instruction may be sent from the first compute node to a third compute node via the second XMPP server to define the virtual machine, at 304. Additionally, the method 300 may comprise, at 306, receiving an indication from the downstream agent component indicating whether the provisioning was successful, and, at 308, sending a message indicating the provision result to the upstream manager component.

[00031] Figure 4 is a flow chart depicting a further method for virtual machine isolation, according to a disclosed implementation. In FIG. 4, as with other figures and processes herein, processing related to performance of the method

steps may be performed by among one or more of the various entities within the system. For example, the method steps or processing related thereto may be performed by one or more processing elements in the system, such as via the federated cluster 110, one or more compute nodes, including but not limited to those within the federated cluster, one or more network administrator components 150, and/or one or more management components 160. According to the exemplary implementation shown in FIG. 4, an illustrative method may comprise, at 402, processing information, related to defining and/or sending a provisioning request to an upstream manager component 130. Then, at 404, processing information related to defining a message, via the upstream manager component 130, including the provision request, and sending the message/provision request to the downstream manager component 112 via an application server/jabber cluster 120. Next, the method may comprise, at 406, processing information related to receipt, at the downstream manager component 112, of the provision request and/or selection of a downstream agent component 114 to host the virtual machine, as well as, at 408, processing information related to transmission of the provision request to the downstream agent component 114 via the application server/jabber cluster 120. Finally, the illustrative method of FIG. 4 may comprise, at 410, processing information related to accepting, via the downstream agent component 114, the provision request and instantiating/provisioning/configuring the virtual machine.

[00032] Figure 5 is a flow chart depicting still another method for virtual machine isolation, according to a disclosed implementation. The processing of the exemplary method of FIG. 5 may be performed by the components and/or distributed components set forth above. According to the exemplary

implementation shown in FIG. 5, another illustrative method may comprise, at 502, processing information regarding transmission, via the downstream agent component 114, of a provision result to the downstream manager component 112. Further, the method may comprise, at 504, processing information regarding polling the upstream manager component 130 for a provision result. Next, the method may comprise, at 506, processing information regarding transmission, via the upstream manager component 130, of a polling result, such as to a user 145, an administrator 150 and/or a management component 160. Then, the method may comprise processing information regarding defining a provision result and transmitting a message, via the downstream manager component 112, indicative of the provision result to the upstream manager component 130 via the application server/jabber cluster 120. Finally, at 510, the method may comprise processing information regarding storing and/or updating records, via the upstream manager component 130.

[00033] FIG. 6 depicts a flow diagram depicting an illustrative method 600 of operating a system for isolating virtual machine hosts, according to a disclosed implementation. Method 600 flows along arrow AA from BB to CC. As shown in method 600, a user accesses a network with a user device, defines a provision request (e.g. instantiating, provisioning, and/or configuring a virtual machine), and sends a provision request via HTTP or REST to the upstream manager, at 602. The upstream manager defines an XMPP message including the provision request, and sends the XMPP provision request to downstream manager via jabber cluster, at 604. In some embodiments, the upstream manager may validate the configuration, including verifying that the request does not violate any user permissions or constraints. Downstream manager then receives the

XMPP provision request and selects a downstream agent to host the virtual machine at 606, and sends the provision request to downstream agent via jabber cluster, at 608. Downstream agent may then accept the provision request and instantiate, provision, and/or configure the virtual machine, at 610. In some embodiments, instantiating, provisioning, and/or configuring the virtual machine may include, for example, unarchiving files, templates, etc. In some embodiments, downstream agent may decline the provision request. In such embodiments, downstream agent may send a decline response to downstream manager, and downstream manager may select an alternative downstream agent. Downstream agent sends a provision result to downstream manager, at 612. The user device polls the upstream manager for a provision result, at 614. The upstream device may then send a polling result to the user device. For example, the polling result may indicate that the provisioning is in-progress, has failed, or is complete, at 616. Downstream manager defines a provision result, at 618, and sends a message indicative of the provision result to the upstream manager via jabber cluster, at 620. The upstream manager stores the result of the provision request and updates user records, for example, virtual machine quotas etc., at 622. The user device may then poll the upstream manager for a provision result, at 624. Finally, the upstream device may send a polling result to the user device, for example, the polling result may indicate that the provisioning is in-progress, has failed, or is complete, at 626.

[00034] As used in this specification, the singular forms “a,” “an” and “the” include plural referents unless the context clearly dictates otherwise. Thus, for example, the term “a downstream agent” is intended to mean a single downstream agent, or a combination of downstream agents.

[00035] Some embodiments described herein relate to a computer storage product with a computer-readable medium (also may be referred to as a processor-readable medium) having instructions or computer code thereon for performing various computer-implemented operations. The media and computer code (also may be referred to as code) may be those designed and constructed for the specific purpose or purposes. Examples of computer-readable media include, but are not limited to: magnetic storage media such as hard disks, floppy disks, and magnetic tape; optical storage media such as Compact Disc/Digital Video Discs (CD/DVDs), Compact Disc-Read Only Memories (CD-ROMs), and holographic devices; magneto-optical storage media such as optical disks; and hardware devices that are specially configured to store and execute program code, such as Application-Specific Integrated Circuits (ASICs), Programmable Logic Devices (PLDs), and Read-Only Memory (ROM) and Random-Access Memory (RAM) devices.

[00036] Examples of computer code include, but are not limited to, micro-code or micro-instructions, machine instructions, such as produced by a compiler, code used to produce a web service, and files containing higher-level instructions that are executed by a computer using an interpreter. For example, embodiments may be implemented using Java, C++, or other programming languages (e.g., object-oriented programming languages) and development tools. Additional examples of computer code include, but are not limited to, control signals, encrypted code, and compressed code.

[00037] While various embodiments have been described above, it should be understood that they have been presented by way of example only, not limitation, and various changes in form and details may be made. Any portion of

the apparatus and/or methods described herein may be combined in any combination, except mutually exclusive combinations. The embodiments described herein may include various combinations and/or sub-combinations of the functions, components and/or features of the different embodiments described. For example, while FIG. 1 depicts a single federated cluster 110 having a single downstream agent 114, in some embodiments, the system 100 may include two or more federated clusters 110 each having one or more downstream agents 114.

Claims:

1. A method for processing information regarding virtual machine host isolation, the method comprising:

receiving, at a first compute node, an instruction from a second compute node via a first application server to define a virtual machine; and

sending, from the first compute node, an instruction to a third compute node via a second application server to define a virtual machine.

2. The method of claim 1 further comprising:

federating a cluster of compute nodes including the first compute node and the second compute node including designating the first compute node as a downstream manager and the second compute node and any remaining compute nodes as downstream agents, wherein the federated cluster is configured to logically appear to the second compute node designated as an upstream manager component as a single compute node.

3. The method of claim 1 further comprising, via the first compute node:

operating an agent emulation module adapted to couple the federated cluster to an upstream manager via an application server cluster or jabber cluster; and/or

presenting the federated cluster to the upstream manager as a single downstream agent including available capacity of an entire cluster of compute nodes.

4. The method of claim 1 further comprising defining, via the first compute node, a provision request associated with one of the downstream agents, and send the provision request to the one downstream agent component.

5. The method of claim 4 wherein the provision request includes a request to instantiate and/or provision a virtual machine, and further comprising, via the first compute node:

receiving an indication from the one downstream agent indicating whether the provisioning was successful; and/or

sending a message indicating the provision result to the second compute node.

6. The method of claim 2 further comprising:

transmitting state information of the federated cluster, via an XMPP message, to the second compute node, such that the second compute node has access to the federated cluster's computing capacity while only updating state information when a state change within the federated cluster occurs, wherein the federated cluster thereby operates with less administrative overhead, making the system and the federated cluster more scalable.

7. A method for processing information regarding virtual machine host isolation, the method comprising:

processing information related to defining and/or sending a provisioning request to an upstream manager component;

processing information related to defining a message, via the upstream manager component, including the provision request, and sending the message/provision request to the downstream manager component via an application server cluster or jabber cluster;

processing information related to receipt, at the downstream manager component, of the provision request and/or selection of a downstream agent component to host the virtual machine;

processing information related to transmission of the provision request to the downstream agent component via the application server cluster or jabber cluster;

and

processing information related to accepting, via the downstream agent component, the provision request and instantiating/provisioning/configuring the virtual machine.

8. The method of claim 7 wherein the instantiating/provisioning/configuring the virtual machine includes unarchiving files and/or templates.

9. The method of claim 7 further comprising processing information regarding decline, by the downstream agent component, of the provision request, wherein the downstream agent component is configured to send a decline response to the downstream manager component, and the downstream manager component is configured to select an alternative downstream agent.

10. The method of claim 7 further comprising processing information regarding validating, via the upstream manager component, the configuration, including verifying that the request does not violate any permissions or constraints.

11. The method of claim 7 further comprising:

processing information regarding transmission, via the downstream agent component, of a provision result to the downstream manager component;

processing information regarding polling the upstream manager component for a provision result;

processing information regarding transmission, via the upstream manager component, of a polling result to a user or a management component;

processing information regarding defining a provision result and transmitting a message, via the downstream manager component, indicative of the provision result to the upstream manager component via the application server cluster or jabber cluster;

processing information regarding storing and/or updating records, via the upstream manager component.

12. The method of claim 11 further comprising processing information regarding:

polling, via the user or the management component, the upstream manager component for a provision result; and

transmission, via the upstream manager component, a polling result to the user or the management component.

13. A system comprising:

a first compute node;

a second compute node operatively coupled to the first compute node via a first application server; and

a third compute node operatively coupled to the first compute node via a second application server;

wherein the first compute node is included within a federated cluster that includes the third compute node, and configured: to receive an instruction from the second compute node via the first application server to define a virtual machine, and to send an instruction to the third compute node via the second application server to define the virtual machine.

14. The system of claim 13 wherein the first application server and the second application server are extensible messaging and presence protocol (XMPP) servers.

15. A system comprising:

a federated cluster component including:

one or more compute nodes including a first compute node, which is connected to a second compute node outside the federated cluster, and a third compute node;

a downstream management component associated with a first compute node; and

a downstream agent component associated with a second compute node;

wherein the second compute node is operatively coupled to the first compute node via one or more application servers;

wherein the first compute node is included within a federated cluster that includes the third compute node, and configured: to receive an instruction from the second compute node via a first server to define a virtual machine, and to send an instruction to the third compute node via a second server to define the virtual machine.

16. A system comprising:

a federated cluster including one or more compute nodes, a downstream management component associated with a first compute node, and a downstream agent component associated with a second compute node;

an application server cluster configured with a messaging or communication protocol that supports communication between upstream and downstream manager-agent components;

a first compute node;

a second compute node operatively coupled to the first compute node via a first application server; and

a third compute node operatively coupled to the first compute node via a second application server;

wherein the first compute node is included within a federated cluster that includes the third compute node, and configured: to receive an instruction from the second compute node via the first application server to define a virtual machine, and to send an instruction to the third compute node via the second application server to define the virtual machine.

17. The system of claim 15 or claim 16 wherein the federated cluster includes a cluster of compute nodes or multiple cooperating compute nodes that are not centrally managed, and wherein one of the compute nodes is designated as a location for the downstream manager component and any remaining compute nodes as the location for downstream agents.

18. The system of claim 15 or claim 16 wherein a compute node within the federated cluster is designated as the location for the downstream manager and is configured to:

operate an agent emulation module adapted to couple the federated cluster to an upstream manager via an application server cluster or jabber cluster; and/or

present the federated cluster to the upstream manager as a single downstream agent including available capacity of an entire cluster of compute nodes.

19. The system of claim 15 or claim 16 wherein the messaging or communication protocol is configured to require low processing and memory resources, is standardized, is customizable, is point-to-point, and/or is configured to send and receive state and/or provisioning messages.

20. The system of claim 15 or claim 16 wherein the application server cluster includes an XMPP application server including one or more extensible messaging and presence protocol (XMPP) servers logically connected and configured to support XMPP communication(s).

21. The system of claim 20 wherein the XMPP application server comprises at least one jabber cluster, and wherein the one or more XMPP servers are configured to support XMPP communication between an upstream manager component and a downstream manager component, and between the downstream manager component and a downstream agent component.

22. The system of claim 20 wherein the XMPP application server comprises two or more jabber clusters, including a first jabber cluster having one or more servers logically connected and configured to support communication between an upstream manager component and a downstream manager component, and a second jabber cluster having one or more servers logically connected and configured to support communication between the downstream manager component and a downstream agent component, such that communication between the downstream manager component and the downstream agent component is secured separately from communication between the upstream manager component and the downstream manager component and is configured to continue in the event of a failure of any portion of the first jabber cluster.

23. The system of claim 15 or claim 16 wherein the downstream manager component is configured to define a provision request associated with the downstream agent component, and send the provision request to the downstream agent component.

24. The system of claim 23 wherein the provision request includes a request to instantiate and/or provision a virtual machine, and the downstream manager component is configured to:

receive an indication from the downstream agent component indicating whether the provisioning was successful; and/or

send a message indicating the provision result to the upstream manager component.

25. The system of claim 15 or claim 16 wherein the system is configured such that state information of the federated cluster is sent, such as via an XMPP message, to upstream manager component, such that the upstream manager component has access to the federated cluster's computing capacity while only updating state information when a state change within the federated cluster occurs, wherein the federated cluster thereby operates with less administrative overhead, making the system and the federated cluster more scalable.

26. The system of claim 15 or claim 16 wherein the system is configured so that the jabber cluster addresses the federated cluster as a single downstream agent, such that any compute node and/or associated downstream agent of the federated cluster can act as a service endpoint, wherein, if any portion of federated cluster becomes unavailable, the upstream manager component is configured to send a provision request to another location, reboot an offline virtual machine in another location, and/or migrate a virtual machine to another location.

27. The system of claim 15 or claim 16 wherein the system is configured in that the downstream manager component receives commands and requests from the upstream manager component, but is not controlled by the upstream manager component, wherein the upstream manager component and associated users have access to the federated cluster's computing capacity without having complete control of the federated cluster's compute capacity.

28. The system of claim 27 wherein the federated cluster, via the downstream manager component, is configured to limit, such as via quotas and/or permissions, computing capacity available to the upstream manager component, providing isolating control of the federated cluster from the upstream manager component and thereby preventing the upstream manager component, or a user of the upstream manager component, from monopolizing the compute capacity of the federated cluster.

29. The invention of any claim herein wherein the upstream manager component is configured to:

receive provision results from the downstream manager component, store those results, and/or send status updates and/or results to the user via network;
and/or

logically interact with the federated cluster as a single compute node.

30. The invention of any claim herein wherein the upstream manager component is configured to process information to treat a first federated cluster including a single downstream manager and a one or more downstream agents as a single compute

node having a first amount of computing capacity, and treat a second federated cluster including a single downstream manager and ten downstream agents as a single compute node have a second amount of computing capacity distinct from the other compute node, such that the upstream manager component need not store state and other information for each individual downstream agent, but instead may only store state and other information for each federated cluster.

31. The invention of any claim herein wherein the upstream manager component is configured to include customer data, such as available capacity, virtual machine quotas, and the like, such that the upstream manager component is configured to accept or decline a request to provision a virtual machine, and/or can determine provisioning option based on that customer information.

32. The invention of any claim herein wherein the upstream manager component is configured to:

receive hypertext transport protocol (HTTP) instructions from a device, a user, a customer, a network administrator, and/or other management component; and/or
define an XMPP instruction(s) to the downstream manager component via the application server cluster and/or the jabber cluster.

33. The invention of any claim herein wherein a compute node within the federated cluster is designated as the location for the downstream manager and is configured to:

operate an agent emulation module adapted to couple the federated cluster to an upstream manager via an application server/jabber cluster; and/or

present the federated cluster to the upstream manager as a single downstream agent including available capacity of an entire cluster of compute nodes.

AMENDED CLAIMS

received by the International Bureau on 29 July 2013 (29.07.2013)

1. A method for processing information regarding virtual machine host isolation, the method comprising:

federating a cluster of compute nodes including a first compute node and a third compute node;

receiving, at the first compute node, an instruction from a second compute node via a first application server to define a virtual machine; and

sending, from the first compute node, an instruction to the third compute node via a second application server to define the virtual machine, wherein the federated cluster is configured to logically appear to the second compute node as a single compute node.

2. The method of claim 1 wherein:

the first compute node acts as a downstream manager, the second compute node and any remaining compute nodes act as downstream agents and the second compute node acts as an upstream agent.

3. The method of claim 1 further comprising, via the first compute node:

operating an agent emulation module adapted to couple the federated cluster to the second compute node via an application server cluster or jabber cluster; and/or

presenting the federated cluster to the second compute node as a single downstream agent including available capacity of an entire cluster of compute nodes.

4. The method of claim 1 further comprising defining, via the first compute node, a provision request associated with the third compute node, and send the provision request to the third compute node.

5. The method of claim 4 wherein the provision request includes a request to instantiate and/or provision a virtual machine, and further comprising, via the first compute node:

receiving an indication from the third compute node indicating whether the provisioning was successful; and/or

sending a message indicating the provision result to the second compute node.

6. The method of claim 1 further comprising:

transmitting state information of the federated cluster, via an XMPP message, to the second compute node, such that the second compute node has access to the federated cluster's computing capacity while only updating state information when a state change within the federated cluster occurs, wherein the federated cluster thereby operates with less administrative overhead, making the system and the federated cluster more scalable.

7. A method for processing information regarding virtual machine host isolation, the method comprising:

processing information related to defining and/or sending a provisioning request to an upstream manager component;

processing information related to defining a message, via the upstream manager component, including the provision request, and sending the message/provision request to the downstream manager component via an application server cluster or jabber cluster;

processing information related to receipt, at the downstream manager component, of the provision request and/or selection of a downstream agent component to host the virtual machine;

processing information related to transmission of the provision request to the downstream agent component via the application server cluster or jabber cluster; and

processing information related to accepting, via the downstream agent component, the provision request and instantiating/provisioning/configuring the virtual machine.

8. The method of claim 7 wherein the instantiating/provisioning/configuring the virtual machine includes unarchiving files and/or templates.

9. The method of claim 7 further comprising processing information regarding decline, by the downstream agent component, of the provision request, wherein the downstream agent component is configured to send a decline response to the downstream manager component, and the downstream manager component is configured to select an alternative downstream agent.

10. The method of claim 7 further comprising processing information regarding validating, via the upstream manager component, the configuration, including verifying that the request does not violate any permissions or constraints.

11. The method of claim 7 further comprising:

processing information regarding transmission, via the downstream agent component, of a provision result to the downstream manager component;

processing information regarding polling the upstream manager component for a provision result;

processing information regarding transmission, via the upstream manager component, of a polling result to a user or a management component;

processing information regarding defining a provision result and transmitting a message, via the downstream manager component, indicative of the provision result to the upstream manager component via the application server cluster or jabber cluster;

processing information regarding storing and/or updating records, via the upstream manager component.

12. The method of claim 11 further comprising processing information regarding:

polling, via the user or the management component, the upstream manager component for a provision result; and

transmission, via the upstream manager component, a polling result to the user or the management component.

13. A system comprising:

a first compute node;

a second compute node operatively coupled to the first compute node via a first application server; and

a third compute node operatively coupled to the first compute node via a second application server;

wherein the first compute node is included within a federated cluster that includes the third compute node, and configured: to receive an instruction from the second compute node via the first application server to define a virtual machine, and to send an instruction to the third compute node via the second application server to define the virtual machine, wherein the federated cluster is configured to logically appear to the second compute node as a single compute node.

14. The system of claim 13 wherein the first application server and the second application server are extensible messaging and presence protocol (XMPP) servers.

15. A system comprising:

a federated cluster component including:

one or more compute nodes including a first compute node, which is connected to a second compute node outside the federated cluster, and a third compute node;

a downstream management component associated with a first compute node; and

a downstream agent component associated with a third compute node;

wherein the second compute node is operatively coupled to the first compute node via one or more application servers;

wherein the first compute node is included within a federated cluster that includes the third compute node, and configured: to receive an instruction from the second compute node via a first server to define a virtual machine, and to send an instruction to the third compute node via a second server to define the virtual machine, and wherein the federated cluster is configured to logically appear to the second compute node as a single compute node.

16. A system comprising:

a federated cluster including one or more compute nodes, a downstream management component associated with a first compute node, and a downstream agent component associated with a third compute node;

an application server cluster configured with a messaging or communication protocol that supports communication between upstream and downstream manager-agent components;

a second compute node operatively coupled to the first compute node via a first application server; and

wherein the first compute node is configured: to receive an instruction from the second compute node via the first application server to define a virtual machine, and to send an instruction to the third compute node via the sec, and wherein the federated cluster is configured to logically appear to the second compute node as a single compute node. second application server to define the virtual machine

17. The system of claim 15 or claim 16 wherein the federated cluster includes a cluster of compute nodes or multiple cooperating compute nodes that are not centrally managed, and wherein one of the compute nodes is designated as a

location for the downstream manager component and any remaining compute nodes as the location for downstream agents.

18. The system of claim 15 or claim 16 wherein a compute node within the federated cluster is designated as the location for the downstream manager and is configured to:

operate an agent emulation module adapted to couple the federated cluster to an upstream manager via an application server cluster or jabber cluster; and/or

present the federated cluster to the upstream manager as a single downstream agent including available capacity of an entire cluster of compute nodes.

19. The system of claim 15 or claim 16 wherein the messaging or communication protocol is configured to require low processing and memory resources, is standardized, is customizable, is point-to-point, and/or is configured to send and receive state and/or provisioning messages.

20. The system of claim 15 or claim 16 wherein the application server cluster includes an XMPP application server including one or more extensible messaging and presence protocol (XMPP) servers logically connected and configured to support XMPP communication(s).

21. The system of claim 20 wherein the XMPP application server comprises at least one jabber cluster, and wherein the one or more XMPP servers are configured to support XMPP communication between an upstream manager component and a

downstream manager component, and between the downstream manager component and a downstream agent component.

22. The system of claim 20 wherein the XMPP application server comprises two or more jabber clusters, including a first jabber cluster having one or more servers logically connected and configured to support communication between an upstream manager component and a downstream manager component, and a second jabber cluster having one or more servers logically connected and configured to support communication between the downstream manager component and a downstream agent component, such that communication between the downstream manager component and the downstream agent component is secured separately from communication between the upstream manager component and the downstream manager component and is configured to continue in the event of a failure of any portion of the first jabber cluster.

23. The system of claim 15 or claim 16 wherein the downstream manager component is configured to define a provision request associated with the downstream agent component, and send the provision request to the downstream agent component.

24. The system of claim 23 wherein the provision request includes a request to instantiate and/or provision a virtual machine, and the downstream manager component is configured to:

receive an indication from the downstream agent component indicating whether the provisioning was successful; and/or

send a message indicating the provision result to the upstream manager component.

25. The system of claim 15 or claim 16 wherein the system is configured such that state information of the federated cluster is sent, such as via an XMPP message, to upstream manager component, such that the upstream manager component has access to the federated cluster's computing capacity while only updating state information when a state change within the federated cluster occurs, wherein the federated cluster thereby operates with less administrative overhead, making the system and the federated cluster more scalable.

26. The system of claim 15 or claim 16 wherein the system is configured so that the jabber cluster addresses the federated cluster as a single downstream agent, such that any compute node and/or associated downstream agent of the federated cluster can act as a service endpoint, wherein, if any portion of federated cluster becomes unavailable, the upstream manager component is configured to send a provision request to another location, reboot an offline virtual machine in another location, and/or migrate a virtual machine to another location.

27. The system of claim 15 or claim 16 wherein the system is configured in that the downstream manager component receives commands and requests from the upstream manager component, but is not controlled by the upstream manager component, wherein the upstream manager component and associated users have access to the federated cluster's computing capacity without having complete control of the federated cluster's compute capacity.

28. The system of claim 27 wherein the federated cluster, via the downstream manager component, is configured to limit, such as via quotas and/or permissions, computing capacity available to the upstream manager component, providing isolating control of the federated cluster from the upstream manager component and thereby preventing the upstream manager component, or a user of the upstream manager component, from monopolizing the compute capacity of the federated cluster.

29. The invention of any claim herein wherein the upstream manager component is configured to:

receive provision results from the downstream manager component, store those results, and/or send status updates and/or results to the user via network;
and/or

logically interact with the federated cluster as a single compute node.

30. The invention of any claim herein wherein the upstream manager component is configured to process information to treat a first federated cluster including a single downstream manager and a one or more downstream agents as a single compute node having a first amount of computing capacity, and treat a second federated cluster including a single downstream manager and ten downstream agents as a single compute node have a second amount of computing capacity distinct from the other compute node, such that the upstream manager component need not store state and other information for each individual downstream agent, but instead may only store state and other information for each federated cluster.

31. The invention of any claim herein wherein the upstream manager component is configured to include customer data, such as available capacity, virtual machine quotas, and the like, such that the upstream manager component is configured to accept or decline a request to provision a virtual machine, and/or can determine provisioning option based on that customer information.

32. The invention of any claim herein wherein the upstream manager component is configured to:

receive hypertext transport protocol (HTTP) instructions from a device, a user, a customer, a network administrator, and/or other management component; and/or

define an XMPP instruction(s) to the downstream manager component via the application server cluster and/or the jabber cluster.

33. The invention of any claim herein wherein a compute node within the federated cluster is designated as the location for the downstream manager and is configured to:

operate an agent emulation module adapted to couple the federated cluster to an upstream manager via an application server/jabber cluster; and/or

present the federated cluster to the upstream manager as a single downstream agent including available capacity of an entire cluster of compute nodes.

200

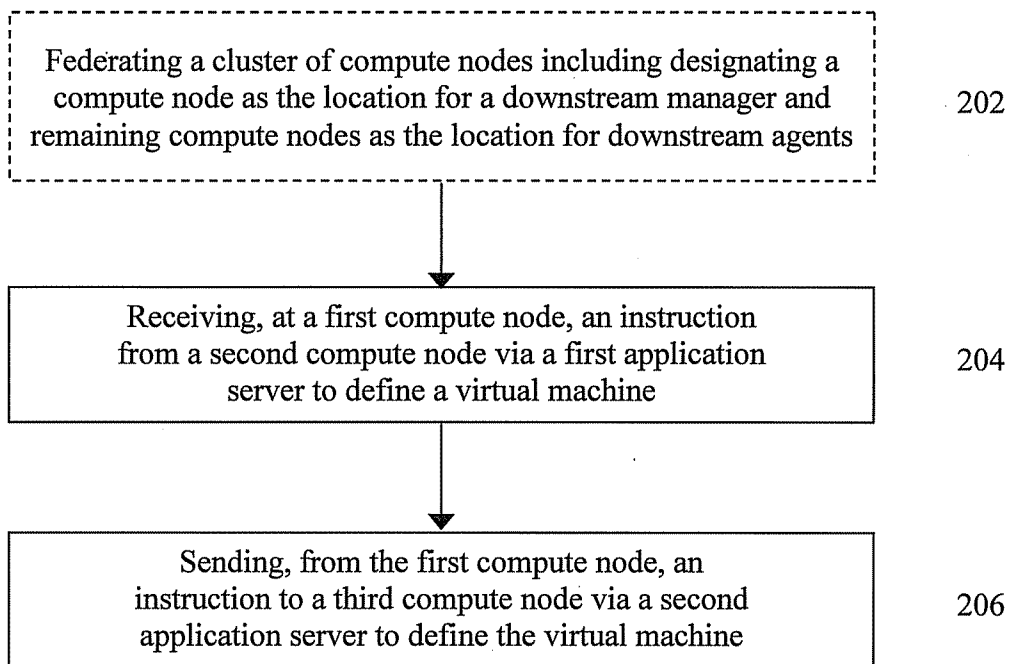


FIG. 2

3/6

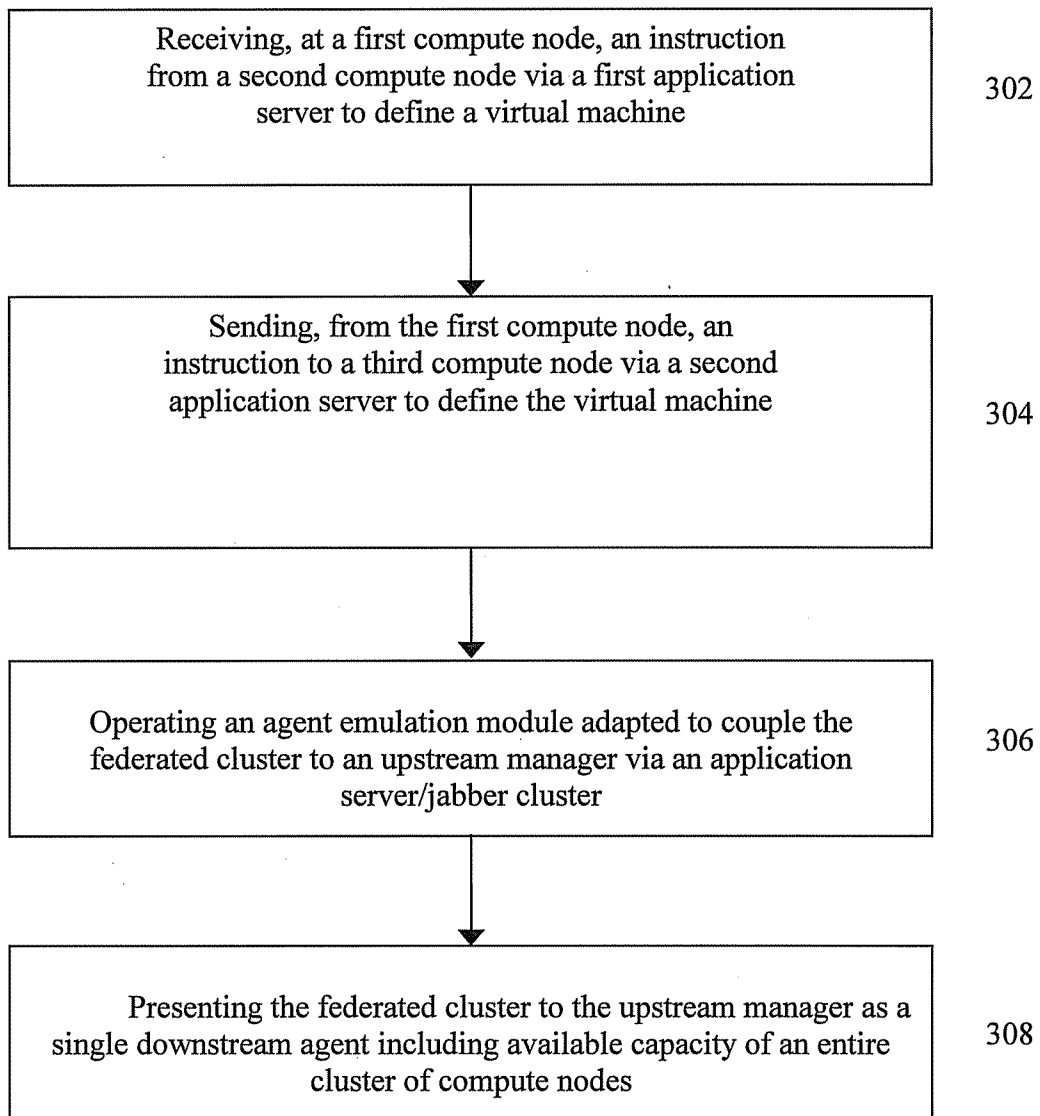
300

FIG. 3

400

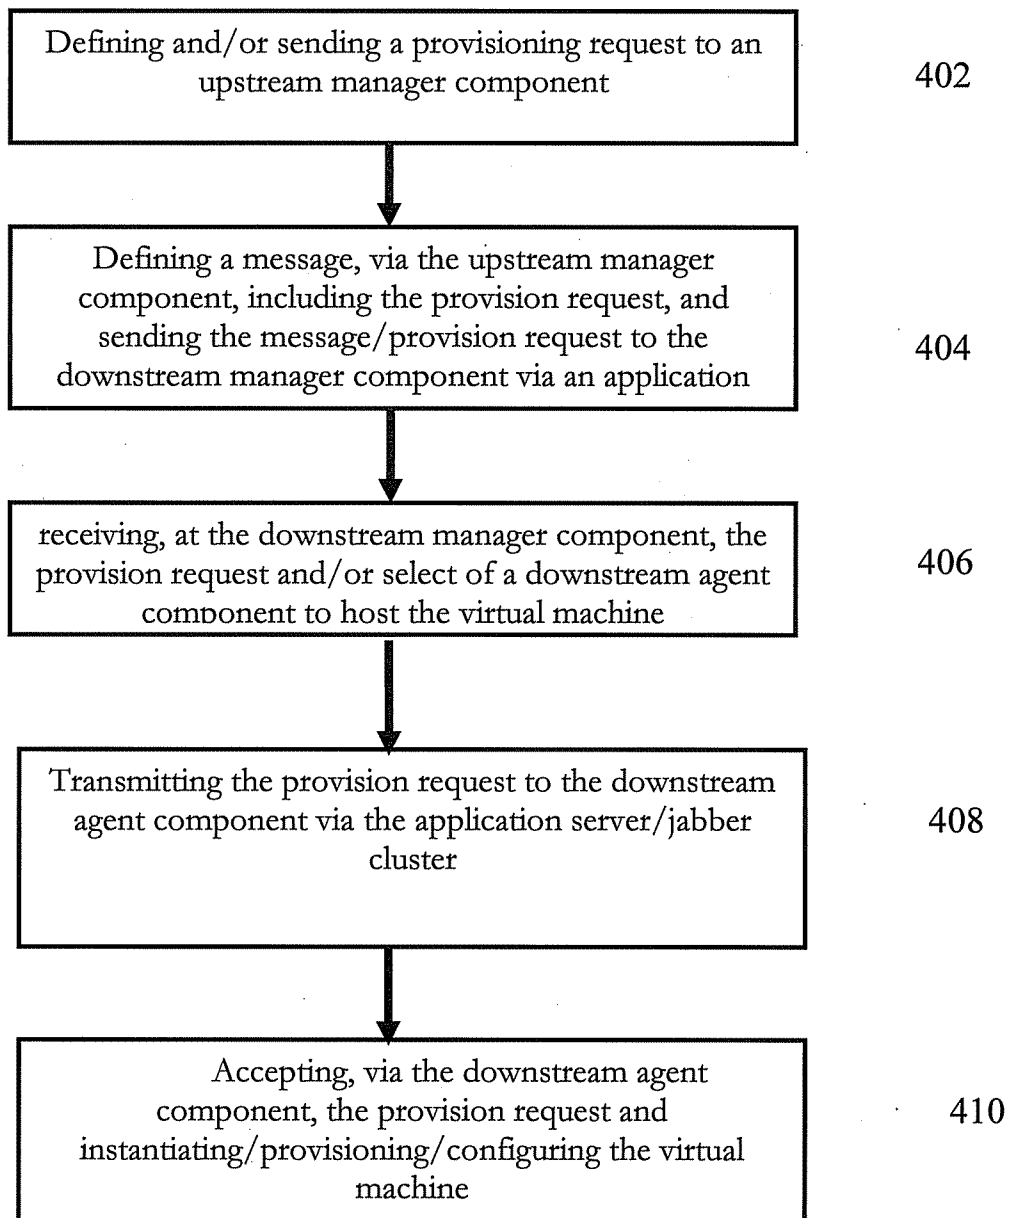


FIG 4

500

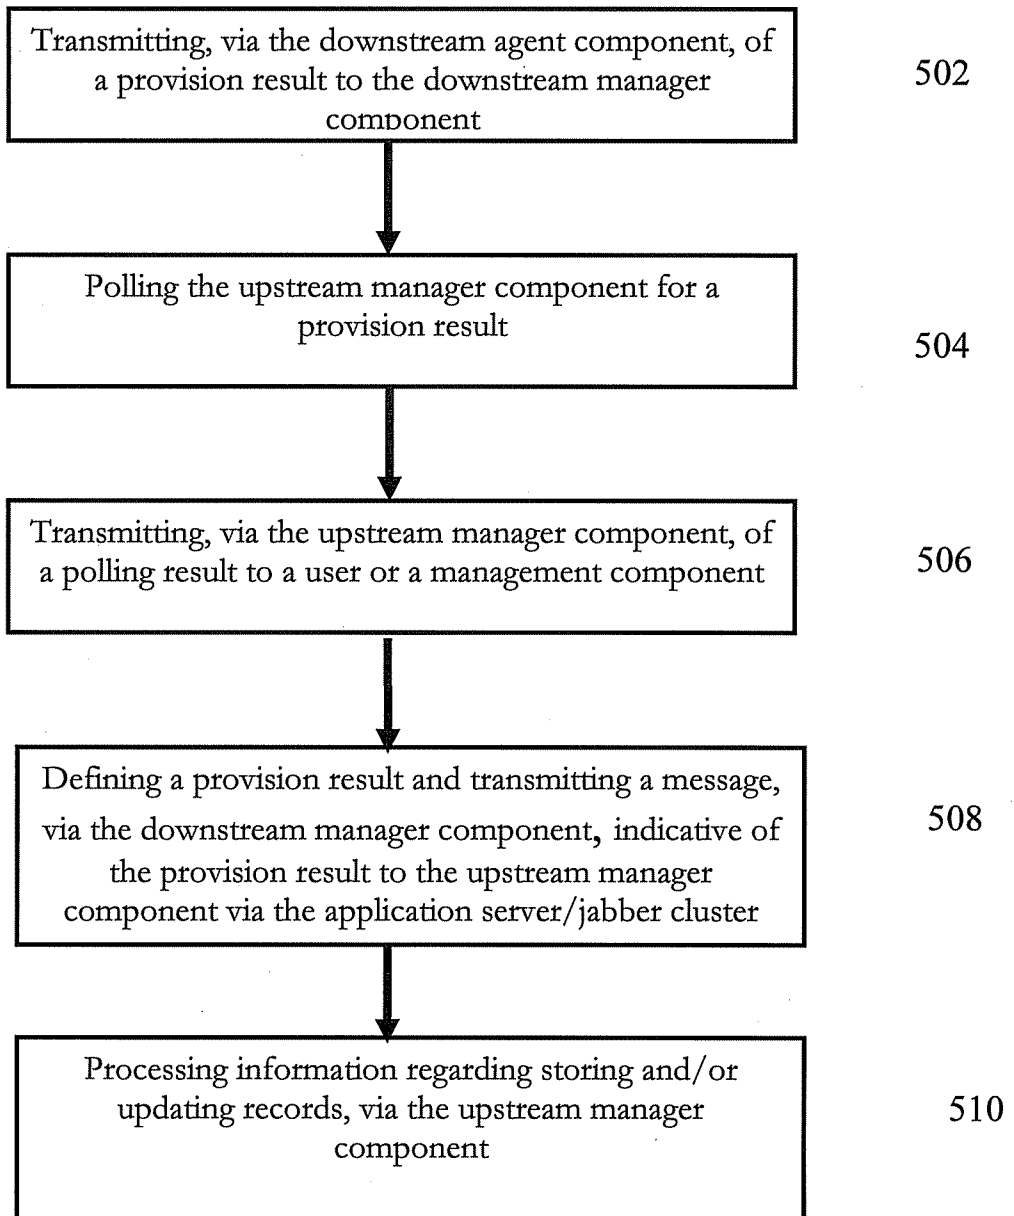


FIG 5




FIG 6

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2013/026907**A. CLASSIFICATION OF SUBJECT MATTER****G06F 15/16(2006.01)i, G06F 9/06(2006.01)i**

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

B. FIELDS SEARCHEDMinimum documentation searched (classification system followed by classification symbols)
G06F 15/16Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean utility models and applications for utility models
Japanese utility models and applications for utility modelsElectronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & Keywords: cluster , first node<or>second node , server, virtual machine, network**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X A A	US 2004-0006587 A1 (DANIEL, RAYMOND MCCONNELL et al.) 08 January 2004 See paragraph 13 - paragraph 23; claims 1-3,10-12 and figures 1-4.	1, 13, 15, 16 2-12, 14, 17-33
A	US 2007-0055781 A1 (CHRISTIAN FLEISCHER et al.) 08 March 2007 See paragraph 90 - paragraph 103; paragraph 129 - paragraph 133; and figures 5,8,9.	1-33
A	US 7890689 B2 (LAM MONICA SIN-LING et al.) 15 February 2011 See column 2, line 28 - column 3, line 62; column 12, line 55 - column 13, line 49; and figures 9,10.	1-33
A	US 7653682 B2 (ERASANI PRANCOOP et al.) 26 January 2010 See column 5, line 5 - column 8, line 31; and claim 1, and figures 1-2.	1-33
A	US 2011-0066672 A1 (ZAMARREO GALDER et al.) 17 March 2011 See paragraph 21 - paragraph 34; and figures 1-4.	1-33

 Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

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

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

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

"&" document member of the same patent family


Date of the actual completion of the international search

30 May 2013 (30.05.2013)

Date of mailing of the international search report

03 June 2013 (03.06.2013)

Name and mailing address of the ISA/KR


 Korean Intellectual Property Office
189 Cheongsu-ro, Seo-gu, Daejeon Metropolitan City,
302-701, Republic of Korea

Facsimile No. 82-42-472-7140

Authorized officer

PARK, Sang Hyun

Telephone No. 82-42-481-8263



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2013/026907

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2004-0006587 A1	08.01.2004	US 2006-0059226 A1	16.03.2006
US 2007-0055781 A1	08.03.2007	US 2011-179133 A1 US 7945677 B2 US 8161169 B2	21.07.2011 17.05.2011 17.04.2012
US 7890689 B2	15.02.2011	US 2005-0125513 A1 US 2008-0215796 A1 US 7373451 B2	09.06.2005 04.09.2008 13.05.2008
US 7653682 B2	26.01.2010	AT471624 T DE 602006014971 D11 EP 1908261 A1 EP 1908261 B1 US 2007-0022138 A1 WO 2007-014079 A1	15.07.2010 29.07.2010 09.04.2008 16.06.2010 25.01.2007 01.02.2007
US 2011-0066672 A1	17.03.2011	None	