



(51) International Patent Classification:

H04L 67/1031 (2022.01) H04L 47/70 (2022.01)

G06F 16/27 (2006.01) H04L 67/60 (2022.01)

G06Q 10/06 (2012.01)

(21) International Application Number:

PCT/US2022/027250

(22) International Filing Date:

02 May 2022 (02.05.2022)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

63/186,906

11 May 2021 (11.05.2021)

US

(71) Applicant: JPMORGAN CHASE BANK, N.A. [US/US];

383 Madison Avenue, New York, New York 10179 (US).

(72) Inventors: SLAVIN, Ilya; c/o JPMorgan Chase Bank,

N.A., 383 Madison Avenue, New York, New York 10179

(US). LEGGE, Mathew Alistair; c/o JPMorgan Chase

Bank, N.A., 383 Madison Avenue, New York, New York

10179 (US). XIE, Xiaolin; c/o JPMorgan Chase Bank,

N.A., 383 Madison Avenue, New York, New York 10179

(US). WILLIG, Stuart; c/o JPMorgan Chase Bank, N.A.,

383 Madison Avenue, New York, New York 10179 (US).

MORRISON, Graeme; c/o JPMorgan Chase Bank, N.A.,
383 Madison Avenue, New York, New York 10179 (US).(74) Agent: MUELLER, Robert W. et al.; GREENBLUM &
BERNSTEIN, P.L.C., 1950 Roland Clarke Place, Reston,
Virginia 20191 (US).(81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ,
CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO,
DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN,
HR, HU, ID, IL, IN, IQ, IR, IS, IT, JM, JO, JP, KE, KG, KH,
KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA,
MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI,
NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU,
RW, SA, SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM,
TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM,
ZW.(84) Designated States (unless otherwise indicated, for every
kind of regional protection available): ARIPO (BW, GH,
GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ,
UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ,
TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK,
EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,
MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,

(54) Title: METHOD AND SYSTEM FOR HOSTING PLATFORM IN PUBLIC CLOUD COMPUTING ENVIRONMENT

400




FIG. 4

(57) Abstract: A method and a system for hosting a market data distribution platform that is widely used by many financial organizations, such as Thomson Reuters Enterprise Platform (TREP), in a public cloud computing environment is provided. The method includes facilitating a communication between the market data distribution platform and at least one server that operates in the public cloud; monitoring an amount of data being distributed on the platform and a number of clients that are accessing the platform; determining a minimum number of cloud servers required for servicing the clients; accessing at least the determined minimum number of servers; and facilitating communications between the servers and the platform. The number of servers and/or load amounts being handled by each server may be adjusted based on varying amounts of data and/or varying numbers of clients.

TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

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

Published:

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

METHOD AND SYSTEM FOR HOSTING PLATFORM IN PUBLIC CLOUD COMPUTING ENVIRONMENT

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims priority from U.S. Provisional Patent Application No. 63/186,906, filed in the U.S. Patent and Trademark Office on May 11, 2021, which is hereby incorporated by reference in its entirety.

BACKGROUND

1. Field of the Disclosure

[0002] This technology generally relates to methods and systems for hosting a platform in a public cloud computing environment, and more particularly to methods and systems for hosting a market data distribution platform that is widely used by many financial organizations in a public cloud computing environment.

2. Background Information

[0003] Thomson Reuters Enterprise Platform (TREP) is an incumbent product that is commonly used by financial organizations of various sizes to distribute market data to a wide range of consumers. This software ecosystem has evolved through several decades of acquisitions and is easily found in every large financial firm and in many medium-sized financial firms.

[0004] The desire of such organizations to move toward operating on modern computing infrastructure is incompatible with requirements that relate to operating TREP on large volumes of data. Further, attempts to expand TREP to a public cloud computing environment have been limited in scope and slow to develop. A lack of progress of developing a cloud-first data approach has given rise to a need to provide an alternate hosting arrangement for TREP.

[0005] Accordingly, there is a need for a mechanism for hosting TREP in a public cloud computing environment that is broadly accessible to financial organizations.

SUMMARY

[0006] The present disclosure, through one or more of its various aspects, embodiments, and/or specific features or sub-components, provides, inter alia, various systems, servers, devices, methods, media, programs, and platforms for hosting a market data distribution platform that is widely used by many financial organizations in a public cloud computing environment.

[0007] According to an aspect of the present disclosure, a method for hosting a market data distribution platform in a public cloud computing environment is provided. The method is implemented by at least one processor. The method includes: accessing, by the at least one processor, a first cloud server that operates in a public cloud computing environment; facilitating, by the at least one processor, a communication between a market data distribution platform and the first cloud server; monitoring, by the at least one processor, an amount of data being distributed on the market data distribution platform and a number of clients that are accessing the market data distribution platform; determining, by the at least one processor based on a result of the monitoring, a minimum number of cloud servers required for servicing the clients; accessing, by the at least one processor, a plurality of cloud servers that includes at least the determined minimum number of cloud servers; facilitating, by the at least one processor, respective communications between the market data distribution platform and each of the plurality of cloud servers; and shifting, by the at least one processor based on a result of the monitoring, a portion of a load from at least a first one of the plurality of cloud servers to at least a second one of the plurality of cloud servers.

[0008] When the minimum number of cloud servers required for servicing the clients is reduced based on a subsequent result of the monitoring, the method may further include reducing a number of cloud servers that are accessed based on the reduced minimum number of cloud servers and shifting a portion of a load from at least a first one of the reduced minimum number of cloud servers to at least a second one of the reduced minimum number of cloud servers.

[0009] When the minimum number of cloud servers required for servicing the clients is increased based on a subsequent result of the monitoring, the method may further include increasing a number of cloud servers that are accessed based on the increased minimum number of cloud servers and shifting a portion of a load from at

least a first one of the increased minimum number of cloud servers to at least a second one of the increased minimum number of cloud servers.

[0010] The facilitating of the respective communications between the market data distribution platform and each of the plurality of cloud servers may include using a multicast communication protocol.

[0011] The market data distribution platform may include a Thomson Reuters Enterprise Platform (TREP).

[0012] The monitoring may be performed continuously.

[0013] A result of the determining of the minimum number of cloud servers required for servicing the clients may increase as a result of a beginning of a trading day and may decrease as a result of an end of the trading day.

[0014] The public cloud computing environment may be administered by Amazon Web Services (AWS).

[0015] According to another aspect of the present disclosure, a computing apparatus for hosting a market data distribution platform in a public cloud computing environment is provided. The computing apparatus includes a processor; a memory; and a communication interface coupled to each of the processor and the memory. The processor is configured to: access a first cloud server that operates in a public cloud computing environment; facilitate a communication between a market data distribution platform and the first cloud server; monitor an amount of data being distributed on the market data distribution platform and a number of clients that are accessing the market data distribution platform; determine, based on a result of the monitoring, a minimum number of cloud servers required for servicing the clients; access a plurality of cloud servers that includes at least the determined minimum number of cloud servers; facilitate respective communications between the market data distribution platform and each of the plurality of cloud servers; and shift, based on a result of the monitoring, a portion of a load from at least a first one of the plurality of cloud servers to at least a second one of the plurality of cloud servers.

[0016] When the minimum number of cloud servers required for servicing the clients is reduced based on a subsequent result of the monitoring, the processor may be further configured to reduce a number of cloud servers that are accessed based on the reduced minimum number of cloud servers and to shift a portion of a load from at

least a first one of the reduced minimum number of cloud servers to at least a second one of the reduced minimum number of cloud servers.

[0017] When the minimum number of cloud servers required for servicing the clients is increased based on a subsequent result of the monitoring, the processor may be further configured to increase a number of cloud servers that are accessed based on the increased minimum number of cloud servers and to shift a portion of a load from at least a first one of the increased minimum number of cloud servers to at least a second one of the increased minimum number of cloud servers.

[0018] The processor may be further configured to use a multicast communication protocol for performing the facilitating of the respective communications between the market data distribution platform and each of the plurality of cloud servers.

[0019] The market data distribution platform may include a Thomson Reuters Enterprise Platform (TREP).

[0020] The processor may be further configured to continuously perform the monitoring of the amount of data being distributed on the market data distribution platform and the number of clients that are accessing the market data distribution platform.

[0021] A result of the determination of the minimum number of cloud servers required for servicing the clients may increase as a result of a beginning of a trading day and may decrease as a result of an end of the trading day.

[0022] The public cloud computing environment may be administered by Amazon Web Services (AWS).

[0023] According to yet another aspect of the present disclosure, a non-transitory computer readable storage medium storing instructions for hosting a market data distribution platform in a public cloud computing environment is provided. The storage medium includes executable code which, when executed by a processor, causes the processor to: access a first cloud server that operates in a public cloud computing environment; facilitate a communication between a market data distribution platform and the first cloud server; monitor an amount of data being distributed on the market data distribution platform and a number of clients that are accessing the market data distribution platform; determine, based on a result of the

monitoring, a minimum number of cloud servers required for servicing the clients; access a plurality of cloud servers that includes at least the determined minimum number of cloud servers; facilitate respective communications between the market data distribution platform and each of the plurality of cloud servers; and shift, based on a result of the monitoring, a portion of a load from at least a first one of the plurality of cloud servers to at least a second one of the plurality of cloud servers.

[0024] When the minimum number of cloud servers required for servicing the clients is reduced based on a subsequent result of the monitoring, the executable code may further cause the processor to reduce a number of cloud servers that are accessed based on the reduced minimum number of cloud servers and to shift a portion of a load from at least a first one of the reduced minimum number of cloud servers to at least a second one of the reduced minimum number of cloud servers.

[0025] When the minimum number of cloud servers required for servicing the clients is increased based on a subsequent result of the monitoring, the executable code may further cause the processor to increase a number of cloud servers that are accessed based on the increased minimum number of cloud servers and to shift a portion of a load from at least a first one of the increased minimum number of cloud servers to at least a second one of the increased minimum number of cloud servers.

[0026] The executable code may further cause the processor to use a multicast communication protocol for the facilitating of the respective communications between the market data distribution platform and each of the plurality of cloud servers.

BRIEF DESCRIPTION OF THE DRAWINGS

[0027] The present disclosure is further described in the detailed description which follows, in reference to the noted plurality of drawings, by way of non-limiting examples of preferred embodiments of the present disclosure, in which like characters represent like elements throughout the several views of the drawings.

[0028] FIG. 1 illustrates an exemplary computer system.

[0029] FIG. 2 illustrates an exemplary diagram of a network environment.

[0030] FIG. 3 shows an exemplary system for implementing a method for hosting a market data distribution platform that is widely used by many financial organizations in a public cloud computing environment.

[0031] FIG. 4 is a flowchart of an exemplary process for implementing a method for hosting a market data distribution platform that is widely used by many financial organizations in a public cloud computing environment.

DETAILED DESCRIPTION

[0032] Through one or more of its various aspects, embodiments and/or specific features or sub-components of the present disclosure, are intended to bring out one or more of the advantages as specifically described above and noted below.

[0033] The examples may also be embodied as one or more non-transitory computer readable media having instructions stored thereon for one or more aspects of the present technology as described and illustrated by way of the examples herein. The instructions in some examples include executable code that, when executed by one or more processors, cause the processors to carry out steps necessary to implement the methods of the examples of this technology that are described and illustrated herein.

[0034] FIG. 1 is an exemplary system for use in accordance with the embodiments described herein. The system 100 is generally shown and may include a computer system 102, which is generally indicated.

[0035] The computer system 102 may include a set of instructions that can be executed to cause the computer system 102 to perform any one or more of the methods or computer-based functions disclosed herein, either alone or in combination with the other described devices. The computer system 102 may operate as a standalone device or may be connected to other systems or peripheral devices. For example, the computer system 102 may include, or be included within, any one or more computers, servers, systems, communication networks or cloud environment. Even further, the instructions may be operative in such cloud-based computing environment.

[0036] In a networked deployment, the computer system 102 may operate in the capacity of a server or as a client user computer in a server-client user network environment, a client user computer in a cloud computing environment, or as a peer computer system in a peer-to-peer (or distributed) network environment. The computer system 102, or portions thereof, may be implemented as, or incorporated into, various devices, such as a personal computer, a tablet computer, a set-top box, a personal digital assistant, a mobile device, a palmtop computer, a laptop computer, a desktop computer, a communications device, a wireless smart phone, a personal trusted device, a wearable device, a global positioning satellite (GPS) device, a web appliance, or any other machine capable of executing a set of instructions (sequential or otherwise) that specify actions to be taken by that machine. Further, while a single computer system 102 is illustrated, additional embodiments may include any collection of systems or sub-systems that individually or jointly execute instructions or perform functions. The term “system” shall be taken throughout the present disclosure to include any collection of systems or sub-systems that individually or jointly execute a set, or multiple sets, of instructions to perform one or more computer functions.

[0037] As illustrated in FIG. 1, the computer system 102 may include at least one processor 104. The processor 104 is tangible and non-transitory. As used herein, the term “non-transitory” is to be interpreted not as an eternal characteristic of a state, but as a characteristic of a state that will last for a period of time. The term “non-transitory” specifically disavows fleeting characteristics such as characteristics of a particular carrier wave or signal or other forms that exist only transitorily in any place at any time. The processor 104 is an article of manufacture and/or a machine component. The processor 104 is configured to execute software instructions in order to perform functions as described in the various embodiments herein. The processor 104 may be a general-purpose processor or may be part of an application specific integrated circuit (ASIC). The processor 104 may also be a microprocessor, a microcomputer, a processor chip, a controller, a microcontroller, a digital signal processor (DSP), a state machine, or a programmable logic device. The processor 104 may also be a logical circuit, including a programmable gate array (PGA) such as a field programmable gate array (FPGA), or another type of circuit that includes discrete gate and/or transistor logic. The processor 104 may be a central processing

unit (CPU), a graphics processing unit (GPU), or both. Additionally, any processor described herein may include multiple processors, parallel processors, or both. Multiple processors may be included in, or coupled to, a single device or multiple devices.

[0038] The computer system 102 may also include a computer memory 106. The computer memory 106 may include a static memory, a dynamic memory, or both in communication. Memories described herein are tangible storage mediums that can store data as well as executable instructions and are non-transitory during the time instructions are stored therein. Again, as used herein, the term “non-transitory” is to be interpreted not as an eternal characteristic of a state, but as a characteristic of a state that will last for a period of time. The term “non-transitory” specifically disavows fleeting characteristics such as characteristics of a particular carrier wave or signal or other forms that exist only transitorily in any place at any time. The memories are an article of manufacture and/or machine component. Memories described herein are computer-readable mediums from which data and executable instructions can be read by a computer. Memories as described herein may be random access memory (RAM), read only memory (ROM), flash memory, electrically programmable read only memory (EPROM), electrically erasable programmable read-only memory (EEPROM), registers, a hard disk, a cache, a removable disk, tape, compact disk read only memory (CD-ROM), digital versatile disk (DVD), floppy disk, blu-ray disk, or any other form of storage medium known in the art. Memories may be volatile or non-volatile, secure and/or encrypted, unsecure and/or unencrypted. Of course, the computer memory 106 may comprise any combination of memories or a single storage.

[0039] The computer system 102 may further include a display 108, such as a liquid crystal display (LCD), an organic light emitting diode (OLED), a flat panel display, a solid state display, a cathode ray tube (CRT), a plasma display, or any other type of display, examples of which are well known to skilled persons.

[0040] The computer system 102 may also include at least one input device 110, such as a keyboard, a touch-sensitive input screen or pad, a speech input, a mouse, a remote control device having a wireless keypad, a microphone coupled to a speech recognition engine, a camera such as a video camera or still camera, a cursor control device, a global positioning system (GPS) device, an altimeter, a gyroscope,

an accelerometer, a proximity sensor, or any combination thereof. Those skilled in the art appreciate that various embodiments of the computer system 102 may include multiple input devices 110. Moreover, those skilled in the art further appreciate that the above-listed, exemplary input devices 110 are not meant to be exhaustive and that the computer system 102 may include any additional, or alternative, input devices 110.

[0041] The computer system 102 may also include a medium reader 112 which is configured to read any one or more sets of instructions, e.g. software, from any of the memories described herein. The instructions, when executed by a processor, can be used to perform one or more of the methods and processes as described herein. In a particular embodiment, the instructions may reside completely, or at least partially, within the memory 106, the medium reader 112, and/or the processor 110 during execution by the computer system 102.

[0042] Furthermore, the computer system 102 may include any additional devices, components, parts, peripherals, hardware, software or any combination thereof which are commonly known and understood as being included with or within a computer system, such as, but not limited to, a network interface 114 and an output device 116. The output device 116 may be, but is not limited to, a speaker, an audio out, a video out, a remote-control output, a printer, or any combination thereof.

[0043] Each of the components of the computer system 102 may be interconnected and communicate via a bus 118 or other communication link. As illustrated in FIG. 1, the components may each be interconnected and communicate via an internal bus. However, those skilled in the art appreciate that any of the components may also be connected via an expansion bus. Moreover, the bus 118 may enable communication via any standard or other specification commonly known and understood such as, but not limited to, peripheral component interconnect, peripheral component interconnect express, parallel advanced technology attachment, serial advanced technology attachment, etc.

[0044] The computer system 102 may be in communication with one or more additional computer devices 120 via a network 122. The network 122 may be, but is not limited to, a local area network, a wide area network, the Internet, a telephony network, a short-range network, or any other network commonly known and

understood in the art. The short-range network may include, for example, Bluetooth, Zigbee, infrared, near field communication, ultraband, or any combination thereof. Those skilled in the art appreciate that additional networks 122 which are known and understood may additionally or alternatively be used and that the exemplary networks 122 are not limiting or exhaustive. Also, while the network 122 is illustrated in FIG. 1 as a wireless network, those skilled in the art appreciate that the network 122 may also be a wired network.

[0045] The additional computer device 120 is illustrated in FIG. 1 as a personal computer. However, those skilled in the art appreciate that, in alternative embodiments of the present application, the computer device 120 may be a laptop computer, a tablet PC, a personal digital assistant, a mobile device, a palmtop computer, a desktop computer, a communications device, a wireless telephone, a personal trusted device, a web appliance, a server, or any other device that is capable of executing a set of instructions, sequential or otherwise, that specify actions to be taken by that device. Of course, those skilled in the art appreciate that the above-listed devices are merely exemplary devices and that the device 120 may be any additional device or apparatus commonly known and understood in the art without departing from the scope of the present application. For example, the computer device 120 may be the same or similar to the computer system 102. Furthermore, those skilled in the art similarly understand that the device may be any combination of devices and apparatuses.

[0046] Of course, those skilled in the art appreciate that the above-listed components of the computer system 102 are merely meant to be exemplary and are not intended to be exhaustive and/or inclusive. Furthermore, the examples of the components listed above are also meant to be exemplary and similarly are not meant to be exhaustive and/or inclusive.

[0047] In accordance with various embodiments of the present disclosure, the methods described herein may be implemented using a hardware computer system that executes software programs. Further, in an exemplary, non-limited embodiment, implementations can include distributed processing, component/object distributed processing, and parallel processing. Virtual computer system processing can be constructed to implement one or more of the methods or functionalities as described

herein, and a processor described herein may be used to support a virtual processing environment.

[0048] As described herein, various embodiments provide optimized methods and systems for hosting a market data distribution platform that is widely used by many financial organizations in a public cloud computing environment.

[0049] Referring to FIG. 2, a schematic of an exemplary network environment 200 for implementing a method for hosting a market data distribution platform that is widely used by many financial organizations in a public cloud computing environment is illustrated. In an exemplary embodiment, the method is executable on any networked computer platform, such as, for example, a personal computer (PC).

[0050] The method for hosting a market data distribution platform that is widely used by many financial organizations in a public cloud computing environment may be implemented by a Platform Hosting in Public Cloud (PHPC) device 202. The PHPC device 202 may be the same or similar to the computer system 102 as described with respect to FIG. 1. The PHPC device 202 may store one or more applications that can include executable instructions that, when executed by the PHPC device 202, cause the PHPC device 202 to perform actions, such as to transmit, receive, or otherwise process network messages, for example, and to perform other actions described and illustrated below with reference to the figures. The application(s) may be implemented as modules or components of other applications. Further, the application(s) can be implemented as operating system extensions, modules, plugins, or the like.

[0051] Even further, the application(s) may be operative in a cloud-based computing environment. The application(s) may be executed within or as virtual machine(s) or virtual server(s) that may be managed in a cloud-based computing environment. Also, the application(s), and even the PHPC device 202 itself, may be located in virtual server(s) running in a cloud-based computing environment rather than being tied to one or more specific physical network computing devices. Also, the application(s) may be running in one or more virtual machines (VMs) executing on the PHPC device 202. Additionally, in one or more embodiments of this technology, virtual machine(s) running on the PHPC device 202 may be managed or supervised by a hypervisor.

[0052] In the network environment 200 of FIG. 2, the PHPC device 202 is coupled to a plurality of server devices 204(1)-204(n) that hosts a plurality of databases 206(1)-206(n), and also to a plurality of client devices 208(1)-208(n) via communication network(s) 210. A communication interface of the PHPC device 202, such as the network interface 114 of the computer system 102 of FIG. 1, operatively couples and communicates between the PHPC device 202, the server devices 204(1)-204(n), and/or the client devices 208(1)-208(n), which are all coupled together by the communication network(s) 210, although other types and/or numbers of communication networks or systems with other types and/or numbers of connections and/or configurations to other devices and/or elements may also be used.

[0053] The communication network(s) 210 may be the same or similar to the network 122 as described with respect to FIG. 1, although the PHPC device 202, the server devices 204(1)-204(n), and/or the client devices 208(1)-208(n) may be coupled together via other topologies. Additionally, the network environment 200 may include other network devices such as one or more routers and/or switches, for example, which are well known in the art and thus will not be described herein. This technology provides a number of advantages including methods, non-transitory computer readable media, and PHPC devices that efficiently implement a method for hosting a market data distribution platform that is widely used by many financial organizations in a public cloud computing environment.

[0054] By way of example only, the communication network(s) 210 may include local area network(s) (LAN(s)) or wide area network(s) (WAN(s)), and can use TCP/IP over Ethernet and industry-standard protocols, although other types and/or numbers of protocols and/or communication networks may be used. The communication network(s) 210 in this example may employ any suitable interface mechanisms and network communication technologies including, for example, teletraffic in any suitable form (e.g., voice, modem, and the like), Public Switched Telephone Network (PSTNs), Ethernet-based Packet Data Networks (PDNs), combinations thereof, and the like.

[0055] The PHPC device 202 may be a standalone device or integrated with one or more other devices or apparatuses, such as one or more of the server devices 204(1)-204(n), for example. In one particular example, the PHPC device 202 may include or be hosted by one of the server devices 204(1)-204(n), and other

arrangements are also possible. Moreover, one or more of the devices of the PHPC device 202 may be in a same or a different communication network including one or more public, private, or cloud networks, for example.

[0056] The plurality of server devices 204(1)-204(n) may be the same or similar to the computer system 102 or the computer device 120 as described with respect to FIG. 1, including any features or combination of features described with respect thereto. For example, any of the server devices 204(1)-204(n) may include, among other features, one or more processors, a memory, and a communication interface, which are coupled together by a bus or other communication link, although other numbers and/or types of network devices may be used. The server devices 204(1)-204(n) in this example may process requests received from the PHPC device 202 via the communication network(s) 210 according to the HTTP-based and/or JavaScript Object Notation (JSON) protocol, for example, although other protocols may also be used.

[0057] The server devices 204(1)-204(n) may be hardware or software or may represent a system with multiple servers in a pool, which may include internal or external networks. The server devices 204(1)-204(n) hosts the databases 206(1)-206(n) that are configured to store data that relates to market information and public cloud architectural data.

[0058] Although the server devices 204(1)-204(n) are illustrated as single devices, one or more actions of each of the server devices 204(1)-204(n) may be distributed across one or more distinct network computing devices that together comprise one or more of the server devices 204(1)-204(n). Moreover, the server devices 204(1)-204(n) are not limited to a particular configuration. Thus, the server devices 204(1)-204(n) may contain a plurality of network computing devices that operate using a master/slave approach, whereby one of the network computing devices of the server devices 204(1)-204(n) operates to manage and/or otherwise coordinate operations of the other network computing devices.

[0059] The server devices 204(1)-204(n) may operate as a plurality of network computing devices within a cluster architecture, a peer-to peer architecture, virtual machines, or within a cloud architecture, for example. Thus, the technology

disclosed herein is not to be construed as being limited to a single environment and other configurations and architectures are also envisaged.

[0060] The plurality of client devices 208(1)-208(n) may also be the same or similar to the computer system 102 or the computer device 120 as described with respect to FIG. 1, including any features or combination of features described with respect thereto. For example, the client devices 208(1)-208(n) in this example may include any type of computing device that can interact with the PHPC device 202 via communication network(s) 210. Accordingly, the client devices 208(1)-208(n) may be mobile computing devices, desktop computing devices, laptop computing devices, tablet computing devices, virtual machines (including cloud-based computers), or the like, that host chat, e-mail, or voice-to-text applications, for example. In an exemplary embodiment, at least one client device 208 is a wireless mobile communication device, i.e., a smart phone.

[0061] The client devices 208(1)-208(n) may run interface applications, such as standard web browsers or standalone client applications, which may provide an interface to communicate with the PHPC device 202 via the communication network(s) 210 in order to communicate user requests and information. The client devices 208(1)-208(n) may further include, among other features, a display device, such as a display screen or touchscreen, and/or an input device, such as a keyboard, for example.

[0062] Although the exemplary network environment 200 with the PHPC device 202, the server devices 204(1)-204(n), the client devices 208(1)-208(n), and the communication network(s) 210 are described and illustrated herein, other types and/or numbers of systems, devices, components, and/or elements in other topologies may be used. It is to be understood that the systems of the examples described herein are for exemplary purposes, as many variations of the specific hardware and software used to implement the examples are possible, as will be appreciated by those skilled in the relevant art(s).

[0063] One or more of the devices depicted in the network environment 200, such as the PHPC device 202, the server devices 204(1)-204(n), or the client devices 208(1)-208(n), for example, may be configured to operate as virtual instances on the same physical machine. In other words, one or more of the PHPC device 202, the

server devices 204(1)-204(n), or the client devices 208(1)-208(n) may operate on the same physical device rather than as separate devices communicating through communication network(s) 210. Additionally, there may be more or fewer PHPC devices 202, server devices 204(1)-204(n), or client devices 208(1)-208(n) than illustrated in FIG. 2.

[0064] In addition, two or more computing systems or devices may be substituted for any one of the systems or devices in any example. Accordingly, principles and advantages of distributed processing, such as redundancy and replication also may be implemented, as desired, to increase the robustness and performance of the devices and systems of the examples. The examples may also be implemented on computer system(s) that extend across any suitable network using any suitable interface mechanisms and traffic technologies, including by way of example only teletraffic in any suitable form (e.g., voice and modem), wireless traffic networks, cellular traffic networks, Packet Data Networks (PDNs), the Internet, intranets, and combinations thereof.

[0065] The PHPC device 202 is described and illustrated in FIG. 3 as including a public cloud platform hosting module 302, although it may include other rules, policies, modules, databases, or applications, for example. As will be described below, the public cloud platform hosting module 302 is configured to implement a method for hosting a market data distribution platform that is widely used by many financial organizations in a public cloud computing environment.

[0066] An exemplary process 300 for implementing a mechanism for hosting a market data distribution platform that is widely used by many financial organizations in a public cloud computing environment by utilizing the network environment of FIG. 2 is illustrated as being executed in FIG. 3. Specifically, a first client device 208(1) and a second client device 208(2) are illustrated as being in communication with PHPC device 202. In this regard, the first client device 208(1) and the second client device 208(2) may be “clients” of the PHPC device 202 and are described herein as such. Nevertheless, it is to be known and understood that the first client device 208(1) and/or the second client device 208(2) need not necessarily be “clients” of the PHPC device 202, or any entity described in association therewith herein. Any additional or alternative relationship may exist between either or both of

the first client device 208(1) and the second client device 208(2) and the PHPC device 202, or no relationship may exist.

[0067] Further, PHPC device 202 is illustrated as being able to access a market information for consumers data repository 206(1) and a public cloud architecture database 206(2). The public cloud platform hosting module 302 may be configured to access these databases for implementing a method for hosting a market data distribution platform that is widely used by many financial organizations in a public cloud computing environment.

[0068] The first client device 208(1) may be, for example, a smart phone. Of course, the first client device 208(1) may be any additional device described herein. The second client device 208(2) may be, for example, a personal computer (PC). Of course, the second client device 208(2) may also be any additional device described herein.

[0069] The process may be executed via the communication network(s) 210, which may comprise plural networks as described above. For example, in an exemplary embodiment, either or both of the first client device 208(1) and the second client device 208(2) may communicate with the PHPC device 202 via broadband or cellular communication. Of course, these embodiments are merely exemplary and are not limiting or exhaustive.

[0070] Upon being started, the public cloud platform hosting module 302 executes a process for hosting a market data distribution platform that is widely used by many financial organizations in a public cloud computing environment. An exemplary process for hosting a market data distribution platform that is widely used by many financial organizations in a public cloud computing environment is generally indicated at flowchart 400 in FIG. 4.

[0071] In process 400 of FIG. 4, at step S402, the public cloud platform hosting module 302 accesses a first cloud server that operates in a public cloud computing environment. In an exemplary embodiment, the public cloud computing environment may be administered by Amazon Web Services (AWS). Fifth generation hardware virtualization in the M5 class of AWS instances facilitates access to many hardware benefits by virtual machines, such as granular metrics collection and elimination of variability in network latencies. Containerization with Docker and

Kubernetes provides an avenue to create a software-defined operating environment that is easily portable between test environments on premises and cloud-based services.

[0072] At step S404, the public cloud platform hosting module 302 facilitates a communication between a market data distribution platform and the first cloud server. In an exemplary embodiment, the market data distribution platform is Thomson Reuters Enterprise Platform (TREP).

[0073] At step S406, the public cloud platform hosting module 302 monitors a data volume, i.e., an amount of data being distributed on the platform, and a number of clients that are accessing the platform. In an exemplary embodiment, due to the cyclical nature of market data, there will be a greater volume of data during daytime hours and a lesser volume of data during nighttime hours for any particular region.

[0074] At step S408, the public cloud platform hosting module 302 uses a result of the monitoring performed in step S406 to determine a minimum number of public cloud servers that are required for servicing the current number of clients. Then, at step S410, the public cloud platform hosting module 302 facilitates communications between the platform and each of a plurality of public cloud servers, the number of which is greater than or equal to the determined minimum number.

[0075] In an exemplary embodiment, these communications are facilitated by using a multicast protocol. The use of multicast, to distribute inside of the platform's components, allows engineers to build an environment that can facilitate elasticity inside the software stack, which is invisible to outside consumers and clients.

[0076] At step S412, the public cloud platform hosting module 302 adjusts the number of public cloud servers being used based on variations in the data volume and/or the number of clients. In this aspect, the monitoring being performed in step S406 is performed continuously so that such variations can be observed. In an exemplary embodiment, by reducing the number of public cloud servers during times of reduced platform demand and increasing the number of public cloud servers during times of heavier platform demand, resources may be efficiently deployed, thereby resulting in significant cost savings. In an exemplary embodiment, Kubernetes provides an orchestration framework that is beneficial to market data-style workloads,

where environment can be grown and shrunk based on time of day to correspond to workloads seen in the financial markets.

[0077] In an exemplary embodiment, the public cloud platform hosting module 302 may also perform a load-shifting operation between public cloud servers being used, based on a result of the monitoring, such that a portion of a load from at least a first one of the cloud servers is shifted to a second one of the cloud servers. For example, when new public cloud servers are provisioned, the public cloud platform hosting module 302 may actively “drain” some of the existing connections to the new servers by detecting lulls in specific sections of the market and then redirecting clients that are interested in those sections to migrate to the new servers in small groups, thereby effectively shifting the computational load among the existing public cloud servers and the newly provisioned cloud servers so as to at least partially balance the loads. This process may be referred to as “load shedding.” Similarly, when the number of public cloud servers is reduced as a result of the monitoring, the public cloud platform hosting module 302 may actively shift the computational loads in order to maintain a relative balance among the remaining public cloud servers.

[0078] Accordingly, with this technology, an optimized process for hosting a market data distribution platform that is widely used by many financial organizations in a public cloud computing environment is provided.

[0079] Although the invention has been described with reference to several exemplary embodiments, it is understood that the words that have been used are words of description and illustration, rather than words of limitation. Changes may be made within the purview of the appended claims, as presently stated and as amended, without departing from the scope and spirit of the present disclosure in its aspects. Although the invention has been described with reference to particular means, materials and embodiments, the invention is not intended to be limited to the particulars disclosed; rather the invention extends to all functionally equivalent structures, methods, and uses such as are within the scope of the appended claims.

[0080] For example, while the computer-readable medium may be described as a single medium, the term “computer-readable medium” includes a single medium or multiple media, such as a centralized or distributed database, and/or associated caches and servers that store one or more sets of instructions. The term “computer-

readable medium” shall also include any medium that is capable of storing, encoding or carrying a set of instructions for execution by a processor or that cause a computer system to perform any one or more of the embodiments disclosed herein.

[0081] The computer-readable medium may comprise a non-transitory computer-readable medium or media and/or comprise a transitory computer-readable medium or media. In a particular non-limiting, exemplary embodiment, the computer-readable medium can include a solid-state memory such as a memory card or other package that houses one or more non-volatile read-only memories. Further, the computer-readable medium can be a random-access memory or other volatile re-writable memory. Additionally, the computer-readable medium can include a magneto-optical or optical medium, such as a disk or tapes or other storage device to capture carrier wave signals such as a signal communicated over a transmission medium. Accordingly, the disclosure is considered to include any computer-readable medium or other equivalents and successor media, in which data or instructions may be stored.

[0082] Although the present application describes specific embodiments which may be implemented as computer programs or code segments in computer-readable media, it is to be understood that dedicated hardware implementations, such as application specific integrated circuits, programmable logic arrays and other hardware devices, can be constructed to implement one or more of the embodiments described herein. Applications that may include the various embodiments set forth herein may broadly include a variety of electronic and computer systems. Accordingly, the present application may encompass software, firmware, and hardware implementations, or combinations thereof. Nothing in the present application should be interpreted as being implemented or implementable solely with software and not hardware.

[0083] Although the present specification describes components and functions that may be implemented in particular embodiments with reference to particular standards and protocols, the disclosure is not limited to such standards and protocols. Such standards are periodically superseded by faster or more efficient equivalents having essentially the same functions. Accordingly, replacement standards and protocols having the same or similar functions are considered equivalents thereof.

[0084] The illustrations of the embodiments described herein are intended to provide a general understanding of the various embodiments. The illustrations are not intended to serve as a complete description of all the elements and features of apparatus and systems that utilize the structures or methods described herein. Many other embodiments may be apparent to those of skill in the art upon reviewing the disclosure. Other embodiments may be utilized and derived from the disclosure, such that structural and logical substitutions and changes may be made without departing from the scope of the disclosure. Additionally, the illustrations are merely representational and may not be drawn to scale. Certain proportions within the illustrations may be exaggerated, while other proportions may be minimized. Accordingly, the disclosure and the figures are to be regarded as illustrative rather than restrictive.

[0085] One or more embodiments of the disclosure may be referred to herein, individually and/or collectively, by the term “invention” merely for convenience and without intending to voluntarily limit the scope of this application to any particular invention or inventive concept. Moreover, although specific embodiments have been illustrated and described herein, it should be appreciated that any subsequent arrangement designed to achieve the same or similar purpose may be substituted for the specific embodiments shown. This disclosure is intended to cover any and all subsequent adaptations or variations of various embodiments. Combinations of the above embodiments, and other embodiments not specifically described herein, will be apparent to those of skill in the art upon reviewing the description.

[0086] The Abstract of the Disclosure is submitted with the understanding that it will not be used to interpret or limit the scope or meaning of the claims. In addition, in the foregoing Detailed Description, various features may be grouped together or described in a single embodiment for the purpose of streamlining the disclosure. This disclosure is not to be interpreted as reflecting an intention that the claimed embodiments require more features than are expressly recited in each claim. Rather, as the following claims reflect, inventive subject matter may be directed to less than all of the features of any of the disclosed embodiments. Thus, the following claims are incorporated into the Detailed Description, with each claim standing on its own as defining separately claimed subject matter.

[0087] The above disclosed subject matter is to be considered illustrative, and not restrictive, and the appended claims are intended to cover all such modifications, enhancements, and other embodiments which fall within the true spirit and scope of the present disclosure. Thus, to the maximum extent allowed by law, the scope of the present disclosure is to be determined by the broadest permissible interpretation of the following claims, and their equivalents, and shall not be restricted or limited by the foregoing detailed description.

What is claimed is:

1. A method for hosting a market data distribution platform in a public cloud computing environment, the method being implemented by at least one processor, the method comprising:

accessing, by the at least one processor, a first cloud server that operates in a public cloud computing environment;

facilitating, by the at least one processor, a communication between a market data distribution platform and the first cloud server;

monitoring, by the at least one processor, an amount of data being distributed on the market data distribution platform and a number of clients that are accessing the market data distribution platform;

determining, by the at least one processor based on a result of the monitoring, a minimum number of cloud servers required for servicing the clients;

accessing, by the at least one processor, a plurality of cloud servers that includes at least the determined minimum number of cloud servers;

facilitating, by the at least one processor, respective communications between the market data distribution platform and each of the plurality of cloud servers; and

shifting, by the at least one processor based on a result of the monitoring, a portion of a load from at least a first one of the plurality of cloud servers to at least a second one of the plurality of cloud servers.

2. The method of claim 1, wherein when the minimum number of cloud servers required for servicing the clients is reduced based on a subsequent result of the monitoring, the method further comprises reducing a number of cloud servers that are accessed based on the reduced minimum number of cloud servers and shifting a portion of a load from at least a first one of the reduced minimum number of cloud servers to at least a second one of the reduced minimum number of cloud servers.

3. The method of claim 1, wherein when the minimum number of cloud servers required for servicing the clients is increased based on a subsequent result of the monitoring, the method further comprises increasing a number of cloud servers that are accessed based on the increased minimum number of cloud servers and shifting a portion of a load from at least a first one of the increased minimum number

of cloud servers to at least a second one of the increased minimum number of cloud servers.

4. The method of claim 1, wherein the facilitating of the respective communications between the market data distribution platform and each of the plurality of cloud servers comprises using a multicast communication protocol.

5. The method of claim 1, wherein the market data distribution platform includes a Thomson Reuters Enterprise Platform (TREP).

6. The method of claim 1, wherein the monitoring is performed continuously.

7. The method of claim 6, wherein a result of the determining of the minimum number of cloud servers required for servicing the clients increases as a result of a beginning of a trading day and decreases as a result of an end of the trading day.

8. The method of claim 1, wherein the public cloud computing environment is administered by Amazon Web Services (AWS).

9. A computing apparatus for hosting a market data distribution platform in a public cloud computing environment, the computing apparatus comprising:

a processor;

a memory; and

a communication interface coupled to each of the processor and the memory, wherein the processor is configured to:

access a first cloud server that operates in a public cloud computing environment;

facilitate a communication between a market data distribution platform and the first cloud server;

monitor an amount of data being distributed on the market data distribution platform and a number of clients that are accessing the market data distribution platform;

determine, based on a result of the monitoring, a minimum number of cloud servers required for servicing the clients;

access a plurality of cloud servers that includes at least the determined minimum number of cloud servers;

facilitate respective communications between the market data distribution platform and each of the plurality of cloud servers; and

shift, based on a result of the monitoring, a portion of a load from at least a first one of the plurality of cloud servers to at least a second one of the plurality of cloud servers.

10. The computing apparatus of claim 9, wherein when the minimum number of cloud servers required for servicing the clients is reduced based on a subsequent result of the monitoring, the processor is further configured to reduce a number of cloud servers that are accessed based on the reduced minimum number of cloud servers and to shift a portion of a load from at least a first one of the reduced minimum number of cloud servers to at least a second one of the reduced minimum number of cloud servers.

11. The computing apparatus of claim 9, wherein when the minimum number of cloud servers required for servicing the clients is increased based on a subsequent result of the monitoring, the processor is further configured to increase a number of cloud servers that are accessed based on the increased minimum number of cloud servers and to shifting a portion of a load from at least a first one of the increased minimum number of cloud servers to at least a second one of the increased minimum number of cloud servers.

12. The computing apparatus of claim 9, wherein the processor is further configured to use a multicast communication protocol for performing the facilitating of the respective communications between the market data distribution platform and each of the plurality of cloud servers.

13. The computing apparatus of claim 9, wherein the market data distribution platform includes a Thomson Reuters Enterprise Platform (TREP).

14. The computing apparatus of claim 9, wherein the processor is further configured to continuously perform the monitoring of the amount of data being

distributed on the market data distribution platform and the number of clients that are accessing the market data distribution platform.

15. The computing apparatus of claim 14, wherein a result of the determination of the minimum number of cloud servers required for servicing the clients increases as a result of a beginning of a trading day and decreases as a result of an end of the trading day.

16. The computing apparatus of claim 9, wherein the public cloud computing environment is administered by Amazon Web Services (AWS).

17. A non-transitory computer readable storage medium storing instructions for hosting a market data distribution platform in a public cloud computing environment, the storage medium comprising executable code which, when executed by a processor, causes the processor to:

access a first cloud server that operates in a public cloud computing environment;

facilitate a communication between a market data distribution platform and the first cloud server;

monitor an amount of data being distributed on the market data distribution platform and a number of clients that are accessing the market data distribution platform;

determine, based on a result of the monitoring, a minimum number of cloud servers required for servicing the clients;

access a plurality of cloud servers that includes at least the determined minimum number of cloud servers;

facilitate respective communications between the market data distribution platform and each of the plurality of cloud servers; and

shift, based on a result of the monitoring, a portion of a load from at least a first one of the plurality of cloud servers to at least a second one of the plurality of cloud servers.

18. The storage medium of claim 17, wherein when the minimum number of cloud servers required for servicing the clients is reduced based on a subsequent result of the monitoring, the executable code further causes the processor to reduce a number of cloud servers that are accessed based on the reduced minimum number of cloud servers and to shift a portion of a load from at least a first one of the reduced minimum number of cloud servers to at least a second one of the reduced minimum number of cloud servers.

19. The storage medium of claim 17, wherein when the minimum number of cloud servers required for servicing the clients is increased based on a subsequent result of the monitoring, the executable code further causes the processor to increase a number of cloud servers that are accessed based on the increased minimum number of cloud servers and to shift a portion of a load from at least a first one of the increased minimum number of cloud servers to at least a second one of the increased minimum number of cloud servers.

20. The storage medium of claim 17, wherein the executable code further causes the processor to use a multicast communication protocol for the facilitating of the respective communications between the market data distribution platform and each of the plurality of cloud servers.




FIG. 1






FIG. 2

**FIG. 3**

400**FIG. 4**

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2022/027250

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - INV. - H04L 67/1031 (2022.01)

ADD. - G06F 16/27; G06Q 10/06; H04L 47/70; H04L 67/60 (2022.01)

CPC - INV. - H04L 67/1031; H04L 67/141; H04L 67/60 (2022.08)

ADD. - G06F 16/283; H04L 47/70 (2022.08)

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

See Search History document

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

See Search History document

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

See Search History document

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	JAISWAL. Load balancing of financial data using machine learning and cloud analytics. 31 May 2020 (31.05.2020) [retrieved on 17-08-2022]. Retrieved from the internet <URL: https://digitalcommons.wku.edu/cgi/viewcontent.cgi?article=4220&context=theses> entire document	1-4, 6, 8-12, 14, 16-20
Y		5, 13
A		7, 15
Y	US 2016/0034588 A1 (THOMSON REUTERS GLOBAL RESOURCES) 04 February 2016 (04.02.2016) entire document	5, 13
A	US 2014/0149784 A1 (DH2I COMPANY) 29 May 2014 (29.05.2014) entire document	1-20
A	US 2015/0234670 A1 (FUJITSU LIMITED) 20 August 2015 (20.08.2015) entire document	1-20
A	US 2015/0347542 A1 (STATE STREET CORPORATION) 03 December 2015 (03.12.2015) entire document	1-20

☐ 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

"D" document cited by the applicant in the international application

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

17 August 2022

Date of mailing of the international search report

SEP 19 2022

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, VA 22313-1450

Facsimile No. 571-273-8300

Authorized officer

Taina Matos

Telephone No. PCT Helpdesk: 571-272-4300