

(51) International Patent Classification:
G06F 15/16 (2006.01) *G06F 9/44* (2006.01)(21) International Application Number:
PCT/US2011/028159(22) International Filing Date:
11 March 2011 (11.03.2011)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
12/724,507 16 March 2010 (16.03.2010) US(71) Applicant (for all designated States except US): **MICROSOFT CORPORATION** [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).(72) Inventors: **BANSAL, Deepak**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **WARKHEDE, Priyank Ramesh**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO,

DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, 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, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

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

Published:

- without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) Title: SHAPING VIRTUAL MACHINE COMMUNICATION TRAFFIC

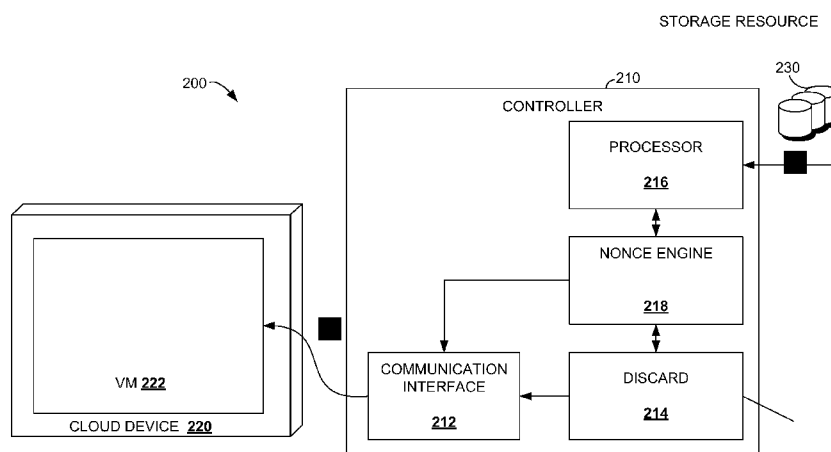


FIG. 2

(57) Abstract: Cloud computing platforms having computer-readable media that perform methods to shape virtual machine communication traffic. The cloud computing platform includes virtual machines and a controller. The controller limits the traffic associated with the virtual machines to enable the virtual machines to achieve desired communication rates, especially when a network servicing the virtual machines is congested. The controller may drop communication messages associated with the virtual machines based on a drop probability evaluated for the virtual machines.

SHAPING VIRTUAL MACHINE COMMUNICATION TRAFFIC

BACKGROUND

[0001] Conventionally, cloud computing platforms host software applications in an Internet-accessible virtual environment. The cloud computing platform allows an organization to use datacenters designed and maintained by third parties. The conventional virtual environment supplies small or large organizations with requested hardware resources, software application resources, network resources, and storage resources. The virtual environment also provides application security, application reliability, application scalability, and application availability.

[0002] The conventional datacenters provide the physical computing resources, physical storage resources, and physical network resources. The physical resources in the data center are virtualized and exposed to the organizations as a set of application programming interfaces. The organizations do not need to maintain their own hardware resources or software resources, or maintain datacenters that are reliable and scalable.

[0003] The organizations may access these physical resources efficiently through the virtual environment without knowing the details of the software or the underlying physical hardware. In a conventional cloud computing platform, the hardware resources and software resources may be shared by organizations who do not trust each other. To prevent denial of service to any one organization, the conventional cloud platforms provide procedures that maintain equitable access to the shared resources. Most of the procedures require a significant amount of state information which increases computational and cost overhead in the cloud computing platform.

[0004] For instance, the sharing procedures may include leaky bucket, token bucket, and fair queuing. The leaky bucket procedure stores communication messages in a queue and transmits the communication messages at a constant rate. When the queue is full, communication messages are discarded. The token bucket procedure stores communication message in a queue and transmits the communication at a rate that is based on the number of tokens associated with the queue. Communication message may be discarded if the communication messages wait in the bucket for a predetermined time.

The fair queuing procedure stores communication message in a queue and transmits the communication messages at a rate the is equal or proportional to the rate of experienced by other queues for other communication sessions. These procedures provide a mechanism to

share resources in the cloud but have high CPU overhead because of the need to maintaining queues and associated state for the queues.

SUMMARY

[0005] Embodiments of the invention relate, in one regard, to cloud computing platforms, computer-readable media, and computer-implemented methods that shape virtual machine traffic in a cloud computing platform. The cloud computing platform includes controllers and virtual machines that are connected to a communication network. The controllers provide a stateless mechanism for shaping communication traffic with minimal overhead.

[0006] The controllers monitor traffic on the communication network. The controllers calculate an average communication rate for virtual machines connected to the communication network. In turn, a desired sending rate for the virtual machines is obtained by the controller. The controller determines a drop probability for the virtual machines based on the average communication rate. Communication messages in the network are dropped based on the drop probability for the virtual machines.

[0007] This summary is provided to introduce a selection of concepts in a simplified form that are further described below in the detailed description. This summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used in isolation as an aid in determining the scope of the claimed subject matter.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] Fig. 1 is a network diagram that illustrates an exemplary cloud computing platform in accordance with embodiments of the invention;

[0009] Fig. 2 is a block diagram that illustrates a communication system having exemplary controllers and virtual machines in the exemplary cloud computing platform in accordance with embodiment of the invention;

[0010] Fig. 3 is a logic diagram that illustrates an exemplary method to manage virtual machine communication traffic in accordance with embodiment of the invention; and

[0011] Fig. 4 is a logic diagram that illustrates an exemplary method to discard virtual machine communication traffic in accordance with embodiment of the invention.

DETAILED DESCRIPTION

[0012] This patent describes the subject matter for patenting with specificity to meet statutory requirements. However, the description itself is not intended to limit the scope of this patent. Rather, the inventors have contemplated that the claimed subject matter might also be embodied in other ways, to include different steps or combinations of steps

similar to the ones described in this patent, in conjunction with other present or future technologies. Moreover, although the terms “step” and “block” may be used herein to connote different elements of methods employed, the terms should not be interpreted as implying any particular order among or between various steps herein disclosed unless and
5 except when the order of individual steps is explicitly described. Further, embodiments are described in detail below with reference to the attached drawing figures, which are incorporated in their entirety by reference herein.

[0013] As utilized herein, the term “component” refers to any combination of hardware, software, and firmware.

10 **[0014]** The cloud computing platform enables sharing of hardware and software resources among virtual machines. In some embodiments, virtual machine traffic is limited to prevent resource starvation in the cloud computing platform. The cloud computing platform execute traffic shaping procedures that limit bandwidth utilized by the virtual machines. Accordingly, the cloud computing platform protects shared resources from
15 malicious virtual machines that attempt to access excess resources to reduce accessibility by other virtual machines. The cloud computing platform also protects shared resources from inadvertent use of excess network resources by non-malicious virtual machines.

[0015] In some embodiments, the cloud computing platform utilizes a queueless and stateless mechanism to implement bandwidth limiting for the virtual machines. This
20 mechanism keeps the overhead for each virtual machine low. In one embodiment, the communication messages for the virtual machines are transmission control protocol (TCP) communication messages, i.e., packets.

[0016] Bandwidth limiting allows an administrator of the cloud computing platform to set the bandwidth available to a given virtual machine based on the relationship with the
25 virtual machine. For instance, a small virtual machine may be limited to 100 Mbps and a large virtual machine may be limited to 1 Gbps. Thus, desired communication rates may be set in service level agreements negotiated with each virtual machine in the cloud computing platform.

[0017] As one skilled in the art will appreciate, the cloud computing platform may include
30 hardware, software, or a combination of hardware and software. The hardware includes processors and memories configured to execute instructions stored in the memories. In one embodiment, the memories include computer-readable media that store a computer-program product having computer-useable instructions for a computer-implemented method. Computer-readable media include both volatile and nonvolatile media,

removable and nonremovable media, and media readable by a database, a switch, and various other network devices. Network switches, routers, and related components are conventional in nature, as are means of communicating with the same. By way of example, and not limitation, computer-readable media comprise computer-storage media and communications media. Computer-storage media, or machine-readable media, include media implemented in any method or technology for storing information. Examples of stored information include computer-useable instructions, data structures, program modules, and other data representations. Computer-storage media include, but are not limited to, random access memory (RAM), read only memory (ROM), electrically erasable programmable read only memory (EEPROM), flash memory or other memory technology, compact-disc read only memory (CD-ROM), digital versatile discs (DVD), holographic media or other optical disc storage, magnetic cassettes, magnetic tape, magnetic disk storage, and other magnetic storage devices. These memory technologies can store data momentarily, temporarily, or permanently.

[0018] In one embodiment, the cloud computing platform includes cloud applications that are available to client devices. The client devices access the cloud computing platform to execute the cloud applications on one or virtual machines. The cloud applications are implemented using storage and processing resources available in the cloud computing platform.

[0019] Fig. 1 is a network diagram that illustrates an exemplary computing system 100 in accordance with embodiments of the invention. In an embodiment, the computing system 100 shown in FIG. 1 is merely exemplary and is not intended to suggest any limitation as to scope or functionality. Embodiments of the invention are operable with numerous other configurations. With reference to FIG. 1, the computing system 100 includes a cloud computing platform 110, cloud applications 120, and client devices 130.

[0020] The cloud computing platform 110 is configured to execute cloud applications 120 requested by the client devices 130. The cloud computing platform 110 maintains computing devices that provide virtual machines, which execute the cloud application 120. The cloud computing platform also includes storage resources that store applications and system information. The cloud computing platform 110 connects to the client devices 130 via a communications network, such as a wireless network, local area network, wired network, or the Internet.

[0021] The cloud applications 120 are available to the client devices 130. The software executed on the cloud computing platform 110 implements the cloud applications 120. In

one embodiment, virtual machines provided by the cloud computing platform 110 execute the cloud applications 120. The cloud applications 120 may include, but are not limited to, editing applications, network management applications, finance applications, or any application requested or developed by the client devices 130. In certain embodiments,

5 some functionality of the cloud application 120 may be executed on the client devices 130.

[0022] The client devices 130 are utilized by a user to interact with cloud applications 120 provided by the cloud computing platform 110. The client devices 130, in some embodiments, must register with the cloud computing platform 110 to access the cloud applications 120. Any client device 130 with an account from the cloud computing
10 platform 110 may access the cloud applications 120 and other resources provided in the cloud computing platform 110. The client devices 130 include, without limitation, personal digital assistants, smart phones, laptops, personal computers, gaming systems, set-top boxes, or any other suitable client computing device. The client devices 130 include user and system information storage to store user and system information on the
15 client devices 130. The user information may include search histories, cookies, and passwords. The system information may include internet protocol addresses, cached Web pages, and system utilization. The client devices 130 communicate with the cloud computing platform 110 to receive results from the cloud applications 120.

[0023] Accordingly, the computing system 100 is configured with a cloud computing
20 platform 110 that provides cloud applications 120 to the client devices 130. The cloud applications 120 remove the burden of updating and managing multiple local client applications on the client devices 130.

[0024] In certain embodiments, the cloud computing platform provides controllers that shape virtual machine communication traffic. The controllers provide a counter for each
25 virtual machine. The counters may track the average sending rate of traffic. The cloud computing platform may evaluate an exponentially weighted average sending rate for the virtual machines. Based on the average sending rate, a drop rate is computed by the controller for each virtual machines. Communication messages are randomly dropped by the controller based on the computed average sending rate. In one embodiment, the drop
30 probability may be set via policy included in the service level agreement governing the virtual machine. Alternatively, the drop probability may be triggered on demand by a central policy based on network congestion detected by the cloud computing platform.

[0025] Fig. 2 is a block diagram that illustrates a communication system 200 having exemplary controllers 210 and virtual machines 222 in the exemplary cloud computing

platform. The controller 210 is connected to a communication network. The communication network connects the cloud devices 220, and storage resources 230.

[0026] The controller 210 monitors communication traffic on the communication network associated with virtual machines 222 that execute the cloud applications. The controller
5 210 includes a communication interface 212, a discard loop 214, processor 216, and nonce engine 218. In certain embodiments, the controller 210 is configured to shape virtual machine communication traffic by discarding a fraction of the outbound traffic based on send rates detected for the virtual machines. The controller is able to shape traffic without maintaining queues for the virtual machines or communication sessions. The controller
10 210 measures a sending rate of each virtual machine and selects one or more communication messages for discarding at a determined probability.

[0027] The communication interface 212 receives the communication messages from the cloud devices 220 executing the virtual machines 222. The communication interface also transmits the communication messages to the cloud device 220 executing the virtual
15 machines 222. The communication interface 212 may include, but is not limited to, Ethernet, or Asynchronous Transfer Mode (ATM) interfaces. The controller 210 may transmit a communication message to the cloud device 220 over the communication interface 212.

[0028] The discard loop 214 is used by the controller 210 to discard communication
20 messages. The communication messages may be discarded by sending the message to an interface connected to ground. The communication messages are sent to the discard loop when the controller 210 determines that a virtual machine 222 has a drop probability above a specified threshold.

[0029] The processor 216 performs calculations for the controller 210. The calculations
25 are used by the controller to determine whether to discard communication messages for a virtual machine 222. The calculations may include evaluating an average sending rate, a drop probability, and a nonce. Based on these calculations, the processor 216 may set a flag to discard one or more communication messages for the virtual machine 222.

[0030] The processor 216 obtains the current sending rate for the virtual machines 222
30 executed by cloud device 220. In another embodiment, the processor 216 may calculate the current sending rate each time a communication message is received from the virtual machines 222 by dividing the size of the communication message by the interval since last communication message was received by the controller 210 from the virtual machine 222. The current sending rate is used to determine a drop probability for the virtual machines

222 when the drop probability is not set in a service level agreement for the virtual machine 222. The service level agreement is stored in storage resource 230, which is accessed by the processor 216 to determine whether the drop probability is set.

[0031] In turn, the processor 216 calculates the average sending rate. In one embodiment, the average sending rate is an exponential average sending rate (ASR). The processor may evaluate $ASR = \alpha * previous_ASR + (1 - \alpha) * current_SR$, where “ α ” is set by the cloud computing platform, the “previous_ASR” is the ASR previously evaluated for the virtual machines, and “current_SR” is the current sending rate detected for the virtual machines. “ α ” is a rational number. The exponential averaging uses less memory than other averaging function. In certain embodiments, the other averaging functions, e.g. mean sending rate, may be used.

[0032] The desired sending rate for the virtual machines 220 is accessed to evaluate the drop probability. The desired sending rate may be set in the service level agreement. In other embodiments, an administrator may set the desired sending rate in real time based on the network congestions. In yet another embodiment, the desired sending rate may be automatically set based on the congestion detected on the network. The desired sending rate may be automatically decreased as congestion on the communication network increases.

[0033] The processor 216 calculates the drop probability (DP). The processor evaluates $DP = \frac{\beta * (ASR - DSR)}{ASR}$, “ β ” is set by the cloud computing platform. β may be set by an administrator to reduce or increase the number of communication messages discarded. β ranges between 0 and 1.

[0034] The nonce engine 218 is a random number generator. It generates a random number for the processor 216. The nonce engine 218 may also provide the processor with the maximum random number. In turn, the processor 216 may generate a nonce (η) that is used as a threshold. The processor 216 evaluates $\eta = \frac{Random_Number}{Maximum_Random_Number}$,

“Random_Number” is the random number generated by the nonce engine 218, and “Maximum_Random_Number” is the maximum random number provided by the nonce engine 218. Both “Random_Number” and “Maximum_Random_Number” are rational numbers.

[0035] The processor 216 compares the nonce and drop probability. When the drop probability is greater than the nonce, the processor 216 may discard the communication message.

[0036] In some embodiments, the drop probability for a virtual machines 222 may be set by the service level agreement based on network congestion. For instance, when there is no network congestion is detected by the controller 210, the service level agreement may indicate that the average sending rate should be set to the desired sending rate. Hence, the drop probability is zero. And the virtual machine 222 is allowed unrestricted resource utilization. But when network congestion is detected by the controller 210, the average sending rate is set to a specific number, e.g., 50 Mbps, 20 Mbps, depending on the congestion level and the number of virtual machines 222 accessing the resources in the cloud computing platform.

[0037] Storage resources 230 store the cloud applications and the service level agreements for each virtual machine 222. In some embodiments, a virtual machine 222 may have multiple service level agreements. Each service level agreement may correspond to different cloud applications executed by the virtual machine. The service level agreements may set desired sending rate for the virtual machines 222, desired sending rates for the cloud applications, and desired sending rates for the cloud device 220. In some embodiments, the service level agreements may also store a drop probability for the virtual machines 222, drop probability for the cloud applications, and drop probability for the cloud device 220. The storage resource 230 may transmit service level agreements to the controller 210 in response to requests for the desired sending rate or drop probability.

[0038] In one embodiment, the cloud computing platform executes cloud applications on virtual machines running on cloud devices. The cloud computing platform shapes communication traffic for the virtual machines. The cloud computing platform monitors communications rates for the virtual machines and discards packets based on the communication rates. The communication rates may include transmission rates and reception rates.

[0039] Fig. 3 is a logic diagram that illustrates an exemplary method to manage virtual machine communication traffic. The method initializes in step 310. The cloud computing platform maintains an average communication rate for the virtual machines in the cloud computing platform, in step 320. The cloud computing platform may use an exponential average or mean average for the virtual machine. In step 330, the cloud computing platform calculates a drop probability for the virtual machines as a function of the average

communication rate and a desired communication rate for the virtual machines. The desired communication rate may be obtained from a storage resource storing service level agreements that specify the desired sending rate for the virtual machines. Alternatively, the desired sending rate may be set by an administrator of the cloud computing platform.

5 **[0040]** The cloud computing platform limits transmission communication messages or reception communication messages associated with the virtual machines to achieve the desired communication rates for virtual machines in the cloud computing platform. In one embodiment, the limiting of transmission communication messages or reception communication messages is triggered when the cloud computing platform detects
 10 congestion on a network servicing the virtual machines. In turn, the cloud computing platform randomly drops a communication message for the virtual machines when a drop probability for the virtual machines is greater than a nonce generated by the cloud computing platform, in step 340. The method terminates in step 350.

[0041] In some embodiments, the cloud computing platform shapes communication traffic
 15 based on sending rates for the virtual machines executing on cloud devices. The cloud computing platform may drop the communication messages to reduce overall network congestion. The cloud computing platform attempts to maintain an agreed service level for the virtual machines when deciding to whether to drop the communication message.

[0042] Fig. 4 is a logic diagram that illustrates an exemplary method to discard virtual
 20 machine communication traffic. The virtual machine communication traffic may be governed by Transmission Control Protocol (TCP). The method initializes in step 410. In step 420, the cloud computing platform calculates an exponentially weighted average for the sending rate of a virtual machine. The exponentially weighted average sending rate (ASR) is evaluated for the virtual machine by the cloud computing platform, where
 25 $ASR = \alpha * previous_ASR + (1 - \alpha) * current_SR$, “ α ” is set by the cloud computing platform, the “previous_ASR” is the ASR previously evaluated for the virtual machines, and “current_SR” is the current sending rate detected by the cloud computing platform for the virtual machines.

[0043] In step 430, the cloud computing platform identifies the desired sending rate for
 30 the virtual machine. The desired sending rate (DSR) may be set in a service level agreement for the virtual machine. Alternatively, the desired sending rate may be set by an administrator of the cloud computing platform.

[0044] The cloud computing platform also evaluates a drop probability for a communication message sent from the virtual machine, in step 440. The drop probability (DP) is evaluated for the virtual machines by the cloud computing platform, where

$$DP = \frac{\beta * (ASR - DSR)}{ASR}, \text{ "}\beta\text{" is set by the cloud computing platform.}$$

5 [0045] In step 450, the cloud computing platform drops a communication message when a nonce generated, by the cloud computing platform, for the virtual machines is greater than the drop probability. The nonce (η) is randomly generated for the virtual machines by the

cloud computing platform, where $\eta = \frac{Random_Number}{Maximum_Random_Number}$,

"Random_Number" is generated by the cloud computing platform, and

10 "Maximum_Random_Number" is set by the cloud computing platform. In an embodiment, the communication message may be a packet. The method terminates in step 460.

[0046] In summary, the cloud computing platform shapes communication traffic associated with virtual machines executing on cloud devices. The cloud computing

15 platform monitors the communication rates and attempts to maintain an appropriate rate of communication for each virtual machine in the cloud computing platform. The service level agreements associated with the virtual machines are processed by the cloud computing platform to ensure an appropriate level of service is maintained for the virtual machines.

20 [0047] The foregoing descriptions of the embodiments of the invention are illustrative, and modifications in configuration and implementation are within the scope of the current description. For instance, while the embodiments of the invention are generally described with relation to FIGS. 1-4, those descriptions are exemplary. Although the subject matter has been described in language specific to structural features or methodological acts, it is
25 understood that the subject matter defined in the appended claims is not necessarily limited to the specific features or acts described above. Rather, the specific features and acts described above are disclosed as example forms of implementing the claims. The scope of the embodiment of the invention is accordingly intended to be limited only by the following claims.

CLAIMS

1. A computer-implemented method to manage virtual machine communication traffic in a cloud computing environment having multiple virtual machines, the method comprising:

5 maintaining an average communication rate for the virtual machines in the cloud computing platform;

computing a drop probability for the virtual machines as a function of the average communication rate and a desired communication rate for the virtual machines; and

10 randomly dropping a communication message for the virtual machines when a drop probability for the virtual machines is greater than a nonce generated by the cloud computing platform.

2. The computer-implemented method of claim 1, further comprising: limiting transmission communication messages or reception communication messages on the virtual machines to achieve the desired communication rates for each virtual machine in
15 the cloud computing platform.

3. The computer-implemented method of claim 2, wherein the limiting of transmission communication messages or reception communication messages is triggered when the cloud computing platform detects congestion on a network servicing the virtual machines.

20 4. The computer-implemented method of claim 3, wherein the limiting of the communication rates for the virtual machines is performed without queues.

5. The computer-implemented method of claim 1, wherein the average communication rate is an average sending rate (ASR) that is evaluated for the virtual machines by the cloud computing platform, wherein

25 $ASR = \alpha * previous_ASR + (1 - \alpha) * current_SR$, “ α ” is set by the cloud computing platform, the “previous_AS” is the ASR previously evaluated for the virtual machines, and “current_SR” is the current sending rate detected by the cloud computing platform for the virtual machines.

6. The computer-implemented method of claim 5, wherein the desired
30 communication rate is a desired sending rate (DSR) that is set in communication policies for the virtual machine.

7. The computer-implemented method of claim 6, wherein the drop probability (DP) is evaluated for the virtual machines by the cloud computing platform, wherein

$$DP = \frac{\beta * (ASR - DSR)}{ASR}, \text{ “}\beta\text{” is set by the cloud computing platform.}$$

8. The computer-implemented method of claim 1, wherein the nonce (η) is randomly
5 generated for the virtual machines by the cloud computing platform, wherein

$$\eta = \frac{Random_Number}{Maximum_Random_Number}, \text{ “Random_Number” is generated by the cloud}$$

computing platform; and “Maximum_Random_Number” is set by the cloud computing platform.

9. One or more computer-readable media storing instructions to perform a method to
10 drop virtual machine communication traffic, the method comprising:

calculating an exponentially weighted average for the sending rate of a virtual machine;

identifying the desired sending rate for the virtual machine;

evaluating a drop probability for a communication message sent from the virtual
15 machine; and

dropping a communication message when a nonce generated for the virtual machines is greater than the drop probability.

10. The computer-readable media of claim 9, wherein the communication message is a packet, and the virtual machine communication traffic is governed by a Transmission
20 Control Protocol (TCP).

11. The computer-readable media of claim 9, wherein the exponentially weighted average sending rate (ASR) is evaluated for the virtual machine by a cloud computing platform, wherein $ASR = \alpha * previous_ASR + (1 - \alpha) * current_SR$, “ α ” is set by the cloud computing platform, the “previous_ASR” is the ASR previously evaluated for the
25 virtual machines, and “current_SR” is the current sending rate detected by the cloud computing platform for the virtual machines.

12. The computer-readable media of claim 11, wherein the desired sending rate (DSR) is set in a service level agreement for the virtual machine.

13. The computer-readable media of claim 12, wherein the drop probability (DP) is
30 evaluated for the virtual machines by the cloud computing platform, wherein

$$DP = \frac{\beta * (ASR - DSR)}{ASR}, \text{ “}\beta\text{” is set by the cloud computing platform.}$$

14. The computer-readable media of claim 1, wherein the nonce (η) is randomly generated for the virtual machines by the cloud computing platform, wherein

$$\eta = \frac{\text{Random_Number}}{\text{Maximum_Random_Number}}, \text{ "Random_Number" is generated by the cloud}$$

computing platform; and "Maximum_Random_Number" is set by the cloud computing platform.

15. A cloud computing platform configured to shape communication traffic for virtual machines, the cloud computing platform comprising:

virtual machines that communicate over a network; and

a controller configured to drop communication messages associated with the

virtual machines when a nonce generated for the virtual machines is greater than the drop probability, wherein the controller evaluates the average sending rate (ASR) for the virtual machine, wherein $ASR = \alpha * \text{previous_ASR} + (1 - \alpha) * \text{current_SR}$, " α " is set by the controller, the "previous_AS" is the ASR previously evaluated for the virtual machines, and "current_SR" is the current sending rate detected by the controller for the virtual

machines, and the controller evaluates the drop probability (DP) for the virtual machines,

$$\text{wherein } DP = \frac{\beta * (ASR - DSR)}{ASR}, \text{ "}\beta\text{" is set by the controller, and desired sending rate}$$

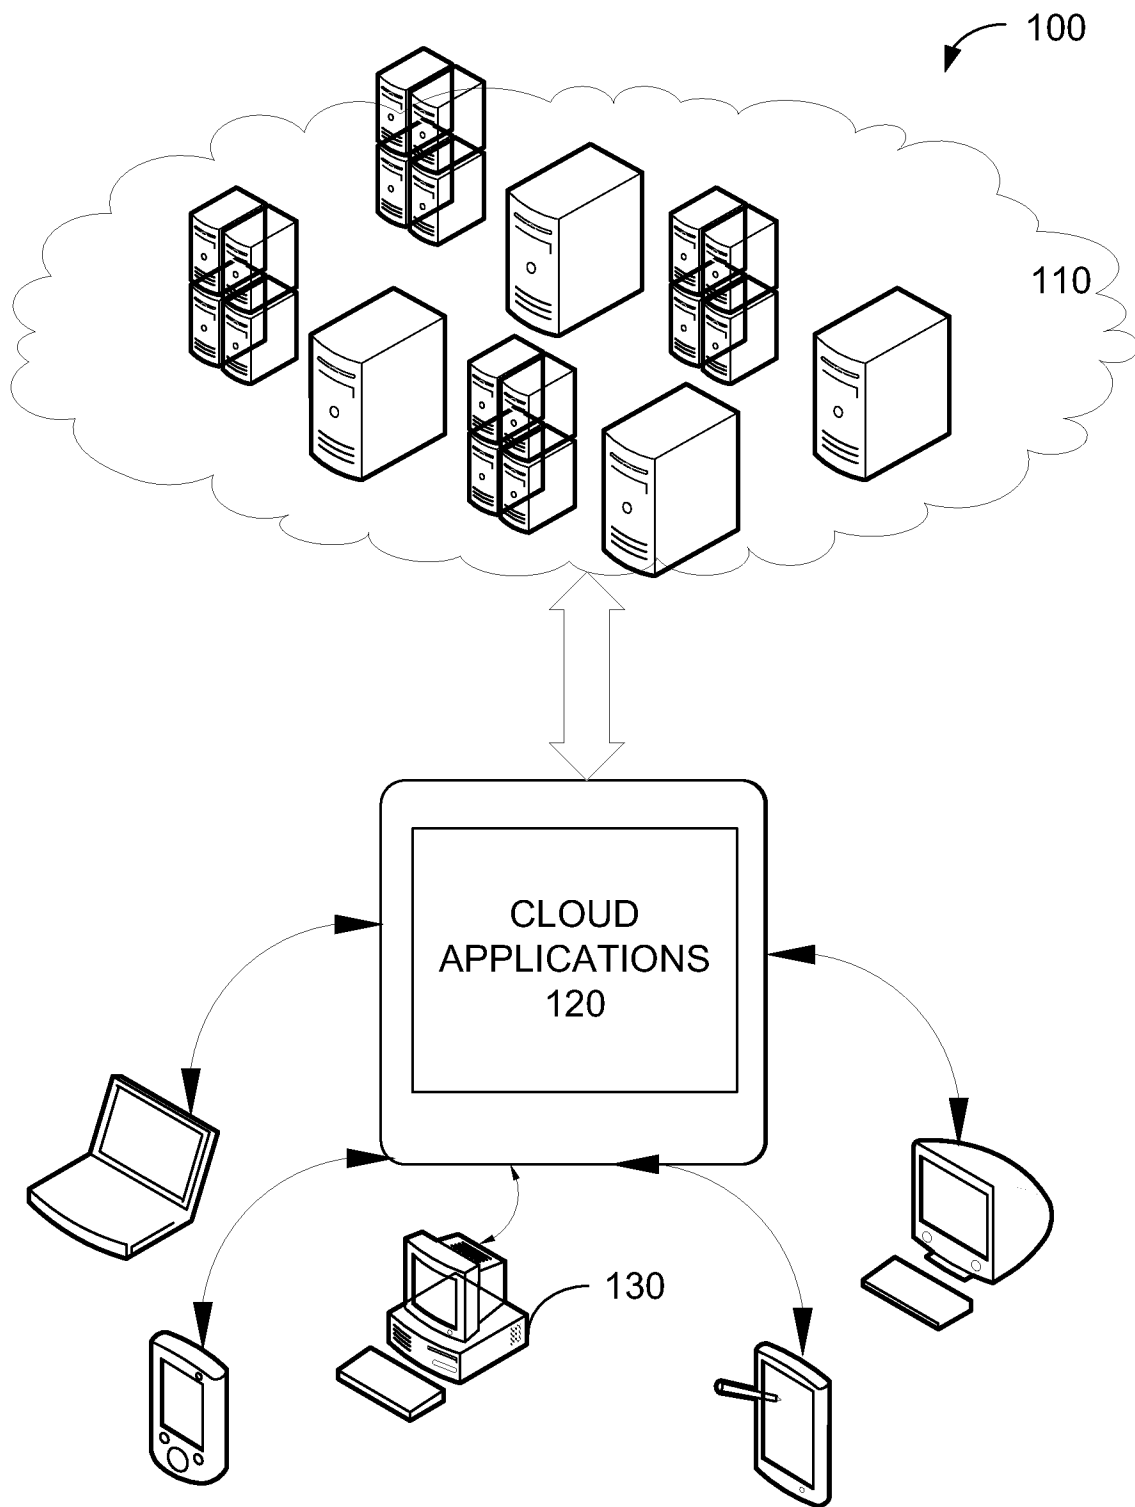
(DSR) is set in a service level agreement for the virtual machines and is used by the

controller if the network servicing the virtual machines is congested, and the controller

randomly generates the nonce (η), wherein $\eta = \frac{\text{Random_Number}}{\text{Maximum_Random_Number}},$

"Random_Number" is generated by the controller; and "Maximum_Random_Number" is set by the controller.

1/4

**FIG. 1**

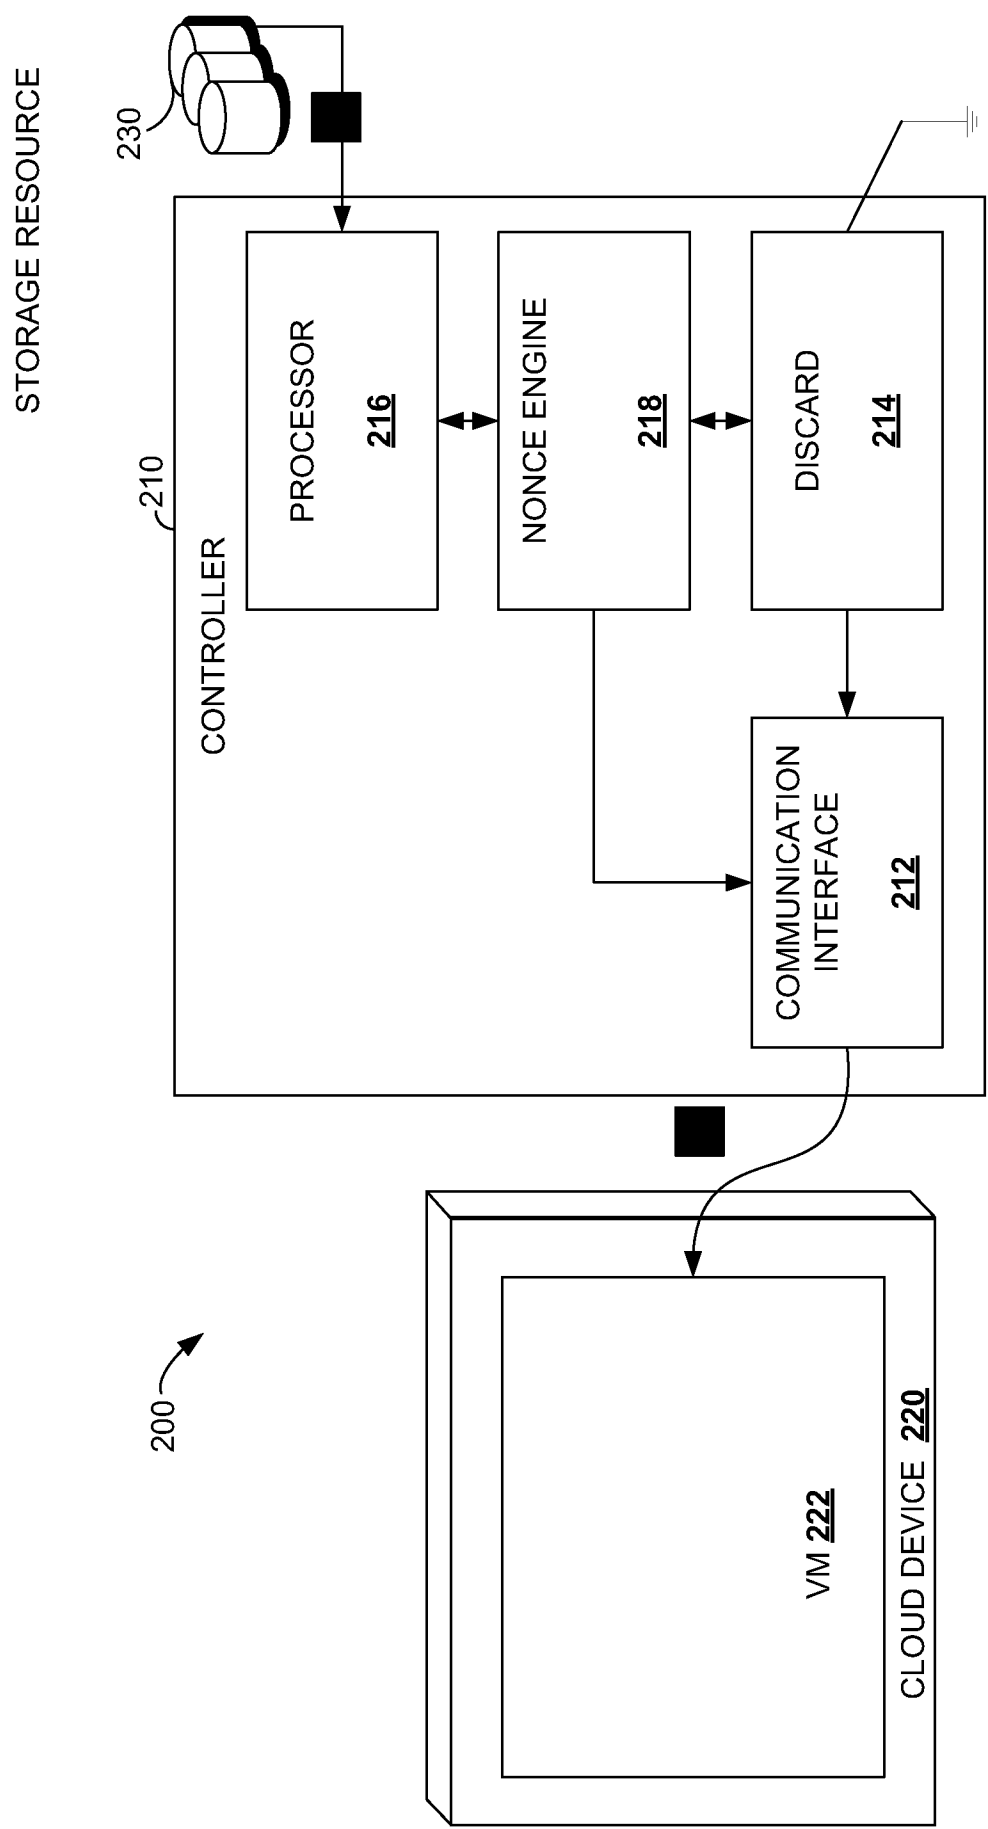
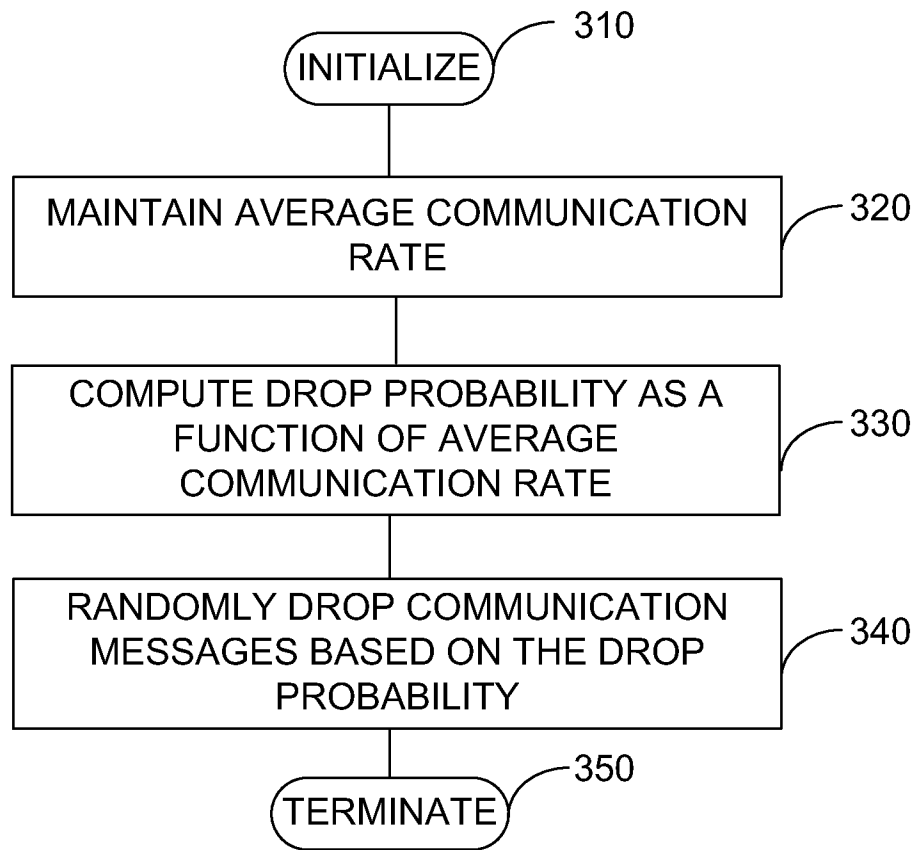
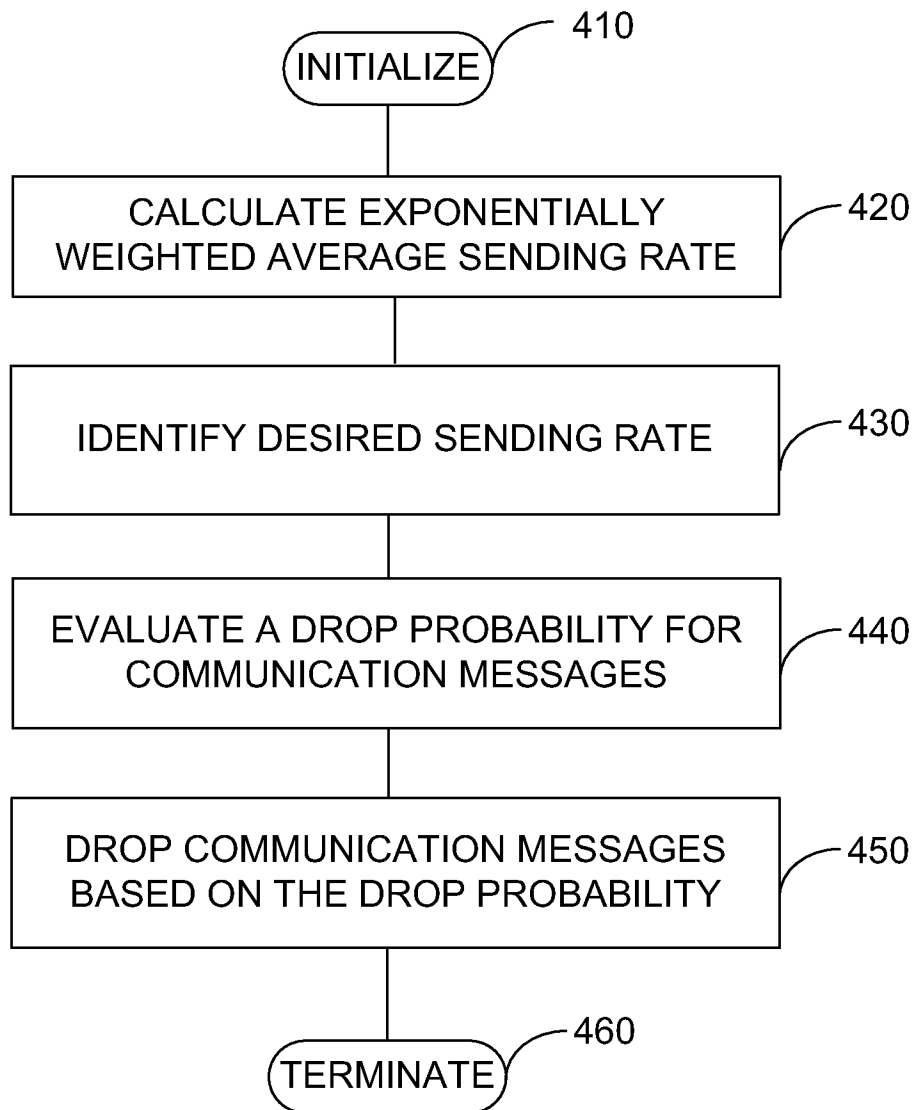


FIG. 2

3/4

**FIG. 3**

4/4

**FIG. 4**