



(51) International Patent Classification:
H04L 67/10 (2022.01) *H04L 12/66* (2006.01)

(21) International Application Number:
PCT/SG2022/050642

(22) International Filing Date:
06 September 2022 (06.09.2022)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
63/242,170 09 September 2021 (09.09.2021) US
17/681,348 25 February 2022 (25.02.2022) US

(71) Applicant: **LEMON INC.** [GB/SG]; P.O. Box 31119
Grand Pavilion, Hibiscus Way, 802 West Bay Road, Grand
Cayman, KY1-1205 (KY).

(72) Inventors: **WANG, Jialin**; 8 Marina View, Level 43,
Asia Square Tower 1, Singapore 018960 (SG). **CHEN,
Fangfei**; 12655 West Jefferson Boulevard, Sixth Floor,
Suite No. 137, Los Angeles, California 90066 (US).
WANG, Zengyao; 8 Marina View, Level 43, Asia Square
Tower 1, Singapore 018960 (SG). **CUI, Zhiqiang**; 12655
West Jefferson Boulevard, Sixth Floor, Suite No. 137, Los
Angeles, California 90066 (US).

(74) Agent: **POH, Chee Kian, Daniel**; Marks & Clerk Singa-
pore LLP, Tanjong Pagar Post Office, P O Box 636, Singa-
pore 910816 (SG).

(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, CV, 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,

(54) Title: MULTIFUNCTIONAL APPLICATION GATEWAY FOR SECURITY AND PRIVACY




Fig. 5

(57) Abstract: Systems and methods directed to an application gateway for security and privacy that supports security and compliance monitoring between production environments and virtual private clouds are described. In examples, the application gateway for security and privacy supports security and compliance logging making such information available to administrators and auditors; accordingly, the administrators and auditors can determine how the application gateway for security and privacy is behaving in a very detailed way. For example, by providing access to security and compliance logs, administrators and auditors can verify that the application gateway is not behaving in a malicious manner, such as but not limited communicating with an unauthorized host. In addition to including a user-friendly management interface that allows a user access to modify existing configurations in real-time, the application gateway for security and privacy may scale in a secure manner to support increasing and decreasing traffic demands.

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

Published:

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

MULTIFUNCTIONAL APPLICATION GATEWAY FOR SECURITY AND PRIVACY

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims priority to U. S. Provisional Application Ser. No. 63/242,170, filed September 9, 2021, and titled “MULTIFUNCTIONAL APPLICATION GATEWAY”, and to U. S. Application Ser. No. 17/681,348, filed February 25, 2022, and titled “MULTIFUNCTIONAL APPLICATION GATEWAY FOR SECURITY AND PRIVACY”, the disclosure of which are incorporated herein by reference in its entirety.

BACKGROUND

[0002] Datacenters are widely used to enhance service quality by running services and caching contents in cloud environments. Many companies adopt datacenters to reduce cost and further improve service quality; however, using datacenters may bring risks in user privacy. For example, user information including, but not limited to, user internet protocol addresses, usernames, passwords, payment information, and other user sensitive data may be identified, tracked, and/or otherwise compromised. Accordingly, such potential security issues may jeopardize the adoption of datacenters in regions governed by national and/or geographic user and data privacy regulations, where user data and user privacy are of top concerns.

[0003] Oftentimes, even if a central datacenter, such as a technology partner appointed by a local jurisdictional entity, is involved and provides a trusted infrastructure for a company to operate on or otherwise use, it is often difficult if not impossible for the technology partner to provide the needed secure infrastructure due to its wide distribution nature and use as well as the needed scalability for handling traffic loads of various sizes. For example, many native layer 7 load balancing solutions employed by many datacenters provide poor functionality when dealing with configuration changes and added capabilities. In examples, a load balancer of a datacenter may be controlled by a configuration file where changes to the file are made by modifying the file and then requiring a restart of the load

balancer - during which, the load balance is unavailable, may cause traffic losses, and may negatively impact business for the datacenter and for businesses utilizing the datacenter's services. Further, traditional load balancers support basic functionalities and may not fulfill the business needs in terms of growth for a particular company or product. Accordingly, a solution is lacking that integrates a scalable and trusted solution when working with trusted technology providers and further while adhering to user and data privacy laws and regulations.

SUMMARY

[0004] In accordance with examples of the present disclosure, a generalized solution and framework are described that provide a scalable application gateway to address secure computing and traffic routing needs of many datacenters. Such framework allows datacenters to continue to provide enhanced service quality in regions where user data and user privacy are of top concerns. Further, an application gateway for security and privacy, as part of a larger datacenter infrastructure, allows the datacenter to employ configuration and change management techniques to a production environment without restarting the application gateway for security and privacy as a whole such that the infrastructure can be updated without requiring much effort. The application gateway for security and privacy also provides additional support for features like service discovery, rate limiting, canary releasing, etc.

[0005] In accordance with examples of the present disclosure, the application gateway for security and privacy supports additional security and compliance features by filtering and modifying traffic using a web application firewall (WAF). The application gateway for security and privacy also supports multiple protocols including but not limited to QUIC, HTTP/2, and WebSocket over TLS encryption adding additional security layers to the communication infrastructure. Additionally, the application gateway for security and privacy also provides a logging and baseline metric system for administrators and users to audit and observe traffic in a comprehensive way.

[0006] In accordance with examples of the present disclosure, the application gateway for security and privacy may include a layer four load balancer, a plurality of processing pods including one or more processing containers, and a management backend configured to

provide a configuration implementation to the one or more processing containers, the configuration implementation specifying one or more processing parameters for each of the processing containers. In some examples, each processing pod of the plurality of processing pods includes one or more processing containers configured to perform layer seven network traffic processing on network traffic received from the layer four load balancer in accordance with the one or more processing parameters.

[0007] In accordance with examples of the present disclosure, a method of providing enhanced security using an application gateway for security and privacy is described. The method may include receiving network traffic at a layer four load balancer of the application gateway for security and privacy, determining, by a layer four load balancer, at least one processing pod to send a portion of the received network traffic for further processing, and routing, by the layer four load balancer, the portion of the received network traffic to the determined at least one processing pod. In examples, each processing pod of the plurality of processing pods includes one or more processing containers that is configured to perform layer seven network traffic processing on network traffic received from the layer four load balancer.

[0008] In accordance with examples of the present disclosure, a computer-readable storage medium is described. The computer-readable storage medium may include instructions, which when executed by a processor, causes the processor to: receive network traffic at a layer four load balancer of an application gateway for security and privacy, determine at least one processing pod to send a portion of the received network traffic for further processing, and route, by a layer four load balancer, the portion of the received network traffic to the determined at least one processing pod. In examples, each processing pod of the plurality of processing pods includes one or more processing containers that is configured to perform layer seven network traffic processing on network traffic received from the layer four load balancer.

[0009] Any of the one or more above aspects in combination with any other of the one or more aspects. Any of the one or more aspects as described herein.

[0010] This Summary is provided to introduce a selection of concepts in a simplified form, which is further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended

to be used to limit the scope of the claimed subject matter. Additional aspects, features, and/or advantages of examples will be set forth in part in the following description and, in part, will be apparent from the description, or may be learned by practice of the disclosure.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] Non-limiting and non-exhaustive examples are described with reference to the following Figures.

[0012] Fig. 1 depicts an example solution and framework that provides enhanced service quality in accordance with examples of the present disclosure.

[0013] Fig. 2 depicts an example of a generalized solution and framework that provides an application gateway for servicing incoming requests for a computing infrastructure in accordance with examples of the present disclosure.

[0014] Fig. 3 depicts additional details of an application gateway in accordance with examples of the present disclosure.

[0015] Fig. 4 depicts additional details of a web application firewall (WAF) in accordance with examples of the present disclosure.

[0016] Fig. 5 depicts a simplified method for routing traffic at a load balancer is described in accordance with examples of the present disclosure.

[0017] Fig. 6 depicts a simplified method for scaling a network processing ability of a load balancer is described in accordance with examples of the present disclosure.

[0018] Fig. 7 depicts a block diagram illustrating physical components (e.g., hardware) of a computing device with which aspects of the disclosure may be practiced.

DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS

[0019] In the following detailed description, references are made to the accompanying drawings that form a part hereof, and in which are shown by way of illustrations specific aspects or examples. These aspects may be combined, other aspects may be utilized, and structural changes may be made without departing from the present disclosure. Aspects may be practiced as methods, systems or devices. Accordingly, aspects may take the form of a hardware implementation, an entirely software implementation, or an implementation combining software and hardware aspects. The following detailed description is therefore

not to be taken in a limiting sense, and the scope of the present disclosure is defined by the appended claims and their equivalents.

[0020] In examples, a virtual private datacenter may include an application gateway for security and privacy to filter, screen, and/or modify one or more requests coming into the virtual private datacenter as well as the associated response that is attempting to leave the virtual private datacenter. The security and compliance features provided by the application gateway for security and privacy makes it easier for an administrator, user, and/or auditor to determine what the application gateway for security and privacy is doing in a very detailed way. The application gateway for security and privacy therefore, provides the users confidence that the configurations work well and provides auditors sufficient information that the application gateway for security and privacy is not doing anything malicious or otherwise servicing malicious requests.

[0021] In addition, the application gateway for security and privacy allows one to make changes to a production environment without restarting the component so the information and configuration used by the application gateway for security and privacy can be updated in a painless manner. The application gateway for security and privacy provides additional support for features including, but not limited to service discovery, rate limiting, canary releasing, etc. Further, the application gateway for security and privacy provides its users and administrators separate GUI control panels, so they won't need to write any configuration files but can make a change to the application gateway for security and privacy firewall with minimal interaction.

[0022] The application gateway for security and privacy allows users to operate a fleet of the resource by interacting with an application gateway for security and privacy manager. The application gateway for security and privacy manager includes a frontend and a backend. Thus, an auditor can access essential information for the application gateway for security and privacy by making a call to a management API or using the application gateway for security and privacy manager. While serving traffic, the application gateway for security and privacy can interact with the trusted virtual private datacenter to stream essential logging/auditing information to the auditors or storage areas associated with a location accessible to auditors.

[0023] The application gateway for security and privacy may include different components including a frontend user interface, a backend manager, a Kubernete or processing container, a controller, the Kubernetes or processing container pods or clusters, a layer 7 load balancer running upon the Kubernetes or processing containers, and a layer 4 load balancer and its associated controller, wherein layer 7 and layer 4 refer to respective layers of the Open Systems Interconnection (OSI) Model. In the OSI Model, communications between computing systems are split into seven different abstraction layers: Physical (layer 1), Data Link (layer 2), Network (layer 3), Transport (layer 4), Session (layer 5), Presentation (layer 6), and Application (layer 7). In examples, a Kubernetes describes a set of building blocks that collectively provide mechanisms to deploy, maintain, and scale applications based on CPU, memory, custom metrics, and combinations thereof.

[0024] The frontend user interface may comprise a platform that allows users and administrators direct access to the one or more configurations. For example, the front end user interface may provide a webpage that enables user the ability to configure the application gateway for security and privacy. The backend interacts with the frontend and the Kubernete or processing container controller; the backend may take the requests from the frontend and convert them into instructions and configurations to be implemented at the processing container controller and layer 7 load balancer. The processing container controller is configured to control the Kubernetes cluster and/or the processing container clusters and/or pods, thereby establishing or controlling the capacity of the application gateway. The processing container controller is responsible for creating pods that contain layer 7 load balancer instances that serve traffic. After the layer 7 load balancers are provisioned, they can fetch configurations and report their readiness to a layer 4 load balancer controller so that they can start serving traffic.

[0025] In examples, an existing layer 7 load balancer may be extended (e.g., via nginx plugin for example) such that each request/response, in plain text, can be processed by custom code or scripts written by a technology partner. This custom code may be in a Turing complete language and may have access to the full request/response including HTTP headers and body, in plain text relevant request metadata (e.g., source IP, timestamp, TCP and TLS session) and resolved routing information (e.g., destination IP, target microservice, etc.). The code may select some requests (and responses) and forward these (along with

other metadata) to an anomaly network traffic detection component for out-of-band analysis. The code may reject/drop some requests/responses; therefore, the code may run before a request is forwarded (but after routing decisions are resolved). In examples, the technology partner may update the code quickly, without restarting the load balancer. The CPU/memory capacity may be allocated for processing at the load balancer.

[0026] In examples, the code may have access to transient storage to allow state to be maintained from one request to the next; in addition, the code may have access to permanent storage for logging and pushing statistics. Further, the code may be able to initiate secure communication with other technology partner components and implement security and compliance features like logging and auditing, including WAF related features for filtering/modifying traffic. In examples, the scalability empowered by the processing container controller allows the layer 7 load balancer to scale up and scale down with ease. Further, a frontend and backend makes the configuration update for one or more components of the load balancer easier for users and administrators. In addition, the described system includes functionality to support multiple protocols including but not limited to QUIC/WebSocket/RPC. In accordance with examples of the present disclosure, the application gateway for security and privacy can be integrated with different layer 4 load balancers.

[0027] Fig. 1 depicts an example solution and framework that provides enhanced service quality in accordance with examples of the present disclosure. The framework depicted in Fig. 1 utilizes a datacenter, such as the virtual private datacenter B 102 to perform an initial processing of requests received from a user, such as user 104. For example, a user 104 may initiate a request for one or more application services 106 residing at a virtual private datacenter A 108. Such request may be sent from the user 104, via the internet 110, and received at the virtual private datacenter B 102. The virtual private datacenter B 102 may utilize an application gateway for security and privacy 112 to restrict, screen, and/or limit information flowing into and out of the virtual private datacenter A 108. In examples, the virtual datacenter A 108 may be a trusted technology provider and may be approved by one or more organizations, government entities, or the like. Alternatively, or in addition, the virtual private datacenter A 108 may be designated as complying with one or more data privacy and/or user privacy regulations or security levels as evidenced by a certification of

such compliance. For example, the virtual private datacenter A 108 may be audited on a regular or otherwise periodic basis to achieve and/or maintain such designation. In examples, the virtual private datacenter A 108 may provide one or more application services 106 which may include compute and/or storage services. As further depicted in Fig. 1, a request initiated from the virtual private datacenter A 108 may be sent through the third-party gateway 114, where the third-party gateway 114 may restrict, screen, and/or limit information flowing out of and into the virtual private datacenter A 108. In examples, the virtual private datacenter A 108 may be owned and/or operated by an entity or organization that is different from the virtual private datacenter B 102.

[0028] Fig. 2 depicts additional details of an example solution and framework that provides enhanced service quality in accordance with examples of the present disclosure. Similar to Fig. 1, the framework depicted in Fig. 2 utilizes a datacenter, such as the production virtual private datacenter 202 to perform an initial processing of requests received from a user, such as user 204. For example, a user 204 may initiate a request for one or more application services 206 residing at a platform virtual private datacenter 208. In examples, the platform virtual private datacenter 208 may be owned and/or operated by an entity or organization that is different from the production virtual private datacenter 202. Such request may be sent from the user 204, via the internet 210, and received at the production virtual private datacenter 202. The production virtual private datacenter 202 may utilize an application gateway for security and privacy 212 to restrict, screen, and/or limit information flowing into and out of the platform virtual private datacenter 208. In examples, the platform virtual datacenter 208 may be a trusted technology provider and may be approved by one or more organizations, jurisdictional entities, or the like. Alternatively, or in addition, the platform virtual private datacenter 208 may be designated as complying with one or more data privacy and/or user privacy regulations or security levels as evidenced by a certification of such compliance. For example, the platform virtual private datacenter 208 may be audited on a regular or otherwise periodic basis to achieve and/or maintain such designation. In examples, the platform virtual private datacenter 208 may provide one or more application services 206 which may include compute and/or storage services.

[0029] The application gateway for security and privacy 212 may include a layer 7 load balancer 218 and a management application programming interface (API) 216. The layer 7 load balancer 218 may be comprised of a scalable plurality of pods including a plurality of processing containers, or Kubernetes configured to filter, or screen, information flowing from the request provided by the user 204 and the platform virtual private datacenter 208. For example, the layer 7 load balancer 218 may scale up or scale down depending on an amount of traffic received by at the production virtual private datacenter 202. Moreover, one or more administrators 222 may initiate a configuration change by interfacing with the application gateway for security and privacy manager 220, where the application gateway manager 220 may cause a new pod configuration or policy to be utilized by the pod including a plurality of processing containers or Kubernetes, comprising the layer 7 load balancer 218.

[0030] As further depicted in Fig. 2, an auditor 224 may have access to logs, configuration files, and/or policy related information concerning the implementation of the application gateway for security and privacy 212. For example, the auditor 224 may view traffic logging information provided by each of the pods including a plurality of processing containers or Kubernetes, comprising the layer 7 load balancer 218. As another example, the auditor 224 may have view logs related to change requests performed at the application gateway for security and privacy manager 220 as well as other information available via logging and/or anomaly detection processing occurring at or otherwise stored at a trusted virtual private datacenter 226. In examples, the layer 7 load balancer 217 may perform aspects of anomaly detection and store such results at the trusted virtual datacenter 226.

[0031] Fig. 3 depicts additional details of the layer 7 load balancer 218 of Fig. 2 together with a change management/scaling example in accordance with examples of the present disclosure. More specifically, the production virtual private datacenter 302 may receive internet traffic, including one or more requests, and route the traffic to one or more components of the application gateway for security and privacy as described in Fig. 2. The application gateway for security and privacy may include a layer 7 load balancer as previously described. The layer 7 load balancer (e.g., 218) may include one or more components to perform layer 4 load balancing (e.g., managing and/or balancing transport-layer protocols, such as but not limited to TCP, UDP, DCCP, and SCTP used to control the

volume of data, where it is sent, and at what rate). For example, a layer 4 load balancer 304 may determine which pod 306A-306n to send received traffic for further processing. Each pod 306A-306n may include one or more Kubernetes, or processing containers 308, that performs traffic processing at the application (e.g., layer 7) level. Although each processing container 308 includes the same reference character (e.g., 308), each processing container may be configured to perform processing specific to one or more requirements of the pod 306A-306n. The number of processing containers 308 and pod to which the processing container belongs may be controlled by the processing container controller 310.

[0032] For example, an administrator 322 may interact with a frontend 318 manager to provide one or more commands to a backend manager 316 located in an internal platform virtual datacenter 314. In examples, and based on one or more parameters, such as but not limited to an amount of traffic, the processing container controller 310 may initiate the creation of a Kubernetes, or processing container 308, to be included in a new or existing pod 306A-306n. Once the processing container is created, the processing container may pull a processing configuration file from the backend manager 316. The processing configuration may include information indicating what kind of domain does the processing container server, what kind of protocol does the processing container server, what kind of origination policy does the processing container implement, and other detailed configuration information. Once created, the layer 4 load balancer controller 312 may register the newly created processing container with the layer 4 load balancer 304 such that the layer 4 load balancer 304 may include the new processing container as a resource to route traffic to.

[0033] In examples, each of the processing containers 308 may be configured to poll or pull the backend manager 316 on a periodic basis to determine if a new configuration file is available. Where a new configuration file is available, each of the processing containers 308 may perform an update at a time that is different from at least one other processing container 308, thereby ensuring a certain processing throughput is available and remains available. In addition, when traffic decreases to the point where not as many processing containers 308 are needed at one time, one or more processing containers maybe suspended or otherwise discarded.

[0034] Fig. 4 depicts additional details of each processing container 308 in accordance with examples of the present disclosure. Each processing container 408 may be the same as

or similar to the processing container 308 as previously described. Each processing container may include a web application firewall 412 configured to filter incoming and outgoing data. For example, the web application firewall may include policy/configuration information 414 indicating a protocol to process, a domain to process, a origin to process, and/or a destination to process. In examples, if an incoming request does not match the protocol, domain, origin, and/or destination, the WAF 412 may cause the request to be logged and then dropped. Alternatively, or in addition, if an incoming request does not match the protocol, domain, origin, and/or destination, the WAF 412 may cause the request to be routed to another pod and/or processing container 408 based on one or more of the protocol, domain, origin, and/or destination.

[0035] Referring now to Fig. 5, a simplified method for routing traffic at a load balancer is described in accordance with examples of the present disclosure. A general order for the steps of a method 500 is shown in Fig. 5. Generally, the method 500 starts at 502 and ends at 510. The method 500 may include more or fewer steps or may arrange the order of the steps differently than those shown in Fig. 5. The method 500 can be executed as a set of computer-executable instructions executed by a computer system and encoded or stored on a computer readable medium. In an illustrative aspect, the method 500 is executed by a computing device associated with an application gateway for security and privacy (e.g., 212). However, it should be appreciated that aspects of the method 500 may be performed by one or more processing devices, such as a computer, server, or other hardware element. Further, the method 500 can be performed by gates or circuits associated with a processor, Application Specific Integrated Circuit (ASIC), a field programmable gate array (FPGA), a system on chip (SOC), a neural processing unit, or other hardware device. Hereinafter, the method 500 shall be explained with reference to the systems, components, modules, software, data structures, user interfaces, etc. described in conjunction with Figs. 1-4.

[0036] The method 500 starts at 502, where flow may proceed to 504. At 504, internet traffic may be received, for example at a production virtual private data center (e.g., 302), where the traffic may include one or more requests, and the traffic may be routed to one or more components of the application gateway for security and privacy as described in Fig. 2. The application gateway for security and privacy may include a layer 7 load balancer as previously described. The layer 7 load balancer (e.g., 218) may include one or more

components to perform layer 4 load balancing (e.g., managing and/or balancing transport-layer protocols, such as but not limited to TCP, UDP, DCCP, and SCTP used to control the volume of data, where it is sent, and at what rate). At 506, the layer 4 load balancer may determine which pod (e.g., 306A-306n) to send received traffic for further processing. Each pod (e.g., 306A-306n) may include one or more Kubernetes, or processing containers 308, that performs traffic processing at the application (e.g., layer 7) level. Each processing container may be configured to perform processing specific to one or more requirements of the pod (e.g., 306A-306n). Thus, at 508, the determined pod may perform processing on the traffic. In examples, the method 500 may end at 510.

[0037] Referring now to Fig. 6, a simplified method for scaling a network processing ability of a load balancer is described in accordance with examples of the present disclosure. A general order for the steps of a method 600 is shown in Fig. 6. Generally, the method 600 starts at 602 and ends at 610. The method 600 may include more or fewer steps or may arrange the order of the steps differently than those shown in Fig. 6. The method 600 can be executed as a set of computer-executable instructions executed by a computer system and encoded or stored on a computer readable medium. In an illustrative aspect, the method 600 is executed by a computing device associated with an application gateway for security and privacy (e.g., 212). However, it should be appreciated that aspects of the method 500 may be performed by one or more processing devices, such as a computer, server, or other hardware element. Further, the method 600 can be performed by gates or circuits associated with a processor, Application Specific Integrated Circuit (ASIC), a field programmable gate array (FPGA), a system on chip (SOC), a neural processing unit, or other hardware device. Hereinafter, the method 600 shall be explained with reference to the systems, components, modules, software, data structures, user interfaces, etc. described in conjunction with Figs. 1-5.

[0038] The method 600 starts at 602, where flow may proceed to 604. At 604, one or more parameters, such as but not limited to an amount of traffic, associated with network traffic may be received. In examples, and based on one or more parameters, such as but not limited to an amount of traffic, the processing container controller may initiate the creation of a Kuberne^{te}, or processing container, to be included in a new or existing pod at 606. Once the processing container is created at 608, the processing container may pull a

processing configuration file from the backend manager. The processing configuration may include information indicating what kind of domain does the processing container server, what kind of protocol does the processing container server, what kind of origination policy does the processing container implement, and other detailed configuration information. Once created, at 610, a load balancer controller may register the newly created processing container with the load balancer such that the load balancer may include the new processing container as a resource to route traffic to. Method 600 may end at 612.

[0039] Fig. 7 is a block diagram illustrating physical components (e.g., hardware) of a processing device 700 with which aspects of the disclosure may be practiced. For example, the processing device 700 may represent one or more components as depicted in Figs. 1-4. In a basic configuration, the processing device 700 may include at least one processing unit 702 and a system memory 704. Depending on the configuration and type of computing device, the system memory 704 may comprise, but is not limited to, volatile storage (e.g., random access memory), non-volatile storage (e.g., read-only memory), flash memory, or any combination of such memories.

[0040] The system memory 704 may include an operating system 705 and one or more program modules 706 suitable for performing the various aspects disclosed herein such. The operating system 705, for example, may be suitable for controlling the operation of the processing device 700. Furthermore, aspects of the disclosure may be practiced in conjunction with other operating systems, or any other application program and is not limited to any particular application or system. This basic configuration is illustrated in Fig. 7 by those components within a dashed line 708. The processing device 700 may have additional features or functionality. For example, the processing device 700 may also include additional data storage devices (removable and/or non-removable) such as, for example, magnetic disks, optical disks, or tape. Such additional storage is illustrated in Fig. 7 by a removable storage device 709 and a non-removable storage device 710.

[0041] As stated above, several program modules and data files may be stored in the system memory 704. While executing on the at least one processing unit 702, the application 707 or the program modules 706 may perform processes including, but not limited to, one or more aspects, as described herein. The application 720 includes a layer 4 load balancer 723, a processing container controller 725, and a layer 4 load balancer

controller 727, as described in more detail with regard to Figs. 1-3. Other program modules that may be used in accordance with aspects of the present disclosure may include electronic mail and contacts applications, word processing applications, spreadsheet applications, database applications, slide presentation applications, drawing or computer-aided application programs, etc., and/or one or more components supported by the systems described herein.

[0042] Furthermore, aspects of the disclosure may be practiced in an electrical circuit comprising discrete electronic elements, packaged or integrated electronic chips containing logic gates, a circuit utilizing a microprocessor, or on a single chip containing electronic elements or microprocessors. For example, aspects of the disclosure may be practiced via a system-on-a-chip (SOC) where each or many of the components illustrated in Fig. 7 may be integrated onto a single integrated circuit. Such an SOC device may include one or more processing units, graphics units, communications units, system virtualization units and various application functionality all of which are integrated (or “burned”) onto the chip substrate as a single integrated circuit. When operating via an SOC, the functionality, described herein, with respect to the capability of client to switch protocols may be operated via application-specific logic integrated with other components of the processing device 700 on the single integrated circuit (chip). Aspects of the disclosure may also be practiced using other technologies capable of performing logical operations such as, for example, AND, OR, and NOT, including but not limited to mechanical, optical, fluidic, and quantum technologies. In addition, aspects of the disclosure may be practiced within a general-purpose computer or in any other circuits or systems.

[0043] The processing device 700 may also have one or more input device(s) 712 such as a keyboard, a mouse, a pen, a sound or voice input device, a touch or swipe input device, etc. The output device(s) 714A such as a display, speakers, a printer, etc. may also be included. An output 714B, corresponding to a virtual display may also be included. The aforementioned devices are examples and others may be used. The processing device 700 may include one or more communication connections 716 allowing communications with other computing or processing devices 750. Examples of suitable communication connections 716 include, but are not limited to, radio frequency (RF) transmitter, receiver, and/or transceiver circuitry; universal serial bus (USB), parallel, and/or serial ports.

[0044] The term computer readable media as used herein may include computer storage media. Computer storage media may include volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information, such as computer readable instructions, data structures, or program modules. The system memory 704, the removable storage device 709, and the non-removable storage device 710 are all computer storage media examples (e.g., memory storage). Computer storage media may include RAM, ROM, electrically erasable read-only memory (EEPROM), flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other article of manufacture which can be used to store information and which can be accessed by the processing device 700. Any such computer storage media may be part of the processing device 700. Computer storage media does not include a carrier wave or other propagated or modulated data signal.

[0045] Communication media may be embodied by computer readable instructions, data structures, program modules, or other data in a modulated data signal, such as a carrier wave or other transport mechanism, and includes any information delivery media. The term “modulated data signal” may describe a signal that has one or more characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication media may include wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, radio frequency (RF), infrared, and other wireless media.

[0046] In addition, the aspects and functionalities described herein may operate over distributed systems (e.g., cloud-based computing systems), where application functionality, memory, data storage and retrieval and various processing functions may be operated remotely from each other over a distributed computing network, such as the Internet or an intranet. User interfaces and information of various types may be displayed via on-board computing device displays or via remote display units associated with one or more computing devices. For example, user interfaces and information of various types may be displayed and interacted with. Interaction with the multitude of computing systems with which embodiments of the invention may be practiced include, keystroke entry, touch screen entry, voice or other audio entry, gesture entry where an associated computing device

is equipped with detection (e.g., camera) functionality for capturing and interpreting user gestures for controlling the functionality of the computing device, and the like.

[0047] The phrases “at least one,” “one or more,” “or,” and “and/or” are open-ended expressions that are both conjunctive and disjunctive in operation. For example, each of the expressions “at least one of A, B and C,” “at least one of A, B, or C,” “one or more of A, B, and C,” “one or more of A, B, or C,” “A, B, and/or C,” and “A, B, or C” means A alone, B alone, C alone, A and B together, A and C together, B and C together, or A, B and C together.

[0048] The term “a” or “an” entity refers to one or more of that entity. As such, the terms “a” (or “an”), “one or more,” and “at least one” can be used interchangeably herein. It is also to be noted that the terms “comprising,” “including,” and “having” can be used interchangeably.

[0049] The term “automatic” and variations thereof, as used herein, refers to any process or operation, which is typically continuous or semi-continuous, done without material human input when the process or operation is performed. However, a process or operation can be automatic, even though performance of the process or operation uses material or immaterial human input, if the input is received before performance of the process or operation. Human input is deemed to be material if such input influences how the process or operation will be performed. Human input that consents to the performance of the process or operation is not deemed to be “material.”

[0050] Any of the steps, functions, and operations discussed herein can be performed continuously and automatically.

[0051] The exemplary systems and methods of this disclosure have been described in relation to computing devices. However, to avoid unnecessarily obscuring the present disclosure, the preceding description omits several known structures and devices. This omission is not to be construed as a limitation. Specific details are set forth to provide an understanding of the present disclosure. It should, however, be appreciated that the present disclosure may be practiced in a variety of ways beyond the specific detail set forth herein.

[0052] Furthermore, while the exemplary aspects illustrated herein show the various components of the system collocated, certain components of the system can be located

remotely, at distant portions of a distributed network, such as a LAN and/or the Internet, or within a dedicated system. Thus, it should be appreciated, that the components of the system can be combined into one or more devices, such as a server, communication device, or collocated on a particular node of a distributed network, such as an analog and/or digital telecommunications network, a packet-switched network, or a circuit-switched network. It will be appreciated from the preceding description, and for reasons of computational efficiency, that the components of the system can be arranged at any location within a distributed network of components without affecting the operation of the system.

[0053] Furthermore, it should be appreciated that the various links connecting the elements can be wired or wireless links, or any combination thereof, or any other known or later developed element(s) that is capable of supplying and/or communicating data to and from the connected elements. These wired or wireless links can also be secure links and may be capable of communicating encrypted information. Transmission media used as links, for example, can be any suitable carrier for electrical signals, including coaxial cables, copper wire, and fiber optics, and may take the form of acoustic or light waves, such as those generated during radio-wave and infra-red data communications.

[0054] While the flowcharts have been discussed and illustrated in relation to a particular sequence of events, it should be appreciated that changes, additions, and omissions to this sequence can occur without materially affecting the operation of the disclosed configurations and aspects.

[0055] Several variations and modifications of the disclosure can be used. It would be possible to provide for some features of the disclosure without providing others.

[0056] In yet another configurations, the systems and methods of this disclosure can be implemented in conjunction with a special purpose computer, a programmed microprocessor or microcontroller and peripheral integrated circuit element(s), an ASIC or other integrated circuit, a digital signal processor, a hard-wired electronic or logic circuit such as discrete element circuit, a programmable logic device or gate array such as PLD, PLA, FPGA, PAL, special purpose computer, any comparable means, or the like. In general, any device(s) or means capable of implementing the methodology illustrated herein can be used to implement the various aspects of this disclosure. Exemplary hardware that can be used for the present disclosure includes computers, handheld devices, telephones (e.g., cellular,

Internet enabled, digital, analog, hybrids, and others), and other hardware known in the art. Some of these devices include processors (e.g., a single or multiple microprocessors), memory, nonvolatile storage, input devices, and output devices. Furthermore, alternative software implementations including, but not limited to, distributed processing or component/object distributed processing, parallel processing, or virtual machine processing can also be constructed to implement the methods described herein.

[0057] In yet another configuration, the disclosed methods may be readily implemented in conjunction with software using object or object-oriented software development environments that provide portable source code that can be used on a variety of computer or workstation platforms. Alternatively, the disclosed system may be implemented partially or fully in hardware using standard logic circuits or VLSI design. Whether software or hardware is used to implement the systems in accordance with this disclosure is dependent on the speed and/or efficiency requirements of the system, the particular function, and the particular software or hardware systems or microprocessor or microcomputer systems being utilized.

[0058] In yet another configuration, the disclosed methods may be partially implemented in software that can be stored on a storage medium, executed on programmed general-purpose computer with the cooperation of a controller and memory, a special purpose computer, a microprocessor, or the like. In these instances, the systems and methods of this disclosure can be implemented as a program embedded on a personal computer such as an applet, JAVA® or CGI script, as a resource residing on a server or computer workstation, as a routine embedded in a dedicated measurement system, system component, or the like. The system can also be implemented by physically incorporating the system and/or method into a software and/or hardware system.

[0059] The disclosure is not limited to standards and protocols if described. Other similar standards and protocols not mentioned herein are in existence and are included in the present disclosure. Moreover, the standards and protocols mentioned herein, and other similar standards and protocols not mentioned herein are periodically superseded by faster or more effective equivalents having essentially the same functions. Such replacement standards and protocols having the same functions are considered equivalents included in the present disclosure.

[0060] The present disclosure, in various configurations and aspects, includes components, methods, processes, systems and/or apparatus substantially as depicted and described herein, including various combinations, subcombinations, and subsets thereof. Those of skill in the art will understand how to make and use the systems and methods disclosed herein after understanding the present disclosure. The present disclosure, in various configurations and aspects, includes providing devices and processes in the absence of items not depicted and/or described herein or in various configurations or aspects hereof, including in the absence of such items as may have been used in previous devices or processes, e.g., for improving performance, achieving ease, and/or reducing cost of implementation.

[0061] The present disclosure relates to an application gateway for security and privacy according to at least the examples provided in the sections below:

[0062] (A1) In one aspect, some examples include an application gateway for security and privacy. The application gateway for security and privacy may include a layer four load balancer; a plurality of processing pods including one or more processing containers; and a management backend configured to provide a configuration implementation to the one or more processing containers, the configuration implementation specifying one or more processing parameters for each of the processing containers, wherein each processing pod of the plurality of processing pods includes one or more processing containers configured to perform layer seven network traffic processing on network traffic received from the layer four load balancer in accordance with the one or more processing parameters.

[0063] (A2) In some examples of A1, the layer four load balancer is configured to receive network traffic, determine at least one processing pod to send a portion of the received network traffic for further processing, and route the portion of the received network traffic to the determined at least one processing pod.

[0064] (A3) In some examples of A1-A2, the application gateway for security and privacy further includes a processing container controller, wherein the processing container controller is configured to receive one or more parameters associated with network traffic and based on the one or more parameters, determine to create a new processing container to be included in a pod.

[0065] (A4) In some examples of A1-A3, the processing container controller is configured to cause a new processing container to be created.

[0066] (A5) In some examples of A1-A4, the application gateway for security and privacy further includes a load balancer controller configured to register the new processing container with the layer four load balancer.

[0067] (A6) In some examples of A1-A5, the one or more processing containers are configured to interrogate the management backend and determine if a new configuration implementation is available, wherein when a new configuration implementation is available, each of the processing containers is configured to perform an update to the new configuration implementation at a time that is different from at least one other processing container.

[0068] (A7) In some examples of A1-A6, the application gateway for security and privacy is included in a system for providing enhanced security at a virtual private datacenter, the system including: a virtual private datacenter including the application gateway for security and privacy; and a technology partner datacenter including application services configured to provide requested data based on a request received at the virtual private datacenter.

[0069] (B1) In one aspect, some examples include a method of providing enhanced security using an application gateway for security and privacy. The method may include receiving network traffic at a layer four load balancer of the application gateway for security and privacy; determining, by a layer four load balancer, at least one processing pod of a plurality of processing pods to send a portion of the received network traffic for further processing; and routing, by the layer four load balancer, the portion of the received network traffic to the determined at least one processing pod, wherein each processing pod of the plurality of processing pods includes one or more processing containers that is configured to perform layer seven network traffic processing on network traffic received from the layer four load balancer.

[0070] (B2) In some examples of B1, the method further includes providing a configuration implementation to the one or more processing containers, the configuration implementation specifying one or more processing parameters for one or more of the processing containers; and performing layer seven network traffic processing on network

traffic received from the layer four load balancer in accordance with the one or more processing parameters.

[0071] (B3) In some examples of B1-B2, the method further includes receiving one or more parameters associated with network traffic; and based on the one or more parameters, determining to create a new processing container to be included in a pod.

[0072] (B4) In some examples of B1-B3, the method further includes causing, by a processing container controller, a new processing container to be created.

[0073] (B5) In some examples of B1-B4, the method further includes registering, by a load balancer controller, the new processing container with the layer four load balancer.

[0074] (B6) In some examples of B1-B5, the method further includes determining if a new configuration implementation is available; and if a new configuration implementation is available, causing each of the processing containers to perform an update to the new configuration implementation at a time that is different from at least one other processing container.

[0075] (B7) In some examples of B1-B6, the processing container is a Kubernetes container.

[0076] In yet another aspect, some examples include a non-transitory computer-readable storage medium storing one or more programs for execution by one or more processors of a storage device, the one or more programs including instructions for performing any of the methods described herein (e.g., B1-B7 described above).

[0077] Aspects of the present disclosure, for example, are described above with reference to block diagrams and/or operational illustrations of methods, systems, and computer program products according to aspects of the disclosure. The functions/acts noted in the blocks may occur out of the order as shown in any flowchart or process. For example, two blocks shown in succession may in fact be executed substantially concurrently or the blocks may sometimes be executed in the reverse order, depending upon the functionality/acts involved.

[0078] The description and illustration of one or more aspects provided in this application are not intended to limit or restrict the scope of the disclosure as claimed in any way. The aspects, examples, and details provided in this application are considered sufficient to

convey possession and enable others to make and use the best mode of claimed disclosure. The claimed disclosure should not be construed as being limited to any aspect, example, or detail provided in this application. Regardless of whether shown and described in combination or separately, the various features (both structural and methodological) are intended to be selectively included or omitted to produce an embodiment with a particular set of features. Having been provided with the description and illustration of the present application, one skilled in the art may envision variations, modifications, and alternate aspects falling within the spirit of the broader aspects of the general inventive concept embodied in this application that do not depart from the broader scope of the claimed disclosure.

WHAT IS CLAIMED IS:

1. An application gateway for security and privacy comprising:
a layer four load balancer;
a plurality of processing pods including one or more processing containers; and
a management backend configured to provide a configuration implementation to the one or more processing containers, the configuration implementation specifying one or more processing parameters for each of the processing containers,
wherein each processing pod of the plurality of processing pods includes one or more processing containers configured to perform layer seven network traffic processing on network traffic received from the layer four load balancer in accordance with the one or more processing parameters.
2. The application gateway for security and privacy of claim 1, wherein the layer four load balancer is configured to receive network traffic, determine at least one processing pod to send a portion of the received network traffic for further processing, and route the portion of the received network traffic to the determined at least one processing pod.
3. The application gateway for security and privacy of claim 1, further comprising a processing container controller, wherein the processing container controller is configured to receive one or more parameters associated with network traffic and based on the one or more parameters, determine to create a new processing container to be included in a pod.
4. The application gateway for security and privacy of claim 3, wherein the processing container controller is configured to cause a new processing container to be created.
5. The application gateway for security and privacy of claim 4, further comprising a load balancer controller configured to register the new processing container with the layer four load balancer.

6. The application gateway for security and privacy of claim 1, wherein the one or more processing containers are configured to interrogate the management backend and determine if a new configuration implementation is available, wherein when a new configuration implementation is available, each of the processing containers is configured to perform an update to the new configuration implementation at a time that is different from at least one other processing container.

7. The application gateway for security and privacy of claim 1, wherein the application gateway for security and privacy is included in a system for providing enhanced security at a virtual private datacenter, the system including:

- a virtual private datacenter including the application gateway for security and privacy; and

- a technology partner datacenter including application services configured to provide requested data based on a request received at the virtual private datacenter.

8. A method of providing enhanced security using an application gateway for security and privacy, the method comprising:

- receiving network traffic at a layer four load balancer of the application gateway for security and privacy;

- determining, by a layer four load balancer, at least one processing pod of a plurality of processing pods to send a portion of the received network traffic for further processing; and

- routing, by the layer four load balancer, the portion of the received network traffic to the determined at least one processing pod,

- wherein each processing pod of the plurality of processing pods includes one or more processing containers that is configured to perform layer seven network traffic processing on network traffic received from the layer four load balancer.

9. The method of claim 8, further comprising:

providing a configuration implementation to the one or more processing containers, the configuration implementation specifying one or more processing parameters for one or more of the processing containers; and

performing layer seven network traffic processing on network traffic received from the layer four load balancer in accordance with the one or more processing parameters.

10. The method of claim 9, further comprising:
receiving one or more parameters associated with network traffic; and
based on the one or more parameters, determining to create a new processing container to be included in a pod.

11. The method of claim 10, further comprising causing, by a processing container controller, a new processing container to be created.

12. The method of claim 11, further comprising registering, by a load balancer controller, the new processing container with the layer four load balancer.

13. The method of claim 8, further comprising:
determining if a new configuration implementation is available; and
if a new configuration implementation is available, causing each of the processing containers to perform an update to the new configuration implementation at a time that is different from at least one other processing container.

14. The method of claim 8, wherein the processing container is a Kubernetes container.

15. A computer-readable storage medium including instructions, which when executed by a processor, causes the processor to:
receive network traffic at a layer four load balancer of an application gateway for security and privacy;

determine, by a layer four load balancer, at least one processing pod of a plurality of processing pods to send a portion of the received network traffic for further processing; and route, by a layer four load balancer, the portion of the received network traffic to the determined at least one processing pod,

wherein each processing pod of the plurality of processing pods includes one or more processing containers that is configured to perform layer seven network traffic processing on network traffic received from the layer four load balancer.

16. The computer-readable storage medium of claim 15, wherein the instructions cause the processor to:

provide a configuration implementation to the one or more processing containers, the configuration implementation specifying one or more processing parameters for one or more of the processing containers; and

perform layer seven network traffic processing on network traffic received from the layer four load balancer in accordance with the one or more processing parameters.

17. The computer-readable storage medium of claim 16, wherein the instructions cause the processor to:

receive one or more parameters associated with network traffic; and

based on the one or more parameters, determine to create a new processing container to be included in a pod.

18. The computer-readable storage medium of claim 17, wherein the instructions cause the processor to create a new processing container.

19. The computer-readable storage medium of claim 18, wherein the instructions cause the processor to register the new processing container with the layer four load balancer.

20. The computer-readable storage medium of claim 15, wherein the instructions cause the processor to:

determine if a new configuration implementation is available; and
if a new configuration implementation is available, cause each of the processing containers to perform an update to the new configuration implementation at a time that is different from at least one other processing container.




Fig. 1




Fig. 2




Fig. 3







Fig. 4

**Fig. 5**

**Fig. 6**

**Fig. 7**