(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
04 July 2019 (04.07.2019)

WIPO | PCT

(10) International Publication Number
WO 2019/133434 A1

(51) International Patent Classification:
H04L 9/32 (2006.01) *H04L 29/06* (2006.01)
H04L 9/08 (2006.01) *H04L 12/931* (2013.01)

(21) International Application Number:
PCT/US2018/066801

(22) International Filing Date:
20 December 2018 (20.12.2018)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/612,093 29 December 2017 (29.12.2017) US

(71) Applicant: **PENSANDO SYSTEMS INC.** [US/US]; 1730 Technology Dr., Ste. 202, San Jose, California 95110 (US).

(72) Inventors: **JAIN, Vipin**; 1730 Technology Dr., Ste. 202, San Jose, California 95110 (US). **GADDE, Ravi Kumar**; 1730 Technology Dr., Ste. 202, San Jose, California 95110 (US). **SCHIATTARELLA, Enrico**; 1730 Technology Dr., Ste. 202, San Jose, California 95110 (US). **HALEMANE, Sukhesh**; 1730 Technology Dr., Ste. 202, San Jose, California 95110 (US).

(74) Agent: **BURKETTE, Scott L.**; WILSON SONSINI GOODRICH & ROSATI, 650 Page Mill Road, Palo Alto, California 94304-1050 (US).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(54) Title: METHODS AND SYSTEMS FOR CRYPTOGRAPHIC IDENTITY BASED NETWORK MICROSEGMENTATION

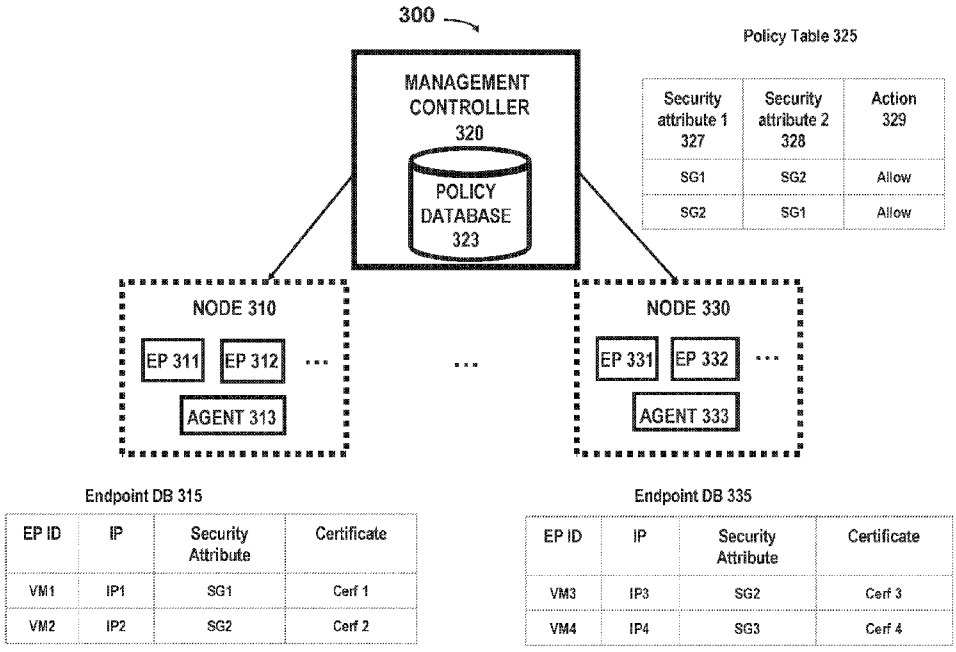

FIG. 3

(57) Abstract: Methods and network interface devices for establishing a secure and authenticated network connection are provided. The method comprises: receiving, from a requesting entity, a destination IP address and a first certificate that is used to establish a secure network connection, wherein the first certificate comprises a first security attribute that is associated with a source destination IP address; identifying, with aid of one or more processors, a stored second security attribute associated with the destination IP address; and determining, with aid of the one or more processors, a policy action based at least in part on the first security attribute and the second security attribute.

[Continued on next page]

WO 2019/133434 A1

(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).

Declarations under Rule 4.17:

— *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

— *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

— *with international search report (Art. 21(3))*

METHODS AND SYSTEMS FOR CRYPTOGRAPHIC IDENTITY BASED NETWORK MICROSEGMENTATION

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the priority and benefit of U.S. Provisional Application No. 62/612,093 filed on December 29, 2017, the entire contents of which are incorporated herein by reference.

BACKGROUND

[0002] Microsegmentaion has been used to enhance security and efficiency of a network. For example, network traffic of a cluster which is trusted, may be segregated from network traffic of another cluster which is untrusted. A cluster may be a physical system such as an organization, department, a node and the like. A cluster may be any virtual network segment. A network device of the cluster can be configured to route network traffic separately for each network segment.

SUMMARY

[0003] Traditionally, network is segmented based on IP address. IP-based microsegmentation may enable forwarding and security policies for network traffic between clusters or sites based on IP addresses. However, this approach requires exchanging IP endpoint information across clusters and sites in a changing environment. For example, data center topologies may be constantly changing: networks are re-numbered, server pools are expanded, or workloads are moved. In this case, the security policy configured when a workload was first deployed may not be enforceable if the definition of the security policy is associated with the IP address, port, or protocol. The difficult challenge is exacerbated when the workloads move from one cluster (e.g., data center) to another.

[0004] In light of the above, it would be desirable to provide a method and system for network micro segmentation with improved scalability and flexibility. The present invention addresses this need and provides related advantages as well.

[0005] Accordingly, in one aspect, a method for establishing a secure and authenticated network connection is provided. The method may comprise: (a) receiving, from a requesting entity, a destination IP address and a first certificate that is used to establish a secure network connection, wherein the first certificate comprises a first security attribute that is associated with a source destination IP address; (b) identifying, with aid of one or more processors, a stored second security attribute associated with the destination IP address; and (c) determining, with aid of the one or more processors, a policy action based at least in part on the first security

attribute and the second security attribute.

[0006] In some embodiments, the method further comprises transmitting, from the requesting entity, a certificate signing request (CSR) to an intermediate certificate authority, wherein the CSR comprising a security attribute and an associated IP address. In some cases, a mapping relationship between the security attribute and the IP address is received from a computing device running the intermediate certificate authority. In some cases, the mapping relationship is at least in part managed by the intermediate certificate authority. In some cases, the method further comprises receiving a certificate from the intermediate certificate authority, wherein the certificate comprises the security attribute. In some cases, the method further comprises storing the certificate in a local database managed by the requesting entity. In some cases, the intermediate certificate authority is verified by a root certificate authority.

[0007] In some embodiments, the method further comprises identifying, with aid of the one or more processors, a second stored certificate associated with the destination IP address. In some cases, the method further comprises transmitting the second stored certificate to the requesting entity. In some cases, the first certificate and the second certificate are issued by the same intermediate certificate authority. Alternatively, the first certificate and the second certificate are issued by different intermediate certificate authorities. In some instances, the different intermediate certificate authorities are verified by the same root certificate authority. In some cases, the second security attribute comprises a security group the destination IP address is associated therewith.

[0008] In some embodiments, the method further comprises executing the determined policy action on a base system logic or by the requesting entity. In some cases, the base system logic comprises a bare metal server, a hypervisor or a docker base for controlling one or more virtual machines or containers. In some embodiments, the policy action comprises at least one of the following: allow, block, rate-limit, mark or alert. In some embodiments, the first security attribute comprises a security group the source IP address is associated therewith.

[0009] In a related yet separate aspect, a network interface device is provided. The device may comprise: (i) a memory for storing a set of instructions , and (ii) one or more processors configured to execute the set of instructions to: (a) receive, from a requesting entity, a destination IP address and a first certificate that is used to establish a secure network connection, wherein the first certificate comprises a first security attribute that is associated with a source IP address; (b) identify a stored second security attribute associated with the destination IP address; and (c) determine a policy action based at least in part on the first security attribute and the second security attribute.

[0010] In some embodiments, the one or more processors are further configured to generate and transmit a certificate signing request (CSR) to an intermediate certificate authority, wherein the CSR comprising a security attribute and an associated IP address. In some cases, a mapping relationship between the security attribute and the IP address is received from a computing device running the intermediate certificate authority. In some cases, the mapping relationship is at least in part managed by the intermediate certificate authority. In some cases, the one or more processors are further configured to receive a signed certificate from the intermediate certificate authority, and wherein the signed certificate comprises the security attribute. For example, the one or more processors are further configured to store the signed certificate in a local database managed by the network interface device. In some cases, the intermediate certificate authority is verified by a root certificate authority.

[0011] In some embodiments, the one or more processors are further configured to identity a second stored certificate associated with the destination IP address. In some cases, the one or more processors are further configured to transmit the second stored certificate to the requesting entity. In some cases, the first certificate and the second certificate are issued by the same intermediate certificate authority. Alternatively, the first certificate and the second certificate are issued by different intermediate certificate authorities. In some cases, the different intermediate certificate authorities are verified by the same root certificate authority.

[0012] In some embodiments, the determined policy action is executed on a base system logic. In some cases, the base system logic comprises a bare metal server, a hypervisor or a docker base for controlling one or more virtual machines or containers and wherein the bare metal server, the hypervisor or the docker base is coupled to the network interface device. In some cases, the policy action is selected from the group comprising: allow, block, rate-limit, mark or alert.

[0013] In some embodiments, the first security attribute comprises a security group the source IP address is associated therewith. In some embodiments, the second security attribute comprises a security group the destination IP address is associated therewith.

[0014] It shall be understood that different aspects of the invention can be appreciated individually, collectively, or in combination with each other. Various aspects of the invention described herein may be applied to any of the particular applications set forth below or for any other types of the network traffic management/security system disclosed herein. Any description herein concerning the network microsegmentation may apply to and be used for any other network microsegmentation situations. Additionally, any embodiments disclosed in the context of the network microsegmentation system are also applicable to the methods disclosed herein.

INCORPORATION BY REFERENCE

[0015] All publications, patents, and patent applications mentioned in this specification are herein incorporated by reference to the same extent as if each individual publication, patent, or patent application was specifically and individually indicated to be incorporated by reference.

BRIEF DESCRIPTION OF THE DRAWINGS

[0016] The novel features of the invention are set forth with particularity in the appended claims. A better understanding of the features and advantages of the present invention will be obtained by reference to the following detailed description that sets forth illustrative embodiments, in which the principles of the invention are utilized, and the accompanying drawings of which:

[0017] **FIG. 1A** is an exemplary block diagram illustrating connections between microsegments over a network, in accordance with embodiments of the invention;

[0018] **FIG. 1B** depicts an exemplary network environment;

[0019] **FIG. 2** depicts an exemplary network environment;

[0020] **FIG. 3** shows an exemplary cryptography based micro-segmentation system for managing network flow;

[0021] **FIG. 4** shows a block diagram for setting up a certificate authority;

[0022] **FIG. 5** depicts an exemplary network environment according to some embodiments;

[0023] **FIG. 6** shows an example of using multiple intermediate CAs in a system for establishing secure connections between endpoints;

[0024] **FIG. 7** shows an exemplary process for establishing a secure and/or authenticated connection, in accordance with embodiments of the invention; and

[0025] **FIG. 8** illustrates an exemplary process performed by the system for mutual authentication between endpoints of a connection.

DETAILED DESCRIPTION OF THE INVENTION

[0026] In the following detailed description, reference is made to the accompanying figures, which form a part hereof. In the figures, similar symbols typically identify similar components, unless context dictates otherwise. The illustrative embodiments described in the detailed description, figures, and claims are not meant to be limiting. Other embodiments may be utilized, and other changes may be made, without departing from the scope of the subject matter presented herein. It will be readily understood that the aspects of the present disclosure, as

generally described herein, and illustrated in the figures, can be arranged, substituted, combined, separated, and designed in a wide variety of different configurations, all of which are explicitly contemplated herein.

Certain definitions

[0027] Unless otherwise defined, all technical terms used herein have the same meaning as commonly understood by one of ordinary skill in the art to which this invention belongs.

[0028] Reference throughout this specification to "some embodiments," or "an embodiment," means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment. Thus, the appearances of the phrase "in some embodiment," or "in an embodiment," in various places throughout this specification are not necessarily all referring to the same embodiment. Furthermore, the particular features, structures, or characteristics may be combined in any suitable manner in one or more embodiments.

[0029] As utilized herein, terms "component," "system," "interface," "unit" and the like are intended to refer to a computer-related entity, hardware, software (e.g., in execution), and/or firmware. For example, a component can be a processor, a process running on a processor, an object, an executable, a program, a storage device, and/or a computer. By way of illustration, an application running on a server and the server can be a component. One or more components can reside within a process, and a component can be localized on one computer and/or distributed between two or more computers.

[0030] Further, these components can execute from various computer readable media having various data structures stored thereon. The components can communicate via local and/or remote processes such as in accordance with a signal having one or more data packets (e.g., data from one component interacting with another component in a local system, distributed system, and/or across a network, e.g., the Internet, a local area network, a wide area network, etc. with other systems via the signal).

[0031] As another example, a component can be an apparatus with specific functionality provided by mechanical parts operated by electric or electronic circuitry; the electric or electronic circuitry can be operated by a software application or a firmware application executed by one or more processors; the one or more processors can be internal or external to the apparatus and can execute at least a part of the software or firmware application. As yet another example, a component can be an apparatus that provides specific functionality through electronic components without mechanical parts; the electronic components can include one or

more processors therein to execute software and/or firmware that confer(s), at least in part, the functionality of the electronic components. In some cases, a component can emulate an electronic component via a virtual machine, e.g., within a cloud computing system.

[0032] Moreover, the word "exemplary" where used herein to means serving as an example, instance, or illustration. Any aspect or design described herein as "exemplary" is not necessarily to be construed as preferred or advantageous over other aspects or designs. Rather, use of the word exemplary is intended to present concepts in a concrete fashion. As used in this application, the term "or" is intended to mean an inclusive "or" rather than an exclusive "or." That is, unless specified otherwise, or clear from context, "X employs A or B" is intended to mean any of the natural inclusive permutations. That is, if X employs A; X employs B; or X employs both A and B, then "X employs A or B" is satisfied under any of the foregoing instances. In addition, the articles "a" and "an" as used in this application and the appended claims should generally be construed to mean "one or more" unless specified otherwise or clear from context to be directed to a singular form.

[0033] Embodiments of the invention may be used in a variety of applications. Some embodiments of the invention may be used in conjunction with various devices and systems, for example, a personal computer (PC), a desktop computer, a mobile computer, a laptop computer, a notebook computer, a tablet computer, a server computer, a handheld computer, a handheld device, a personal digital assistant (PDA) device, a handheld PDA device, a wireless communication station, a wireless communication device, a wireless access point (AP), a modem, a network, a wireless network, a local area network (LAN), a virtual local area network (VLAN), a wireless LAN (WLAN), a metropolitan area network (MAN), a wireless MAN (WMAN), a wide area network (WAN), a wireless WAN (WWAN), a personal area network (PAN), a wireless PAN (WPAN), a virtual private network (VPN), a storage area network (SAN), a frame relay connection, an Advanced Intelligent Network (AIN) connection, a synchronous optical network (SONET) connection, devices and/or networks operating in accordance with existing IEEE 802.11, 802.11a, 802.11b, 802.11e, 802.11g, 802.11h, 802.11i, 802.11n, 802.16, 802.16d, 802.16e standards and/or future versions and/or derivatives and/or long term evolution (LTE) of the above standards, units and/or devices which are part of the above networks, one way and/or two-way radio communication systems, cellular radio-telephone communication systems, a cellular telephone, a wireless telephone, a personal communication systems (PCS) device, a PDA device which incorporates a wireless communication device, a multiple input multiple output (MIMO) transceiver or device, a single input multiple output (SIMO) transceiver or device, a multiple input single output (MISO)

transceiver or device, or the like.

[0034] It is noted that various embodiments can be used in conjunction with one or more types of wireless or wired communication signals and/or systems, for example, radio frequency (RF), infrared (IR), frequency-division multiplexing (FDM), orthogonal FDM (OFDM), time-division multiplexing (TDM), time-division multiple access (TDMA), extended TDMA (E-TDMA), general packet radio service (GPRS), extended GPRS, code-division multiple access (CDMA), wideband CDMA (WCDMA), CDMA 2000, multi-carrier modulation (MDM), discrete multi-tone (DMT), Bluetooth®, ZigBee™, or the like. Embodiments of the invention may be used in various other devices, systems, and/or networks.

[0035] While portions of this disclosure, for demonstrative purposes, refer to wired and/or wired communication systems or methods, embodiments of the invention are not limited in this regard. As an example, one or more wired communication systems, can utilize one or more wireless communication components, one or more wireless communication methods or protocols, or the like.

[0036] Although some portions of the discussion herein may relate, for demonstrative purposes, to a fast or high-speed interconnect infrastructure, to a fast or high-speed interconnect component or adapter with OS bypass capabilities, to a fast or high-speed interconnect card or Network Interface Card (NIC) with OS bypass capabilities, or to a to a fast or high-speed interconnect infrastructure or fabric, embodiments of the invention are not limited in this regard, and may be used in conjunction with other infrastructures, fabrics, components, adapters, host channel adapters, cards or NICs, which may or may not necessarily be fast or high-speed or with OS bypass capabilities. For example, some embodiments of the invention may be utilized in conjunction with InfiniBand (IB) infrastructures, fabrics, components, adapters, host channel adapters, cards or NICs; with Ethernet infrastructures, fabrics, components, adapters, host channel adapters, cards or NICs; with gigabit Ethernet (GEth) infrastructures, fabrics, components, adapters, host channel adapters, cards or NICs; with infrastructures, fabrics, components, adapters, host channel adapters, cards or NICs that have OS with infrastructures, fabrics, components, adapters, host channel adapters, cards or NICs that allow a user mode application to directly access such hardware and bypassing a call to the operating system (namely, with OS bypass capabilities); with infrastructures, fabrics, components, adapters, host channel adapters, cards or NICs; with infrastructures, fabrics, components, adapters, host channel adapters, cards or NICs that are connectionless and/or stateless; and/or other suitable hardware.

[0037] In one aspect, methods and systems of the present disclosure are provided for

establishing a secure and/or authenticated network connection. The secured network connection may be provided by means of cryptography based microsegmentation. In some embodiments, a certificate is used for establishing a secure network connection. The certificate may be generated and issued to each workload. In some embodiments, the certificate may comprise information related to policy and/or security attributes of ends of connection for establishing the connection. The described system and method may utilize a cryptographic identity of ends of a connection that is consistent across clusters, thereby providing improved scalability and flexibility of network microsegmentation.

Cryptography

[0038] In some embodiments, the system and method may utilize cryptographic technologies such as Transport Layer Security (TLS) and its predecessor, Secure Sockets Layer (SSL). TLS and SSL are cryptographic protocols or encryption protocols used to provide secure connections over the Internet. SSL utilizes X.509 certificates, certificate authorities, and a public key infrastructure to verify relation between a certificate and its owner, as well as to generate, sign, and administer the validity of certificates. According to SSL protocols, session information between endpoints of a connection such as an SSL client and an SSL server are negotiated through a handshake phase and the identity of the SSL server is verified by the SSL client. In the server-client example, the session information may include a session ID, peer certificates, the cipher specification to be used, the compression algorithm to be used, and shared secrets that are used to generate symmetric cryptographic keys. The SSL client may encrypt a premaster secret with a public key from the SSL server's certificate and transmits the premaster secret to the server. Then, both parties compute the master secret locally and derive the session key from it. After the handshake phase, a secure socket may be established, and application data encrypted by the session key can be securely transmitted between the client and server.

[0039] Certificates may be used to create secure connections between endpoints of a network. For example, trusted certificates can be used to create secure connections to a server via the Internet. A certificate is essential in order to circumvent a malicious party which happens to be on the route to a target server which acts as if it were the target. Such a scenario is commonly referred to as a man-in-the-middle attack. The client uses the certificate authority (CA) certificate to authenticate the CA signature on the server certificate, as part of the authorizations before launching a secure connection. A certificate authority or certification authority (CA) is an entity that issues digital certificates. A digital certificate may certify the ownership of a public key by the named subject of the certificate. This allows others (relying parties) to rely upon signatures or on assertions made about the private key that corresponds to the certified

public key. A CA acts as a trusted third party—trusted both by the subject (owner) of the certificate and by the party relying upon the certificate.

Microsegmentation

[0040] In some cases, a microsegment may correspond to a virtual network. In some cases, a microsegment may correspond to an endpoint group. In some instances, the endpoint group may refer to a security group and the definition of a security group can be defined by a network administer such that the definition can be changed or modified. In some cases, a microsegment may correspond to a cluster (e.g., data center, subnet, node device) that may be a physical system or a virtual network segment. The microsegment may be dynamic or ever changing. Because the certificate is tied to a workload rather than the IP address, the provided method and system may be well adapted to the dynamic environment. In some embodiments, endpoint can be referred to as workload. In this case, a certificate may be issued to an endpoint for establishing a secure connection with other endpoints.

[0041] In some cases, endpoint groups can be used in network environment for mapping applications to the network. In particular, endpoint groups can use a grouping of application endpoints in the network to apply connectivity and policy to the group of applications. Endpoint groups can act as a container for buckets or collections of applications, or application components, and tiers for implementing forwarding and policy logic. Endpoint groups may also allow separation of network policy, security, and forwarding from addressing by instead using logical application boundaries. For example, each endpoint group may connect to network fabric via leaf switches. In some embodiments, endpoints belong to the same endpoint group may be considered as sharing security attributes. In some embodiments, an endpoint group that an endpoint belongs to may be a security attribute of the endpoint. Details regarding endpoint are described later herein.

[0042] **FIG. 1A** is an exemplary block diagram illustrating connections between microsegments over a network 101, in accordance with embodiments of the invention. As mentioned above, a microsegment may correspond to an endpoint group such as security group SG 103, SG 105, SG 107, SG 109. A security group may comprise one or more endpoints. An endpoint may belong to a security group. In the illustrated example, each endpoint belongs to a respective security group. However, any number of endpoints that may or may not be associated with the same node can be assigned to the same security group. Endpoints can be grouped regardless of the physical system or infrastructure that the endpoints are attached to.

[0043] Network 101 may be a telecommunications network that allows computers to exchange

data. For example, in network 101, networked computing devices pass data to each other along data connections (e.g., network links). Data can be transferred in the form of packets. The connections between nodes may be established using either cable media or wireless media. The network 101 may, for example, include a wireless network, a local area network (LAN), a virtual local area network (VLAN), a wireless LAN (WLAN), a metropolitan area network (MAN), a wireless MAN (WMAN), a wide area network (WAN), a wireless WAN (WWAN), a personal area network (PAN), a wireless PAN (WPAN), a virtual private network (VPN), a storage area network (SAN), a frame relay connection, an Advanced Intelligent Network (AIN) connection, a synchronous optical network (SONET) connection, devices and/or networks operating in accordance with existing IEEE 802.11, 802.11a, 802.11b, 802.11e, 802.11g, 802.11h, 802.11i, 802.11n, 802.16, 802.16d, 802.16e standards and/or future versions and/or derivatives and/or long term evolution (LTE) of the above standards. The network can be used in conjunction with one or more types of wireless or wired communication signals and/or systems, for example, radio frequency (RF), infrared (IR), frequency-division multiplexing (FDM), orthogonal FDM (OFDM), time-division multiplexing (TDM), time-division multiple access (TDMA), extended TDMA (E-TDMA), general packet radio service (GPRS), extended GPRS, code-division multiple access (CDMA), wideband CDMA (WCDMA), CDMA 2000, multi-carrier modulation (MDM), discrete multi-tone (DMT), Bluetooth®, ZigBee™, or the like.

[0044] In an example, the network 101 may be Internet that connects disparate networks throughout the world, providing global communication between nodes on various networks. The nodes typically communicate over the network by exchanging discrete frames or packets of data according to predefined protocols, such as the Transmission Control Protocol/Internet Protocol (TCP/IP). In this context, a protocol can refer to a set of rules defining how the nodes interact with each other. Computer networks may be further interconnected by an intermediate network node, such as a router, to extend the effective size of each network.

[0045] An endpoint can be any network entity, component or communication device. For instance, an endpoint can be a physical server, process (e.g., function running on a virtual machine), external network, a virtual partition, legacy mainframes, modem, hub, bridge, switch, router, server, workstation, desktop computer, laptop computer, tablet, mobile phone, desk phone, wearable device, or other network or electronic device. A virtual partition may be an instance of a virtual machine (VM), sandbox, container, or any other isolated environment that can have software operating within it. In some example embodiments, endpoints can include a server, hypervisor, process, or switch configured with virtual tunnel endpoint (VTEP) functionality which connects an overlay network with network fabric. The

overlay network may allow virtual networks to be created and layered over a physical network infrastructure. Overlay network protocols, such as Virtual Extensible LAN (VXLAN), Network Virtualization using Generic Routing Encapsulation (NVGRE), Network Virtualization Overlays (NVO3), and Stateless Transport Tunneling (STT), can provide a traffic encapsulation scheme which allows network traffic to be carried across L2 and L3 networks over a logical tunnel. Such logical tunnels can be originated and terminated through VTEPs. The overlay network can host physical devices, such as servers, applications, endpoint groups, virtual segments, virtual workloads, etc. In addition, endpoints can host virtual workload(s), clusters, and applications or services, which can connect with network fabric or any other device or network, including an internal or external network. For example, endpoints can host, or connect to, a cluster of load balancers or an endpoint group of various applications.

[0046] Endpoints belong to the same endpoint group or security group may or may not be of the same type. In the illustrated example of **FIG. 1A**, the endpoints are shown as virtual machines (VMs). However, endpoints of different types such as VM and container can be assigned to the same endpoint group. In some cases, endpoints may be grouped according to a virtual network segment. For instance, VMs connected to the same physical hosts may be grouped together as a security group or endpoint group. In other examples, VMs connected to different physical hosts may be grouped together. Endpoints may be grouped based on IP address or subnet they belong to. Alternatively, endpoints may be grouped regardless of their IP addresses. In some embodiments, endpoints may be grouped by user defined rules. For example, user may define a security group or endpoint group based on one or more attributes of endpoints that may or may not relate to IP address. For instance, the attributes can be a VM-name, operating system (OS)-name, host-name or a fully qualified domain name (FQDN).

[0047] **FIG. 1A** depicts an exemplary virtual environment. The one or more endpoints may be virtual machines. The described system or method may be used to provide secure network connection between endpoints using the cryptographic identity of the endpoint. In some embodiments, the cryptographic identity of an endpoint may comprise the security attribute of the endpoint such as the security group the endpoint belongs to. It should be noted that a security group can be defined regardless of the physical system or IP addresses. For instance, a security group may correspond to multiple nodes, a subset of endpoints connected to a node, multiple endpoints connected to multiple nodes, a single node or a single endpoint.

[0048] The depicted virtual environment may comprise one or more physical hosts hosting one or more endpoints. A physical host may be a computing device that acts as computing server such as a blade server or bare-metal server. In some cases, a physical host may be part of a cloud

computing environment. By way of further non-limiting examples, a physical host can host different combinations and permutations of virtual and container environments.

[0049] In the illustrated example, a virtual environment of the node may comprise hardware, host operating system (not shown), hypervisor 131, 133, and virtual machines 111, 113, 115, 117. The host operating system can run on hardware and can also be referred to as the host kernel. In some embodiments, hypervisor 131, 133 may comprise a virtual switch 121. The virtual machine (i.e., endpoint 111, 113, 115, 117) may each include a respective one of operating systems (OS) and applications (APP). In the illustrated example, a node may comprise a network interface device 151, 153 configured to enable communication with another system or device over the network 101.

[0050] In some cases, an agent 141, 143 may run on the network interface device. The agent may be configured to operate in conjunction with a certificate authority 100 to facilitate a secure network connection. The agent may be configured for managing security of an endpoint. For example, the agent at a sender's side or source side may send a request to the certificate authority (CA) for signing a certificate of a secure connection associated with the source and a security attribute. In some cases, the agent may verify if the certificate is a valid certificate based on the digital signature included in the certificate. The agent may then store and manage the signed certificate associated with the endpoint in a local database. For example, a certificate issued by a certificate authority may have a validity period based on the certificate management protocol. If a new certificate is generated or the certificate is renewed and pushed to the agent, the agent may delete the old certificate from the client machine and install the new certificate to the trusted certificate store of the client machine.

[0051] In some embodiments, the agent may initiate a connection to a destination endpoint using the stored certificate. In some cases, the agent at the receiver's side or destination side may check a policy to allow or deny the connection based on the security attribute of the source and a security attribute of the destination. In some cases, the agent may also be configured for managing the security of traffic pass through the network. For instance, the agent may execute an action according to the policy determined based on the source and destination security attribute to, for example, permit or allow the flow described in the policy (i.e., forward the communication), block or deny the flow described in the policy (i.e., drop the communication), limit the bandwidth consumed by the flow (i.e., rate-limit), log the flow, "mark" the flow for quality of service (QoS) (e.g., set a lower or higher priority for the flow), redirect the flow (e.g., to avoid critical paths), copy the flow, and various other actions. Alternatively, such operations may be performed by other modules of the endpoint or the node.

[0052] The agent may be local to the node of the network or the endpoint. In the illustrated example, the agent is running on the network interface device 151 such as NIC. Alternatively or additionally, the agent may reside on any component of the node such as virtual partition, hypervisor, physical server, switch, router, gateway, or other independent systems or devices operably or communicatively coupled to the node. As mentioned elsewhere herein, the virtual partition may be an instance of a virtual machine (VM), sandbox, container, virtual switch, or any other isolated environment that can have software operating within it. The agent may be implemented by software, hardware or a combination of both.

[0053] Hypervisor (also known as a virtual machine monitor (VMM)) 131, 133, is software running on at least one of physical hosts, and the hypervisor 131, 133, runs VMs. A physical host on which hypervisor is running one or more virtual machines 111, 113 is also referred to as a host machine. Each VM 111, 113 can also be referred to as a gust machine. For example, hypervisor 131 may allow multiple operating systems to share a single physical host.

[0054] Hypervisor may also include a virtual switch 121, 123. A virtual switch is a logical switching fabric for networking VMs. For example, virtual switch 121 may be a program running on a physical host that allows VM 111, 113 to communicate with another VM.

[0055] The certificate authority 100 may be implemented by any suitable systems such as a server. The certificate authority 100 as described elsewhere herein may issue certificate in response to a request submitted by the agent. In some cases, the certificate authority 100 may be verified by another certificate authority (e.g., root certificate authority).

[0056] **FIG. 1B** depicts another exemplary network environment. The network environment may comprise hardware, host operating system, container engine 221, 223, containers (e.g., endpoint 211, 213, 215, 217), virtual switch 121, 123 and agent 141, 143.

[0057] Host operating system may allow for multiple (instead of just one) isolated user-space instances (e.g., containers 211, 213, 215, 217) to run in host operating system (e.g., a single operating system instance). In some cases, a host operating system may include a container engine.

[0058] Container engine 221, 223 may be configured to create and manage containers 211, 213, 215, 217, for example, using an (high-level) application programming interface (API). By way of non-limiting example, container engine 221, 223 may be at least one of Docker®, Rocket (rkt), Solaris Containers, and the like. For example, container engine 221may create a container (e.g., one of containers 211, 213) using an image. An image can be a (read-only) template comprising multiple layers and can be built from a base image (e.g., for host operating system)

using instructions (e.g., run a command, add a file or directory, create an environment variable, indicate what process (e.g., application or service) to run, etc.). Each image may be identified or referred to by an image type. In some embodiments, images (e.g., different image types) are stored and delivered by a system (e.g., server side application) referred to as a registry or hub.

[0059] Container engine 221, 223 may allocate a filesystem of host operating system to the container and add a read-write layer to the image. Container engine can create a network interface that allows the container to communicate with hardware (e.g., talk to a local host). Container engine can set up an Internet Protocol (IP) address for the container (e.g., find and attach an available IP address from a pool). Container engine can launch a process (e.g., application or service) specified by the image (e.g., run an application, such as one of APPs 231, 233, described further below). Container engine can capture and provide application output for the container (e.g., connect and log standard input, outputs and errors). The above examples are only for illustrative purposes and are not intended to be limiting.

[0060] Containers 211, 213, 215, 217 can be created by container engine 221, 223. In some embodiments, containers are each an environment as close as possible to an installation of host operating system, but without the need for a separate kernel. For example, containers 211, 213, may share the same operating system kernel with each other and with host operating system 320. Each container of containers 211, 213 can run as an isolated process in user space on host operating system. Shared parts of host operating system can be read only, while each container of containers 211, 213 can have its own mount for writing.

[0061] Containers 211, 213, 215, 217 can include one or more applications (APPs) 231, 233 (and all of their respective dependencies). APPs 231, 233 can be any application or service. By way of non-limiting example, APPs can be a database (e.g., Microsoft® SQL Server®, MongoDB, HTFS, etc.), email server (e.g., Sendmail®, Postfix, qmail, Microsoft® Exchange Server, etc.), message queue (e.g., Apache® Qpid™, RabbitMQ®, etc.), web server (e.g., Apache® HTTP Server™, Microsoft® Internet Information Services (IIS), Nginx, etc.), Session Initiation Protocol (SIP) server (e.g., Kamailio® SIP Server, Avaya® Aura® Application Server 5300, etc.), other media server (e.g., video and/or audio streaming, live broadcast, etc.), file server (e.g., Linux server, Microsoft® Windows Server®, etc.), service-oriented architecture (SOA) and/or microservices process, object-based storage (e.g., Lustre®, EMC® Centera®, Scality® RING®, etc.), directory service (e.g., Microsoft® Active Directory®, Domain Name System (DNS) hosting service, etc.), and the like. The APPs can also be any endpoint protection software that may perform security action determined based on cryptographic identity of the endpoints.

[0062] Virtual switch 121,123 may be a logical switching fabric for networking containers. For example, virtual switch 121 may allow a container (of containers 211, 213) to communicate with another container. For example, containers 211, 213 can communicate with other devices such as VMs and bare-metal. In some embodiments, virtual switch 121, 123 may execute as a part of host operating system.

[0063] **FIG. 2** depicts another exemplary network environment. The network environment may comprise computing servers (e.g., endpoint 221, 223, 225, 227), network interface devices 151, 153, 155 and agent 141, 143.

[0064] The computing servers may be physical hosts communicate with other devices or systems via the network interface device 151, 153, 155. The computing servers may be, for example, blade servers or bare-metal servers. In some cases, the physical hosts may host different combinations and permutations of virtual and container environments as described above.

[0065] **FIG. 3** shows an exemplary cryptography based micro-segmentation system for managing network flow. A network flow is conventionally characterized as one or more packets sharing certain attributes that are sent within a network within a specified period of time. Packet attributes can include a network source address (e.g., Internet Protocol (IP) address, Media Access Control (MAC) address, Domain Name System (DNS) name, or other network address), source port, destination address, destination port, protocol type, class of service, among other characteristics. The network source address may correspond to a first endpoint (e.g., modem, hub, bridge, switch, router, server, workstation, desktop computer, laptop computer, tablet, mobile phone, desk phone, wearable device, or other network or electronic device) of the network, and the network destination address may correspond to a second endpoint of the network. In some cases, a switch or router interface can also be a packet attribute used to distinguish network flows.

[0066] In some embodiments, the system 300 may comprise a management controller 320 and one or more agents 313, 333. The management controller and the agent may be implemented by hardware, software or a combination of both. The management controller may work in conjunction with the one or more agents to provide secure and/or authenticated connection between endpoints of a network. The system may comprise any number of management controllers or agents.

[0067] In some embodiments, the management controller 320 may operate as a certificate authority. In some cases, the certificate authority can be the same as the certificate authority as

described in **FIG. 1A** or elsewhere herein. In some cases the certificate authority may be an intermediate CA. The intermediate CA may be verified by a trusted root certificate authority. In this case, the system may comprise multiple intermediate CAs that may or may not be verified by the same trusted root certificate authority.

[0068] In some embodiments, the management controller may comprise or coupled to a policy database 323. The policy database 323 may server as a repository for network policy or security policy. The policy database can utilize various database structures such as a normalized relational database or NoSQL database.

[0069] Network policies can determine whether a particular flow is allowed or denied by the network as well as a specific route by which a packet traverses the network. Policies can also be used to mark packets so that certain kinds of traffic receive differentiated service when used in combination with queuing techniques such as those based on priority, fairness, weighted fairness, token bucket, random early detection, round robin, among others. A rule, for example, allows or denies a connection to a specific group, allows connection between groups or denies a connection between groups, redirects a connection from one IP address to another IP address, logs communications to and/or from a specific IP address, allows or denies a connection for rate-limit purpose, marks a suspicious flow, generates alert and various others. As described elsewhere herein, an IP address may be virtual, physical, or both.

[0070] Network policy can be stored in any suitable data structure. In the depicted example, the network policy may be stored in a table 325. For instance, a connection may be allowed or denied based on a network policy or rule set stored in the table 325. A policy may specify security attributes 327, 328 such as source security group, destination security group, and action 329 to be applied to a packet when the packet matches security attributes. In an example, once a source security group (e.g., SG1) and a destination security group (e.g., SG2) are identified, the action (e.g., Allow) can be determined.

[0071] Action 329 can be the action that is applied to a communication when the communication matches a corresponding policy. For example, action 329 can be to permit or allow a flow (i.e., forward the communication), block or deny a flow (i.e., drop the communication), limit the bandwidth consumed by the flow, log the flow, "mark" the flow for quality of service (QoS) (e.g., set a lower or higher priority for the flow), redirect the flow (e.g., to avoid critical paths), copy the flow, etc. In some example embodiments, action 329 can have an expiration time or date. For example, it can only take the designated action (e.g., allow, block, mark, etc.) for a certain amount of time before the communication is dropped. Similarly, action can have designated times of applicability, for example only during peak hours. A policy

can be over-inclusive or under-inclusive. For example, in certain situations, a whitelist policy may allow communications that are potentially harmful to the network while a blacklist policy can block communications that are permitted by the network.

[0072] Policies can be established external to a network system. In some cases, a network administrator can manually change the policies. Policies can dynamically change and be conditional on events. In some cases, the policy table may be pushed to some or all of the nodes of a network that is managed by the management controller.

[0073] The management controller 320 may be in communication with one or more nodes 310, 330 over the network. The node 310, 330 can be the same as the node as described in connection with **FIG. 1A**. In some embodiments, each node may comprise or coupled to an agent 313, 333. The agent may be in communication with the management controller 320. The agent can be the same as the agent as described in **FIG. 1A** or elsewhere herein.

[0074] Each of the nodes may be associated with one or more endpoints. In some embodiments, an endpoint may be associated with a security attribute. Endpoints share the same security attribute may be defined as a microsegment. In some cases, endpoints share the same security attribute may be defined as an endpoint group or security group (SG).

[0075] Endpoints associated with the same node may or may not share the same security attribute. For example, endpoint 311 may have security attribute as SG1 and endpoint 312 may have security attribute as SG2. Endpoints associated with different nodes may or may not share the same security attribute. For example, endpoint 312 associated with node 310 may share the attribute (e.g., SG2) of the endpoint 331that is associated with a different node 330.

[0076] In some embodiments, an endpoint may be mapped to security attribute by one or more other attributes of the endpoint. In some cases, the one or more other attributes may be attributes of a packet transmitted between endpoints. A packet attribute can be a description of a certain characteristic that can be matched with a communication (e.g., a subnet or port range). For instance, source and destination packet attributes may include a MAC address, IP address, endpoint ID, user, process (e.g., name, PID), subnet, geographical location and the like, or any combination of the above.

[0077] In some embodiments, the mapping relationship may be stored in a database such as the endpoint database 315, 335. For example, endpoint 311 may be a virtual machine (e.g., VM1) with IP address of IP1 and the security attribute of endpoint 311 is identified as SG1 according to the mapping relationship stored in the endpoint DB 315. In some cases, the endpoint database may be a local database to the node. The endpoint database local to each node

may not be synchronized across the nodes. For instance, endpoint databases associated with different nodes may contain local data that may not be the same across the nodes. For example, endpoint database 315 may comprise mapping information related to endpoints 311, 312 (e.g., VM1, VM2) and endpoint database 335 may comprise mapping information related to endpoints 331, 332 (e.g., VM3, VM4).

[0078] In some embodiments, the endpoint database 315, 335 may further store certificates issued to each endpoint. The certificate may be issued by a certificate authority operating on the management controller 320. For example, the certificate may be TLS certificate (e.g., Cerf1) associated with each workload or endpoint (e.g., VM1). Details regarding the certificates are described later herein.

[0079] As mentioned above, a certificate authority may be verified by another certificate authority such as root certificate authority prior to setting up a connection. **FIG. 4** shows a block diagram for setting up a certificate authority 320. As shown in the example, a certificate authority 410 may be provided by a management controller 320. The certificate authority 410 may be configured for issuing certificate for endpoints. In an example, Secure Sockets Layer (SSL) which is a well-known cryptographic protocol may be used to secure communications over networks such as the Internet. Cryptographic protocols such as SSL are often based on public key cryptographic systems, such as the RSA (Rivest, Shamir and Adelman) encryption algorithm. For a traditional RSA-based SSL session, the two sides of a connection agree upon a "pre-master secret" (PMS) which is used to generate the parameters for the remainder of the session. Typically, the two sides use RSA asymmetric encryption to establish the pre-master secret without exchanging the actual value in plaintext. In operation, the SSL client generates the pre-master secret and encrypts it with the SSL server's publicly available RSA key. This generates an encrypted pre-master secret (ePMS), which is then provided to the SSLserver. The SSL server has a private decryption key, which is then used to decrypt the encrypted pre-master secret. At this point, both the client and the server have the original pre-master secret and can use it to generate the symmetric key used for actual encrypted and secure data exchange.

[0080] In some cases, a management controller or certificate authority may have a CA certificate and keypair used for communication with nodes or endpoints. When a new management controller 320 or certificate authority 410 is provisioned, the certificate authority 410 may generate a CA certificate and CSR (Certificate Signing Request) 401 and send the latter to a third party CA 420. The third party CA 420 may be a trusted CA such as a root CA, a CA a level higher in the certificate chain than the CA 410, or any other CA that may be further certified by another CA. in some cases, the CSR 401 preferably may contain a pair of identifiers,

such as a unique serial number for the management controller or the CA 410, as well as a serial number for the CSR version. The third party CA 420 may sign the CSR and send the certificate 400 to the CA 410. In some cases, the third party CA 420 may send a message containing the certificate details (serial number pair) to other intermediate CA that may communicate with the CA 410.

[0081] The CA certificate 400 signed by the third party certificate authority may comprise information related to the CA that signed the certificate (e.g., signed by foo.com) and the certified entity (e.g., intermediate CA 410 with name Venice.dc1). In the case when SSL is used, the certificate may comply with the X.509 standard. For example, the name of the intermediate CA may be a Distinguished Name (DN) uniquely identifies a CA in the X.509 certificate.

[0082] In some embodiments, an end of a connection such as a node may comprise a key repository such as a key store 430 for storing private keys of the certificate authority 410. The private key may be obtained from the CA 410 corresponding to the CA certificate. In some cases, the key repository may also contain one or more certificates (e.g., TLS/SSL certificate) issued for each workload or endpoint, one or more private keys associated with the TLS/SSL certificate, or certificates request that to be signed by the intermediate CA. In some cases when mutual authentication is required, the key repository may further store certificates issued to another endpoint or workload. In some instances, the key repository may also comprise the endpoint database such as endpoint database 315, 335 as described in **FIG. 3**. Alternatively, the endpoint database may be stored in a separate location from the key repository. The key repository may or may not store the policy table that has been described elsewhere herein.

[0083] The key repository may reside on a memory unit. The memory unit may be local to the node 310. The memory unit may be coupled to the agent 313 or accessed by the agent 313. In some cases, the memory unit may be a component of a network interface infrastructure such as NIC of the node 310. In some embodiments, the NIC may comprise a write port or communication interface allowing for memory read/write operations. For instance, the communication interface may support packets written to or read from the memory unit such as an external memory (e.g., high bandwidth memory (HBM) of a host device) or an internal static random access memory (SRAM). The communication interface may employ any suitable protocol such as Advanced Microcontroller Bus Architecture (AMBA) Advanced extensible Interface (AXI) protocol. AXI is a bus protocol for a high-speed/high-end on-chip bus protocol and has channels associated with read, write, address, and write response, which are respectively separated, individually operated, and have transaction properties such as multiple-outstanding address or write data interleaving. The AXI interface may include features that support for

unaligned data transfers using byte strobes, burst based transactions with only start address issued, separate address/control and data phases, issuing of multiple outstanding addresses with out of order responses, and easy addition of register stages to provide timing closure. Though packet data is transferred according the AXI protocol in the packet data communication on-chip interconnect system according to the present exemplary embodiment in the present specification, it can also be applied to a packet data communication on-chip interconnect system operating by other protocols supporting a lock operation, such as Advanced High-performance Bus (AHB) protocol or Advanced Peripheral Bus (APB) protocol in addition to the AXI protocol.

[0084] **FIG. 5** depicts an exemplary network environment according to some embodiments. A secure or authenticated connection may be provided by a system. The system may comprise at least a management controller operating a certificate authority and an agent associated with an end of a connection. In some cases, endpoints may be the source and destination of a connection 510. The connection 510 between a source and destination endpoints may be secured and authenticated utilizing certificates associated with each workload or the endpoints. Source node and destination node can be endpoints. For example, source node 310 may communicate (i.e., initiate a flow) with destination node 330 by sending data (e.g., packets) through a network (not shown). Alternatively, the connection 510 may be between an endpoint (e.g., EP 311, EP 312) supported by a source node and an endpoint (e.g., EP 331, EP 332) supported by the destination node 330. In some embodiments, each node of the network may comprise a key repository. The key repository associated with each node may comprise local data as described above. For instance, the key store 430 associated with node 310 may store certificates issued to endpoints 311, 312 and security attributes thereof, whereas the key store 510 associated with node 330 may store certificates issued to endpoints 331, 332 and security attributes thereof.

[0085] The certificate 500 associated with an endpoint may be issued by the certificate authority 301. The certificate 500 may be an endpoint certificate or workload certificate that identifies the endpoint and/or the CA that signs the certificate. In the illustrated example, the certificate issued to endpoint 311 may be a TLS/SSL certificate that comprises security attribute of the endpoint 311 and the CA that signs the certificate (e.g., CA with name Venice.dc1).

[0086] In some embodiments, the security attribute of an endpoint may be included in the certificate 500. The security attribute may be embedded in any field of the certificate. For example, as part of the X.509 certificate standard, the security attribute may be included in the Subject Alternative Name (SAN) field of the certificate. As shown in the example, the security attribute may be the security group (e.g., security group with name or identifier WebGroup.dc1) that endpoint 311 belongs to and is included in the certificate as SAN extension.

[0087] An endpoint may be created by obtaining a certificate issued to the endpoint. The endpoint certificate may be signed or issued by the certificate authority in response to receiving a CSR from the agent. For example, agent 313 may generate and send a CSR to the certificate authority 301 for creating endpoint 311. The CSR may comprise security attribute (e.g., security group) and other packet attributes such as IP address uniquely associated with the endpoint. The certificate authority 301 may sign the certificate and send it to the agent 313. The agent 313 may then manage and store the certificate associated with the endpoint 311 in a local database such as the endpoint database 315. The certificate associated with a workload or network flow may be identified by the packet attribute (e.g., endpoint ID, IP) as described elsewhere herein. For instance, a packet with source IP address (e.g., IP1) may be identified to be issued certificate Cerf1 having security attribute SG1 in the SAN field of the certificate.

[0088] For establishing a connection, the agent at the source endpoint side may be configured to inspect the packet attributes in a packet in order to identify the certificate associated with the endpoint. For example, agent 313 may see the IP address of the source endpoint 311 (e.g., IP1) then retrieve the Cerf1 from the endpoint database 315. Next the agent may initiate a connection with the certificate (e.g., Cerf1). For example, the agent 313 may continue with a SSL handshake with the destination endpoint (e.g., EP 331) using Cerf1 that has the security attribute (e.g., SG1) of the source endpoint in the SAN field.

[0089] In some cases, the agent at the destination endpoint may check the certificate and determine the security attributes of the source endpoint and security attributes of the destination endpoint. In some cases, the agent at the destination endpoint may determine the security attribute of the source endpoint by extracting the information from the certificate. The agent at the destination endpoint may determine the security attributes of the destination endpoint by mapping the destination IP to the security attributes based on the endpoint database local to the destination. For example, agent 333 may receive the certificate then examine the certificate (e.g., Cerf1) transmitted from the agent 313, and extract SG1 that is associated with the source endpoint 311 from the SAN field of Cerf1. The agent 333 may also determine the security group associated with the destination endpoint as specified by the packet (e.g., destination IP) based on a mapping relationship stored in the local endpoint database coupled to the agent 333.

[0090] In some cases, the agent at the destination side may also check policy database to determine an action according to the security attributes of the source and the destination endpoints. In some cases, the agent at the destination side may retrieve a certificate issued to the destination endpoint then transmitted the certificate to the agent at the source side for mutual authentication.

[0091] In some cases, the source endpoint certificate and destination endpoint certificate may be signed by different CAs. Different CAs may be intermediate CAs certified by a third party CA such as a root CA. **FIG. 6** shows an example of using multiple intermediate CAs 610, 620 in a system for establishing secure connections between endpoints.

[0092] In the illustrated example, multiple intermediate CAs such as cluster certificate authority 610 and cluster certificate 620 may obtain CA certificate respectively signed by an external certificate authority 640. The external CA may be a third party CA as described elsewhere herein. The third party CA may have a certificate 630 signed by a trusted CA such as itself if it is the root CA or another trusted CA.

[0093] An intermediate CA may be signed by the external certificate authority. An intermediate CA may be issued a certificate (e.g., certificate 400) by the external certificate authority 640. An intermediate CA may be configured to manage network security of a subnet. For example, cluster certificate authority 610 (i.e., intermediate CA) may be configured to sign certificates for nodes 611, 612 and associated endpoints 613, 614, 615, 616, cluster certificate authority 620 may be configured to sign certificates for nodes 621, 622 and associated endpoints 623, 624, 625, 626. The intermediate CA may then work in conjunction with an agent at a node to issue a certificate for a workload or endpoint (e.g., certificate 510). A connection between endpoints may then be allowed or denied based on a policy (e.g., policy 510) determined by security attributes of the endpoints as described elsewhere herein.

[0094] **FIG. 7** shows an exemplary process 700 for creating a secure and authenticated connection, in accordance with embodiments of the invention. The process 700 may begin by receiving a packet traversing a network (not shown). An agent at the source node or source endpoint may extract information related to the source endpoint and identify a certificate associated with the source endpoint 702. The information related to the source endpoint may be packet attribute related to the source endpoint such as IP address. The certificate may be determined by mapping the information related to the source endpoint to the certificate according to the mapping relationship stored in the endpoint database. The endpoint database may be a local database associated with the source node as described elsewhere herein.

[0095] Next, the source node may initiate a connection with the identified certificate 704. For example, the agent at the source node may initiate a handshake phase with the destination node or destination endpoint with the certificate issued to the source endpoint. In response to receiving the handshake request and the source endpoint certificate, the agent at the destination node may determine the security attribute associated with the destination endpoint 706. In some embodiments, the security attribute of the destination endpoint may be determined using the

mapping relationship stored in the endpoint database at the destination node. The agent at the destination node may further extract the security attribute contained in the certificate sent from the source node. Then the agent at the destination node may look up a policy database and determine a policy or action based on the security attributes of the source and destination endpoints 708.

[0096] Although **FIG. 7** shows a method in accordance with some embodiments a person of ordinary skill in the art will recognize that there are many adaptations for various embodiments. For example, the operations can be performed in any order. Some of the operations may be precluded, some of the operations may be performed concurrently in one step, some of the operations repeated, and some of the operations may comprise sub-steps of other operations. The method may also be modified in accordance with other aspects of the disclosure as provided herein

[0097] **FIG. 8** illustrates another process 800 performed by the system for mutual authentication between endpoints of a connection. Similar to the aforementioned process, the mutual authentication process 800 may begin by receiving a packet traversing a network (not shown). An agent at the source node or source endpoint may extract information related to the source endpoint and identify a certificate associated with the source endpoint 802. The information related to the source endpoint may be packet attribute related to the source endpoint such as IP address. The certificate may be determined by mapping the information related to the source endpoint to the certificate based on the endpoint database. The endpoint database may be a local database associated with the source node.

[0098] Next, the source node may initiate a connection with the identified certificate 804. For example, the agent at the source node may initiate a handshake phase with the destination node or destination endpoint using the certificate issued to the source endpoint. In response to receiving the handshake request and the source endpoint certificate, the agent at the destination node may determine the security attribute associated with the destination endpoint 806. In some embodiments, the security attribute of the destination endpoint may be determined using the mapping relationship stored in the endpoint database at the destination node. In some cases, the agent at the destination node may further extract the security attribute contained in the certificate sent from the source node to verify the authenticity of the source endpoint. Next, the agent at the destination node may identify a certificate issued to the destination endpoint and send it back to the source node 808. Upon receiving the destination endpoint certificate, the agent at the source node may extract the security attribute of the destination endpoint from the certificate. Then the agent at the source node may look up a policy database and determine a policy or action based

on the security attributes of the source and destination endpoints 810.

[0099] The agent can have one or more processors and at least one memory for storing program instructions. The processors can be part of the network interface system or device. Alternatively or additionally, the processors can be part of the host system. The processor(s) can be a single or multiple microprocessors, field programmable gate arrays (FPGAs), or digital signal processors (DSPs) capable of executing particular sets of instructions. Computer-readable instructions can be stored on a tangible non-transitory computer-readable medium, such as a flexible disk, a hard disk, a CD-ROM (compact disk-read only memory), and MO (magneto-optical), a DVD-ROM (digital versatile disk-read only memory), a DVD RAM (digital versatile disk-random access memory), or a semiconductor memory. Alternatively, the agent can be implemented in hardware components (e.g., ASICs, special purpose computers, or general purpose computers), software or combinations of hardware and software.

[0100] Methods according to the above-described examples can be implemented using computer-executable instructions that are stored or otherwise available from computer readable media. Such instructions can comprise, for example, instructions and data which cause or otherwise configure a general purpose computer, special purpose computer, or special purpose processing device to perform a certain function or group of functions. Portions of computer resources used can be accessible over a network. The computer executable instructions may be, for example, binaries, intermediate format instructions such as assembly language, firmware, or source code. Examples of computer-readable media that may be used to store instructions, information used, and/or information created during methods according to described examples include magnetic or optical disks, flash memory, USB devices provided with non-volatile memory, networked storage devices, and so on.

[0101] The instructions, media for conveying such instructions, computing resources for executing them, and other structures for supporting such computing resources are means for providing the functions described in these disclosures.

[0102] While preferred embodiments of the present invention have been shown and described herein, it will be obvious to those skilled in the art that such embodiments are provided by way of example only. Numerous variations, changes, and substitutions will now occur to those skilled in the art without departing from the invention. It should be understood that various alternatives to the embodiments of the invention described herein may be employed in practicing the invention. It is intended that the following claims define the scope of the invention and that methods and structures within the scope of these claims and their equivalents be covered thereby.

CLAIMS

WHAT IS CLAIMED IS:

1. A method for establishing a secure and authenticated network connection, the method comprising:

   a) receiving, from a requesting entity, a destination IP address and a first certificate that is used to establish a secure network connection, wherein the first certificate comprises a first security attribute that is associated with a source destination IP address;

   b) identifying, with aid of one or more processors, a stored second security attribute associated with the destination IP address; and

   c) determining, with aid of the one or more processors, a policy action based at least in part on the first security attribute and the second security attribute.

2. The method of claim 1, further comprising transmitting, from the requesting entity, a certificate signing request (CSR) to an intermediate certificate authority, wherein the CSR comprising a security attribute and an associated IP address.

3. The method of claim 2, wherein a mapping relationship between the security attribute and the IP address is received from a computing device running the intermediate certificate authority.

4. The method of claim 3, wherein the mapping relationship is at least in part managed by the intermediate certificate authority.

5. The method of claim 2, further comprising receiving a certificate from the intermediate certificate authority, wherein the certificate comprises the security attribute.

6. The method of claim 2, further comprising storing the certificate in a local database managed by the requesting entity.

7. The method of claim 2, wherein the intermediate certificate authority is verified by a root certificate authority.

8. The method of claim 1, further comprising identifying, with aid of the one or more processors, a second stored certificate associated with the destination IP address.

9. The method of claim 8, further comprising transmitting the second stored certificate to the requesting entity.

10. The method of claim 8, wherein the first certificate and the second certificate are issued by the same intermediate certificate authority.

11. The method of claim 8, wherein the first certificate and the second certificate are issued by different intermediate certificate authorities.

12. The method of claim 11, wherein the different intermediate certificate authorities are verified by the same root certificate authority.

13. The method of claim 1, further comprising executing the policy action determined in (c) on a base system logic or by the requesting entity.

14. The method of claim 13, wherein the base system logic comprises a bare metal server, a hypervisor or a docker base for controlling one or more virtual machines or containers.

15. The method of claim 1, wherein the policy action comprises at least one of the following: allow, block, rate-limit, mark and alert.

16. The method of claim 1, wherein the first security attribute comprises a security group the source IP address is associated therewith.

17. The method of claim 11, wherein the second security attribute comprises a security group the destination IP address is associated therewith.

18. A network interface device comprising: a memory for storing a set of instructions and one or more processors configured to execute the set of instructions to perform at least the following:

a) receive, from a requesting entity, a destination IP address and a first certificate that is used to establish a secure network connection, wherein the first certificate comprises a first security attribute that is associated with a source IP address;

b) identify a stored second security attribute associated with the destination IP address; and

c) determine a policy action based at least in part on the first security attribute and the second security attribute.

19. The network interface device of claim 18, wherein the one or more processors further configured to generate and transmit a certificate signing request (CSR) to an intermediate certificate authority, wherein the CSR comprising a security attribute and an associated IP address.

20. The network interface device of claim 19, wherein a mapping relationship between the security attribute and the IP address is received from a computing device running the intermediate certificate authority.

21. The network interface device of claim 20, wherein the mapping relationship is at least in part managed by the intermediate certificate authority.

22. The network interface device of claim 19, wherein the one or more processors further configured to receive a signed certificate from the intermediate certificate authority, and wherein the signed certificate comprises the security attribute.

23. The network interface device of claim 22, wherein the one or more processors further configured to store the signed certificate in a local database managed by the network interface device.

24. The network interface device of claim 19, wherein the intermediate certificate authority is verified by a root certificate authority.

25. The network interface device of claim 18, wherein the one or more processors further configured to identity a second stored certificate associated with the destination IP address.

26. The network interface device of claim 25, wherein the one or more processors further configured to transmit the second stored certificate to the requesting entity.

27. The network interface device of claim 25, wherein the first certificate and the second certificate are issued by the same intermediate certificate authority.

28. The network interface device of claim 25, wherein the first certificate and the second certificate are issued by different intermediate certificate authorities.

29. The network interface device of claim 28, wherein the different intermediate certificate authorities are verified by the same root certificate authority.

30. The network interface device of claim 18, wherein the policy action determined in (c) is executed on a base system logic.

31. The network interface device of claim 30, wherein the base system logic comprises a bare metal server, a hypervisor or a docker base for controlling one or more virtual machines or containers and wherein the bare metal server, the hypervisor or the docker base is coupled to the network interface device.

32. The network interface device of claim 18, wherein the policy action is selected from the group comprising: allow, block, rate-limit, mark and alert.

33. The network interface device of claim 18, wherein the first security attribute comprises a security group the source IP address is associated therewith.

34. The network interface device of claim 18, wherein the second security attribute comprises a security group the destination IP address is associated therewith.

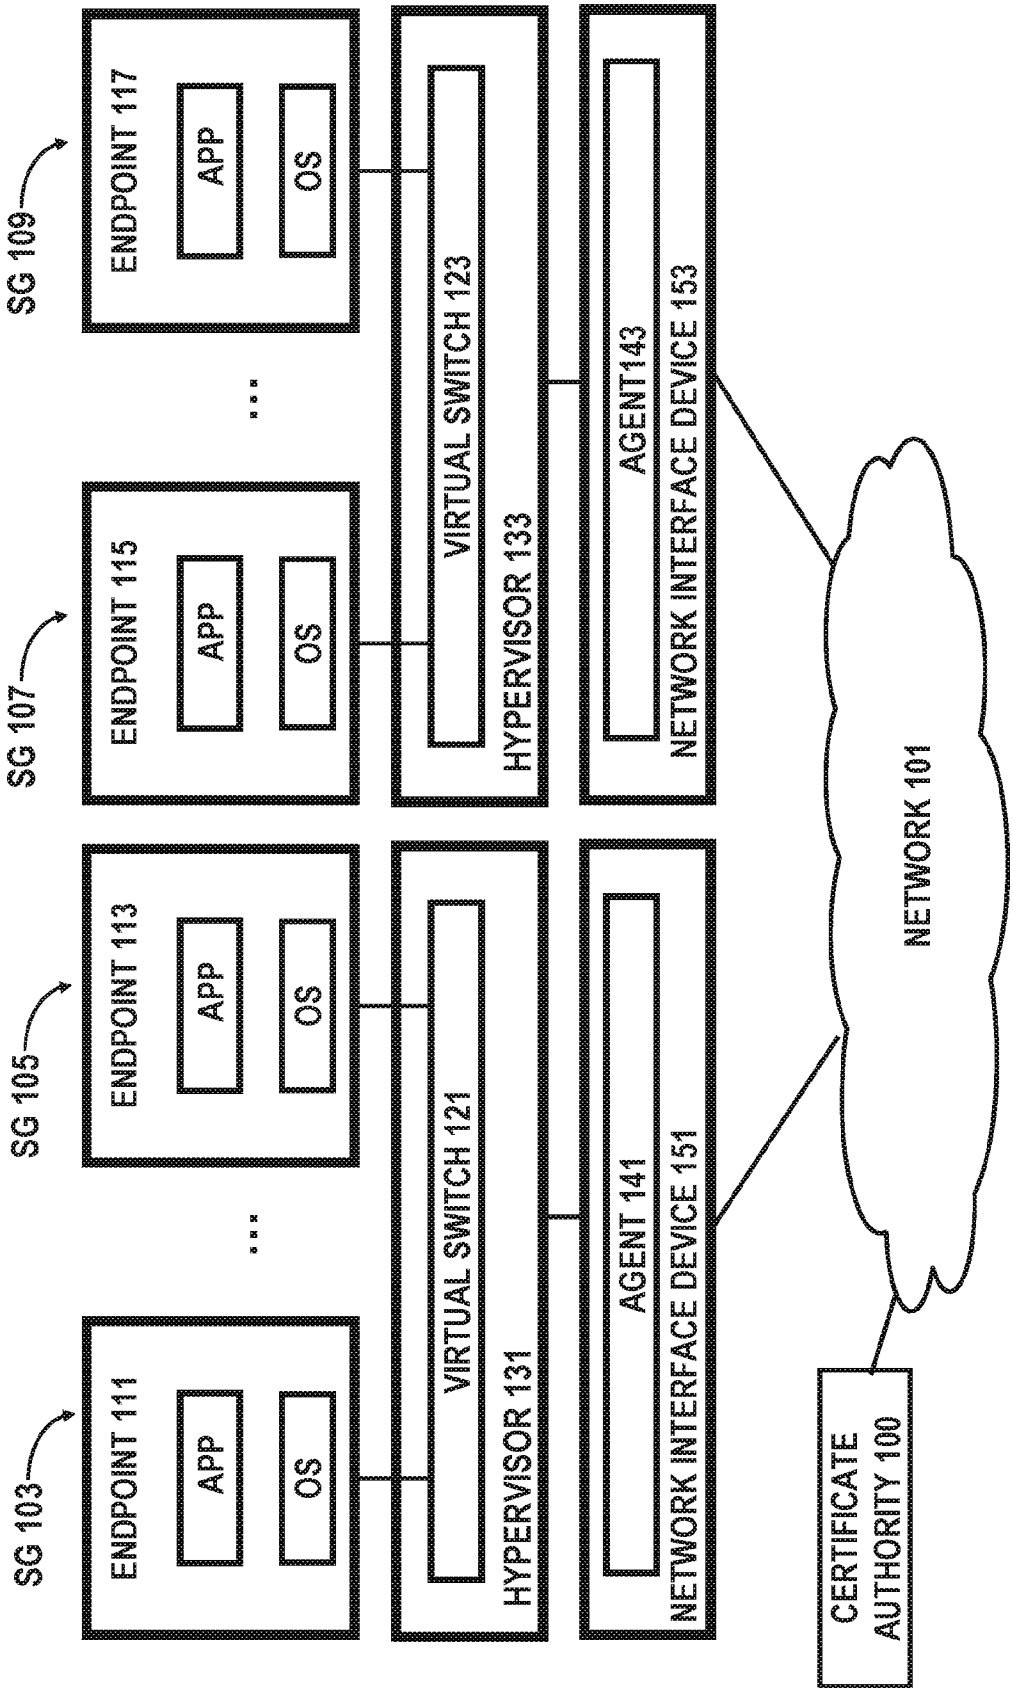
SG 103
ENDPOINT 111
APP
OS
SG 105
ENDPOINT 113
APP
OS
SG 107
ENDPOINT 115
APP
OS
SG 109
ENDPOINT 117
APP
OS
VIRTUAL SWITCH 121
HYPERVISOR 131
VIRTUAL SWITCH 123
HYPERVISOR 133
AGENT 141
NETWORK INTERFACE DEVICE 151
AGENT143
NETWORK INTERFACE DEVICE 153
NETWORK 101
CERTIFICATE AUTHORITY 100

FIG. 1A

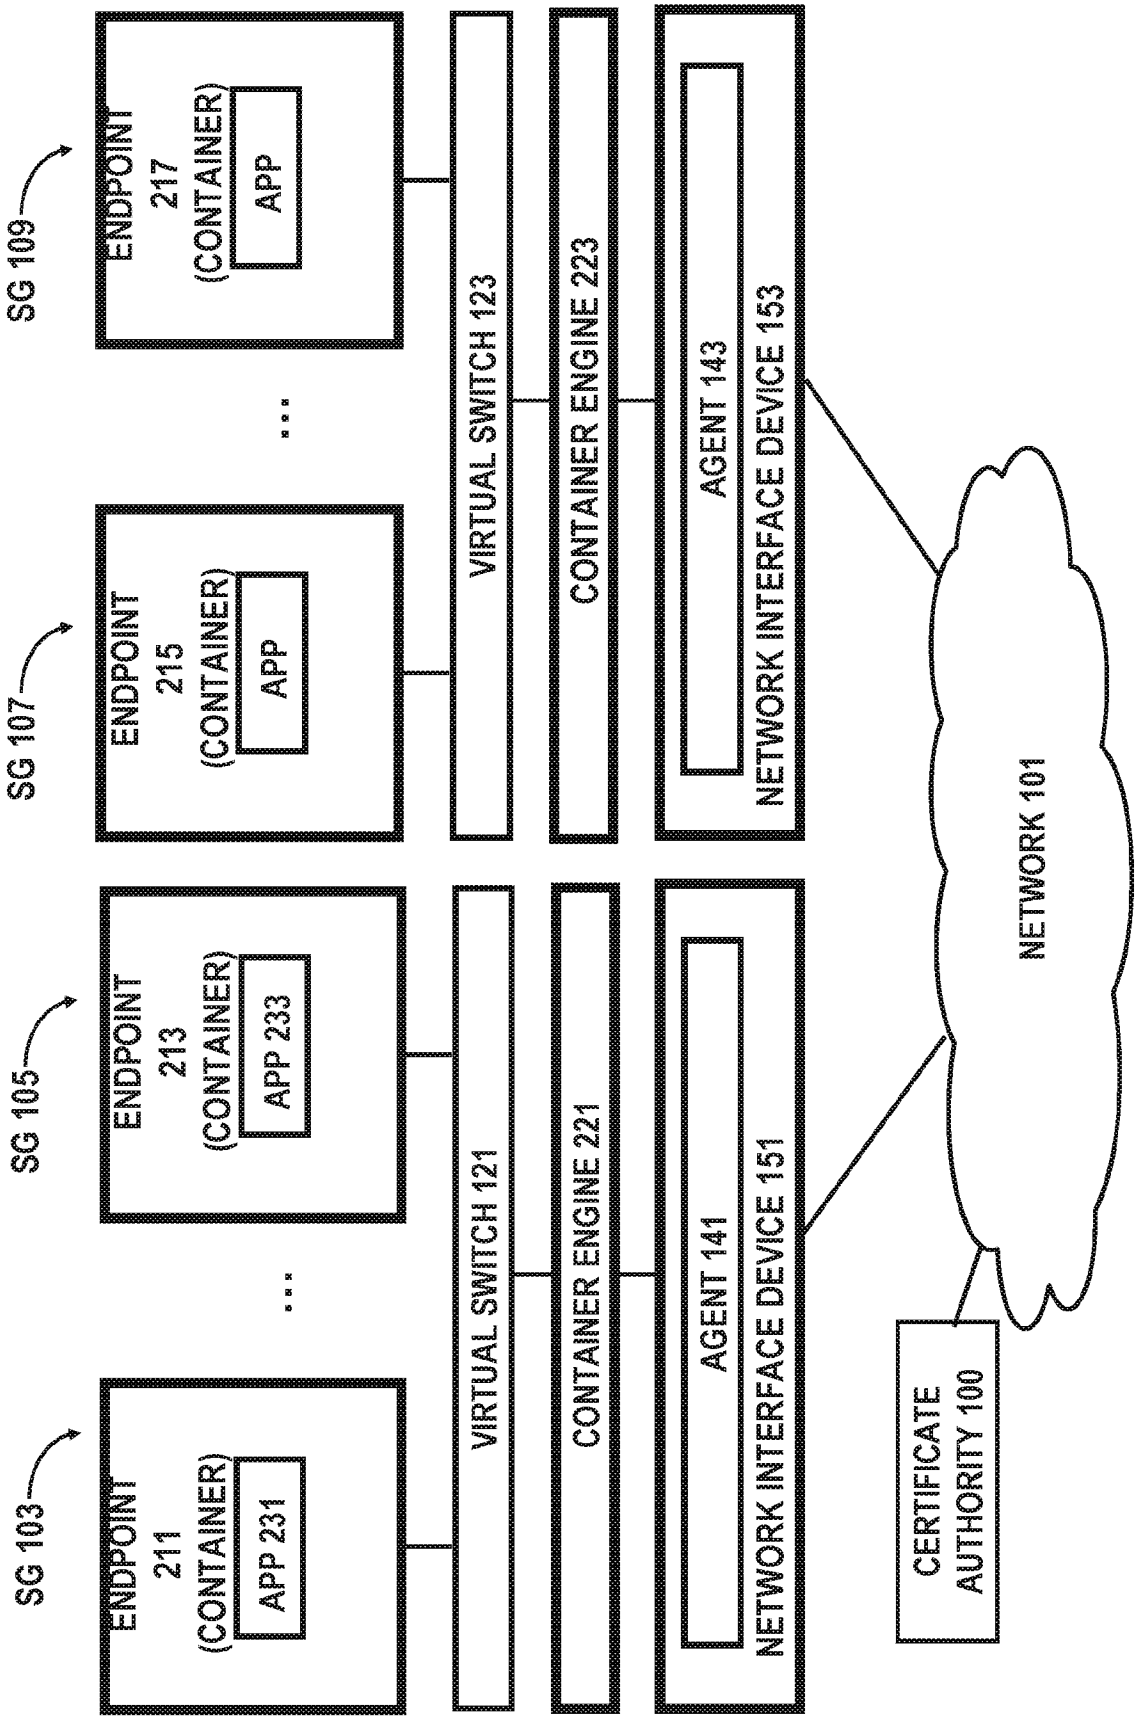
SG 103
ENDPOINT 211 (CONTAINER)
APP 231
SG 105
ENDPOINT 213 (CONTAINER)
APP 233
SG 107
ENDPOINT 215 (CONTAINER)
APP
SG 109
ENDPOINT 217 (CONTAINER)
APP
VIRTUAL SWITCH 121
VIRTUAL SWITCH 123
CONTAINER ENGINE 221
CONTAINER ENGINE 223
AGENT 141
AGENT 143
NETWORK INTERFACE DEVICE 151
NETWORK INTERFACE DEVICE 153
NETWORK 101
CERTIFICATE AUTHORITY 100

FIG. 1B

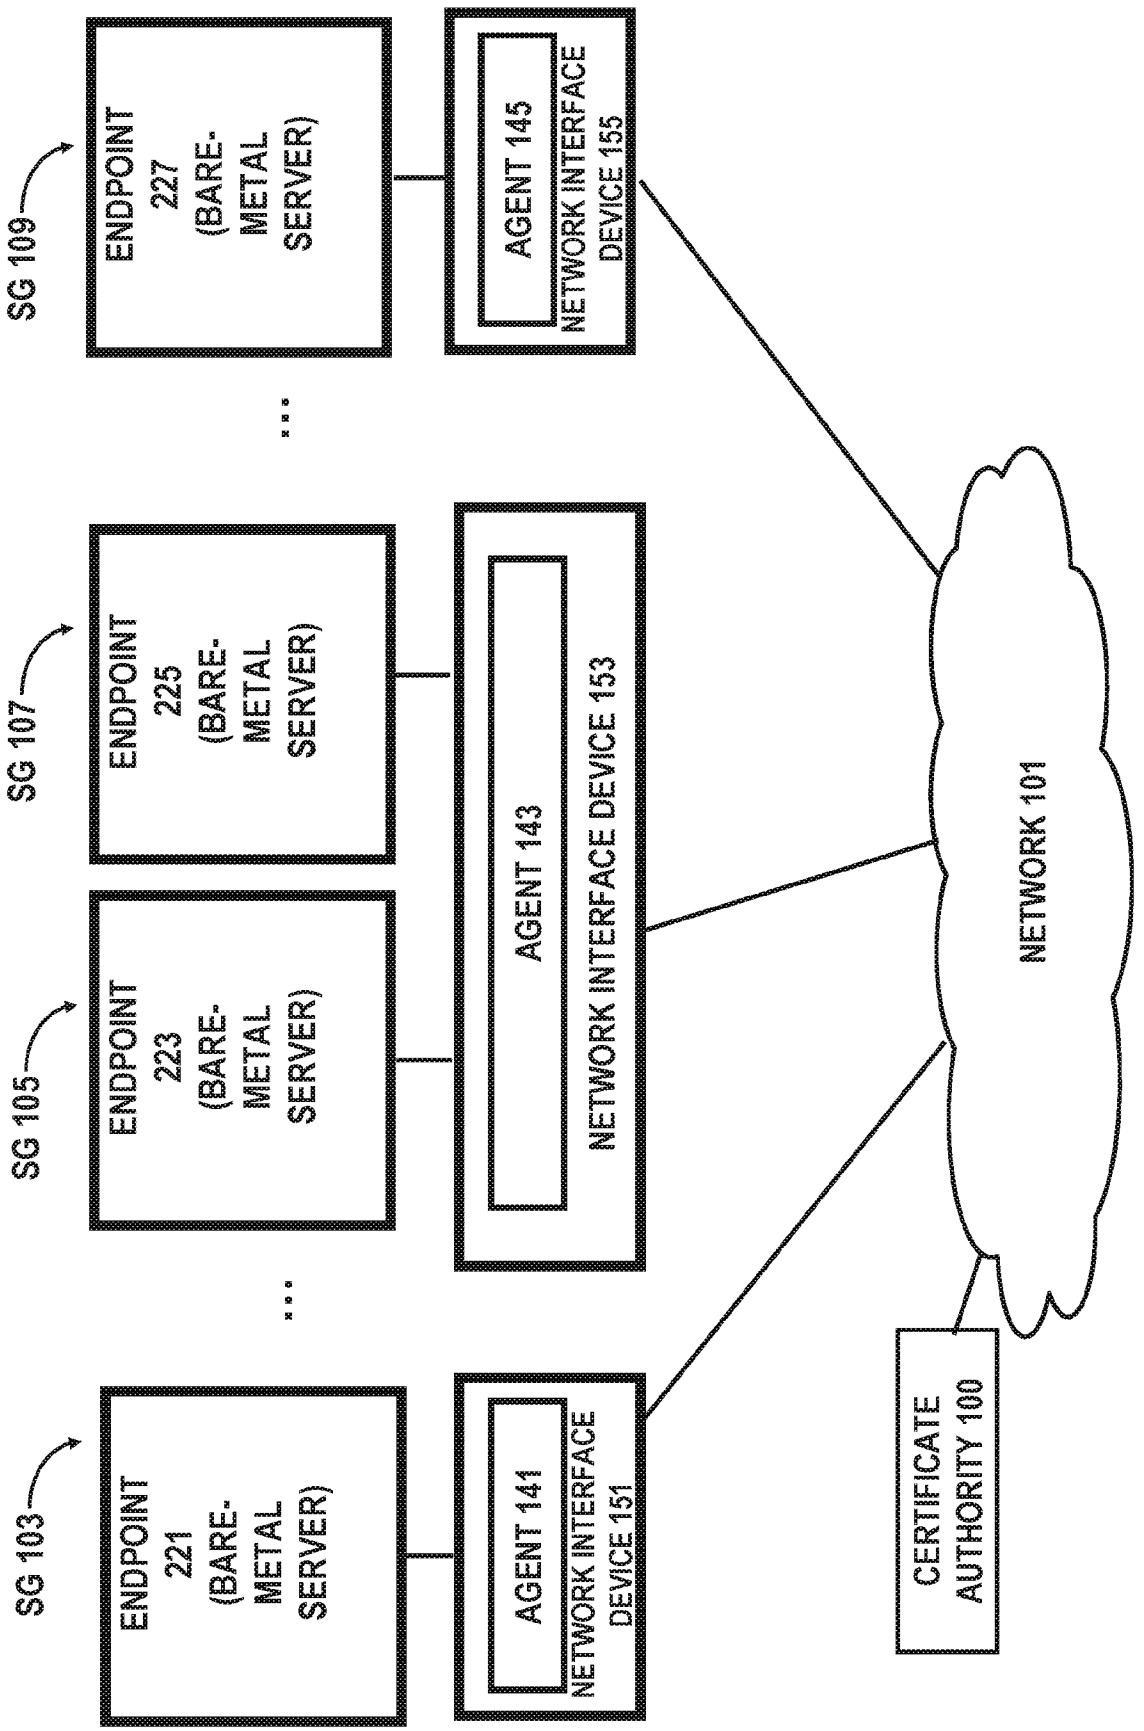
SG 109
ENDPOINT 227 (BARE-METAL SERVER)
AGENT 145
NETWORK INTERFACE DEVICE 155
SG 107
ENDPOINT 225 (BARE-METAL SERVER)
SG 105
ENDPOINT 223 (BARE-METAL SERVER)
AGENT 143
NETWORK INTERFACE DEVICE 153
SG 103
ENDPOINT 221 (BARE-METAL SERVER)
AGENT 141
NETWORK INTERFACE DEVICE 151
NETWORK 101
CERTIFICATE AUTHORITY 100

FIG. 2

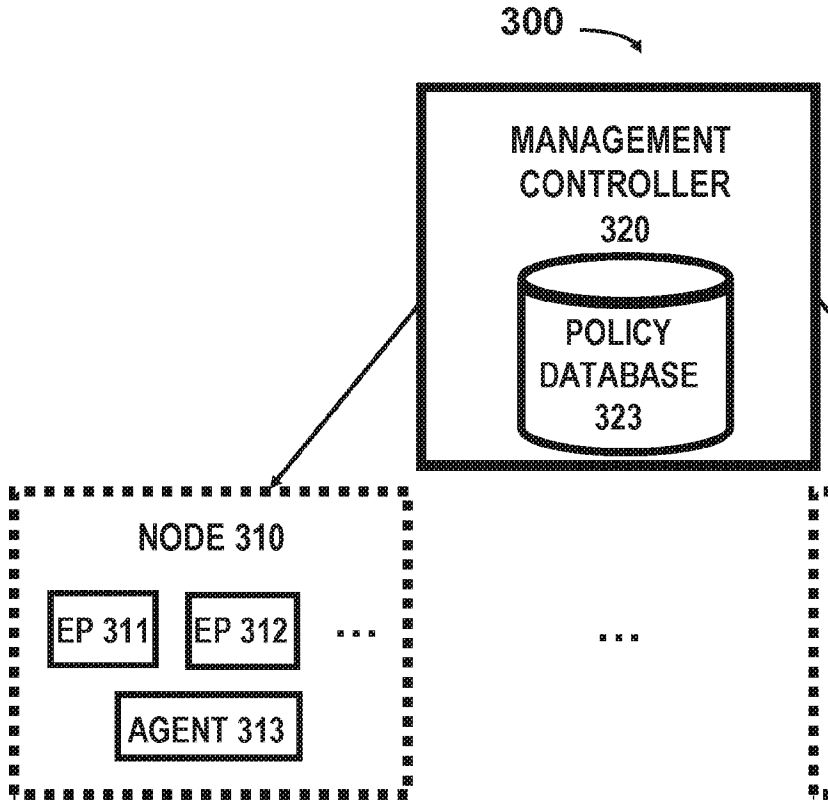

Policy Table 325

Security attribute 1 327	Security attribute 2 328	Action 329
SG1	SG2	Allow
SG2	SG1	Allow

Endpoint DB 315

EP ID	IP	Security Attribute	Certificate
VM1	IP1	SG1	Cerf 1
VM2	IP2	SG2	Cerf 2

Endpoint DB 335

EP ID	IP	Security Attribute	Certificate
VM3	IP3	SG2	Cerf 3
VM4	IP4	SG3	Cerf 4

FIG. 3

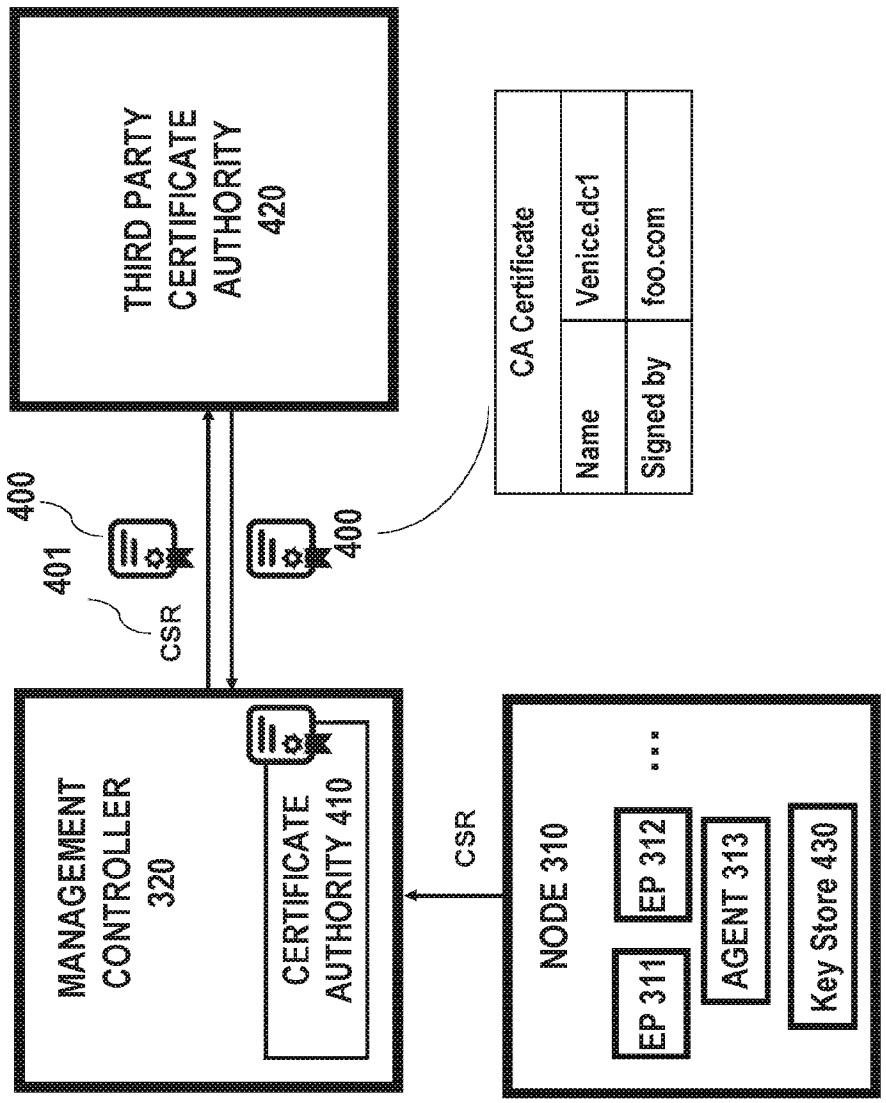
THIRD PARTY CERTIFICATE AUTHORITY 420
CA Certificate
Name
Venice.dc1
Signed by
foo.com
400
401
CSR
400
MANAGEMENT CONTROLLER 320
CERTIFICATE AUTHORITY 410
CSR
NODE 310
EP 311
EP 312
...
AGENT 313
Key Store 430

FIG. 4

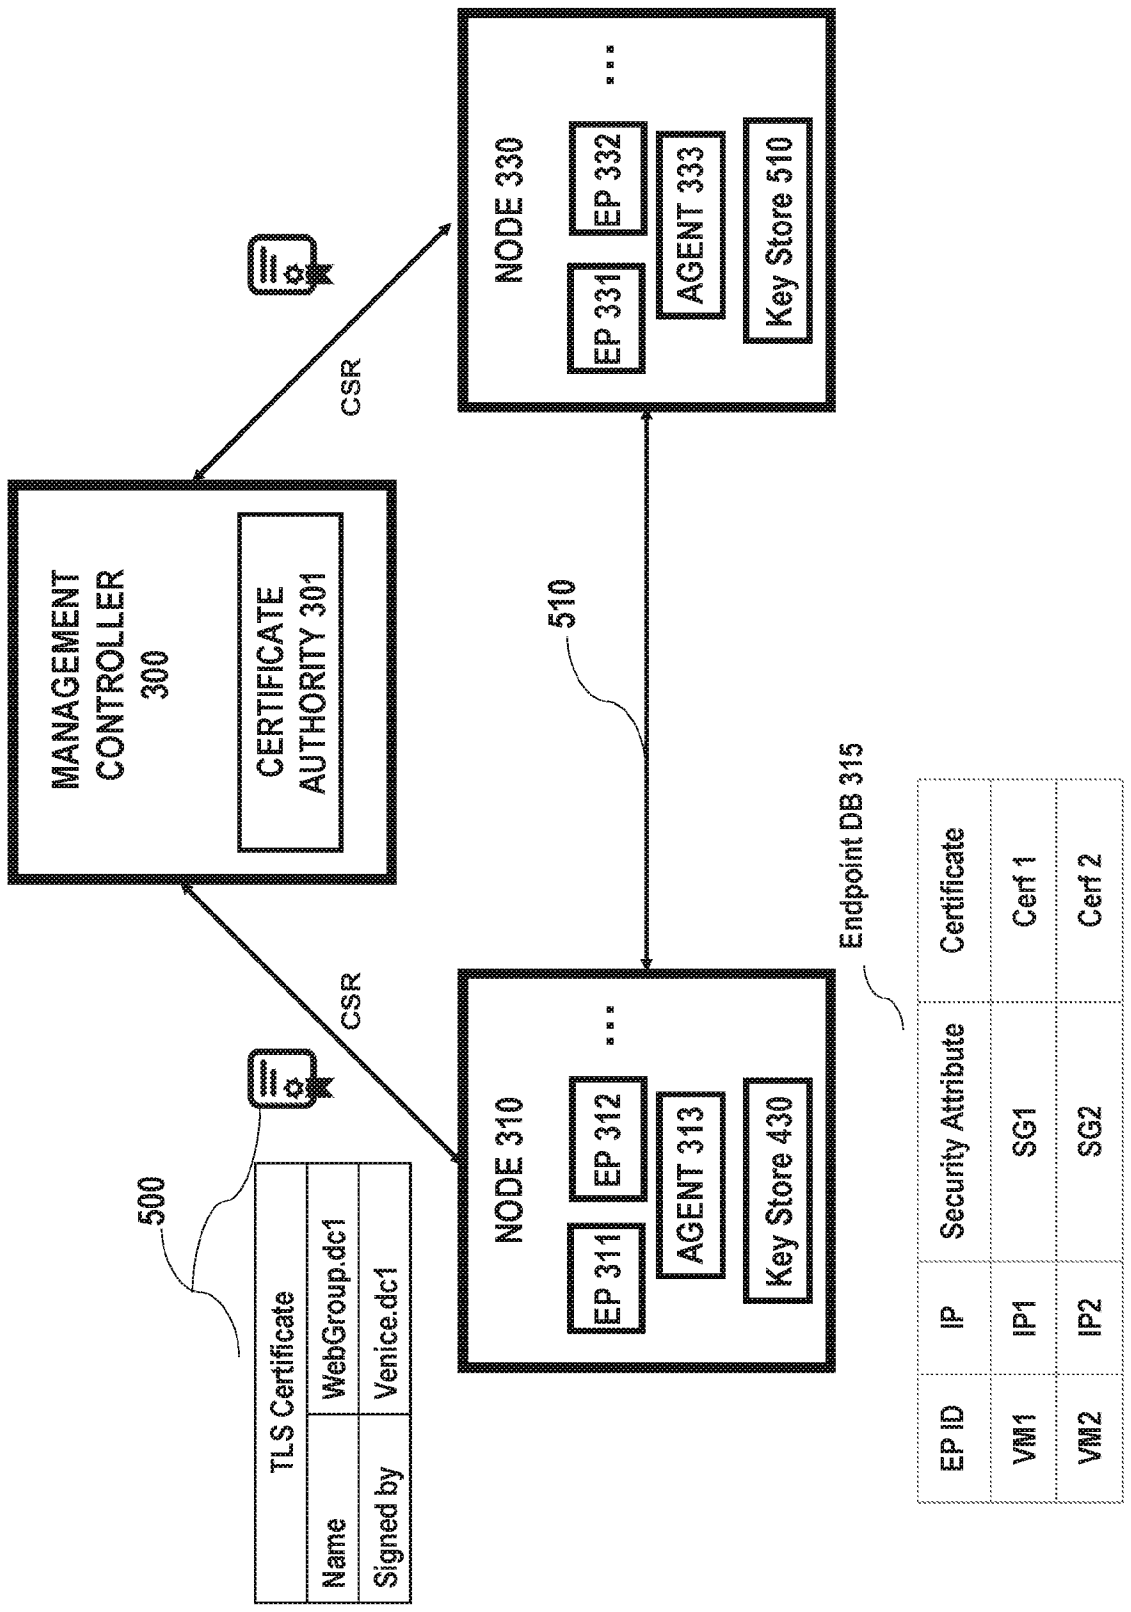
MANAGEMENT CONTROLLER 300
CERTIFICATE AUTHORITY 301
CSR
CSR
NODE 330
EP 331
EP 332
AGENT 333
Key Store 510
NODE 310
EP 311
EP 312
AGENT 313
Key Store 430
510
500
TLS Certificate
Name
WebGroup.dc1
Signed by
Venice.dc1
Endpoint DB 315
EP ID
IP
Security Attribute
Certificate
VM1
IP1
SG1
Cert 1
VM2
IP2
SG2
Cert 2

FIG. 5

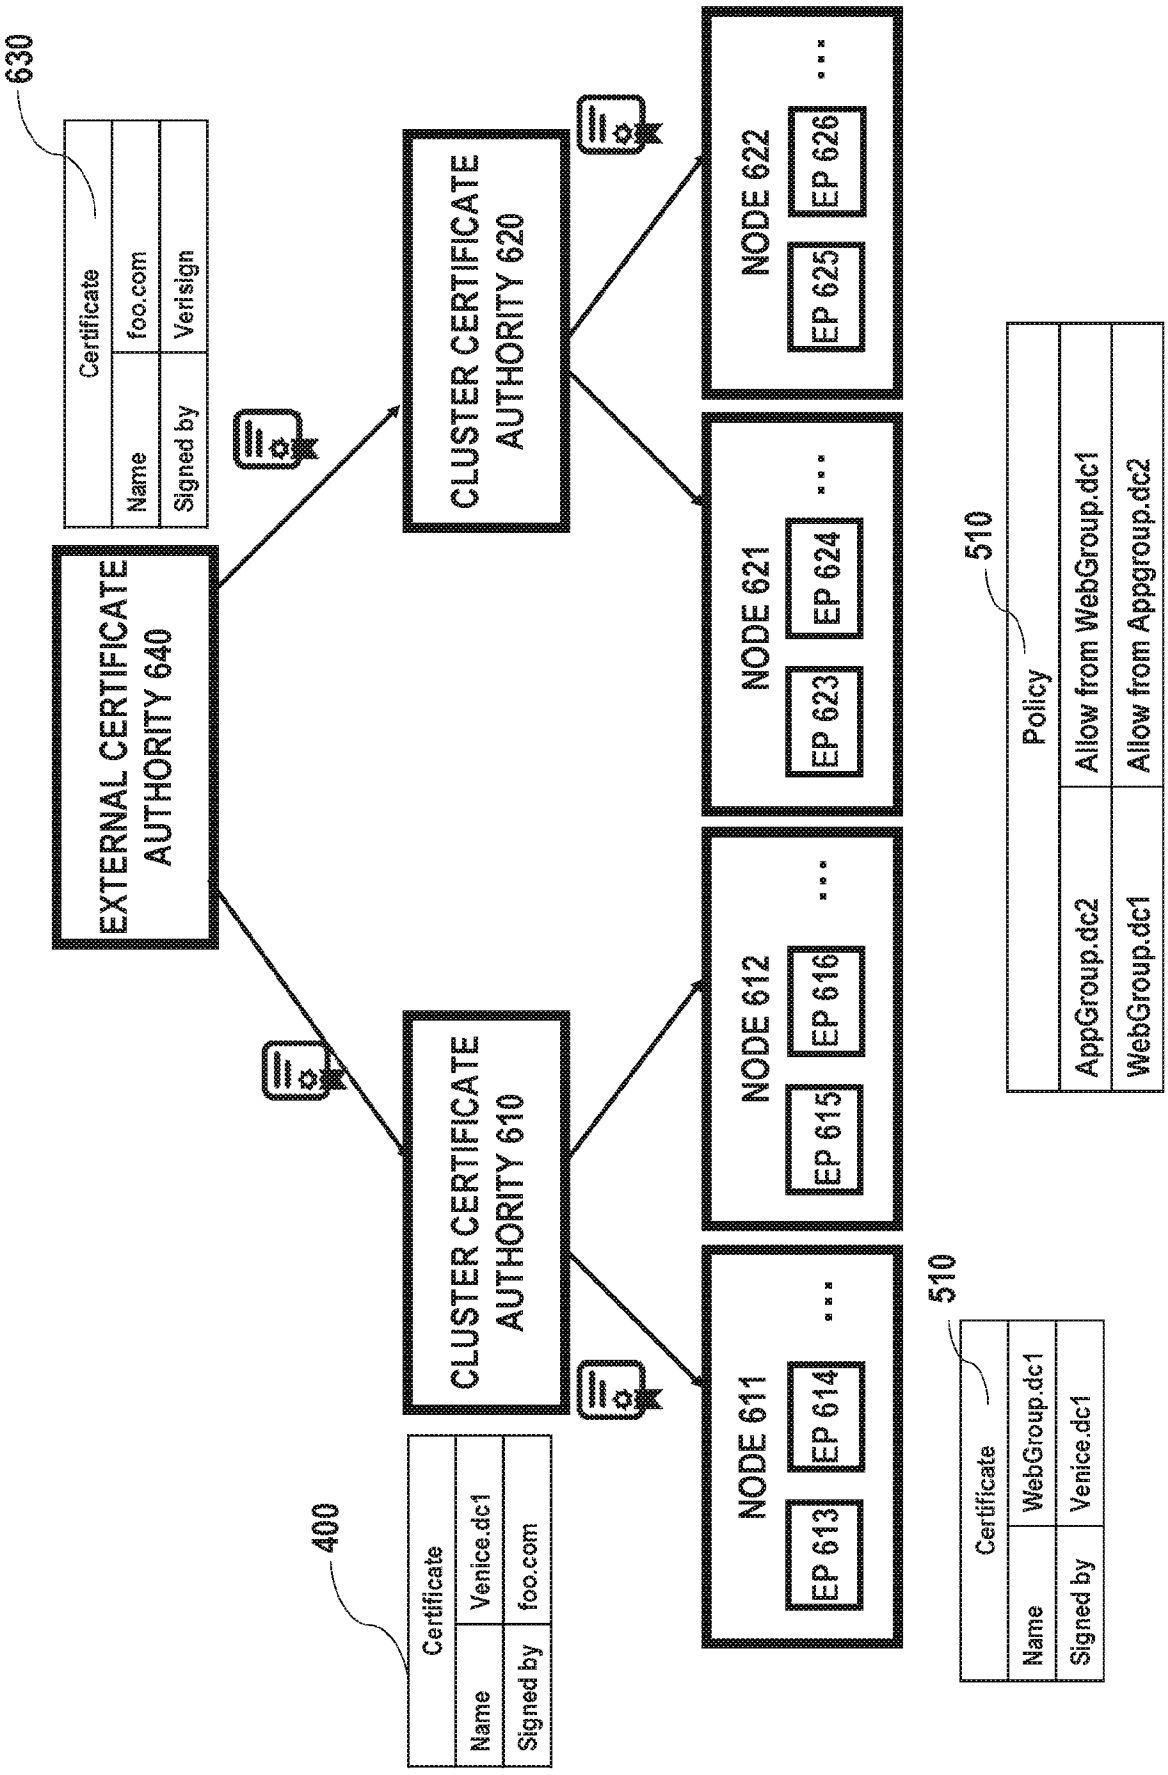
630
Certificate
Name
foo.com
Signed by
Verisign
EXTERNAL CERTIFICATE AUTHORITY 640
CLUSTER CERTIFICATE AUTHORITY 620
CLUSTER CERTIFICATE AUTHORITY 610
NODE 622
EP 625
EP 626
NODE 621
EP 623
EP 624
NODE 612
EP 615
EP 616
NODE 611
EP 613
EP 614
510
Policy
AppGroup.dc2
Allow from WebGroup.dc1
WebGroup.dc1
Allow from Appgroup.dc2
400
Certificate
Name
Venice.dc1
Signed by
foo.com
510
Certificate
Name
WebGroup.dc1
Signed by
Venice.dc1

FIG. 6

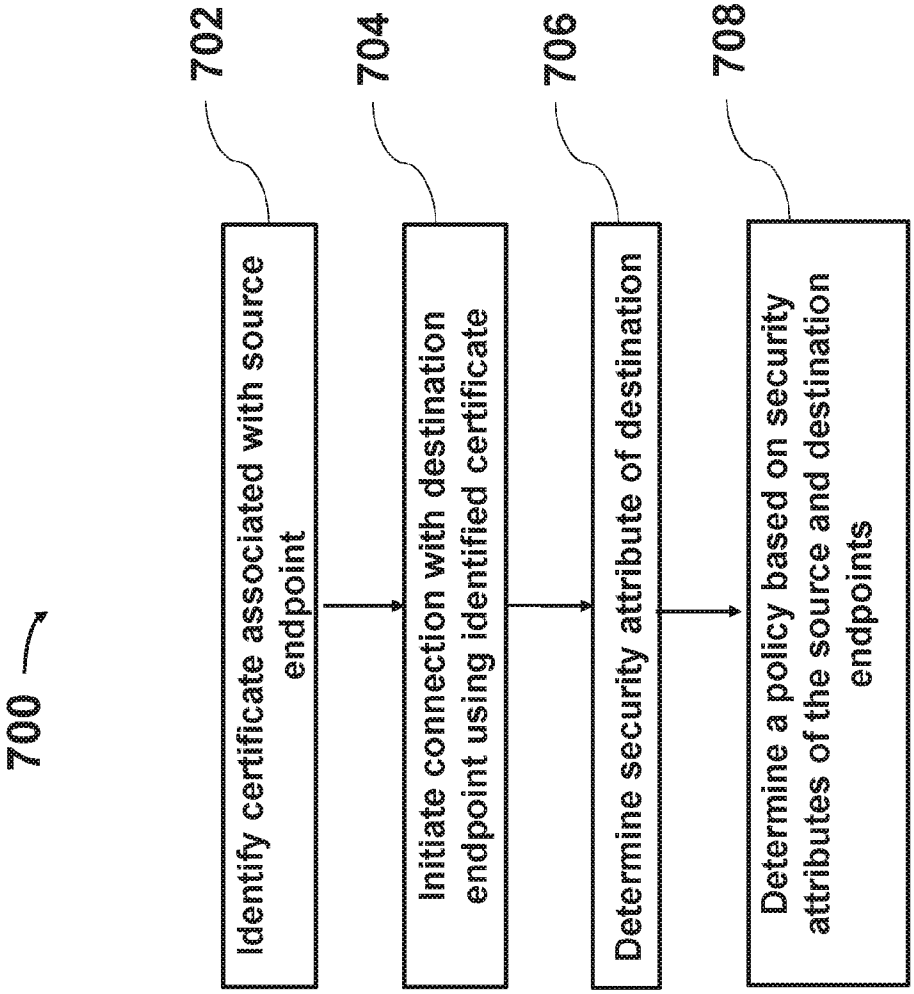
700
702
Identify certificate associated with source endpoint
704
Initiate connection with destination endpoint using identified certificate
706
Determine security attribute of destination
708
Determine a policy based on security attributes of the source and destination endpoints

FIG. 7

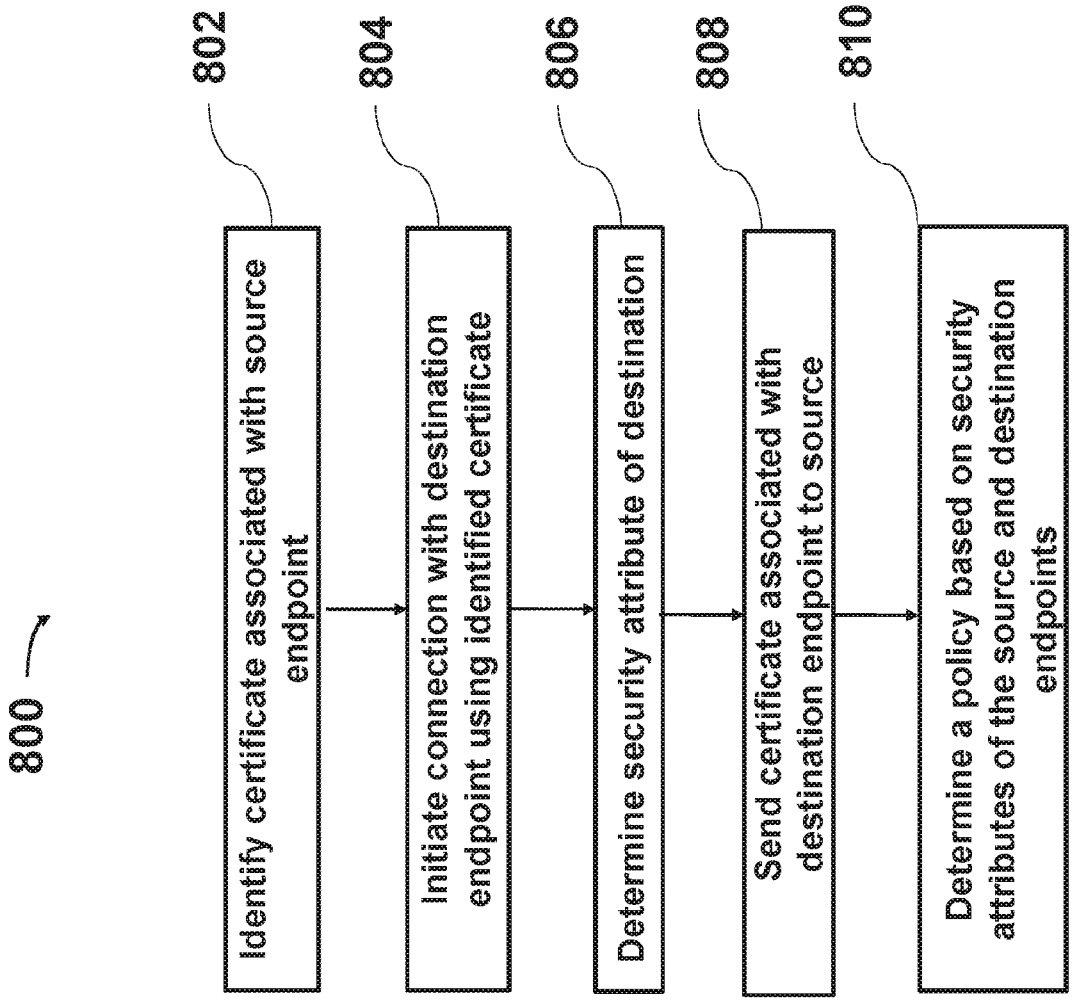
800
802
Identify certificate associated with source endpoint
804
Initiate connection with destination endpoint using identified certificate
806
Determine security attribute of destination
808
Send certificate associated with destination endpoint to source
810
Determine a policy based on security attributes of the source and destination endpoints

FIG. 8

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2018/066801

A. CLASSIFICATION OF SUBJECT MATTER

H04L 9/32(2006.01)i, H04L 9/08(2006.01)i, H04L 29/06(2006.01)i, H04L 12/931(2013.01)i

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
H04L 9/32; G06F 11/00; G06F 12/14; G06F 15/16; G06F 17/00; H04L 29/06; H04L 9/08; H04L 12/931

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean utility models and applications for utility models
Japanese utility models and applications for utility models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & Keywords: requesting entity, destination IP address, source IP address, first certificate, first security attribute, second security attribute, policy action

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2010-0306816 A1 (DAVID MCGREW et al.) 02 December 2010 See paragraphs [0017], [0026]-[0029], [0041]; claims 1-4, 6, 14, 17-18; and figures 1, 3.	1-34
Y	US 2016-0036778 A1 (A10 NETWORKS, INC.) 04 February 2016 See paragraph [0090]; and claim 1.	1-34
A	US 2007-0261112 A1 (JOHN TODD et al.) 08 November 2007 See paragraphs [0113]-[0123]; and figure 6.	1-34
A	US 2010-0131646 A1 (DEAN DRAKO) 27 May 2010 See paragraph [0068]; and figure 6.	1-34
A	US 8413238 B1 (MICHAEL SUTTON) 02 April 2013 See column 11, line 35 - column 13, line 9; and figure 4.	1-34

☐ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents:
"A" document defining the general state of the art which is not considered to be of particular relevance
"E" earlier application or patent but published on or after the international filing date
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
"O" document referring to an oral disclosure, use, exhibition or other means
"P" document published prior to the international filing date but later than the priority date claimed
"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents,such combination being obvious to a person skilled in the art
"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
01 April 2019 (01.04.2019)	**02 April 2019 (02.04.2019)**

Name and mailing address of the ISA/KR
International Application Division
Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea
Facsimile No. +82-42-481-8578

Authorized officer
KIM, Seong Woo
Telephone No. +82-42-481-3348

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2018/066801

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2010-0306816 A1	02/12/2010	US 8806572 B2	12/08/2014
US 2016-0036778 A1	04/02/2016	CN 101361037 A	04/02/2009
		CN 101361037 B	13/07/2011
		CN 102098316 A	15/06/2011
		CN 102098316 B	16/09/2015
		CN 102123156 A	13/07/2011
		CN 102123156 B	05/11/2014
		CN 102918801 A	06/02/2013
		CN 102918801 B	25/05/2016
		EP 2577910 A2	10/04/2013
		EP 2577910 A4	16/12/2015
		JP 2013-528330 A	08/07/2013
		JP 5946189 B2	05/07/2016
		US 2008-0148357 A1	19/06/2008
		US 2010-0217819 A1	26/08/2010
		US 2010-0235880 A1	16/09/2010
		US 2011-0239289 A1	29/09/2011
		US 2012-0216266 A1	23/08/2012
		US 2014-0059702 A1	27/02/2014
		US 2015-0312237 A1	29/10/2015
		US 2016-0050233 A1	18/02/2016
		US 2016-0105395 A1	14/04/2016
		US 2016-0105446 A1	14/04/2016
		US 2016-0119382 A1	28/04/2016
		US 2016-0182456 A1	23/06/2016
		US 2016-0261642 A1	08/09/2016
		US 2017-0041350 A1	09/02/2017
		US 2017-0289106 A1	05/10/2017
		US 2017-0295185 A1	12/10/2017
		US 7716378 B2	11/05/2010
		US 7979585 B2	12/07/2011
		US 8312507 B2	13/11/2012
		US 8423676 B2	16/04/2013
		US 8584199 B1	12/11/2013
		US 8595383 B2	26/11/2013
		US 8595791 B1	26/11/2013
		US 8813180 B1	19/08/2014
		US 8826372 B1	02/09/2014
		US 8868765 B1	21/10/2014
		US 9060003 B2	16/06/2015
		US 9219751 B1	22/12/2015
		US 9253152 B1	02/02/2016
		US 9270705 B1	23/02/2016
		US 9294467 B2	22/03/2016
		US 9350744 B2	24/05/2016
		US 9356910 B2	31/05/2016
		US 9497201 B2	15/11/2016
		US 9661026 B2	23/05/2017

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.
PCT/US2018/066801

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
		US 9712493 B2	18/07/2017
		US 9954868 B2	24/04/2018
		US 9954899 B2	24/04/2018
		WO 2008-067013 A2	05/06/2008
		WO 2008-067013 A3	04/09/2008
		WO 2011-149796 A2	01/12/2011
		WO 2011-149796 A3	19/04/2012
US 2007-0261112 A1	08/11/2007	CA 2587867 A1	08/11/2007
		CA 2587867 C	23/06/2015
		US 7890612 B2	15/02/2011
US 2010-0131646 A1	27/05/2010	US 8447856 B2	21/05/2013
US 8413238 B1	02/04/2013	None	