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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
20 September 2012 (20.09.2012)

WIPO | PCT

(10) International Publication Number
WO 2012/125137 A1

(51) International Patent Classification:
H04L 12/26 (2006.01) *H04L 29/02* (2006.01)
H04L 12/56 (2006.01)

(21) International Application Number:
PCT/US2011/028043

(22) International Filing Date:
11 March 2011 (11.03.2011)

(25) Filing Language: English

(26) Publication Language: English

(71) Applicant *(for all designated States except US)*: **HEWLETT-PACKARD DEVELOPMENT COMPANY, L.P.** [US/US]; 11445 Compaq Center Drive W., Houston, Texas 77070 (US).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **HALL, Matthew, Richard, Thomas** [US/US]; 2527 Marchese Wy., Santa Clara, California 95051 (US). **KOORNSTRA, Reinoud, Jelmer, Jeroen** [NL/US]; Hewlett-Packard Co., 8000 Foothills Blvd. Stop 5541, Roseville, California 95747 (US). **WORTH, Kevin, M.** [US/US]; Hewlett-Packard Co., 8000 Foothills Blvd. Stop 5541, Roseville, California 95747 (US).

(74) Agents: **VOISINET, Catherine, M.** et al.; Hewlett-Packard Company, Intellectual Property Administration, 3404 E. Harmony Road, Mail Stop 35, Fort Collins, Colorado 80528 (US).

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

[Continued on next page]

(54) Title: SAMPLING NETWORK TRAFFIC

(57) Abstract: Sampling network traffic includes: loading a packet sampling module (215) into a processor-based network device (201) coupled to a network (235); determining with the packet sampling module (215) if a network packet addressed to or from the network device (235) is selected for sampling; and transmitting data from the network packet over the network (235) to a monitoring device (240) external to the network device (201) if the network packet is selected for sampling.

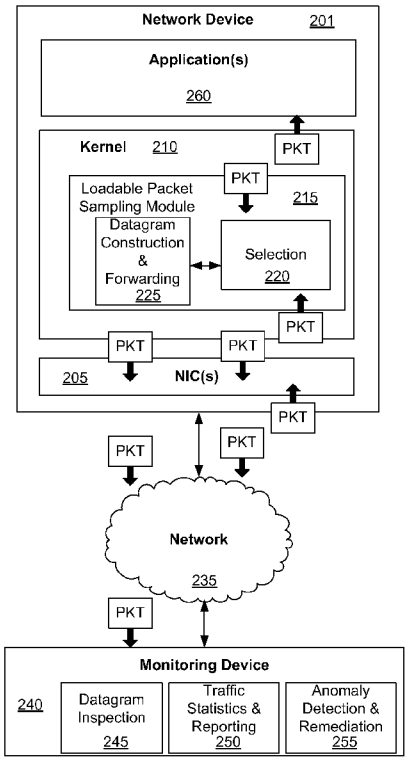

Fig. 2A

WO 2012/125137 A1

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

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

Declarations under Rule 4.17:

— *as to the identity of the inventor (Rule 4.17(i))*

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

Published:

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

Sampling Network Traffic

BACKGROUND

[0001] Organizations continue to rely on networks of interconnected devices to exchange information and provide services. Accordingly, the size of many computer networks continues to grow, along with the amount of data exchanged over the networks. With this growth come increased threats to network security and network efficiency. These threats may include malicious network traffic designed to exploit vulnerabilities in network devices to compromise network security and unnecessary or unwanted network traffic that consumes resources and degrades network performance.

[0002] To detect such threats and manage network traffic flow generally, a network may utilize network traffic sampling to obtain a view of the overall health of the network. One popular method of network traffic sampling involves the installation of specialized packet sampling software on switches used by the network to deliver packets. This software samples network packets passing through the switches en route to their destinations and transmits a portion of each sampled network packet to a monitoring appliance. However, this method of sampling network traffic has its drawbacks. For example, the network switches used by the network must be capable of supporting the packet sampling software to employ sampling. Additionally, sampling packets

at the switches provides no visibility into encrypted packets or traffic exchanged between virtual machines implemented by the same virtual host.

BRIEF DESCRIPTION OF THE DRAWINGS

[0003] The accompanying drawings illustrate various embodiments of the principles described herein and are a part of the specification. The illustrated embodiments are merely examples and do not limit the scope of the claims.

[0004] Fig. 1 is a block diagram of an illustrative network device, according to one example of principles described herein.

[0005] Figs. 2A, 2B, and 2C are block diagrams of network traffic sampling in an illustrative network, according to various examples of principles described herein.

[0006] Fig. 3 is a block diagram of an illustrative network system, according to one example of principles described herein.

[0007] Figs. 4A, 4B, and 4C are diagrams of illustrative sample reporting packets derived from sampled packets, according to one example of principles described herein.

[0008] Fig. 5 is a flowchart diagram of an illustrative method of sampling network traffic, according to one example of principles described herein.

[0009] Fig. 6 is a flowchart diagram of an illustrative analyzing network traffic sampled by loadable kernel modules in multiple network devices, according to one example of principles described herein.

[0010] Figs. 7A and 7B are flowchart diagrams illustrative methods of analyzing network traffic sampled by loadable kernel modules in multiple network devices, according to examples of principles described herein.

[0011] Fig. 8 is a flowchart diagram of an illustrative method of analyzing network traffic sampled by loadable kernel modules in multiple network devices, according to one example of principles described herein.

[0012] Throughout the drawings, identical reference numbers designate similar, but not necessarily identical, elements.

DETAILED DESCRIPTION

[0013] The present specification describes methods, systems, and computer program products which use loadable modules in the source and/or destination of network packets to accomplish network traffic sampling without the need for sampling support from network switches. By conducting network traffic sampling within the kernel of a network device which sends and receives packets on the network, a network administrator can gain visibility into encrypted traffic and traffic within virtualized environments that would not otherwise be visible through switch-based sampling.

[0014] In particular, the present specification describes a method of sampling network traffic in an operating system kernel that includes: loading a packet sampling module into a processor-based network device coupled to a network; determining with the packet sampling module if a network packet addressed to or from the network device is selected for sampling; and transmitting data from the network packet over the network to a monitoring device external to the network device if the network packet is selected for sampling.

[0015] Additionally, the present specification describes a method of sampling network traffic that includes: selecting a number of processor-based devices in a network for packet sampling; loading a packet sampling module into an operating system kernel for each selected network device; receiving data contained in sampled network packets from the packet sampling modules over the network; and compiling the data to determine a health of the network.

[0016] The present specification also describes a network device which includes a processor communicatively coupled to a memory. The processor executes operating system kernel code stored on the memory, which causes the processor to: determine in the operating system kernel if a network packet addressed to or from the network device is selected for sampling; and

transmit data from the network packet over a network to a monitoring device external to the network device if the network packet is selected for sampling

[0017] As used in the present specification and in the appended claims, the word "packet" means a block of data formatted for transmission to an addressable entity over a network.

[0018] As used in the present specification and in the appended claims, the word "kernel" means a central component of an operating system which controls access to hardware resources associated with a processor executing the operating system.

[0019] As used in the present specification and in the appended claims, the word "external," when describing a computer-implemented machine or device, refers to a machine or device that is implemented by a physically distinct processor. For example, a security device that is external to a virtualized host is implemented by a processor that is physically distinct from the processor(s) used to implement the virtualized host.

[0020] As used in the present specification and in the appended claims, the word "processor" refers to a hardware apparatus capable of executing code. A processor may include multiple central processing units.

[0021] In the following description, for purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present systems and methods. It will be apparent, however, to one skilled in the art that the present apparatus, systems and methods may be practiced without these specific details. Reference in the specification to "an example" or similar language means that a particular feature, structure, or characteristic described in connection with the example is included in at least that one example, but not necessarily in other examples. The various instances of the phrase "in one example" or similar phrases in various places in the specification are not necessarily all referring to the same example.

[0022] Referring now to the Figures, Fig. 1 shows a block diagram of an illustrative network device (100) which may send and receive data over a network. The illustrative network device (100) may implement, for example, an

addressable device on a computer network, such as a server device or a client computer. The illustrative network device (100) includes a hardware platform (105) made up of at least one processor (110), computer memory (115), a network interface card (NIC) (120), and other hardware devices (125). A motherboard may interconnect some or all of the hardware platform devices. The other hardware devices (125) may include, but are not limited to, peripheral input/output devices, storage devices, and any other hardware devices that may be suitable for a particular application of the principles described in the present specification.

[0023] The processor (110) executes code stored by the main memory (115). In certain examples, the processor (110) may include at least one multi-core processor having multiple independent central processing units (CPUs), with each CPU having its own L1 cache and all CPUs sharing a common bus interface and L2 cache. Additionally or alternatively, the processor (110) may include at least one single-core processor.

[0024] The main memory (115) stores code which is executed by the processor (110) to implement an operating system kernel (130). The operating system kernel (130) initializes and manages the devices of the hardware platform (105), and serves as a bridge between the hardware platform (105) and higher-level applications (135).

[0025] As shown, the operating system kernel (130) may include modules for CPU management (140), memory management (145), network communications management (150), and other device management (160). The operating system kernel (130) may also be extensible through the use of one or more loadable kernel modules. A loadable kernel module is an object file that contains code to extend the functionality of the base operating system kernel (130). Functionality may be added to the operating system kernel (130) by selectively activating a loadable kernel module implementing the desired functionality to be added. Similarly, functionality may be removed from the operating system kernel (130) by selectively deactivating or removing a loadable kernel module from the operating system kernel (130).

[0026] The loadable packet sampling module (155) in the operating system kernel (130) of Fig. 1 is one such loadable kernel module. The loadable packet sampling module (155) causes packet monitoring and sampling operations to be performed from within the operating system kernel (130) of a host device or client device on the network. The use of a loadable kernel module in host and client devices on the network to perform packet monitoring and sampling provides a number of benefits over traditional approaches which employ packet sampling in network switches.

[0027] One of the benefits associated with the use of a loadable packet sampling kernel module (155) in addressed network devices (100) is the fact that the loadable packet sampling kernel module (155) can be customized to the network device (100). For example, it may be desirable to sample more packets from the network traffic through a first network device and fewer packets from the network traffic passing through a second network device. In this case, the loadable packet sampling kernel module (155) for the first network device can be customized to perform select more network packets for sampling while the loadable packet sampling kernel module (155) for the second network device can be customized to select fewer packets for sampling. This added degree of flexibility in monitoring and sampling network traffic may allow for the most efficient and beneficial use of processing resources in an external network monitoring appliance.

[0028] Another benefit associated with the use of a loadable packet sampling kernel module (155) in network devices (100) is the ability to gain visibility into traffic passing between virtual machines in a virtualized environment. For example, in traditional systems where packet sampling occurs at network switches, it may be difficult to sample packets transmitted between two virtual machines hosted by the same host device, as this traffic may never pass through a physical network switch. By contrast, in the present system packet sampling occurs within the operating system kernel (130) of the host device itself, thereby enabling the examination and sampling of network traffic between the virtual machines.

[0029] Yet another benefit associated with the use of a loadable packet sampling kernel module (155) in network devices (100) is the ability to selectively activate and deactivate the packet monitoring and sampling functionality in real-time without interrupting the flow of network traffic. Loadable kernel modules may be loaded to and removed from the operating system kernel (130) while the operating system kernel (130) is running and without disrupting system operations. In this way, network traffic monitoring and sampling may be selectively activated or deactivated for each network device (100) for which a loadable packet sampling kernel module is available. Network traffic monitoring and sampling may be dynamically switched in on or off in one or more machines to conserve processing resources in specific devices and/or to focus network monitoring and sampling operations on one or more specific devices. Alternatively, the sampling function in a loadable packet sampling kernel module may be selectively disabled or enabled in real time without removing the kernel module from the kernel.

[0030] Still another benefit associated with the use of a loadable packet sampling kernel module (155) in a network device (100) is that of security. The operating system kernel (130) is typically very secure and less likely to fall prey to attacks from foreign applications or processes. Thus, it is less likely that an external process or malicious user without root access would be able to compromise packet monitoring and sampling operations in the network device (100).

[0031] Many of the same benefits described above with respect to a loadable packet sampling kernel module may also be achieved using a userspace sampling dameon, driver, or other machine-readable instructions that run within an application server and/or above a microkernel. While for the sake of clarity the present specification primarily describes examples using loadable packet sampling kernel modules to sample packets in a device that is the originator or final recipient of network packets, it should be understood that many of these principles may also be applied to userspace sampling daemons, drivers, or other machine-readable instructions running within an application server and/or above a microkernel.

[0032] Figs. 2A-2C show block diagrams of network traffic sampling in network devices using loadable packet sampling modules according various examples of the principles of the present specification. In each of Figs. 2A-2C, a network device (201, 202, 203, respectively) includes one or more network interface controllers (NICs) (205) and an operating system kernel (210). Other elements of the network device (201, 202, 203), including the hardware platform and various elements of the operating system and operating system kernel (210) are omitted in Figs. 2A-2C for clarity. The operating system kernel (210) for each network device (201, 202, 203) includes a loadable packet sampling module (215), consistent with the explanation given above with reference to Fig. 1.

[0033] The loadable packet sampling module (215) includes a selection submodule (220) and a datagram construction and forwarding submodule (225). Of course, while the functionality of the loadable packet sampling module (215) is shown in these figures using two submodules (220, 225), this same functionality may be divided up into more or fewer submodules as may suit a particular application of the principles described herein.

[0034] The selection submodule (220) monitors network packets (PKT) passing through the network device (201, 202, 203) and determines whether each packet is selected for sampling. These network packets may be packets sent by the network device (202) over a network (235), received by the network device (202) from the network (235), and/or packets to or from virtual machines (230-1, 230-2, Fig. 2C) hosted by the network device (201, 202, 203). The selection process may be based on a formula used to sample an average of every *n* packets passing through the network device (201, 202, 203). If the packet is selected for sampling, the datagram construction and forwarding submodule (225) creates a datagram containing data from the sampled packet and sampling statistics for the network device (201, 202, 203), and forwards the datagram in a packet over the network (235) to an external monitoring device (240).

[0035] The external monitoring device (240) of the present example is a processor-based network apparatus that includes a datagram inspection

module (245), a traffic statistics and reporting module (250), and an anomaly detection and remediation module (255). For each packet received by the external monitoring device (240) from a loadable packet sampling module (215), the datagram inspection module (245) retrieves the sampled packet data and sampling statistics from the application-layer datagram.

[0036] The traffic statistics and reporting module (250) updates compiled traffic statistics for the network (235) and makes the statistics available to a network administrator or other authorized entity. The traffic statistics and reporting module (250) may report the statistics as raw data and/or in a summarized form. Additionally or alternatively, the traffic statistics and reporting module (250) may make conclusions regarding the health of the network (235) from the compiled traffic statistics and provide an indication of network health based on the compiled traffic statistics.

[0037] The anomaly detection and remediation module (255) may examine the compiled traffic statistics and/or data from individual sampled packets to detect anomalies. Examples of such anomalies include network security issues or events (e.g., software vulnerability exploitations, malware, resource attacks, traffic to or from prohibited entities), overly burdened network devices, network errors, unusual or unexpected network traffic characteristics, and the like.

[0038] In certain examples, the anomaly detection and remediation module (255) may take action to inform a network administrator or other entity of the detected anomaly through an appropriate medium (e.g., alarm, email, textual message, etc.). Additionally or alternatively, the anomaly detection and remediation module (255) may take automatic action to directly remediate or alleviate the anomaly. For example, if the monitoring device (240) determines from the compiled traffic statistics that a certain network device (201, 202, 203) is overburdened with traffic, the monitoring device (240) may take steps to divert some of the network traffic from the overburdened network device (201, 202, 203) to an underutilized network device (201, 202, 203). In another example, if traffic from a prohibited entity is detected on the network (235), the

network monitoring device (240) may adjust routing tables in network routers to foreclose the prohibited traffic.

[0039] In alternate examples, some or all of the functionality of the network monitoring device (240) may be performed within the network device (201) itself. The network device (201) may have processing resources which are allocable to the inspection and analysis of packets selected by the loadable packet sampling module (215). Thus, in some examples the network device (201) may inspect the sampled packets, gather traffic statistics for the network device, and detect and remediate anomalies from the traffic statistics without the aid of an external monitoring device (240). Alternately, the network device (201) may perform some inspection and analysis of packets sent from and received by the network device and forward only some of the sampled packets to the external monitoring device (240) for use in gathering traffic statistics and detecting anomalies in the network as a whole.

[0040] Figs. 2A, 2B, and 2C illustrate the functionality of the loadable packet sampling module (215) and the monitoring device (240) in different contexts. In Fig. 2A, the loadable packet sampling module (215) samples packets as the packets pass through the network device (201) between the network (235) and one or more applications (260) executed by the network device (201).

[0041] In Fig. 2B, the loadable packet sampling module (215) works in conjunction with a packet sampling module (265) in an application (270) implemented by the network device (202) in order to sample and report application-level data. This approach may prove particularly useful in the monitoring of encrypted application-level data. Under one potential scenario, the selection module (220) of the loadable packet sampling module (215) may select an encrypted packet addressed to the application (270) for sampling and indicate the selection to the packet sampling module (265) of the application (270). When the application (270) receives and decrypts the data from the selected packet, the packet sampling module (265) of the application (270) may provide at least a portion of the decrypted data from the selected packet to the

loadable packet sampling module (215), which includes the decrypted data in the datagram sent to the monitoring device (240) for the selected packet.

[0042] Thus, where prior approaches to network traffic sampling are unable to effectively sample encrypted application-level data from network packets, the present system provides an efficient solution to sampling this type of high-level data. This ability can prove invaluable to network security, as the monitoring device (240) may be able to detect and remediate against malware or other problematic data transmitted to a network device (202). Additionally, the monitoring device (240) may be able to compile a more accurate and complete view of network health (235) and traffic trends by including a view of application-level data in its analysis.

[0043] In Fig. 2C, the network device (203) is a virtualized host which executes a hypervisor (275) to implementing multiple virtual machines (230-1, 230-2). These virtual machines (230-1, 230-2) may transmit data to each other using a virtualized network switch implemented by the hypervisor (275). The loadable packet sampling module (215) may communicate with the hypervisor (275) to select certain packets transmitted between the virtual machines (230-1, 230-2) for sampling. In this way, even data from traffic between virtual machines that would ordinarily never go beyond the hypervisor (275) may be transmitted to the monitoring device (240) for inspection and reporting.

[0044] Fig. 3 is a block diagram of an illustrative system (300) including multiple network devices (301-1 to 301-3) having respective loadable packet sampling modules (305-1 to 305-3) in their respective operating system kernels (310-1 to 310-3). Each of the network devices (301-1 to 301-3) is communicatively coupled to a network (315). Each of the loadable packet sampling modules (305-1 to 305-3) samples packets passing through its respective network device (301-1 to 301-3) and transmits data from sampled packets together with sampling statistics to a monitoring device (320) over the network (315), consistent with the details described previously.

[0045] Because the monitoring device (320) can receive sampled network traffic data from each of the network devices (301-1 to 301-3) on the network (315), the monitoring device (320) may compile network traffic statistics

for the entire system (300). However, another feature of this example is the fact that the loadable packet sampling modules (305-1 to 305-3) may be selectively loaded to or removed from their respective operating system kernels (310-1 to 310-2, 310-3) in real-time without rebooting their respective network devices (301-1 to 301-3). In alternative examples, the loadable packet sampling modules (305-1 to 305-3) may continuously run in the kernel and the sampling functionality of the kernels may be selectively enabled or disabled.

[0046] Thus, if one or more network devices (301-1 to 301-3) become overburdened, the packet sampling module (305-1 to 305-3) for that network device (301-1 to 301-3) may be removed to free up computing resources. Additionally, if the monitoring device (320) or a network administrator (325) elects to sample network traffic from only a subset of network devices (301-1 to 301-3), the packet sampling modules (305-1 to 305-3) for those network devices (301-1 to 301-3) not in the subset may be removed.

[0047] Conversely, if the monitoring device (320) of the network administrator (325) chooses to begin or resume sampling traffic from a particular network device (301-1 to 301-3), the packet sampling module (305-1 to 305-3) for that network device (301-1 to 301-3) may be loaded and reactivated.

[0048] The selective loading or removing of the packet sampling modules (305-1 to 305-3) in network devices (301-1 to 301-3) may in some examples occur by way of a command from the monitoring device (320) or another administrative device connected to the network (315). This command may occur as a result of dynamic decisions automatically made by the monitoring device (320) or another administrative device to enforce network policy.

[0049] Additionally or alternatively, the administrator (325) may manually load and remove the packet sampling modules (305-1 to 305-3) in the network devices (301-1 to 301-3) directly, using the monitoring device (320), or by taking other administrative action that may better suit a specific application of the principles described herein.

[0050] Additionally or alternatively, the network devices (301-1 to 301-3) themselves may be configured to automatically load and remove the packet sampling modules (305-1 to 305-3) from their respective operating system kernels (310-1 to 310-3) based on detected events, conditions or triggers. For example, if a network device (301-1) detects a utilization of processor resources beyond a predefined threshold, the network device (301-1) may automatically remove the packet sampling module (305-1) to free up processing resources. Conversely, if the network device (301-1) detects that resource utilization drops below a certain threshold, the network device (301-1) may automatically reload the packet sampling module (305-1) into its operating system kernel (310-1).

[0051] In the same way that the loadable packet sampling modules (305-1 to 305-3) may be selectively loaded and removed from their respective kernels (310-1 to 310-3), the sampling parameters of the packet sampling modules (305-1 to 305-3) may be dynamically updated as may suit a particular situation or network policy. For example, it may be desirable to sample more packets from a subset of the network devices (301-1 to 301-3) and fewer packets from the remaining network devices (301-1 to 301-3). In this case, the sampling parameters of the packet sampling modules (305-1 to 305-3) in the selected network devices (301-1 to 301-3) may be automatically updated by an administrative device on the network, by the devices (301-1 to 301-3), or manually by an administrator (325) to increase the number of packets selected for sampling. Likewise, the sampling parameters of the network devices (301-1 to 301-3) not in the selected subset may remain the same or be updated to decrease the number of packets selected for sampling.

[0052] Similarly, in some examples the sampling functionality of the loadable packet sampling modules (305-1 to 305-3) may be selectively disabled without removing the loadable packet sampling modules (305-1 to 305-3) from their respective kernels (310-1 to 310-3). In this way, packet sampling effectuated by the operating system kernels (310-1 to 310-3) may be turned on and off through a simple application programming interface (API) call to the kernel without expending the processing resources to load and remove the modules (305-1 to 305-3) whenever sampling functionality is desired.

[0053] Figs. 4A-4C show various examples of the composition of sampling packets sent to a monitoring device from a loadable packet sampling kernel module according to the principles described above. In each of Figs. 4A-4C, a network-layer Internet Protocol (IP) packet (405) is sampled by the packet sampling kernel module, and an IP packet (410, 415, 420) is sent to the monitoring device with sampling parameters and data from the sampled IP packet (405). Each IP packet (405, 410, 415, 420) includes an IP Packet Header for delivery to an IP address, a User Datagram Protocol (UDP) header with application-layer delivery information, and a UDP datagram containing the application-layer payload data. It will be understood that while the examples of Figs. 4A-4C show IP type packets, any type of packet may be sampled according to the principles described herein. Examples of packets that may be sampled using the principles described herein include, but are not limited to Transmission Control Protocol (TCP) packets, Internet Control Message Protocol (ICMP) packets, Address Resolution Protocol (ARP) packets, and the like.

[0054] In the example of Fig. 4A, the UDP datagram of the IP packet (410) sent to the monitoring device includes sampling parameters from the loadable packet sampling kernel module and the IP Packet Header from the sampled IP packet (405). In the example of Fig. 4B, the UDP datagram of the IP packet (415) sent to the monitoring device includes sampling parameters and the UDP header from the sampled IP packet (405). In the example of Fig. 4C, the UPD datagram of the IP packet (420) sent to the monitoring device includes sampling parameters and a specified number (*n*) of bytes from the sampled IP packet (405). Of course, any other arrangement of data may be used in the packet sent to the monitoring device. For example, some or all of the packets sent to the monitoring device may omit the sampling parameters. Additionally or alternatively, the UDP datagram of the IP packet sent to the monitoring device may include the entire sampled IP packet (405) or the entire UDP datagram of the sampled IP packet (405). Any suitable data arrangement may be used to report data from a sampled packet and/or sampling parameters to a

monitoring device, as may best suit a particular application of the principles described herein.

[0055] Fig. 5 is a flowchart diagram of an illustrative method (500) of sampling network traffic. According to the method (500), a packet sampling module is loaded (block 505) into a kernel of an operating system executed by a processor-based network device. In certain examples, the packet sampling module may be loaded into the kernel of the network device while the kernel is running. A determination is then made (block 510) by the packet sampling module as to whether a packet received or transmitted by the network device is selected for sampling. This determination may be made using, for example, a pseudo-random formula which results in an average selection of a certain percentage or ratio of the total packets transmitted through the network device.

[0056] If the packet is selected for sampling (block 510, YES), data from the selected network packet is transmitted (block 515) over a network to a monitoring device external to the network device. The packet is directed (block 525) or delivered to its intended destination.

[0057] In certain examples, the method (500) may further include loading a second packet sampling module in an application executed by the network device. In these examples, the second packet sampling module may cooperate with the kernel packet sampling module to sample application-layer data (e.g., encrypted application data) as described above.

[0058] Fig. 6 is a flowchart diagram of an illustrative method (600) of sampling network traffic which may be performed by an external monitoring device. In this method (600), the external monitoring device receives (block 605) from loadable kernel module in a network device a datagram including at least a portion of a sampled packet. The external monitoring device uses the data from the datagram to update (block 610) statistics for the network, compiles the statistics for the network (block 615), and reports (block 620) the compiled statistics to an administrator of the network.

[0059] Figs. 7A and 7B are flowchart diagrams of related illustrative methods (700, 750) of analyzing network traffic sampled by loadable kernel modules in multiple network devices. In each of the methods, the external

monitoring device receives (block 705) from loadable kernel module in a network device a datagram including at least a portion of a sampled packet. The external monitoring device uses the data from the datagram to update (block 710) statistics for the network, compiles the statistics for the network (block 715), and then makes a determination (block 720) from the statistics for the network and/or the data from the individual datagram whether an anomaly exists in the network. In the method (700) of Fig. 7A, the external monitoring device provides (block 725) an indication of any anomaly detected to an administrator. An additional or alternative course of action is provided in the method (750) of Fig. 7B, in which the external monitoring device takes action to automatically remediate (block 730) any detected anomaly.

[0060] Fig. 8 is a flowchart diagram of another illustrative method (800) of analyzing network traffic sampled by loadable kernel modules in multiple network devices, according to one example of the principles of the present specification. In the method (800) of Fig. 8, a number of processor-based addressable devices in a network is selected (block 805) for packet sampling. For each selected network device, a packet sampling kernel module is loaded (block 810) into the operating system kernel for that device. Data is then received (block 815) from the packet sampling modules over the network and the received data is compiled (block 820) to determine a health of the network. In certain examples, the method (800) may further include determining whether any network device not selected for packet sampling has a loaded sampling kernel module in its operating system kernel, and removing the sampling kernel module from the operating system kernel of any such network device. Additionally, the method may include detecting anomalies in the network from the received data, providing an indication of any detected anomaly to a network administrator, and/or automatically performing a remedial action to correct the anomaly, as described above.

[0061] The preceding description has been presented only to illustrate and describe examples of the principles described. This description is not intended to be exhaustive or to limit these principles to any precise form

disclosed. Many modifications and variations are possible in light of the above teaching.

CLAIMS

WHAT IS CLAIMED IS:

1. A method of sampling network traffic, comprising:

loading a packet sampling module (215) into a processor-based network device (201) coupled to a network (235);

determining with said packet sampling module (215) if a network packet addressed to or from said network device (201) is selected for sampling; and

transmitting data from said network packet over said network (235) to a monitoring device (240) external to said network device (201) if said network packet is selected for sampling.

2. The method according to claim 1, wherein said packet sampling module (215) is loaded into a kernel (210) of an operating system executed by said processor-based network device (201).

3. The method according to claim 2, further comprising loading said packet sampling module (215) into said kernel (210) while said kernel is running.

4. The method according to any of the above claims, wherein transmitting said data from said network packet to said monitoring device (240) comprises transmitting a datagram to said monitoring device (240), said datagram comprising said data from said network packet and sampling data corresponding to a sampling of said network packet.

5. The method of claim according to any of the above claims, further comprising loading a second packet sampling module (265) in an application (270) executed by said network device (202).

6. The method according to claim 5, wherein said data from said network packet comprises application-level data collected by said second packet sampling module (265) in said application (270).

7. The method according to any of the above claims, wherein said data from said network packet comprises data decrypted by said network device (202) from an encrypted portion of said network packet.

8. The method according to any of the above claims, wherein said data from said network packet comprises at least a portion of said network packet.

9. A method, of sampling network traffic, comprising:

selecting a number of processor-based devices (301-1 to 301-3) in a network (315) for packet sampling;

loading a packet sampling module (305-1 to 305-3) into an operating system kernel (310-1 to 310-3) for each selected network device (301-1 to 301-3);

receiving data contained in sampled network packets from said packet sampling modules (305-1 to 305-3) over said network (315); and

compiling said data to determine a health of said network (315).

10. The method according to claim 9, further comprising detecting an anomaly in said network (315) from said data.

11. The method according to any of claims 9 or 10, further comprising automatically performing a remedial action to correct said anomaly.

12. The method according to any of claims 9, 10, or 11, further comprising determining whether a said packet sampling kernel module (310-1 to 310-3) has been loaded into an operating system kernel (310-1 to 310-3) of a network device (301-1 to 301-3) not selected for packet sampling.

13. The method according to claim 12, further comprising removing said sampling kernel module (305-1 to 305-3) from said operating system kernel (310-1 to 310-3) of said network device (301-1 to 301-3) not selected for packet sampling.

14. A network device (100), comprising:

a processor (110) communicatively coupled to a memory (115), said processor (110) executing operating system kernel (130) code stored on said memory (115) which causes said processor (110) to:

determine in said operating system kernel (130) if a network packet addressed to or from the network device (100) is selected for sampling; and

transmit data from said network packet over a network (235) to a monitoring device (240) external to said network device (100) if said network packet is selected for sampling.

15. The network device according to claim 14, wherein said data from said network packet comprises data decrypted by said network device (100) from an encrypted portion of said network packet.

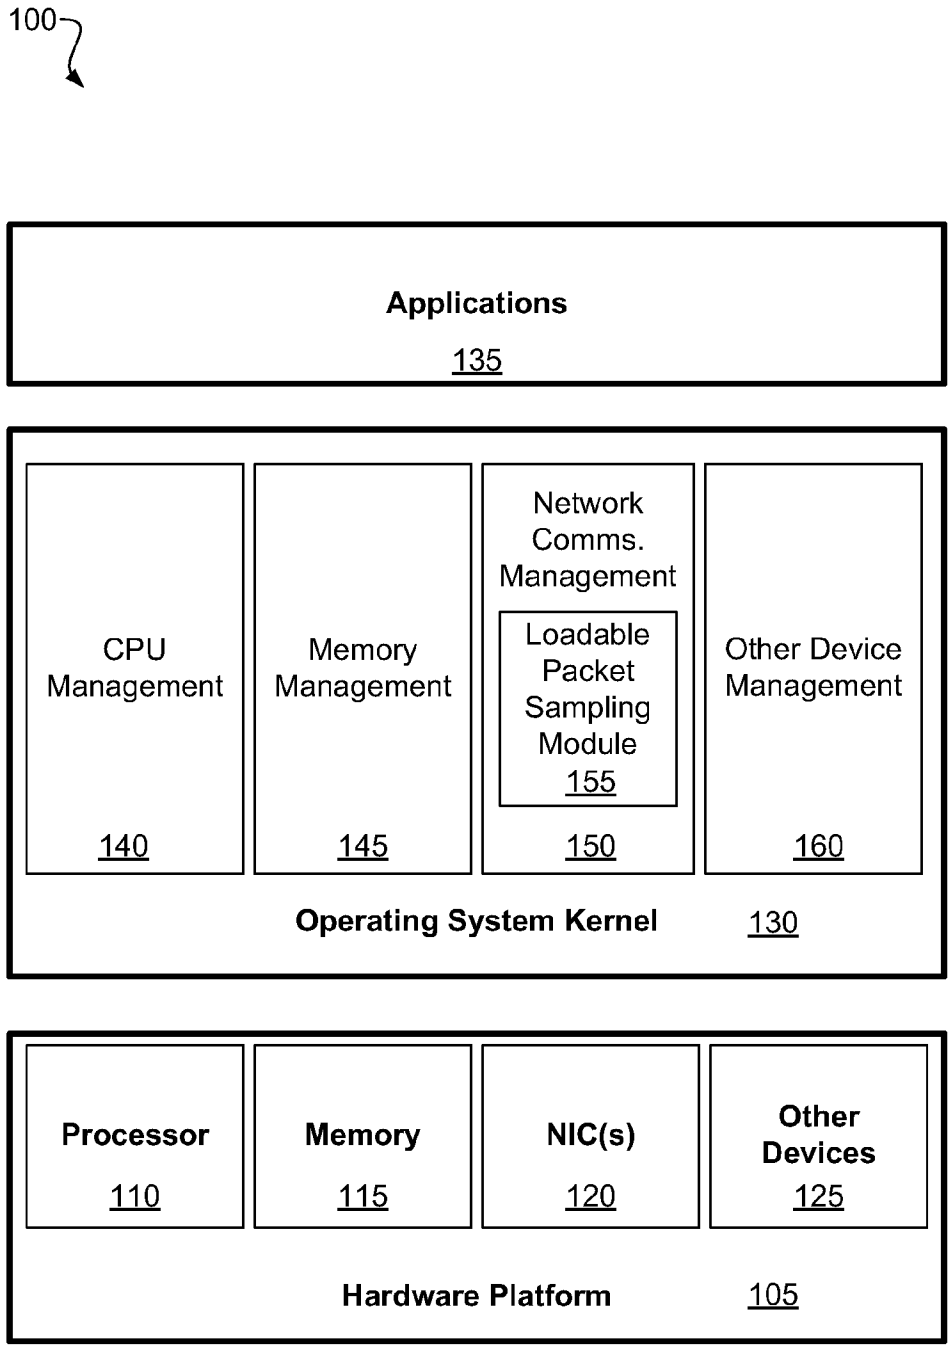
100
Applications
135
CPU Management
140
Memory Management
145
Network Comms. Management
Loadable Packet Sampling Module
155
150
Other Device Management
160
Operating System Kernel
130
Processor
110
Memory
115
NIC(s)
120
Other Devices
125
Hardware Platform
105

Fig. 1

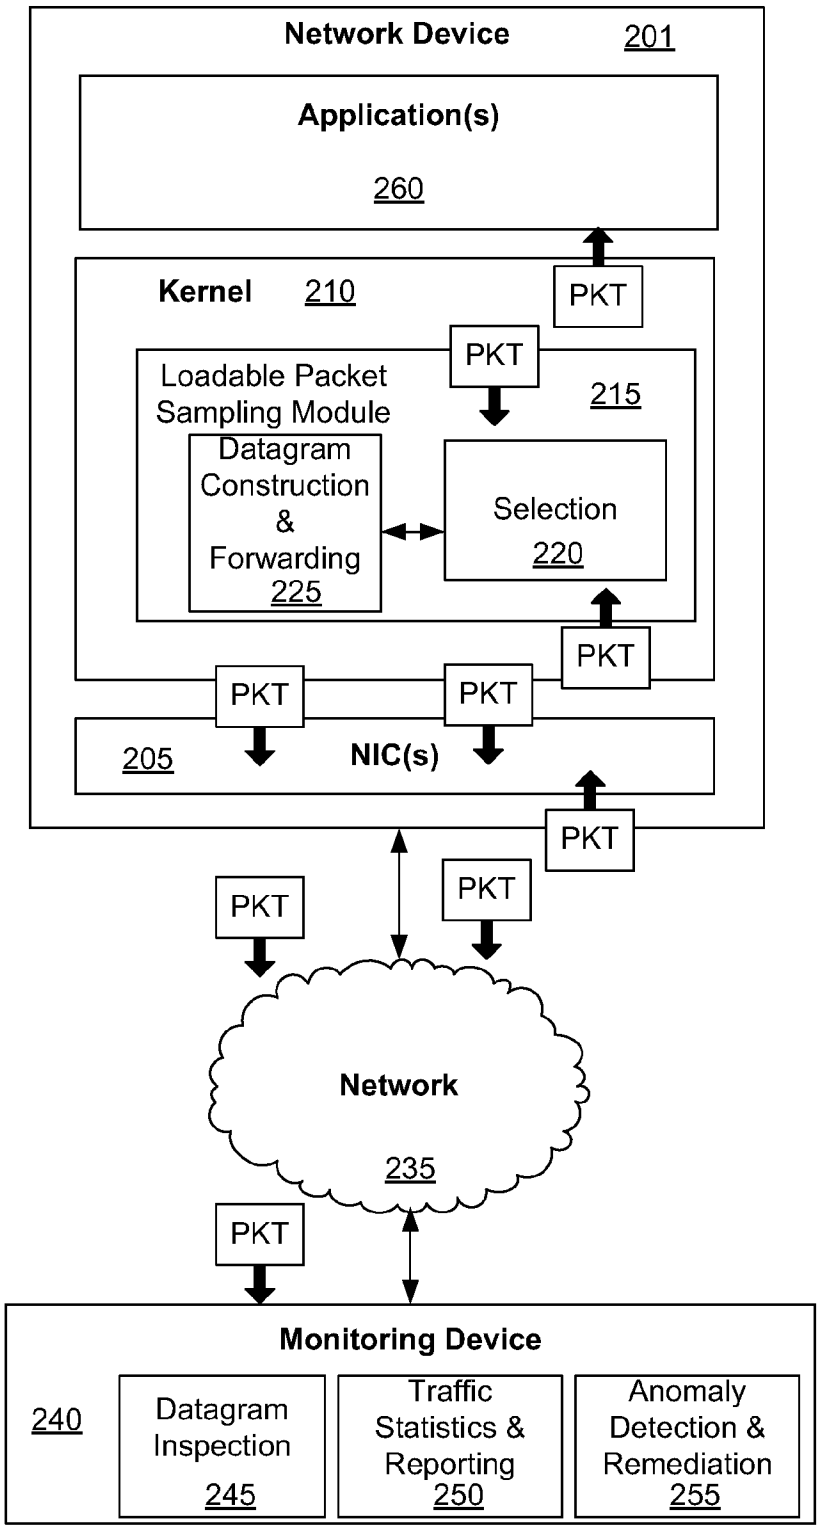
Network Device
201
Application(s)
260
Kernel
210
PKT
PKT
Loadable Packet Sampling Module
215
Datagram Construction & Forwarding
225
Selection
220
PKT
PKT
PKT
205
NIC(s)
PKT
PKT
PKT
Network
235
PKT
Monitoring Device
240
Datagram Inspection
245
Traffic Statistics & Reporting
250
Anomaly Detection & Remediation
255

Fig. 2A

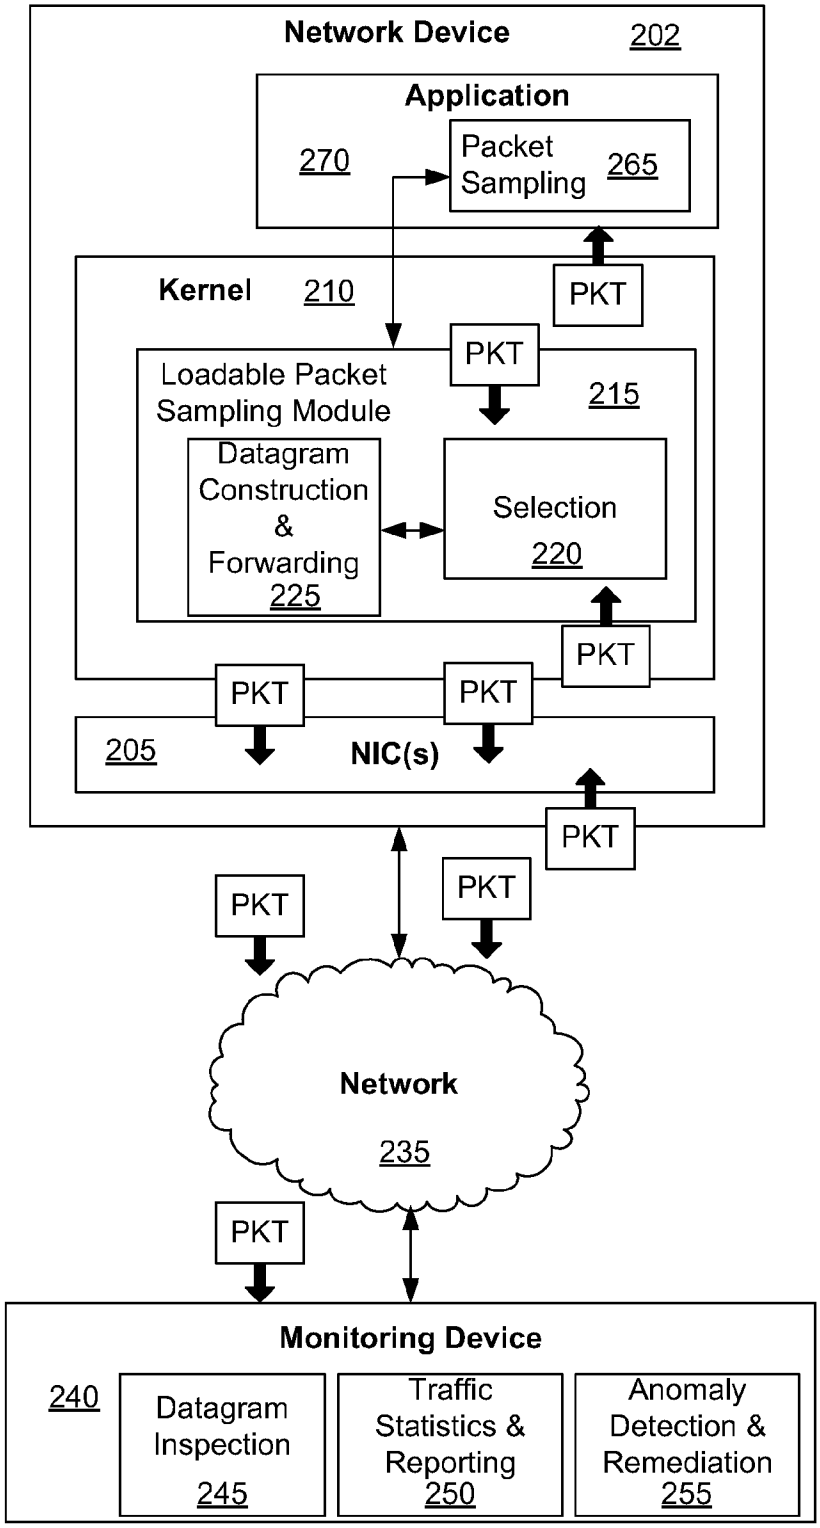
Network Device
202
Application
270
Packet Sampling
265
Kernel
210
PKT
PKT
Loadable Packet Sampling Module
215
Datagram Construction & Forwarding
225
Selection
220
PKT
PKT
PKT
205
NIC(s)
PKT
PKT
PKT
Network
235
PKT
Monitoring Device
240
Datagram Inspection
245
Traffic Statistics & Reporting
250
Anomaly Detection & Remediation
255

Fig. 2B

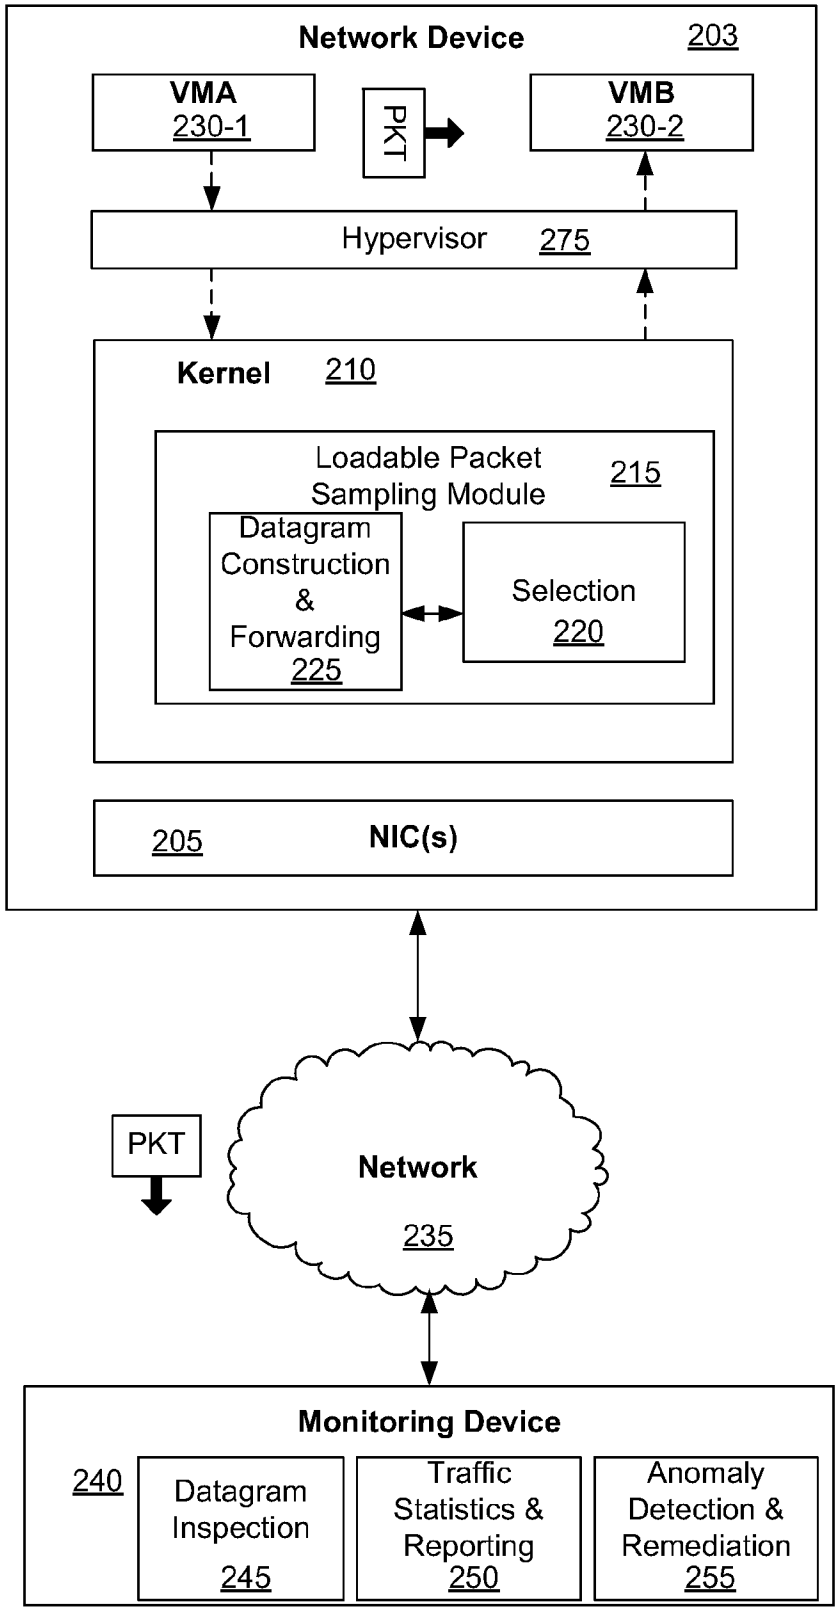
Network Device
203
VMA
230-1
PKT
VMB
230-2
Hypervisor
275
Kernel
210
Loadable Packet
Sampling Module
215
Datagram
Construction
&
Forwarding
225
Selection
220
205
NIC(s)
PKT
Network
235
Monitoring Device
240
Datagram
Inspection
245
Traffic
Statistics &
Reporting
250
Anomaly
Detection &
Remediation
255

Fig. 2C

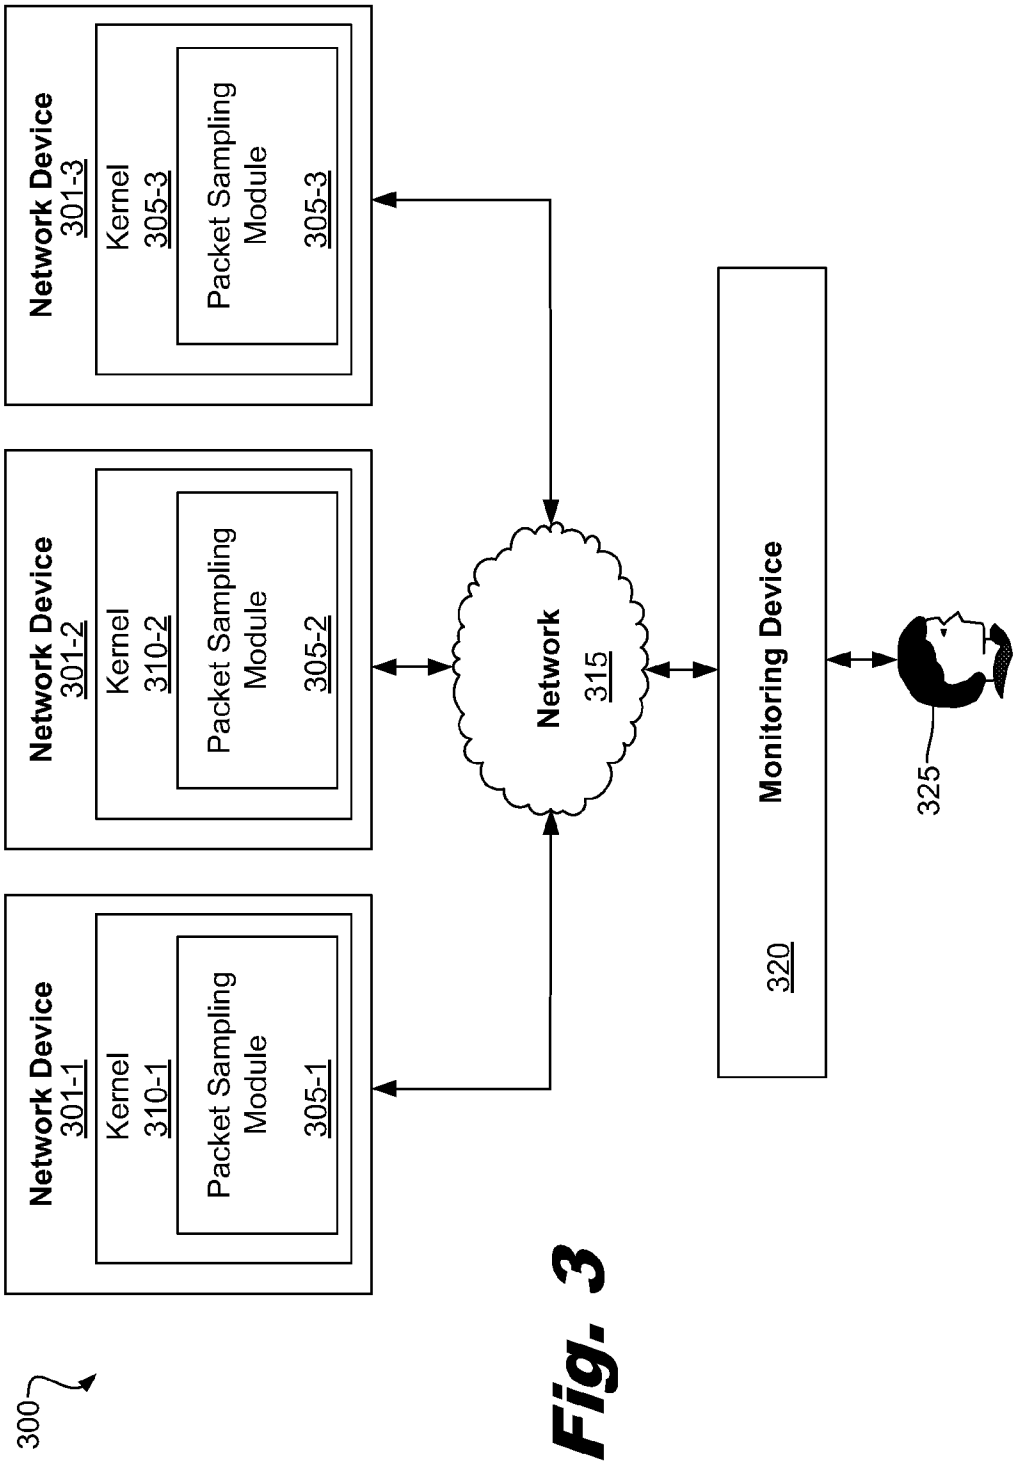
Network Device
301-1
Kernel
310-1
Packet Sampling
Module
305-1
Network Device
301-2
Kernel
310-2
Packet Sampling
Module
305-2
Network Device
301-3
Kernel
305-3
Packet Sampling
Module
305-3
Network
315
Monitoring Device
320
325
300

Fig. 3

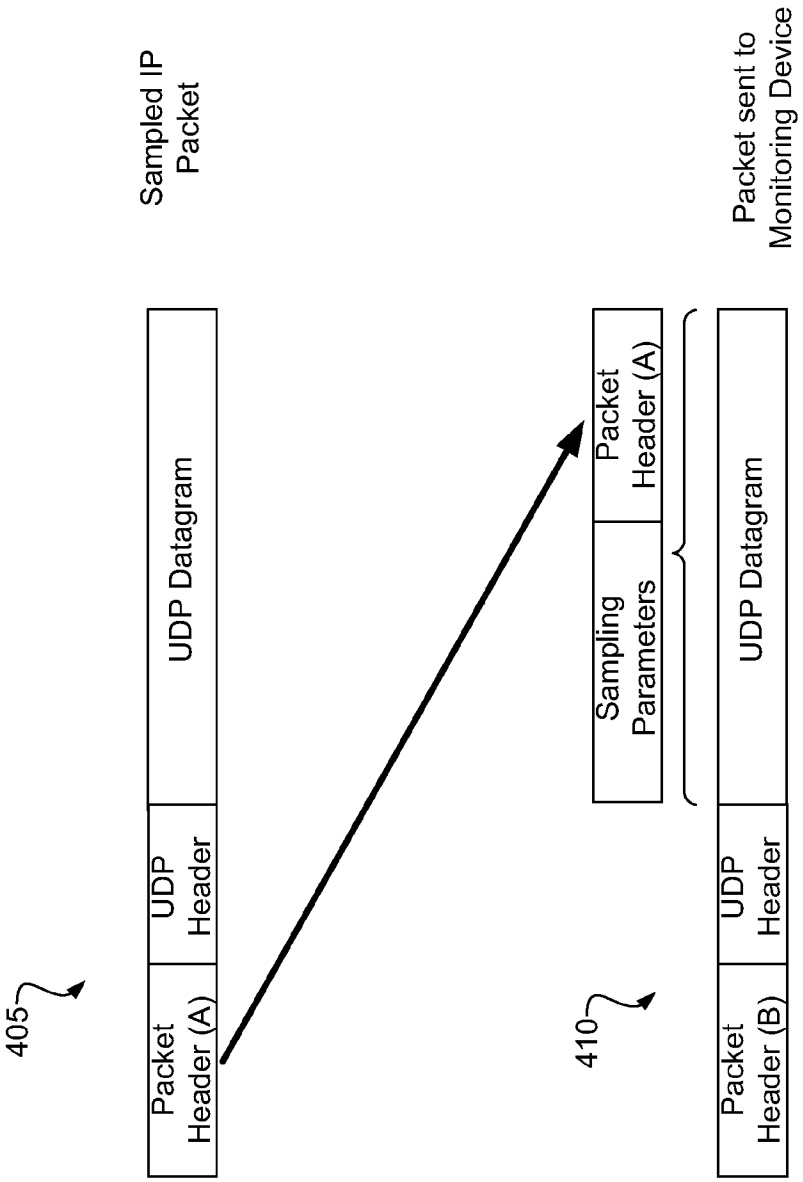
Sampled IP Packet
Packet sent to Monitoring Device
405
410
Packet Header (A)
UDP Header
UDP Datagram
Sampling Parameters
Packet Header (A)
Packet Header (B)
UDP Header
UDP Datagram

Fig. 4A

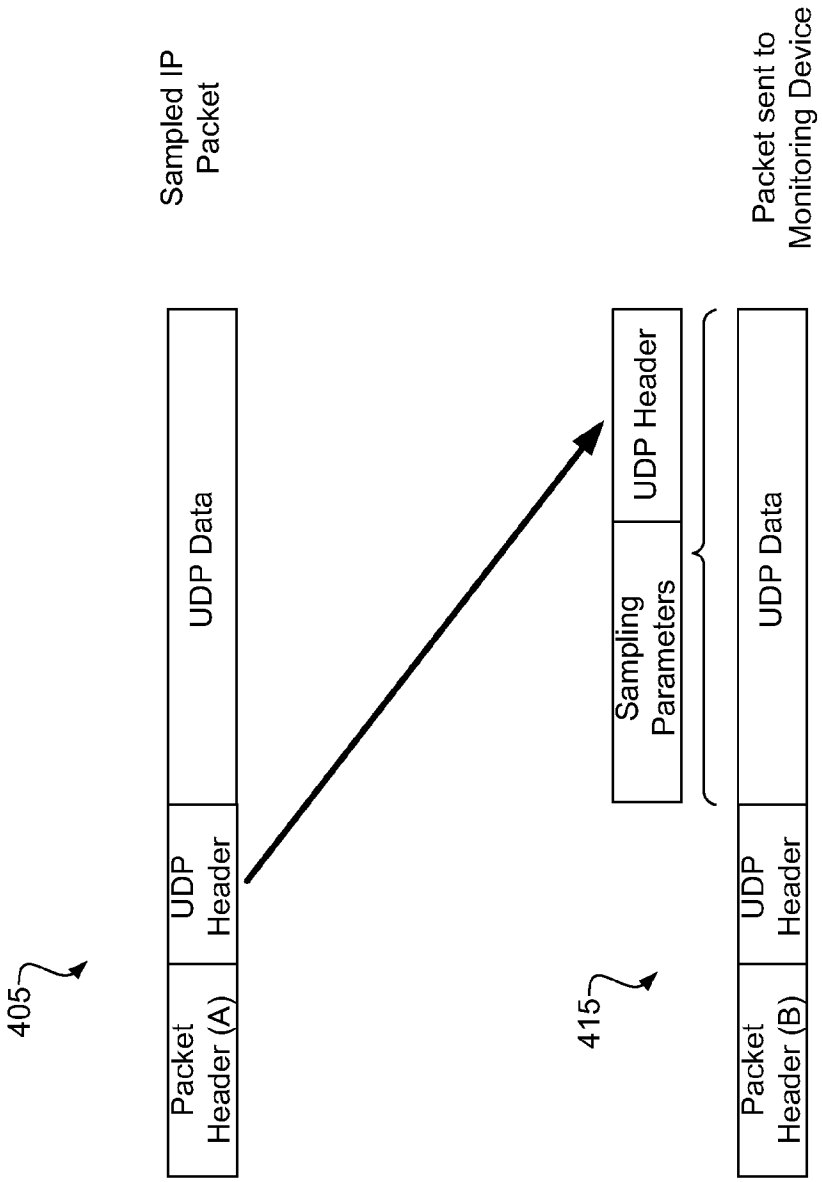
Sampled IP Packet
Packet sent to Monitoring Device
405
Packet Header (A)
UDP Header
UDP Data
415
Sampling Parameters
UDP Header
Packet Header (B)
UDP Header
UDP Data

Fig. 4B

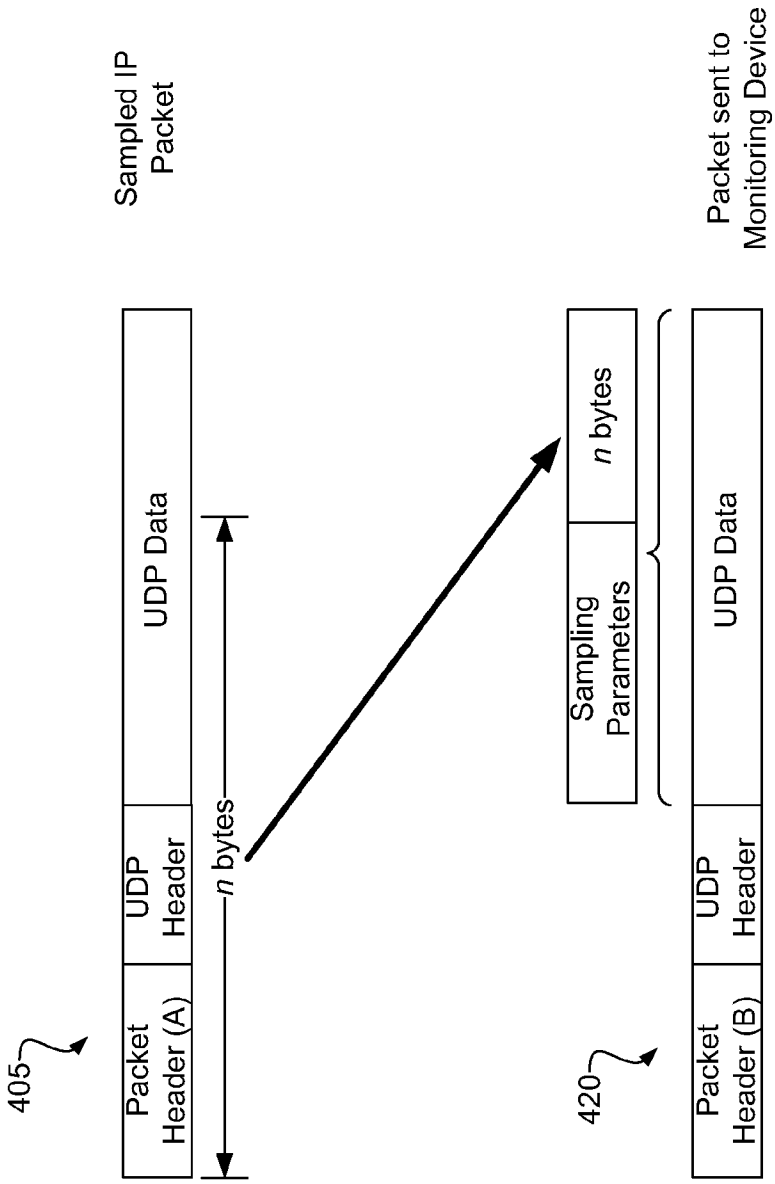
Sampled IP Packet
Packet sent to Monitoring Device
405
Packet Header (A)
UDP Header
UDP Data
n bytes
420
Sampling Parameters
n bytes
Packet Header (B)
UDP Header
UDP Data

Fig. 4C

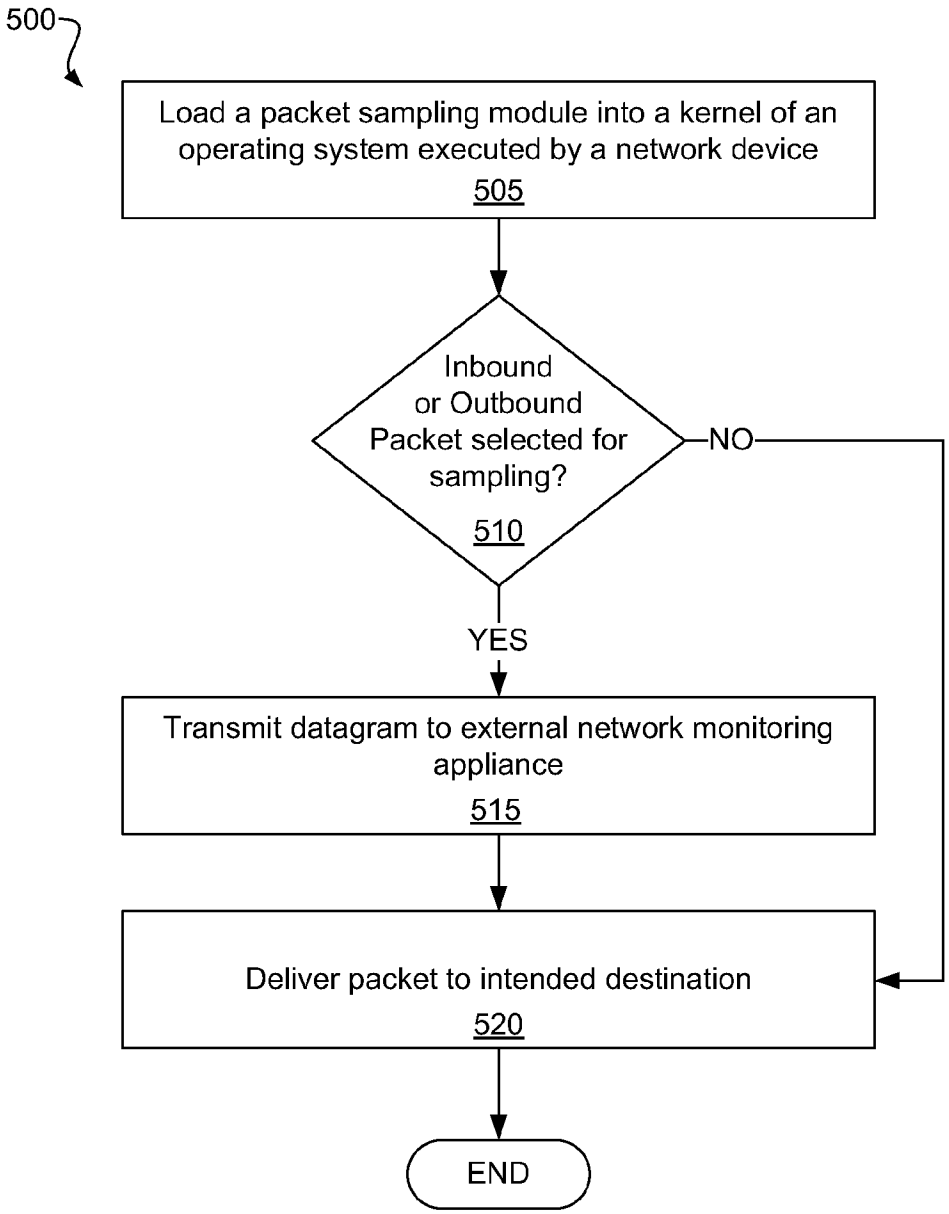
500
Load a packet sampling module into a kernel of an operating system executed by a network device
505
Inbound or Outbound Packet selected for sampling?
510
NO
YES
Transmit datagram to external network monitoring appliance
515
Deliver packet to intended destination
520
END

Fig. 5

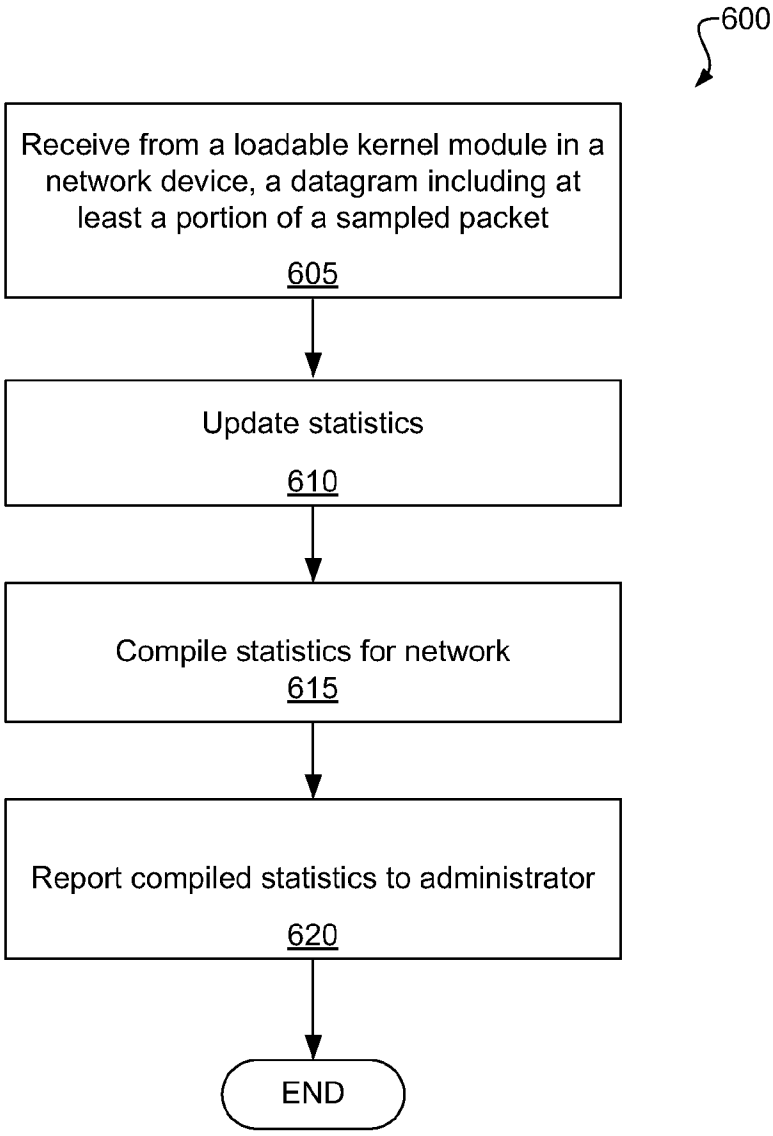
600
Receive from a loadable kernel module in a network device, a datagram including at least a portion of a sampled packet
605
Update statistics
610
Compile statistics for network
615
Report compiled statistics to administrator
620
END

Fig. 6

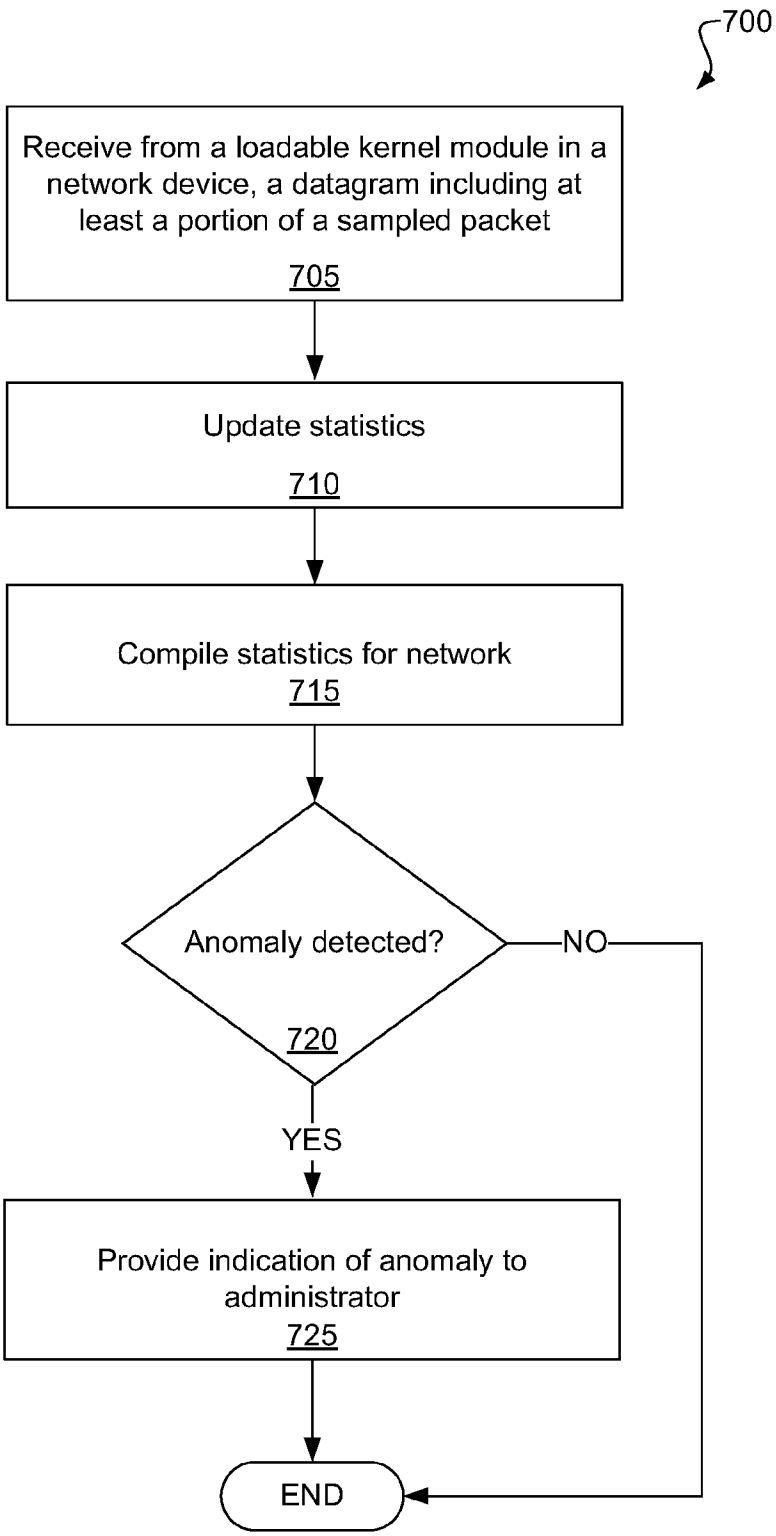
700
Receive from a loadable kernel module in a network device, a datagram including at least a portion of a sampled packet
705
Update statistics
710
Compile statistics for network
715
Anomaly detected?
720
NO
YES
Provide indication of anomaly to administrator
725
END

Fig. 7A

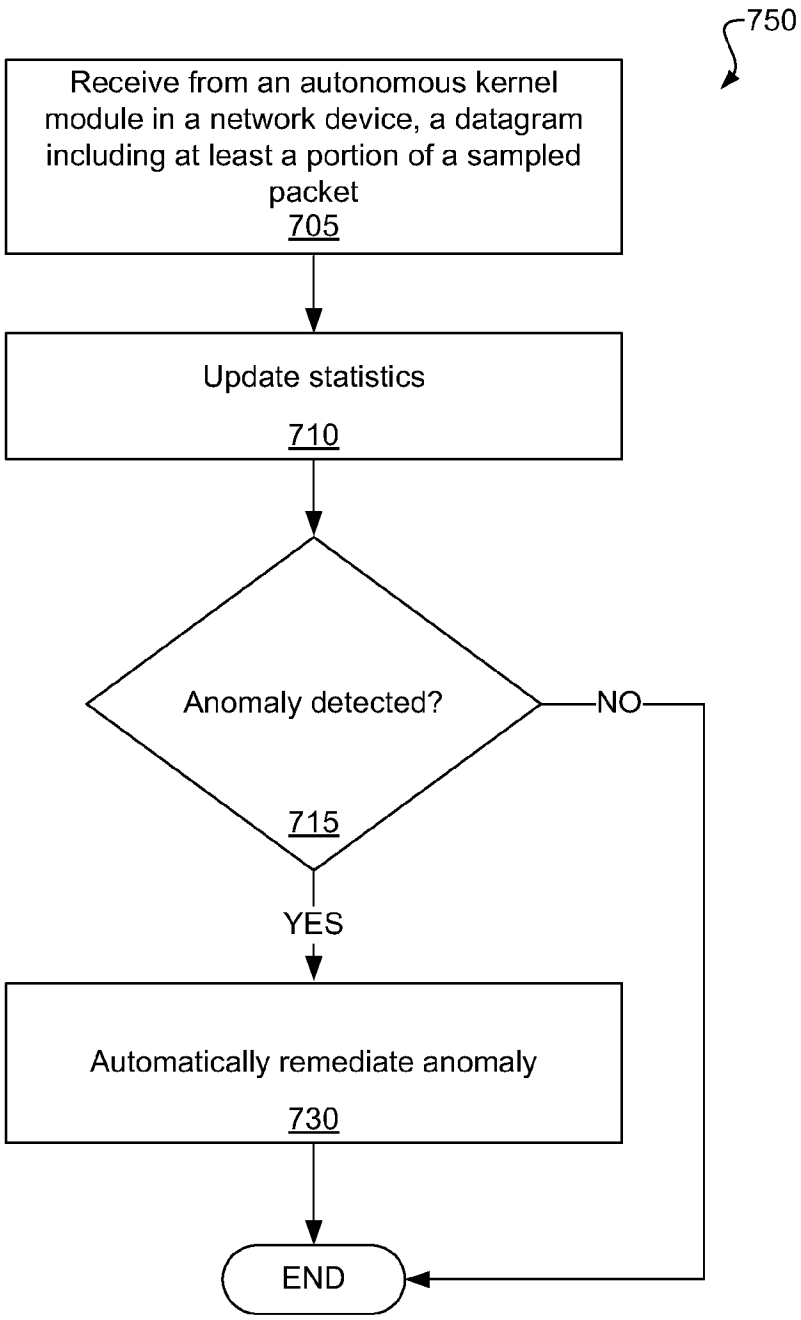
750
Receive from an autonomous kernel module in a network device, a datagram including at least a portion of a sampled packet
705
Update statistics
710
Anomaly detected?
715
NO
YES
Automatically remediate anomaly
730
END

Fig. 7B

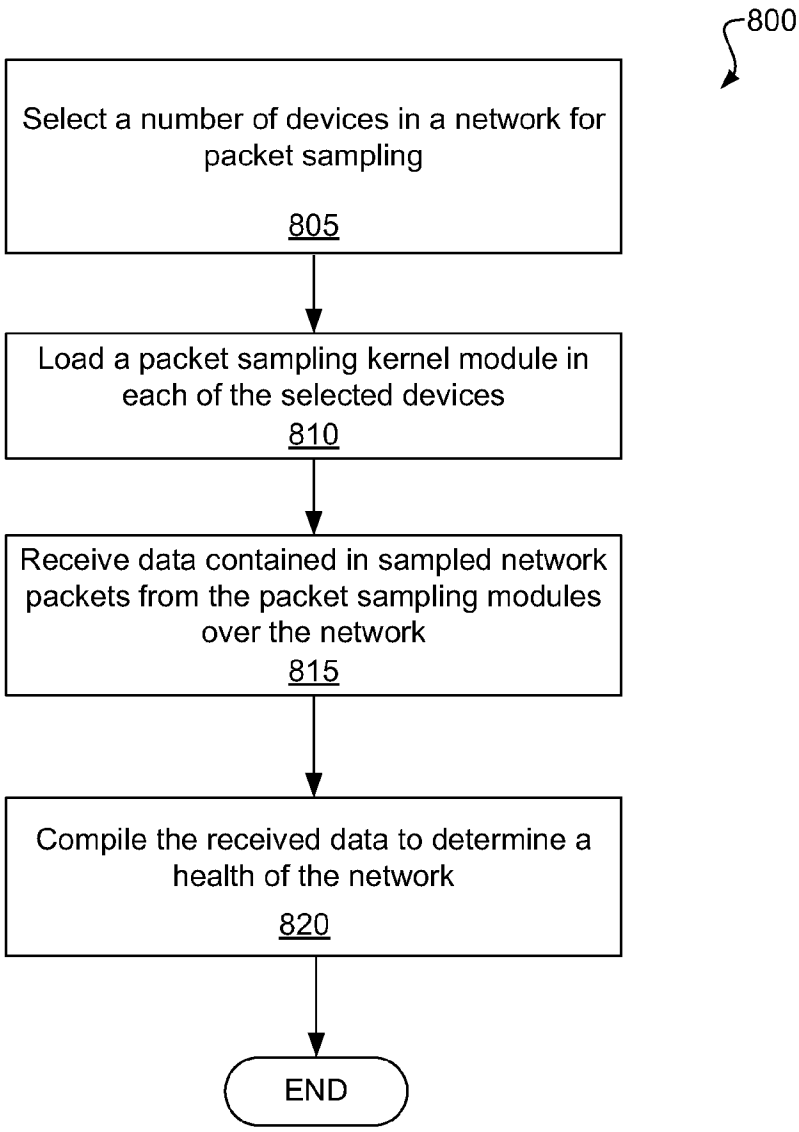
800
Select a number of devices in a network for packet sampling
805
Load a packet sampling kernel module in each of the selected devices
810
Receive data contained in sampled network packets from the packet sampling modules over the network
815
Compile the received data to determine a health of the network
820
END

Fig. 8

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2011/028043

A. CLASSIFICATION OF SUBJECT MATTER

H04L 12/26(2006.01)i, H04L 12/56(2006.01)i, H04L 29/02(2006.01)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 12/26; H04J 1/16

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: network, traffic, sampling, monitoring.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	"Traffic monitoring using sflow", sflow.org, 2003. Retrieved from http://www.sflow.org/sFlowOverview.pdf See pages 1-4; figures 2, 3.	1-15
X	"sFlow and Benefits", sflow.org, 2004. Retrieved from http://www.sflow.org/about/benefits.ppt See pages 6-9.	1-15
X	US 2005-0190695 A1 (PETER PHAAL) 01 September 2005 See abstract; figures 1-6, 9; paragraphs [0060]-[0074].	1-15
A	US 2006-0159028 A1 (MARTIN CURRAN-GRAY et al.) 20 July 2006 See abstract; figures 1, 2, 7; claims 1-16.	1-15

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

☒ See patent family annex.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
30 NOVEMBER 2011 (30.11.2011)	**30 NOVEMBER 2011 (30.11.2011)**

Name and mailing address of the ISA/KR

Korean Intellectual Property Office
Government Complex-Daejeon, 189 Cheongsa-ro,
Seo-gu, Daejeon 302-701, Republic of Korea

Facsimile No. 82-42-472-7140

Authorized officer

Kim Sun Jong

Telephone No. 82-42-481-8260

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2011/028043

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2005-0190695 A1	01.09.2005	US 6894972 B1	17.05.2005
		US 7164657 B2	16.01.2007
US 2006-0159028 A1	20.07.2006	EP 1684463 A1	26.07.2006
		GB 0501174 D0	02.03.2005
		GB 2422505 A	26.07.2006