




(86) Date de dépôt PCT/PCT Filing Date: 2002/09/26
 (87) Date publication PCT/PCT Publication Date: 2003/04/10
 (85) Entrée phase nationale/National Entry: 2004/03/25
 (86) N° demande PCT/PCT Application No.: US 2002/030835
 (87) N° publication PCT/PCT Publication No.: 2003/030004
 (30) Priorité/Priority: 2001/09/28 (09/967,878) US

(51) Cl.Int.⁷/Int.Cl.⁷ G06F 15/16, G06F 15/173, G06F 11/00
 (71) Demandeur/Applicant:
NETSCREEN TECHNOLOGIES, INC., US
 (72) Inventeurs/Inventors:
MAO, YU MING, US;
LIAN, ROGER JIA-JYI, US;
HUANG, GUANGSONG, US;
CHEUNG, LEE CHIK, US
 (74) Agent: SMART & BIGGAR

(54) Titre : PROCEDE ET DISPOSITIF PERMETTANT L'IMPLEMENTATION D'UN PARE-FEUDE COUCHE 3/COUCHE 7 DANS UN DISPOSITIF L2
 (54) Title: METHOD AND APPARATUS FOR IMPLEMENTING A LAYER 3/LAYER 7 FIREWALL IN AN L2 DEVICE



(57) **Abrégé/Abstract:**

Methods and apparatus for transferring packets in a packet switched communication system. A system (100) is provided that includes an L2 device (106) including a controller (204) determining for each packet received whether the received packet is to be inspected, an inspection device (210) operable to inspect and filter packets identified by the controller (204) including using a zone specific policy and an L2 controller (230) for transferring inspected packets in accordance with L2 header information using L2 protocols.

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

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
10 April 2003 (10.04.2003)

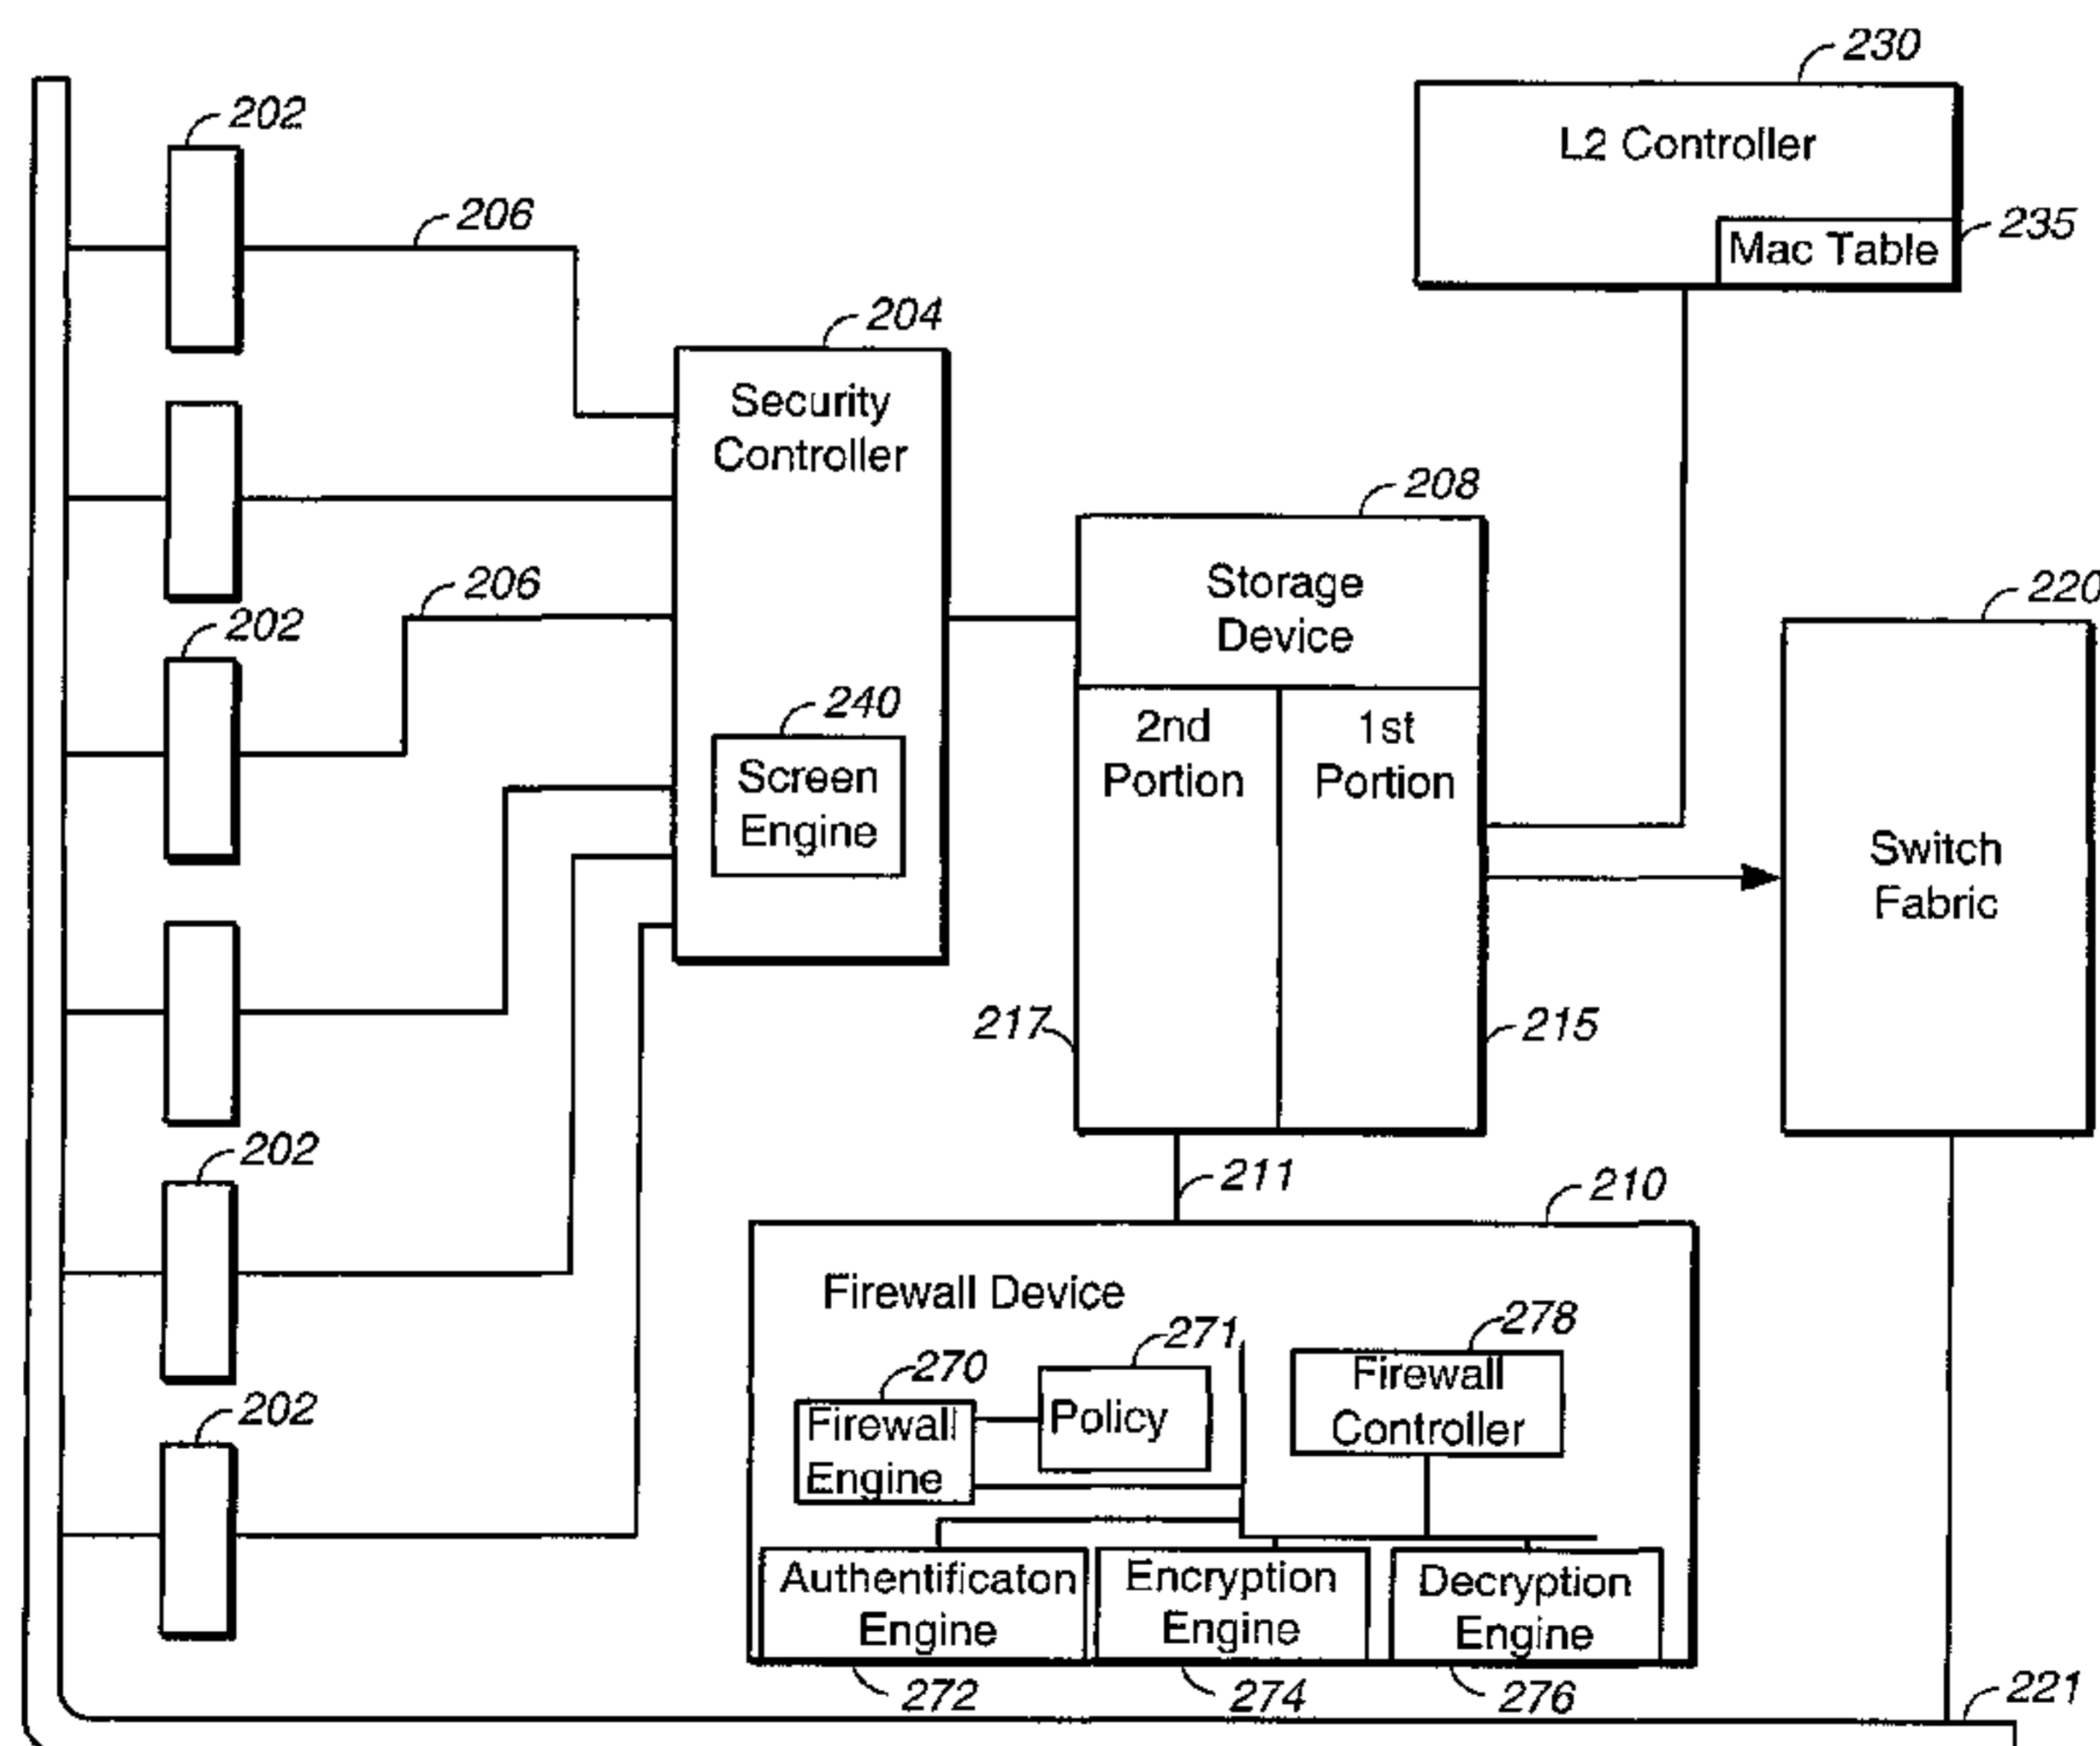
PCT

(10) International Publication Number
WO 03/030004 A1

- (51) International Patent Classification⁷: G06F 15/16, 15/173, 11/00 CHEUNG, Lee, Chik [US/US]; 1020 Belevedere Lane, San Jose, CA 95129 (US).
- (21) International Application Number: PCT/US02/30835 (74) Agents: KIRKLAND, Mark, D. et al.; Fish & Richardson, P.C., 500 Arguello Street, Suite 500, Redwood City, CA 94063 (US).
- (22) International Filing Date: 26 September 2002 (26.09.2002) (81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.
- (25) Filing Language: English (84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- (26) Publication Language: English
- (30) Priority Data: 09/967,878 28 September 2001 (28.09.2001) US
- (71) Applicant (for all designated States except US): NETSCREEN TECHNOLOGIES, INC. [US/US]; 350 Oakmead Parkway, Sunnyvale, CA 94085 (US).
- (72) Inventors; and (75) Inventors/Applicants (for US only): MAO, Yu, Ming [CN/US]; 350 Oakhurst Way, Milpitas, CA 95035 (US). LIAN, Roger, Jia-Jyi [US/US]; 1559 Sun Lane, San Jose, CA 95132 (US). HUANG, Guangsong [CN/US]; 930 W. Remington Drive, Apt. 3A, Sunnyvale, CA 94087 (US).
- Published: — with international search report

[Continued on next page]

(54) Title: METHOD AND APPARATUS FOR IMPLEMENTING A LAYER 3/LAYER 7 FIREWALL IN AN L2 DEVICE



(57) Abstract: Methods and apparatus for transferring packets in a packet switched communication system. A system (100) is provided that includes an L2 device (106) including a controller (204) determining for each packet received whether the received packet is to be inspected, an inspection device (210) operable to inspect and filter packets identified by the controller (204) including using a zone specific policy and an L2 controller (230) for transferring inspected packets in accordance with L2 header information using L2 protocols.

 WO 03/030004 A1

WO 03/030004 A1



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

**METHOD AND APPARATUS FOR IMPLEMENTING
A LAYER 3/LAYER 7 FIREWALL IN AN L2 DEVICE**

BACKGROUND

5 The present invention relates generally to data routing systems, and more particularly to methods and apparatus for providing secure communications on a network.

 A packet switch communication system includes a network of one or more switches or routers connecting a plurality of users. A packet is the fundamental unit
10 of transfer in the packet switch communication system. A user can be an individual user terminal or another network.

 A layer 2 (L2) switch is a switching device which receives packets containing data or control information on one port, and based on a media access connection (MAC) address contained within the packet, switches the packet out another port.
15 Conventional L2 switches perform this switching function by evaluating layer 2 (L2) header information contained within the packet in order to determine the proper output port for a particular packet. The L2 switch includes a table that maps MAC addresses with output ports. If a MAC address is unknown (i.e., there is no corresponding entry in the table), then the corresponding packet is broadcast to all
20 output ports with the hope that another component in the packet switched communication system will recognize the MAC address (and pass back information to the forwarding L2 switch to update its table). Other types of L2 devices include bridges.

 A router is a switching device which receives packets containing data or
25 control information on one port, and based on destination information contained within the packet, routes the packet to a next hop to/toward the destination. Conventional routers perform this switching function by evaluating layer 3 (L3) header information contained within the packet in order to determine a next hop for a particular packet. The layer 3 information includes an IP address associated with the
30 intended destination (as well as source address) for the packet.

The network coupling the users can be an intranet, that is, a network connecting one or more private servers such as a local area network (LAN). Alternatively, the network can be a public network, such as the Internet, in which data packets are passed over untrusted communication links. The network configuration
5 can include a combination of public and private networks. For example, two or more LAN's with individual terminals can be coupled together using a public network such as the Internet. Data security issues can arise when public and private networks are linked or when distinct networks are coupled. For example, conventional packet
10 switched communication systems that include links between public and private networks typically include security measures for assuring network access control and data integrity.

In order to assure individual packet security, packet switched communication systems can include encryption/decryption services. Prior to leaving a trusted network (or portion of a network), individual packets can be encrypted to minimize
15 the possibility of data loss while the packet is transferred over an untrusted (e.g., public) network (or portion thereof). Upon receipt at a destination or another trusted portion of the communication system (e.g., at a firewall just before the destination), the packet can be decrypted and subsequently delivered to its intended destination. The use of encryption and decryption allows for the creation of a virtual private
20 network (VPN) between users separated by untrusted communication links.

In addition to security concerns for the data transferred over the public portion of the communications system, the private portions of the network must safeguard against intrusions through the gateway provided at the interface of the private and the public networks. A firewall is a device that can be coupled in-line between a public
25 network and private network for screening packets received from the public network. A firewall is a particular type of L3/L4 device that can be used to enforce policy and filtering functions. A firewall can include one or more engines for inspecting, filtering, authenticating, encrypting, decrypting and otherwise manipulating received packets. Conventional firewalls use L3 and L4 header information including IP
30 addresses associated with the source and destination of a given packet being processed. Received packets are inspected and thereafter forwarded or dropped in accordance with the policies associated with the given domain.

SUMMARY

In one aspect, the invention provides an L2 device in a packet switched communication system. The packet switched communication system has plural zones and each zone represents a distinct security domain and has an associated policy for use in inspecting packets entering/exiting an associated zone. The L2 device includes at least one port coupled to a terminal unit included in a first security zone, at least one port coupled to a terminal unit included in a second security zone, a controller determining for each packet received whether the received packet is destined for another zone, a firewall engine operable to inspect and filter inter-zone packets using a zone specific policy and an L2 switching engine. The L2 switching engine is operable to immediately route to a port all intra-zone packets passing through the L2 device using a table of MAC addresses and corresponding ports, and only route to a port inter-zone packets that are retained after the inspection by the firewall engine.

In another aspect, the invention provides an L2 device in a packet switched communication system. The L2 device includes a controller determining for each packet received whether the received packet is to be transferred intra-zone or inter-zone, a firewall engine operable to inspect and filter inter-zone packets using a zone specific policy and an L2 switching engine operable to immediately route to a port all intra-zone packets passing through the L2 device using a table of MAC addresses and corresponding ports and only route to a port inter-zone packets that are retained after the inspection by the firewall engine.

In another aspect, the invention provides an L2 device in a packet switched communication system including a controller determining for each packet received whether the received packet is to be transferred inter-zone and a firewall engine operable to inspect and filter inter-zone packets using a zone specific policy prior to routing using L2 protocols.

In another aspect, the invention provides an L2 device in a packet switched communication system including a controller determining for each packet received whether the received packet is to be transferred inter-zone and an inspection device operable to inspect and filter inter-zone packets using a zone specific policy prior to routing using L2 protocols.

In another aspect, the invention provides an L2 device in a packet switched communication system including a controller determining for each packet received whether the received packet is to be inspected, an inspection device operable to inspect and filter packets identified by the controller including using a zone specific policy and an L2 controller for transferring inspected packets in accordance with L2 header information using L2 protocols.

Aspects of the invention can include one or more of the following features. The inspection device can be a firewall including a layer 3 firewall device, a layer 4 firewall device and a layer 7 firewall device. The inspection device can be a firewall that filters based on layer information other than layer 2 header information. The controller can determine each packet that is to pass between security zones and the inspection device only processes inter-zone traffic. The controller can determine each packet that is to remain in a single security zone and the inspection device immediately routes intra-zone packets. The device can route traffic using the MAC address in the layer 2 header of a given packet to determine an egress port on the device to which the packet is to be routed.

The device can include a storage element for storing packets that are to be inspected and an L2 controller for transferring packets through the device including determining an egress port for transferring a given packet using a destination MAC address in the given packet and a MAC address table that includes a mapping of MAC addresses and associated egress nodes. The memory element can include a first and second portion. The first portion can store packets to be transferred through the device and the second portion can store packets waiting for inspection. The device can be a L2 switch or an L2 bridge.

In another aspect, the invention provides a method for transferring packets in a communication network including receiving a packet at an L2 device, determining whether the received packet is to be transferred inter-zone and inspecting and filtering inter-zone packets using a zone specific policy prior to routing using L2 protocols.

In another aspect, the invention provides a method for transferring packets in a communication network including receiving a packet at an L2 device, determining whether the received packet is to be inspected and inspecting and filtering identified

packets using a zone specific policy prior to transferring the packet through the L2 device using L2 protocols.

In another aspect, the invention provides a method for switching packets in a communication network including receiving a packet at an interface of an L2 device,
5 determining if a destination MAC address associated with the received packet is known and, if not, holding the received packet a predetermined amount of time without transferring the packet to any port of the L2 device, creating a probe packet that includes the unknown MAC address and broadcasting the probe packet to all interfaces except the receiving interface.

10 Aspects of the invention can include one or more of the following features. The probe packet can include a time to life (TTL) field in a IP header and the method can include setting a value of the TTL field such that a downstream node having the unknown MAC address and receiving the probe cell will return an expired message to the L2 device. The method can include dropping the packet after the expiration of the
15 predetermined amount of time. The packet can be dropped if the MAC address is unknown. The method can include receiving a response from on one of the broadcast interfaces and updating a table indicating a previously unknown MAC address is associated with the responding interface.

In another aspect, the invention provides method of providing secure
20 communications between users without requiring encryption and decryption services at a respective user. The method includes identifying first and second users, coupling the first and second users through two or more L2 devices over a communication network and specifying a virtual private network for communications between the first and second users. The virtual private network is defined between a first and
25 second L2 device in the network. The method includes receiving a packet at either the first or the second L2 device, determining whether the received packet is associated with the virtual private network and encrypting and decrypting as appropriate identified packets using local encryption and decryption services prior to transferring the packet through the L2 device using L2 protocols.

Aspects of the invention can include one or more of the following features. The step of determining can include using a destination MAC address associated with the packet to identify a virtual private network.

In another aspect, the invention provides a virtual private network for
5 providing secure communications between users without requiring encryption and decryption services at a respective user. The virtual private network includes first and second L2 devices coupling first and second users over a communication network where each of the first and second L2 devices includes a screening mechanism
10 determining whether a received packet is associated with the virtual private network and encryption and decryption services operating on packets associated with the virtual private network prior to a transfer of the packet through the L2 device using L2 protocols.

Aspects of the invention can include one or more of the following advantages. A packet switched communication system is provided that allows for the creation of
15 plural security zones within a single device without requiring changes to the network or terminal configuration. Between each zone, a terminal unit can communicate with other terminal units without the knowledge of, yet receiving the benefits of, L2 switching and up to layer 7 security filtering as discussed below. A packet switched communication system is provided that includes L2 switch and firewall functionality.
20 The packet switched communication system acts as an IEEE 802.1Q VLAN L2 conventional switch forwarding/ filtering based on MAC-address for all intra-zone communications. The packet switched communication system allows L2 switching among multiple ports inside a given security zone. The L2 switch also provides up to layer 7 security firewall protections as appropriate for inter-zone or intra-zone traffic
25 including TCP stateful inspection, syn-attack guard, policy-based control, load balancing and other functionalities on each data stream. In one implementation, the packet switched communication system can be configured to include multiple IEEE 802.1Q VLAN based L2 transparent domains. A user can create multiple VLANs, each having its own policy for firewall control. In addition, methods are provided for
30 VPN tunnel capability to connect remote clients to the L2 domain. Methods are provided to guard against broadcasting information throughout the zones and violating one or more security constraints when a MAC address that is being

processed is not recognized. The methods include the broadcast of probe packets to discover topology information for unknown MAC destinations.

The details of one or more embodiments of the invention are set forth in the accompanying drawings and the description below. Other features, objects, and advantages of the invention will be apparent from the description and drawings, and
5 from the claims.

DESCRIPTION OF DRAWINGS

FIG. 1 is a block diagram of a packet switched communication system including an L2 firewall enabled switch.

10 FIG. 2a is a schematic view of an L2 firewall enabled switch.

FIG. 2b shows an exemplary communication network including plural zones partitioned by a single security switch.

FIG. 3 is a flow diagram of a method for processing packets in the security switch of FIG. 2a.

15 FIG. 4 is a flow diagram for a method for processing un-recognized packets in the security switch of FIG. 2a.

Like reference symbols in the various drawings indicate like elements.

DETAILED DESCRIPTION

Referring now to FIG. 1, a packet switch communication network 100 includes
20 a plurality of terminal units 102 configured in a plurality of zones 104 and coupled by one or more switches 106.

In one implementation, each terminal unit 102 is of the form of a standalone computer (e.g., a personal computer, a laptop or workstation). Alternatively, one or more terminal units may be of the form of a personal digital assistant (PDA), Web
25 pad, two-way pager, cellular handset, or other termination or remote device in a communication or computing environment. In one implementation, each terminal is a gateway to another network or group of terminal units (e.g., to a LAN or a pool of servers).

Each zone 104 embodies a security domain in the communication system. Each security domain can include separate policy, traffic management, accounting and administrative definitions and functions. Security policies, traffic management and other filtering functions can be enforced among and within zones. In one
5 implementation, security policies are enforced between zones, while intra-zone communications are not subject to the security constraints. In one implementation, zones overlap. When zones overlap, policies associated with a parent zone can be a superset of the policies associated with one or more sub-zones (each including a subset of the overall policies). Alternatively, the policies associated with the parent
10 zone may be separate and distinct from the policies of each sub-zone. For example, in one implementation, a zone can include one or more sub-zones, each including a separate set of policies.

In one implementation, each zone is associated with physical boundaries or other segmentation in the communication network. Alternatively, the assignment of
15 particular terminal units to zones may represent groupings or combinations in a business structure (e.g., zones used to separate different functional entities in a business organization). Alternatively, the zones have no particular relation to physical boundaries. Communication between terminal units in each zone and among terminal units within a zone are controlled in accordance with protocols described
20 below in association with switch 106.

Switch 106 may be of different types. In one implementation, each switch 106 is configured as a layer 2 (L2) device and includes a plurality of ports on which packets from the communication network are received and transferred in accordance with L2 protocols. Each switch 106 includes a media access connection (MAC) table
25 for use in determining switching of received packets. The MAC table associates MAC addresses with ports of the switch 106. Packets are processed as they arrive at the ports of each switch 106 in accordance with L2 header information contained within a given packet. Depending on the MAC address, packets are switched to an appropriate output port as specified in the MAC table.

30 One or more of switches 106 are configured to enforce security domain constraints. For example, one or more of switches 106 is configured as an L2 firewall

enabled security switch (hereinafter “security switch”). Referring now to FIG. 2, a security switch 200 includes a plurality of ports 202, a switch fabric 220 and an L2 controller 230. Each port 202 is coupled to a security controller 204 by a bus 206. The security controller 204 is coupled to one or more storage elements 208. In one
5 implementation (not shown), each port 202 is associated with a separate security controller 204 and storage element 208. Alternatively, the security controller functionality can be combined in a single (as shown) or lesser number of individual security controller units. In addition, packets associated with all ports 202 can be stored in a single memory element 208 (as shown). Security switch 200 also includes
10 a firewall device 210 that is coupled to (each) storage element 208 by a security bus 211.

L2 controller 230 supports L2 switching protocols. Packets are either directly processed (e.g., intra-zone packets) or processed after a security screening (e.g., for inter-zone packets) as discussed in greater detail below. Associated with L2 controller
15 230 is a MAC table 235. MAC table 235 includes plural entries each of which includes a MAC address and an indicator of a port 202 associated therewith. Switch fabric 220 is used to route traffic from storage element 208 to a respective port 202 under the control of L2 controller 230 using bus 221.

Storage element 208 is partitioned into two portions. A first portion 215 is
20 used to store packets received from a port 202 that are not subject to security screening. For example, in one implementation, packets received from a terminal unit in a same security zone (e.g., intra-zone traffic) are not subject to security screening. Un-screened packets are processed directly by L2 controller 230 and forwarded out a designated port in accordance with L2 protocols as specified in MAC table 235.
25 Second portion 217 is used to store packets to be screened by firewall device 210.

Security controller 204 includes a screening engine 240. Screening engine 240 examines each packet received from a respective port 202 and determines whether security screening is to be performed. In one implementation, screening engine 240 examines the L2 header for each packet, and based on the screening, either
30 forwards the packet to the first or second portion 215 and 217, respectively, of storage element 208. The L2 header includes a destination MAC address that can be mapped

to an egress port on the device using the MAC table 235. Associated with each ingress and egress port is a security zone identifier. Security zone identifiers can be stored in a table of zone identifiers (not shown) that is indexed by port identifier (id). Screening engine 240 compares the security zone identifier associated with the packet being processed (determined from the identification of the egress port from the MAC table using the destination MAC address in the header of the packet being processed) with the security zone identifier associated with the port on which the packet was received in the device. Based on the comparison, screening engine 240 can determine whether the packet is destined for another zone (i.e., constitutes intra-zone or inter-zone communication).

The screening of packets can be with or without the knowledge of the individual terminal units. Associated with security switch 200 is a user interface (not shown) and associated management tools (not shown) for constructing one or more security zones. In one implementation, the security zones are determined based on the destination MAC address included in the L2 header of the packet received. More specifically, each egress port can be assigned to a security zone and have an associated security zone identifier associated therewith. Alternatively, the security zones can be created for plural users coupled to different ports of the security switch 200. For example, security switch 200 can be configured to include three ports, where terminal units associated with a first two of the ports are assigned to a first zone, while terminal units associated with the third port are assigned to a second zone. Other configurations are possible. Zone assignments and partitions are discussed in greater detail below. The user interface allows an administrator or user to configure the security switch 200. The security switch 200 can be configured to create plural security zones and associate one or more interfaces with each zone. Thereafter, policies can be established for inspecting or otherwise screening packets as they traverse the security switch 200.

Firewall device 208 includes plural engines for performing packet screening prior to routing packets through security switch 200. Firewall device 208 includes a firewall engine 270 and associated policies 271, authentication engine 272, encryption engine 274, decryption engine 276 and a firewall controller 278.

Firewall controller 278 extracts packets from second portion 217 of storage element 208. Firewall controller 278 oversees the distribution of packets within the firewall device as well as the coordination among the respective engines. Each packet is evaluated and processed in accordance with policies based on one or more
5 considerations. For example, packets can be screened based on source, destination or both. One or more policies 271 are retrieved and used by firewall engine 270 to inspect the packet. Packet inspection may also require encryption, decryption and authentication services. One or more of the encryption 272, decryption 274 and authentication 276 engines can be invoked by the firewall controller 278 as part of the
10 inspection processes. In addition, other services can be provided including virtual private network termination services, session set-up and various other traffic management and security related functions. Examples of screening services are discussed in greater detail below. After the inspection, packets can be forwarded in the network or dropped as appropriate. In one implementation, packets that are to be
15 forwarded (e.g., pass the inspection) are prepared as appropriate (e.g., encrypted) then forwarded to the first portion 215 of storage element 208. Alternatively, the packets may be returned to the second portion 217 of storage element 208 and marked as having been screened. In one implementation, screened packets are forwarded to a queue for processing by L2 controller 230. Screened packets are then processed by
20 L2 controller 230 and switched to an appropriate output port in accordance with conventional L2 processing protocols.

Each security switch 200 can be configured to create plural security zones. For example, a communications network having a security switch 200 is shown in FIG. 2b. The communications network is a VLAN structure that includes 3 zones.
25 Security switch 200 includes a user interface and administrative control mechanisms for creating each of the security zones, specifying policies and other criteria for defining and managing each zone. The security zones enforced by the security switch 200 can be transparent to the end users. That is, the security zones can be established at the security switch 200 including the specification of all operating parameters
30 associated with the security domain. Users in each zone may be unaware of the zone structure and may communicate with other users in a conventional manner. For example, a virtual private network can be created between users including encryption

and decryption services without requiring the actual encryption and decryption support in the respective end users (e.g., encryption and decryption services can be provided in secure switches disposed between the two users). Accordingly, a system administrator can create a virtual private network between a remote user in one security zone and another user in a second security zone where the individual users are unaware of the VPN services and are not required to include encryption or decryption services locally. In one implementation, the administrator provisioned VPN services are specified for remote users in a same zone.

Alternatively, the users may be aware of the security structure and include indicators (e.g., zone identifiers) in packets transferred to other users. Each user may define their own custom L2 zone and an inter-zone policy for their network security requirements. For example, security switch 200 shown in FIG. 2b embodies a VLAN that includes v1-trust, v1-untrust and v1-dmz zones. V1-trust defines a zone that includes two users including user 291 and user 292. V1-untrust defines a zone that includes a single user 293. V1-dmz defines a zone that includes three users, users 291, 292 and user 294. Separate policies can be enforced for communications between the three zones. For example, communications that are intra-zone between user 291 and user 292 will not require inspection, and as such are handled by security switch 200 in accordance with conventional L2 protocols. Communications from user 291 to user 293 will invoke an inspection process as defined by the security system architect (e.g., user 291 or 292 or an administrator for such) for communications between V1-trust and V1-untrust. Similarly, communications between user 294 and user 291 will invoke an inspection process (e.g., a potentially lesser screen) for communications between V1-dmz and V1-trust.

Multiple interfaces are allowed inside each zone. For intra-zone traffic, security switch 200 behaves like a tradition L2 bridge forwarding a given packet based on the destination MAC-address. In one implementation, no firewall protection mechanisms are applied for the intra-zone traffic.

For inter-zone traffic, standard firewall inspections (including policy inspection, TCP stateful inspection, etc. as described above) are performed for each

incoming packet. In all cases, the egress interface is determined by the learned destination MAC address on the interface.

Packet Flow

Referring now to FIG. 3, a method 300 is shown, as invoked by the security switch 200, for processing packets. The method described is made with no particular reference to the specific hardware elements performing the steps. An exemplary hardware configuration is given above. The method can however be implemented in L2 switches having other configurations. The method begins with the receipt of a packet (302). The packet is evaluated to determine whether the packet is to be inspected (304). If so, the packet is pre-processed as appropriate (305) and one or more policies are retrieved (306). The pre-processing of the packet can include decryption and authentication services. The retrieval of a policy includes the identification of the zone to which the packet is being transferred. Packets traveling between zones can be inspected using a security policy. Intra-zone communications may not be inspected. In one implementation, policies can be enforced on intra-zone communications. The retrieval of a policy includes a MAC look-up for the MAC destination address in a received packet in the MAC table to determine an egress port associated with the MAC address and necessarily a security zone. The security zones associated with the packet's ingress and egress ports are compared to determine if the packet is passing to another zone. Assuming that an inspection is to occur, an appropriate policy is retrieved (i.e., based on the ingress port and egress port identifiers and their respective security zones). Thereafter, the packet is inspected (308). Packet inspection can include screening and dropping the packet as required. If the packet is to be forwarded on the network (309), post-processing operations are invoked as appropriate (310). Alternatively, the packet is dropped (311). The post processing operations can include session set-up, encryption and other functions. Thereafter the packet is processed in accordance with conventional L2 protocols starting at step 312.

At step 312, either a packet has passed inspection or did not require inspection. In either case, L2 header information is extracted to determine a MAC address associated with the packet. A look-up of the MAC address is performed

(314) and the packet is then routed to an appropriate output port (316). Thereafter the process ends.

Referring again to FIG. 2, the process steps are described with reference to one hardware implementation of the invention. Packets are received at a port 202. Each
5 packet is transferred on bus 205 to, and routed through, security controller 204 and stored in storage element 208 via a storage bus 209. Security controller 204 evaluates each packet to determine if inspection is required and forwards the packets to an appropriate portion of storage device 208. Packets that are not to be inspected (i.e., packets stored in first portion 215 of storage device 208) are processed by L2
10 controller 230. When L2 controller 230 is available, packets are fetched and processed to determine a port to which the packet should be forwarded. L2 controller 230 evaluates the MAC address associated with the packet, and using MAC table 235, determines a port for routing. After processing by the L2 controller 230, the packet is forwarded to an appropriate link into switch fabric 220 for routing to a determined
15 output port 202.

Packets that are to be inspected are transferred by security controller 204 into second portion 217 of storage element 208. When firewall engine 210 is available, a packet is fetched and processed to determine a security policy to be used in inspecting the packet. Firewall engine 270 evaluates IP address(es) associated with the packet
20 and implements traffic control and management functions as appropriate. Packets that are to be forwarded (i.e., pass inspection) are returned to storage element 208. Thereafter, the packet can be forwarded to an appropriate link into switch fabric 220 for routing to a determined output port 202. Other packets are dropped or otherwise handled in accordance with the policies defined for the given security zones.

25 As discussed above, all packets that pass the inspection in the firewall device 210 as well as all packets that are not required to be inspected, are processed by L2 controller 230 in accordance with conventional L2 protocols. In one implementation, the processing of packets by L2 controller is modified to maintain security zones. More specifically, as discussed above, conventional L2 switches broadcast on all ports
30 a packet that has a MAC address that is not recognized. This type of broadcast may well violate one or more security policies in place for given zones in the

communication network. Accordingly, in one implementation a test packet is broadcast to each port. The broadcasting of test packets is described in more detail in association with FIG. 4.

Referring now to FIG. 4, a method 400 is shown for handling packets by the L2 controller and includes receiving a packet to be processed (402). The MAC address for the packet is extracted (404). A check is made to locate an entry in a MAC address table that corresponds to the extracted MAC address (406). If a match is located (407), the packet is routed to an output port associated with the matching entry (408). If no match is located, the packet is dropped (410). In one implementation, the packet is merely held for a predetermined amount of time in hope of receiving information regarding the non-matching MAC address. If no match is located, a probe packet is created (412). The probe packet includes the MAC address associated with the packet being processed (i.e., the original ingress packet). In one implementation, the probe packet is an "ICMP PING" packet with an IP TTL field set to 1. Each packet includes the same MAC addresses (L2) and source /destination IPs (L3) as the ingress packet whose MAC address could not be located. The probe packet is then broadcast to all ports (414). A check is made to determine if a response is received on any of the security device's ports (416). The ICMP PING packet will cause the right gateway, which was to receive and forward the original ingress packet, to respond to the L2 controller in the device with an "ICMP TTL expired" message packet. From the expired packet, the system can identify the proper egress port/zone associated with the received MAC address. This method guarantees that no information in the original ingress packet will be leaked out. If a response is received (indicating that a device coupled to the receiving port is configured to process packets having the identified MAC address), then the MAC table is updated to include an entry having the MAC address and a port identifier indicating the port on which the response was received (418). Thereafter the process ends.

A number of embodiments of the invention have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention. For example, the firewall device has been described in terms of screening at the L3 layer level. Alternatively, other screening can be invoked at other levels including layers up to and including layer 7 (L7)

processing. Accordingly, other embodiments are within the scope of the following claims.

CLAIMS

1. A layer 2 device in a packet switched communication system, the packet switched communication system having plural zones, each zone representing a distinct security domain and having an associated policy for use in inspecting packets
5 entering/exiting an associated zone, the a layer 2 device comprising:

at least one port coupled to a terminal unit included in a first security zone;

at least one port coupled to a terminal unit included in a second security zone;

10 a controller determining for each packet received whether the received packet is destined for another zone;

a firewall engine inspecting and filtering inter-zone packets using a zone specific policy; and

15 a layer 2 switching engine immediately transferring to a port all intra-zone packets passing through the layer 2 device using a table of MAC addresses and corresponding ports, and

only transferring to a port inter-zone packets that are retained after the inspection by the firewall engine.

2. A layer 2 device in a packet switched communication system, the packet switched communication system having plural zones, each zone representing a distinct security domain and having an associated policy for use in inspecting packets
entering/exiting an associated zone, the layer 2 device comprising:

20 a controller determining for each packet received whether the received packet is to be transferred intra-zone or inter-zone;

a firewall engine inspecting and filtering inter-zone packets using a zone specific policy; and

25 a layer 2 switching engine operable to immediately route to a port all intra-zone packets passing through the layer 2 device using a table of MAC addresses and corresponding ports, and

only route to a port inter-zone packets that are retained after the inspection by the firewall engine.

3. A layer 2 device in a packet switched communication system, the packet switched communication system having plural zones, each zone representing a distinct security domain, the layer 2 device comprising:

5 a controller determining for each packet received whether the received packet is to be transferred inter-zone; and

a firewall engine inspecting and filtering inter-zone packets using a zone specific policy prior to routing using layer 2 protocols.

10 4. A layer 2 device in a packet switched communication system, the packet switched communication system having plural zones, each zone representing a distinct security domain, the layer 2 device comprising:

a controller determining for each packet received whether the received packet is to be transferred inter-zone; and

an inspection device inspecting and filtering inter-zone packets using a zone specific policy prior to routing using layer 2 protocols.

15 5. A layer 2 device in a packet switched communication system, the packet switched communication system having plural zones, each zone representing a distinct security domain, the layer 2 device comprising:

a controller determining for each packet received whether the received packet is to be inspected;

20 an inspection device inspecting and filtering packets identified by the controller including using a zone specific policy; and

a layer 2 controller transferring inspected packets in accordance with layer 2 header information using layer 2 protocols.

6. The device of claim 5 wherein the inspection device is a firewall.

25 7. The device of claim 5 wherein the inspection device is a layer 3 firewall device.

8. The device of claim 5 wherein the inspection device is a layer 4 firewall device.

9. The device of claim 5 wherein the inspection device is a layer 7 firewall device.

5 10. The device of claim 5 wherein the inspection device is a firewall that filters based on layer information other than layer 2 header information.

11. The device of claim 5 wherein the controller determines each packet that is to pass between security zones and the inspection device only processes inter-zone traffic.

10 12. The device of claim 5 wherein the controller determines each packet that is to remain in a single security zone and the inspection device immediately transfers intra-zone packets.

13. The device of claim 12 wherein the device uses a MAC address in the layer 2 header of a given packet to determine an egress port on the device to which
15 the packet is to be transferred.

14. The device of claim 5 further comprising a storage element for storing packets that are to be inspected and a layer 2 controller transferring packets through the device including determining an egress port for transferring a given packet using a destination MAC address in the given packet and a MAC address table that includes a
20 mapping of MAC addresses and associated egress nodes.

15. The device of claim 14 wherein the memory element includes a first and second portion, the first portion storing packets to be transferred through the device, and the second portion storing packets waiting for inspection.

16. The device of claim 5 wherein the device is a layer 2 switch.

25 17. The device of claim 5 wherein the device is a layer 2 bridge.

18. A method for transferring packets in a communication network, the communication network including plural zones, each zone representing a distinct security domain, the method comprising:

receiving a packet at a layer 2 device;

determining whether the received packet is to be transferred inter-zone;
and

inspecting and filtering inter-zone packets using a zone specific policy
prior to routing using layer 2 protocols.

5 19. A method for transferring packets in a communication network, the
communication network including plural zones, each zone representing a distinct
security domain, the method comprising:

receiving a packet at a layer 2 device;

determining whether the received packet is to be inspected; and

10 inspecting and filtering identified packets using a zone specific policy
prior to transferring the packet through the layer 2 device using layer 2 protocols.

20. A method for switching packets in a communication network the
communication network including plural zones, each zone representing a distinct
security domain, the method comprising:

15 receiving a packet at an interface of a layer 2 device;

determining if a destination MAC address associated with the received
packet is known; and

if not,

20 holding the received packet a predetermined amount of time without
transferring the packet to any port of the layer 2 device,

creating a probe packet that includes the unknown MAC address, and

broadcasting the probe packet to all interfaces except the receiving
interface.

25 21. The method of claim 20 wherein the probe packet includes a time to life
(TTL) field in an IP header and the method includes setting a value of the TTL field
such that a downstream node having the unknown MAC address and receiving the
probe packet will return an expired message to the layer 2 device.

22. The method of claim 20 further comprising dropping the packet after the
expiration of the predetermined amount of time.

23. The method of claim 20 wherein the packet is dropped if the MAC address is unknown.

24. The method of claim 20 further comprising receiving a response from one
5 of the broadcast interfaces and updating a table indicating a previously unknown MAC address is associated with the responding interface.

25. A method of providing secure communications between users without requiring encryption and decryption services at a respective user, the method comprising:

10 identifying first and second users;

coupling the first and second users through two or more layer 2 devices over a communication network;

specifying a virtual private network for communications between the first and second users, the virtual private network defined between a first and second layer
15 2 device in the network;

receiving a packet at either the first or the second layer 2 device;

determining whether the received packet is associated with the virtual private network; and

20 encrypting and decrypting as appropriate identified packets using local encryption and decryption services prior to transferring the packet through the layer 2 device using layer 2 protocols.

26. The method of claim 25 wherein the step of determining includes using a destination MAC address associated with the packet to identify a virtual private network.

25 27. A virtual private network for providing secure communications between users without requiring encryption and decryption services at a respective user, the virtual private network comprising:

first and second layer 2 devices coupling first and second users over a communication network where each of the first and second layer 2 devices includes

a screening mechanism determining whether a received packet is associated with the virtual private network, and

5 encryption and decryption services operating on packets associated with the virtual private network prior to a transfer of the packet through the layer 2 device using layer 2 protocols.

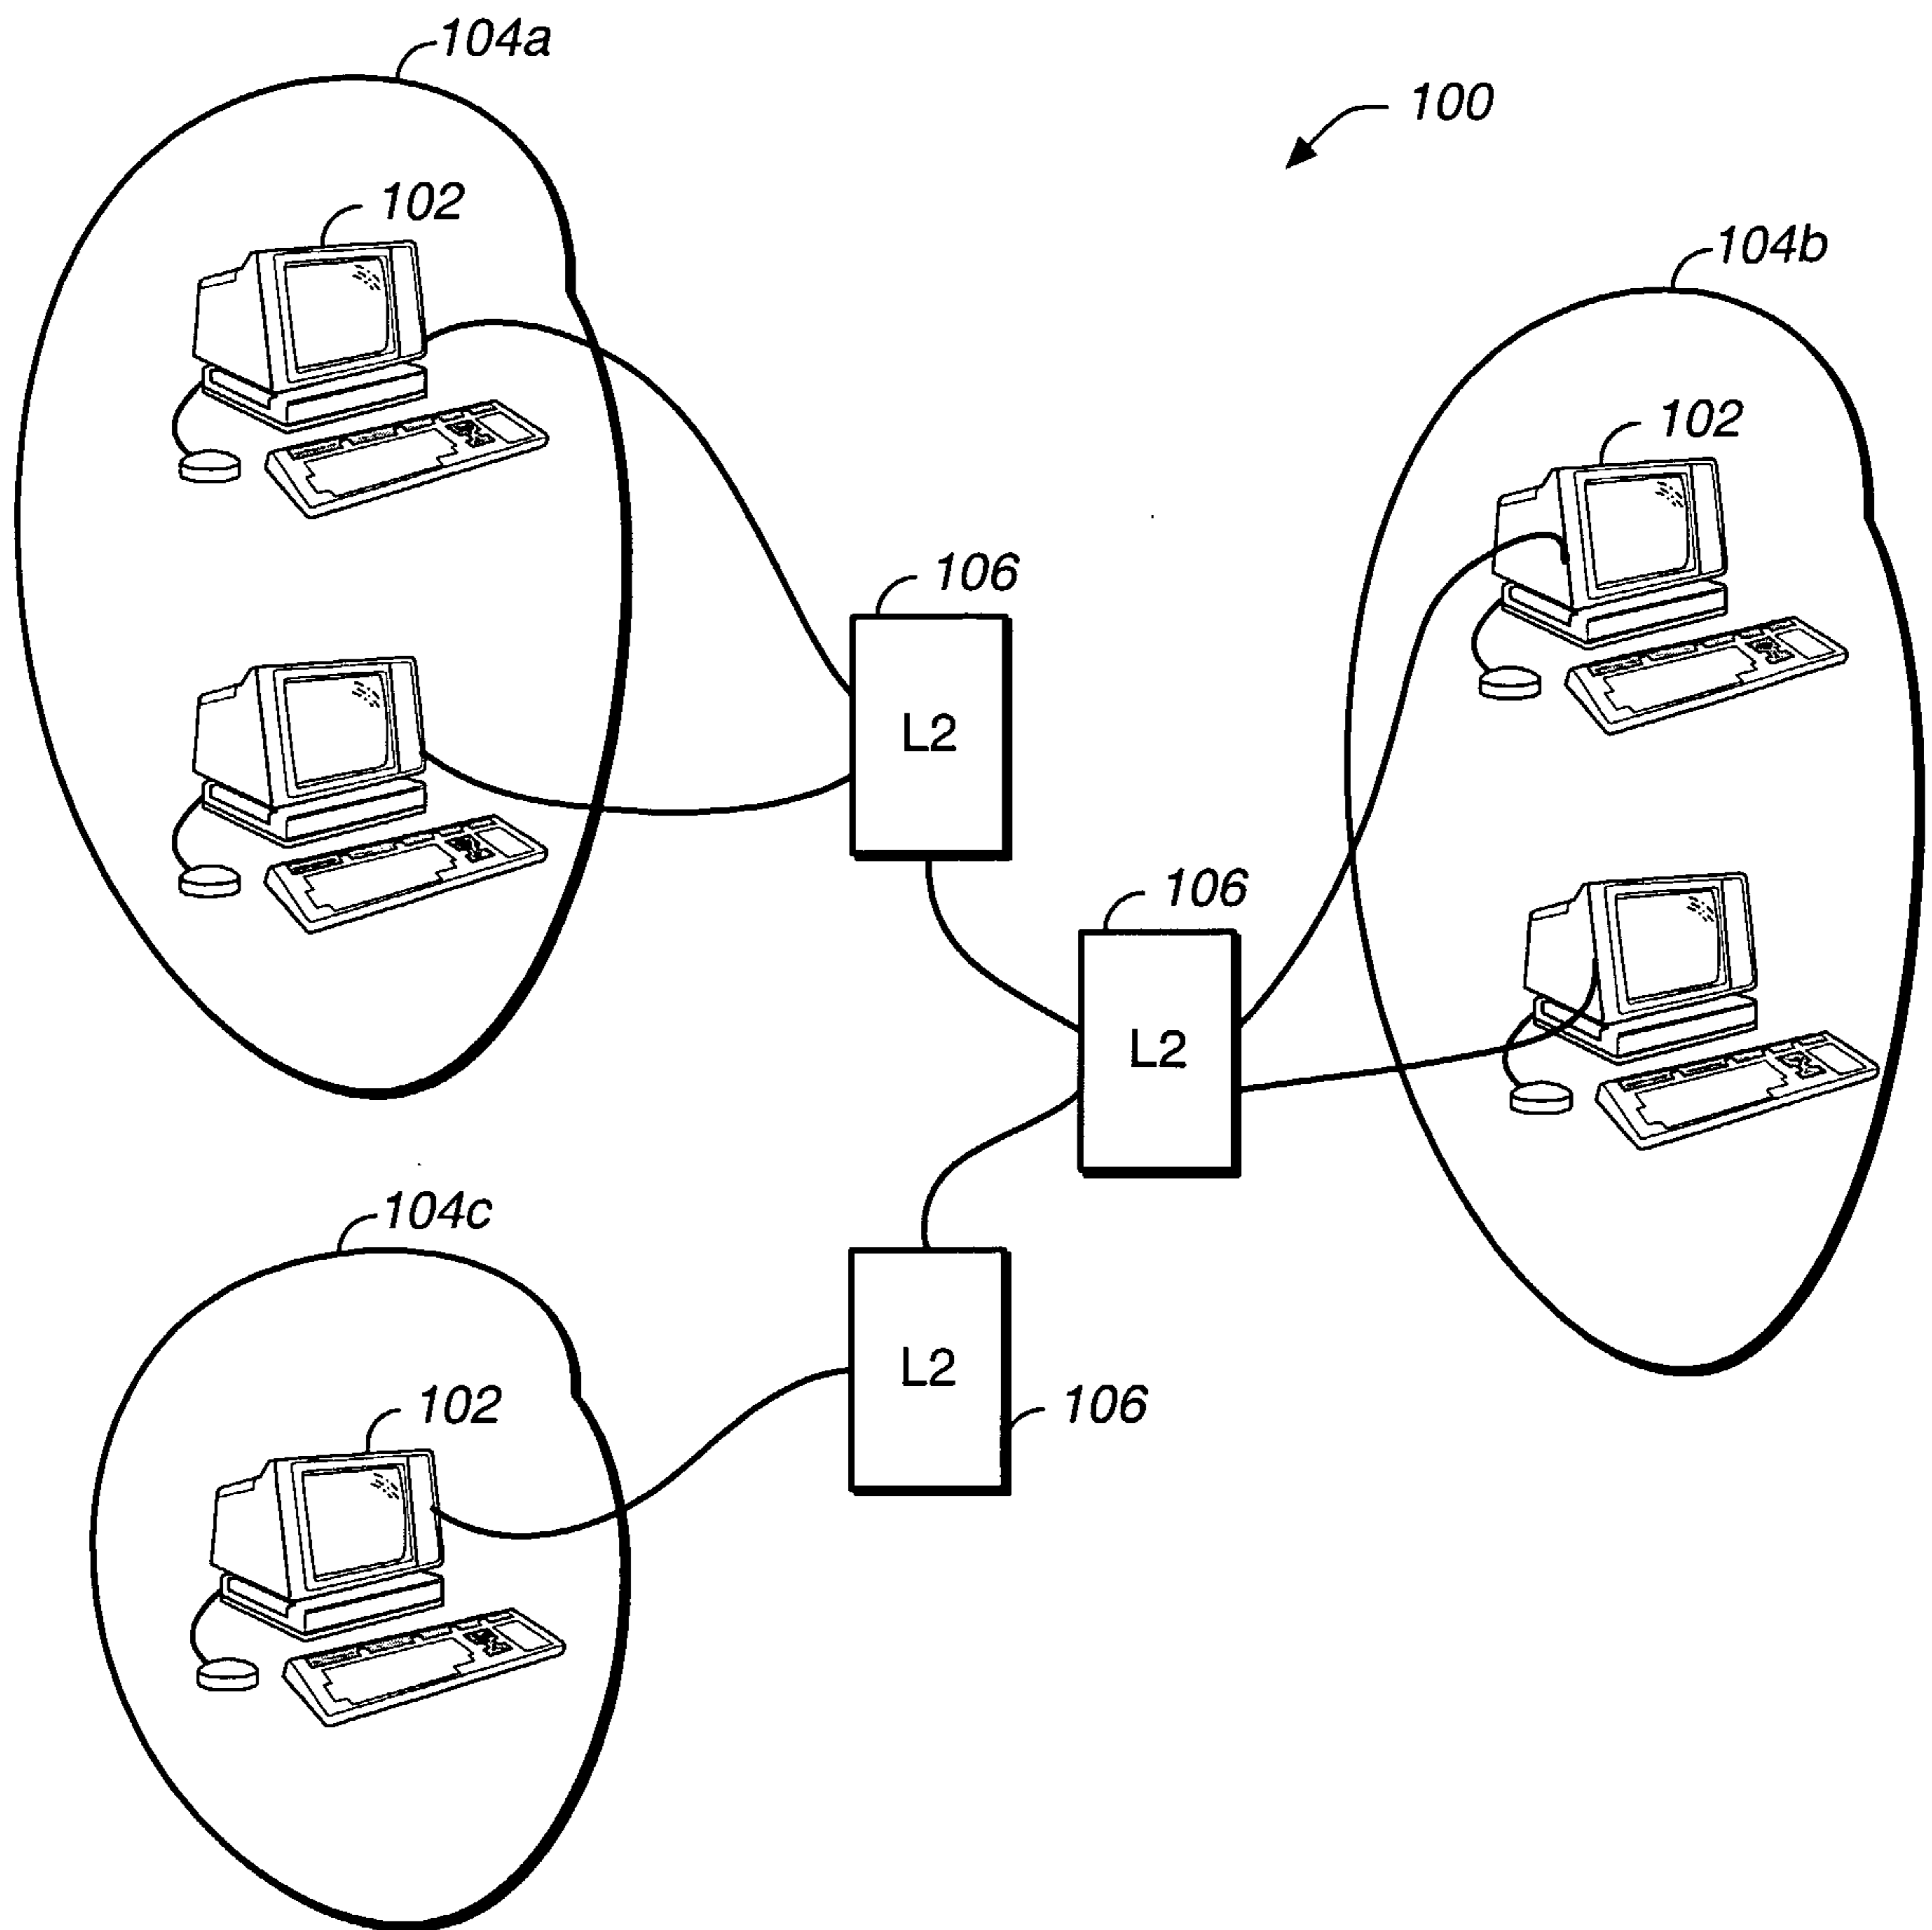


FIG. 1




FIG. 2a




FIG. 2b

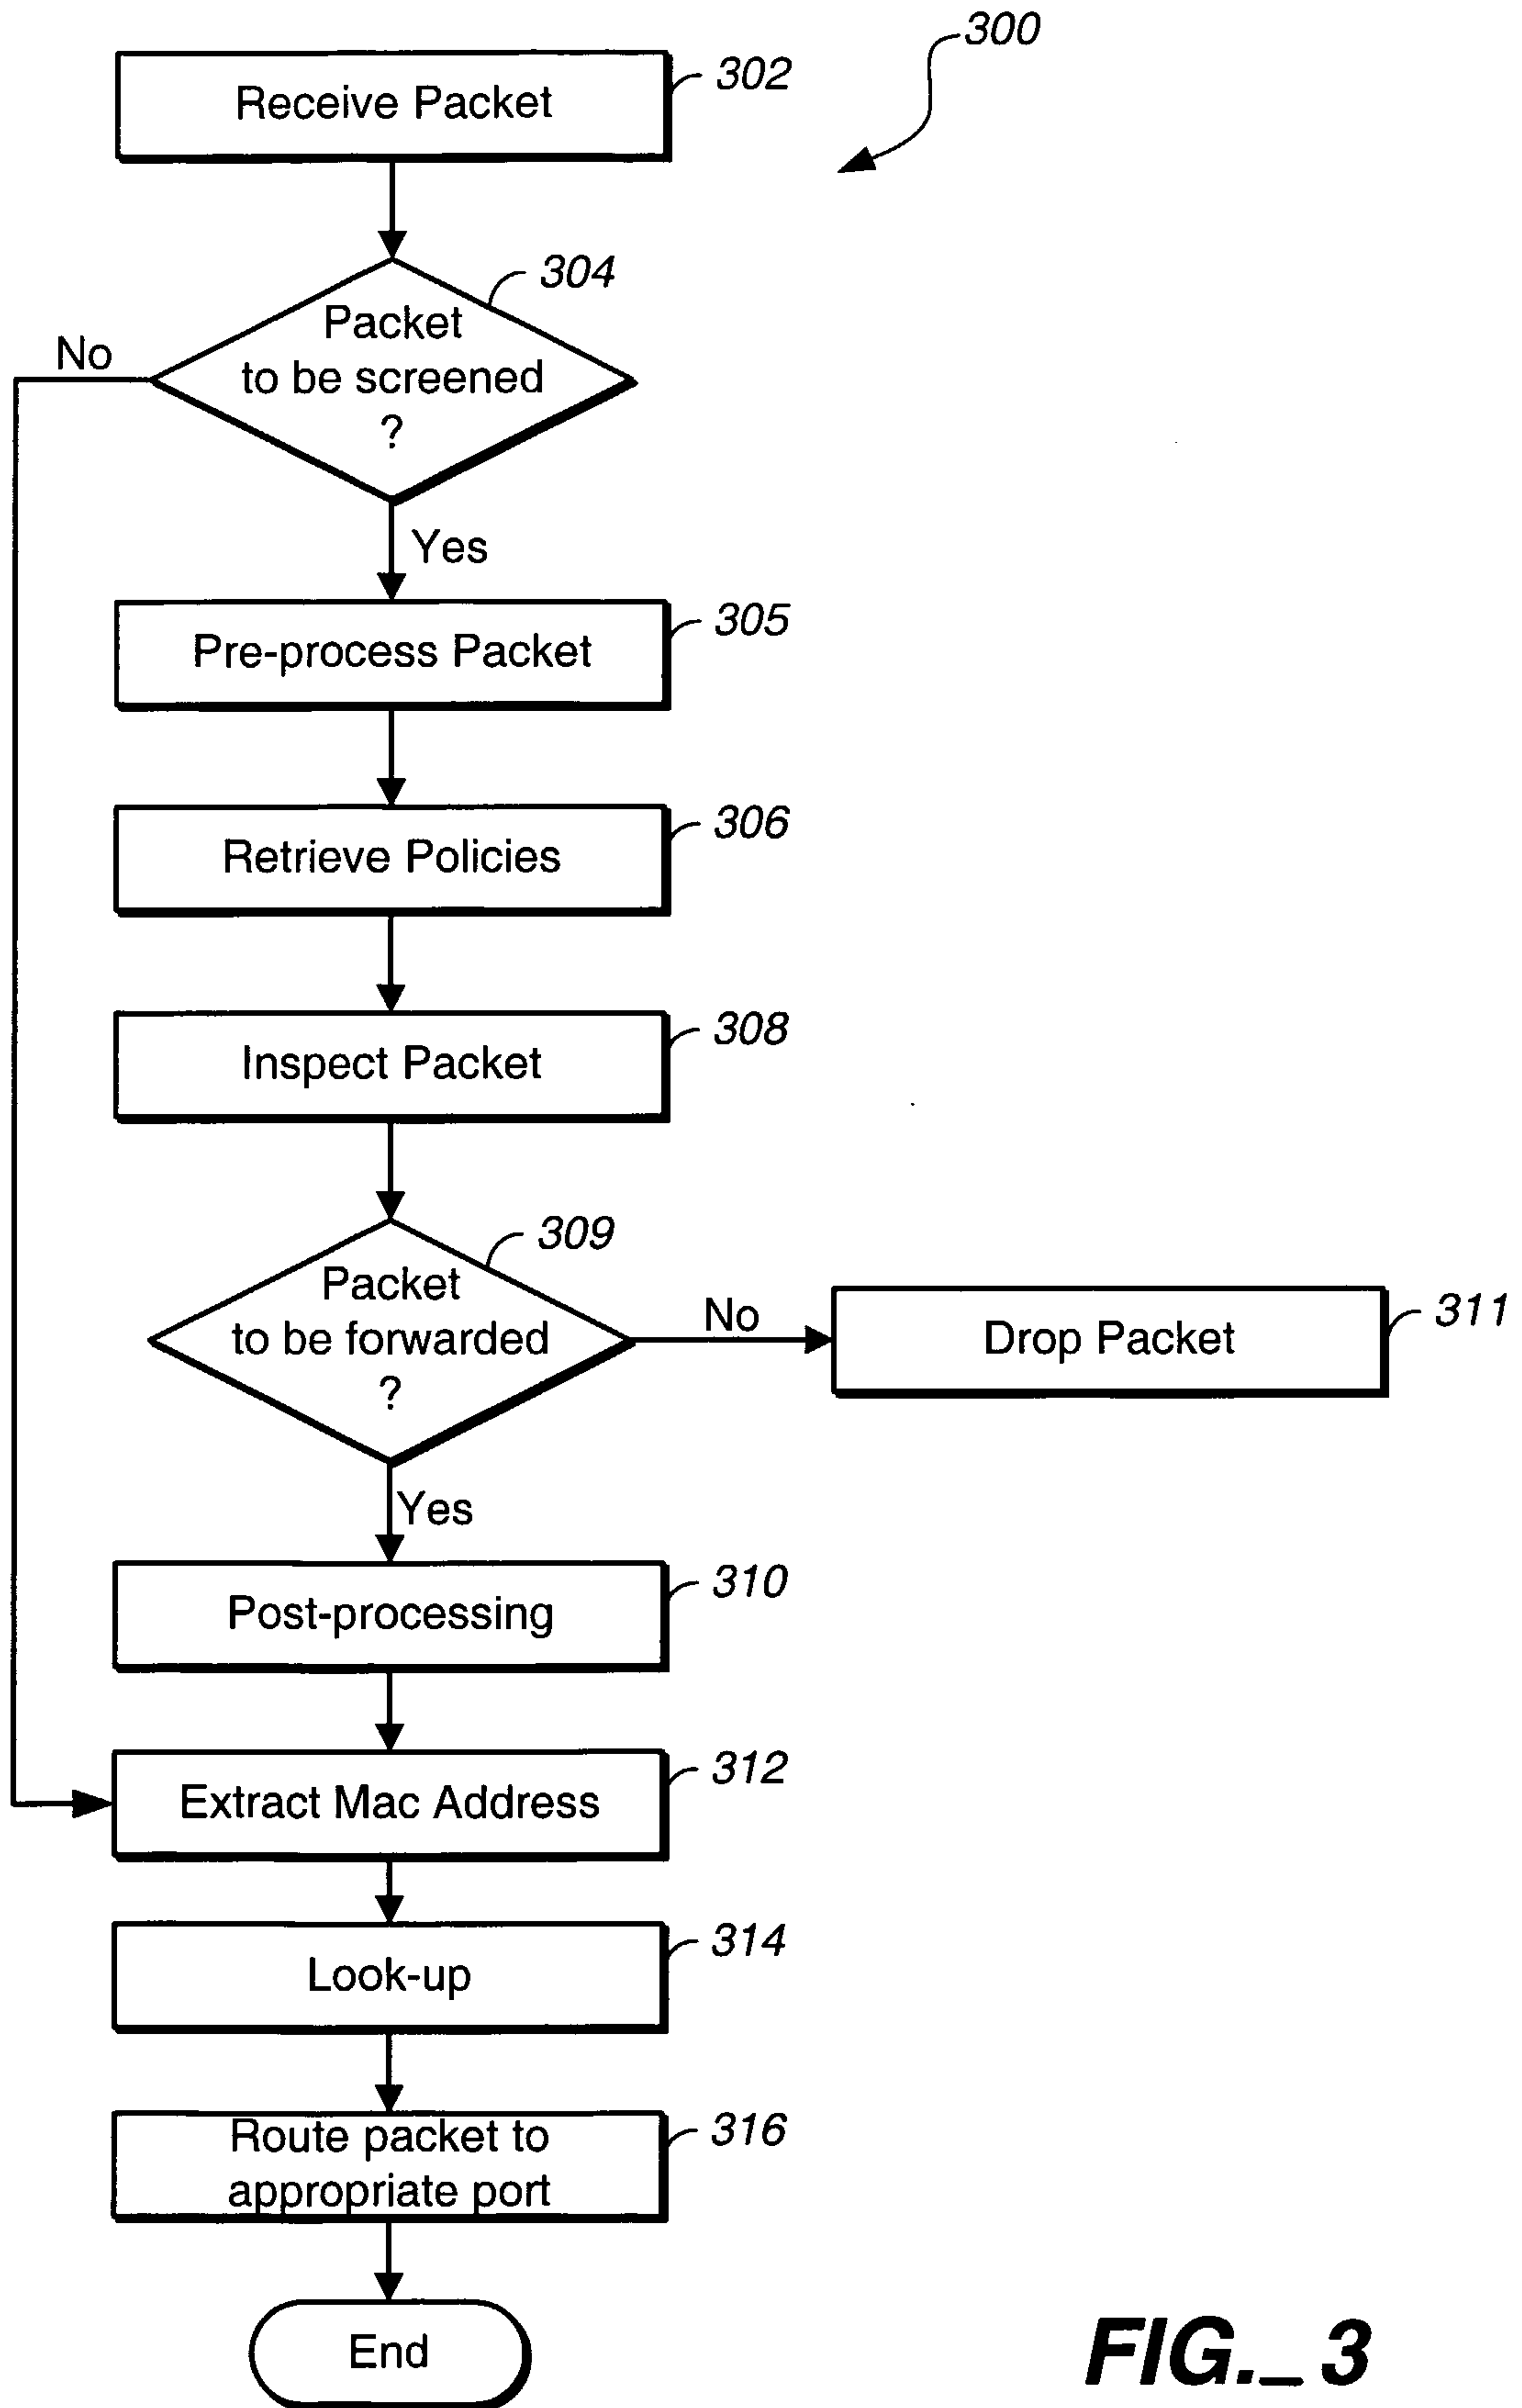
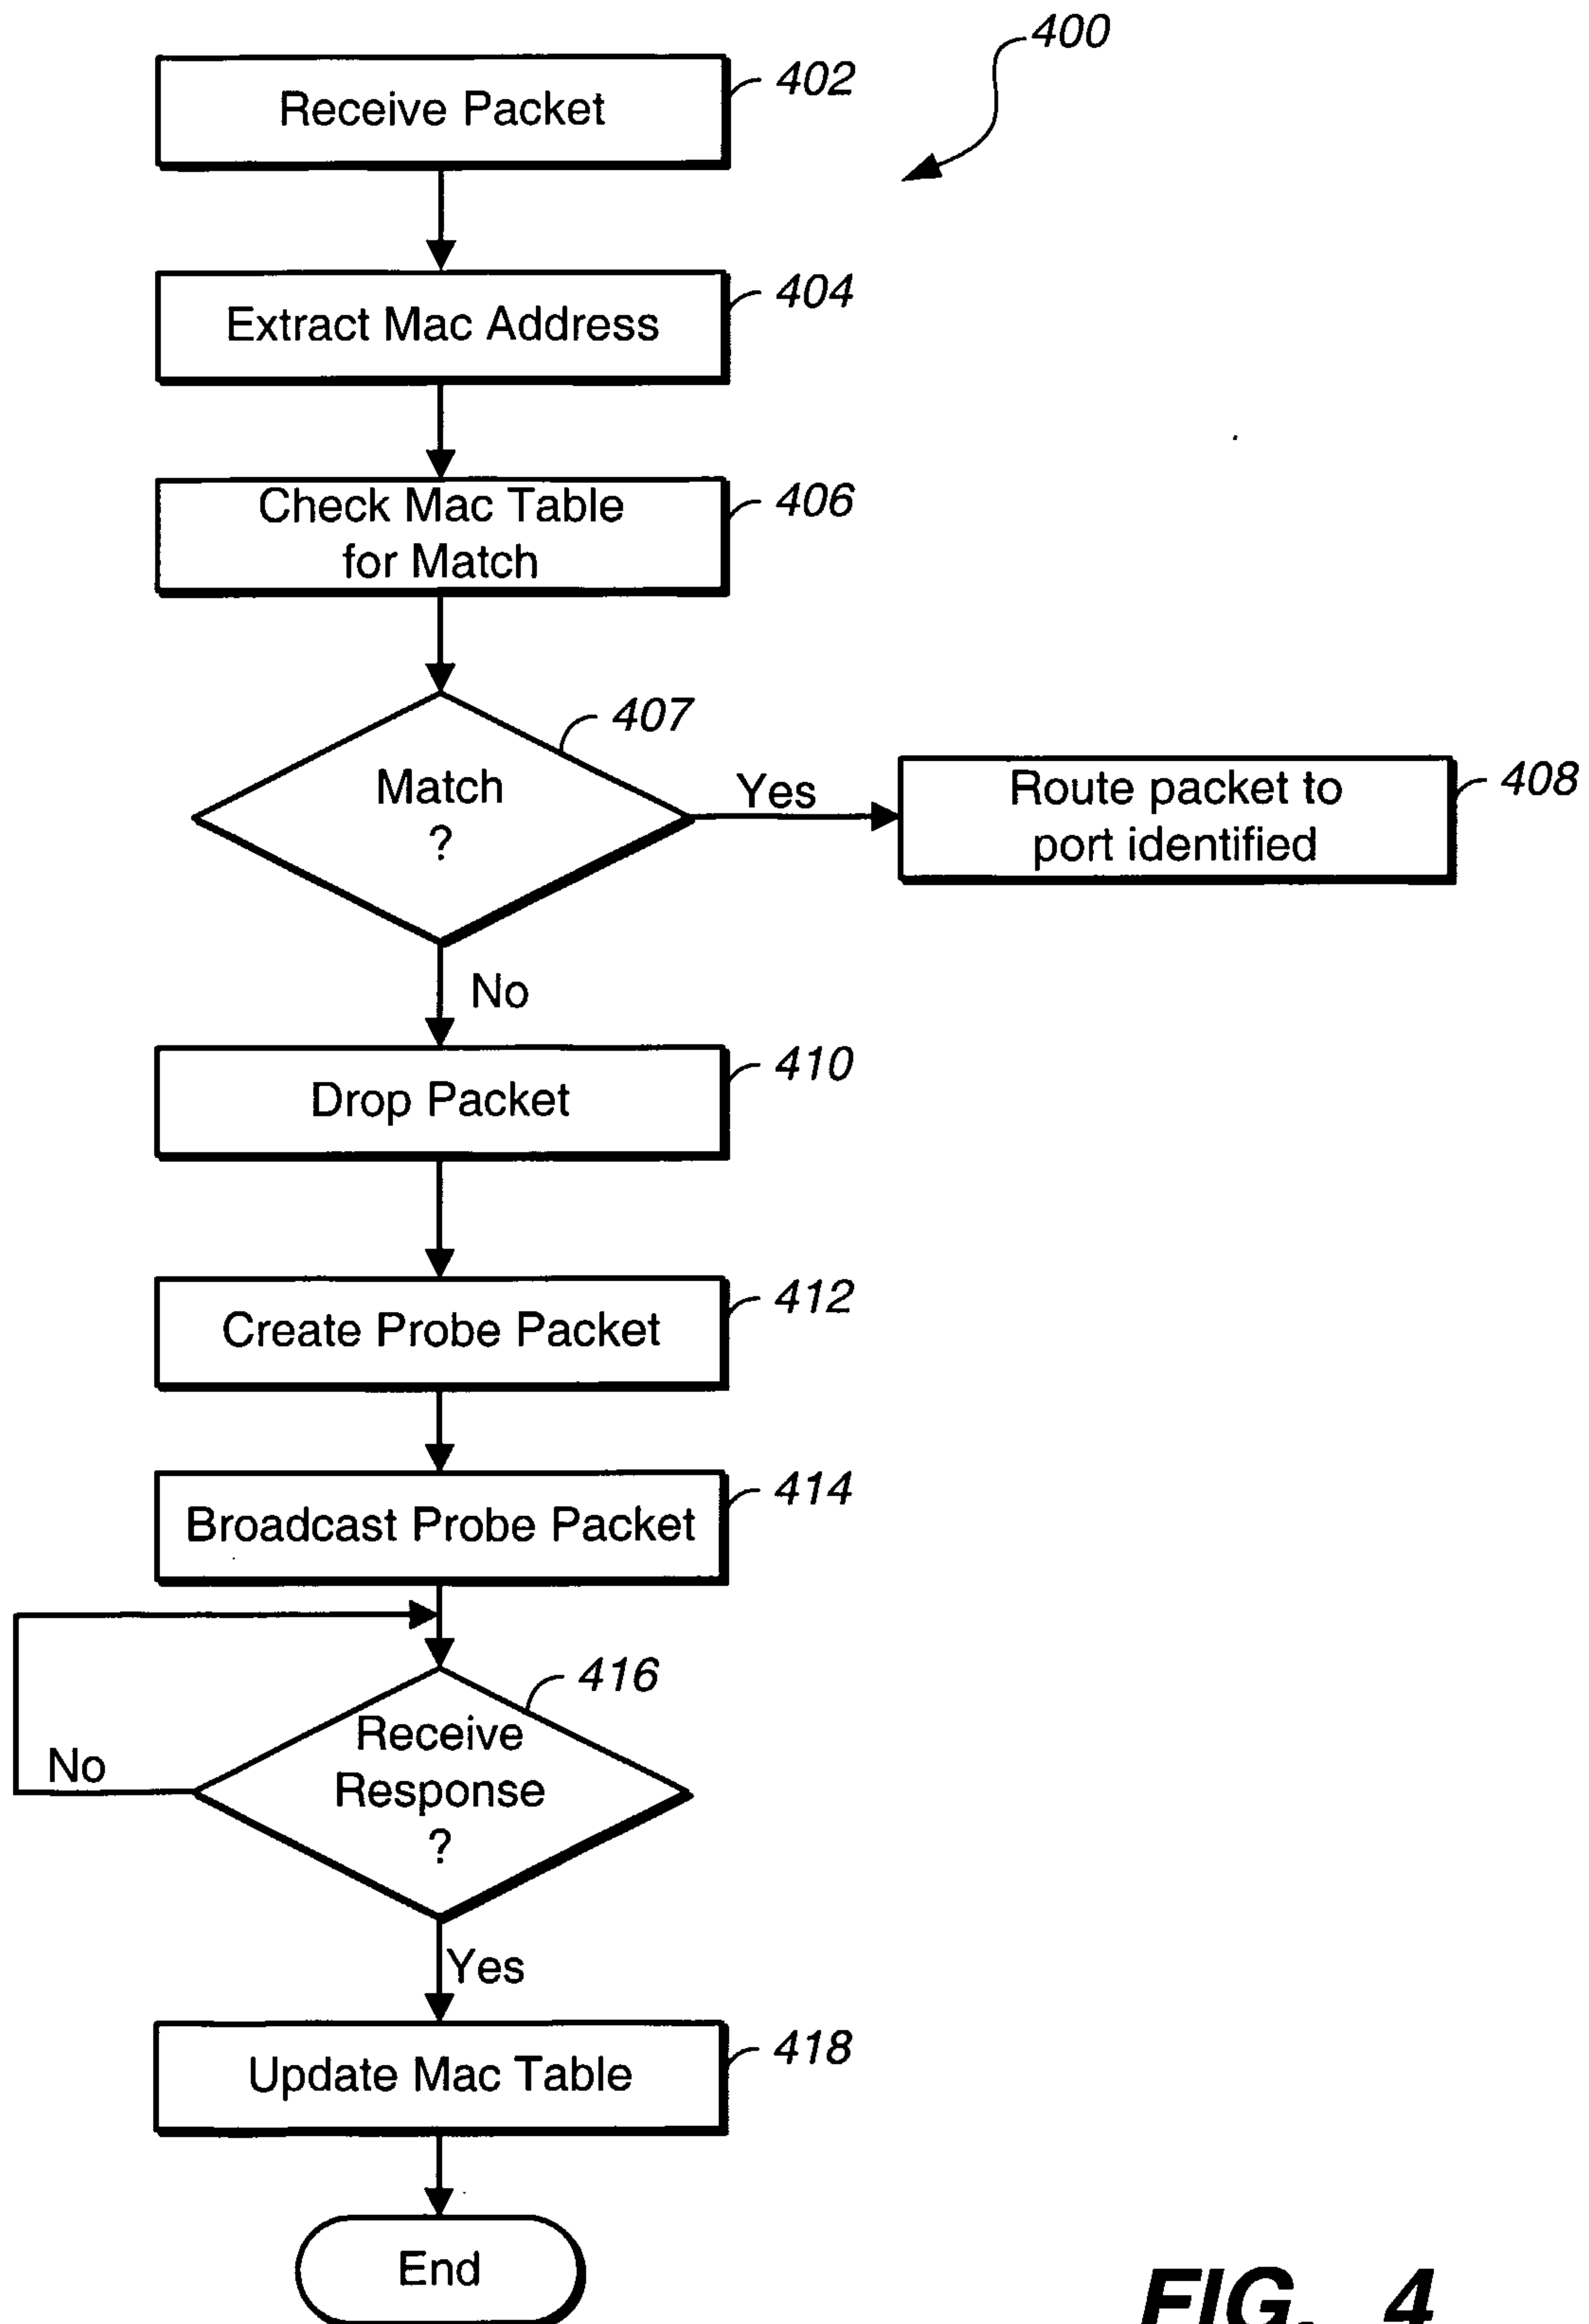


FIG. 3

5 / 5

**FIG. 4**

