

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
21 February 2002 (21.02.2002)

PCT

(10) International Publication Number
WO 02/15514 A2

(51) International Patent Classification⁷: **H04L 29/00**

(21) International Application Number: PCT/US01/25277

(22) International Filing Date: 10 August 2001 (10.08.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/638,351 15 August 2000 (15.08.2000) US

(71) Applicant: **CYBER IQ SYSTEMS** [US/US]; 225 Bay-pointe Parkway, San Jose, CA 95134 (US).

(72) Inventors: **BOMMAREDDY, Satish**; 5843 Comanche Drive, San Jose, CA 95123 (US). **KALE, Makarand**; 1235 Wildwood Avenue, #61, Sunnyvale, CA 94089 (US). **CHAGANTY, Srinivas**; 2180 Bellington Court, San Jose, CA 95138 (US).

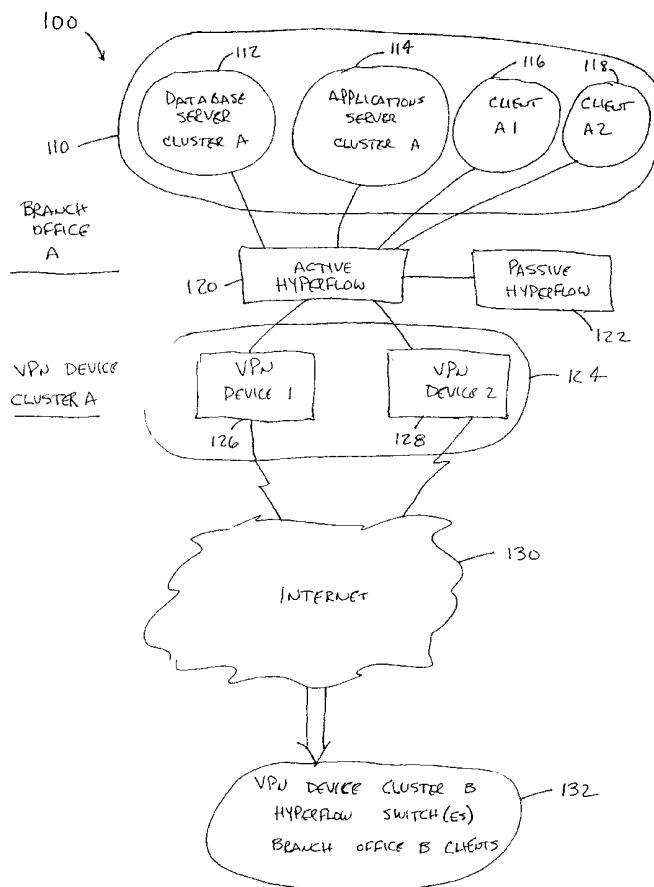
(74) Agents: **MACPHERSON, Alan, H.** et al.; Skjerven Morrill MacPherson LLP, 25 Metro Drive, Suite 700, San Jose, CA 95110 (US).

(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, EC, 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, UZ, VN, YU, ZA, ZW.

(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: VPN DEVICE CLUSTERING USING A NETWORK FLOW SWITCH



(57) Abstract: A VPN device clustering system connects two or more VPN devices on one side of a virtual private network to a similarly clustered system of two or more VPN devices on the other side of a virtual private network. The VPN device clustering system typically includes a plurality of clustering units for redundancy that avoids difficulties that arise with a single point of failure. For example two clustering units may be used in an active-passive high-availability configuration. A VPN device cluster creator creates or configures a VPN device cluster. To create a VPN device cluster, an administrator assigns to the cluster a logical Internet protocol (IP) address IP_{vpn} and specifies VPN devices that are members of the cluster.

WO 02/15514 A2



Published:

— without international search report and to be republished upon receipt of that report

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.

VPN DEVICE CLUSTERING USING A NETWORK FLOW SWITCH

BACKGROUND OF THE INVENTION5 Field of the Invention

The present invention relates generally to computer networks and more specifically, to virtual private networks.

Description of the Related Art

10 Computer networking is a widespread and constantly expanding approach to the sharing of data and software among users with a common interest in such resources. Virtually every business, governmental, or other organization with more than a very few computers has those computers networked so that individual workstations can share the resources of one or more common processors or servers. Within a single building or a 15 relatively small geographic area, the network computers can be connected through some form of Local Area Network (LAN).

There is an increasing need for remote access capability between computers and computer networks over larger and larger geographic areas. It is essential for companies with branch offices to have the capability to share computer resources between offices. 20 As more and more employees do substantial work from home, or as they travel away from company offices, there is a need to provide them with access to the company's computer network with minimal inconvenience while still providing security for data access and transfer. Companies may be in partnership with other companies where there is a desire to share at least some computer resources. It may be expensive, difficult, and 25 perhaps even impossible to network such far-flung computers using traditional approaches.

One solution to the problem of interconnecting remote computers is the use of owned or leased telecommunications lines dedicated to the sole use of a single company to service its remote computing sites. This technique, called a Wide Area Network 30 (WAN), can be expensive depending upon how far and how extensively the lines need to run, and is wasteful of resources since the telecommunications lines may have relatively limited use or, correlatively, substantial unused capacity. In addition, there may be considerable organizational overhead associated with the establishment, expansion, maintenance, and administration of the WAN.

The concept of a virtual private network (VPN) has been developed to satisfy the need for lower cost, efficient networking of dispersed computers. A virtual private network is a private data network that makes use of the public telecommunications infrastructure, maintaining privacy through the use of a tunneling protocol and security

5 procedure. VPNs extend the corporate network out to distant offices, home workers, salespeople, and business partners. VPNs use worldwide IP network services, including the Internet service provider's backbones. Remote users can make a local Internet call instead of dialing in at long distance rates. Alternatively, other types of public network connections can be used, such as a frame relay.

10 One of the keys to a VPN system is its ability to "tunnel" through public telecommunications lines so that data or applications are passed only between authorized users. Tunnels are virtual point-to-point connections that offer authentication, encryption, and access control between tunnel endpoints. Tunnels can exist at several protocol layers. Also called "encapsulation," tunneling or "IP Tunneling" encloses one

15 type of data packet into the packet of another protocol, usually TCP/IP. With VPN tunneling, before encapsulation takes place, the packets are encrypted so the data is unreadable to outsiders. The encapsulated packets travel through the internet until they reach their intended destination, then they are separated and returned to their original format. Authentication technology is employed to make sure the client has authorization

20 to contact the server.

VPNs may be either hardware or software based. A hardware based system consists of a dedicated processor running any of a number of commercially available or proprietary VPN software packages that perform the necessary VPN functions, such as encryption/decryption and authentication. Hardware based systems are most appropriate

25 for larger firms because they offer tighter security, and the ability to handle larger volumes of traffic with a dedicated VPN processor. To process even larger volumes of traffic, with greater speed, scalability, redundancy, and reliability, large VPN users can employ multiple VPN devices.

30 SUMMARY OF THE INVENTION

The present invention provides a VPN network flow switch and a method of operation thereof for connecting two or more VPN devices on one side of a virtual private network (VPN) to the authorized servers or users at that network site. A similar clustering arrangement is provided on the other side of the VPN. The clustered VPN

devices share a single IP address, without requiring translation of the IP address, and providing bi-directional clustering. The clustering unit, by operating transparently at the ISO layers 2 and 3, enables cross-platform clustering of VPN devices. This means the VPN devices within any single cluster can come from any manufacturer of such hardware or software.

5

The VPN device clustering system typically includes a plurality of clustering units for redundancy to avoid difficulties that arise with a single point of failure. For example, two clustering units may be used in an active-passive high-availability configuration.

The clustering system operates on outgoing data packets before they go through 10 the transmitting VPN device. Similarly, the clustering system operates on incoming data packets after processing by the VPN device. Thus, the VPN device clustering system operates in a manner that is independent of the VPN hardware and software. The clustering system can therefore operate with any VPN hardware or software configuration without affecting the VPN authentication, security, or “tunneling” functions.

15 In some embodiments, the VPN network flow switch, in addition to routing of the packets, performs load balancing and fault tolerance functions. In these embodiments, a processor of the VPN network flow switch periodically executes a load balancing routine to determine the relative workload of each of the VPN devices. When the VPN network flow switch receives a packet destined to the cluster of VPN devices, the packet is routed 20 to the VPN device with an optimal workload, so as to ensure that the workload is evenly distributed among the VPN devices. In addition, if a failure of a VPN device is detected, a packet addressed to that VPN device is re-routed to a different VPN device by re-writing the Data Link Layer (MAC) destination address of the packet. Since the VPN network flow switch continuously monitors the status of the VPN devices, no lengthy 25 time delay is introduced in point-to-point communications when a VPN device is disabled.

Since the cluster IP header is not modified, the VPN network flow switch of the present invention operates on packets encoded according to any VPN protocol. In addition, the VPN network flow switch can handle re-routing, load balancing and fault tolerance of encrypted packets transparently to users on both sides of the VPN.

30

BRIEF DESCRIPTION OF THE DRAWINGS

The features of the described embodiments believed to be novel are specifically set forth in the appended claims. However, embodiments of the invention relating to both

structure and method of operation, may best be understood by referring to the following description and accompanying drawings.

FIGURE 1 is a schematic block diagram that illustrates an embodiment of a VPN device clustering system connecting two or more VPN devices on one side of a virtual private network to a similar VPN device clustering system on the other side of the virtual private network.

FIGURE 2 is a schematic flow chart that depicts operations of a VPN device cluster creator.

FIGURE 3 is a schematic flow diagram showing operations of a traffic 10 distributor.

FIGURE 4 is a schematic block diagram and associated transition tables that illustrate a technique for transferring a packet between two authorized users with a VPN device clustering system.

FIGURE 5 is a flow diagram that illustrates a further implementation of a traffic 15 distribution method.

DETAILED DESCRIPTION OF THE INVENTION

Referring to **FIGURE 1**, a schematic block diagram illustrates an embodiment of a VPN device clustering system **100** that connects two or more VPN devices, for example 20 VPN device1 **126** and VPN device2 **128**, to the Internet **130** in an arrangement with complete high-availability, scalability, and traffic distribution. Although not illustrated in detail, it is to be understood that on the other side of the Internet **130** from the 25 components shown in detail in **Fig. 1** is located a similar configuration of VPN devices, clustering units, and peer-to-peer devices **132**. In the illustrative VPN device clustering system **100**, a network flow controller **120**, or “hyperflow,” includes a processor (not shown) and storage (not shown) that execute special-purpose software for control of the network flow controller **120**. The network flow controller **120** arranges the two or more 30 VPN devices, VPN device1 **126** and VPN device2 **128**, in a VPN device cluster **124** to connect to the other end of a VPN “tunnel” through the Internet **130**. The network flow controller **120** also arranges the branch office servers and other client devices **112, 114, 116, 118** in a cluster **110** on one side of the VPN tunnel for secure communication with similar peer devices at the other end of the VPN tunnel.

The VPN device clustering system **100** includes a plurality of clustering units, **120** and **122**, which operate as an active flow controller **120** and a passive flow controller **122**,

for redundancy that avoids difficulties that arise with a single point of failure. The network flow controller 120 and the passive flow controller 122 are used in an active-passive high-availability configuration.

Outgoing traffic from the branch office (or client) A cluster 110 that is destined 5 for branch office (or client) B cluster devices (included in 132) on the Internet 130 is distributed among the two or more VPN devices, VPN device1 126 and VPN device2 128. The VPN device clustering system 100 distributes traffic based on the destination cluster IP address of the packet, supporting all IP-based protocols. A single cluster IP address is assigned to all VPN devices in the respective (client A or client B) clusters.

10 Additional VPN devices may be seamlessly added to the cluster 124 to supply additional bandwidth and greater fault tolerance.

The network flow controller 120 operates independently of the hardware and software that are arranged in the VPN device clusters. For example, various combinations of VPN devices can exist in the cluster 124 as long as the VPN devices 15 have the same connectivity.

The VPN device clustering system 100 includes multiple control processes that execute on the network flow controller 120 and the passive flow controller 122. One control process is a VPN device cluster creator that creates or configures the VPN device cluster 124.

20 **FIGURE 2**, best understood in conjunction with **FIGURE 1**, is a schematic flow chart depicting operations of a VPN device cluster creator software routine 200. To create the VPN device cluster 124, in a *cluster IP and VPN device assignment operation 210*, an administrator assigns to the cluster a logical Internet protocol (IP) address IP_{vpn} and specifies VPN devices, for example VPN device1 126 and VPN device2 128, that are 25 members of the VPN device cluster 124. In a *begin monitoring VPN device health operation 212*, the network flow controller 120 begins to monitor health of the VPN devices, VPN device1 126 and VPN device2 128, typically using a health check operation at a configured polling interval. In a *configure VPN device cluster address operation 214*, the logical cluster address IP_{vpn} is configured on the client A devices 110.

30 In a *respond to ARP request operation 216*, the network flow controller 120 responds to an Address Resolution Protocol (ARP) request from the servers in the client A device cluster 110 to identify a Media Access Control (MAC) address associated with the VPN device cluster 124. Associating the MAC address with the VPN device cluster 124 ensures that the client A devices 110 send all outbound traffic to the VPN device

cluster 124 for forwarding on to the corresponding VPN device cluster on the Internet 130.

Another control process is a traffic distributor that distributes outbound traffic destined for the Internet 130 among the VPN devices, for example, VPN device1 126 and 5 VPN device2 128. Referring to **FIGURE 3** in combination with **FIGURE 1**, a schematic flow diagram shows operations of a traffic distributor 300. The traffic distributor executes from the network flow controller 120. The traffic distributor 300, in a *select VPN device for outbound traffic operation 310*, determines which VPN device is to forward the outbound traffic based on the packet destination IP address, which will be the 10 cluster IP address of the corresponding VPN device cluster at the receiving end of the VPN “tunnel”. Usage of the destination cluster IP address ensures that, for a given flow designating a particular VPN tunnel connection, the same VPN device is used for every outbound packet so long as the VPN device remains operational. Since flow is based on the destination cluster IP address, measurement and analysis operations by the network 15 flow controller 120 are reduced since measurements of parameters such as load on the VPN device is not necessary. Accordingly, VPN device load sharing is on a probabilistic or statistical basis, which may result in slightly unbalanced loading. The probabilistic loading presumes that VPN devices in the VPN device cluster 124 have similar forwarding power.

20 Internally, in a *maintain list of operational VPN devices operation 312* the traffic distributor 300 maintains a list of operational VPN devices. Fields from the packet are used to compute the index into this list, identifying the active VPN devices. In a distribute traffic to appropriate VPN device cluster operation 314, traffic is directed to the appropriate peer VPN cluster based on that cluster's assigned IP address.

25 The network flow controller 120 has a particular MAC address that identifies the traffic distributor. The traffic distributor replaces the packet destination MAC address, which previous to replacement is the MAC address of the traffic distributor, with the MAC address of the VPN device handling the flow.

30 Each VPN device, VPN device1 126 or VPN device2 128, has an equal probability of assignment for an outbound flow forwarding since the traffic distributor uses only information in the packet IP header to select between VPN devices. Processing load or potential processing power of the VPN device is not analyzed as part of the selection process.

The VPN device cluster 124 does not affect the processing performed by the network flow controller 120 for inbound traffic coming from the Internet 130. Traffic destined for any VPN device cluster 124 continues to be distributed among the client A operational servers and other devices 110 defined in the VPN device cluster 124. At most 5 only a single VPN device cluster 124 is supported for inbound traffic.

Another control process is a VPN device monitor that monitors “health” of the VPN devices. In some implementations, the VPN device clustering system 100 monitors VPN device health using a configured polling interval and health check method. The health probe authenticates connectivity of a flow across a VPN. In one example the 10 network flow controller 120 periodically sends a Ping packet to VPN device1 126, using ICMP extension to confirm that the flow is operative. VPN device1 126 responds on the same port since there is a one-to-one correlation between VPN devices and individual ports of the network flow controller 120.

In some implementations, the VPN device clustering system 100 continually 15 monitors the operational health of the VPN devices and associated wide area network (WAN) links.

In some implementations, the VPN device clustering system 100 detects one or more of various failure conditions. Failures can occur in the VPN device to LAN 20 interface and link, or in the VPN device itself due to power outage, software malfunction, hardware malfunction, or other condition. Failures also can occur in the VPN device to WAN interface and link. When the VPN device clustering system 100 detects a failure, traffic is automatically forwarded to the remaining operational VPN device or devices. The VPN device clustering system does not require manual intervention at the client A servers to bypass the failed VPN devices.

25 Referring to **FIGURE 4**, a schematic block diagram and associated transition tables depicts a technique for transferring a packet between a client A device 410 and a client B device (not shown) that is assigned to use VPN device1 414 by the VPN device clustering system. The outbound traffic 416 has a destination MAC address designating the MAC address of the traffic distributor, but has a destination IP address that designates 30 neither the traffic distributor nor any cluster supported by the traffic distributor. VPN device cluster traffic has no unique attribute other than destination VPN cluster IP address so that designation of the destination IP address effectively limits the current traffic distributor to support only a single VPN device cluster 418. Although only a single VPN device cluster 418 is included in the VPN device clustering system 400, the

VPN device cluster 418 typically includes a plurality of VPN devices, here shown as VPN device1 414 and VPN device2 415.

The limitation to a single VPN device cluster 418 further extends to limit the VPN device clustering system 400 to a single cluster that performs a routing function.

5 Other implementations of a VPN device clustering system that supports multiple MAC addresses can support additional VPN device clusters.

The cluster IP address of the VPN device cluster 418 does not appear in the packet since the VPN device cluster 418 is only a gateway on the path to an actual end destination.

10 A network flow controller 420 uses ARP probe methods to monitor the VPN devices in the VPN device cluster 418. Software executing in the network flow controller 420 uses the Address Resolution Protocol (ARP) to probe for an unused IP address. If a VPN device responds to the ARP probe, the software tries the next IP address. If, after several tries, no response is elicited from an ARP probe, software uses that address as the 15 IP address of the VPN device.

Referring to **FIGURE 5**, a flow diagram illustrates a traffic distribution method 500. In a *check destination IP address operation 510*, a traffic distributor checks the destination IP address of a packet to determine whether the destination IP address is a cluster address.

20 If the *check destination IP address operation 510* determines that the destination IP address is not a cluster address then, in a *test destination MAC address operation 512*, the traffic distributor checks to determine whether the destination MAC address is a cluster address. The destination MAC address matches the cluster address when a Proxy ARP is used to indicate to attached VPN devices that the MAC address of the network 25 flow controller is used when sending packets to any of the configured cluster IP addresses.

If the *test destination MAC address operation 512* determines that the MAC address is not a cluster address then, in a *VPN health test operation 514*, the traffic distributor performs a performance test on the VPN devices in the cluster.

30 A first redirection operation is a *set VPN device cluster identifier operation 516* in which the cluster address in the form of either the MAC address or the destination IP address is set to identify the cluster data structure. A *bucket check operation 518* determines whether at least one bucket exists in a cluster data structure. If not, one is created in a *create bucket operation 520*. A *load balancing operation 522* retrieves an

appropriate bucket that attains load balancing.

A *flow test operation 524* determines whether the flow is assigned to the bucket and, if not, performs a *flow assignment operation 526* that assigns buckets to a server.

The traffic distributor *executes a forward packet to assigned VPN device cluster member 532* with the buckets used to forward data requests from client A to client B.

Further details of a traffic distribution and load balancing system are disclosed and claimed in co-pending application Serial Number 09/540,296, entitled "Router Clustering for Multiple Network Service", incorporated herein in full.

While the invention has been described with reference to various embodiments, it 10 will be understood that these embodiments are illustrative and that the scope of the invention is not limited to them. Many variations, modifications, additions and improvements of the embodiments described are possible. For example, those skilled in the art will readily implement the steps necessary to provide the structures and methods disclosed herein, and will understand that the process parameters, materials, and dimensions 15 are given by way of example only and can be varied to achieve the desired structure as well as modifications which are within the scope of the invention. Variations and modifications of the embodiments disclosed herein may be made based on the description set forth herein, without departing from the scope and spirit of the invention as set forth in the following claims.

20 In the claims, unless otherwise indicated the article "a" is to refer to "one or more than one".

WHAT IS CLAIMED IS:

1. A method of routing message traffic on a virtual private network (VPN) between a first plurality of users and a second plurality of users, the method comprising:
 - 5 creating a cluster containing a plurality of VPN devices, the cluster being addressed by a logical Internet protocol (IP) address that is distinct from the unique IP addresses of VPN devices contained within the cluster; and distributing traffic between the first plurality of users and the second plurality of users via a VPN device selected from among the VPN devices contained in the cluster, the VPN device being selected on the basis of a packet destination IP address.
- 10 2. A method according to Claim 1 wherein the creating operation further comprises:
 - assigning a single MAC address to the VPN device cluster.
3. A method according to Claim 1 wherein the creating operation further comprises:
 - 15 assigning a unique MAC address to each VPN device cluster of a plurality of VPN device clusters.
4. A method according to Claim 1 wherein the creating operation further comprises:
 - assigning a logical IP address to the VPN device cluster.
- 20 5. A method according to Claim 1 wherein the creating operation further comprises:
 - assigning, by an administrator, a logical IP address to the VPN device cluster.
6. A method according to Claim 1 wherein the creating operation further comprises:
 - 25 specifying VPN devices that are contained within the VPN device cluster.
7. A method according to Claim 1 further comprising:
 - monitoring the operational health of the VPN devices.

8. A method according to Claim 1 wherein the distributing traffic operation further comprises:

5 selecting a VPN device from among the plurality of VPN devices contained within the cluster for distributing outbound traffic from one VPN user to another VPN user based on the packet IP destination address.

9. A method according to Claim 1 wherein the distributing traffic operation further comprises:

10 selecting a VPN device from among the plurality of VPN devices contained within the cluster for distributing outbound traffic from one VPN user to another VPN user so that for any given VPN user-to-user connection flow the same VPN device is used for every outbound packet so long as the flow remains operational.

11. A method according to Claim 1 wherein the distributing traffic operation further comprises:

15 selecting a VPN device from among the plurality of VPN devices contained within the cluster for distributing outbound traffic from VPN user to user so that the probability of any particular VPN device being used for forwarding is the same.

20 11. A method according to Claim 1 further comprising:
 maintaining a list of operational flows.

12. A computer readable storage medium for execution on a processor for routing message traffic on a virtual private network (VPN) between a first plurality of users and a second plurality of users via VPN devices, the medium comprising:

25 an encoding defining:
 a cluster containing a plurality of VPN devices, the cluster being addressed by a logical Internet protocol (IP) address that is distinct from the unique IP addresses of VPN devices contained within the cluster;
 a list of the VPN devices contained within the cluster; and
 a redirecting VPN device for redirecting traffic when an active VPN
30 device fails.

13. A computer readable storage medium according to Claim 12 further comprising:
 - a cluster type designator that classifies a cluster type.
14. A computer readable storage medium according to Claim 12 further comprising:
 - a cluster type designator that classifies a cluster type into a VPN device cluster type and a firewall cluster type.
15. A computer readable storage medium according to Claim 12 further comprising:
 - 10 a traffic distributor that distributes traffic between a user of the first plurality of VPN users and a user of the second plurality of VPN users via a VPN device selected from among the VPN devices contained in the cluster, the VPN device being selected on the basis of a packet destination IP address.
16. A computer readable storage medium according to Claim 12 further comprising:
 - 15 a single MAC address assigned to the cluster.
17. A computer readable storage medium according to Claim 12 further comprising:
 - 20 a unique MAC address assigned to each cluster of a plurality of clusters.
18. A computer readable storage medium according to Claim 12 further comprising:
 - an operational health probe manager that determines health of the VPN devices in the cluster.

1/4

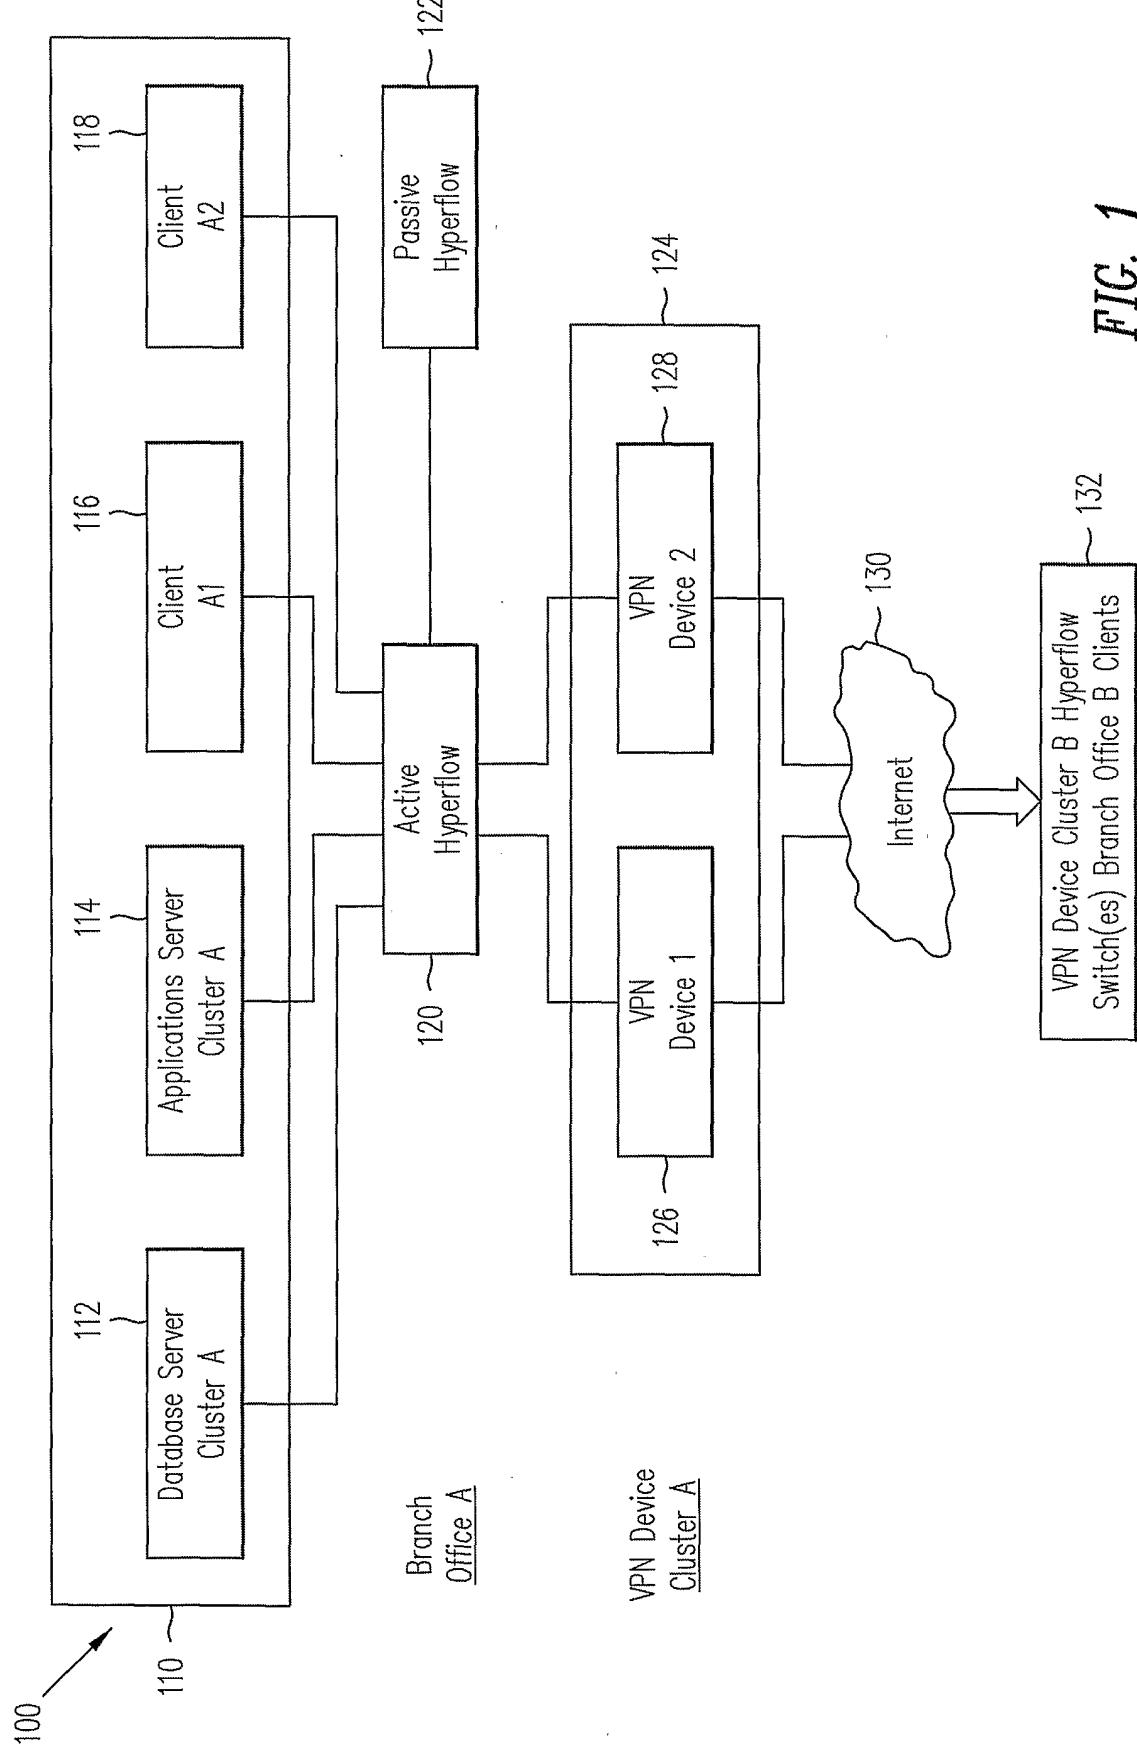
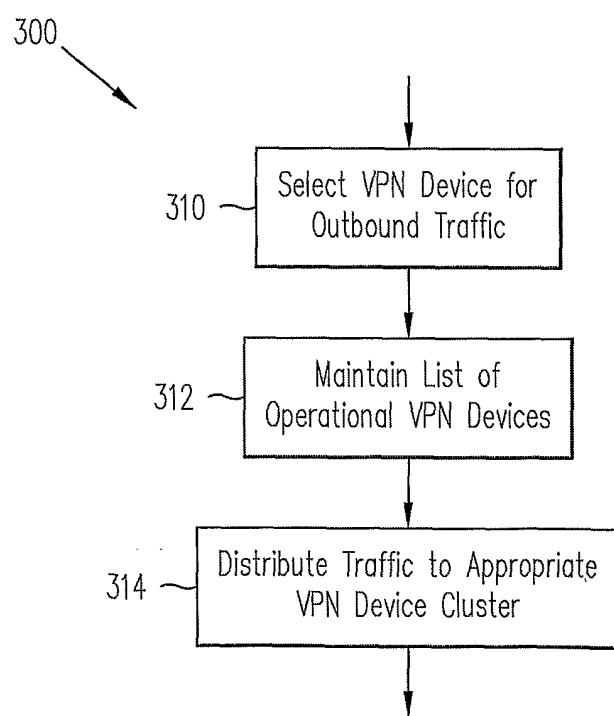
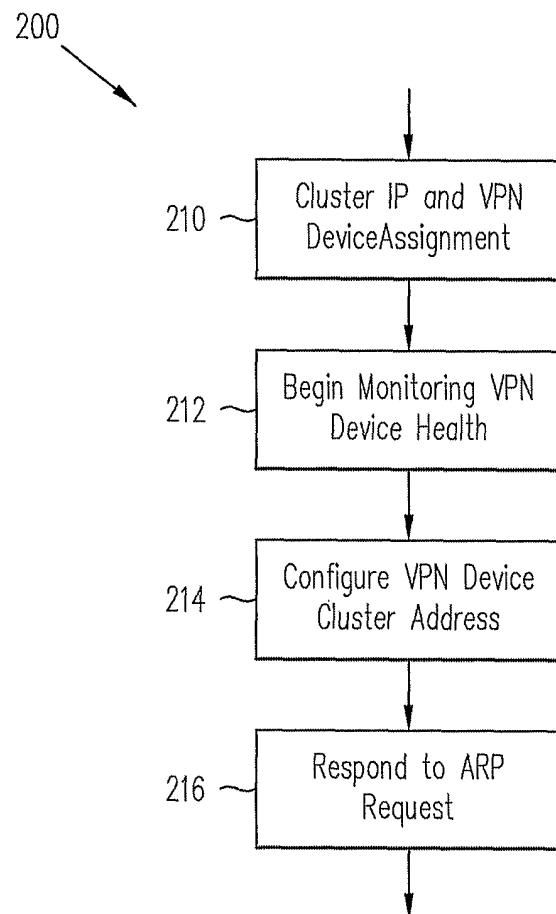


FIG. 1

2/4



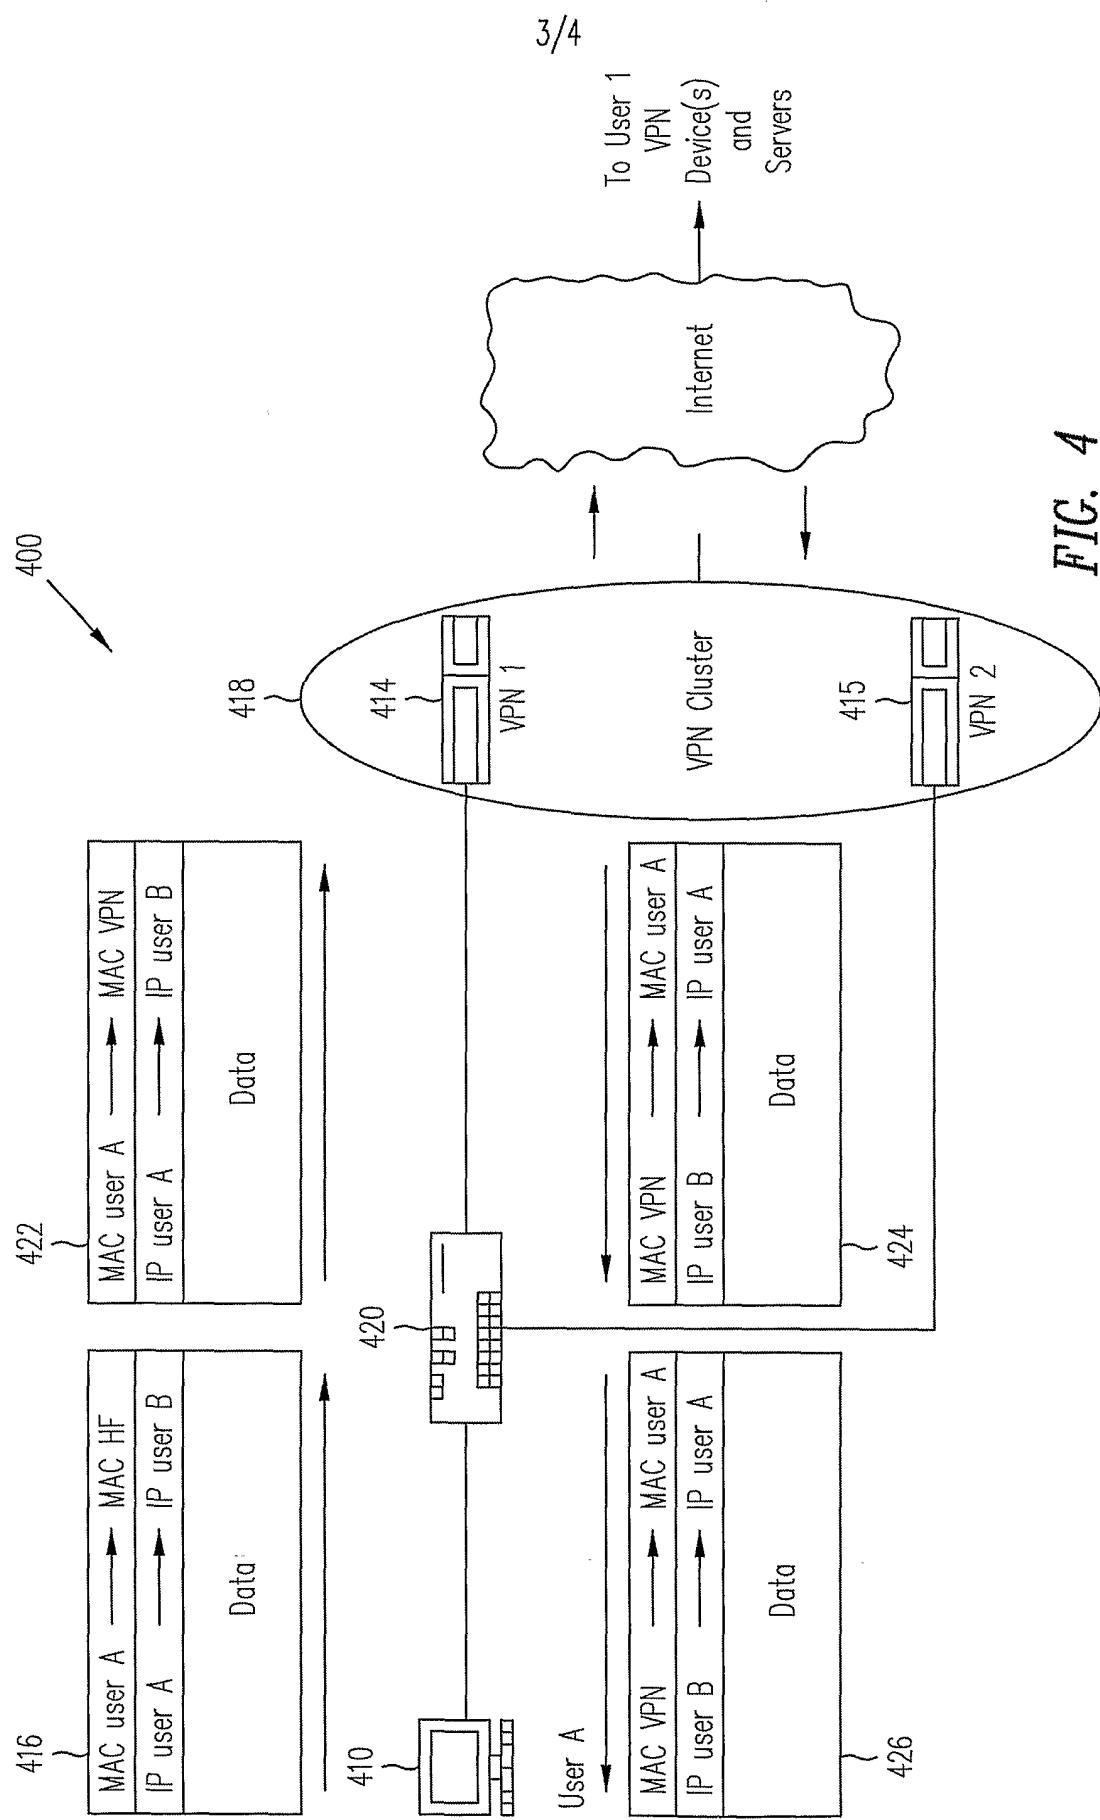


FIG. 4

4/4

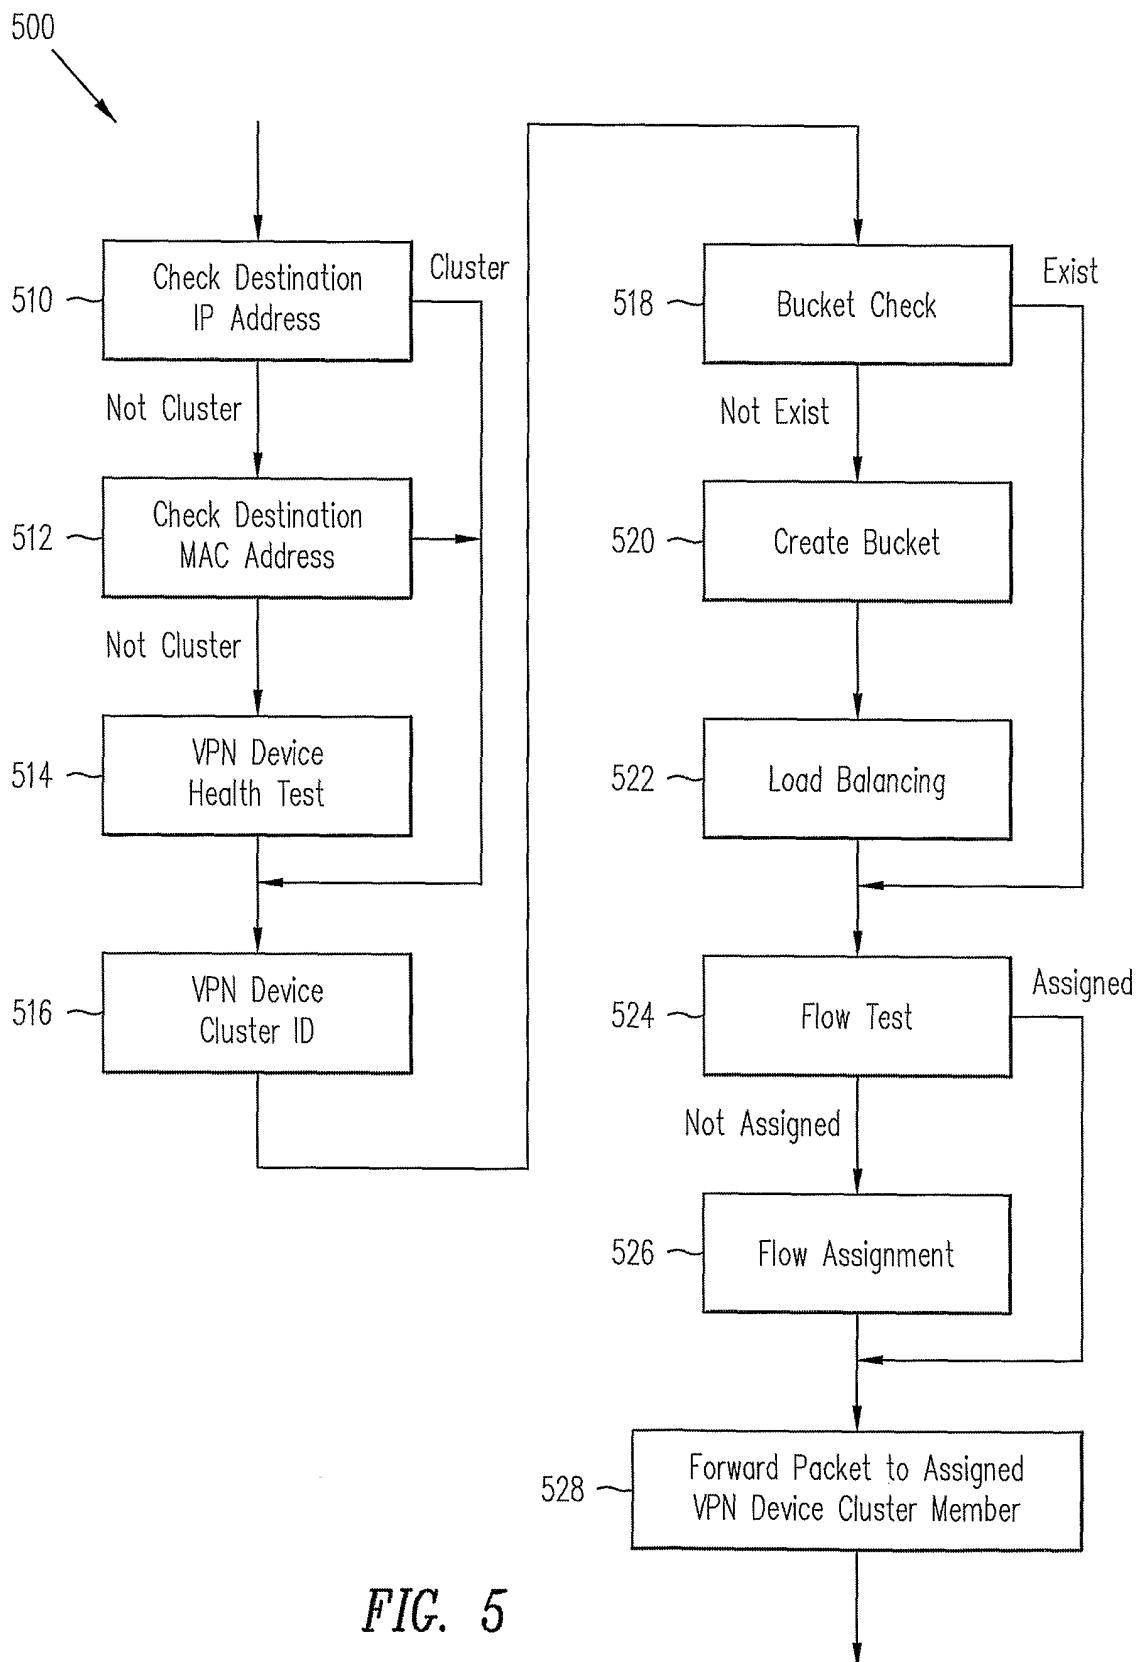


FIG. 5