EX 1002

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

FORTINET, INC.,
Petitioner,

v.

NETSKOPE, INC.,
Patent Owner.

Case No. IPR2026-00025

U.S. Patent No. 8,117,639

DECLARATION OF DR. KEVIN C. ALMEROTH IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW U.S. PATENT NO. 8,117,639

TABLE OF CONTENTS

I. INTRODUCTION 1

II. QUALIFICATIONS 1

III. LEGAL STANDARDS 14

A. Anticipation 14

B. Obviousness 15

C. Level of Ordinary Skill in the Art 18

D. Claim Construction 19

IV. CHALLENGED CLAIMS AND GROUNDS OF UNPATENTABILITY 19

V. TECHNOLOGY BACKGROUND 20

VI. THE '639 PATENT 23

A. Overview of the Specification 23

B. The Challenged Claims 26

C. Level Ordinary Skill in the Art 26

VII. CLAIM INTERPRETATION 27

VIII. PRIOR ART OVERVIEW 28

A. Overview of Richmond 28

B. Overview of Wood 34

C. Overview of Teraslinna 35

IX. GROUNDS 1 & 2 – CLAIMS 1-8, 10-15, 17-20, 22-25 ARE ANTICIPATED OR RENDERED OBVIOUS BY RICHMOND 38

A. Independent Claim 1 38

B. Dependent Claim 2 67

C. Dependent Claim 3 71

D. Dependent Claim 4 (Ground 2 Only) 74

E. Dependent Claim 5 77

F. Dependent Claim 6 80

G. Dependent Claim 7 81

H. Dependent Claim 8 83

I. Dependent Claim 9 85
J. Independent Claim 10 87
K. Dependent Claim 11 90
L. Dependent Claim 12 90
M. Dependent Claim 13 90
N. Dependent Claim 14 93
O. Dependent Claim 15 93
P. Dependent Claim 16 93
Q. Independent Claim 17 93
R. Dependent Claim 18 97
S. Dependent Claim 19 97
T. Dependent Claim 20 97
U. Dependent Claim 22 98
V. Dependent Claim 23 99
W. Dependent Claim 24 100
X. Dependent Claim 25 101
Y. Dependent Claim 26 101
Z. Independent Claim 27 101

X. GROUND 3 – CLAIMS 9, 16, 26, AND 27 ARE RENDERED OBVIOUS BY RICHMOND IN LIGHT OF WOOD 103

A. Dependent Claim 9 103
B. Dependent Claim 16 108
C. Dependent Claim 26 108
D. Independent Claim 27 108

XI. GROUND 4 – CLAIM 21 IS RENDERED OBVIOUS BY RICHMOND IN LIGHT OF TERASLINNA 110

A. Dependent Claim 21 110

XII. CONCLUSION 114

TABLE OF ABBREVIATIONS

Abbreviation	Full Name
'639 Patent	U.S. Patent No. 8,117,639
Challenged Claims	Claims 1-27 of U.S. Patent No. 8,117,639
Patent Owner	Patent Owner Netskope, Inc.
Petitioner	Petitioner Fortinet, Inc.
POSITA	Person of Ordinary Skill In The Art

FORTINET EX1002 Page 4

EXHIBIT LIST

Exhibit No.	Document
1001	U.S. Patent No. 8,117,639 ("the '639 Patent")
1002	Declaration of Kevin C. Almeroth.
1003	WIPO Publication WO 01/11452A2 ("Wood")
1004	U.S. Patent Publication No. 2003/0152067 ("Richmond")
1005	U.S. Patent No. 5,623,492 ("Teraslinna")
1006	Enterasys Personalized User Network Whitepaper ("Enterasys Whitepaper")
1007	Excerpts from File History of Richmond
1008	Curriculum Vitae of Kevin Almeroth
1009	Excerpts from A. Tannenbaum, *Computer Networks* (3d ed., 1996)

All emphasis appearing herein is added unless otherwise noted.

I, Kevin C. Almeroth, declare as follows:

I. INTRODUCTION

1. I have been retained on behalf of Petitioner Fortinet, Inc. for the above captioned *inter partes* review of claims 1-27 ("the Challenged Claims") of U.S. Patent No. 8,117,639 ("the '639 Patent"), which I understand is assigned to Patent Owner Netskope, Inc. I am being compensated for my time at $850 per hour. My compensation is in no way dependent upon my testimony or the outcome of the IPR.

2. I have reviewed and am familiar with the specification of the '639 Patent, which is attached as Exhibit 1001. I have also reviewed and am familiar with the file history for the '639 Patent.

3. I have been asked to provide my technical review, analysis, insights and opinions regarding the '639 Patent and the references that form the basis for the grounds of rejection set forth in the petition for *inter partes* review of the '639 Patent.

II. QUALIFICATIONS

4. I am currently a Professor Emeritus in the Department of Computer Science at the University of California, Santa Barbara (UCSB). While at UCSB, I held faculty appointments and was a founding member of the Computer Engineering (CE) Program, Media Arts and Technology (MAT) Program, and the Technology Management Program (TMP). I also served as the Associate Director of the Center for Information Technology and Society (CITS) from 1999 to 2012. I have been a

faculty member at UCSB since July 1997.

5. I hold three degrees from the Georgia Institute of Technology: (1) a Bachelor of Science degree in Information and Computer Science (with minors in Economics, Technical Communication, and American Literature) earned in June 1992; (2) a Master of Science degree in Computer Science (with specialization in Networking and Systems) earned in June 1994; and (3) a Doctor of Philosophy (Ph.D.) degree in Computer Science (Dissertation Title: Networking and System Support for the Efficient, Scalable Delivery of Services in Interactive Multimedia System, minor in Telecommunications Public Policy) earned in June 1997. During my education, I took a wide variety of courses as demonstrated by my minors. My undergraduate degree also included a number of courses more typical of a degree in electrical engineering including digital logic, signal processing, and telecommunications theory.

6. One of the major concentrations of my research over the past 30+ years has been the delivery of multimedia content and data between computing devices, including through various network architectures. In my research, I have studied large-scale content delivery systems, and the use of servers located in a variety of geographic locations to provide scalable delivery to hundreds or thousands of users simultaneously. I have also studied smaller-scale content delivery systems in which content is exchanged between individual computers and portable devices. My work

has emphasized the exchange of content more efficiently across computer networks, including the scalable delivery of content to many users, mobile computing, satellite networking, delivering content to mobile devices, and network support for data delivery in wireless networks.

7. In 1992, the initial focus of my research was on the provision of interactive functions (e.g., VCR-style functions like pause, rewind, and fast-forward) for near video-on-demand systems in cable systems; in particular, how to aggregate requests for movies at a cable head-end and then how to satisfy a multitude of requests using one audio/video stream broadcast to multiple receivers simultaneously. This research has continually evolved and resulted in the development of techniques to scalably deliver on-demand content, including audio, video, web documents, and other types of data, through the Internet and over other types of networks, including over cable systems, broadband telephone lines, and satellite links.

8. An important component of my research has been investigating the challenges of communicating multimedia content, including video, between computers and across networks including the Internet. Although the early Internet was used mostly for text-based, non-real time applications, the interest in sharing multimedia content, such as video, quickly developed. Multimedia-based applications ranged from downloading content to a device to streaming multimedia

content to be instantly used. One of the challenges was that multimedia content is typically larger than text-only content, but there are also opportunities to use different delivery techniques since multimedia content is more resilient to errors. I have worked on a variety of research problems and used a number of systems that were developed to deliver multimedia content to users. One content-delivery method I have researched is the one-to-many communication facility called "multicast," first deployed as the Multicast Backbone, a virtual overlay network supporting one-to-many communication. Multicast is one technique that can be used on the Internet to provide streaming media support for complex applications like video-on-demand, distance learning, distributed collaboration, distributed games, and large-scale wireless communication. The delivery of media through multicast often involves using Internet infrastructure, devices and protocols, including protocols for routing and TCP/IP.

9. Starting in 1997, I worked on a project to integrate the streaming media capabilities of the Internet together with the interactivity of the web. I developed a project called the Interactive Multimedia Jukebox (IMJ). Users would visit a web page and select content to view. The content would then be scheduled on one of a number of channels, including delivery to students in Georgia Tech dorms delivered via the campus cable plant. The content of each channel was delivered using multicast communication.

10. In the IMJ, the number of channels varied depending on the capabilities of the server including the available bandwidth of its connection to the Internet. If one of the channels was idle, the requesting user would be able to watch their selection immediately. If all channels were streaming previously selected content, the user's selection would be queued on the channel with the shortest wait time. In the meantime, the user would see what content was currently playing on other channels, and because of the use of multicast, would be able to join one of the existing channels and watch the content at the point it was currently being transmitted.

11. The IMJ service combined the interactivity of the web with the streaming capabilities of the Internet to create a jukebox-like service. It supported true Video-on-Demand when capacity allowed, but scaled to any number of users based on queuing requested programs. As part of the project, we obtained permission from Turner Broadcasting to transmit cartoons and other short-subject content. We also connected the IMJ into the Georgia Tech campus cable television network so that students in their dorms could use the web to request content and then view that content on one of the campus's public access channels.

12. More recently, I have also studied issues concerning how users choose content, especially when considering the price of that content. My research has examined how dynamic content pricing can be used to control system load. By

raising prices when systems start to become overloaded (i.e., when all available resources are fully utilized) and reducing prices when system capacity is readily available, users' capacity to pay as well as their willingness can be used as factors in stabilizing the response time of a system. This capability is particularly useful in systems where content is downloaded or streamed on-demand to users.

13. As a parallel research theme, starting in 1997, I began researching issues related to wireless devices and sensors. In particular, I was interested in showing how to provide greater communication capability to "lightweight devices," *i.e.*, small form-factor, resource-constrained (*e.g.*, CPU, memory, networking, and power) devices. Starting in 1998, I published several papers on my work to develop a flexible, lightweight, battery-aware network protocol stack. The lightweight protocols we envisioned were similar in nature to protocols like Bluetooth, Universal Plug and Play (UPnP) and Digital Living Network Alliance (DLNA).

14. From this initial work, I have made wireless networking—including ad hoc, mesh networks and wireless devices—one of the major themes of my research. My work in wireless networks spans the protocol stack from applications through to the encoding and exchange of data at the data link and physical layers.

15. At the application layer, even before the large-scale "app stores" were available, my research looked at building, installing, and using apps for a variety of purposes, from network monitoring to support for traditional computer-based

applications (*e.g.*, content retrieval) to new applications enabled by ubiquitous, mobile devices. For example, my research has looked at developing applications for virtually exchanging and tracking "coupons" through "opportunistic contact" among mobile wireless devices (*i.e.*, communication among devices moving into communication range with each other). In many of the courses I have taught there is a project component. Through these projects I have supervised numerous efforts to develop new "apps" for download and use across a variety of mobile platforms.

16. Toward the middle of the protocol stack, my research has also looked to build wireless infrastructure support to enable communication among a set of mobile devices unaided by any other kind of network infrastructure. These kinds of networks are useful either in challenged network environments (*e.g.*, when a natural disaster has destroyed existing infrastructure) or when suitable support for network communication never existed. The deployment of such networks (or even the use of traditional network support) are critical to support services like disaster relief, catastrophic event coordination, and emergency services deployment.

17. Yet another theme is monitoring wireless networks, in particular different variants of IEEE 802.11 compliant networks, to (1) understand the operation of the various protocols used in real-world deployments, (2) use these measurements to characterize use of the networks and identify protocol limitations and weaknesses, and (3) propose and evaluate solutions to these problems. I have

successfully used monitoring techniques to study wireless data link layer protocol operation and to improve performance by enhancing the operation of such protocols. For wireless protocols, this research includes functions like network acquisition and channel bonding.

18. One theme in my wireless network research has been cross-layer solutions and innovations. As mentioned above, with greater wireless device use and network support, we envisioned new application paradigms and services, for example, when mobile devices come into contact with each other. Instead of relying on existing infrastructure to relay communication, the devices are able to discover each other and communicate directly. Other examples include discovering and using location information to enhance users' experiences. Network support and novel applications use a variety of network architectures supporting users on foot, in vehicles, and across varying terrains and environments. Finally, we studied how communication efficiency can be supported through intelligent handoffs as well as location and movement prediction.

19. Protecting networks, including their operation and content, has been an underlying theme of my research almost since the beginning of my research career. Starting in 2000, I have been involved in several projects that specifically address security, network protection, and firewalls. After significant background work, a team on which I was a member successfully submitted a $4.3M grant proposal to the

Army Research Office (ARO) at the Department of Defense to propose and develop a high-speed intrusion detection system. Key aspects of the system included associating streams of packets and analyzing them for viruses and other malware. Once the grant was awarded, we spent several years developing and meeting the milestones of the project. A number of my students worked on related projects and published papers on topics ranging from intrusion detection to developing advanced techniques to be incorporated into firewalls. I have also used firewalls, including their associated malware detection features, in developing techniques for the classroom to ensure that students are not distracted by online content.

20. My recent work ties some of the various threads of my past research together. I have investigated content delivery in online social networks and proposed reputation management systems in large-scale social networks and marketplaces. On the content delivery side, I have looked at issues of caching and cache placement, especially when content being shared and the cache have geographical relevance. We were able to show that effective caching strategies can greatly improve performance and reduce deployment costs. Our work on reputation systems showed that reputations have economic value, and as such, creates a motivation to manipulate reputations. In response, we developed a variety of solutions to protect the integrity of reputations in online social networks. The techniques we developed for content delivery and reputation management were

particularly relevant in peer-to-peer communication and recommendations for downloadable "apps."

21. As an important component of my research program, I have been involved in the development of academic research into available technology in the market place. One aspect of this work is my involvement in the Internet Engineering Task Force (IETF). The IETF is a large and open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet. I have been involved in various IETF groups including many content delivery-related working groups like the Audio Video Transport (AVT) group, the MBone Deployment (MBONED) group, Source Specific Multicast (SSM) group, the Inter-Domain Multicast Routing (IDMR) group, the Reliable Multicast Transport (RMT) group, the Protocol Independent Multicast (PIM) group, etc. I have also served as a member of the Multicast Directorate (MADDOGS), which oversaw the standardization of all things related to multicast in the IETF. Finally, I was the Chair of the Internet2 Multicast Working Group for seven years.

22. My involvement in the research community extends to leadership positions for several academic journals and conferences. I am the co-chair of the Steering Committee for the ACM Network and System Support for Digital Audio and Video (NOSSDAV) workshop and on the Steering Committees for the

International Conference on Network Protocols (ICNP), ACM Sigcomm Workshop on Challenged Networks (CHANTS), and IEEE Global Internet (GI) Symposium. I have served or am serving on the Editorial Boards of IEEE/ACM Transactions on Networking, IEEE Transactions on Mobile Computing, IEEE Network, ACM Computers in Entertainment, AACE Journal of Interactive Learning Research (JILR), and ACM Computer Communications Review. I have co-chaired a number of conferences and workshops including the IEEE International Conference on Network Protocols (ICNP), IEEE Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON), International Conference on Communication Systems and Networks (COMSNETS), IFIP/IEEE International Conference on Management of Multimedia Networks and Services (MMNS), the International Workshop On Wireless Network Measurement (WiNMee), ACM Sigcomm Workshop on Challenged Networks (CHANTS), the Network Group Communication (NGC) workshop, and the Global Internet Symposium, and I have served on the program committees for numerous conferences.

23. Furthermore, in the courses I taught at UCSB, a significant portion of my curriculum covered aspects of the Internet and network communication including the physical and data link layers of the Open System Interconnect (OSI) protocol stack, and standardized protocols for communicating across a variety of physical media such as cable systems, telephone lines, wireless, and high-speed

Local Area Networks (LANs). The courses I have taught also cover most major topics in Internet communication, including data communication, multimedia encoding, and mobile application design. My research and courses have covered a range of physical infrastructures for delivering content over networks, including cable, Integrated Services Digital Network (ISDN), Ethernet, Asynchronous Transfer Mode (ATM), fiber, and Digital Subscriber Line (DSL). For a complete list of courses I have taught, see my curriculum vitae (CV).

24. In addition, I co-founded a technology company called Santa Barbara Labs that was working under a sub-contract from the U.S. Air Force to develop very accurate emulation systems for the military's next generation internetwork. Santa Barbara Labs' focus was in developing an emulation platform to test the performance characteristics of the network architecture in the variety of environments in which it was expected to operate, and, in particular, for network services including IPv6, multicast, Quality of Service (QoS), satellite-based communication, and security. Applications for this emulation program included communication of a variety of multimedia-based services, including video conferencing and video-on-demand.

25. In addition to having co-founded a technology company myself, I have worked for, consulted with, and collaborated with companies for nearly 30 years. These companies range from well-established companies to start-ups and include

IBM, Hitachi Telecom, Turner Broadcasting System (TBS), Bell South, Digital Fountain, RealNetworks, Intel Research, Cisco Systems, and Lockheed Martin.

26. Through my graduate education, leadership with CITS, involvement in TMP, role in the development of the Internet2 infrastructure, and consulting with ISPs, I have gained a strong understanding in the role of the Internet in our society and the challenges of deploying large-scale production networking infrastructure. CITS, since its inception, has looked at the role of the Internet in society, including how the evolution of technology has created communication opportunities and challenges, including, for example, through disruptive technologies like P2P. TMP looks to focus on non-purely technical issues, including, for example, state-of-the-art business methods, strategies for successful technology commercialization, new venture creation, and best practices for fostering innovation. Through my industry collaborations and Internet2 work, I have developed significant experience in the challenges of deploying, monitoring, managing, and scaling communication infrastructure to support evolving Internet services like streaming media, conferencing, content exchange, social networking, and e-commerce.

27. I am a Member of the Association of Computing Machinery (ACM) and a Fellow of the Institute of Electrical and Electronics Engineers (IEEE).

28. Additional details about my employment history, fields of expertise, courses taught, and publications are further included in my CV attached as Ex1006.

III. LEGAL STANDARDS

29. In forming my opinions and considering the subject matter of the patent and its claims in light of the prior art, I am relying on certain legal principles that counsel in this case explained to me. My understanding of these concepts is summarized below.

30. I understand that, in an *inter partes* review, earlier publications and patents may act to render a patent unpatentable for one of two reasons: (1) anticipation, and (2) obviousness.

A. Anticipation

31. I understand that the claims of a patent are anticipated by a prior art reference if each and every element of the claim is found either explicitly or inherently in the reference. I understand that inherency requires a showing that the missing descriptive matter in the claim is necessarily present in the allegedly anticipating reference, and that it would have been so recognized by a person of ordinary skill in the art ("POSITA").

32. I understand that when a challenged claim covers several structures, either generically or as alternatives, the claim is deemed anticipated if any of the structures within the scope of the claim is found in the prior art reference.

33. Although anticipation typically involves the analysis of a single prior art reference, I understand that additional references may be used to show that the

prior art reference has enabling disclosure (*i.e.*, allows a POSITA to make the invention without undue experimentation), to explain the meaning of a term used in the prior art reference, and/or to show that a characteristic is inherent in the prior art reference.

B. Obviousness

34. I understand that a claim is invalid as obvious if it would have been obvious to a person of ordinary skill in the art at the time the alleged invention was made. This means that even if all of the elements of the claim cannot be found in a single prior art reference that would anticipate the claim, a person of ordinary skill in the art who was aware of the prior art would have been able to come up with the claimed invention. This may be the case, for example, where the missing element represents only an insubstantial difference over the prior art or a reconfiguration of a known system. I understand that in an obviousness determination, the person of ordinary skill in the art is presumed to have knowledge of all material prior art.

35. I understand that an obviousness analysis requires an understanding of the scope and content of the prior art, any differences between the alleged invention and the prior art, and the level of ordinary skill in evaluating the pertinent art.

36. I understand that when a product is available, design incentives and other market forces can prompt variations of it, either in the same field or a different one. If a person of ordinary skill in the art can implement a predictable variation,

obviousness likely bars its patentability. For the same reason, if a technique has been used to improve one device and a person of ordinary skill in the art would recognize that it would improve similar devices in the same way, using the technique would have been obvious.

37. I understand that whether a prior art reference renders a patent claim unpatentable as obvious is determined from the perspective of a person of ordinary skill in the art at the time of the alleged invention. I have been told that there is not a requirement that the prior art contain an express suggestion to combine known elements to achieve the claimed invention, but a suggestion to combine known elements to achieve the claimed invention may come from the prior art, as filtered through the knowledge of one skilled in the art. In addition, I have been told that the inferences and creative steps a person of ordinary skill in the art would employ are relevant to the determination of obviousness.

38. I understand that one may consider, e.g., whether (1) the change was merely the predictable result of using prior art elements according to their known functions, or whether it was the result of true inventiveness; (2) there is some teaching or suggestion in the prior art to make the modification or combination of elements claimed in the patent; (3) the claimed innovation applies a known technique that had been used to improve a similar device or method in a similar way; (4) the claimed invention would have been obvious to try, meaning that the claimed

innovation was one of a relatively small number of possible approaches to the problem with a reasonable expectation of success by those skilled in the art; (5) the invention merely substituted one known element for another known element in order to obtain predictable results; (6) the invention merely applies a known technique to a known device, method, or product to yield predictable results; or (7) known work in one field of endeavor may have prompted variations of it for use in either the same field or a different one based on design incentives or other market forces that would have been predictable to a person of ordinary skill in the art.

39. I further understand that certain factors may support or rebut the obviousness of a claim. I understand that such secondary considerations include, among other things, commercial success of the patented invention, skepticism of those having ordinary skill in the art at the time of the invention, unexpected results of the invention, any long-felt but unsolved need in the art that was satisfied by the alleged invention, the failure of others to make the alleged invention, praise of the alleged invention by those having ordinary skill in the art, and copying of the alleged invention by others in the field. I understand that there must be a nexus—that is, a connection—between any such secondary considerations and the alleged invention. I also understand that contemporaneous and independent invention by others is a secondary consideration tending to show obviousness.

40. Additionally, I understand that in considering obviousness, it is

important not to use the benefit of hindsight derived from the patent under consideration.

C. Level of Ordinary Skill in the Art

41. When interpreting a patent, I understand that it is important to identify the relevant art pertaining to that patent, as well as the level of ordinary skill in that art at the time of the claimed invention. The "art" is the field of technology to which the patent is related.

42. I am informed and understand that the POSITA is a hypothetical person who is presumed to know the relevant prior art. I understand that the actual inventor's skill is not determinative of the level of ordinary skill. I further understand that factors that may be considered in determining level of skill include: (i) the types of problems encountered in the art; (ii) prior art solutions to those problems; (iii) the rapidity with which innovations are made; (iv) the sophistication of the technology; and (v) the educational level of active workers in the field. I understand that not all such factors may be present in every case, and one or more of them may predominate.

43. I understand that a POSITA is one who is presumed to be aware of all pertinent art, thinks along conventional wisdom in the art, and is a person of ordinary creativity.

D. Claim Construction

44. I understand that, in this IPR, the claims are interpreted consistent with the plain and ordinary meaning of the term as it would have had to a POSITA at the time of the alleged invention. I understand that, in the pending district court proceeding, Patent Owner has contended that the term "global stage" (as used in claims 3, 12, and 19) should be construed as "a global stage, distinct from the client discrimination stage and distinct from the user specific stage." However, regardless of whether this construction is adopted, it is my opinion that this limitation is disclosed and rendered obvious by the prior art.

IV. CHALLENGED CLAIMS AND GROUNDS OF UNPATENTABILITY

45. The '639 patent claims priority to a provisional application filed October 10, 2002. I am not providing any opinion as to whether the '639 Patent is in fact entitled to that priority date, but I am assuming that the October 10, 2002 priority date applies for the purpose of the IPR proceedings. My opinions that the Challenged Claims are unpatentable are based on the following references. I understand all of these references would be prior art even assuming that the October 10, 2002 priority date applies. None of these references were cited during the prosecution of the '639 Patent.

Exhibit	Reference	Reference Date	Prior Art Status
Ex1003	U.S. Patent Publication No. 2003/0152067 ("Richmond")	Sept. 20, 2002 (filing date)	§102(e)

Ex1004	WIPO Publication WO 01/11452A2 ("Wood")	Feb. 15, 2001 (publication date)	§102(a), (b)
Ex1005	U.S. Patent No. 5,623,492 ("Teraslinna")	Mar. 24, 1995 (filing date) Apr. 22, 1997 (issue date)	§102(a), (b), (e)

46. I conclude that the Challenged Claims are invalid based on the following grounds of unpatentability.

Ground	Claims	Basis	Prior Art
1	1-3, 5-8, 10-15, 17-20, 22-25	§102	Richmond
2	1-8, 10-15, 17-20, 22-25	§103	Richmond
3	9, 16, 26-27	§103	Richmond, Wood
4	21	§103	Richmond, Teraslinna

V. TECHNOLOGY BACKGROUND

47. The '639 Patent generally rates to the process of controlling users' access to packet-switched computer networks. In such a network, the flow of data to or from any connected device is divided into small, distinct units known as "packets," each of which are transmitted through the network independently of one another. A network device sends a packet to network hardware, such as a switch or router, which then passes it along to the next piece of equipment in line until it arrives at its intended destination. It is possible for two packets from the same origin, heading to the same destination, to take different paths, especially if a switch's route is congested due to heavy traffic.

48. To allow a switch to identify where a packet should go and enable the

recipient to recognize the packet's source, each packet carries unique addresses for both the sending and receiving devices. This information is stored in set positions at the start of the packet, within a section dubbed the "header." The header also includes additional information that the network infrastructure may use for switching or routing purposes. Additionally, the packet holds the data being sent, termed the "payload."

49. One of the most common protocols for packet-based networking is defined in the IEEE 802.3 Ethernet standard. Each Ethernet device is assigned a distinct 6-byte Media Access Control ("MAC") address. Ex1009, 280. When transmitting data, an Ethernet device constructs a packet with a header showing its own MAC address and that of the intended recipient. (Note that many sources instead refer to an Ethernet "frame" instead of a packet. Strictly speaking, the Ethernet frame is a subset of the packet, omitting the 7-byte preamble and 1-byte start-of-frame delimiter.) The data to be transmitted, typically up to 1500 bytes, is then loaded into the payload area. This packet/frame is then sent to a switch, which examines the header to decide the proper destination device and forwards it accordingly.

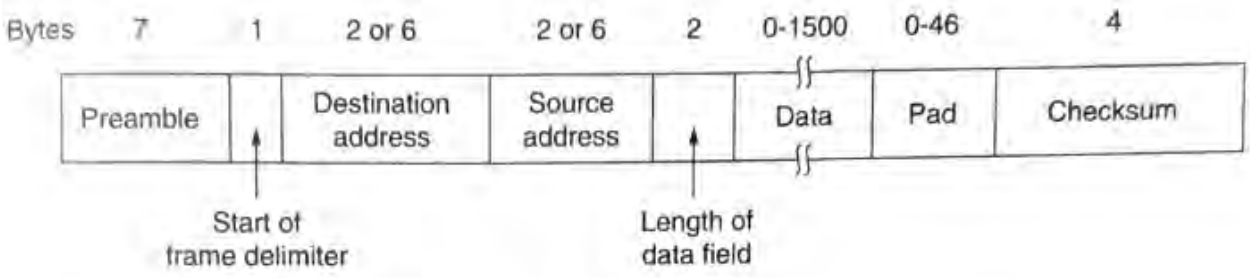

Fig. 4-21. The 802.3 frame format.

Ex1009, 281.

50. The payload of an Ethernet packet often includes a packet from a higher-level protocol. One higher-level protocol that is commonly used in conjunction with Ethernet is the Internet Protocol ("IP"), which underpins the Internet. Every IP network device is allocated a 4-byte (or in some instances, 16-byte) IP address. Ex1009, 416. Each IP packet has a header with both source and destination IP addresses, along with additional fields. The IP payload comes after this header. As a result, in a network employing IP over Ethernet, every packet a device receives will include both a MAC address and an IP address within the Ethernet and IP headers, respectively.

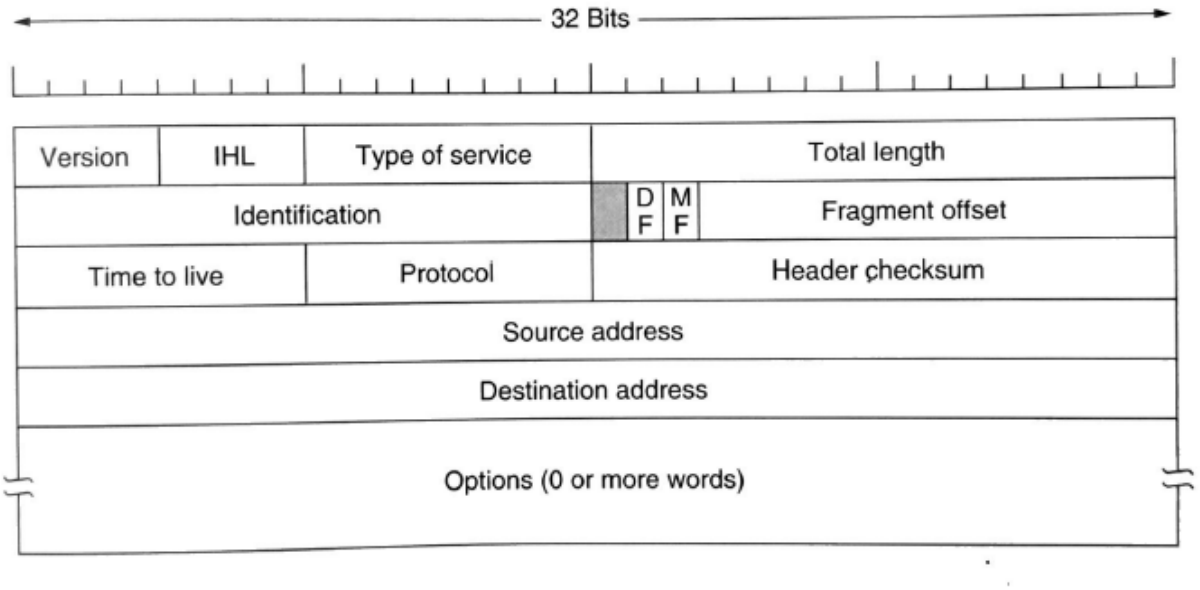

Fig. 5-45. The IP (Internet Protocol) header.

Ex1009, 413.

VI. THE '639 PATENT

A. Overview of the Specification

51. The '639 Patent is titled "System and Method for Providing Access Control," and issued on February 14, 2012. The application that led to the '639 Patent, U.S. Patent Application No. 10/683,317, was filed on October 10, 2003. The '639 Patent claims priority to U.S. Provisional Application No. 60/417,674, which was filed October 10, 2002. It is my understanding that, because the application leading to the '639 Patent was filed before March 16, 2013, the '639 Patent is subject to the first-to-invent system that was in place prior to the America Invents Act ("AIA").

52. The stated goal of the '639 Patent is to address bandwidth

overutilization in a shared network. For instance, it points out that "many cafes have public wireless networks to allow patrons to access the Internet," while warning that without limiting bandwidth usage, "one user connected to a particular router can consume a majority of the bandwidth (e.g., downloading pictures from the Internet), slowing down the wireless network for other users connected to that router." Ex1001, 1:63-2:4. The patent acknowledges that authenticating users prior to granting network access was a standard practice in earlier systems. *See id.*, 1:39-44 ("To access a LAN and, in turn, the Internet, many prior art access control systems require a user to … enter a user name and password."). However, it contends that these older access control mechanisms "typically assume that a user is tied to a particular physical port," such that supposedly "provisioning of bandwidth is done on a per port rather than a per user basis." *Id.*, 1:45-52.

53. The disclosed solution of the '639 Patent to this purported problem is a control device 20 that receives network traffic from a user device 14 and controls the user device's access to controlled network 18. *Id.*, 4:13-27. In reaching these access control decisions, the control device 20 can consult an authentication database 24, which maintains a list of authorized users along with their respective access rights. *Id.*, 4:28-33.

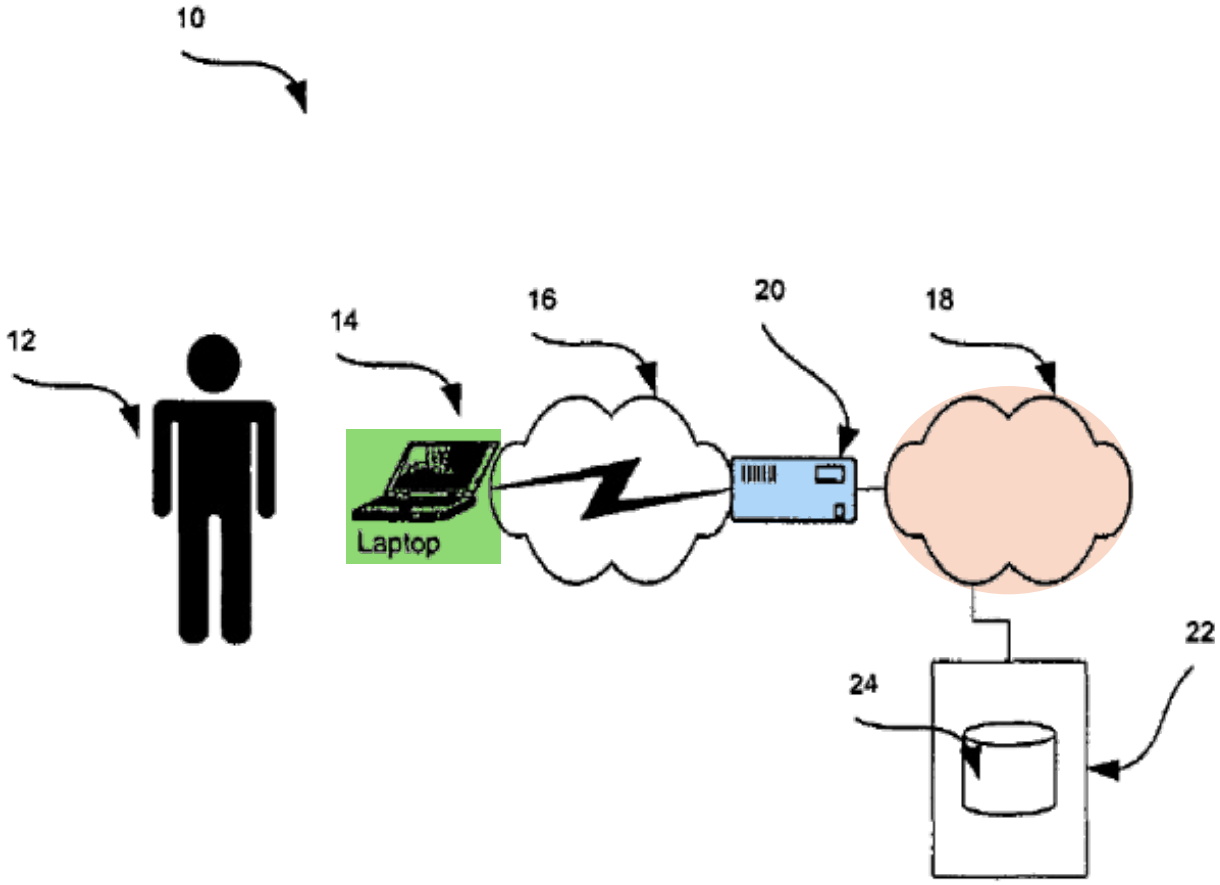

Ex1001, Fig. 1 (annotated).

54. The control device 20 functions by analyzing packets received from user devices, extracting fields such as MAC and IP addresses to establish user identity. *Id.*, 4:54-62. Should these addresses not match a previously authenticated user, the device will prompt the user to verify their identity, potentially by redirecting them to a login webpage. *Id.*, 4:63-5:13, 6:24-56. Once authentication is complete, the control device stores the user device's MAC and IP addresses for future recognition. *Id.*, 5:32-35. For users who are authenticated, the device regulates their network access based on rules associated with that user (e.g., limits on upload or

download bandwidth, or restrictions on certain sites or applications). *Id.*, 7:4-21.

B. The Challenged Claims

55. The '639 Patent contains 27 claims. Claims 1, 10, 17, and 27 are independent claims. Each of the independent claims are similar. Claim 1 is directed to "[a] method for network access control" performed "at a control device." Independent claims 10 is directed to a "computer program product" that is configured to perform, when executed, steps that are substantively the same as those of method claim 1. Independent claim 17 is directed to a "control device for network access control" that performs substantively the same steps as claims 1 and 17. Independent claim 27 is (like claim 17) directed to a "control device for network access control." It differs from claim 17, however, in that it adds limitations that an unauthenticated user is redirected to a login webpage at which to provide credentials.

C. Level Ordinary Skill in the Art

56. In my opinion, based on my review of the intrinsic record and my professional experience in the field, a person of ordinary skill in the art ("POSITA") in the field of the '639 Patent would have had at least a bachelor's degree in computer science, computer engineering, or related field, with two to three years working with computer networking. It is my opinion that additional graduate education could substitute for professional experience, while significant experience in the field might substitute for formal education. However, minor variations of this definition would

not have an impact on the conclusions. As I have stated above, I am assuming for the purpose of this petition that the '639 Patent is entitled to claim priority to the October 10, 2002 filing date of the provisional application (though I am not offering an opinion on this issue). However, my conclusions would not change if the October 10, 2003 non-provisional filing date were applicable, as all the references I have cited would be prior art under either date, and the understanding of a POSITA would not have meaningfully differed between those dates.

VII. CLAIM INTERPRETATION

57. I understand that claims in an IPR proceeding are interpreted using the plain-and-ordinary meaning standard that was set out by the Federal Circuit in *Phillips v. AWH Corp.*, 415 F.3d 1303, 1312 (Fed. Cir. 2005). As I understand it, although the patent is currently the subject of litigation between Petitioner and Patent Owner, the Court has not yet construed any claims. I further understand that neither party has (at least as of this date) proposed that any terms should be construed by the court, or argued that any terms are used in the patent with a meaning other than their plain and ordinary meaning to a POSITA. Accordingly, I am applying the plain meaning of claim terms as would have been understood by a POSITA at the time of alleged invention. Moreover, it is my opinion that the '639 Patent would be invalid under any reasonable constructions of the terms. I reserve the right to provide further reasoning should Patent Owner propose narrowing constructions and then argue that

the prior art does not meet the claim limitations under those narrow constructions.

VIII. PRIOR ART OVERVIEW

A. Overview of Richmond

58. U.S. Patent Publication No. 2003/0152067 ("Richmond") is a publication of U.S. Patent Application No. 10/251,140, filed on September 20, 2002. Even if it is assumed that the '639 Patent is entitled to the October 10, 2002 filing date of its provisional application, Richmond predates it and thus qualifies as prior art under 35 U.S.C. §102(e). Richmond is also a continuation-in-part of U.S. Patent Application No. 10/071,873, submitted on February 8, 2002. Richmond is titled "Controlling Concurrent Usage of Network Resources by Multiple Users at an Entry Point to a Communications Network Based on Identities of the Users." Ex1003, cover. This publication details a system where user devices link to a network, including the Internet, via "network entry devices," as illustrated in the figure below.

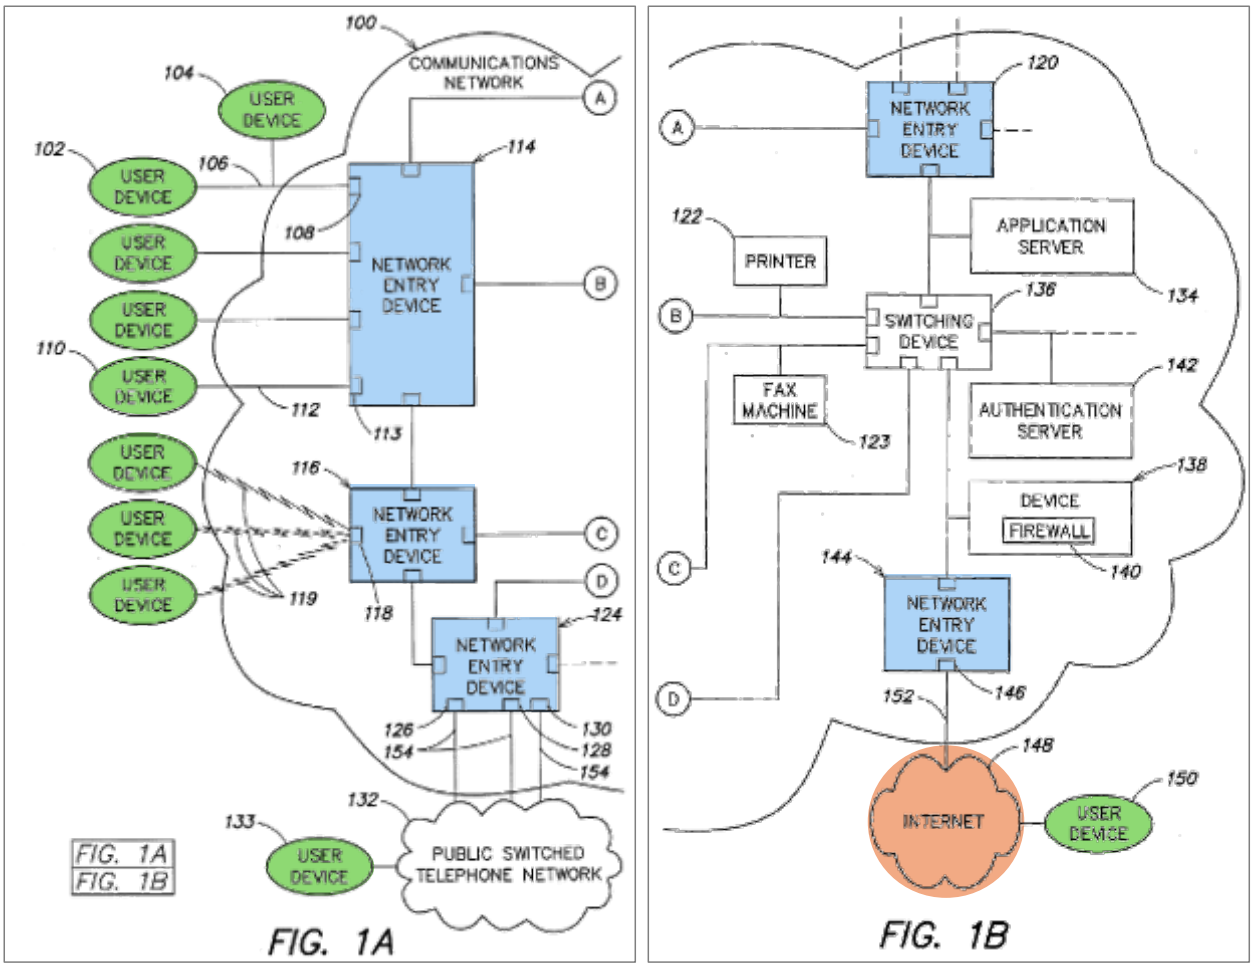

Ex1003, Figs. 1A-1B (annotated).

59. As is typical in Internet communications, Richmond's user devices exchange data on the network by sending and receiving a sequence of packets, each containing a header containing source and destination addresses as well as a data payload. Ex1003, [0010]. The network entry devices apply a set of "packet rules" to the packets that pass through them. "[E]ach packet rule include[es] a condition and action to be taken if a packet received at a device satisfies the condition." *Id.*, [0082]. For example, the sample packet rule depicted in Figure 5 indicates (in field

508) that "[the] IP addresses of a packet are to be examined," and specifies (in field 510) a comparison of the packet's IP address against "10.20.54.0." *Id.*, [0162]-[0163]. If a packet matches this rule, the rule directs the assignment of priority 7 (item 516), "representing a Quality of Service (QoS)/Class of Service (CoS)" for the packet. *Id.*, [0156], [0165]. Additionally, the rule instructs not to drop the packet (item 581) and not to apply rate limiting (item 520).

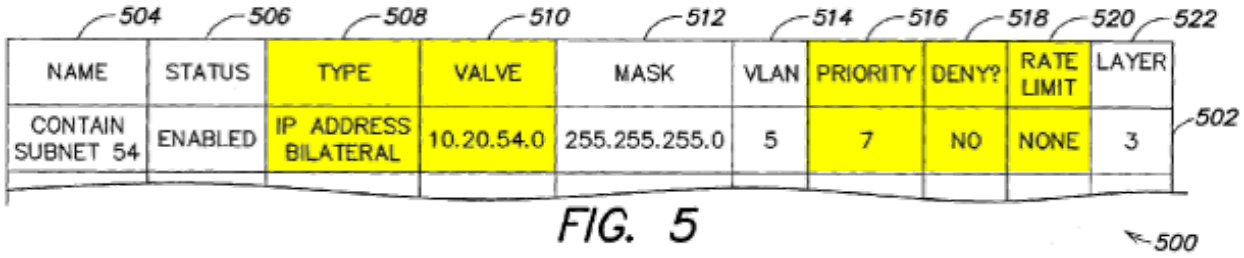

Id., Fig. 5 (annotated).

60. Richmond's packet rules can be assigned to "service abstractions," which bundle multiple rules to establish service classes for specific applications. *Id.*, [0116]-[0118]. As shown in Figure 2, various rules might be grouped to create a "high-priority" class of service (items 208 and 212) that favors traffic for particular applications (e.g., SAP and Citrix). On the other hand, other service abstractions might restrict certain applications (e.g., service 210 labeled "NO SNMP" blocks SNMP traffic). *Id.*, [0141], Fig. 2:

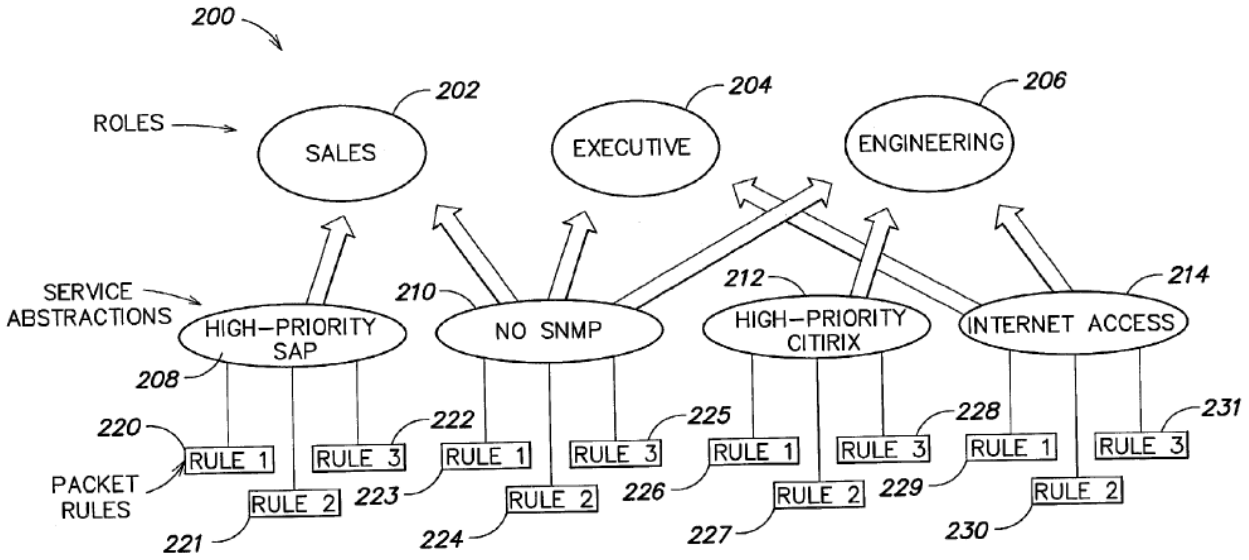

FIG. 2

61. Figure 2 also demonstrates that service abstractions (along with their associated packet rules) can be allocated to user categories, referred to as "roles." *Id.*, [0121], [0132], [0139]. For instance, the Sales role 202 is granted access to the "High-Priority SAP" service 208 but denied general "Internet Access" 214. In contrast, users in the Executive and Engineering roles (items 204 and 206) are granted Internet access but are excluded from SAP. *Id.*, [0141], Fig. 2. All three roles are prohibited from using the SNMP network management protocol, owing to their assignment to the "NO SNMP" service abstraction 210. *Id.*, [0141].

62. In Richmond's system, the enforcement and application of packet rules—along with determining which rules apply based on roles and service abstractions—are handled by network entry devices, specifically their entry port modules, working in tandem with an authentication module. *Id.*, [0230], [0233],

[0238], [0244], [0249]. These components are depicted in the figures below:

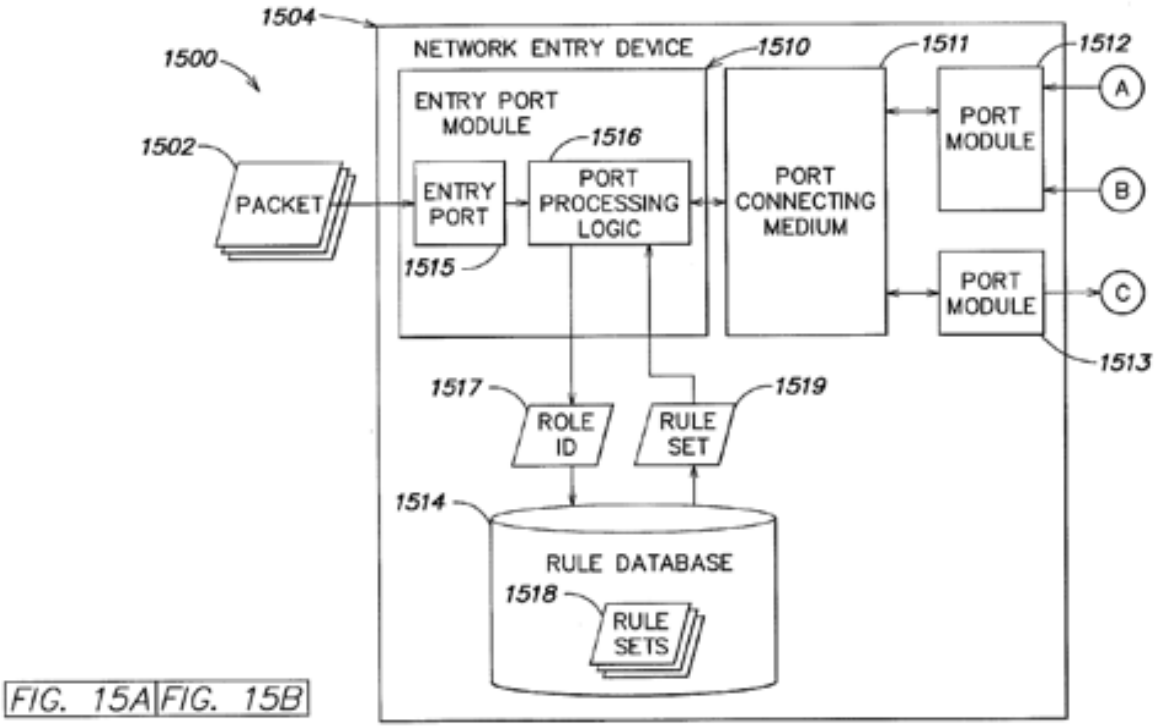

FIG. 15A

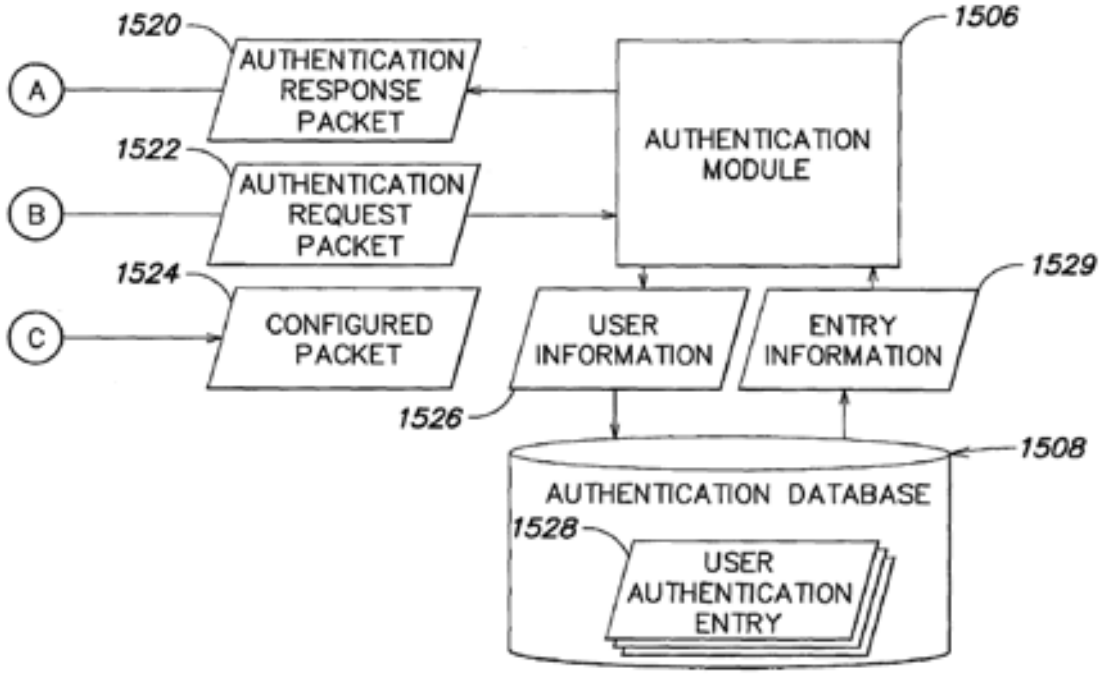

FIG. 15B

63. When a user attempts to connect to a network managed by Richmond's network entry devices, their device transmits a login packet to the entry port module 1510 of the network entry device 1504, containing information sufficient to verify the user's identity, such as a username and password. *Id.*, [0217], Fig. 13A(2) (act 1310). The network entry device then relays this identity information to an authentication module, which maintains a database holding the credentials and roles of each user. *Id.*, [0273]-[0275]. If the authentication module reports that the credentials are invalid, the network entry device will block the user's access to the network. *Id.*, [0222], Fig. 13A(2) (acts 1314 and 1316). Conversely, if the credentials are valid, the authentication module notifies the network entry device of the user's roles, prompting the device to configure its entry port modules with user-specific packet rules that correspond to the user's identity, roles, and service abstractions. *Id.*, [0223], [0277]-[0279].

64. Afterward, when the user sends packets for transmission on the network, the network entry device processes them according to these pre-configured user-based rules. *Id.*, [0280]-[0281], Fig. 13B (act 1327). For example, one type of rule directs the device to discard all packets aimed at a specific destination, such as to prevent access to non-work-related sites or applications. *Id.*, [0157]. Another type of rule restricts the rate of a user's bandwidth usage, in order to avoid any single user or small group from dominating the available upstream bandwidth. *Id.*, [0158]-

[0159].

65. For simplicity, Richmond permits the use of default rules when user-specific rules do not address all fields. *Id.*, [0212], [0231], [0252]. Once processed (and if not dropped), the device forwards the packet through another port module. *Id.*, [0282]. In cases where rate limiting applies, the packets are transmitted at the designated rate. *Id.*, [0251], Fig. 14B (act 1430).

B. Overview of Wood

66. WIPO Publication WO 01/11452A2 to Wood was published on February 15, 2001. Wood therefore is prior art under §102(a) and (b). Wood discloses a system in which access to a network is controlled by a "Gatekeeper and entry handler component 110," which "provides an entry point for external client applications requesting access." Ex1004, 5:10-15. The gatekeeper interacts with authorization component 130 to determine whether a user is authenticated. *Id.*, 6:3-8.

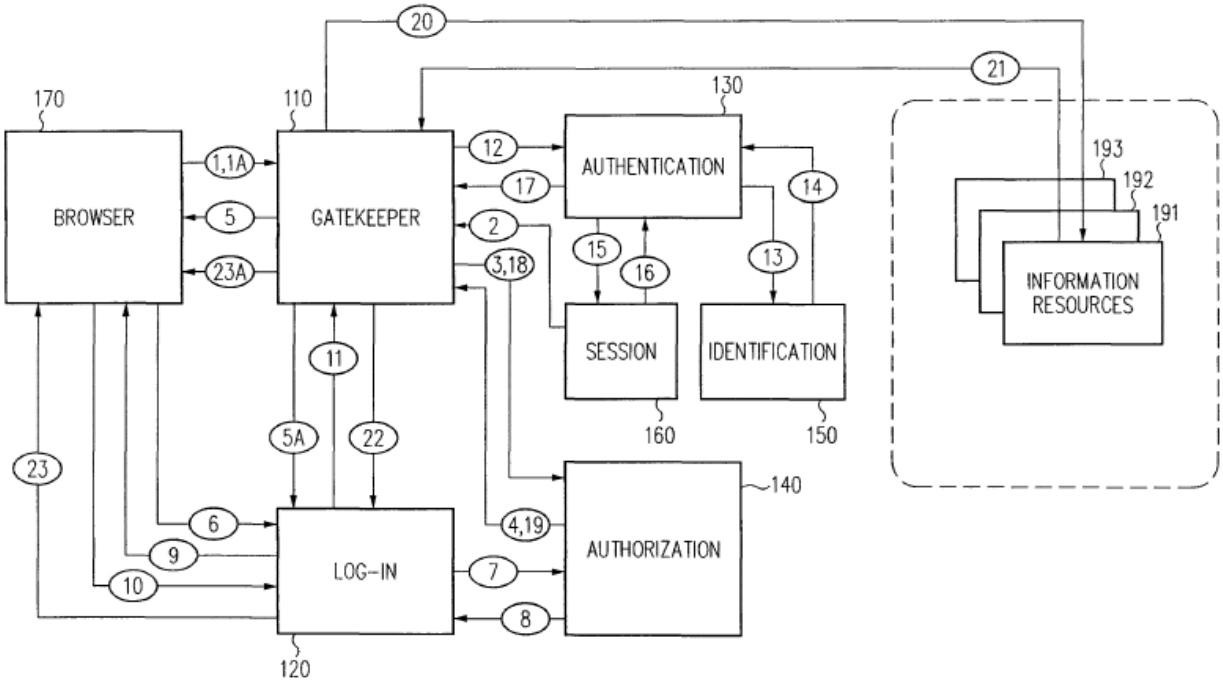

FIG. 1

67. If the user has not authenticated, the authorization component informs the gatekeeper, which then redirects the user to a login component 120. *Id.*, 6:7-9, 12:1-26. The login component then redirects the user's browser to an authentication webpage that requests the user to supply his or her credentials. *Id.,* 14:3-10. The system then uses the received credentials to authenticate the user. *Id.*, 14:22-25.

C. Overview of Teraslinna

68. U.S. Patent No. 5,623,492 ("Teraslinna") was filed on March 24, 1995, and issued on April 22, 1997. Therefore, it is prior art to the '639 Patent under 35 U.S.C. §102(a), (b), and (e). Teraslinna is titled "Methods and Systems for Managing Bandwidth Resources in a Fast Packet Switching Network." Teraslinna relates to "managing bandwidth resources within a switching communication

network having at least one fast packet switch." Ex1005, 1:8-13. A POSITA would recognize that Teraslinna uses the term "fast packet switching" to refer simply to a packet-switched network. Ex1005, 1:20-25 (defining "fast packet switching" as "a technique wherein a message is communicated between a source endpoint and a destination endpoint of a switching network using a series of information passing elements, i.e. packets, each of which contain both a payload field and an address"). The "fast" in Teraslinna's "fast packet switching" is by comparison to circuit-switched networks, which are typically slower than packet-switched networks because they need to establish a dedicated fixed path (or "circuit") between two devices before communication can occur between them.

69. Teraslinna presents a framework to regulate packet flow from user devices ("subscriber equipment 135") through a customer's network ("customer premises equipment (CPE) [130]") and a telecom carrier's network ("fast packet switching network 134") to destinations ("subscriber equipment 140"). Ex1005, 15:27-39. It emphasizes controlling bandwidth to enforce service agreements with telecom providers. *Id.*, 1:58-2:7. This is accomplished through a "traffic flow processor" linking the networks. *Id.*, 15:40-42, Fig. 11.

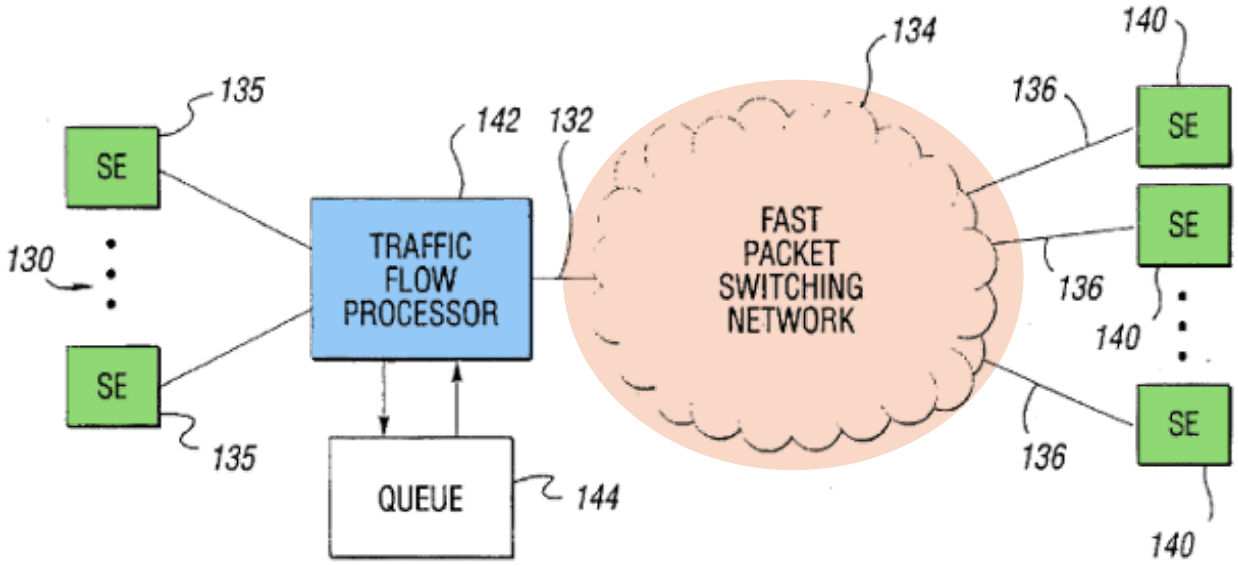

Id., Fig. 11 (annotated).

70. Teraslinna's traffic flow processor 142 connects to queue 144, which "queues all incoming packets from the customer premises equipment 130 for traffic shaping by the traffic flow processor." *Id.*, 16:5-8. The processor then applies an algorithm to determine the forwarding rate, ensuring compliance with bandwidth restrictions. *Id.*, 16:9-17:18, Fig. 12.

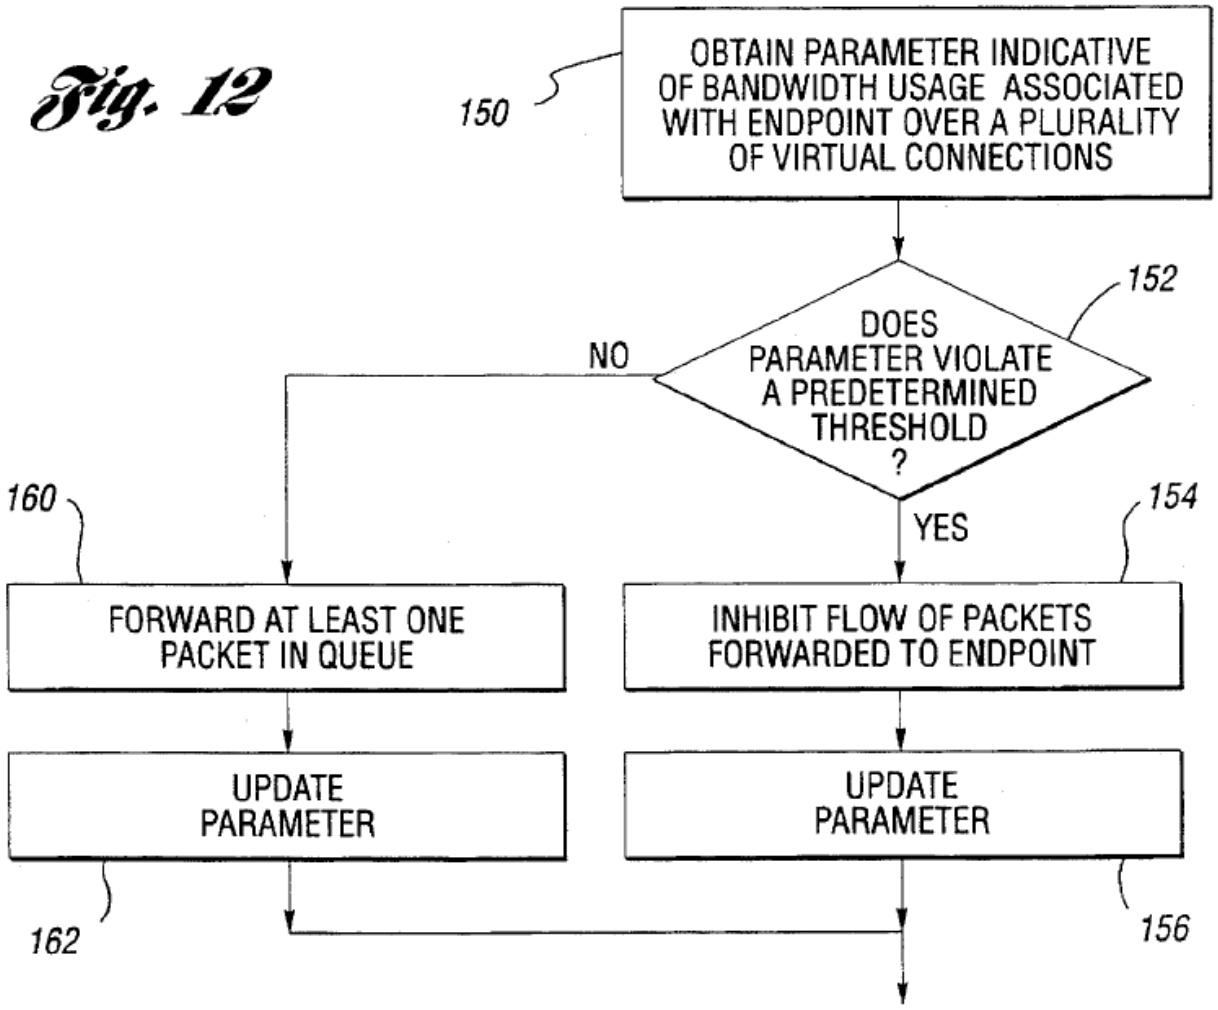

Id., Fig. 12.

IX. GROUNDS 1 & 2 – CLAIMS 1-8, 10-15, 17-20, 22-25 ARE ANTICIPATED OR RENDERED OBVIOUS BY RICHMOND

A. Independent Claim 1

71. In my opinion, Richmond discloses and renders obvious all limitations of claim 1.

1. A method for network access control, comprising

72. In my opinion, to the extent the preamble is limiting, Richmond discloses a method for network access control.

73. Richmond's "Background" section highlights that "[a] common issue in [communications] networks is how to control usage of network resources by users," referring to access control, and its specification overall provides a framework for addressing this issue. Ex1003, [0003]. The Abstract of Richmond further clarifies that, within its disclosed system, "[e]ach users' usage of [i.e., access to] network resources is controlled, after such user has been authenticated," through the use of "[a]n entry port module of a network entry device," where "each packet received from each user may be examined to control usage [i.e., access] of network resources." *Id.*, Abstract.

74. The specific methods performed by the entry port module are disclosed in methods 1300 and 1400, which "configur[e] an entry port module of a network device … to control usage of network resources." *Id.*, [0206]. Below I discuss further details of these methods, including how they meet each element of claim 1.

2. at a control device, receiving a packet originating from a user device in a first network, wherein the first network is connected to a second network via the control device, and wherein the user device is associated with a user;

75. Richmond discloses, at a control device (*e.g.*, network entry devices 114, 116, 120, 124, 144, and 1504), receiving a packet originating from a user device (*e.g.*, user devices 102, 104, 110, 133, 150, and 1502) in a first network (*e.g.*, Network 100), wherein the first network is connected to a second network (*e.g.*, the Internet 148) via the control device, and wherein the user device is associated with

a user.

76. Richmond discloses a "Network 100" (serving as an example of a first network) that "include[es] … network entry devices 114, 116, 120, 124, and 144" (which are exemplary control devices within the meaning of the claim term). Ex1003, [0012]. Richmond further discloses that "[a] network entry device serves as an entry point to the network for at least one user." *Id.*, [0016]. Additionally, Richmond's "network entry device … reside[s] at an edge or boundary of the communications network and provide[s] connectivity between network resources of the communications network and devices located external to the communications network," such as user devices. *Id.*, [0016].

77. Richmond describes a "user device" as "a device used by a user of a communications network to … receive a packet from the network [and/or] send a packet to the network," encompassing items like "workstations, terminals, personal computers, laptops, telephones, [and] pagers." *Id.*, [0016]. The overall layout of Richmond's disclosed system, featuring both network entry devices and user devices, is illustrated in Figures 1A-1B.

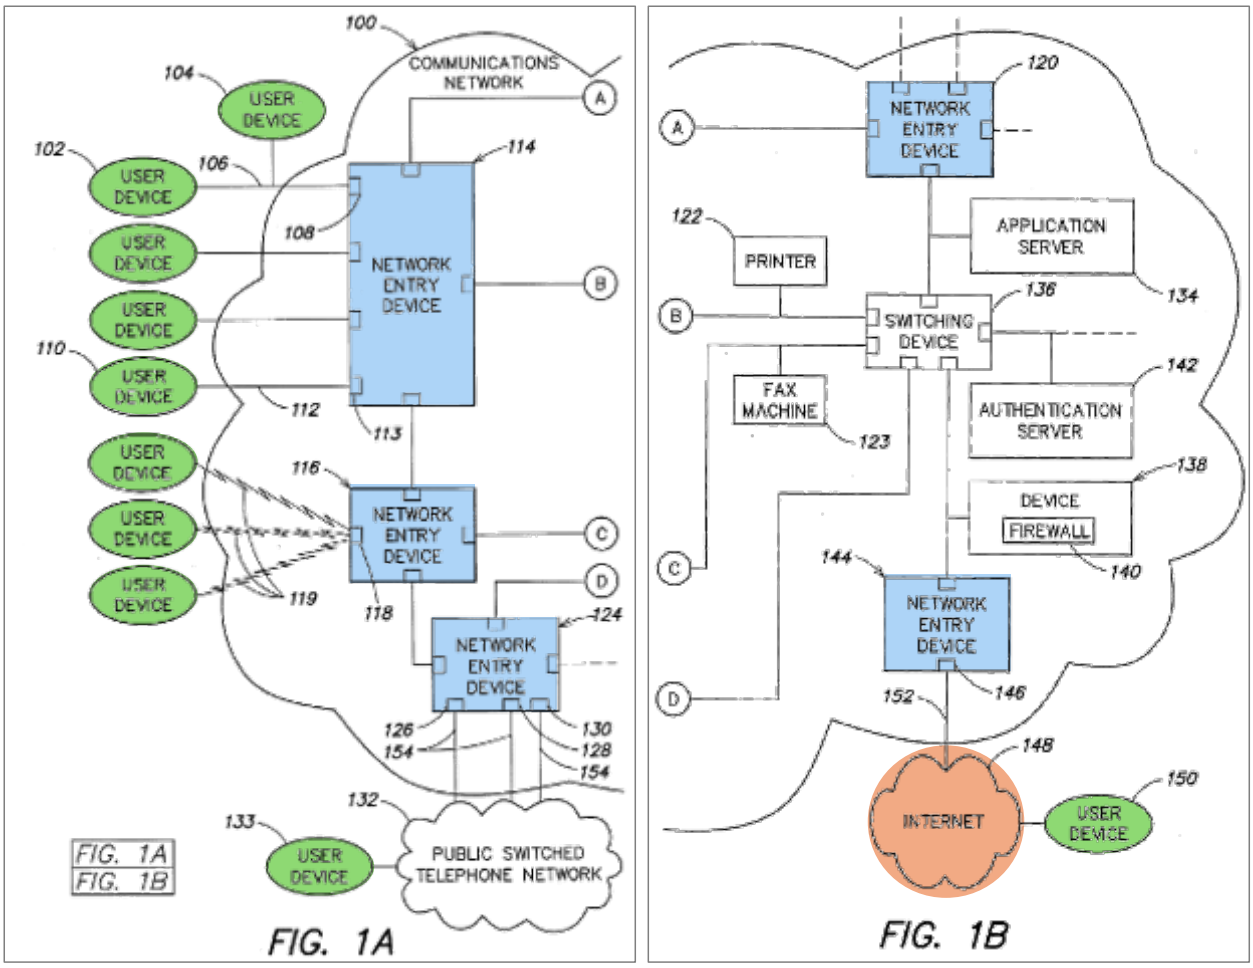

Id., Figs. 1A-1B (annotated).

78. The component of the network entry devices to which user devices and other devices connect is referred to as a "port module." *Id.*, [0013]. A port module "includes a physical port of the device and port processing logic associated with the physical port." *Id.*, [0014]. The port processing logic of a port module "include[s] switching logic, memory, and one or more processors for configuring the port module and for processing packets sent to and received at the port module." *Id.* An "entry port module" is a particular type of port module that "serves as a user's entry

point into a network." *Id.*, [0015]. "[T]o communicate with devices on the network, a user's device may transmit one or more packets to an entry port ." *Id.*

79. In Figure 15A, Richmond illustrates the configuration of its network entry devices (labeled as item 1504). *Id*., [0261]-[0262]. Packets 1502, which originate from a user device, are received by the entry port module 1515 within the network entry device. *Id*., [0265], [0269], [0271]. The network entry device also comprises additional port modules, such as item 1513, through which packets from the user device can be relayed, thus linking the network entry device and the user device to the second network. *Id*., [0270], [0282]. The network entry device functions as an exemplary control device by determining whether the traffic from the user device is permitted to access the network.

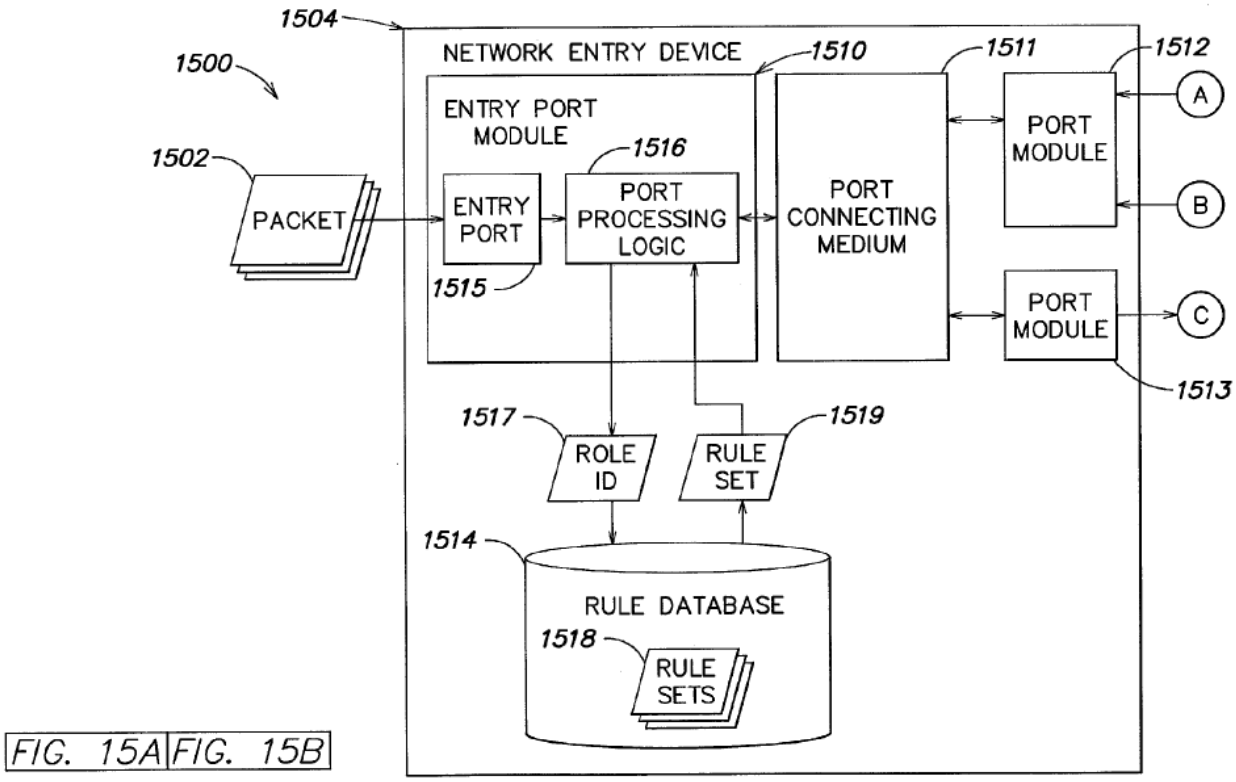

FIG. 15A

Id., Fig. 15A.

80. User devices have the ability to connect to network entry devices through a range of connectivity options, including dedicated or shared wire-based transmission mediums (such as wired Ethernet), wireless transmission mediums (such as IEEE 802.11 Wifi connections, which are typically shared), or the Internet. *Id.*, [0021]-[0025]. Each user device is "associated with a user" (meaning the operator or owner of that device) as required by the claim. *See, e.g., id.*, [0028], [0035].

81. As depicted in Fig. 1B, the first network (for instance, Network 100) is linked to the Internet 148 (serving as an exemplary second network) through the

network entry devices (e.g., item 144, an exemplary control device). Alternatively, if the Internet 148 is considered the first network, it is similarly connected to the Network 100 (in this alternative identification, an exemplary second network) via the network entry devices.

3. processing the packet according to a plurality of stages, including a client discrimination stage and a user specific rule stage;

82. Richmond discloses the step of processing the packet according to a plurality of stages, including a client discrimination stage (*e.g.*, acts 1310 and 1312 in Fig. 13A(2) and act 1401 in Fig. 14A(1)) and a user-specific rule stage (*e.g.*, acts 1402 and later in method 1400).

83. Richmond depicts a series of flow diagrams in Figures 13 and 14, illustrating methods 1300 and 1400, that show that the acts performed by Richmond's system proceed in stages, including a client discrimination stage and user specific rule stage.

84. **Client discrimination stage**. The '639 Patent includes in the specification examples of the acts that can occur in a client discrimination stage, for example that:

> At client discrimination stage 102, firewall module 82 can read the packet header of packet 80 to extract information to associate packet 80 with a user. Continuing with the previous example, if firewall module 82 extracts MAC address 08:00:69:02:01:FF and the IP address 100.100.100, firewall module 82 can associate packet 80 with user specific rule 92.

Ex1001, 10:56-61. These examples are like that which are disclosed in Richmond. Exemplary client discrimination stages are illustrated in method 1300 of Figure 13, in which Richmond's system encompasses various steps that are related to client discrimination. For instance, the network entry device accepts a login packet from the user that contains "identification information of the user" (act 1310), followed by the device "authenticate[s] the received identification information" (act 1312). *Id*., [0217], [0220], Fig. 13A(2). Depending on whether "the identification information is valid" (act 1314), the network entry device will either "deny the user access to the communications network," or alternatively determine the user's role (act 1318) and configure the entry port module accordingly (acts 1319 and 1320). *Id*., [0221]-[0225], Fig. 13B.

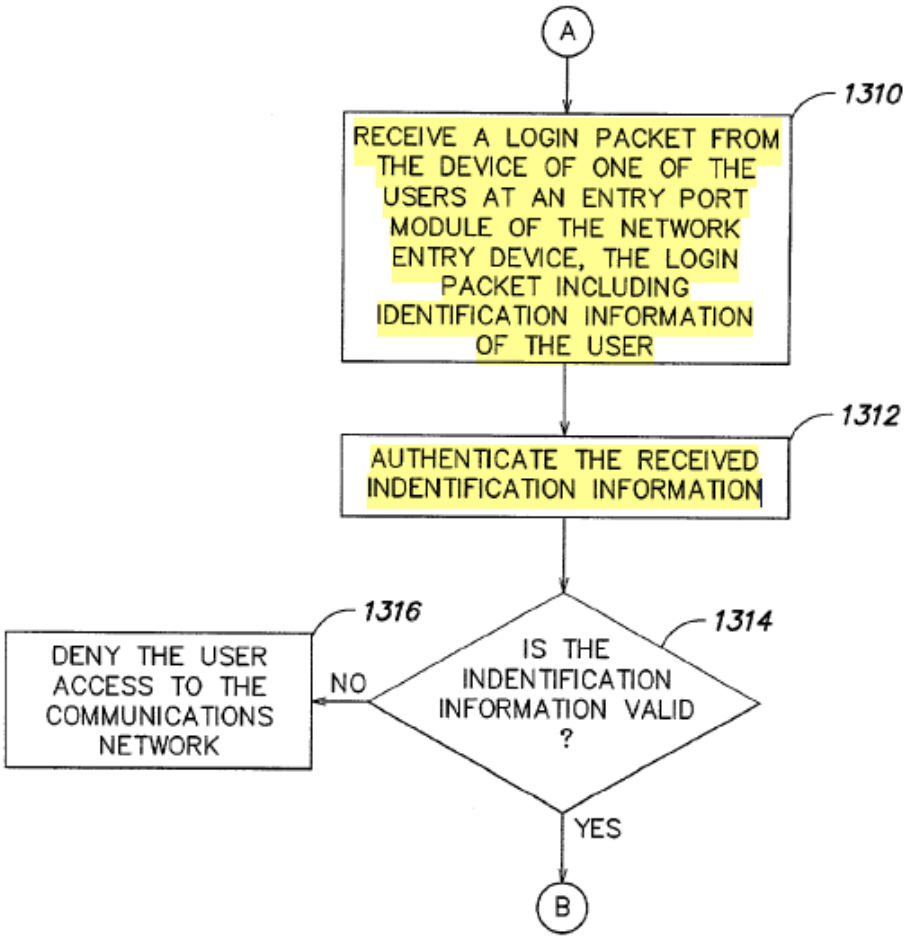
A
1310
RECEIVE A LOGIN PACKET FROM THE DEVICE OF ONE OF THE USERS AT AN ENTRY PORT MODULE OF THE NETWORK ENTRY DEVICE, THE LOGIN PACKET INCLUDING IDENTIFICATION INFORMATION OF THE USER
1312
AUTHENTICATE THE RECEIVED INDENTIFICATION INFORMATION
1314
IS THE INDENTIFICATION INFORMATION VALID ?
NO
1316
DENY THE USER ACCESS TO THE COMMUNICATIONS NETWORK
YES
B

FIG. 13A(2)

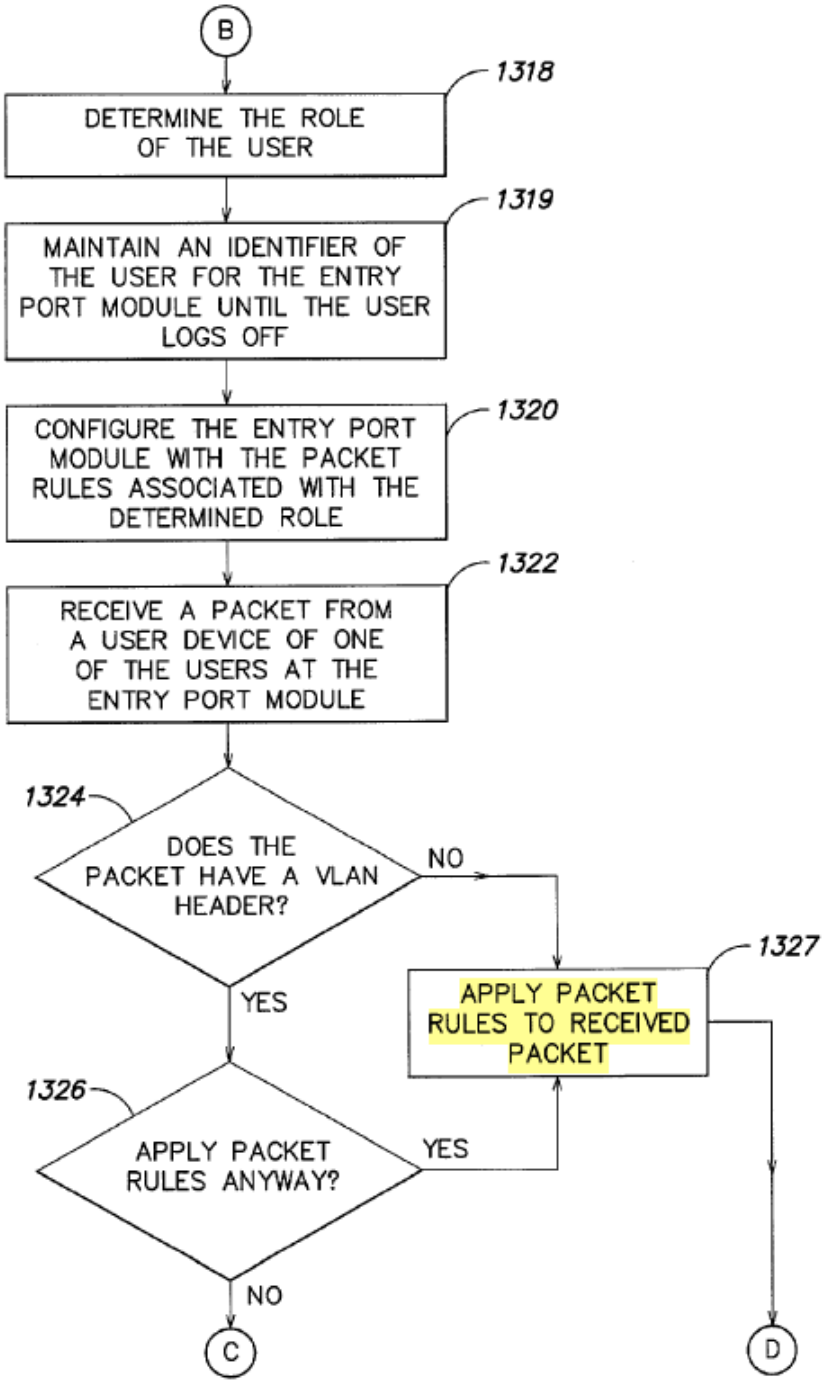

FIG. 13B

Id., Figs. 13A(2) & 13B (annotated).

85. Method 1400 (which Richmond explains is an implementation of block 1327 in Fig. 13B, *see id.*, [0243]) further illustrates actions related to client

discrimination, which are thus a client discrimination stage. For example, the initial action involves "determin[ing] an identity of a user of the user device from which the packet was received" (act 1401), which is client discrimination. *Id*., [0224], Fig. 14A(1).

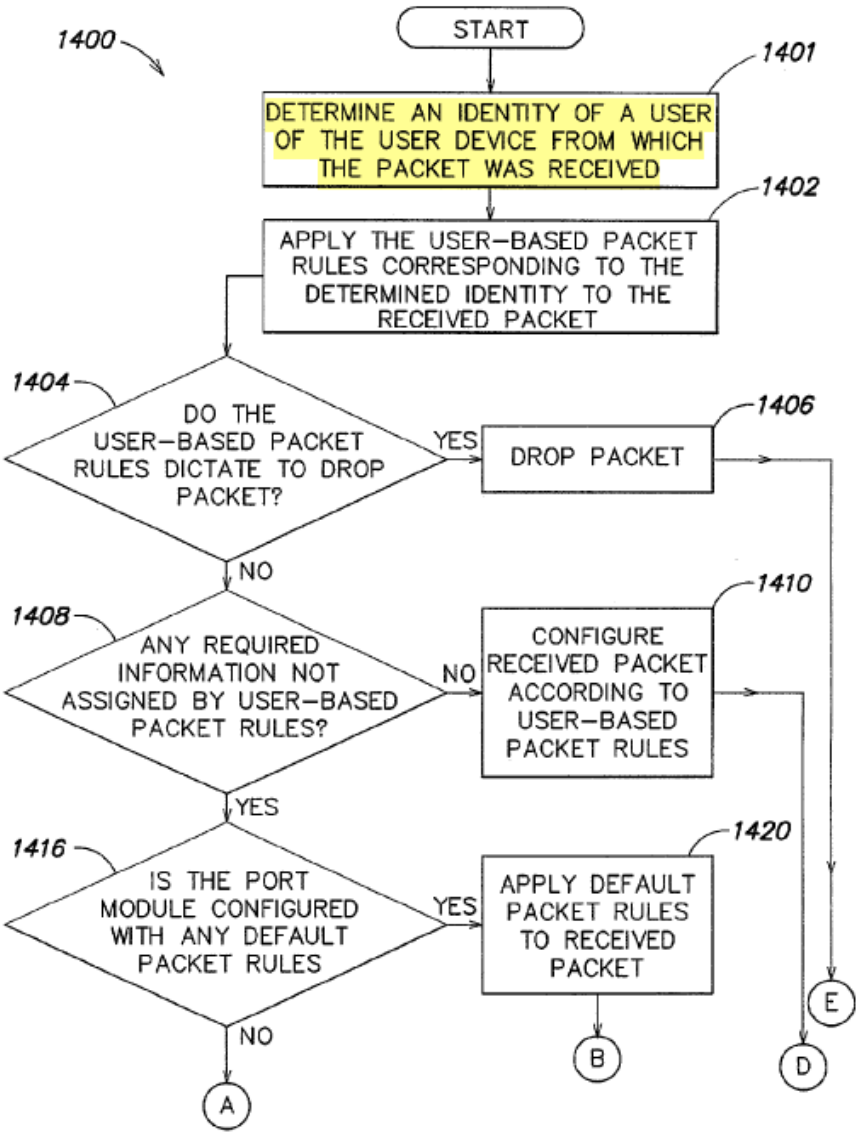

FIG. 14A(1)
FIG. 14A(2)
FIG. 14B

FIG. 14A(1)

Id., Fig. 14A(1) (annotated).

86. **User-specific rule stage**. The '639 Patent includes in the specification examples of the acts that can occur in a client discrimination stage, for example that:

> At user specific rule stage 104, firewall module 82 can access user specific rule 92 based, for example, on the extracted MAC address and IP address. Based on user specific rule 92, firewall module 82 can access firewall parameters 94 and restricted web site list 96 and can process packet 80 accordingly. It should be noted that firewall parameters 94 can include any arbitrary set of parameters that can be used by a firewall to control traffic flow. For example, firewall parameters 94 can specify whether the associated user (e.g., user 12 of FIG. 1) can use voice over IP applications, video conferencing services, transient VPN applications or other applications.

Ex1001, 10:61-11:6. These examples are like that which are disclosed in Richmond. As demonstrated in methods 1300 and 1400, Richmond's system also encompasses various steps associated with the user-specific processing of packets, which are examples of a "user specific rule stage." For instance, if the outcome of the client discrimination stage indicates that the user is not authenticated, the network entry device will enforce the user-specific rule of "deny[ing] the user access to the communications network" (acts 1314 and 1316). *Id*., [0222], Fig. 13A(2). Method 1300 additionally applies user-specific rules by, after establishing the identity and role of the authenticated user, "configur[ing] the entry port module with the packet rules associated with the determined role" (act 1320). *Id*., [0225], [0227], [0278], Fig. 13B. Lastly, the network entry device performs a user-specific rule stage by applying, on packets received at act 1322, "the packet rules corresponding to the user

from which the packet was received (i.e., user-based packet rules)." *Id.*, [0231], Fig. 13B.

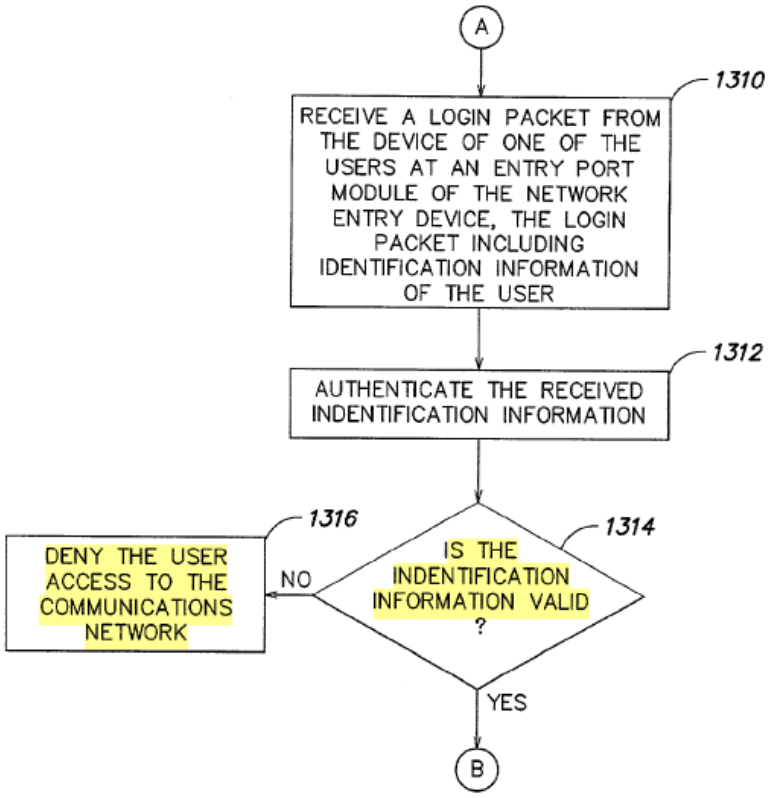

FIG. 13A(2)

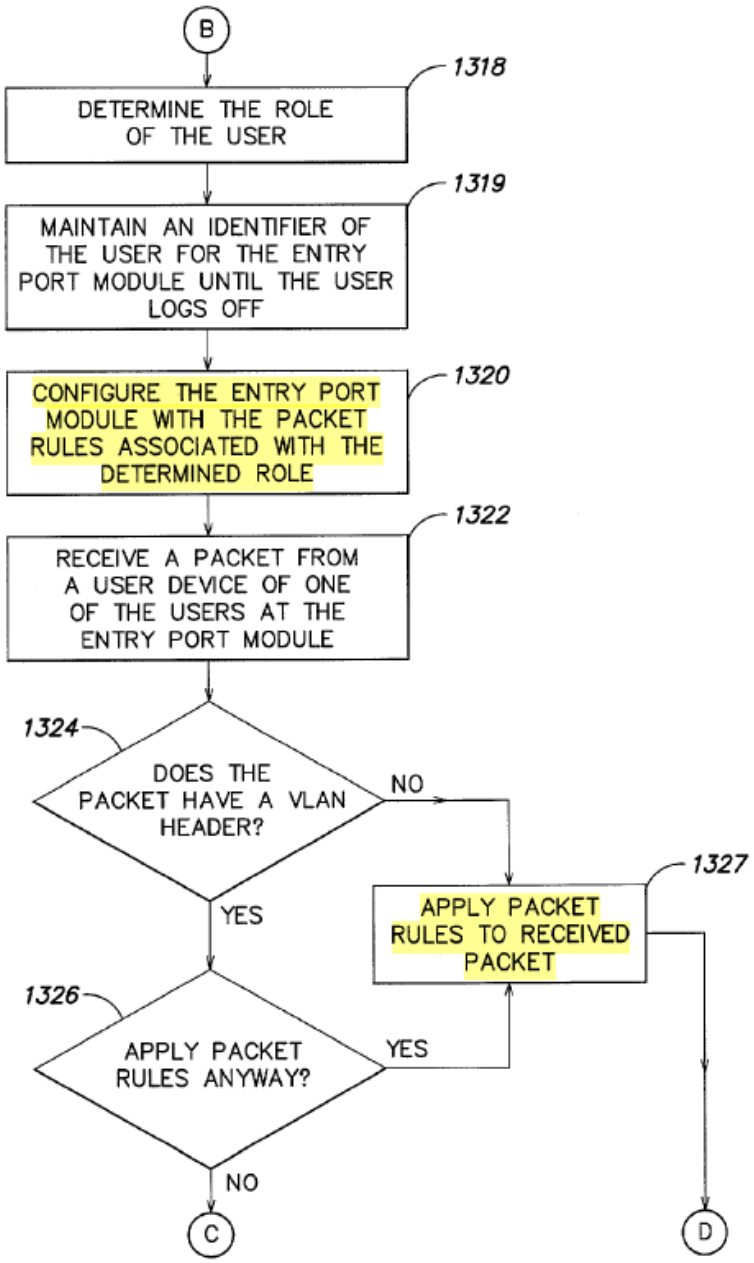

FIG. 13B

Id., Figs. 13A(2) & 13B (annotated).

87. Method 1400 encompasses further user-specific actions for processing a packet, which together form part of a user-specific rule stage. *Id.*, [0243], Fig. 14A(1)-14A(2), 14B. For example, these user-specific actions involve deciding

whether "user-based packet rules dictate to drop [the] packet" (acts 1404 and 1406), determining whether to configure the "received packet …according to user-based packet rules" (acts 1408 and 1410) or based on a combination of user-based rules and other rules (acts 1418, 1422, 1424, and 1426), and assessing whether to apply "rate limiting to the received packet" (acts 1428 and 1430). *Id*., [0247]-[0251]. The user-specific rule stage is explored in greater detail in Section XI.A.4, below.

4. at the client discrimination stage, extracting information associated with the user device from a header of the packet and associating the packet with user specific traffic control rules and user specific firewall rules; and

88. Richmond discloses, at the client discrimination stage (*e.g.*, acts 1310 and 1312 in Fig. 13A(2) and act 1401 in Fig. 14A(1)), the step of extracting information associated with the user device from a header of the packet (*e.g.*, MAC address and IP address) and associating the packet with user-specific traffic control rules and user-specific firewall rules (*e.g.*, Richmond's user-based packet rules).

89. As an illustration of how Richmond satisfies this limitation, in acts 1310 and 1312, when a user attempts to gain access to the network controlled by a network entry device, the user's device transmits a "login packet" to a port module of that network entry device. Ex1003, [0217]. "[T]he login packet includ[es] identification information of [i.e., associated with] the user" who sent it, such as "a user name, password, and other credential information corresponding to the user." *Id*., [0215], [0217], [0219]. The login packet also incorporates "location identifiers

such as a MAC address or IP address or a device identifier." *Id*., [0224]. These location identifiers are linked to the user device, as they serve to identify the user's device within the network.

90. A POSITA would recognize that the portion of the packet containing the MAC address or IP address is to the packet header (either the Ethernet header or IP header, respectively). As I have discussed above in Section III, packet-switched networking protocols like Ethernet and IP mandate that a source device place the source and destination addresses in fixed positions within the packet header. Consequently, when Richmond indicates that its network entry devices may "examine the MAC source address and/or IP source address included in each received packet" (Ex1003, [0230]), a POSITA would understand this as disclosing that the examined addresses are extracted from the header of the packet.

91. The network entry device extracts this user-identifying information (such as the MAC or IP address) from the packet header to "determine an identity of the user of the user device from which the packet was received." *Id*., [0224]. The network entry device then issues a request to search "an authentication database 1508 … for a user authentication entry 1528 corresponding to the identification information received in the login packet" (act 1312). *Id*., [0220]. If the authentication database determines that "the identification information [is not] valid" (act 1314), the network entry device will "deny the user access to the

communications network." *Id*., [0221]-[0222], Fig. 13B. Conversely, if the information is valid and the user is authenticated, the authentication database will retrieve the role ID corresponding to the user and provide this to the network entry device (act 1318). *Id*., [0223].

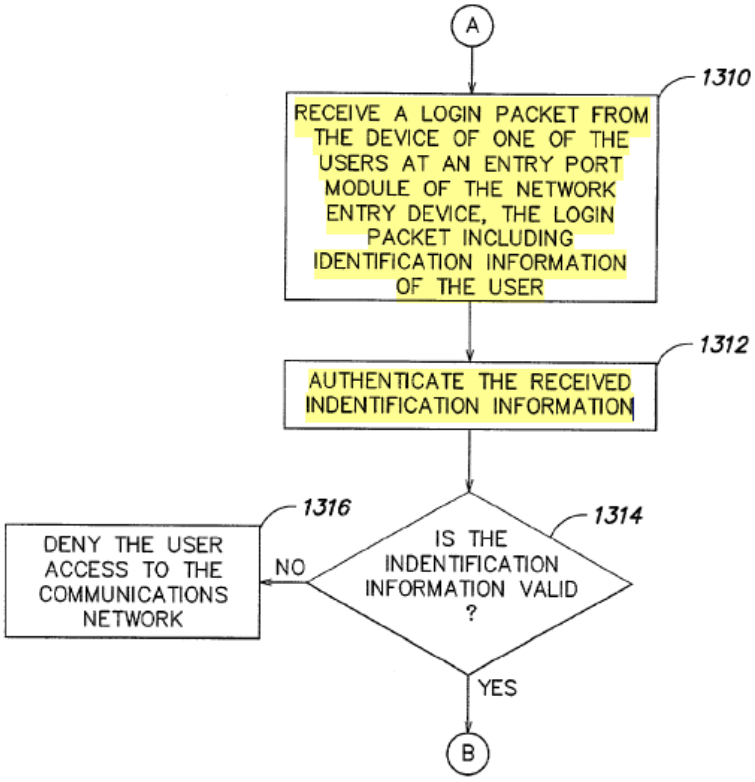

FIG. 13A(2)

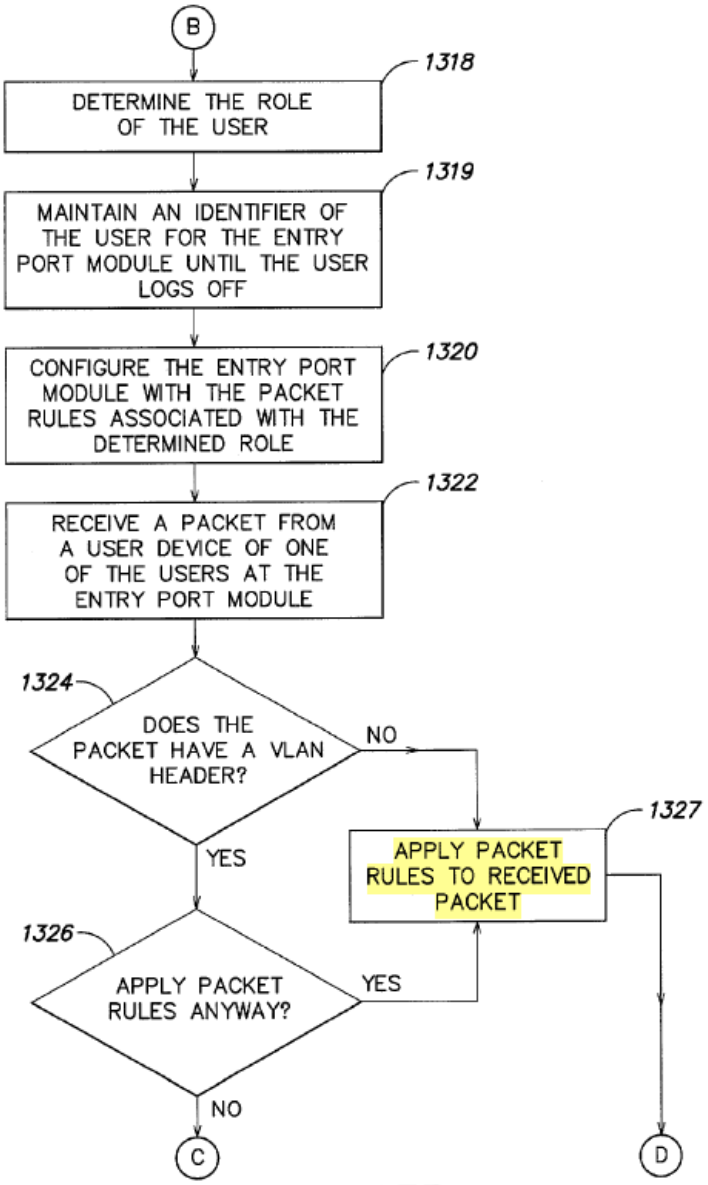

FIG. 13B

Id., Figs. 13A(2) & 13B (annotated).

92. Once the authentication database has authenticated a user and thus verified the user's identity, the network entry device retains a user identifier for that user (act 1319) until the user logs off, such that for later "packets received at the entry port module … the user identifier may be accessed to determine an identity of

the user of the user device from which the packet was received." *Id*., [0224]. Richmond teaches that "the user identifier [that is stored by the network entry device] should include information that can be compared against information included in packets received at the entry port module, for example, a location identifier such as a MAC address or IP address or a device identifier." *Id*., [0224]; *see also id*., [0244] ("Act 1319 may include storing a MAC address and/or an IP address of a user device of a user, where such information was obtained [*i.e.*, extracted] from an authentication packet received from the user device of the user.").

93. Once the network entry device has been provided with a role ID corresponding to a valid user, as retrieved from the authentication database, the network entry device proceeds to associate the packet with user-specific traffic control rules and user-specific firewall rules by utilizing the role linked to that user to search for matching user-based rules within a rule database 1514. *Id*., [0267]. Specifically, "the port configuration logic [of the entry port module 1510] selects and extracts a rule set 1518 from rule database 1514 using the role ID 1517 or the one or more service abstraction IDs" that are tied to the user. *Id*., [0278]. "As a result, port configuration logic receives rule set 1519 and configures port module 1510 accordingly." *Id*., [0278]. A POSITA would recognize that, by setting up port module 1510 according to rule set 1519 based on the received login packet, Richmond's system associates the packet (along with subsequent packets received

at that port module from the user) with the rules contained within rule set 1519.

* * *

94. As a further illustration of how Richmond fulfills this limitation, once subsequent packets are received from a user device (act 1322, *see id.*, [0229]-[0230]), the network entry device will ascertain the "identity of a user of the user device from which the packet was received" (act 1401). *Id.*, [0244].

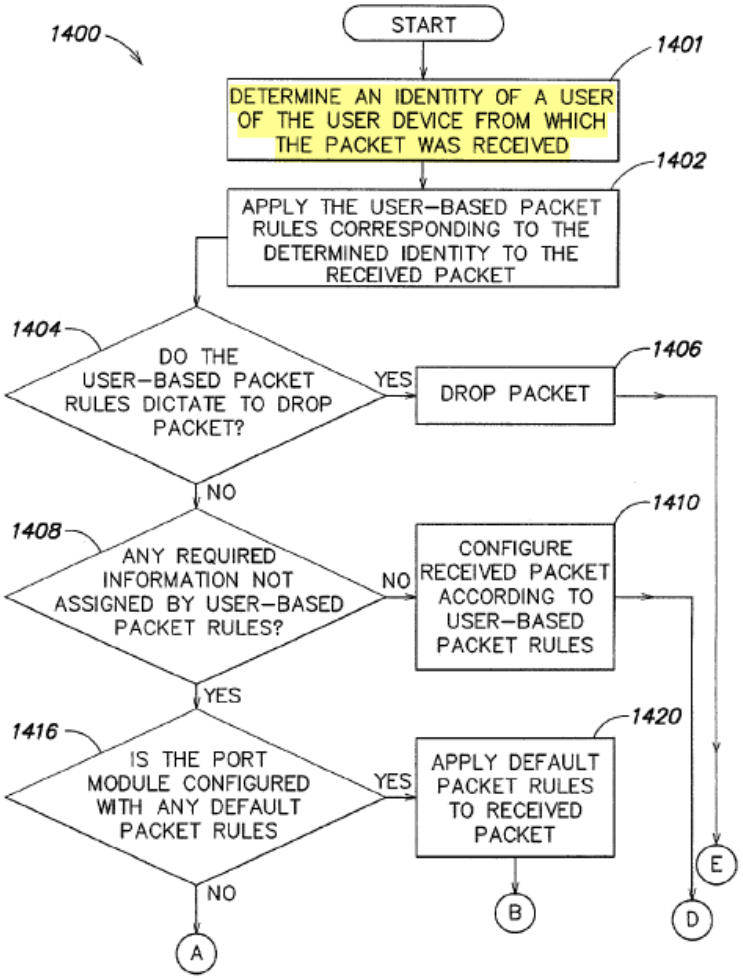

FIG. 14A(1)
FIG. 14A(2)
FIG. 14B

FIG. 14A(1)

Ex1003, Fig. 14A(1) (annotated).

95. Richmond instructs that one approach to achieving this is by "extracting the MAC address and/or IP address of the received packet and comparing it to one or more of the MAC addresses and/or IP addresses maintained on the entry port module." *Id*., [0244]. A POSITA would recognize that this step involves extracting information associated with the user device from the header(s) of the packet, given that the MAC address and IP address are located in the headers, and because these identifiers pinpoint the user device on the Ethernet and IP networks.

96. A POSITA would understand that Richmond is teaching to collect user identifying information is the packet's header, as this is where the MAC and IP addresses (both of which identify a user) are located. For a network that uses IP over Ethernet (as in Richmond), the header of every packet contains both the MAC and IP address, as was well known to a POSITA. Accordingly, a POSITA implementing Richmond's system would understand that it is a straightforward design choice and a straightforward implementation detail to extract information identifying the user from the packet's header.

97. The network entry device then associates the received packet with user-specific traffic control rules and user-specific firewall rules, specifically the user-based rules that correspond to the user's role. *See* Ex1003, [0246] ("The user-based packet rules are packet rules corresponding to an identity of the user, for example,

the one or more packet rules associated with the role or one or more service abstractions corresponding to the user."). As I discuss in greater detail in the next subsection, the processing logic will handle the packets in accordance with the associated user-based rules.

* * *

98. A POSITA would understand Richmond's user-based packet rules to encompass both "user specific traffic control rules" and "user specific firewall rules." Instances of user-specific traffic control rules include those that enforce bandwidth rate limitations on users.

99. As examples of what would be included within each type of rule, the '639 patent at least describes examples of each. First, the specification provides examples of what a "user specific traffic control rule" might be. *See, e.g.*, Ex1001, 8:13-15 ("A traffic rule can include any globally applicable traffic rule: for example, that the total used bandwidth of control device 20 must not exceed a particular amount."), 10:30-37 ("[U]ser specific rule 92 can point to traffic control rule 94 to govern bandwidth provisioning…. As an example, user profile 76 can specify that user 12 is entitled to 128 kbps up and 512 kbps down."), 11:66-67 ("Bandwidth limits can be enforced based on user specific traffic control rules.").

100. Richmond discloses that its user-specific rules can enforce bandwidth rate limitations on users. *See* Ex1003, [0010] ("Controlling usage of network

resources may include … regulating bandwidth on the network consumed by packets received from the user, for example, by assigning priorities to received packets or applying rate limiting to received packets…"), [0159] ("[A]n entry of [packet rule] data structure 500 may specify a value such as 1 MB/sec, for which a network device may be configured to interpret as meaning that the number of bytes that may be transmitted in packets that satisfy the conditions of 508-512 during one second cannot exceed 1 MB…. The consumption of bandwidth by a user also may be limited by the priority specified in field 516."). For example, Richmond indicates that rule entry 502 in its rule data structure 500 may include a "[r]ate limit field 520" that "store[s] a value representing a rate limit to be applied to a packet if it satisfies the rule." *Id*., [0158].

101. The '639 Patent further explains that examples of "user-specific firewall rules" include those that prevent traffic to or from specific destinations or sources. *See, e.g.*, Ex1001, 8:10-12 ("A firewall rule can include any firewall rule known in the art, such as, for example, that all packets destined for a particular web site will be blocked.").

102. Consistent with this description, Richmond describes that its user-specific rules may encompass firewall rules designed to block traffic to or from particular destinations or sources. *See, e.g.*, Ex1003, [0010] ("Controlling usage of network resources may include, … denying one or more packets access to any

network resources beyond a network device (i.e., dropping the packet); … [and] denying access to certain network resources …"), [0035]-[0036] ("The firewall may be configured to intercept such packets and determine whether to allow the packets to proceed on to the Internet…"), [0112] ("Such packet rules may be provisioned to a number of network entry devices and may serve as a distributed firewall for users of a network, as opposed to a centralized firewall. Such distributed firewall reduces use of network resources from users exchanging packets with resources external to the network by applying packet rules at each user's entry point to the network…"), [0230] (disclosing rule that "drop[s] any received packets that are not from a specific user's device"). This is consistent with a POSITA's understanding of a firewall. *See* Ex1009, 410-411 (explaining that prior art firewalls comprised "packet filters" that "allows every incoming or outgoing packet to be inspected," wherein "[p]ackets meeting some criterion are forwarded normally," whereas "[t]hose that fail the test are dropped"). Richmond discloses that its rule data structure 500 may feature a "Deny field 518" that "may specify that a packet is to be dropped." *Id*., [0157]. "For example, a packet rule may specify that if a logical port number identified in the packet corresponds to a particular application, the packet is to be dropped." *Id*., [0157]. "Using field 518, a network administrator may configure a packet rule to prevent certain users from using particular applications such as sensitive business applications running on the network or games residing on the network." *Id*., [0157].

5. at the user specific rule stage, accessing the user specific traffic control rules and user specific firewall rules based on the extracted information associated with the user device and applying the user specific traffic control rules and the user specific firewall rules to the packet as governed by at least one user specific class of service rule associated with the user on the user device in the first network.

103. Richmond discloses, at the user specific rule stage (*e.g.*, acts 1402 and later in method 1400), accessing the user specific traffic control rules (*e.g.*, Richmond's user-based packet rules that limit a user's bandwidth allocation) and user specific firewall rules (*e.g.*, Richmond's user-based packet rules that limit access to specific websites) based on the extracted information associated with the user device (*e.g.*, MAC or IP address) and applying the user-specific traffic control rules and the user-specific firewall rules to the packet as governed by at least one user-specific class of service rule associated with the user (*e.g.*, service abstraction assigned to the user's role(s)) on the user device in the first network (*e.g.*, Network 100).

104. This limitation involves executing the steps based on the previous limitations. As I discussed above in Section IX.A.4, once a user submits a login packet and is successfully authenticated, the entry port module is configured with user-based packet rules that correspond to that user's associated roles and service abstractions. Afterward, the user's device transmits additional packets, which are received by the entry port module of the network entry device (act 1322). Ex1003, [0229]. Within Richmond's system, the port modules of the network entry devices

(serving as exemplary control devices) contain "processing logic … for processing packets sent to and received at the port module." *Id*., [0014]; *see also id*., [0270]-[0271] (same).

105. Upon receiving a packet, the port processing logic 1516 of the network entry device retrieves those user-based packet rules to apply them to the incoming packets. *Id.*, [0270], [0231]. The port processing logic identifies which rules to implement by comparing the extracted information linked to the user device (such as the packet's source or destination MAC address or IP address) to fields within the rule. *See, e.g., id.*, [0033], [0224], [0230], [0241]. "For example, one of the packet rules associated with the role of such user may be defined to examine the MAC source address and/or IP source address included in each received packet and to drop the packet if it is not equal to the MAC address and/or IP address, respectively, of the user…." *Id*., [0230].

106. As depicted in acts 1402 and beyond in Figures 14A(1), 14A(2), and 14B, the port processing logic applies the various user-based packet rules to the incoming packet. As one instance of applying a user-based packet rule, the logic will assess whether the user-based packet rules mandate dropping the packet (act 1404), and if so, it will drop the packet (act 1406), resulting in the blocking of that network traffic. *Id*., [0247].

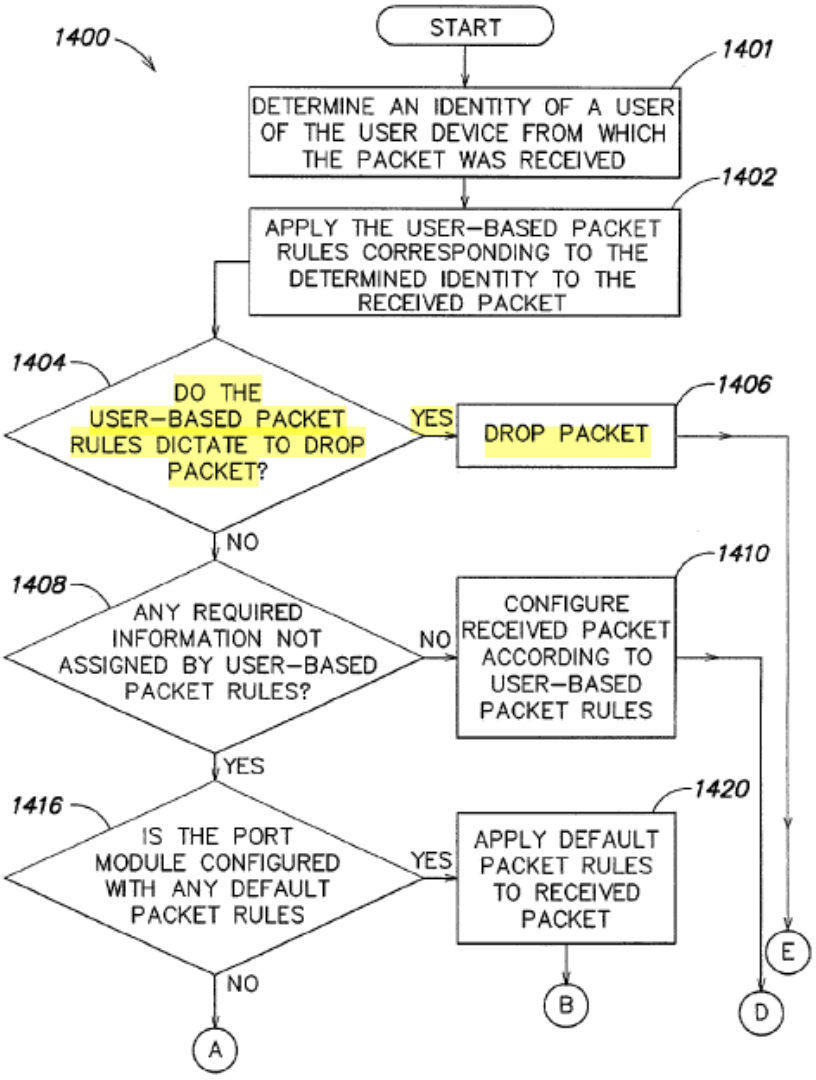

Id., Fig. 14A(1) (annotated).

107. A user-based rule that drops packets has its "deny field 518" configured to "deny access to the resources of [the] communications network by a packet that satisfies [the] conditions" set out in the remaining parts of the rule, such as the source or destination address or source or destination port numbers. *Id.*, [0157]; *see also*

id., [0166]. The deny field can be utilized to block a user from "using particular applications such as sensitive business applications … or games." *Id.*, [0157]. This type of rule qualifies as a "user specific firewall rule."

108. As another example of applying a user-based packet rule, the processing logic will evaluate whether the user-based packet rules indicate that rate limiting should be imposed on the received packet (act 1428), and if they do, the network entry device will apply rate limiting to reduce the rate of network traffic (act 1430) rather than promptly forwarding the packet to its next destination (act 1432). *Id.*, [0159], [0249]-[0251].

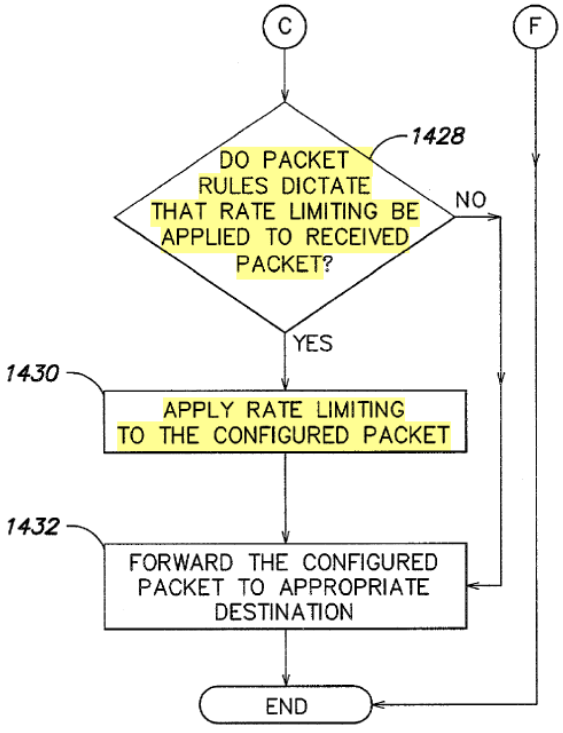

FIG. 14B

Id., 14B (annotated).

109. Such user-specific rules may specify, for instance, "a threshold volume of bytes that may be received [from a user device] during a specified temporal interval (e.g., one second)." *Id.*, [0158]; *see also id.*, [0159] (providing an exemplary traffic control rule that specifies "the number of bytes that may be transmitted … during one second cannot exceed 1 MB"). When this type of rule is in effect, the network control device may "drop some or all of the bytes of a packet … that exceeds the threshold amount during the unit interval." *Id.*, [0159]. These rules qualify as user-specific traffic control rules, as they regulate the amount of traffic and the rate of traffic that a user can send or receive over the network.

110. Richmond's user-specific rules (both firewall and traffic control rules) are directed by service abstractions, which serve as exemplary user-specific class of service rules associated with the user. As discussed, Richmond's service abstractions enable a network administrator to define a class of service for applications or users by compiling a group of individual packet rules. *Id.*, [0116]-[0118]. For example, as depicted in Figure 2, various packet rules can be allocated to create a "high-priority" class of service (items 208 and 212) prioritizes traffic for specific applications. On the other hand, other classes of service might restrict certain applications (e.g., service 210 labeled "NO SNMP" blocks SNMP traffic). *Id.*, [0141].

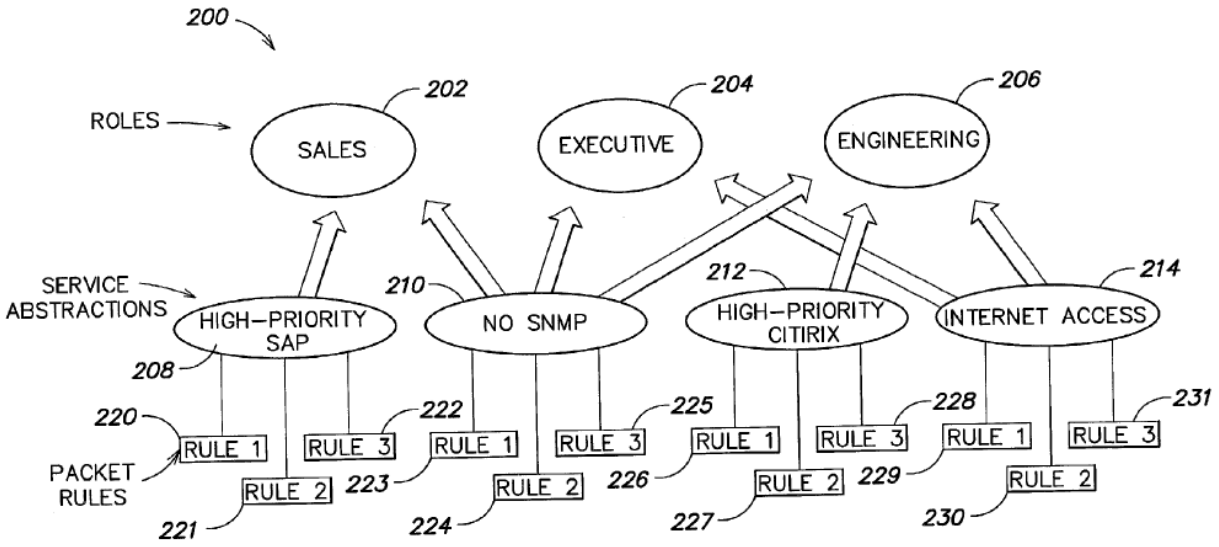

FIG. 2

Id., Fig. 2.

111. The service abstractions (and thus the packet rules) that are used to process a packet will be based on the role of the user. *See id.*, [0121], [0132], [0139]. Thus, the application of the corresponding rules is user specific.

B. Dependent Claim 2

1. The method according to claim 1, wherein the plurality of stages further comprises an interface specific stage, further comprising: at the interface specific stage, applying interface specific rules to the packet based on a type of a network interface over which the packet is received.

112. Richmond discloses and renders obvious claim 1, as discussed. Richmond further discloses and renders obvious that the plurality of stages further comprises an interface-specific stage, and that the method further comprises, at the interface-specific stage, applying interface-specific rules (*e.g.*, authentication states, default rules, other rules) to the packet based on a type of network interface over

which the packet is received (*e.g.*, dedicated/shared wired or wireless).

113. Within Richmond's system, each network interface is managed by a "port module" corresponding to that specific interface. Ex1003, [0014]. Each port module (including "entry port modules" that receive packets from user devices) can be configured independently to enforce a set of rules that are not tied to individual users on traffic passing through that network interface. *Id.*, [0014], [0023], [0041]. In other words, these are deemed "interface specific rules" rather than "user specific rules." *See id.*, [0231] (distinguishing "user-based packet rules" from "default rules"), [0246] ("A user-based packet rule is distinguished from a default packet rule or any other type of rule that is not based on an identity of a user.").

114. As one example of interface-specific rules, each packet at a port module might also have "other additional rules with which the entry port module is configured" applied to it. *Id.*, [0253]; *see also id.*, [0256] (same), [0253] (such rules may include "rules in accordance with … IEEE 802.1D and/or 802.1Q"), [0257] ("other rules may have applied rate limiting to the packet"). Again, these are applied irrespective of the user that sent the packet, such that they are not user-specific rules, but instead interface-specific rules.

115. As another instance of interface-specific rules, a port module (which is specific to an interface) can be configured with a state labeled "authentication off/port on" (which enforces a rule permitting all packets received at the port

module's network interface to pass irrespective of authentication), "authentication on/port off" state (which enforces a rule blocking packets received at the port module's network interface if not authenticated), or "authentication on/port on with default behavior" (which enforces a rule allowing unauthenticated packets to pass for "some basic network services, such as access to the Internet or access to a word processing application," or "assign[ing] a lower priority"). *Id.*, [0208]-[0210]. A POSITA would understand that it is logical to choose this setting (and therefore set the interface specific rule to be applied at the interface specific stage) according to the interface type. For example, authentication is more significant for a wireless connection, where any user that is in the general vicinity of the wireless access point (even outside of the building) can potentially gain access to the network, as compared to a wired access point which requires physical access to the network port.

116. As a third example of interface-specific rules, a port module may be configured with a "default rule," which applies rules to every packet received at that port module (independent of the particular user that sent the packet). *See, e.g., id.*, [0252] (default rule may "assign a default priority to all packets received at the entry port module"), [0254] (act 1416 applying default rules).

117. For any of the aforementioned interface-specific rules, the stage of processing dedicated to them constitutes an "interface specific stage." *See, e.g., id.*, Fig. 14A(1) (act 1420), Fig. 14A(2) (acts 1418, 1424, 1426). Richmond discloses

that the specific interface-specific rules to be applied can differ depending on the type of network interface, such as a dedicated wired, shared wired, or shared wireless communication media. *See, e.g., id.*, [0021]-[0023] (explaining that port modules for wired media and for wireless media process packets differently), [0207], [0263]-[0264], [0270], [0288].

118. To the extent that the claim limitation is interpreted to require different rules are to be applied to different interfaces based on the different protocols/media used for each interface (*e.g.*, one set of interface-specific rules for dedicated wired interfaces and another set of interface-specific rules for shared wireless interfaces), such an implementation would have been obvious to a POSITA. A POSITA would have known about the need to deal with different scenarios that arise for different types of network protocols, based on the unique characteristics of each protocol. Indeed, Richmond contains express teaching about the benefits of applying different sets of rules based on the type of network protocol. For example, Richmond teaches that where an "entry port module [is] connected to a plurality of user devices by a shared transmission medium [*e.g.*, WiFi, as opposed to a dedicated transmission medium] … it may be desirable to configure the port module to drop any received packets that are not from a specific user's device until such user logs off of the port module." Ex1003, [0230]. This consideration is specific to shared transmission media (such as WiFi) because such connections are not limited to the specific user

that has physical access to a specific network port. In a dedicated wired transmission media, for example an Ethernet connection that connects to a single port in a secured office room, it can be assumed that the only user who is likely able to access to that connection is the individual employee that has physical access to the office. However, a WiFi transmission medium can be accessed by any user who is in the general proximity of the WiFi access point, even outside of the building. Thus, a POSITA would understand the need to apply different rules to different types of interfaces based on the type of media, at least so as to account for different security risks. Richmond's system needs no modification in order to apply such rules; Richmond allows for different port modules to be configured differently on a port-by-port basis, such that the network administrator only needs to set the interface specific rules for each port based on the type of network protocol or medium.

C. Dependent Claim 3.

1. The method according to claim 2, wherein the plurality of stages further comprises a global stage, further comprising: at the global stage, applying global firewall rules to the packet.

119. Richmond discloses and renders obvious claim 2, as discussed. Richmond further discloses and renders obvious that the plurality of stages further comprises a global stage, and that the method further comprises, at the global stage, applying global firewall rules to the packet. My opinion would not change if Patent Owner's proposed construction of "global stage" ("a global stage, distinct from the

client discrimination stage and distinct from the user specific stage") is adopted.

120. Richmond discloses that all packets passing through a network entry device is received at and processed by a port module, which, as outlined for claim 2, will enforce a series of interface-specific rules on all traffic entering that port module (e.g., authentication state, default rules, other rules). *See, e.g.*, Ex1003, [0208]-[0210], [0253]-[0257]. As I have discussed for claim 1 (see Section IX.A.4, supra), the rules may encompass firewall rules, such as restricting certain applications or devices for unauthenticated users. *See id.*, [0210], [0212]. A POSITA would recognize that when all port modules are set up with the same firewall rule (e.g., when all port modules are configured to the “authentication on/port off" blocking traffic for unauthenticated users, *see id.*, [0209]), this represents a global firewall rule, and the mechanism applying it constitutes a global stage. Richmond describes these authentication-stage rules as being distinct from user-based packet rules. Thus, a POSITA would understand that their application occurs in a stage distinct from the client-discrimination stage (*e.g.*, act 1401) and the user-specific rule stage (*e.g.*, acts 1402 and later). Accordingly, Richmond meetings the limitation under Patent Owner's construction.

121. Richmond describes the implementation of a global firewall to enforce identical rules on all packets, noting that this approach was standard in prior art. *See id.*, [0035] ("Use of a firewall (e.g., firewall 140) is another example of a technique

for controlling usage of network resources…"), [0036]-[0037]. A POSITA would understand that a conventional firewall imposes global rules on each packet, in contrast to user-specific or interface-specific rules. *See* Ex1009, 411 ("Each packet filter [in a firewall] … allows every incoming or outgoing packet to be inspected. Packets meeting some criterion are forwarded normally. Those that fail the test are dropped…. Packet filters are typically driven by tables configured by the system administrator. These tables list list sources and destinations that are acceptable, sources and destinations that are blocked, and default rules about what to do with packets coming from or going to other machines."). A POSITA would additionally appreciate that a firewall offers considerable versatility and flexibility in the kinds of rules that can be established, and that these would include global rules (i.e., those enforced on all packets, as was conventional in prior art, rather than Richmond's user-specific or interface-specific rules). A POSITA would further understand that the traditional firewall disclosed in Richmond processes its global rules in a stage that is distinct from the stages that process Richmond's user-based rules.

122. Furthermore, Richmond renders obvious that its network entry devices incorporate a global stage for enforcing global firewall rules. A POSITA would be inclined to incorporate traditional global firewall rules alongside Richmond's user-specific and interface-specific rules, since global rules facilitate easier network configuration. For instance, if a network administrator aims to restrict all traffic to

specific websites or applications, irrespective of the receiving interface, it would be more straightforward to establish a single global firewall rule rather than configuring interface-specific firewall rules for every port module. This modification would represent a combination of prior art elements (e.g., Richmond's network entry device and traditional firewalls) that produce predictable outcomes. Additionally, Richmond indicates that its teachings are suitable for traditional firewalls, and that traditional firewalls can coexist with Richmond's disclosed network entry devices. Ex1003, [0205]-[0206], [0262]. A POSITA would also recognize that it would be an obvious implementation detail to process these global rules in a separate stage from the user-specific rules, because such rules do not need to discriminate clients or otherwise access user-specific rules or other user-specific information. Ex1002 ¶122. It would be straightforward for a POSITA to implement such a system with distinct stages, with a reasonable expectation of success, because this distinct-stage implementation is the same type of implementation used in traditional firewalls that do not discriminate clients.

D. Dependent Claim 4 (Ground 2 Only)

1. The method according to claim 1, wherein the packet is processed utilizing a provisioning module running on the control device, wherein the provisioning module comprises a first table for interface specific rules, a second table for global rules, and a third table for user specific rules.

123. Richmond discloses and renders obvious claim 1, as discussed.

Richmond further renders obvious that the packet is processed utilizing a provisioning module running on the control device (*e.g.*, network interface device), wherein the provisioning module comprises a first table for interface-specific rules, a second table for global rules, and a third table for user-specific rules.

124. Richmond discloses that its rules are stored in a "packet rule data structure 500," which "may be a table of a relational database." Ex1003, [0150], Fig. 5. A POSITA would understand that a relational database comprises an arbitrary number of tables, each with its own fixed set of columns (individual data fields) and variable number of rows (data entries). *Id.*, Figs. 5-7. Richmond provides evidence of this understanding, by disclosing that the service abstractions and roles may be stored in separate tables within the relational database. Ex1003, [0170], [0175]., Figs. 6-7. Moreover, it was well-known to a POSITA that prior art firewalls used multiple tables to store packet rules. Ex1009, ("Packet filters are typically driven by ***tables*** [*i.e.*, plural] configured by the system administrator. These tables list sources and destinations that are acceptable, sources and destinations that are blocked, and default rules about what to do with packets coming from or going to other machines.").

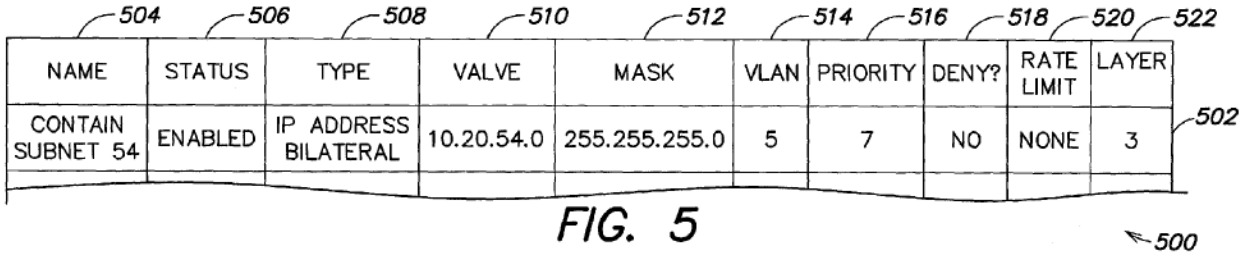

504 NAME	506 STATUS	508 TYPE	510 VALVE	512 MASK	514 VLAN	516 PRIORITY	518 DENY?	520 RATE LIMIT	522 LAYER
CONTAIN SUBNET 54	ENABLED	IP ADDRESS BILATERAL	10.20.54.0	255.255.255.0	5	7	NO	NONE	3

502

FIG. 5

500

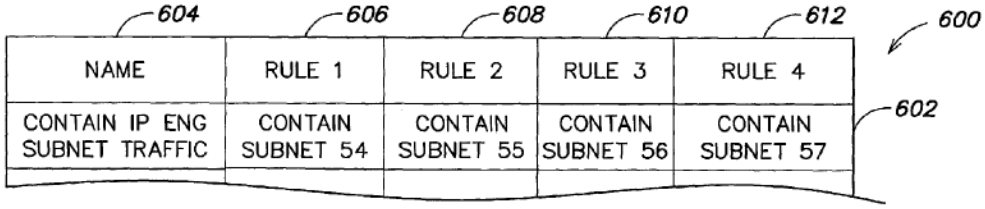

604 NAME	606 RULE 1	608 RULE 2	610 RULE 3	612 RULE 4
CONTAIN IP ENG SUBNET TRAFFIC	CONTAIN SUBNET 54	CONTAIN SUBNET 55	CONTAIN SUBNET 56	CONTAIN SUBNET 57

600

602

FIG. 6

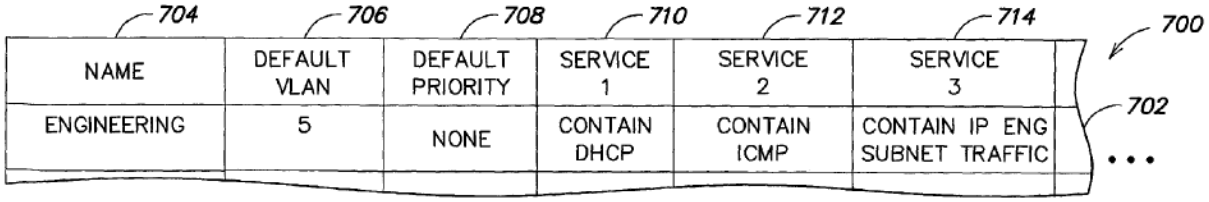

704 NAME	706 DEFAULT VLAN	708 DEFAULT PRIORITY	710 SERVICE 1	712 SERVICE 2	714 SERVICE 3
ENGINEERING	5	NONE	CONTAIN DHCP	CONTAIN ICMP	CONTAIN IP ENG SUBNET TRAFFIC

700

702

FIG. 7

125. It would be an obvious and routine implementation detail in light of Richmond's disclosures for a POSITA to implement Richmond's system by storing each type of rule (interface-specific, global, and user-specific) in its own table of a relational database. As an initial matter, a POSITA would understand that Richmond's rule sets could be stored in the database in a variety of ways; there is nothing particular about Richmond's system that would require that all types of rules be stored in a single table. Instead, Richmond leaves this as an implementation detail. Moreover, a POSITA that is looking to implement a network entry device according to Richmond's disclosure would understand that it would be more efficient to store each type of rule in independent tables (consistent with the claim), because

not all of the rules require the same information. For example, a POSITA would recognize that global rules and interface-specific rules do not need to store a user's MAC or IP address (*e.g.*, field 510), because those fields are used in Richmond to identity the user that sent a packet in order to apply user-specific rules, but the global rules and interface-specific rules do not depend on the user's identity. Moreover, a POSITA would recognize that a global rule does not need to store the interface with which it is associated, because it is applied on all packets regardless of which interface they are received on. Thus, by using separate tables in a relational database, the POSITA implementing the system can optimize storage by defining each table to just store the columns/fields that are appropriate for that rule type. A POSITA would recognize that the substitution of three relational database tables in place of a single relational database table would be a simple substitution of known elements. A POSITA would have a reasonable expectation of success, because a POSITA knows that it is simple to create and access additional tables in a relational database (such as that disclosed in Richmond) according to well-known, conventional techniques.

E. Dependent Claim 5

1. The method according to claim 1, wherein the user device is assigned with an Internet Protocol (IP) address in the first network and wherein the at least one user specific rule is associated with the IP address.

126. Richmond discloses and renders obvious claim 1, as discussed.

Richmond further discloses and renders obvious that the user device is assigned with an Internet Protocol (IP) address in the first network and wherein the at least one user-specific rule is associated with the IP address.

127. As an initial matter, it is the case in all Internet Protocol-based networks that each networked user device is allocated an IP address, a fact well-known to a POSITA. This would be the case both prior to the priority date of the '639 Patent, and through to the present date. It is essential for networked devices to possess an IP address to enable its identification within that network. Every IP packet transmitted to or received from the networked device will include the device's IP address in its destination or source field, as is necessary for the packet to be correctly routed and recognized.

128. Richmond's user devices send and receive packets that contain the IP address assigned to them, aligning with the POSITA's understanding of what is inherent in all IP-based networks. See Ex1003, [0224] (explaining that "information included in packets received at the entry port module" includes "for example, a location identifier such as a MAC address or IP address or a device identifier"), [0230] (disclosing that there is a "MAC Source address and/or IP Source address included in each received packet," which can be compared to "the MAC address and/or IP address, respectively, of [i.e., assigned to] the user").

129. In one embodiment of Richmond's system, user-based packet rules are

defined in packet rule data structure 500, which contains entries 502 corresponding to individual rules, and in which each entry contains various fields. *Id.*, [0150]-[0151]. Richmond provides an example rule entry in Figure 5. In that example rule, field 508 is set to "IPAddress Bilateral," which "specif[ies] that the source and destination (i.e., bilateral) IP addresses of a packet are to be examined" in order to see if the rule should be applied. *Id.*, [0162]. As also shown in Figure 5, field 510 "specifies that the value to be compared to the IP addresses of the received packet [in order to determine whether to apply the rule] is '10.20.54.0.'" *Id.*, [0163]; *see also id.*, [0154]. Thus, the rule shown in Figure 5 is associated with a particular IP address.

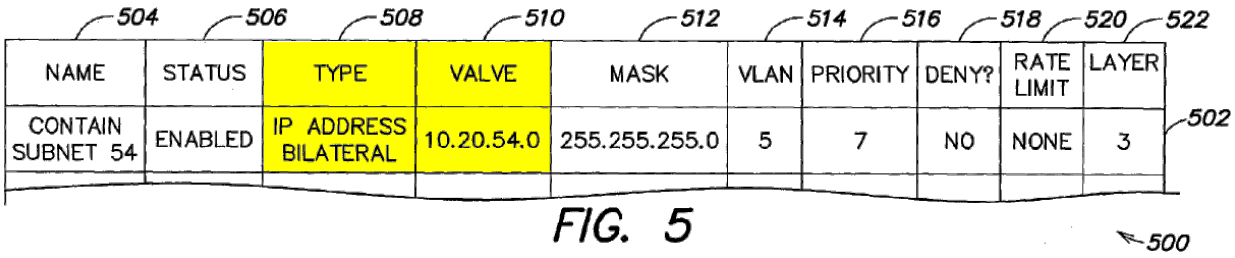

FIG. 5

Ex1003, Fig. 5 (annotated).

130. Richmond offers further instances of rules associated with a user's IP address; for example, "one of the packet rules associated with the role of such user may be defined to examine the MAC Source address and/or IP Source address included in each received packet and to drop the packet if it is not equal to the MAC address and/or IP address, respectively, of the user." *Id.*, [0230]. Richmond further

provides another example with a similar rule that involves comparing the address to those previously stored from the authentication packet. *Id.*, [0244].

F. Dependent Claim 6

1. The method according to claim 5, wherein the user device is associated with a media access control (MAC) address and wherein the at least one user specific rule is further associated with the MAC address.

131. Richmond discloses and renders obvious claim 5, as discussed. Richmond further discloses and renders obvious that the user device is associated with a media access control (MAC) address, and wherein the at least one user-specific rule is further associated with the MAC address.

132. As an initial matter, it is the case in all Ethernet-based networks that each networked user device is assigned a MAC address, a fact well-known to a POSITA. This would be the case both prior to the priority date of the '639 Patent, and through to the present date. It is essential for networked devices to possess a MAC address to ensure its unique identification within that network. Ethernet packets transmitted to or received from the networked device will include the device's MAC address in its destination or source field, in order to properly direct and recognize the packet.

133. Richmond's user devices send and receive packets that contain the MAC address assigned to them, aligning with the POSITA's understanding of what is inherent in all Ethernet-based networks. *See* Ex1003, [0224] (explaining that

"information included in packets received at the entry port module" includes "for example, a location identifier such as a MAC address or IP address or a device identifier"), [0230] (disclosing that there is a "MAC Source address and/or IP Source address included in each received packet," which can be compared to "the MAC address and/or IP address, respectively, of [i.e., assigned to] the user").

134. Richmond provides examples of rules associated with a user's MAC address; for example, "one of the packet rules associated with the role of such user may be defined to examine the MAC Source address and/or IP Source address included in each received packet and to drop the packet if it is not equal to the MAC address and/or IP address, respectively, of the user." *Id*., [0230]. Richmond further provides another example with a similar rule that involves comparing the address to those previously stored from the authentication packet. *Id.*, [0244].

G. Dependent Claim 7

1. The method according to claim 1, wherein the user specific traffic control rules and the user specific firewall rules are established for the user based on a user profile that includes parameters that determine an extent of the user's access to the second network, services available to the user, or a combination thereof.

135. Richmond discloses and renders obvious claim 1, as discussed. Richmond further discloses and renders obvious that the user specific traffic control rules and the user specific firewall rules are established for the user based on a user profile (*e.g.*, the information in relationship database 410 regarding the user's role

and associated service abstractions) that includes parameters that determine an extent of the user's access to the second network, services available to the user, or a combination thereof.

136. As I have discussed above in Section VI.A, the users within Richmond's system are assigned to roles, each of which is linked to associated service abstractions. *Id*., [0116]-[0118]. This data is maintained in relationship database 410. *Id*., [0149]-[0150], [0170], [0175]. The details recorded about a user, encompassing the user's roles and service abstractions, serve as an exemplary user profile. These profiles contain parameters that dictate a user's access to the second network (e.g., the Internet 148). For instance, as illustrated in Fig. 2, users assigned to sales role 202 might be excluded from the "Internet Access" 214 service abstraction, thereby preventing their access to the Internet. In contrast, users in executive role 204 or engineering role 206 are not barred from Internet access. Even for users whose profile permits general Internet access, their connection to that network may be restricted by rules that prohibit certain websites or applications, as shown:

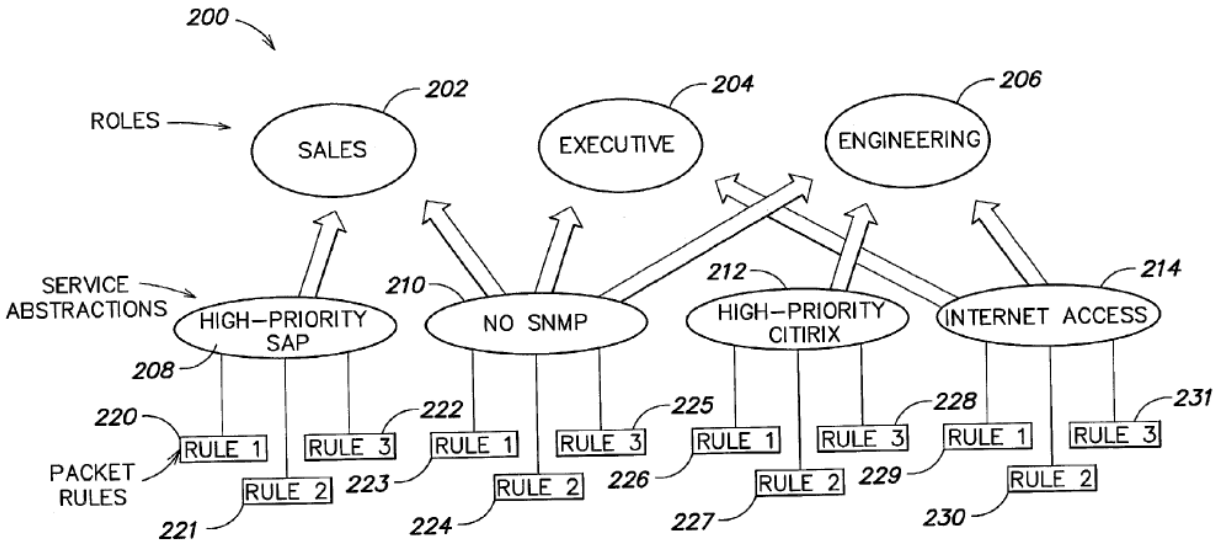

FIG. 2

137. The user profiles additionally contain parameters that determine services available to the user. As also shown in Fig. 2, the "High-Priority SAP" service and "High-Priority Citirix [*sic*]" service may be limited to users in the sales role 202 or engineering role 206, respectively. Additionally, all non-administrators may be blocked from using the SNMP service. *Id.*, [0141], Fig. 2.

H. Dependent Claim 8

1. A method according to claim 1, further comprising determining whether the user has authenticated.

138. Richmond discloses and renders obvious the method according to claim 1, as I have discussed above. Richmond further discloses a method further comprising the step of determining whether the user has authenticated.

139. For example, Richmond discloses its system has an "authentication module 1506," shown in Fig. 15B, and further explains that the authentication

module "may be configured to perform authentication in accordance with one or more authentication technologies." Ex1003, [0274].

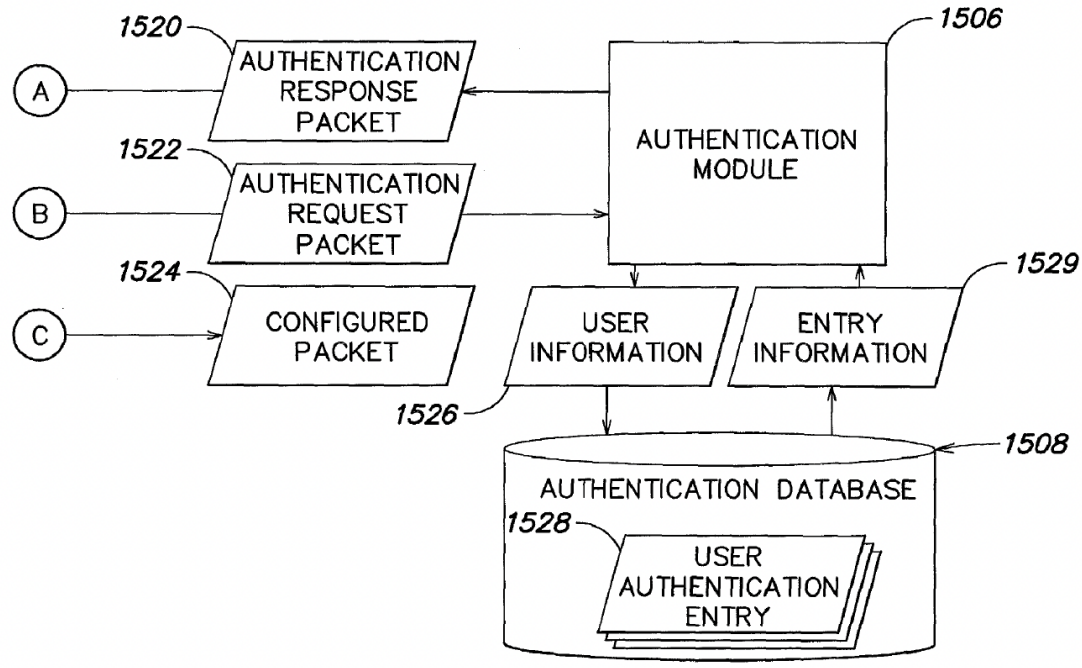

FIG. 15B

140. The network entry device transmits identity information to the authentication module 1506, which maintains a database holding the credentials of each user along with their respective roles. *Id*., [0273]-[0275]. Should the authentication module indicate that the credentials are invalid, the network entry device will block the user's access. *Id*., [0222], Fig. 13A(2) (acts 1314 and 1316). If the authentication module determines that the credentials are valid (meaning that the user has authenticated), the authentication module notifies the network entry device of the user's roles, prompting the device to set up its entry port modules with

user-based packet rules corresponding to the user's identity and the user's roles and service abstractions. *Id*., [0223], [0277]-[0279]. Thereafter, when the user sends packets, the network entry device handles these packets in accordance with the pre-configured user-based rules. *Id*., [0280]-[0281], Fig. 13B (act 1327).

141. Richmond further discloses that "[a] port module may be configured to have a state of 'authentication on/port off,' where a user must authenticate before using any network resources," or alternatively, "a state of 'authentication on/port on with default behavior,' where a user may be allowed to use some network resources before authentication or in the event of a failed authentication." *Id.*, [0209]-[0210]. A POSITA would understand that the network entry device must determine whether the user has authenticated in order to apply the correct behavior according to these states.

142. Additionally, Richmond discloses a packet rule that filters packets based on whether their MAC address and/or IP address matches those of authenticated users. *Id.*, [0230]. Such rule acts to determine whether the user has been authenticated and drops packets of those users that are not.

I. Dependent Claim 9

1. A method according to claim 8, wherein the user has not authenticated, further comprising: redirecting the user to a login web page that requests user credentials; and using the

user credentials received from the user to authenticate the user.

143. Richmond discloses and renders obvious claim 8, as discussed. Richmond further discloses that when the user has not authenticated, redirecting the user to a login web page that requests user credentials and using the user credentials received from the user to authenticate the user.

144. According to Richmond, in its method 1300, the network entry devices will receive "identification information … includ[ing] a user name, password and other credential information corresponding to the user," and that "the received identification information may be authenticated" using "a variety of authentication techniques." Ex1003, [0219]-[0220]. Richmond further states that "[i]n an embodiment of method 1300, various acts of method 1300 … are performed as described in *Enterasys User Personalized Network*, a white paper available from Enterasys Networks, Inc., available on the filing date of this application at http://www.enterasys.com/products/whitepapers, the entire contents of which is hereby incorporated by reference." *Id.*, [0239]. I understand that, because of the incorporation by reference, the disclosures of the Enterasys Whitepaper are considered to be part of the disclosure of Richmond. I understand that Ex1007 is an excerpt of the Richmond file history, which includes a copy of the whitepaper that was submitted to the USPTO on Sept. 20, 2002 along with the Richmond application, as well as an Information Disclosure Statement stating that the

whitepaper had been publicly available as of February 12, 2001. *See* Ex1007, 1-3. Ex1006 is a color copy of the whitepaper.

145. The Enterasys Whitepaper discloses that its system "provid[es] two primary mechanisms for authenticating users" that have not already authenticated, one of which is "the use of web-based authentication." Ex1007, 15. The whitepaper discloses that, for web-based authentication, "[u]sers who want to receive network [s]ervices access [*i.e.*, are redirected to] a secure web page … from a browser [and] are asked to provide username and credentials." *Id.* This is an exemplary login webpage that requests user credentials as required by the claim. The whitepaper further discloses that, after the user has submitted credentials, "[t]hese credentials are then forwarded on to a RADIUS server, which either authenticates them, or forwards the request to the appropriate authenticator," *Id.* Thus, the whitepaper discloses that Richmond's system uses the received user credentials to authenticate the user as required by the claim.

J. Independent Claim 10

146. In my opinion, Richmond discloses and renders obvious all limitations of claim 10. I note that the limitations of claim 10 are nearly identical to those of claim 1. The only possibly substantive difference between claims 1 and 10 are that, while claim 1 recites a method, claim 10 recites a computer program that executes that method when run on a processor. Consequently, in my analysis for this claim,

I largely incorporate by reference the discussion I have provided above for claim 1.

1. A computer program product comprising at least one non-transitory computer-readable storage medium storing computer instructions translatable by a processor of a control device to perform:

147. To the extent the preamble is limiting, it is my opinion that Richmond discloses a computer program product comprising at least one non-transitory computer-readable storage medium storing computer instructions translatable by a processor of a control device to perform the recited steps of claim 10.

148. Richmond discloses that its methods (including methods 1300 and 1400) "may be implemented as a computer program product tangibly embodied as computer-readable signals on a computer-readable medium, for example, a non-volatile recording medium, an integrated circuit memory element, or a combination thereof." *Id.*, [0242]. Richmond further discloses that signals may "define instructions, for example, as part of one or more programs, that, as a result of being executed by a computer, instruct the computer to perform one or more of the methods or acts described [in Richmond]." *Id.*, [0242]; *see also id.* [0260] (disclosing the same for Method 1400), [0047], [0147], [0062] & [0085] (disclosing that various embodiments of Richmond's system can be implemented as computer program products consisting of instructions stored on a storage medium). Richmond further specifies that "[t]he computer-readable medium on which such instructions are stored may reside on one or more of the components of network 100 described above

in relation to FIG. 1 and/or system 1500 described below in relation to FIG. 15 [*i.e.*, the network entry device], and may be distributed across one or more of such components." *Id.*, [0242]. Based on these disclosures, a POSITA would understand Richmond to disclose that the functionality of its network entry devices may be implemented by means of a computer program, which is consistent with a POSITA's general understanding at the time.

2. receiving a packet originating from a user device in a first network, wherein the first network is connected to a second network via the control device, and wherein the user device is associated with a user;

149. I incorporate by reference my discussion of the corresponding limitation of claim 1. *See* Section IX.A.2.

3. processing the packet according to a plurality of stages, including a client discrimination stage and a user specific rule stage;

150. I incorporate by reference my discussion of the corresponding limitation of claim 1. *See* Section IX.A.3.

4. at the client discrimination stage, extracting information associated with the user device from a header of the packet and associating the packet with user specific traffic control rules and user specific firewall rules; and

151. I incorporate by reference my discussion of the corresponding limitation of claim 1. *See* Section IX.A.4.

5. at the user specific rule stage, accessing the user specific traffic control rules and user specific firewall rules based on the extracted information associated with the user device and

applying the user specific traffic control rules and the user specific firewall rules to the packet as governed by at least one user specific class of service rule associated with the user on the user device in the first network.

152. I incorporate by reference my discussion of the corresponding limitation of claim 1. *See* Section IX.A.5.

K. Dependent Claim 11

1. The computer program product of claim 10, wherein the plurality of stages further comprises an interface specific stage and wherein the computer instructions are further translatable by the processor to perform: at the interface specific stage, applying interface specific rules to the packet based on a type of a network interface over which the packet is received.

153. I incorporate by reference my discussion of the corresponding limitations of claim 2. *See* Section IX.B.1.

L. Dependent Claim 12

1. The computer program product of claim 10, wherein the plurality of stages further comprises a global stage and wherein the computer instructions are further translatable by the processor to perform: at the global stage, applying global firewall rules to the packet.

154. I incorporate by reference my discussion of the corresponding limitations of claim 3. *See* Section IX.C.1.

M. Dependent Claim 13

1. The computer program product of claim 10, wherein the computer instructions are further translatable by the processor to implement a control program on the control device, wherein the control program comprises a

provisioning module, wherein the provisioning module comprises a traffic conditioning module for user specific allocation of bandwidth.

155. Richmond discloses and renders obvious claim 10, as discussed above. Richmond further discloses that the computer instructions are further translatable by the processor to implement a control program on the control device, wherein the control program comprises a provisioning module, wherein the provisioning module comprises a traffic conditioning module for user-specific allocation of bandwidth.

156. As discussed for claim 10, Richmond's network device (which is an exemplary control device) may be implemented using a computer program. Ex1003, [0047], [0147], [0062], [0085], [0242], [0260]. Such a program constitutes a control program within the meaning of the claim, as it is a program that controls the transmission of packets.

157. As discussed for claim 1 (*see* Sections IX.A.4 and IX.A.5, *supra*), Richmond's user-based packet rules include rules that enforce bandwidth rate limitations on users, which are based on the limits to which the user is allocated. *See* Ex1003, [0010], [0158] (disclosing that user-based packet rule has a "[r]ate limit field 520 … [that] store[s] a value representing a rate limit to be applied to a packet"), [0159] ("[A]n entry of [packet rule] data structure 500 may specify a value such as 1 MB/sec, for which a network device may be configured to interpret as meaning that the number of bytes that may be transmitted in packets that satisfy the conditions

of 508-512 during one second cannot exceed 1 MB…. Such rate limiting technique may be used to limit the amount of bandwidth that a user may consume….").

158. The control program of the network entry device applies rate limiting according to the user-specific rules. *Id.*, [0158]-[0159], [0249]-[0251]. A POSITA would understand that the provisioning module (with the traffic conditioning module within it) is the hardware and the portion of the software program that implements the corresponding functionality.

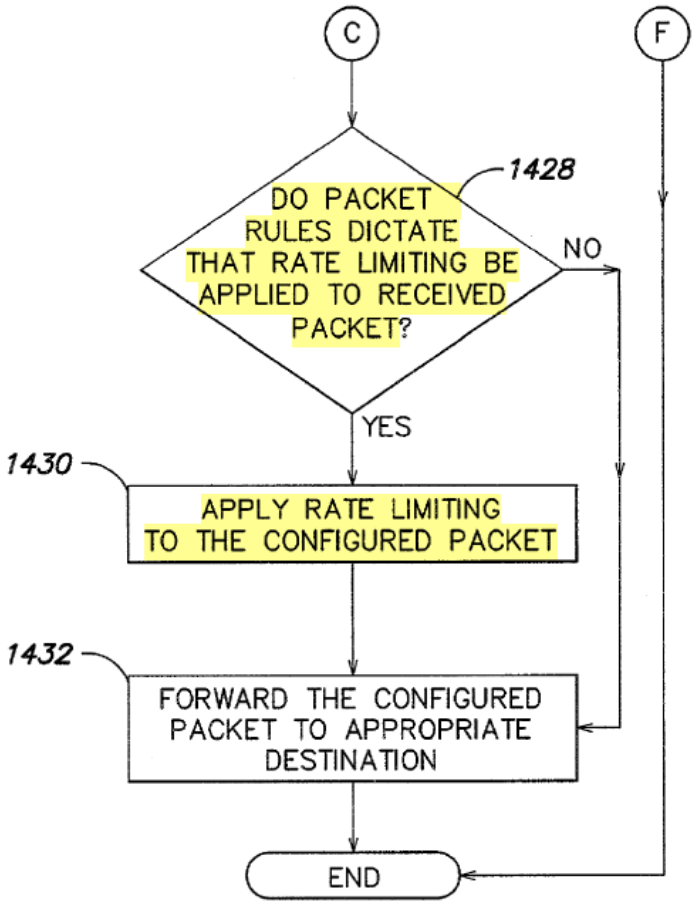

FIG. 14B

N. Dependent Claim 14

1. The computer program product of claim 10, wherein the computer instructions are further translatable by the processor to establish the user specific traffic control rules and the user specific firewall rules for the user based on a user profile that includes parameters that determine an extent of the user's access to the second network, services available to the user, or a combination thereof.

159. I incorporate by reference my discussion of the corresponding limitations of claim 7. *See* Section IX.G.1.

O. Dependent Claim 15

1. The computer program product of claim 10, wherein the computer instructions are further translatable by the processor to determine whether the user has authenticated.

160. I incorporate by reference my discussion of the corresponding limitations of claim 8. *See* Section IX.H.1.

P. Dependent Claim 16

1. The computer program product of claim 15, wherein the user has not authenticated and wherein the computer instructions are further translatable by the processor to perform: redirecting the user to a login web page that requests user credentials; and using the user credentials received from the user to authenticate the user.

161. I incorporate by reference my discussion of the corresponding limitations of claim 9. *See* Section X.A.1.

Q. Independent Claim 17

162. In my opinion, Richmond discloses and renders obvious all limitations of claim 17. The limitations of claim 17 are nearly identical to those of claim 1, with

the only possibly substantive difference being that claim 17 recites a control device containing a storage medium that (like that of the computer program product of claim 10) stores computer instructions that, when executed, perform the method of claim 1. Accordingly, the analysis for this claim primarily incorporates by reference the discussion above for claims 1 and 10.

1. A control device for network access control, comprising:

163. In my opinion, to the extent the preamble is limiting, Richmond discloses a control device for network access control (*e.g.*, network entry devices 114, 116, 120, 124, 144, and 1501). Richmond teaches that "[a] network entry device serves as an entry point to the network for at least one user," and that it "reside[s] at an edge or boundary of the communications network and provide[s] connectivity between network resources of the communications network and devices located external to the communications network," such as user devices. Ex1003, [0016]. As I have discussed above (*see* Section VII.A), Richmond's network entry device (including its entry port modules) controls the traffic that is sent from user devices to the network and therefore Richmond's network entry device acts as a control device for network access control.

2. a processor;

164. Richmond discloses that its control device (*e.g.*, network access device) comprises a processor. For example, Richmond discloses that network access device

1504 comprises one or more entry port modules 1510, and that each entry port module contains port processing logic 1516. *See, e.g.*, Ex1003, [0014] ("a 'port module' of a switching device includes … port processing logic associated with the physical port."), [0270]. Richmond further discloses that "[p]ort processing logic 1516 may include switching logic, memory, and ***one or more processors*** (not shown) for configuring the port module and for processing packets sent to and received at the port module." *Id.*, [0270].

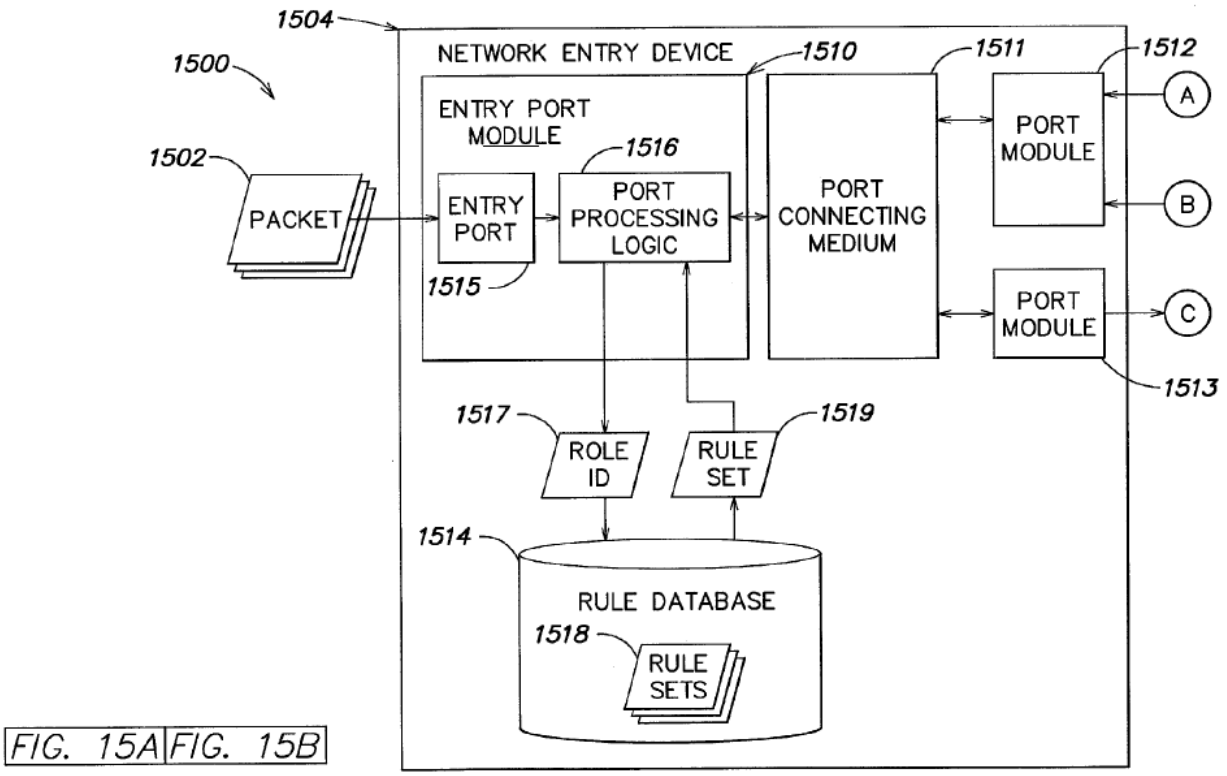

FIG. 15A

3. at least one computer-readable storage medium storing computer instructions translatable by the processor to perform:

165. I incorporate by reference my discussion of the preamble of claim 10. *See* Section IX.J.1.

4. receiving a packet originating from a user device in a first network, wherein the first network is connected to a second network via the control device, and wherein the user device is associated with a user;

166. I incorporate by reference my discussion of the corresponding limitation of claim 1. *See* Section IX.A.2.

5. processing the packet according to a plurality of stages, including a client discrimination stage and a user specific rule stage;

167. I incorporate by reference my discussion of the corresponding limitation of claim 1. *See* Section IX.A.3.

6. at the client discrimination stage, extracting information associated with the user device from a header of the packet and associating the packet with user specific traffic control rules and user specific firewall rules; and

168. I incorporate by reference my discussion of the corresponding limitation of claim 1. *See* Section IX.A.4.

7. at the user specific rule stage, accessing the user specific traffic control rules and user specific firewall rules based on the extracted information associated with the user device and applying the user specific traffic control rules and the user specific firewall rules to the packet as governed by at least

one user specific class of service rule associated with the user on the user device in the first network.

169. I incorporate by reference my discussion of the corresponding limitation of claim 1. *See* Section IX.A.5.

R. Dependent Claim 18

1. The control device of claim 17, wherein the plurality of stages further comprises an interface specific stage and wherein the computer instructions are further translatable by the processor to perform: at the interface specific stage, applying interface specific rules to the packet based on a type of a network interface over which the packet is received.

170. I incorporate by reference my discussion of the corresponding limitations of claim 2. *See* Section IX.B.1.

S. Dependent Claim 19

1. The control device of claim 17, wherein the plurality of stages further comprises a global stage and wherein the computer instructions are further translatable by the processor to perform: at the global stage, applying global firewall rules to the packet.

171. I incorporate by reference my discussion of the corresponding limitations of claim 3. *See* Section IX.C.1.

T. Dependent Claim 20

1. The control device of claim 17, wherein the computer instructions are further translatable by the processor to implement a control program on the control device, wherein the control program comprises a provisioning module, wherein the provisioning module comprises a traffic

conditioning module for user specific allocation of bandwidth.

172. I incorporate by reference my discussion of the corresponding limitations of claim 13. *See* Section IX.M.1.

U. Dependent Claim 22

1. The control device of claim 20, wherein the provisioning module further comprises a firewall module for associating the packet with the user specific firewall rules.

173. Richmond discloses and renders obvious claim 20, as discussed. Richmond further discloses and renders obvious that the provisioning module (*e.g.*, the hardware of the network entry device and portion of the software that implements the corresponding functionality) further comprises a firewall module for associating the packet with the user-specific firewall rules.

174. As discussed further in connection with claim 1 (Section IX.A.4, *supra*), Richmond's network entry device utilizes an authentication database to authenticate a user (Ex1003, [0224]). Richmond's network entry device then associates that user with user-specific rules (including firewall rules) by searching in the rule database 1514 for those matching the role of a user (*id.*, [0267], [0278]). The port module is then configured with the rules, and subsequently, all packets received at that port module from the user are associated with the previously configured rules. *See id.*, [0244], [0246]. The functionality in the port module (*e.g.*, the hardware and portion of the control program software implementing the

functionality) that applies the firewall rules constitutes a "firewall module."

V. Dependent Claim 23

1. The control device of claim 20, wherein the provisioning module further comprises an authentication module for authenticating the user.

175. Richmond discloses and renders obvious claim 20, as discussed. Richmond further discloses and renders obvious that the provisioning module further comprises an authentication module for authenticating the user. Richmond discloses its system has an "authentication module 1506," shown in Fig. 15B, and further discloses that the authentication module "may be configured to perform authentication in accordance with one or more authentication technologies." Ex1003, [0274]. As I have discussed further in connection with claim 8 (*see* Section IX.H), the network entry device sends a user's identity information to the authentication module for authentication. *Id.*, [0273]-[0275].

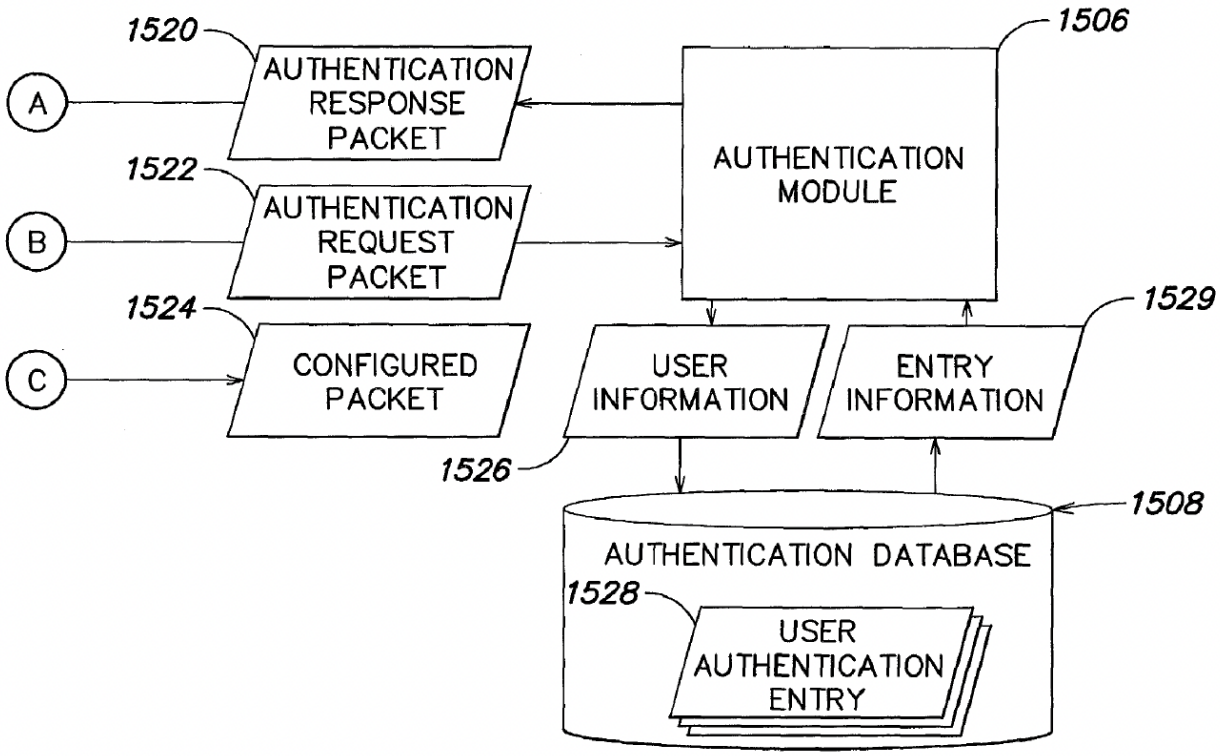

FIG. 15B

W. Dependent Claim 24

1. The control device of claim 17, wherein the computer instructions are further translatable by the processor to establish the user specific traffic control rules and the user specific firewall rules for the user based on a user profile that includes parameters that determine an extent of the user's access to the second network, services available to the user, or a combination thereof.

176. I incorporate by reference my discussion of the corresponding limitations of claim 7. *See* Section IX.G.1.

X. Dependent Claim 25

1. The control device of claim 17, wherein the computer instructions are further translatable by the processor to determine whether the user has authenticated.

177. I incorporate by reference my discussion of the corresponding limitations of claim 8. *See* Section IX.H.1.

Y. Dependent Claim 26

1. The control device of claim 25, wherein the user has not authenticated and wherein the computer instructions are further translatable by the processor to perform: redirecting the user to a login web page that requests user credentials; and using the user credentials received from the user to authenticate the user.

178. I incorporate by reference my discussion of the corresponding limitations of claim 9. *See* Section X.A.1.

Z. Independent Claim 27

179. In my opinion, Richmond Wood renders obvious all limitations of claim 27. The limitations of claim 27 are nearly identical to those of claim 17 combined with its dependent claims 25 and 26. Moreover, claims 17, 25, and 26 are apparatus claims that are substantively similar to method claims 1, 8, and 9. Accordingly, my analysis for this claim incorporates by reference the discussion above for these claims.

1. A control device for network access control, comprising:

a processor;

at least one computer-readable storage medium storing computer instructions translatable by the processor to perform:

180. I incorporate by reference my discussion of the preamble and corresponding limitations of claim 17. *See* Section IX.Q.1-3.

2. receiving a packet originating from a user device in a first network, wherein the first network is connected to a second network via the control device, and wherein the user device is associated with a user;

181. I incorporate by reference my discussion of the corresponding limitation of claim 1. *See* Section IX.A.2.

3. determining whether the user has authenticated;

182. I incorporate by reference my discussion of the corresponding limitation of claim 8. *See* Section IX.H.1.

4. if the user has not authenticated: redirecting the user to a login web page that requests user credentials; and using the user credentials received from the user to authenticate the user;

183. I incorporate by reference my discussion of the corresponding limitation of claim 9. *See* Sections X.A.1.

5. processing the packet according to a plurality of stages, including a client discrimination stage and a user specific rule stage;

at the client discrimination stage, extracting information associated with the user device from a header of the packet

and associating the packet with user specific traffic control rules and user specific firewall rules; and

at the user specific rule stage, accessing the user specific traffic control rules and user specific firewall rules based on the extracted information associated with the user device and applying the user specific traffic control rules and the user specific firewall rules to the packet as governed by at least one user specific class of service rule associated with the user on the user device in the first network.

184. I incorporate by reference my discussion of the corresponding limitations of claim 1. *See* Sections IX.A.3-5.

X. GROUND 3 – CLAIMS 9, 16, 26, AND 27 ARE RENDERED OBVIOUS BY RICHMOND IN LIGHT OF WOOD

A. Dependent Claim 9

1. A method according to claim 8, wherein the user has not authenticated, further comprising: redirecting the user to a login web page that requests user credentials; and using the user credentials received from the user to authenticate the user.

185. Richmond discloses and renders obvious claim 8, as discussed. Richmond in light of Wood renders obvious that when the user has not been successfully authenticated, is redirected to a login web page that requests user credentials. The user credentials received from the user are then used to authenticate the user.

186. Richmond explains in its Abstract that its focus is on "[c]ontrolling a user's usage of network resources, ***after the user has been authenticated***, without using any network resources beyond the user's entry point to the network." Ex1003,

cover. Richmond's primary focus is thus on what occurs after authentication, and thus does not dictate a particular method of authentication. A POSITA would understand, in light of Richmond's disclosure of authentication, that some form of authentication would be needed in conjunction with Richmond's system. Further, authentication is a widely recognized function, and a POSITA would know that many methods existed for authenticating users.

187. A POSITA would thus look to other prior art references to better understand the known ways in which authentication could be performed. One such reference is Wood, which, like Richmond and the '639 Patent, are in the same technical field of controlling access to network resources.

188. In Wood's system, user devices (e.g., "a browser 170 operated by a user") interact with the network via a "Gatekeeper and entry handler component 110," which "provides an entry point for external client applications requesting access." Ex1004, 5:10-15. Wood discloses that when the gatekeeper/entry handling component 110 receives packets from a user device attempting to access a particular website or application, the "gatekeeper/entry handler component 110 queries authorization component 140 to obtain authorization for access" to that application. *Id.*, 6:3-8.

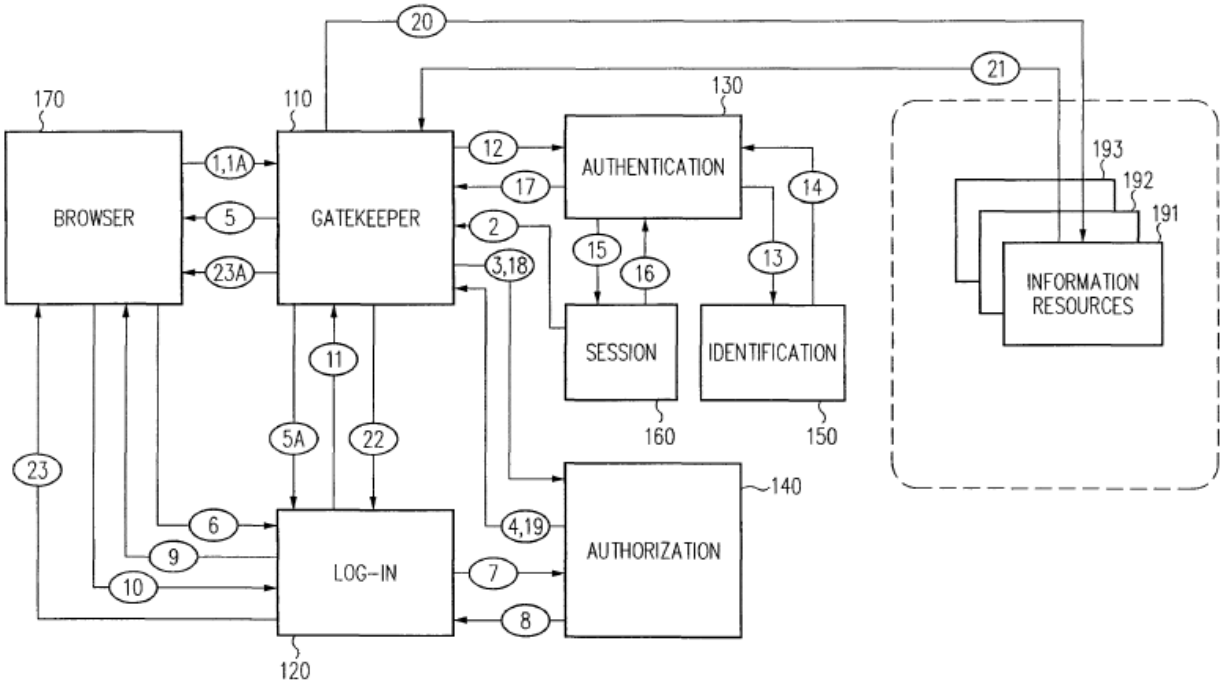

FIG. 1

189. If the user has not authenticated (or has an insufficient authentication level), "authorization component 140 may indicate that the access request is to be redirected to login component 120 so that login credentials may be obtained and authenticated to a particular trust level." *Id.*, 6:7-9; *see also id.*, 12:1-26 (disclosing that the authentication module 130 sends a "REDIRECT response" indicating "to forward the access request to login component 120 so that sufficient login credentials may be obtained and authenticated").

190. Upon receiving a REDIRECT response, the gatekeeper/entry handler "redirects (5) browser 170 to login component," for example, by "using a client-side redirect via HTTP location directive to forward the request to login component 120." *Id.*, 12:23-29. The login component 120 then "supplies browser 170 with a login

page (e.g., HTML) that prompts the user for an application specific user ID and a choice of authentication schemes." *Id.*, 13:37-38:1. Wood discloses that this exemplary login webpage then requests user credentials, explaining that "login component 120 may interactively allow a user to select from amongst the credential types (e.g., using any HTML-based dialog)," including "username/password pairs, onetime passwords, enigma challenge, etc.," which a user may supply through "entry into an HTML form." *Id.,* 14:3-9. Subsequently, the credentials received from the user are "supplied (11) to gatekeeper/entry handler component 110 for authentication," in order to authenticate the user. *Id.*, 14:22-25.

191. A POSITA would understand that the combination of Richmond and Wood would achieve the benefits of Richmond's system for user-specific rules controlling network access, at the same time as achieving the benefits of Wood's user-friendly authentication method. A POSITA would understand that Wood's web-based authentication method is simple for network administrators to set up. Web pages and web servers were readily available at the time, such that a network administrator could use existing tools to design a login page, and could use existing web servers to provide the page. Thus, as a POSITA would readily recognize, a network administrator using a combination of Richmond and Wood would avoid the burden of designing a custom server for accepting user credentials. A POSITA would also understand that end users would find Wood's web-browser-based system

easy to use, because web browsers were ubiquitous on users' computers at the time. Accordingly, users would not need to install new software. Additionally, webpage-based login systems were widely used and thus familiar to the end user. A POSITA would thus recognize that the end user would readily be able to use Wood's authentication system.

192. The use of Wood's authentication method in Richmond's system would be a combination of prior art elements according to known methods. The combination would be straightforward for a POSITA to implement: users authenticate with the system (as explained in Wood) and then their traffic is monitored and controlled (as explained in Richmond). These two functions are separate and discrete: how the authentication is performed does not affect Richmond's operation and the functions of Richmond do not dictate a particular type of authentication. Thus, a POSITA would recognize that they could be easily combined. Any modifications would be straightforward and not involve substantial alteration to Richmond's system. That is, Richmond's network entry device (as modified) would simply send an HTTP redirect (as in done in Wood) to the user upon determining (in the steps discussed for claim 9) that the user has not yet authenticated. Accordingly, a POSITA would have a reasonable expectation of success.

B. Dependent Claim 16

1. The computer program product of claim 15, wherein the user has not authenticated and wherein the computer instructions are further translatable by the processor to perform: redirecting the user to a login web page that requests user credentials; and using the user credentials received from the user to authenticate the user.

193. I incorporate by reference my discussion of the corresponding limitations of claim 9. *See* Section X.A.1.

C. Dependent Claim 26

1. The control device of claim 25, wherein the user has not authenticated and wherein the computer instructions are further translatable by the processor to perform: redirecting the user to a login web page that requests user credentials; and using the user credentials received from the user to authenticate the user.

194. I incorporate by reference my discussion of the corresponding limitations of claim 9. *See* Section X.A.1.

D. Independent Claim 27

195. In my opinion, Richmond in light of Wood renders obvious all limitations of claim 27. The limitations of claim 27 are nearly identical to those of claim 17 combined with its dependent claims 25 and 26. Moreover, claims 17, 25, and 26 are apparatus claims that are substantively similar to method claims 1, 8, and 9. Accordingly, my analysis for this claim incorporates by reference the discussion above for these claims.

1. A control device for network access control, comprising:

a processor;

at least one computer-readable storage medium storing computer instructions translatable by the processor to perform:

196. I incorporate by reference my discussion of the preamble and corresponding limitations of claim 17. *See* Section IX.Q.1-3.

2. receiving a packet originating from a user device in a first network, wherein the first network is connected to a second network via the control device, and wherein the user device is associated with a user;

197. I incorporate by reference my discussion of the corresponding limitation of claim 1. *See* Section IX.A.2.

3. determining whether the user has authenticated;

198. I incorporate by reference my discussion of the corresponding limitation of claim 8. *See* Section IX.H.1.

4. if the user has not authenticated: redirecting the user to a login web page that requests user credentials; and using the user credentials received from the user to authenticate the user;

199. I incorporate by reference my discussion of the corresponding limitation of claim 9. *See* Sections X.A.1.

5. processing the packet according to a plurality of stages, including a client discrimination stage and a user specific rule stage;

at the client discrimination stage, extracting information associated with the user device from a header of the packet

and associating the packet with user specific traffic control rules and user specific firewall rules; and

at the user specific rule stage, accessing the user specific traffic control rules and user specific firewall rules based on the extracted information associated with the user device and applying the user specific traffic control rules and the user specific firewall rules to the packet as governed by at least one user specific class of service rule associated with the user on the user device in the first network.

200. I incorporate by reference my discussion of the corresponding limitations of claim 1. *See* Sections IX.A.3-5.

XI. GROUND 4 – CLAIM 21 IS RENDERED OBVIOUS BY RICHMOND IN LIGHT OF TERASLINNA

A. Dependent Claim 21

1. The control device of claim 20, wherein the traffic conditioning module further comprises an interface master queue for controlling a flow of network traffic over a particular network interface.

201. As I have discussed above, Richmond discloses and renders obvious the control device of claim 20. It is further my opinion that Richmond in light of Teraslinna renders obvious that the traffic conditioning module further comprises an interface master queue (*e.g.*, Teraslinna's queue 144) for controlling a flow of network traffic over a particular network interface.

202. Richmond discloses that its user-specific rules include rules that enforce bandwidth rate limitations on users (which a POSITA would also refer to as traffic shaping or rate limiting). *See* Ex1003, [0010], [0158]-[0159]. Richmond

further discloses that the network entry device will apply rate limiting (*e.g.*, act 1430) according to these user-specific rules. *Id.*, [0158]-[0159], [0249]-[0251].

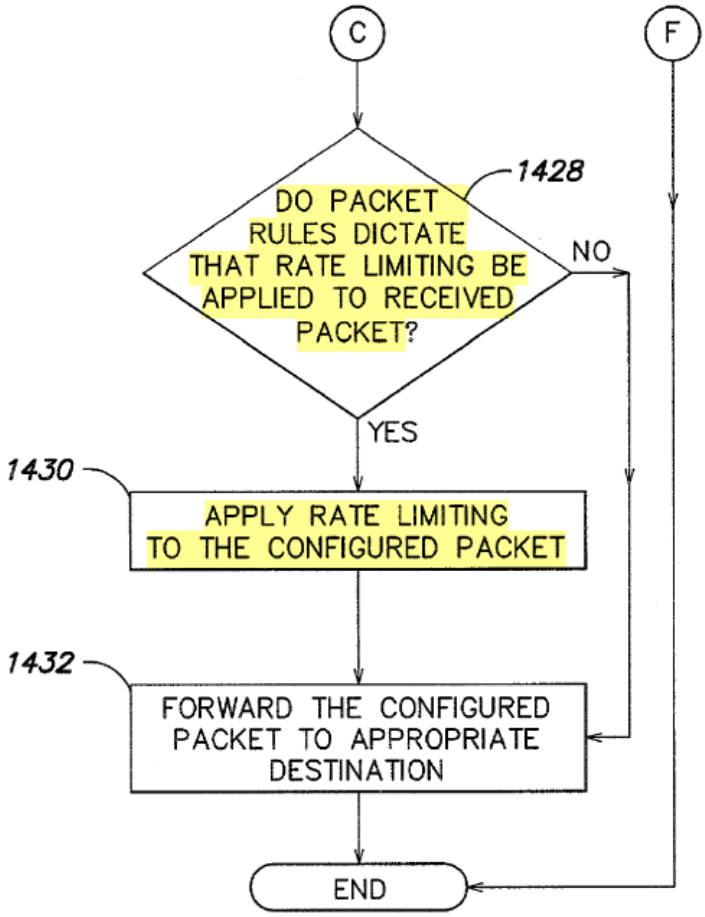

FIG. 14B

203. A POSITA reading Richmond's disclosure of traffic shaping would naturally look to find a reference that discloses an implementation of traffic shaping, of which there were a variety of known ways in the prior art, including methods using queues. *See, e.g.*, Ex1009, 379-384 (discussing several methods). One example of a method to implement is disclosed in Teraslinna. A POSITA would

look to Teraslinna because it is in the same field as Richmond and the '639 Patent, namely, systems for controlling network access, and in particular, systems for controlling bandwidth allocations and performing traffic conditioning in packet-switched networks.

204. Teraslinna discloses that traffic from a subscriber equipment 135 (which is analogous to Richmond's user devices) passes through a traffic flow processor 142 (which is analogous to Richmond's network entry device), in order to reach a fast packet switching network 134 (which includes the Internet, as disclosed in Richmond). Teraslinna further discloses that the traffic flow processor 142 is coupled to queue 144, which "queues all incoming packets from the customer premises equipment 130 for traffic shaping by the traffic flow processor." Ex1005, 16:5-8.

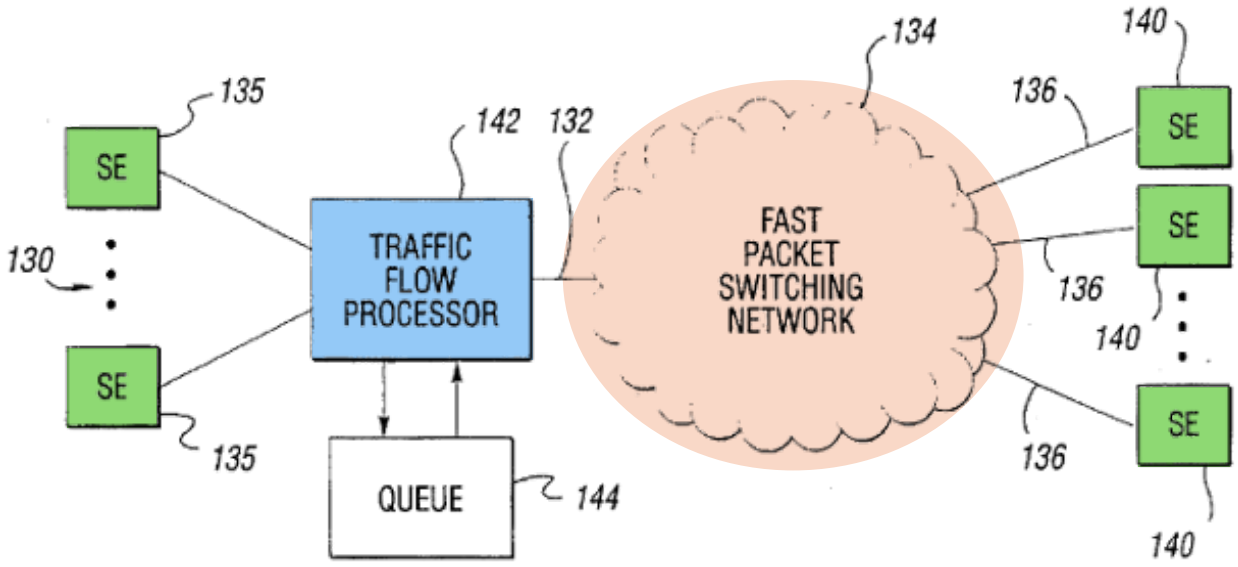

Ex1005, Fig. 11 (annotated).

205. A POSITA would understand queue 144 to be an exemplary interface

master queue, because it receives all traffic that is received at the interfaces of traffic flow processor 142 and controls all traffic on the interface to network endpoint 132. Teraslinna uses the queue 144 to control the flow of network traffic across the interface to network endpoint 132, according to the bandwidth limiting algorithm also disclosed in Teraslinna. *Id.*, 16:9-17:18, Fig. 12.

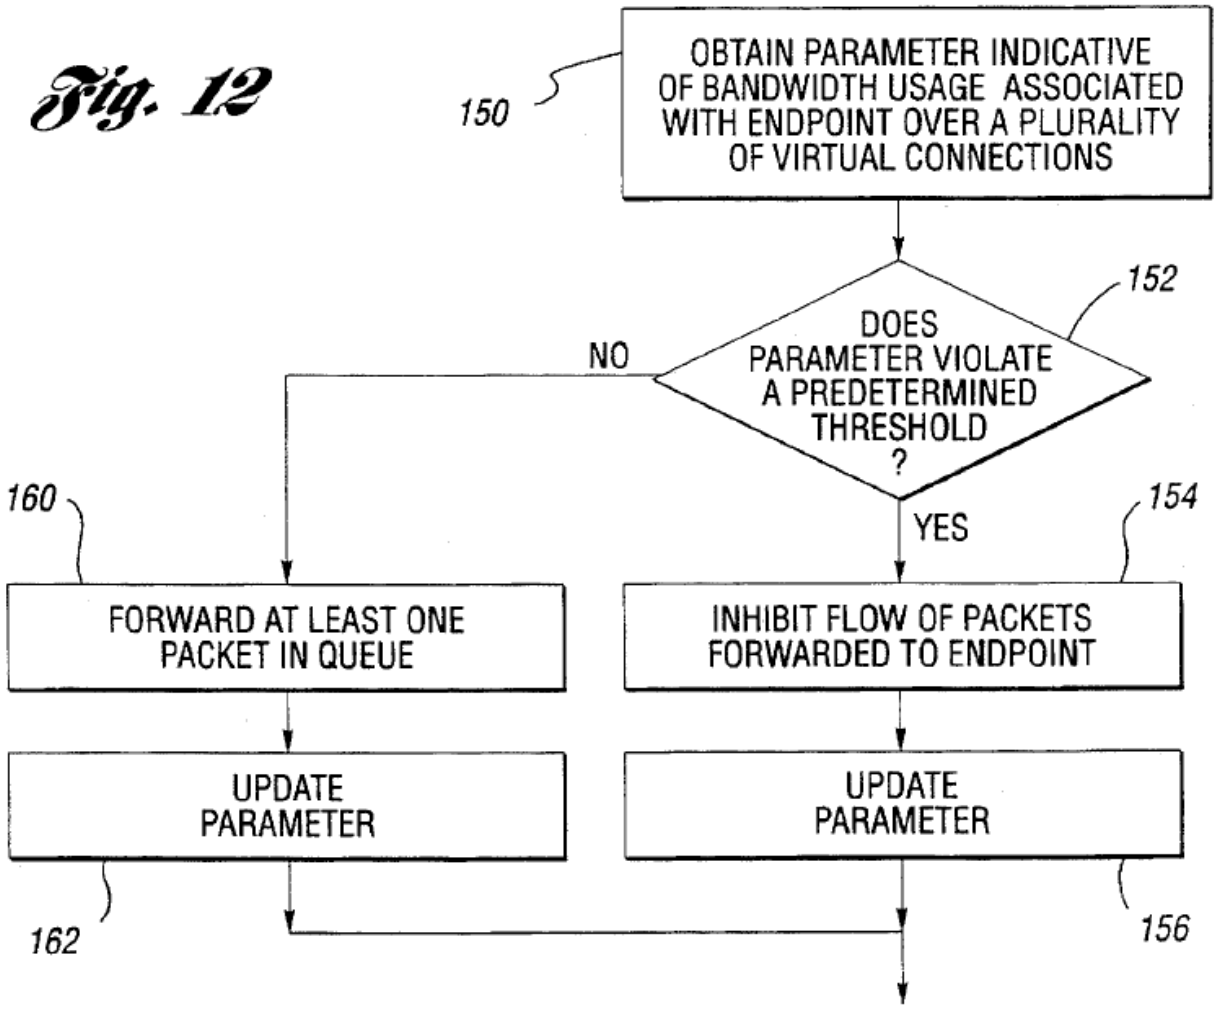

Id., Fig. 12.

206. A POSITA would find that the combination of Richmond with Teraslinna's traffic shaping method (including queue 144 that Tereslinna's method

is based on) provides for cost-efficient bandwidth allocation, in a way that "reduce[s] bandwidth requirements" and "reduce[s] the cost of implementing a virtual private network." Ex1005, 2:43-48.

207. A POSITA would understand that the combination of Teraslinna's queue 144 and traffic shaping method would be a combination of prior art elements according to known methods to yield predictable results. Richmond's and Teraslinna's systems are highly analogous. For example, Teraslinna's traffic flow processor 142 plays the same role of traffic shaping in its system as Richmond's network entry device.

208. Thus, a POSITA would be motivated to add the traffic-shaping functionality of Teraslinna to Richmond's network entry device. To make the combination, a POSITA would simply add one of Teraslinna's queues to each of Richmond's port modules and then add Teraslinna's algorithm to the port module's port processing logic, in order to process the stream of packets. Because of the straightforward nature of this combination, a POSITA would have a reasonable expectation of success.

XII. CONCLUSION

209. In signing this declaration, I recognize that the declaration will be filed as evidence in a contested case before the Patent Trial and Appeal Board of the United States Patent and Trademark Office. I also recognize that I may be subject to

cross-examination in the case and that cross-examination will take place within the United States. If cross-examination is required of me, I will appear for cross-examination within the United States during the time allotted for cross-examination.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements

Kevin C. Almeroth, Ph.D.