

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

RED HAT, INC.
Petitioner

v.

COMPETITIVE ACCESS SYSTEMS, INC.
Patent Owner

Case No. IPR2025-01371
U.S. Patent No. 10,868,908

**DECLARATION OF DR. STAN MCCLELLAN IN SUPPORT OF PATENT
OWNER PRELIMINARY RESPONSE**

Competitive Access Systems EX2015
Red Hat, Inc. v. Competitive Access
Systems, Inc.
IPR2025-01371

TABLE OF CONTENTS

	<u>Page</u>
I. Background.....	3
II. Qualifications.....	4
III. Compensation.....	6
IV. Materials Considered	6
V. Legal Principles	6
VI. Overview of the '908 Patent and the Invention Claimed in the '908 Patent	9
VII. The Prior Art	26
A. Kotzin	26
B. Phatak	28
VIII. Claim Construction.....	29
A. "Multilink Connection"	29
IX. Claim Construction and Opinions on Ground 1 (Kotzin Alone or in Combination with Held, Peirce, and Decasper).....	34
A. "request for a multilink communication"/"multilink request"	34
B. "multilink session identification"	39
X. Claim Construction and Opinions on Ground 2 (Phatak Alone or in Combination with Held).....	42
A. "request for a multilink communication"/"multilink request"	42
B. "multilink session identification"	46

I, Dr. Stan McClellan, hereby declare as follows:

I. Background

1. I have been retained by Patent Owner Competitive Access Systems, Inc. ("CAS") to provide my opinions in the above-captioned *inter partes* review ("IPR") proceeding involving U.S. Patent No. 10,868,908 (the "'908 Patent," EX1007), which is titled "Devices and Methods for Multipath Communications." I have been asked to evaluate whether certain references disclose or suggest the features recited in Claims 4, 7-12, and 17 of the '908 Patent (the "Challenged Claims"). I understand this declaration will be entered in the above-captioned IPR proceeding, as an exhibit with Patent Owner's Preliminary Response to the Petition for Inter Partes Review (the "Preliminary Response") submitted by CAS.

2. The '908 Patent is a continuation of U.S. Patent Nos. 9,350,649 (the "'649 Patent"), 8,861,349 (the "'349 Patent"), 8,228, 801 (the "'801 Patent"), and 7,606,156 (the "'156 Patent"). In addition to being retained in the present matter, I was also retained to provide my opinions as to the validity of the claims of the '156 Patent '801 Patent, '349 Patent, and '649 Patent, as well as the validity of the claims of U.S. Patent Nos. 11,418,641 (the "'641 Patent") and 11,582,343 (the "'343 Patent"), which are each also continuations of the '156 Patent. I have and/or will submit declarations as part of the '156 Patent, '801 Patent, '349 Patent, '649 Patent, '641 Patent, and '343 Patent IPR proceedings.

II. Qualifications

3. I currently serve as the Director of the JSC Engineering, Technology and Science Program (JETS) at Texas State University, a partnership with Jacobs Technology Group and NASA's Johnson Space Center (JSC). Previously, I held the following positions at Texas State:

- a. Professor in the Ingram School of Engineering from 2013-2023;
- b. Co-Director of the Connected Infrastructure initiative for Education, Demonstration, and Research (CIEDAR) from 2019 to 2023;
- c. Director of the Ingram School of Engineering from 2013 to 2018;
- d. Associate Director of the Ingram School of Engineering from 2011 to 2013; and
- e. Associate Professor in the Ingram School of Engineering from 2008 to 2013.

4. My research activities at Texas State have involved the development and management of FCC-licensed private 5G and LTE wireless networks, and the management of a regional LoRaWAN wireless network for low-rate, real-time data collection in multiple projects at CIEDAR. During my time as a Professor, I developed performance enhancements for Internet transport protocols and packet

handling based on signal processing and machine learning algorithms, and designed a multi-building "5G-capable" core network to support multiple faculty re-search projects as well as academic priorities. During my time as an Associate Professor, I investigated various Smart Grid systems and advanced technologies for broadband networking (e.g. 40/100Gb Ethernet), developed novel PHY layer for low-frequency communication via power lines, and developed performance enhancements for IP transport protocols.

5. I have prior academic experience as an Associate Professor of Electrical and Computer Engineering at the University of Alabama at Birmingham from 1995 to 2002.

6. I also have technical experience from my time in the private sector. I previously held executive or senior engineering positions at companies including Power Tagging Technologies, ZNYX Networks, SBE, Hewlett Packard, General Dynamics and LTV Aerospace and Defense Company. I have also served as a technical consultant for companies such as Cisco, American Express, 3COM, BellSouth, N.E.T., Alcatel, AT&T, Research in Motion (RIM), F5 Networks, LSU Medical Center, and the U.S. National Science Foundation.

7. I have published dozens of books and peer-reviewed research papers in this field, and have given numerous related technical talks world-wide. I am a

named inventor on fourteen patents and have served as a technical expert witness in numerous patent litigation cases.

8. I received my Ph.D. from Texas A&M University (1995) in Electrical Engineering with a specialization in communication and data processing. I received a M.S. in Electrical Engineering (1991) and a B.S. in Electrical Engineering (1986), both from Texas A&M as well.

9. My qualifications and experience are detailed further in my Curriculum Vitae, which is EX2016.

III. Compensation

10. I am being compensated at the rate of \$350 per hour for the work I conduct related to this proceeding. My compensation does not depend on the content of my opinions or the outcome of the proceeding.

IV. Materials Considered

11. In providing the testimony set out in this declaration, I have reviewed the specification and claims of the '908 Patent, the file history of the '908 Patent, Red Hat's Petition for Inter Partes Review of the '908 Patent (the "Petition"), and the prior art, declarations, and other documents attached to the Petition.

V. Legal Principles

12. I am not an attorney and offer no legal opinions. In forming my opinions, I have been informed by counsel about certain aspects of the law that are relevant to my analysis. My understanding of these principles is summarized below.

13. I understand that in patent law a patent's claims define the scope of protection for the claimed invention.

14. I understand that the first step in an invalidity analysis involves construing the claims, as necessary, to determine their scope. I have also been informed by counsel that claims are generally given their ordinary and custom meaning as understood by a person of ordinary skill in the art ("POSITA") at the time of the invention, in light of the patent's claims, specification and file history. I have additionally been informed by counsel that claim construction is a matter of law and that the final claim constructions for this proceeding will be determined by the Patent Trial and Appeal Board ("PTAB").

15. I understand that the validity of a patent's claims is determined on a claim-by-claim basis.

16. I understand that in this proceeding, Petitioner has the burden of proving that the challenged claims would have been unpatentable over the prior art by a preponderance of the evidence.

17. Ground 1 of the Petition argues that the Challenged Claims of the '908 Patent should be invalidated as disclosed or rendered obvious by the "Kotzin" prior art reference alone or in combination with the "Held," "Peirce," and "Decasper" prior art references. As pertinent to my analysis, these references are described further below.

18. Ground 2 of the Petition argues that the Challenged Claims of the '908 Patent should be invalidated as disclosed or rendered obvious by the "Phatak" prior art reference alone or in combination with the "Held" prior art reference. As pertinent to my analysis, these references are described further below.

19. I understand that a claim is unpatentable as anticipated if a single prior art reference teaches each limitation of the claim.

20. I understand that a claim is unpatentable as obvious if the differences between the claimed subject matter and the prior art are such that the subject matter, as a whole, would have been obvious to a POSITA at the time of the invention.

21. I understand that the question of obviousness is resolved based on underlying factual determinations, including: (1) the scope and content of the prior art; (2) any differences between the claimed subject matter and the prior art; (3) the level of ordinary skill in the art; and (4) objective evidence of non-obviousness.

22. I further understand that for each of its obviousness grounds, Petitioner must demonstrate a POSITA would have: (1) a motivation to combine the relevant references and (2) a reasonable expectation of success of developing the claimed invention.

23. In forming my opinions, I have applied the understanding of one of ordinary skill in the art working at the time of the invention claimed in the '980 Patent, i.e., prior to January 2001. In my opinion, a person of ordinary skill in the

art at this time would have an acquired body of knowledge gained through formal education, or practical experience providing an equivalent to, a Bachelor of Science in Electrical Engineering or Computer Science, or a related/equivalent field, and at least 3 years of work experience in the field of computer networking. I understand that the date of conception of the '343 Patent may be contested. My identification as to the level of skill of the person of ordinary skill in the art would not vary between the date of the conception or the October 15, 2002 filing date of Provisional Application No. 60/418,521.


VI. Overview of the '908 Patent and the Invention Claimed in the '908 Patent

24. The '908 Patent pertains to devices and methods to improve the effective bandwidth between a remote sending device (e.g., a server) and a receiving device by leveraging the bandwidth of at least two separate communication paths to receive the data sent from a remote device.

25. Although specific claims or claim limitations may be directed to only certain aspects of the configuration or to separate, but related embodiments, Figure 7 of the specification of the '908 Patent shows an example network configuration.


26. In this configuration, a request is sent from a computer network device (e.g., phone, laptop, server, etc.) over a first path (e.g., Path 1) to another computer network device, such as, for example, a remote server computer, for participation in a multilink session. To this end, as depicted and annotated below, the

specification states: "Referring to FIG. 7, an RCG-85 that decides to initiate a multilink PPP session with a remote computer i.e. internet URL, requests a multilink PPP session with that URL resource via Link Control Protocol (LCP) session management[.]" EX1007 at 12:2–6.




27. The two devices negotiate the pertinent multilink parameters. For example here, as depicted and annotated below, the specification states: "The initiating RCG-85 and the remote server computer (not shown) will negotiate the

pertinent multilink parameters such as compression, packet size, sync/async, link quality monitoring, bit stuffing etc." EX1007 at 12:6–9.




28. The multilink session setup is completed including that the remote server computer sends a multilink session ID to the computer network device (e.g., phone, laptop, etc.), which is used to identify this particular multilink session and is also used for adding or removing links from the multilink session. As is also depicted and annotated below, the specification states: "As part of a successful

multilink PPP session setup, the initiating RCG-85 receives a multilink PPP session ID, or "magic number", which is used to identify this particular multilink PPP session and is also used for adding or removing links from the multilink bundle-87." EX1007 at 12:9–14.




29. The computer network device (e.g., phone, laptop, server etc.) then sends a multilink request over a second path (e.g., Path 2) to one or more remote communications devices such as, for example, a router or base station to participate

in the multilink session. As is relevant here, and as depicted and annotated below, the specification states: "The initiating RCG-85 then looks up the addresses of up to 31 (in this embodiment) other RCGs in its router tables to use for the multilink PPP transfer. (If a Neighborhood Access Point is available the RCG could request service through it instead of setting up a Multilink PPP connection—see FIG. 8.) The initiating RCG-85 determines the number of remote RCGs-86 that it wants to use to get the desired bandwidth for the transfer and sends multilink PPP service requests via the 802.11b/g wireless connections-88 to the selected remote RCGs-86." EX1007 at 12:16–26.




30. The remote communications devices respond back to the computer network device (e.g., phone, laptop, etc.) with acceptance or denial. For example, as depicted and annotated below, the specification states: "The selected remote RCGs-86 respond back to the initiating RCG with either an acceptance or denial of multilink PPP service request. This decision is made based upon the remote RCG's currently available POTS bandwidth and the current local demands for that bandwidth. The remote RCG-86 makes a decision to participate in a multilink PPP service for an initiating RCG If the remote RCG-86 accepts the multilink PPP

service request from the initiating RCG-85 it also sends its currently available bandwidth on its POTS line-91 that the initiating RCG-85 will have available to it for the multilink PPP transfers." EX1007 at 12:26–32; 12:46–51.




31. The computer network device (e.g., phone, laptop, server, etc.) sends a multilink initialization packet containing the address of the destination as well as the multilink session ID (magic number) and any other pertinent information required for the remote communications device to join in the existing multilink session. Here, as depicted and annotated below, the specification states: "The


initiating RCG then selects which of the responding remote RCGs-86 it will use for the multilink PPP bundle-87 and sends them a multilink initialization packet containing the URL of the destination as well as the multilink PPP session ID (magic number) and any other pertinent information required for the remote RCG to join in the existing multilink PPP bundle-87. Each remote RCG POTS interface-91 participating in a multilink PPP bundle, represents a logical "link" in the initiating RCG's multilink PPP bundle-87." EX1007 at 12:51-63.




32. Each of the participating remote communications devices sends a multilink packet to the remote server computer with the multilink session ID and requests to be added to the multilink session. For example, as depicted and annotated below, the specification states: "The remote RCGs-86 then send a multilink PPP packet (over their respective POTS lines-91) to the destination URL with the multilink PPP session ID and request to be added to that multilink PPP bundle-87." EX1007 at 12:60–63.



33. The remote server computer accepts the request and begins sending data packets to each of the remote communications devices over their respective paths or links. As an example, as depicted and annotated below, the specification states: "The destination URL accepts the request of each of the remote RCGs-86 and begins sending data packets to each over their respective POTS lines-91." EX1007 at 12:63–66.



34. The remote communications device relays the data packets to the computer network device (e.g., phone, laptop, etc.). In this instance, as depicted and annotated below, the specification states: "When the remote RCGs-86 receive data packets sent over their respective POTS lines-91, that are destined for the multilink PPP bundle-87, they simply relay the data packets over their 802.11b/g interface-88 to the initiating RCG-85." EX1007 at 12:66–13:3.



35. The use of the additional links as claimed in the '908 Patent thus results in increased bandwidth, data throughput, quality of service, and/or connection reliability. For example, as explained in the specification: "An RCG that is transferring a very large file can request the services of 31 (in this embodiment) other RCGs to assist in the file transfer using their respective POTS lines. A single RCG can therefore gain the combined bandwidth of up to 32 POTS lines transferring data at as much as 56 Kbps each to transfer information. This is roughly equivalent to a transfer rate of 1.79 Mbps in the inbound direction (towards the RCG) and approximately 1.075 Mbps in the outbound direction (towards the PSTN)." EX1007 at 11:47–56.

36. The claims of the '908 Patent are limited to particular configuration and/or operation of computer network devices involved in an exchange of data designed to enable technological improvements such as increasing bandwidth, data throughput, quality of service, and/or connection reliability.

37. Specifically, claim 1 of the '908 Patent (from which Claim 7 depends) recites the following:

A client device for multilink data packet communication with a server, the client device comprising:

a first network interface having a first data transfer rate associated with a first path from the server to the first network interface of the client device;

a second network interface having a second data transfer rate associated with a second path from the server to the second network interface of the client device; and

at least one processor configured to:


transmit a first message through the first network interface to the server, wherein the first message includes a multilink request a destination URL;

transmit a second message through the second network interface to the server, wherein the second message includes a multilink session identification the destination URL; and

receive data packets from the server through the first and second network interfaces of the client device during the multilink session wherein the first data transfer rate and the second data transfer rate provide an increased data transfer rate of the first path and the second path from the server to the client device.

EX1007 at 16:60–17:14.

38. Using Figure 7 as an exemplary, non-exclusive example, this claimed configuration can be depicted as shown below:



39. Claim 8 of the '908 Patent recites the following:

A server for multilink communication with a device, the server comprising:

at least one server network interface; and

one or more processors coupled to the at least one server network interface and configured to:

receive a first multilink packet over the at least one server network interface, the first multilink packet comprised of a request for a multilink communication and first information associated with first network communication interface of the device that is capable of being used by the device to receive data packets;


sending a data packet including the multilink session identification to the first network communication interface of the device;

receive a second multilink packet over the at least one server network interface, the second multilink packet comprised of the multilink session identification and second information associated with a second network communication interface of the device that is capable of being used by the device to receive data packets; and

send different groups of data packets over the at least one server network interface to the first network communication interface of the device and to the second network communication interface of the device, wherein each of the different groups of data packets includes the multilink session identification.

EX1007 at 17:44–18:4.

40. Using Figure 7 as an exemplary, non-exclusive example, this claimed configuration can be depicted as shown below:



41. Claim 17 of the '908 Patent recites the following:

A client device for multilink communication, the client device comprising:

- a first network interface having a first data transfer rate associated therewith;

- a second network interface having a second data transfer rate associated therewith; and

- at least one processor and configured to:


 - transmit a first message through the first network interface to at least one server, wherein the first message includes a multilink request and a destination URL;

 - transmit a second message through the second network interface to the at least one server, wherein the second message includes a multilink session identification that is unique to the multilink session and the destination URL; and

 - receive concurrently the different data packets from the server through the first network interface and the second network interface of the client device during the multilink session, wherein the first data transfer rate and the second data transfer rate provide an increase in the data transfer rate of the client device.

EX1007 at 19:16–20:18.

42. Using Figure 7 as an exemplary, non-exclusive example, this claimed configuration can be depicted as shown below:



43. The inventions recited in the claims of the '908 Patent and its family members also improve upon the prior art. The typical and representative prior art configurations and computer network device operations involved in exchanging data included the following: a Plain Old Telephone Service ("POTS") lines, cable lines, or DSL lines to connect two computer network devices.

44. The problems with the prior art configurations and methods included low bandwidth, data throughput, quality of service, and reliability and substantial costs.

45. For example, as highlighted in the specification, DSL and cable lines require tremendous infrastructure costs, suffer from high costs associated with

purchase, installation, and maintenance, required expensive modems at the customer's end, were difficult to deploy, were too complicated for the average user, and were not available to many customers. EX1007 at 1:65–67 ("DSL or Cable services [] require tremendous infrastructure costs on behalf of the CLEC."); *id.* at 2:9–14 ("The high-speed services require expensive infrastructure enhancements to the local telephone companies' facilities and cable plants, as well as expensive modems at the customer's end ...").


46. As another example, POTS lines had very limited bandwidth. *See also* EX1007 at 11:5–15 ("... standard POTS lines can only guarantee a maximum of 56 Kbps transfer rate under the best conditions ... DSL, however, suffers from the very high costs associated with the purchase, installation, maintenance, floor lease costs, etc. ...").

47. In addition, prior art configurations and methods also suffered from poor quality of service. EX1007 at 2:16–18 ("Internet based VoIP services suffer from Internet delays caused by a lack of true Quality of Service (QOS) capabilities across the Internet.").

VII. The Prior Art

A. Kotzin

48. U.S. Patent Application Publication No. 2003/0026221 ("Kotzin") (EX1018) discloses "[a] method and apparatus splits information into subsets of information to be communicated for a wireless apparatus." EX1018 at Abstract. EX1018 at FIG. 4.



49. Kotzin also discloses in its Figure 5 the following sequence by which participation of wireless units are requested and the data is sent from the network communication system through the multiple wireless units to the recipient wireless unit.




FIG. 5

EX1018 at FIG. 5.

B. Phatak

50. IEEE INFOCOM 2002 ("Phatak") (EX1045), "A Novel Mechanism for Data Streaming Across Multiple IP Links for Improving Throughput and Reliability in Mobile Environments," proposes a "mechanism to aggregate the bandwidth of multiple IP links by splitting a data flow across multiple network interfaces at the IP level" as depicted in its Figure 1. EX1045 at 773.



Fig. 1

COMMUNICATING HOSTS WITH MULTIPLE NETWORK CONNECTIONS
(MULTI-HOMED).

EX1045 at FIG. 1.

51. Phatak describes HOST A, depicted in Figure 1 above, as, e.g., "a laptop at an airport with access to the Internet via an 802.11 LAN (say interface A_1), a Bluetooth scatternet (A_2), and a cellular phone modem (A_3)." EX1045 at 773.

52. Phatak discloses: "Our solution to these problems is to employ *tunneling*. At the sender side, A , the transport layer assembles all packets as if they were going through A_1 and addressed to B_1 . The packets going out on interface A_2 then get encapsulated in a new IP packet with an extra header having destination address B_1 and source address A_2 ." (emphasis in original). This suggests the use where the packets' source and destination addresses are encapsulated by an outer header.

VIII. Claim Construction

A. "Multilink Connection"

53. It is my opinion that in the context of the '908 Patent that claim term "multilink connection" should be construed to mean "transfer of data during a data

transfer session using at least two network interfaces to form links between two devices, where each link is associated with the same session identification for assisting in the data transfer, and at least one of the links connects the two devices through a remote communications device having available bandwidth to assist with the data transfer."

54. The claim language supports that "multilink connection" comprises the transfer of data during a data transfer session" portion of the proposed construction. For example, in Claim 8 once the multilink communication is established it is described as "send[ing] different groups of data packets [i.e., a transfer of data] over the at least one server network interface to the first network communication interface of the device and to the second network communication interface of the device, wherein each of the different groups of data packets includes the multilink session identification [i.e., the transfer of data was within a single data transfer session]." EX1007 at 17:66-18:4.

55. The specification also supports the "transfer of data during a data transfer session" portion of the proposed construction. EX1007 at 11:57-67 ("The multilink PPP connection-87 is initiated by an RCG when it determines that the file (or streaming service) that it is to download (receive) is of sufficient size or requires a minimum amount of bandwidth to justify the use of a multilink PPP connection. This determination can be made on many different criteria such as the file size, type

of data being transmitted, the type of compression if any, a Quality-of-Service requirement or any other pertinent information. (Since a certain amount of overhead is required to set up a Multilink PPP connection, it doesn't make sense to use it for small or low data rate transfers.); 12:2-5 ("Referring to FIG. 7, an RCG-85 that decides to initiate a multilink PPP session with a remote computer i.e. internet URL, requests a multilink PPP session with that URL resource").

56. The claim language supports that "multilink connection" comprises "using at least two network interfaces to form links between two devices" portion of the proposed construction.

57. For example, in Claim 8 once the multilink communication is established it is described as "send[ing] different groups of data packets over the at least one server network interface to the first network communication interface of the device and to the second network communication interface of the device, wherein each of the different groups of data packets includes the multilink session identification." EX1007 at 17:66-18:4.

58. The specification also supports this construction. EX1007 at 3:44-50 ("Broadband networking for home computers and appliances as well as wide area networking using a combination of a plurality POTS lines connected to a plurality of RCGS and/or a direct wireless connection to a wireless access point. (These interfaces allow smart devices to communicate to other devices within the home, or

over the Internet through the RCG.)"); 2:65-3:1 ("The 802.11b/g interface is for wireless communications between the RCG and PCs, other RCGs as well as any other types of compatible wireless devices."); 11:31-36 ("An 802.11b/g interface-87 is shown as the preferred embodiment, since its 11/54 Mbps interface is considered sufficient 10 bandwidth at this time also. its transmission distances are sufficient for implementation of this device, as opposed to other shorter ranged wireless transmission schemes"); 12:66-13:3 ("When the remote RCGs-86 receive data packets sent over their respective POTS lines 91, that are destined for the multilink PPP bundle-87, they 45 simply relay the data packets over their 802.11b/g interface88 to the initiating RCG—85."); 11:25-28 ("Referring to FIG. 7, the RCG implements broadband over POTS by employing an 802.11 b/g wireless interface to configure a multilink PPP bundle87 (per RFC 1990) comprised of the POTS lines connected to up to 32 RCGs."); 12:14-19 ("Once the multilink bundle is setup, the initiating RCG-85 begins data transfer across its POTS interface 90 with the far end computer system. The initiating RCG—85 then looks up the addresses of up to 31 (in this embodiment) other RCGs in its router tables to use for the multilink PPP transfer."); 12:66-13:3 ("When the remote RCGs-86 receive data packets sent over their respective POTS lines 91, that are destined for the multilink PPP bundle-87, they 45 simply relay the data packets over their 802.11b/g interface88 to the initiating RCG—85."); 11:47-56 ("An RCG that is transferring a very large file can

request the services of 31 (in 25 this embodiment) other RCGs to assist in the file transfer using their respective POTS lines. A single RCG can therefore gain the combined bandwidth of up to 32 POTS lines transferring data at as much as 56 Kbps each to transfer information. This is roughly equivalent to a transfer rate of 30 1.79 Mbps in the inbound direction (towards the RCG) and approximately 1.075 Mbps in the outbound direction (towards the PSTN)."); 12:57-60 ("Each remote RCG POTS interface 91 participating in a multilink PPP bundle, represents a logical "link" in the initiating RCG's multilink PPP bundle-87.").

59. The specification supports that "multilink communication" comprises "where each link is associated with the same session identification for assisting in the data transfer and at least one of the links connects the two devices through a remote communications device having available bandwidth to assist with the data transfer" portion of the proposed claim construction. '908 Patent, Fig. 7; 11:47-56 ("An RCG that is transferring a very large file can request the services of 31 (in 25 this embodiment) other RCGs to assist in the file transfer using their respective POTS lines. A single RCG can therefore gain the combined bandwidth of up to 32 POTS lines transferring data at as much as 56 Kbps each to transfer information. This is roughly equivalent to a transfer rate of 30 1.79 Mbps in the inbound direction (towards the RCG) and approximately 1.075 Mbps in the outbound direction (towards the PSTN)."); 12:57-60 ("Each remote RCG POTS interface 91

participating in a multilink PPP bundle, represents a logical "link" in the initiating RCG's multilink PPP bundle-87."); 12:2-4 ("Referring to FIG. 7, an RCG-85 that decides to initiate a multilink PPP session with a remote computer i.e. internet URL,"); 12:9-14 ("As part of a successful multilink PPP session setup, the initiating RCG-85 receives a multilink PPP session ID, or "magic number", which is used to identify this particular multilink PPP session and is also used for adding or removing links from the multilink bundle 87"); 12:51-57 ("The initiating RCG then selects which of the responding remote RCGs-86 it will use for the multilink PPP bundle 87 and sends them a multilink initialization packet containing the URL of the destination as well as the multilink PPP session ID (magic number) and any other pertinent information required for the remote RCG to join in the existing multilink PPP bundle-87.").

IX. Claim Construction and Opinions on Ground 1 (Kotzin Alone or in Combination with Held, Peirce, and Decasper)

A. "request for a multilink communication"/"multilink request"

60. Both Claims 8 and 17 require "a first multilink packet"/"a first message" to be sent from the device to the server that includes a "multilink request"/"request for multilink communication":

[8.c] receive a first multilink packet over the at least one server network interface, the first multilink packet comprised of a request for a multilink communication and first information associated with first network communication interface of the device that is capable of being used by the device to receive data packets;

[17.d] transmit a first message through the first network interface to at least one server, wherein the first message includes a multilink request and a destination URL;

61. It is my understanding that the claimed "request for a multilink communication" is a device-originated message to a server that initiates a multilink session. That message conveys information needed to establish a single, virtualized data transfer session spanning at least two network interfaces.

62. As explained above, the person of ordinary skill in the art would understand "multilink communication" (as used in both Claims 8 and 17) to mean "transfer of data during a data transfer session using at least two network interfaces to form links between two devices, where each link is associated with the same session identification for assisting in the data transfer, and at least one of the links connects the two devices through a remote communications device having available bandwidth to assist with the data transfer."

63. At least one of the links connects the two devices through a remote communications device having available bandwidth to assist with the data transfer. This construct requires a coordinated, session-bound relationship across multiple paths, bundled into a single logical end-to-end connection.

64. I understand Red Hat argues that Kotzin's bandwidth distribution message 508 (broken into 508a–508d) qualifies as a "request for a multilink

communication" because it requests that the Internet Web server send data for the same session to multiple units.

65. Based on my review, message 508 is a single message that identifies participating wireless units to the WAN system, not multiple messages 508a–508d. It is a single point-to-point message from Unit 1 to the WAN identifying "volunteers."

66. In my opinion, message 508 does not request that the server engage in a multilink session. Nor does it solicit the server's participation in an end-to-end virtualized connection under a common session identification. It neither identifies nor requests the use of a multilink session identifier. In fact, subsequent message pairs, including "control message" 512A and "data message" 512B, are independent, end-to-end communications between a single "participant" such as Unit 2, and the WAN system.

67. I understand Red Hat argues that "Messages 512 are also 'multilink requests messages,'" because they request that the Internet Web server send a specific amount of data to each participating wireless unit, and that if Unit 1 were itself receiving data directly, a message 512 from Unit 1 would meet the limitation.

68. In my opinion, the messages labeled 512 in Kotzin describe point-to-point communications between the WAN and individual wireless units. In those communications, the WAN delivers specific data portions to those units as

independent single connections, and any subsequent relays by units to a secondary recipient device occur over separate local links which are "appended" to the WAN/Unit-N connections, and are not coordinated by the WAN as part of a single multilink session.

69. I understand Red Hat also argues that a person of ordinary skill in the art would have understood messages 508 and 512 to include first information associated with the device's interface (e.g., an IP address) and that using TCP/IP for packet-based communication would have been obvious. I disagree. The use of TCP/IP for messages between the WAN system and participating Units 2-N is neither suggested nor implied in Kotzin. Instead, the communications are simply referred to as "messages" which would be appropriately interpreted as any type of message-based communication technology, not "packets" as might indicate a TCP/IP network. In fact, the word "packet" does not appear in Kotzin at all, but the term "any other suitable communication system" is used when describing the components of "communication system 100" [0022] .

70. In my opinion, the WAN's interactions are segmented and independent, not bundled across multiple paths to the same endpoints under a common session. A "multilink" connection is defined by interdependent logical paths between the same two endpoints, all bound to the same data transfer session and virtualized into one logical connection.

71. More specifically, distribution message 508 is single point-to-point message informing the WAN system of "recruits" who have volunteered to host independent, separate connections from the WAN system. Notably, the distribution message 508 is not requesting the server's capability to participate in a multilink communication. Nor is the distribution message 508 requesting or identifying a session identification to use. It is merely informing the WAN system of the list of "volunteers" which may or may not participate in future activities.

72. The 512 messages are likewise, point-to-point connections between the "server" (WAN) and the "client volunteers" (Units 2-N) which happens to transport one or more "tranches" of data from the server to each of the "client volunteers." The independent pieces of this "tranche" of data may be further forwarded or relayed to a different, single "client" device by the initial "client volunteers" using some other network connection, but the WAN system is uninvolved in this secondary part of the process. Any secondary connection(s) between the volunteer Units 2-N and a different client are "appended" to each of the end-to-end connections between the WAN system and each of Units 2-N.

73. Kotzin presents independent server-to-helper-unit transfers followed by separate unit-to-end-device relays that the server does not orchestrate as a single session. Accordingly, Kotzin does not disclose or suggest a device-originated "multilink request" to the server as required.

B. "multilink session identification"

74. Both claims 8 and 17 require the use of a "multilink session identification." As explained in the specification and Claim 17, the "multilink session identification" is an identifier for the multilink session:

[8.d] sending a data packet including the multilink session identification to the first network communication interface of the device

[8.f] send different groups of data packets over the at least one server network interface to the first network communication interface of the device and to the second network communication interface of the device, wherein each of the different groups of data packets includes the multilink session identification.

[17.e] transmit a second message through the second network interface to the at least one server, wherein the second message includes a multilink session identification that is unique to the multilink session and the destination URL[.]

75. In my opinion, a "multilink session identification" is an identifier that binds multiple concurrent links into a single virtualized end-to-end connection between two endpoints.

76. In a multilink configuration, multiple constituent end-to-end logical connections are bundled, and data for the same session may traverse separate parallel paths. A session identifier creates a common identity shared across those connections and is used to coordinate distribution, reassembly, quality-of-service, and link membership.

77. I understand Red Hat argues that session identifications were known or obvious, including the use of arbitrary or randomly generated numbers in single-

path protocols such as Real Time Streaming Protocol (RTSP, RFC2326), Session Initiation Protocol (SIP, EX1046 (RFC 2543)), and Dynamic Host Configuration Protocol from 1997 (DHCP, EX1047 (RFC 2131), §2, to track packets within a single end-to-end flow.

78. Those identifiers do not create a relationship among multiple concurrent end-to-end connections that together form a single virtualized session.

79. I understand Red Hat further argues that TCP/IP would most likely be used as the communication protocol—and that if wasn't used, it would have been obvious as taught by Held—such that it would make use of "a tuple comprised of source IP address/port number and destination IP address/port."

80. Real Time Streaming Protocol (RTSP, RFC2326), Session Initiation Protocol (SIP, EX1046 (RFC 2543)), Dynamic Host Configuration Protocol from 1997 (DHCP, EX1047 (RFC 2131), §2 are protocols for tracking packets in a single end-to-end session, not a multi-link or multi-path session.

81. TCP establishes a single, reliable, end-to-end connection between communicating endpoints using a three-way handshake process (SYN, SYN-ACK, ACK) and terminates connections with a four-step handshake (FIN, ACK, FIN, ACK). TCP packets consist of a header which contains multiple elements of metadata.

82. IP delivers packets of higher-layer protocol data units that consist of a header, containing source and destination addresses, packet length, and other metadata. The headers of IP packets are used to route data which transits systems on the Internet.

83. Once IP packets are delivered to endpoint systems, the headers of TCP packets are used to deliver data between corresponding processes on the endpoint systems.

84. A "tuple" is a set of parameters from the packet's metadata which is generally used to identify an end-to-end network connection or a network flow. For example, a TCP "4-tuple" combines the source and destination TCP port numbers with the source and destination IP addresses of a packet flow. The "4 tuple" is used by the operating systems at each end of the connection to direct data to the correct corresponding application.

85. This collection of metadata may be unique to a particular TCP session or end-to-end connection, but each corresponding node has a limited number of non-reserved TCP port numbers which must be re-used between TCP sessions. In a "multilink" or virtualized connection between two corresponding endpoints, each constituent TCP connection within the virtualized connection will use different source/destination TCP port identifiers for each endpoint and each connection, and will have the same source/destination IP addresses for the endpoints. As a result,

the "session identifier" would be an additional piece of metadata necessary to create a relationship between the separate TCP/IP connections and produce a single virtualized connection.

86. Based on my review, Kotzin does not disclose the use of such a "session identifier." Kotzin also does not teach providing such an identifier to or from the server to coordinate multilink membership. In fact, Kotzin specifically leverages independent, end-to-end connections between the WAN system and participating Units 2-N, which are established via "control" and "data" message pairs such as 512A/B (WAN and Unit 2) and 512C/D (WAN and Unit 3).

X. Claim Construction and Opinions on Ground 2 (Phatak Alone or in Combination with Held)

A. "request for a multilink communication"/"multilink request"

87. As set forth above with respect to Red Hat's Ground 1, claims 8 and 17 both require "a first multilink packet"/"a first message" to be sent from the device to the server that includes a "multilink request"/"request for multilink communication":

[8.c] receive a first multilink packet over the at least one server network interface, the first multilink packet comprised of a request for a multilink communication and first information associated with first network communication interface of the device that is capable of being used by the device to receive data packets[.]

[17.d] transmit a first message through the first network interface to at least one server, wherein the first message includes a multilink request and a destination URL[.]

88. I understand Red Hat relies on Phatak's identification of TCP/IP to attempt to establish a "multilink request"/"request for multilink communication." Specifically, Red Hat argues that in the context of Figure 1, "when HOST B wishes to stream video from HOST A over multiple interfaces at HOST B, it can send an initial request to HOST A for streaming using TCP/IP, such as a SYN packet" and that the "HOST B will notify HOST A of other IP addresses/interfaces where it can receive streaming data by using unused TCP/IP header fields."




Fig. 1
COMMUNICATING HOSTS WITH MULTIPLE NETWORK CONNECTIONS
(MULTI-HOMED).

89. In my opinion, and as explained above, TCP/IP does not implement multilink communication. TCP/IP is specifically not a multilink protocol. Instead, TCP creates a single reliable connection using a three-way handshake process (SYN, SYN-ACK, and ACK), and terminates connections with a four-step handshake (FIN, ACK, FIN, ACK).

90. IP delivers packets of higher-layer protocol data units that consist of a header for use to route data which transits systems on the Internet.

91. When a user requests data (e.g., a web page), the data may be divided into smaller groupings by both the TCP and IP layers of the protocol stack. These groupings of data, which become "segments" as sequences of data bytes are divided or "packets" as header and other data is appended to them, generally take different routes to the destination as they transit the Internet. If they are delivered out of order, TCP packets can be reassembled in the correct order via a sequence number embedded in the TCP header. The receiver in a TCP connection acknowledges the successful receipt of sequence of data bytes, or "segments" of transmitted data, which may have been transmitted using one or more "packets," and any missing or corrupted data is retransmitted.

92. As TCP/IP connections were end-to-end, singly-specified relationships between communicating applications or processes, they did not comprise or implement a multilink protocol.

93. Phatak describes multiple potential links between two hosts and proposes tunneling to address splitting flows across interfaces.

94. In tunneling, an outer header encapsulates an inner packet destined for a particular interface address. This encapsulation acts as an envelope around another envelope and does not provide the constituent links with a common session identity at the transport or higher logical layer. The tunnels do not cause the underlying paths to recognize or operate as a single, session-bound multilink bundle. The

sender assembles packets as if destined for one interface and externally encapsulates packets for transmission on other interfaces, creating multiple tunnel flows.

95. In my opinion, this design does not disclose a device-originated multilink request to the server that establishes a virtualized, session-identified bundle across multiple paths.

96. Phatak's configuration presumes multi-homed endpoints—hosts with multiple network interfaces associated with distinct IP addresses. For example, for Host A, interface A1 is connected to ISPA1 and interface A2 is connected to ISPA2, and similarly for Host B. However, Phatak's configuration cannot function when all tunnels have a common endpoint on a single host with a single network interface connected to a single ISP, as would be common in a multilink communication between two conventional endpoints.

97. The claimed invention of the '908 Patent does not require a "multihomed" configuration but instead employs a "session identifier" to relate multiple simultaneous connections between endpoints which are not multi-homed. Thus in the claimed invention, the "session identifier" is required to create the relationship or "virtual logical connection" involving multiple, independent TCP/IP connections established between communicating endpoints.

98. Accordingly, in my opinion, a TCP/IP-based initial request as described in Phatak does not disclose or suggest the required "multilink request."

B. "multilink session identification"

99. Both claims 8 and 17 require the use of a "multilink session identification." Similar to its Kotzin ground, Red Hat is relying on the TCP/IP tuple as the multilink session identification:

[8.d] sending a data packet including the multilink session identification to the first network communication interface of the device[.]

[8.f] send different groups of data packets over the at least one server network interface to the first network communication interface of the device and to the second network communication interface of the device, wherein each of the different groups of data packets includes the multilink session identification.

[17.e] transmit a second message through the second network interface to the at least one server, wherein the second message includes a multilink session identification that is unique to the multilink session and the destination URL[.]

100. I understand Red Hat argues, similar to its Kotzin ground, that the TCP/IP tuple serves as the "multilink session identification" in the Phatak combination. In my opinion, the required multilink session identification is not satisfied by TCP/IP tuples or by the multi-interface tunneling described by Phatak.

101. The TCP/IP tuple identifies a single transport-layer flow and is limited by port number to reuse across sessions. It does not provide a cross-flow identifier that binds multiple simultaneous connections into a single multilink session. In conventional use, the "tuple" is used to discriminate or differentiate between TCP/IP connections or flows, not to combine or commingle them.

102. In a true multilink configuration, multiple concurrent connections between the same endpoints will each have distinct transport port combinations. A tuple is therefore insufficient to serve as a shared session identifier across those connections, and could not be used for such a purpose without additional information (e.g. a "session identifier") to link the disparate connections into a single, logical connection.

103. A multilink session requires an explicit identifier included in messages and data packets across all participating connections. That identifier allows endpoints—and, where applicable, remote communications devices—to associate packets from separate, independent connections with the same logical session, perform reassembly, manage connection membership, and maintain session continuity.

104. Based on my review, Phatak does not disclose using such a "session identifier" to create a virtualized single connection across multiple end-to-end paths. Its tunneling encapsulation performs per-interface encapsulation without unifying transport flows under a common session identity.

105. In my opinion, this falls short of the required "multilink session identification" that is unique to the multilink session and carried in communications to bind multiple links into a single end-to-end session.

I declare under penalty of perjury that the foregoing is true and correct.

Executed this 6th day of November, 2025.

By Dr. Stan McClellan: 