

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

ARISTA NETWORKS, INC.,

Petitioner,

v.

CORRIGENT CORPORATION,

Patent Owner.

Case IPR2023-00839

U.S. Patent No. 9,118,602

REPLY DECLARATION OF PAUL MIN, PH.D.

TABLE OF CONTENTS

	Page
I. INTRODUCTION	1
II. MATERIALS RELIED ON IN FORMING MY OPINIONS.....	2
III. LEVEL OF ORDINARY SKILL IN THE ART AND PRIORITY DATE.....	2
IV. CLAIM CONSTRUCTION.....	3
V. THE '602 PATENT	4
VI. GROUND 1: RFC 3209 IN COMBINATION WITH RAZ RENDERS CLAIMS 1-26 OBVIOUS.....	4
A. The Combination of RFC 3209 and Raz Discloses or Renders Obvious the Claimed LAG.....	5
B. The Combination of RFC 3209 and Raz Discloses or Renders Obvious the Assignment and Use of a LAG Port to “Send[] Said Data Packet from Said First MPLS/LAG Switch to a Second MPLS/LAG Switch Via Said Single Physical Tunnel Port” (Limitation [1.5]).....	23
C. RFC 3209 and Raz Teach the Same Signaling Protocol.....	27
D. “Reasonable Expectation of Success” to Modify RFC 3209 Based on Raz.....	28
VII. THE APPLICATION OF RAZ AND FERGUSON TO RFC 3209 IN GROUND 2 RENDERS CLAIMS 1-26 OBVIOUS AS I ASSERTED IN MY FIRST DECLARATION	30
A. Ferguson Discloses a LAG	30
B. A POSITA Would Have Been Motivated to Incorporate the Teachings of Raz and Ferguson into the RFC 3209 System with a Reasonable Expectation of Success	32
1. Ground 2 relies on the Combination of RFC 3209, Raz and Ferguson	32
2. A POSITA would have had a Reasonable Expectation that the Combination would Result in a Successful and Functional System.....	33
C. The Combination of RFC3209, Raz and Ferguson Renders the Claimed Inventions of the '602 Patent Obvious.....	34
VIII. Conclusion	40

I, Paul Min, Ph.D., hereby declare as follows:

I. INTRODUCTION

1. I have been retained on behalf of Arista Networks, Inc. (“Arista”) to provide my technical review, analysis, insights, and opinions concerning the validity of the claims of U.S. Patent No. 9,118,602 (EX1001; “the ’602 Patent”) entitled “Tunnel Provisioning with Link Aggregation.” I understand that the patent is assigned to Corrigent Corporation (“Corrigent”).

2. I am the same Paul Min who provided a declaration on April 18, 2023, in support of Arista’s Petition for *Inter Partes* Review in this proceeding. See EX1003. I maintain the opinions that were set forth in that previous declaration. I provide this reply declaration to respond to certain opinions provided by Dr. George N. Rouskas in a declaration (the “Rouskas Declaration”) submitted in support of the Patent Owner Response filed by Patent Owner. The Rouskas Declaration is marked as Exhibit 2007 in this proceeding.

3. My background and qualifications were set forth in Paragraphs 1-13 of my original declaration in this proceeding, as well as my curriculum vitae marked as Exhibit 1003, and I incorporate that information by reference in this reply declaration. In this reply declaration, I apply the same understanding of the governing law as set forth in Paragraphs 14-21 of my original declaration.

II. MATERIALS RELIED ON IN FORMING MY OPINIONS

4. In forming my opinions in this reply declaration, I have relied on the '602 Patent's claims, specification and prosecution history, on the prior art exhibits to the IPR Petition (Paper 2), any other materials cited in this reply declaration, and my own knowledge, experience and expertise, and the knowledge of a person of ordinary skill in the art ("POSITA") in the relevant timeframe. I have also reviewed and relied upon the materials cited in my original declaration and the materials cited in the Rouskas Declaration. I have also reviewed the Decision on Institution of Inter Partes Review provided by the U.S. Patent and Trademark Office Patent Trial and Appeal Board ("PTAB") in this proceeding. *See* Paper 9.

III. LEVEL OF ORDINARY SKILL IN THE ART AND PRIORITY DATE

5. In this reply declaration, I apply the same definition of a POSITA as set forth in Paragraphs 60-61 of my original declaration. I note that Dr. Rouskas used the same definition of a POSITA in providing his opinions. EX2007, ¶¶21-24. In this reply declaration, I apply the same priority date for the '602 Patent of May 6, 2005, as was set forth in Paragraph 56 of my original declaration. Dr. Rouskas agrees on this priority date. EX2007, ¶21.

IV. CLAIM CONSTRUCTION

6. Paragraph 82 of my original declaration provides my opinion that the terms in the claims of the '602 Patent do not require construction in order to evaluate the prior art discussed in this proceeding. In the Institution Decision, the PTAB determined that there was “no need for express construction of any term” at that stage of the proceedings. Paper 9 at 18. In the Rouskas Declaration, Dr. Rouskas subsequently “accept[ed] and adopt[ed]” the position that no claim terms required construction. EX2007, ¶¶35-36. As such, I maintain my opinion that no claim term requires construction.

7. I note that the Rouskas Declaration points out that “a district court has previously construed” the term RSVP-TE found in Claims 3 and 15 of the '602 Patent to mean “RSVP-TE protocol in IETF RFC 3209,” and the same district court found the preamble of Claim 15 of the '602 Patent not to be limiting. EX2007, ¶37; *see also* EX2001. While I have not construed RSVP-TE as part of providing my opinions on the patentability of the claims of the '602 Patent, my opinions would not change if the PTAB adopted the definition used by the district court. Further, my opinions on the patentability of the claims of the '602 Patent assume the preamble of Claim 15 is limiting for the purposes of this proceeding, but my opinions would not change if the PTAB determines that the preamble of Claim 15 is not limiting.

V. THE '602 PATENT

8. Paragraphs 29-58 of my original declaration provide a description of the specification disclosure and claims of the '602 Patent, as well as the prosecution history for the patent application that resulted in the issuance of the '602 Patent. In Paragraphs 38-41 of the Rouskas Declaration, Dr. Rouskas provides his description of the specification and independent Claims 1, 15 and 26 of the '602 Patent. I generally agree with the description of the '602 Patent provided by Dr. Rouskas, and I reserve the right to rely upon the statements in these Paragraphs of the Rouskas Declaration. However, I disagree with Dr. Rouskas's opinion that "the claimed inventions of the '602 Patent were new and inventive at the time of the priority date" of the '602 Patent. EX2007, ¶38. It is my opinion that the claimed inventions of the '602 Patent were not new and were not inventive at the time of the priority date of the '602 Patent for the reasons stated in my original declaration and this reply declaration. I note that the U.S. Patent and Trademark Office did not consider any of the RFC 3209, Raz, or Ferguson prior art references during the prosecution of the application that resulted in the '602 Patent.

VI. GROUND 1: RFC 3209 IN COMBINATION WITH RAZ RENDERS CLAIMS 1-26 OBVIOUS

9. Paragraphs 42-147 of the Rouskas Declaration contain Dr. Rouskas's opinions that the combination of RFC 3209 and Raz does not render Claims 1-26 of

the '602 Patent obvious. I disagree for the reasons stated in my original declaration and the reasons set forth below.

10. Paragraph 42 of the Rouskas Declaration asserts Dr. Rouskas's opinion that "the alleged combination of RFC 3209 and Raz does not render obvious claims 1-26, including because (i) neither RFC 3209 nor Raz discloses or renders obvious the claimed "LAG," (ii) a POSA would not have combined RFC 3209 and Raz as Petitioner alleges, (iii) Petitioner has failed to establish a motivation to modify RFC 3209 with Raz to arrive at the claimed inventions of the '602 patent, and (iv) Petitioner has failed to establish a reasonable expectation of success for any such modifications. I disagree for the reasons stated in my original declaration and the reasons set forth below.

A. The Combination of RFC 3209 and Raz Discloses or Renders Obvious the Claimed LAG

11. Dr. Rouskas asserts that "neither RFC 3209 nor Raz discloses or renders obvious the claimed 'LAG'" (EX2007, ¶43), but misapplies the Link Aggregation standard (802.3ad) and ignores what a POSITA would have understood from Raz. LAGs were well-known in the art and widely used to communicate packet streams before the priority date for the '602 Patent, as demonstrated by Raz (EX1006, Col. 9:9-28), Ferguson (EX1007, Cols. 1:57-2:3, 1:28-42, 3:36-4:19,

Figures 2A-2B). *See* EX1036, 21:19-22:6, 27:20-28:16 (deposition testimony from Dr. Rouskas acknowledging the LAGs were well-known and used in the art).

12. Dr. Rouskas argues that “a POSA would have understood that a LAG is a group of network connections (or links) that are **between** two network components.” EX2007, ¶46. The basis for Dr. Rouskas’ opinion that Raz does not disclose or render obvious the claimed “LAG” appears to be a very limited conception of what a LAG is. *See* EX1036, 32:17-18 (testifying his understanding that “‘link aggregation group’ has a very specific meaning.”). I disagree that the concept of LAG is so limited, and I believe that Dr. Rouskas’ opinions conflict with the 802.3ad standard. A Link Aggregation Group, or “LAG,” refers to the aggregation of *ports*. The Ethernet LAG standards-defining document, 802.3ad, which a POSITA would have considered an authority on LAGs, states that “Link Aggregation allows a MAC Client to treat a set of one or more ports as if it were a single port.” EX1033, §43.2.1. To create a LAG, “[t]he Aggregator binds to one or more ports within a System.” EX1033, §43.2.1. It is the binding of ports to an aggregator that creates a LAG. EX1033, §43.2.1, §43.3.7 (describing how “[e]ach port is selected for membership in the Link Aggregation Group.”). When this is done, “[e]ach port is assigned a unique, globally administered MAC address.” *Id.*, §43.2.1.

13. To understand how the LAGs are formed, Figure 43-1 of the 802.3ad standard, shown below, illustrates that “Link Aggregation comprises an optional sublayer between a MAC Client and the MAC.” This means that Link Aggregation occurs at the MAC layer, which is part of Layer 2 of the Open Systems Interconnect (OSI) standard (see EX1035) residing *inside of* a switch, as opposed to the Physical Layer or Layer 1 of the OSI, which interfaces with the external connection. See EX1033, Fig. 43-2, §43.1.4.




Figure 43-1 – Architectural positioning of Link Aggregation sublayer

EX1033, §43.1.3, Fig. 43-1 (illustrating the “Architectural positioning of Link Aggregation sublayer” at the “data link” MAC layer, separate from the “physical” layer).

14. Figure 43-2 of the 802.3ad standard, shown below, further illustrates the details of the “Link Aggregation Sublayer,” where the link aggregation takes place and where the Link Aggregation Groups (LAGs) are formed.




Figure 43-2—Link Aggregation sublayer block diagram

EX1033, §43.1.4, Fig. 43-2 (illustrating the details of the “Link Aggregation sublayer block diagram.”).

15. As shown in Figure 43-2 above, data frames from a MAC client are distributed by the Frame Distribution function to one or more of the Aggregator Parser/Multiplex functions within the Aggregator. The 802.3ad standard explains that “[a] given port will bind to (at most) a single Aggregator at any time,” and that “[t]he binding of ports to Aggregators within a System is managed by the Link Aggregation Control function for that System, which is responsible for determining which links may be aggregated, aggregating them, binding the ports within the System to an appropriate Aggregator, and monitoring conditions to determine when a change in aggregation is needed.” EX1033, §43.2.1. This means that each of the Aggregators has one or more ports bound by the Link Aggregation Control function and according to the LAGs that have been formed.

16. Everything I have just described occurs within the Link Aggregation sublayer. By the time that the aggregated data frames are passed to the MAC (at the bottom of Figure 43-2) via the optional MAC Control function, the LAGs have been formed.

17. To be sure, the physical links—the physical connections such as Ethernet cables which carry traffic between components—occur at the Physical Layer of the OSI as shown at the bottom of the figure above. *Id.* By the time any of the data frames reach the Physical Layer, these data frames would have already been sorted into the LAGs to which they belong. As explained above, link aggregation is

implemented by grouping the *ports*, which takes place at the MAC layer above and separate from the Physical Layer. This means that no matter the form with which a particular Physical Layer may be implemented (e.g., multiple 1 Gigabit Ethernet cables or a single 10 Gigabit Ethernet cable) it is still the case that certain ports have been aggregated to become part of a LAG. EX1033, §43.1.3 (“ . . . implement the optional Link Aggregation sublayer for some *ports* within a System . . . it is not necessary for all *ports* in a System to be subject to Link Aggregation.”), §1.4 (describing an “Aggregation Key” as a “parameter . . . identifying those *ports* that can be aggregated together”), §1.4 (describing an “Aggregation System” as a “uniquely identifiable entity comprising (among other things) an arbitrary grouping of one or more *ports* for the purpose of aggregation. An instance of an aggregated link always occurs between exactly two Aggregation Systems.”), §43.3.5 (“the ability of one *port* to aggregate with another . . .”). 802.3ad describes such a port as an “Aggregation Port,” which it defines as “[a]n instance of a MAC-Physical Layer entity within an Aggregation System.” EX1033, §1.4. I explained this during my deposition taken in connection with this proceeding. EX2013, 45:13-22 (“Link aggregation group happens at the layer 2 of the stack, the MAC layer. The MAC layer is what determines the link aggregation group by assigning ports in the LSR to belong to a particular link aggregation group, and this is so according to 802.3ad. It

is -- that does not concern how the physical layer implements or transmit between the routers.”).

18. Dr. Rouskas’ assertion that a LAG refers exclusively to “links” is overly limited because it does not appear that he is taking into consideration the necessary *port* component of the LAG. 802.3ad states that “the term *link* refers to an *Aggregation Link*” (EX1033, §43.1.1 (emphasis in original)) and defines an “Aggregation Link” as “[a]n instance of a MAC-Physical Layer-Medium[-]Physical Layer-MAC entity between a pair of Aggregation Systems.” EX1033, §1.4. This says that the “link” of link aggregation is an instance of an entity that spans the entire path from one Aggregation System to another, starting with an aggregated group of *ports* at the MAC layer and ending at the paired aggregated group of *ports*, also at the MAC layer. As I have explained, the aggregation of these *ports* creates a LAG. The path of a data packet bound for each *port* of the group follows the “MAC-Physical Layer-Medium[-]Physical Layer-MAC” instance of the “Aggregation Link” definition. *Id.*

19. It is my opinion that Raz describes an implementation of LAG. I have testified that a POSITA would understand that the subports are configured in a LAG, and I stand by that testimony. Raz teaches provision of subports at both ends of a high-speed 10G PHY link. EX1006, 7:7-16. The “PHY” refers to the physical layer, which interfaces with the MAC layer at the LSR and with the subports at this MAC

layer. *Id.* Raz describes that the high-speed traffic comes into one of the paired LSRs on an interface which carries “the ten individual 1G packet streams” bound for the ten subports, and that Raz “trunks” that port into the subports which are logically associated together. EX1006, 7:14-16. It is my opinion that this describes the aggregation of ports and is an implementation of link aggregation. I also explained this during my deposition for this proceeding. EX2013, 41:22-42:7, 43:22-44:12, 47:5-8. As I understand that Dr. Rouskas also acknowledged in a proceeding involving the same patent, there are many implementations of LAG. EX1036, 17:10-13.

20. As further support, I explained in my original declaration that Raz teaches that the subports on either end of the high-speed PHY are paired, as in link aggregation. EX1003, ¶¶67-68. Referring to Raz’s Figure 3, at either end of high-speed PHY 60 are aggregated subports 56 (at LSR A) and 62 (at LSR B). These aggregated ports are paired together such that “[t]he selected subport is the same for both sides of a 10G link.” EX1006, 7:20.




FIG.3

EX1006, Fig. 3 (showing subport 0 provisioned on both ends of PHY 60). Raz discloses the assignment of an LSR's "receiving" subport to an LSP, which Dr. Rouskas does not appear to dispute. EX1006, 7:26-29 ("it is the downstream (i.e. receiving) [LSR] that determines the subport it wants to receive the LSP on."). Raz also discloses that when a subport is selected by an LSR, "[t]he selected subport is the same for both sides of a 10G link." EX1003, ¶67 (citing EX1006, 7:20). Raz discusses this as follows in connection with Figure 3: "LSR B decides to accept the LSP on subport 3 and thus assigns a label comprising an indication that subport 3 was selected. The label is communicated to LSR A which now knows that the output subport should be subport 3 (i.e., the same subport is provisioned)." EX1006, 7:33-38, Fig. 3. At LSR A, "[t]he output of all the subports is aggregated and combined

into a single high speed packet stream and transmitted over the link.” EX1006, Abstract. Then, at LSR B, “[i]nbound packets are received by the high speed PHY and forwarded to one of a plurality of subports.” *Id.* A POSITA reading Raz’s disclosure would have concluded that the pairing of subports on LSRs A and B, on both sides of the high-speed PHY link, such that the receiving subport “ha[s] the same subport index as the one through which the packet was sent out over at the transmitting node,” refers to an implementation of link aggregation. This is because Raz’s subports provide multiple 1G pathways for handling packet streams that are logically associated and treated as a single aggregated 10G pathway. EX1006, 7:7-16 (“10G interface that is composed of ten individual 1G packet streams”).

21. Explaining further, Raz describes that the motivation for aggregating its subports is consistent with the LAG motivation described by 802.3ad: that ports combined in parallel increase available bandwidth beyond what a single port was capable of alone. Pet., 6, 17-18, EX1006, 1:40-59 (“A problem arises[] when the aggregate link speeds exceed the capacity of the [internal processing. . . . One solution[] is to split the input packet stream into smaller bandwidth links. . . . There is thus a need for a forwarding and provisioning mechanism that enables a high speed packet stream to be processed by multiple processors that operate at a slower rate tha[n] the inbound packet stream.”); see EX1003, ¶¶64, 85, 96. Consistent with 802.3ad, Raz teaches “assigning a unique Media Access Control (MAC) address to

each subport.” EX1006, 2:52-56; *see* EX1033, §43.2.1 (“[e]ach port is assigned a unique, globally administered MAC address.”).

22. That Raz teaches the selection of a *port* of an aggregated set of ports is consistent with the claims of the ’602 Patent, which require the selection of a “single physical tunnel port of a LAG.” *See, e.g.*, EX1001, Claim 1. The ’602 Patent’s claims further describe the selected “physical data port *in* an Ethernet Link Aggregation Group (LAG).” *Id.* That the claims of the instant patent themselves describe ports as *in* and *of* a LAG further supports my opinion that a LAG is a collection of ports. Raz also describes LAG in terms of its ports. EX1006, 9:12-15 (“defining its *ports* ... as a single LAG”). Prior art reference Ferguson, discussed below, does as well. EX1007, 1:47-52 (“aggregated *ports*”). Reinforcing this point, the ’602 Patent incorporates the LAG standards document, 802.3ad, by reference. EX1001, at 2:1-17.

23. Dr. Rouskas asserts that Raz’s subports are not a LAG because they are implemented “within” an LSR and not “between” two LSRs. I believe that he takes an overly narrow approach to LAG and misunderstands Raz’s disclosure. That Raz’s subports are physically located within an LSR does not disqualify them from being a LAG. As I explained above, the 802.3ad standard describes the positioning of link aggregation within the *MAC* layer, which is not “between” LSRs but operationally within LSRs. EX1033, Fig. 43-2, §43.1.4. Additionally, 802.3ad states that “[a]n

instance of an aggregated link always occurs between exactly two Aggregation Systems,” where “a physical device may comprise a single Aggregation System *or more than one Aggregation System.*” EX1033, §1.4. I understand this to teach that even under Dr. Rouskas’ limited interpretation of LAG as a link, a LAG may be implemented wholly within a physical device such as an LSR. Confusingly, Dr. Rouskas apparently agreed that Raz properly taught use of a LAG within an LSR. *Id.* at 33:19-24 (“**[Raz] gives, you know, the definition of, ‘LAG (Link Aggregation Group).’** So when he—and—**and by the way, he is using that within the switch.** So now he’s taking a LAG switch and he’s repurposing it to use for his embodiment inside the switch.”); 34:18-19 (“Raz uses the term ‘LAG’ in a way that is normally understood in the[] arts.”).

24. A LAG is not defined by the characteristics of the link or links that are connected to its ports, and it is not correct to say that a LAG must be “between” LSRs. While a POSITA understands that ports are connected to links, the ‘group’ of a link aggregation group properly refers to the *ports*, and ports do not sit “between” LSRs. This is because those skilled in the art generally understand that what is actually happening when interfaces are bundled together is happening *to the ports*. The ‘bundling’ of any connected links is a consequence of grouping the ports. EX2013, 43:22-44:12 (“Multiple physical streams combining to a single physical stream, that’s—that’s a consequence of grouping ports [] as a single -- link

aggregation group.”). It may be the case, as a POSITA would have understood Raz to describe, that the output of LAG ports goes through an aggregator before it is sent out on a single link to another LSR. As I explained during my deposition, “as for how they’re [(the packet streams)] actually being submitted across and between the two switches, two LSR switches, that is not part of the definition of link aggregation group.” EX2013, 43:22-44:12.

25. A POSITA looking at Figure 6 would understand that Raz describes an ordinary ethernet switch with some ports collected in a LAG group connected, via a high-speed link, to a grouping of parallel ports in another LSR. In comparison, Figure 2 shows an ordinary ethernet switch with *all* ports shown as part of the LAG group. Thus I disagree with Dr. Rouskas’ opinion expressed in his declaration, as I believe that Raz discloses or renders obvious the claimed “LAG.”





FIG.6

EX1006, Fig. 6.



EX1006, Fig. 2.

26. Even under Dr. Rouskas' limited implementation of a LAG, it is my opinion that the combination of RFC 3209 and Raz discloses or renders the claimed LAG obvious. Functionally, Raz describes the hop-by-hop reservation of bandwidth to an LSP, because Raz teaches that reserving bandwidth on the previous LSR's subport is the same as to reserving bandwidth on the choosing LSR's subport. This reservation spans the entire hop from the previous LSR's output subport to the receiving subport on the receiving LSR. A POSITA looking at these disclosures would have considered this a description of reservation of resources on a LAG operating "between" LSRs, effectively reserving the entire "link" of a plurality of "links" between them.

27. As I also pointed out in my initial declaration (EX1003, ¶64), Raz expressly discloses LAG and specifically discloses “defining its ports leading to the different subports as a single LAG.” EX1006, 9:9-15. This passage explains that a LAG may lead to the described subports, meaning that in order to be forwarded to a particular selected subport, the LSR needs to know which LAG port connects to which subport. EX1006, 9:64-65. Choosing a subport in this embodiment necessarily requires identification of the LAG port leading to that particular subport in order to forward the data packet to the appropriate subport. Thus, even assuming arguendo that Raz’s subports are not a LAG, Raz discloses ports leading to the different subports which explicitly are. Dr. Rouskas acknowledged this during his deposition, agreeing that “[the inventor of Raz] knew how to repurpose a LAG switch to use inside [] his invention,” and “[Raz]’s talking about Layer 2 LAG that he proposes to use as part of one embodiment.” EX1036, 33:6-8, 33:21-34:5. EX1006, 9:18-22, 9:41-10:4.

28. Dr. Rouskas asserts that if the inventor of Raz was aware of LAG, and considered his invention’s subports to be a LAG, he would have called them a LAG. EX2007, ¶50. I disagree. A POSITA would consider Raz’s subports to be a LAG for the reasons discussed above.

29. I additionally do not agree with Dr. Rouskas’ attempt to distinguish Raz’s subports due to the differences between “multiplexing” and “inverse

multiplexing.” EX2007, ¶49. First, Dr. Rouskas appears to misapply the principals of multiplexing in the context of Raz. Raz teaches multiplexing in conjunction with link aggregation. “Aggregation” to create a LAG is the logical association of the transport entities—it is the fact that the ports are assigned to a particular trunk. Multiplexing is done on the *traffic itself*; it is concerned with collecting the packets that are output by the transport entities, and with packing those packets into the stream that is to be transported to the PHY link. These concepts are not incompatible. 802.3ad itself describes that “multiplexing” is performed within the Link Aggregation sublayer.




Figure 43-2—Link Aggregation sublayer block diagram

EX1033, Fig. 43-2, §43.1.4.

30. Second, Dr. Rouskas’ opinion appears to conflict with the authority he quotes. He states that “combining or multiplexing of subport streams into a single stream is the opposite of aggregating multiple parallel links into a group to form a *LAG (known as inverse multiplexing)*.” EX2007, ¶49. But Dr. Rouskas goes on to

quote that “[w]hile multiplexing combines multiple signals into a single signal, ***[inverse multiplexing] divides a data stream into multiple concurrent streams that are transmitted at the same time across separate channels***”). EX2007, ¶49 (quoting EX2011). This is not inconsistent with Raz. Raz teaches that the receiving LSR “divides a data stream” (the high speed 10G PHY comprised of multiple 1G data streams) “into multiple concurrent streams” (the multiple component 1G data streams) “that are transmitted at the same time across separate channels” (separate subports at the receiving LSR). *Id*; see EX1006, 7:7-11 (“[T]he example LSR shown comprises a 10G interface that is composed of ten individual 1G packet streams. The ten subports are provided wherein each is adapted to handle a 1G packet stream. The overall aggregate speed, however, is 10G.”), 7:49-52. Dr. Rouskas himself says that this is “inverse multiplexing” and an implementation of a LAG. On my review, all of the authority quoted by Dr. Rouskas in paragraph 49 of his declaration supports my understanding that Raz’s subports would be considered by a POSITA to be a LAG.

31. Finally, I understand that Dr. Rouskas argues that “Raz describes its subports as internal network processors, not links between two endpoints of a network” (EX2007, ¶47), but this misunderstands the relationship between subports and processors. Raz describes the subports as 1G interfaces trunked from a 10G interface, which a POSITA would have understood refers to ports. EX1006, 7:7-16

(“the high speed interface (e.g., 10G) may be trunked into any number of individual subports.”). Ports contain or are served by processors; they are not themselves solely processors. Raz itself notes that there are network processors which handle the traffic through its subports: “each subport [may be] handled by a different network processor;” alternatively, it may be “the case that network processors are used to handle two subports each.” EX1006, 11:24-29. A POSITA would have understood that subports are ports that are served by network processors, not simply network processors themselves.

B. The Combination of RFC 3209 and Raz Discloses or Renders Obvious the Assignment and Use of a LAG Port to “Send[] Said Data Packet from Said First MPLS/LAG Switch to a Second MPLS/LAG Switch Via Said Single Physical Tunnel Port” (Limitation [1.5])

32. It is my opinion that the combination of RFC 3209 with Raz render obvious the claimed assignment and use of a single physical tunnel port of a LAG that is used for sending, as recited by limitations [1.1] “assigning a single physical tunnel port of a LAG to a network tunnel” and [1.5] “sending to a second MPLS/LAG switch via the single physical tunnel port.” EX1001, Claim 1.

33. Dr. Rouskas argues that “a POSA would not and could not have implemented Raz’s subport-related teachings to apply to communications between two network devices in RFC 3209.” EX2007, ¶59. To the extent Dr. Rouskas argues that the combination does not teach or suggest the assignment and use of a LAG

physical tunnel port to transmit data packets from one LSR to another LSR, I disagree for the reasons set out in my original declaration and for the following reasons.

34. I have already addressed above why a POSITA would not consider a LAG to be “between” network devices and why, in any case, Raz already relates to communications “between” LSRs using the claimed LAG. It is my opinion that the combination discloses or renders obvious the selection of a LAG port and assignment of that port to an LSP. Raz discloses the assignment of a receiving LSR’s “receiving” subport and a transmitting LSR’s “output” subport to an LSP, which Dr. Rouskas does not appear to dispute. EX1006, 7:26-29 (“it is the downstream (i.e. receiving) [LSR] that determines the subport it wants to receive the LSP on.”), EX1006, 7:20 (“[t]he selected subport is the same for both sides of a 10G link.”), *see also* EX1003, ¶67. Raz discusses this as follows in connection with Figure 3: “LSR B decides to accept the LSP on subport 3 and thus assigns a label comprising an indication that subport 3 was selected. The label is communicated to LSR A which now knows that the output subport should be subport 3 (i.e., the same subport is provisioned).” EX1006, 7:33-38, Fig. 3. It is my opinion that Raz discloses assigning an LSR’s sending port (i.e., outgoing interface) to an LSP. Referring to Figure 3, Raz teaches that LSR B assigns LSR A’s output subport to the LSP. EX1006, 7:33-38, Fig. 3. As I discussed above, a POSITA would have understood that Raz

describes hop-by-hop reservation of bandwidth to an LSP. Because Raz teaches reserving bandwidth on the transmitting LSR's subport in addition to reserving bandwidth on the receiving LSR's subport, the reservation spans the entire hop from the transmitting LSR's subport, across the PHY, and through the incoming subport on the receiving LSR. A POSITA looking at these disclosures would have considered this a description of reservation of resources on an aggregated link operating "between" LSRs.

35. Additionally, Raz's disclosure of reserving bandwidth across one of the hops in the LSP path is not different from RFC 3209's disclosure of "next hop" selection. The difference in the LSP setup process between RFC 3209's and Raz's teachings regarding interface selection is that Raz teaches an LSR's selection of *the previous LSR's* outgoing interface (the upstream hop with respect to the LSR), while the portions relied upon in RFC 3209 refer to an LSR's selection of *its own* outgoing interface (the downstream hop with respect to the LSR). This difference is inconsequential. This difference is not material to the proposed combination: a teaching from Raz that is proposed to be combined with RFC 3209 is that an LSR is capable of selecting a LAG subport and assigning it to an LSP. RFC 3209 discloses setting up label-switched paths using basic, well-known traffic engineering functions of reserving bandwidth and allocating labels identifying the entire path. EX1003, ¶¶130-31. Raz confirms that LSP setup is "conventional" for an LSR.

EX1006, 7:65-8:1; EX1003, ¶64 (“Raz discloses an LSR provisioning a label-switched path (LSP) and in particular, assigning a subport to the LSP.”) (citing EX1006, 7:65-8:13). RFC 3209 teaches the assignment of an outgoing interface of an LSR to an LSP, as claim limitations [1.0] and [1.1] require. Raz teaches the concept of applying LSP setup with traffic engineering to environments employing aggregated links and specifically the concept of setting up an LSP whereby a LAG subport is chosen as part of that LSP. Thus, it would have been obvious to a POSITA for RFC 3209’s selected outgoing interface to be a LAG subport. There is motivation to apply Raz’s allocation of a specific LAG subport to other points along the LSP, in addition to aggregation which is “internal” to an LSR. EX1003, ¶90.

36. I believe this difference is inconsequential for the additional reason that Raz is not concerned with which LSR (downstream or upstream) makes the subport determination for a particular LSP hop. Raz describes that “[t]he subports used on both sides of the next hop, however, are independent of the current hop.” EX1006, 6:43-44. Raz explains that this is the case because “the next link may have other connections in addition to the 10G link and therefore it may be the case that a different subport is the only subport (or the subport best suited) that can accept the LSP with the required bandwidth.” EX1006, 7:20-26. Raz is concerned, primarily, with ensuring that there is sufficient bandwidth reserved on each “hop” on the LSP, recognizing that an LSR may have other connections such that the same subport

cannot be used for every hop. I note that it *could* be the case that the same subport is used for every hop, depending on the bandwidth requirements of the LSP. But in general, Raz is concerned with ensuring, at each hop, that the LSP bandwidth considerations are met. *Id.* This goal does not depend on whether the transmitting LSR or the receiving LSR evaluates the hop between the two LSRs and performs the subport selection and assignment method. In fact, we know that both LSRs have the same information about available bandwidth on each side of the 10G PHY, because Raz teaches that there is a one-to-one mapping between the sets of subports. EX1006, 7:33-38, Fig. 3. As an example, subport 0 on LSR B would always have the same bandwidth available as subport 0 on LSR A. Thus, both LSR A and LSR B would always have the same information. Raz's stated goals are served regardless of which LSR performs the bandwidth allocation method. It is therefore my opinion that the choice for the transmitting LSR to make the subport selection would have been an obvious design choice. It is further my opinion that Raz's teachings regarding subport selection would have readily and easily been implemented into the system described by RFC 3209.

C. RFC 3209 and Raz Teach the Same Signaling Protocol

37. Dr. Rouskas argues that "the assignment and labeling methods in RFC 3209 and Raz are incompatible." I disagree. RFC 3209 and Raz are fully compatible at least because they use same labeling protocols. Raz describes that its LSR

“performs the conventional LSR functionality including the provisioning of LSPs. In particular, the CPU [of the LSR] functions to perform the assignment of subports to LSPs in response to LSP setup requests.” EX1006, 7:65-8:13. All MPLS switches employing RSVP traffic engineering are capable of responding to setup requests and provisioning LSPs. As I explained in my original declaration, RSVP dictates the process by which LSPs are set up according to reservation of path resources. One method of accomplishing this, as described by RFC 3209, is according to decisions made at each LSR and communicated to other LSRs along the path by allocating and transmitting labels. These labels are appended to data packets bound for that LSP so that an LSR needs only to examine the label to know where to forward the packet.

38. Dr. Rouskas argues that “RFC 3209’s assignment of an outgoing interface” is different from “Raz’s assignment of a subport is done by a downstream node” (EX2007, ¶72), but I have already explained above why this difference is inconsequential in view of the actual disclosures relied upon from Raz. *Supra* §VI.B.

D. “Reasonable Expectation of Success” to Modify RFC 3209 Based on Raz

39. Dr. Rouskas argues that “the Petition does not contain persuasive argument or evidence that a POSA would have had a reasonable expectation of success for modifying Raz to meet the challenged claims.” EX2007, ¶85. I disagree. As I explained in Paragraphs 14 and 17 of my original declaration, it is my

understanding that the obviousness analysis requires that, where there is a reason to modify or combine the prior art to achieve the claimed invention, there must also be a reasonable expectation of success in so doing to render the claimed invention obvious. However, it is my understanding that the Petition and my supporting opinions do not have to specifically use the words “reasonable expectation of success” in order to establish that a modification of the prior art will be successful. Further, it is my opinion that the Petition, my original declaration and this reply declaration all provide sufficient explanation why a POSITA reasonably would have expected success in applying the teachings of Raz to RFC 3209 so that LSPs “(1) can pass through aggregated links (such as LAG) and still employ traffic engineering functions of reserving bandwidth of the aggregated link to the LSP, and (2) support traffic engineering so that LSPs can handle high speed packet streams and wide connections,” as well as other recognized improvements. Pet., 13 (citing EX1003, ¶90); see Pet., 11-12 (citing EX1003, ¶88); EX1003, ¶¶88-91.

40. Dr. Rouskas also asserts his opinion that “a POSA would not have had a reasonable expectation of success of modifying RFC 3209 based on Raz to meet the claims, including because a POSA would not have even had a reasonable expectation of success of somehow applying Raz’s mechanism to communications between two network devices of RFC 3209 to begin with.” EX2007, ¶89. Dr. Rouskas follows this with an unsupported assertion that “modifying Raz to apply to

such communications between network devices would render Raz inoperable for its intended purpose.” EX2007, ¶89. I disagree with Dr. Rouskas for the reasons stated in my original declaration and throughout this reply declaration. I disagree at least because the proposed combination does not “modify[] Raz”—rather, teachings from Raz are incorporated into the system of RFC 3209. I disagree additionally because Raz’s invention would work as designed when features of it are incorporated into RFC 3209: both RFC 3209 and Raz describe the use of the same devices (LSRs) provisioning the same type of network path (LSP) and using the same labeling protocol (MPLS with RSVP Traffic Engineering extensions) to perform LSP bandwidth reservations on LSR interfaces. *See supra* §VI.B; EX1003, ¶¶90-91, 101-03. Furthermore, as I have already explained above, a POSITA would already consider Raz’s teachings to apply to communications between network devices, as Raz teaches reservation of resources along an entire hop from one LSR’s subports to the next LSR’s subports, which functionally, if not explicitly, applies to communications between the two LSRs. *Supra*, §VI.B.

VII. THE APPLICATION OF RAZ AND FERGUSON TO RFC 3209 IN GROUND 2 RENDERS CLAIMS 1-26 OBVIOUS AS I ASSERTED IN MY FIRST DECLARATION

A. Ferguson Discloses a LAG

41. As I noted in my first declaration, Ferguson explicitly discloses a LAG. EX1003, ¶¶ 243, 246, 249, 251. For example, in discussing the background of

network data communication, Ferguson explicitly identifies the relevant specification for LAG, IEEE 802.3ad, in discussing the nature of a LAG. EX1007, 1:27-49. This highlights that for a POSITA, as of the filing date for Ferguson, November 12, 2002, LAGs were a common part of the state of the technology for network data communication. It also confirms that IEEE 802.3ad defines the nature and requirements of LAGs, as previously discussed. *Supra* §VI.A.

42. Ferguson also discusses LAG in the context of the claimed inventions. Ferguson discloses that routers can be connected to enable communication “via **two or more physical interconnects that are logically associated as an aggregated link.**” EX1007, 3:39-44. Additional disclosures note the use of MPLS labels across label switched paths via LAGs, noting routers connected “via an **aggregated link.**” EX1007, 3:47-57. Ferguson further describes how “devices may make use of [the MPLS] protocol information to facilitate distribution of the MPLS packets across multiple paths, e.g., **multiple physical interconnects of an aggregated link.**” EX1007, 4:16-19. Further disclosures illustrate the distribution of MPLS packets across physical interconnects of a LAG, noting the aggregate interface (EX1007, 6:15-23), and the forwarding of the MPLS packets along “one of the paths” or “**physical interconnects of an aggregated link.**” EX1007, 6:52-55.

43. As I confirmed in my first declaration (EX1003, ¶¶ 243, 246, 249), and have discussed further here, Ferguson describes LAGs extensively, including in the

context of sending and receiving MPLS data packets over label switched paths of a LAG in an MPLS network.

B. A POSITA Would Have Been Motivated to Incorporate the Teachings of Raz and Ferguson into the RFC 3209 System with a Reasonable Expectation of Success

1. *Ground 2 relies on the Combination of RFC 3209, Raz and Ferguson*

44. Dr. Rouskas asserts that “Ground 2 includes no discussion regarding the specific teachings of primary reference RFC3209 or any modifications thereof to meet the claims.” EX2007, ¶ 150. This is inaccurate. Dr. Rouskas also asserts that Ground 2 relates to Raz and Ferguson, not primary reference RFC 3209. Ex 2007, ¶¶ 148-151. This is also inaccurate.

45. As I stated in my first declaration, “a POSITA would have understood that the techniques disclosed in Raz for using subports to transmit data packets within a router or switch would also have been applicable to an Ethernet physical data port in a LAG that communicates data packets between routers or switches in an MPLS network and readily implementable and compatible with Raz’s disclosures *when each of Raz and Ferguson’s teachings are put into RFC 3209’s system framework.*” EX1003, ¶ 246. I then went on to describe the system in RFC 3209 (EX1003, ¶ 245), the disclosures of Raz and Ferguson and how those disclosures would be applied to the system in RFC 3209 (EX1003, ¶¶ 244 – 246). In providing my opinions for Ground 2, asserting claims 1 – 26 were obvious over RFC 3209 in

view of Raz and in further view of Ferguson, I provided a claim limitation by claim limitation analysis, incorporating my earlier analysis regarding RFC3209 and Raz and focusing the additional analysis on Ferguson, where necessary. EX1003, ¶¶ 247 – 327. Contrary to Dr. Rouskas’s assertion (2007, ¶ 148), Ground 2 relates to primary reference RFC 3209 and explicitly discusses RFC 3209, where relevant to the analysis. *See, e.g.*, EX1003, ¶¶ 246, 262, 293. A POSITA would have been motivated to apply the teachings of Raz and Ferguson to RFC 3209 to arrive at the claimed inventions of the ’602 Patent. EX1003, ¶¶ 243 – 246.

2. *A POSITA would have had a Reasonable Expectation that the Combination would Result in a Successful and Functional System*

46. Dr. Rouskas repeatedly asserts that a POSITA would not have had a reasonable expectation of success in combining RFC 3209, Raz and Ferguson. EX2007, ¶¶180-83. I disagree. In my first declaration I highlighted that RFC 3209, Raz and Ferguson are focused on solving the same problem of ensuring enough bandwidth to communicate high speed packet streams through an MPLS network. EX1003, ¶¶ 244 -246. In addition, RFC 3209, Raz and Ferguson are all related to the use of an MPLS network to communicate data packets (EX1003, ¶¶ 243), and Raz and Ferguson discuss the use of a LAG in an MPLS network to communicate data packets (EX1003, ¶ 243). For at least these reasons and others identified in my first declaration, a POSITA would have been motivated to combine the references,

and because of these commonalities across the references would have had a reasonable expectation of success.

47. The nature and substance of these references would motivate a POSITA to combine them and would also suggest to a POSITA that there would be a reasonable likelihood that such a combination would be successful. It is my understanding that the reasonable expectation of success may be implied and the words ‘reasonable expectation of success’ do not need to be explicitly recited to communicate an opinion regarding a reasonable expectation of success. My first declaration in providing the motivations to combine RFC 3209, Raz and Ferguson for Ground 2, also indicated the reasons why a POSITA would reasonably expect the combination to be successful.

C. The Combination of RFC3209, Raz and Ferguson Renders the Claimed Inventions of the ’602 Patent Obvious

48. I reviewed Dr. Rouskas’s opinions with respect to Ground 2. The crux of Dr. Rouskas opinions arguing that a POSITA would not have combined Raz and Ferguson to arrive at the claimed inventions are based on a mischaracterization of the proposed combination for Ground 2, and a persistent assertion contrary to the explicit disclosures of Raz and Ferguson, as well as the knowledge of a Person of Ordinary Skill in the Art, that a POSITA would not have applied techniques

disclosed as operational within a network device to connections between network devices.

49. First, as discussed in VII.B.i, the combination in question is not simply a combination of Raz and Ferguson, or a combination of Ferguson to a system created by a combination of RFC 3209 and Raz, instead, the combination in question applies the teachings from Raz, and the teachings from Ferguson, into the system disclosed in RFC 3209. EX1003, ¶ 246.

50. Second, Dr. Rouskas asserts that a POSITA would not have understood that Raz's provisioning technique to subports within an LSR could be applied the same way between LSRs. EX2007, ¶¶ 152-153, 178. Dr. Rouskas asserts that Raz teaches away from applying its method between two network devices and doing so would have made the method inoperable or unsatisfactory for its intended purpose. EX2007, ¶ 153. Dr. Rouskas makes this assertion without explaining why and how Raz teaches away from applying its teachings regarding subports within an LSR to communications between an LSR, and why the application of these methods would be inoperable.

51. I disagree that the Raz technique disclosed and used for a LAG inside a network device cannot be applied to a LAG between network devices. Dr. Rouskas makes this assertion even as he acknowledges that Ferguson explicitly teaches a POSITA to apply the techniques used within a router, between routers. EX2007 ¶

157. However, Dr. Rouskas asserts that “Ferguson’s teaching about applying techniques within a router is limited to Ferguson’s own techniques and does not apply to Raz’s techniques.” *Id.* Again, I disagree. Raz teaches a method for selection and provisioning of a subport, selecting from an internal LAG, or aggregation of subports, the subport that will then be used to send and receive MPLS data packets as part of an LSP. EX1003, ¶¶ 85–86. This internal provisioning of one of an aggregation of subports located internally within an LSR, would readily be applied by a POSITA to the selection of one of an aggregation of links between LSRs, especially in view of Ferguson’s suggestion indicating that techniques used externally between switches can be utilized internally within LSRs. EX1007, 5:61 – 6:9. Ferguson notes that the techniques disclosed can be applied “in a variety of environments.” The techniques in question are techniques for the distribution of MPLS packets across physical data paths, load balancing the packets across an aggregated link. EX1007, Abstract. Similarly, Raz provides a mechanism for provisioning protection tunnels and allocating subport bandwidth to subports for MPLS packets. EX1006, Abstract. It would have been obvious for a POSITA to apply the technique in Raz for the distribution of MPLS packets within a network device to distribution of packets between network devices, having received guidance from Ferguson that techniques for the distribution of MPLS packets can be applied interchangeably within network devices and between network devices.

52. As I had previously asserted, Ferguson, like RFC 3209 and Raz focus on improving the use of bandwidth in an MPLS network. EX1003 ¶ 246. RFC 3209, Raz and Ferguson also discuss the use of MPLS labels for selecting path for communicating data packets. *Id.* Raz teaches the provisioning of a path through LSRs, choosing which subport among an aggregation of subports, the data packet will be sent and received. 1003 ¶¶ 244 - 245. Ferguson teaches using one of the physical data ports of a LAG for communicating a high-speed packet between routers. 1003 ¶¶ 249. It would have been obvious, and in fact elementary, to apply the provisioning method employed within an LSR of an MPLS/LAG switch to selecting and reserving physical data ports of LAGs between an MPLS/LAG switch.

53. Dr. Rouskas's analysis reaches an erroneous conclusion because he focuses on the aggregation of physical links or physical connections (such as the Ethernet cables). As I have explained, the link aggregation takes places at the MAC Layer above and separate from the Physical Layer, where the specific form of the physical links or the physical connections are defined. *Supra* §VI.A. For example, in smaller routers (such as those typically found in a small office environment) where the individual physical link may not have high link speed, it would nonetheless be beneficial to aggregate multiple MAC Layer ports that are connected to each of the physical links to form a higher speed LAG. This is the case in Ferguson, where multiple physical links are bundled together to become a LAG by the bundling of

their connected ports. But the fact that the LAG is a bundling of multiple physical links is nothing more than a coincidence by the fact that each physical link has exactly one port at the MAC layer in Ferguson. As I have explained abundantly based on the 802.3ad standard, it is the bundling of the ports that makes a LAG. On the other hand, when you have a single high-speed physical link (i.e., 10G Ethernet cable) as is the case in Raz, you have a many-to-one mapping of bundled subports to physical link (i.e., a 10 to 1 mapping of ten 1G subports bundled into a LAG whereby the data frames are transmitted over a 10G Ethernet cable). EX1006, 7:14-16. Both situations describe a LAG. Whether or not the ten 1G subports were individually transmitted through ten separate 1G Ethernet cables, the fact that the ten 1G subports were bound together to be a LAG does not change.

54. To the extent Dr. Rouskas attempts to explain why a POSITA would not have applied the techniques in Raz used for subports of a LAG within a network device to LAGs between network device, he focuses on how packets are processed and move within a switch, arguing that this is different from how packets are communicated between switches. EX2007, ¶¶178, 179. Dr. Rouskas argues that a POSITA would have understood that Raz's subports are not for communicating but internal processing. EX2007, ¶178. This is inaccurate. The subports in Raz, as the specification in Raz made it clear, are for communicating, and a POSITA would have reached this conclusion after reading Raz. A POSITA would understand the

subports, as disclosed by Raz, to be components of LAGs that are used for communicating or transporting MPLS packets. Raz states “[o]**utbound packets** are **assigned a subport** based on their MPLS labels. . . . **outputs of all the subports** are aggregated and combined into a single high-speed packet stream and transmitted over the link.” EX1006, 2:9-13. Raz further notes “[t]he invention also comprises a mechanism for provisioning **protection tunnels** wherein the **outbound subport** is always set equal to the **inbound subport**. In addition, the invention comprises an algorithm for allocating **subport bandwidth to LSPs**.” EX1006, 2:21-25. Raz further discloses that “the **assignment of traffic to the subports** is based on LSPs where the MPLS label is used to **assign** an **inbound packet** to one of the **subports**.” EX1006, 8:56-58. Raz further discloses “**outbound subport**” and “**subport for receiving**.” EX1006, 9:50-53. Raz specifically identifies that the purpose of the subports are as a communication channel, tunnel or path for traffic comprised of MPLS packets to be transported. *See supra* §VI.A.

55. Lastly, Dr. Rouskas argues that how packets are processed and move within a switch differ from how they are communicated between switches, asserting three technical reasons why packets are not communicated the same way within and between LSRs and why as a result a POSITA would not be motivated to modify RFC 3209 with Raz and Ferguson. EX2007, ¶¶ 178, 179. It has already been demonstrated above that the subports in Raz are used for data communication in

addition to data processing. Further, while Dr. Rouskas's analysis of the process executed for data packet is not technically inaccurate, these issues are irrelevant to my analysis. That the data packets may take a particular form when traveling through the subports to the Physical Layer does not affect the fact that the subports are aggregated into a LAG.

VIII. Conclusion

56. This declaration and my opinions herein are made to the best of my knowledge and understanding, and based on the material available to me, at the time of signing this declaration. I declare that all statements made herein on my own knowledge are true and that all statements made on information and belief are believed to be true, and further, that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under 18 U.S.C. § 1001.

Executed on this 28th day of May, 2024.

Respectfully submitted,

A handwritten signature in black ink, appearing to read "Paul Min", written over a horizontal line.

Paul Min, Ph.D.