

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

DELL TECHNOLOGIES INC.,
DELL INC.,
and
CISCO SYSTEMS, INC.,
Petitioners

v.

CORRIGENT CORPORATION,
Patent Owner.

Case No. IPR2023-00370

U.S. Patent No. 7,593,400

Petition for *Inter Partes* Review of U.S. Patent No. 7,593,400

Mail Stop **PATENT BOARD**
Patent Trial and Appeal Board
U.S. Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

TABLE OF CONTENTS

	<u>Page</u>
I. INTRODUCTION	1
II. MANDATORY NOTICES	2
A. Real Parties-in-Interest	2
B. Related Matters.....	2
C. Lead and Backup Counsel	3
D. Service Information	4
III. PAYMENT OF FEES	4
IV. REQUIREMENTS OF <i>INTER PARTES</i> REVIEW	5
A. Standing.....	5
B. Identification of Challenge and Relief Requested	5
1. Smith	5
2. Ishimori	6
3. Edsall.....	6
4. Zelig	6
5. IEEE 802.1Q Standard	6
C. How the Challenged Claims Are To Be Construed Under 37 C.F.R. § 42.104(b)(3)	7
D. How the Challenged Claims Are Unpatentable Under 37 C.F.R. § 42.104(b)(4)	8
E. Supporting Evidence Under 37 C.F.R. § 42.104(b)(5).....	8
V. THRESHOLD REQUIREMENT FOR <i>INTER PARTES</i> REVIEW	9
VI. PATENT OVERVIEW.....	9
A. Person of Ordinary Skill in the Art.....	9
B. Background in the Relevant Art	9
1. Computer Networks	10
2. Switches and Learning MAC Tables.....	10
3. Link Aggregation	11

C.	Summary of the Alleged Invention of the '400 Patent	11
D.	Summary of the Prosecution History of the '400 Patent	14
VII.	GROUND OF UNPATENTABILITY	15
A.	Ground 1: Claims 1, 3, 6, 11, 13, and 16 would have been obvious over Smith.....	15
1.	Overview of Ground 1.....	15
a)	Smith.....	15
2.	Analysis of Ground 1	17
a)	Claim 1[pre]: A method for communication, comprising:	17
b)	Claim 1[a]: configuring a network node having a plurality of ports, and at least first and second line cards with respective first and second ports, to operate as a distributed media access control (MAC) bridge in a Layer 2 data network;.....	17
c)	Claim 1[b]: configuring a link aggregation (LAG) group of parallel physical links between two endpoints in said Layer 2 data network joined together into a single logical link, said LAG group having a plurality of LAG ports and a plurality of conjoined Member Line Cards;.....	19
d)	Claim 1[c]: providing for each of said member line cards a respective forwarding database (FDB) to hold records associating MAC addresses with ports of said plurality of ports of said network node;.....	21
e)	Claim 1[d]: receiving a data packet on an ingress port of said network node from a MAC source address, said data packet specifying a MAC destination address on said Layer 2 data network;.....	22
f)	Claim 1[e]: conveying, by transmitting said data packet to said MAC destination address via said first port, said received data packet in said network node to at least said first line card for transmission to said MAC destination address;	23

- g) Claim 1[f]: if said MAC destination address does not appear in said FDB, flooding said data packet via one and only one LAG port of said plurality of LAG ports; 24
- h) Claim 1[g]: checking said MAC source address of the data packet against records in said FDB of said first line card; and 26
- i) Claim 1[h]: if said FDB of said first line card does not contain a record of an association of said MAC source address with said ingress port, creating a new record of said association, adding said new record to the FDB of said first line card, and sending a message of the association to each member line card and said plurality of member line cards. 26
- j) Claim 3: The method according to claim 1, and comprising receiving the message at the second line card, and responsively to the message, adding the record of the association to the FDB of the second line card if the record does not already exist in the FDB of the second line card. 27
- k) Claim 6: The method according to claim 1, wherein sending the message comprises transmitting a synchronization packet from the first line card via switching core of the network node to at least the second line card. 28
- l) Claim 11[pre] 11: A node for network communication, comprising: 29
- m) Claim 11[a]: a switching core; 29
- n) Claim 11[b]: a plurality of ports; 29
- o) Claim 11[c]: a plurality of member line cards conjoined in a link aggregation (LAG) group of parallel physical links between two endpoints in a Layer 2 data network joined together into a single logical link, having a plurality of LAG ports to forward packets through said switching core so that the node operates as a virtual media access

- control (MAC) bridge in said Layer 2 data network, said plurality of member line cards including at least first and second line cards, each line card having respective ports and having a respective forwarding database (FDB) to hold records associating MAC addresses with said respective ports of said line cards; 29
- p) Claim 11[d]: wherein said line cards are arranged so that upon receiving a data packet on an ingress line card from a MAC source address, said data packet specifying a MAC destination address, said ingress line card conveys said data packet via said switching core to at least said first line card for transmission to said MAC destination address, whereupon said first line card checks said MAC source address of said data packet against records in said FDB of said first line card, and if said FDB database of said first line card does not contain a record of an association of said MAC source address with said ingress port, adds said record to the FDB of said first line card and sends a message to at least said second line card informing said second line card of said association, and arranged, when said MAC destination address does not appear in said FDB, to flood said data packet via one and only one of said LAG ports. 30
- q) Claim 13: The node according to claim 11, wherein responsively to the message, the second line card adds the record of the association to the MAC database of the second line card if the record does not already exist in the FDB of the second line card..... 31
- r) Claim 16: The node according to claim 11, wherein the message comprises a synchronization packet, which is transmitted from the first line card via the switching core to at least the second line card..... 31

B.	Ground 2: Claims 1–3 and 11–13 would have been obvious over Smith in view of Ishimori.....	31
1.	Overview of Ground 2.....	31
a)	Ishimori.....	31
b)	Motivation to combine Smith and Ishimori	34
2.	Analysis of Ground 2	35
a)	Claim 1[c]	35
b)	Claim 1[f].....	36
c)	Claim 1[g].....	37
d)	Claim 1[h].....	37
e)	Claim 2: The method according to claim 1, wherein sending the message comprises sending messages periodically at predefined times to inform at least the second line card of new associations between the MAC address and the respective ports.	38
f)	Claim 3.....	39
g)	Claim 11[c]	40
h)	Claim 11[d].....	40
i)	Claim 12: The node according to claim 11, wherein at least the first line card is adapted to send messages periodically at predefined times to inform at least the second line card of new associations between the MAC addresses and the respective ports.	40
j)	Claim 13.....	40
C.	Ground 3: Claims 1, 4–7, 10–11, 14–17, and 20 would have been obvious over Smith in view of Ishimori in further view of Edsall.....	41
1.	Overview of Ground 3.....	41
a)	Edsall	41
b)	Motivation to Combine Smith and Ishimori with Edsall	42

2.	Analysis of Ground 3	45
a)	Claim 1[g].....	45
b)	Claim 1[h].....	45
c)	Claim 4: The method according to claim 3 and comprising marking the records in the respective FDB of each line card to distinguish a first type of the records, which are added in response to data packets transmitted via a port of the line card, from a second type of the records, which are added in response to messages received from another of the line cards.	46
d)	Claim 5[a]: The method according to claim 4, and comprising: associating a respective aging time with each of the records;	47
e)	Claim 5[b]: refreshing the records in the FDB responsively to further packets transmitted by the line cards; and	48
f)	Claim 5[c]: removing the records from the respective FDB if the records are not refreshed within the respective aging time.	49
g)	Claim 6.....	49
h)	Claim 7: The method according to claim 6, wherein sending the synchronization packet comprises, if the record in the FDB associates the MAC source address with a port different from the one of the ports on which the data packet was received, changing the record in the FDB of the first line card and sending a synchronization update packet to at least the second line card to indicate that the record has been changed.....	49
i)	Claim 10[a]: The method according to claim 1, and comprising: conveying a further data packet, received from a further MAC source address, to the second line card for transmission over the network;	51

j)	Claim 10[b]: checking the further MAC source address against the records in the FDB of the second line card; and.....	53
k)	Claim 10[c] responsively to the further data packet, adding a further record with respect to the MAC source address to the FDB of the second line card and sending a further message to inform at least the first line card of the further record.	54
l)	Claim 11[a]	55
m)	Claim 11[d].....	55
n)	Claim 14: The node according to claim 13, wherein the records in the respective FDB of each line card are marked to distinguish a first type of the records, which are added in response to data packets transmitted via a port of the line card, from a second type of the records, which are added in response to messages received from another of the line cards.	55
o)	Claim 15: The node according to claim 14, wherein a respective aging time is associated with each of the records, and wherein the line cards are operative to refresh the records in the FDB responsively to further packets transmitted by the line cards, and to remove the records from the respective FDB if the records are not refreshed within the respective aging time.	56
p)	Claim 16.....	56
q)	Claim 17: The node according to claim 16, wherein the line cards are operative so that if the record in the FDB associates the MAC source address with a port different from the one of the ports on which the data packet was received, the record in the FDB of the first line card is changed, and the synchronization packet comprises a synchronization update packet, which indicates to at least the second line card to indicate that the record has been changed.	56

r)	Claim 20: The node according to claim 11, wherein the line cards are adapted to forward a further data packet, received from a further MAC source address, to the second line card for transmission over the network, whereupon the second line card checks the further MAC source address against the records in the FDB of the second line card, and responsively to the further data packet, adds a further record with respect to the MAC source address to the FDB of the second line card and sends a further message to inform at least the first line card of the further record.	56
D.	Ground 4: Claims 8–9 and 18–19 would have been obvious over Smith in view of Ishimori in further view of Zelig	57
1.	Overview of Ground 4.....	57
a)	Zelig	57
b)	Motivation to combine Smith and Ishimori with Zelig	58
2.	Analysis of Ground 4	60
a)	Claim 8: The method according to claim 1, wherein the network node is configured to operate as multiple virtual MAC bridges in a Layer 2 virtual private network (VPN), wherein each virtual MAC bridge is configured to serve a respective VPN instance, and wherein the records associating the MAC addresses with the respective ports are maintained independently for each of the VPN instances.....	60
b)	Claim 9: The method according to claim 8, wherein the VPN instance is a VPLS instance among multiple VPLS instances served by the network node, and wherein sending the message comprises identifying the VPLS instance in the message so as to inform all the line cards that serve the VPLS instance.	60
c)	Claim 18: The node according to claim 11, wherein at least some of the line cards are	

	configured so that the node operates as multiple virtual MAC bridges in a Layer 2 virtual private network (VPN), wherein each virtual MAC bridge is configured to serve a respective VPN instance, and wherein the records associating the MAC addresses with the respective ports are maintained independently for each of the VPN instances.	61
d)	Claim 19: The node according to claim 18, wherein the VPN instance is a VPLS instance among multiple VPLS instances served by the network node, and wherein the VPLS instance is identified in the message so as to inform all the line cards that serve the VPLS instance of the association.....	62
E.	Ground 5: Claims 9 and 19 would have been obvious over Smith in view of Ishimori, Zelig, and 802.1Q.....	62
1.	Overview of Ground 5.....	62
a)	802.1Q.....	62
b)	Motivation to combine Smith, Ishimori, and Zelig with 802.1Q.....	62
2.	Analysis of Ground 5	63
a)	Claim 9.....	63
b)	Claim 19.....	63
VIII.	PTAB DISCRETION SHOULD NOT PRECLUDE INSTITUTION.....	63
A.	PTAB should not exercise its discretion to deny institution under <i>Fintiv</i>	63
1.	Factor 1: Institution will increase the likelihood of stay	63
2.	Factor 2: District Court schedule.....	64
3.	Factor 3: Petitioners’ investment in IPR outweighs forced investment in litigation to date	65
4.	Factor 4: The Petition raises unique issues.....	65
5.	Factor 5: Whether the Petitioners and Defendants in the parallel litigation are the same party	66
6.	Factor 6: Other circumstances support institution	66

B.	PTAB should not exercise its discretion to deny institution under <i>Becton</i> and <i>Advanced Bionics</i>	66
C.	Discretionary denial under <i>General Plastic</i> is not appropriate.....	67
IX.	CONCLUSION	68
	APPENDIX: CHALLENGED CLAIM LISTING.....	72

TABLE OF AUTHORITIES

CASES

<i>Abbot Vascular, Inc. v. Flexstent</i> , IPR2019-882, Paper 48 (PTAB Oct. 2, 2020)	46
<i>Advanced Bionics LLC v. MED-EL Elektromedizinische Gerate GmbH</i> , IPR2019-01469, Paper 6 (PTAB Feb. 13, 2020)	66
<i>In re Apple Inc.</i> , 979 F.3d 1332 (Fed. Cir. 2020)	64
<i>Becton, Dickinson & Co. v. B. Braun Melsungen AG</i> , IPR2017-01586, Paper 8 (PTAB Dec. 15, 2017).....	66
<i>Gen. Plastic Indus. Co., Ltd. v. Canon Kabushiki Kaisha</i> , IPR2016-01357, Paper 19, 16 (PTAB Sept. 6, 2016)	67
<i>Haas Automation, Inc. v. Olati LLC</i> , IPR2021-00146, Paper 11 (PTAB May 18, 2021).....	26, 32
<i>HP Inc. v. Slingshot Printing LLC</i> , IPR2020-01084, Paper 13 (PTAB Jan. 14, 2021).....	65, 66
<i>HTC Corp. v. Cellular Commc’ns Equip., LLC</i> , 877 F.3d 1361 (Fed. Cir. 2017)	8
<i>Intel Corp. v. Qualcomm Inc.</i> , 21 F.4th 801 (Fed. Cir. 2021)	26, 32
<i>Laird Techs., Inc. v. Garftech Int’l Holdings, Inc.</i> , IPR2014-24, Paper 46 (PTAB Mar. 25, 2015)	46
<i>Oticon Med. AB v. Cochlear Ltd.</i> , IPR2019-00975 Paper 15, 20 (PTAB Oct. 16, 2019)	67
<i>PEAG LLC v. Varta Microbattery GMBH</i> , IPR2020-01214, Paper 8, 17 (PTAB Jan. 6, 2021).....	64
<i>Sand Revolution II, LLC v. Continental Intermodal Grp.</i> , IPR2020-01393, Paper 24 (PTAB June 16, 2020).....	64

<i>Verizon v. Huawei,</i> IPR2020-01079, Paper 10 (PTAB Jan. 14, 2021).....	67
--	----

<i>VMware Inc. v. Intellectual Ventures II LLC,</i> IPR2020-00859, Paper No. 13 (PTAB Nov. 5, 2020)	63
--	----

REGULATIONS

37 C.F.R. § 42.100(b) (November 13, 2018).....	7
--	---

83 F.R. 51358 (2018).....	7
---------------------------	---

PETITIONERS' EXHIBIT LIST

Exhibit No.	Description
1001	U.S. Patent No. 7,593,400 (“the ’400 patent”)
1002	Copy of Prosecution History of the ’400 patent
1003	Declaration of Dr. Nicholas Bambos
1004	U.S. Patent Application Publication No. 2005/0198371 (“Smith”)
1005	Certified English Translation of Japanese Patent Application No. 2005086668 (“Ishimori”)
1006	U.S. Patent No. 6,735,198 (“Edsall”)
1007	U.S. Patent Application Publication No. 2004/0133619 (“Zelig”)
1008	IEEE Standard 802.1Q, Virtual Bridged Local Area Networks, 1998 Edition (“802.1Q”)
1009	Declaration of Dr. Mary K. Bolin
1010	<i>Corrigent Corp. v. Dell Technologies et al.</i> , No. 1:22-cv-00496 (D. Del.), Dkt. 1
1011	<i>Corrigent Corp. v. Arista Networks, Inc.</i> , No. 1:22-cv-00497 (D. Del.), Dkt. 1
1012	<i>Corrigent Corp. v. Cisco Systems, Inc.</i> , No. 6:22-cv-00396 (W.D. Tex.), Dkt. 1
1013	<i>Corrigent Corp. v. Cisco Systems, Inc.</i> , No. 6:22-cv-00396 (W.D. Tex.), Dkt. 42
1014	<i>Corrigent Corp. v. Cisco Systems, Inc.</i> , No. 6:22-cv-00396 (W.D. Tex.), Dkt. 46
1015	IEEE 802.1D, Media Access Control (MAC) Bridges, 2004 Edition
1016	IEEE Standard 802.1Q, Virtual Bridged Local Area Networks, 2005 Edition
1017	IEEE 802.3 Standard for Local and metropolitan area networks: Specific requirements, 2002 Edition

Exhibit No.	Description
1018	Kompella et al., <i>Virtual Private LAN Service</i> , IETF (December 2005)
1019	Lasserre et al., <i>Virtual Private LAN Services over MPLS</i> , IETF (November 2005)
1020	Martini et al., <i>Encapsulation Methods for Transport of Ethernet Over MPLS Networks</i> , IETF (November 2005)
1021	U.S. Patent No. 6,917,986
1022	U.S. Patent No. 7,974,223
1023	Western District of Texas Statistics (Docket Navigator)
1024	District of Delaware Statistics (Docket Navigator)
1025	Japanese Patent Application No. 2005086668 (“Ishimori”)

I. INTRODUCTION

Dell Technologies Inc., Dell Inc., and Cisco Systems, Inc. hereby petition for *inter partes* review of U.S. Patent No. 7,593,400 and seek cancellation of claims 1–20 (“the Challenged Claims”) because they are unpatentable under 35 U.S.C. § 103.

The ’400 patent relates to communication networks, specifically synchronizing learned MAC addresses to bridge ports in a distributed or virtual bridge in a network. The patent relies on a known technique by which a group of parallel physical links between two endpoints in a network are joined together in what is called a link aggregation (LAG) group. *Id.*, 2:37–40. The LAG group is seen by the user as a single link. *Id.*, 2:31–33. The disclosed bridges also use a known concept called a MAC forwarding database (FDB) to store associations of MAC addresses and the connected ports. *Id.*, 3:7–11. When a line card—which receives and forwards packets within a bridge—receives a packet with a destination MAC address that appears in the FDB, it forwards the packet to the port associated with the MAC address in the FDB. *Id.* On the other hand, if the destination MAC address does not appear in the FDB, the line card floods the packet to all the ports. *Id.* This was all admittedly known in the art. *Id.*, 1:33–2:56.

The patent’s alleged improvement is a synchronization method that purportedly avoids the need for excessive flooding when LAGs are used. According to the patent, when a data packet arrives at a line card, if the port of the line card on

which the packet was received belongs to a LAG group, the line card sends a synchronization message to the other line cards in the same LAG group. In this way, the patent suggests that other line cards in a LAG group can learn MAC addresses before they have received packets from the MAC address in question. *Id.*, 3:46–53.

However, this alleged advance is not new. To the contrary, both the problem and solution were disclosed in the prior art, as discussed herein. For these reasons, Petitioners respectfully request the Board institute IPR proceedings and cancel the Challenged Claims.

II. MANDATORY NOTICES

A. Real Parties-in-Interest

The real parties-in-interest are Dell Technologies Inc., Dell Inc., and Cisco Systems, Inc. (collectively, “Petitioners”). No unnamed entity is funding, controlling, or directing this Petition, or otherwise has had an opportunity to control or direct this Petition or Petitioners’ participation in any resulting IPR.

B. Related Matters

The ’400 patent is the subject of this IPR Petition is also the subject of patent litigation suits brought by Patent Owner Corrigent Corporation, including against Petitioners:

- *Corrigent Corp. v. Dell Technologies et al.*, No. 1:22-cv-00496 (D. Del.) (EX1010);

- *Corrigent Corp. v. Arista Networks, Inc.*, No. 1:22-cv-00497 (D. Del.) (EX1011); and
- *Corrigent Corp. v. Cisco Systems, Inc.*, No. 6:22-cv-00396 (W.D. Tex.) (EX1012).

C. Lead and Backup Counsel

Pursuant to 37 C.F.R. §§ 42.8(b)(3), 42.8(b)(4), and 42.10(a), Petitioners provide the following designation of counsel:

Lead Counsel	Back-up Counsel
Stuart Rosenberg (Reg. No. 60,772) Gibson, Dunn & Crutcher LLP 1881 Page Mill Road Palo Alto, CA 94304-1211 Tel: 650-849-5389 SRosenberg@gibsondunn.com	<u>First Back-Up Counsel</u> Brian Rosenthal (<i>pro hac vice</i> motion requested) Gibson, Dunn & Crutcher LLP 200 Park Avenue New York, NY 10166-0193 Tel: 212-351-2339 BARosenthal@gibsondunn.com <u>Additional Backup Counsel</u> Ryan Iwahashi (Reg. No. 63,378) Gibson, Dunn & Crutcher LLP 1881 Page Mill Road Palo Alto, CA 94304-1211 Tel: 650-849-5367 RIwahashi@gibsondunn.com

Petitioners respectfully request authorization to file a motion for Brian Rosenthal to appear before the USPTO *pro hac vice*. Mr. Rosenthal is an

experienced litigating attorney and is currently serving as lead counsel for Petitioners in related matters *Corrigent Corp. v. Dell Technologies et al.*, No. 1:22-cv-00496 (D. Del.) and *Corrigent Corp. v. Cisco Systems, Inc.*, No. 6:22-cv-00396 (W.D. Tex.). Mr. Rosenthal has established familiarity with the subject matter at issue in this proceeding. Petitioners intend to file a motion to appear *pro hac vice* under 37 C.F.R. § 42.10. Pursuant to 37 C.F.R. § 42.10(b), powers of attorney accompany this Petition.

D. Service Information

Service via hand delivery or postal mail may be made at the addresses of the lead and back-up counsel above. Petitioners hereby consent to electronic service, and service via electronic mail may be made at the email addresses provided above for the lead and back-up counsel.

III. PAYMENT OF FEES

Pursuant to 37 C.F.R. §§ 42.103 and 42.15(a), the required fee is being submitted herewith. The Office is authorized to charge any fee deficiency, or credit overpayment, to deposit account no. 50-1408. Any additional fees due in connection with this Petition may be charged to the foregoing account.

IV. REQUIREMENTS OF *INTER PARTES* REVIEW

A. Standing

Pursuant to 37 C.F.R. § 42.104(a), Petitioners certify that the '400 patent is available for IPR and that no Petitioner is barred from requesting an IPR on the grounds identified in this Petition. Specifically, Petitioners certify that: Petitioners have not filed a civil action challenging the validity of the '400 patent; this petition is filed not more than one year from April 19, 2022, the date on which the Petitioners were served with the complaint alleging infringement of the '400 patent; the estoppel provisions of 35 U.S.C. § 315(e)(1) do not prohibit this IPR; and this petition is filed after the later of (a) the date that is nine months after the date of the grant of the '400 patent or (b) the termination of any post-grant review of the '400 patent.

B. Identification of Challenge and Relief Requested

Pursuant to 37 C.F.R. § 42.104(b), Petitioners request the Board institute IPR of claims 1–20 of the '400 patent under pre-AIA 35 U.S.C. § 103 on the prior art references and grounds described below:

1. Smith

U.S. Patent Application Publication No. 2005/0198371 (“Smith”) (EX1004) qualifies as prior art under 35 U.S.C. § 102(a) and/or (e). Smith was filed on February 19, 2004 and published on September 8, 2005.

2. Ishimori

Japanese Patent Application No. 2005086668 (“Ishimori”) (EX1005, EX1025) qualifies as prior art under 35 U.S.C. § 102(a) and/or (b). Ishimori was filed on September 10, 2003 and published on March 31, 2005.

3. Edsall

U.S. Patent No. 6,735,198 (“Edsall”) (EX1006) qualifies as prior art under 35 U.S.C. § 102(a), (b), and/or (e). Edsall was filed on December 21, 1999 and issued on May 11, 2004.

4. Zelig

U.S. Patent Application Publication No. 2004/0133619 (“Zelig”) (EX1007) qualifies as prior art under 35 U.S.C. § 102(a), (b), and/or (e). Zelig was filed on January 7, 2003 and published on July 8, 2004.

5. IEEE 802.1Q Standard

The IEEE 802.1Q-1998 Standard (“802.1Q”) (EX1008) qualifies as prior art under 35 U.S.C. § 102(b). Specifically, 802.1Q was published on March 8, 1999 and was publicly available no later than December 20, 2001. EX1009 ¶¶17–28.

Smith, Ishimori, Zelig, and 802.1Q were not cited during the prosecution of the ’400 patent. Edsall was cited during prosecution; however, Petitioners rely on Edsall only for disclosures of claim limitations that were found to be present in Edsall during prosecution—findings that the applicant did not challenge.

In this IPR, Petitioners apply the above references and assert the following grounds of rejection under 35 U.S.C. § 103:

Ground	Claims	Basis for Rejection ¹
1	1, 3, 6, 11, 13, 16	Obvious over Smith
2	1–3, 11–13	Obvious over Smith in view of Ishimori
3	1, 4–7, 10–11, 14–17, 20	Obvious over Smith in view of Ishimori in further view of Edsall
4	8–9, 18–19	Obvious over Smith in view of Ishimori in further view of Zelig
5	9, 19	Obvious over Smith in view of Ishimori in further view of Zelig and 802.1Q

C. How the Challenged Claims Are To Be Construed Under 37 C.F.R. § 42.104(b)(3)

A claim subject to *inter partes* review “shall be construed using the same claim construction standard that would be used to construe a claim in a civil action under 35 U.S.C. § 282(b), including construing the claim in accordance with the ordinary and customary meaning of such claim as understood by one of ordinary skill in the art and the prosecution history pertaining to the patent.” 37 C.F.R. § 42.100(b) (November 13, 2018); 83 F.R. 51358 (2018). Petitioners do not believe

¹ All obviousness grounds include knowledge of a POSITA.

(Cont’d on next page)

any constructions have an impact on the invalidity analyses set forth herein.² See *HTC Corp. v. Cellular Commc'ns Equip., LLC*, 877 F.3d 1361, 1367–68 (Fed. Cir. 2017) (affirming the Board absent an express construction of a term).

D. How the Challenged Claims Are Unpatentable Under 37 C.F.R. § 42.104(b)(4)

The following sections explain how the Challenged Claims are unpatentable under the statutory grounds identified above, including where each element of the claim is found in the prior art patents or printed publications.

E. Supporting Evidence Under 37 C.F.R. § 42.104(b)(5)

The exhibit numbers of the supporting evidence relied upon and the relevance of the evidence to the Challenged Claims, including an identification of specific portions of the evidence that support the challenge, are provided below. The technical information and grounds for unpatentability are further supported by the Declaration of Dr. Nicholas Bambos (EX1003). A List of Exhibits is included in this paper pursuant to 37 C.F.R. § 42.63(e).

² Claim construction has begun in the *Cisco* litigation. Petitioners will request leave to submit the district court's claim construction order when available, so it is timely made of record and can be considered by the Board. 37 C.F.R. §42.100(b).

V. THRESHOLD REQUIREMENT FOR *INTER PARTES* REVIEW

Under 35 U.S.C. § 314(a), institution of *inter partes* review requires “a reasonable likelihood that the petitioners would prevail with respect to at least one of the claims challenged in the petition.” This petition meets this threshold for each ground of unpatentability.

VI. PATENT OVERVIEW

A. Person of Ordinary Skill in the Art

The alleged invention relates to communication networks, specifically learning MAC addresses in a distributed or virtual bridge. EX1001, Abstract. A person of ordinary skill in the art (“POSITA”) at the time of the alleged invention (May 2006) would have had a degree in electrical engineering or a similar discipline, with at least two years of relevant industry or research experience (or additional education). EX1003 ¶17. The relevant experience would include a working understanding of networking systems, including distributed bridges, MAC forwarding tables, LAG groups, and virtual LAN services. *Id.* Lack of professional experience can be remedied by additional education, and vice versa. *Id.*

B. Background in the Relevant Art

All concepts in this section were well known and widely used by POSITAs at least as of May 19, 2006. EX1003 ¶¶36–53.

1. Computer Networks

Computer networks are comprised of communications links that connect to ports of communication nodes. *Id.* ¶37. Data packets or “frames” transmit information between the nodes based on MAC (Media Access Control) addresses, which are unique identifiers assigned to network hardware. *Id.* ¶¶37,44–45. An L2 switching node maintains a MAC address table, which maps the output port for a particular frame based on its destination MAC address. *Id.* ¶¶47–50.

2. Switches and Learning MAC Tables

MAC addresses are used by L2 switches (also called bridges) to switch packets at the network node from an input port to an output port. *Id.* ¶¶44–46. A MAC bridge is a network device connects two physically separated LANs into a single logical LAN. *Id.* ¶44. To forward data packets, called “frames,” a MAC bridge maintains a forwarding database (FDB) of MAC addresses and corresponding bridge ports of packets that have been received at the bridge. *Id.* ¶¶45–47. When a bridge receives a packet, it checks the destination MAC address of the packet. The bridge checks a MAC address table (or forwarding database) to determine whether it knows where to forward the message. *Id.* ¶47. If the destination MAC address of the packet is in the FDB, the bridge forwards the packet to the port associated with that destination MAC address. *Id.* If not, the bridge may send the message over every port to ensure it is received at the intended destination. *Id.* This is called

“flooding.” *Id.* When packets are received, the bridge also “learns” information about the *source* MAC address of each packet and, if the source MAC address is not in the FDB, updating the database to reflect the association between the incoming port and the source MAC address. *Id.*


3. Link Aggregation

In the 1990s, link aggregation was used to increase bandwidth, traffic load balancing, scalability, and reliability by grouping multiple links from a switch to form a link aggregate group, or “LAG.” *Id.* ¶51. These multiple links in a LAG are treated as a single logical link. *Id.* The Layer 2 Link Aggregation Sublayer combines a number of individual physical links into a single logical link, which presents a single MAC interface to the MAC client. *Id.* (citing EX1017, Section 3, Fig. 43-1).

C. Summary of the Alleged Invention of the '400 Patent

The '400 patent relates to communication networks, specifically learning MAC addresses in a distributed bridge. The patent acknowledges it was already well known that computing systems connect to LANs at OSI Layer 2. EX1001, 1:13–19. As the patent acknowledges, “a number of authors have described methods for creating a virtual private LAN service (VPLS), which links different LANs together over an IP network.” *Id.*, 2:9–17. The patent explains a single provider can participate in multiple VPLSs. *Id.*, 2:18–30. Treating both virtual and physical

interfaces identically gives the illusion, from the user's perspective, that the provider network is a single LAN domain. *Id.*, 2:31–33. This VLAN functionality, which was known in the art, allows end points to be joined using link aggregation (LAG) as if they were part of the same local network even if they are not physically close to each other. *Id.*, 2:33–40 (citing EX1017). For example, as shown in Fig. 1, although users of **terminal 22 and 24** are not on the same LAN domain (terminal 22 is connected to a LAN 26 and terminal 24 is connected to a wide area network (WAN) 28), the VPLS “permits the users” to “communicate with one another as though they were connected to the same LAN domain.” *Id.*, 5:51–62; Fig. 1.



A **network node 30** “comprises multiple **line cards 32**, linked by a switching core 34,” and each line card has **ports 36**, which connect to other nodes in LAN 26 and WAN 28. *Id.*, 6:8–12. The line cards may serve as both ingress and egress with a MAC forwarding database (FDB), which stores learned MAC addresses and which

ports they were learned on. *Id.*, 3:7–11. When an ingress line card receives an incoming packet, it consults the FDB for the MAC destination address of the packet to determine which line card and port to forward the packet to. *Id.*, 3:11–17. When the MAC destination address does not appear in the FDB, the line card floods the packet to all of the ports, as in the prior art. *Id.* According to the patent, however, there was a need for better MAC address learning at the bridges within a VLAN to prevent the need to flood (or broadcast) messages when the destination MAC address is unknown. *Id.*, 3:34–53.

To address this, the purported invention provided “improved methods for MAC learning and network nodes that implement such methods.” *Id.*, 2:60–62. It describes a synchronization method that allegedly avoids the need for constant flooding. *Id.*, 3:41–45. When the transmitting line card sends the data packet via a port that belongs to a LAG group, it sends a synchronization message so the other line cards in the same LAG group can learn the MAC address association even when those other line cards have not yet received packets from that MAC address. *Id.*, 3:46–53. Therefore, upon receiving an incoming packet, “[i]f the record is found, the packet processor adds a tag to the packet indicating the egress port through which the packet should be forwarded.” *Id.*, 7:31–34. If there is “no corresponding record in the database however, it tags the packet for flooding.” *Id.*, 7:42–44. While the “switching core 34 will pass the packet for transmission via all ports” that “are used

by this VPLS instance,” “[f]or each LAG group serving the VPLS instance, however, the flooded packet is transmitted ***via only one port in the group.***” *Id.*, 7:47–49 (emphasis added).

Additionally, these synchronization messages (“SYNC”) are sent at regular intervals “to report each SELF entry [*e.g.*, entries that are learned by the packet processor on the line card itself] that it has created in the FDB 58 to the other line cards 32 in node 30.” *Id.*, 8:17–22. The patent thus discloses a record is removed from the database “if a predetermined aging time elapses following the timestamp without a further packet having been received with the same key” and “free[s] up space for new records.” *Id.*, 9:4–11. If, on the other hand, the current packet matches a record in the FDB, “the packet processor refreshes the timestamp of the record,” and “forwards the packet to the appropriate output port.” *Id.*, 9:29–31.

D. Summary of the Prosecution History of the ’400 Patent

The ’400 patent’s application was filed on May 19, 2006. EX1001, [22]. The Examiner rejected all pending claims as anticipated by Edsall, or obvious over Edsall and U.S. Patent No. 6,788,681 (“Hurren”). EX1002, 81–85. During an Examiner interview, it was discussed if the limitation “link aggregation group having a plurality of ports” in claim 1 were amended to recite the claim term as defined by the specification at 2:37–40, then the amended claim would overcome Edsall. *Id.*, 112. The applicant therefore amended the claim limitation to add “a link aggregation

(LAG) group of parallel physical links between two endpoints in said Layer 2 data network joined together into a single logical link, said LAG group having a plurality of LAG ports and a plurality of conjoined member line cards.” *Id.*, 114. The applicant made no other amendments to overcome the rejections over Edsall. The Examiner allowed the claims, and the patent issued on September 22, 2009. *Id.*, 128, 160.

VII. GROUNDS OF UNPATENTABILITY

A. Ground 1: Claims 1, 3, 6, 11, 13, and 16 would have been obvious over Smith.

Smith renders obvious claims 1, 3, 6, 11, 13, and 16.

1. Overview of Ground 1

a) Smith

Smith is directed to a virtual network device comprising interface bundles, which are managed as a single logical interface. The virtual network devices contain multiple “sub-units,” which “collectively operate as a single logical network device.” EX1004, Abstract, ¶34. To achieve this, the “system includes a virtual link bundle” with “several communications links,” and the “communications links are configured to be managed as a single link.” *Id.* ¶9.

As depicted below, **virtual network device 202** is coupled to other network devices 120(1)–120(3). *Id.* ¶44. The virtual network device consists of virtual network device sub-units 122(1) and 122(2), which includes several **line cards**

304(1)–304(4). *Id.* ¶46. The line cards include a “forwarding engine” and “interfaces.” *Id.* ¶¶46–47. When a packet is received at an uplink interface, the virtual network device sub-unit can learn “the sending device’s [MAC] address” by “associating the MAC address with the logical identifier of the uplink interface.” *Id.* ¶54. The sub-unit then informs each forwarding engine of this association. *Id.* “[P]ackets addressed to that MAC address will be sent from an uplink interface having the associated logical identifier.” *Id.*




FIG. 3

Smith uses “MAC notification frames” in order to “keep the content of the L2 tables in the virtual network device sub-unit 122(1) synchronized with the content

of the L2 tables in virtual network device sub-unit 122(2) and vice versa.” *Id.* ¶62. If a “forwarding table already includes an entry associating the destination address with a port of one of the network devices,” that forwarding engine will generate “a MAC notification identifying this association” to “any other forwarding engines within” the virtual network device sub-unit. *Id.* ¶63. “If there is no hit in the forwarding table,” “the virtual network device sub-unit floods the packet to all egress ports and/or uplink interfaces in the incoming VLAN” excluding the “interface that the packet arrived on.” *Id.* ¶66. As a result of the virtual link bundling, the data packet is sent via only “one of the communication links.” *Id.* ¶9.

2. Analysis of Ground 1

a) Claim 1[pre]: A method for communication, comprising:

To the extent the preamble is limiting, Smith discloses claim 1[pre]. Smith’s “system includes a virtual link bundle, which includes several communication links,” which constitute methods for communication. *Id.* ¶9; EX1003 ¶98.

b) Claim 1[a]: configuring a network node having a plurality of ports, and at least first and second line cards with respective first and second ports, to operate as a distributed media access control (MAC) bridge in a Layer 2 data network;

Smith discloses claim 1[a]. Smith discloses a **virtual network device 202**, which is a network node. EX1004 ¶36. The virtual network device “includes **several cards**” such as 304(1) and 304(4) (*e.g.*, first and second line cards). *Id.* ¶46.

These line cards include **interfaces**, for example, 320(5), 320(7), and 320(9) on line card 304(1), and 320(12), 320(14), and 320(16) on line card 304(4) (e.g., a plurality of ports). *Id.* ¶47. Thus, Smith’s first and second line cards have respective first and second ports. EX1003 ¶99.




FIG. 3

An “interface” and “port” are synonymous and Smith treats them interchangeably, and thus, discloses the first and second ports. EX1004 ¶63 (stating “port or uplink interface”). Smith teaches the virtual link bundles, which are a link aggregation (LAG) group of parallel physical links, may provide Layer 2 forwarding, and are “managed as a single link.” *Id.* ¶¶6,9,30. The virtual network device 202 “route[s] and forward[s] packets to and from network devices 120(1)–

120(3)” by associating the MAC address of a received data packet with the logical identifier of the uplink interface. *Id.* ¶54. Smith’s virtual network device 202 therefore operates as a distributed MAC bridge in a Layer 2 data network. EX1003 ¶¶100–102.

- c) **Claim 1[b]: configuring a link aggregation (LAG) group of parallel physical links between two endpoints in said Layer 2 data network joined together into a single logical link, said LAG group having a plurality of LAG ports and a plurality of conjoined Member Line Cards;**

Smith discloses claim 1[b]. As stated above, Smith’s virtual link bundles, which are a LAG group of parallel physical links, may provide Layer 2 forwarding, and are “managed as a single link.” EX1004 ¶¶6,9,30. Smith discloses **network device 120(2)** (*e.g.*, one endpoints) “is coupled to **virtual network device 202**” (*e.g.*, another endpoint) “by **virtual link bundle 250(2)**” (*e.g.*, a LAG group) as shown in the annotated figure below. *Id.* ¶44. The virtual link bundle 250(2) consists of **two uplinks** (*e.g.*, a plurality of LAG ports). *Id.* ¶51.




FIG. 3

Smith further specifies the virtual network devices provide Layer 2 “forwarding and routing.” *Id.* ¶30. Because virtual network device sub-units 122(1) and 122(2) “can coordinate their behavior such that they appear to be a single virtual device,” the remaining line cards 304(2) and 304(3) comprise the conjoined member line cards. Smith therefore discloses configuring a link aggregation (LAG) group of parallel physical links (*e.g.*, virtual link bundle 250(2)) between two endpoints (*e.g.*, virtual network device 202 and network device 120(2)) in said Layer 2 data network joined together into a single logical link, said LAG group having a plurality of LAG

ports (*e.g.*, interfaces 320(8) and 320(11)) and a plurality of conjoined Member Line Cards (*e.g.*, line cards 304(2) and 304(3)). EX1003 ¶¶103–105.

- d) **Claim 1[c]: providing for each of said member line cards a respective forwarding database (FDB) to hold records associating MAC addresses with ports of said plurality of ports of said network node;**

Smith discloses claim 1[c]. The member line cards in Smith include a **forwarding engine** and **interfaces** (*e.g.*, ports). *Id.* ¶47.




FIG. 3

When a packet is received on a particular uplink interface, the virtual network device learns the sending device's MAC address by "associating the MAC address with the logical identifier of [the] uplink interface" (*e.g.*, associating MAC addresses

with ports of said plurality of ports of said network node). *Id.* ¶54. The “forwarding engines” use the association between a packet and a particular logical identifier to route and forward packets to the network devices. *Id.* ¶54. Smith teaches this information may be used to “set up or modify lookup tables” (*e.g.*, the forwarding database (FDB)). *Id.* ¶¶57,61. Thus, Smith teaches the lookup table (*e.g.*, FDB) may store records associating MAC addresses with the logical identifier of the uplink interface (*e.g.*, ports of said plurality of ports) of said network node. EX1003 ¶¶106–107.

e) Claim 1[d]: receiving a data packet on an ingress port of said network node from a MAC source address, said data packet specifying a MAC destination address on said Layer 2 data network;

Smith discloses claim 1[d]. Smith discloses the uplink interface (*e.g.*, ingress port) receives a data packet with the “sending device’s MAC address” (*e.g.*, a MAC source address). EX1004 ¶54. Based on Figure 3, one such ingress port from a data packet received from network device 120(1) (*e.g.*, MAC source address) would be interface 320(9).




FIG. 3

Smith therefore teaches receiving a data packet on an ingress port (*e.g.*, interface 320(9)) of said network node (*e.g.*, virtual network device 202) from a MAC source address (*e.g.*, network device 120(1)). Smith further discloses this data packet has a “destination logical identifier,” such as the MAC destination address of network device 120(3). *Id.* ¶¶60,62. Smith specifies these devices provide Layer 2 forwarding. *Id.* ¶30. Thus, Smith teaches this limitation. EX1003 ¶¶108–111.

- f) **Claim 1[e]: conveying, by transmitting said data packet to said MAC destination address via said first port, said received data packet in said network node to at least said first line card for transmission to said MAC destination address;**

Smith discloses claim 1[e]. EX1003 ¶¶112–113. Smith teaches the data packet is forwarded based on “the association between a packet and a particular logical identifier.” EX1004 ¶54. Smith additionally teaches its system “favor[s] local interfaces.” *Id.* ¶56. Again, referring to Figure 3, in order to transmit the data packet to the MAC destination address (*e.g.*, network device 120(3)), the data packet would be conveyed from interface 320(9) to interface 320(7) (*e.g.*, said first port), because the system favors local interfaces. EX1003 ¶112. Interface 320(7), which is on line card 304(1) (*e.g.*, said first line card), would then transmit the data packet to the MAC destination address (*e.g.*, network device 120(3)).




FIG. 3

- g) **Claim 1[f]: if said MAC destination address does not appear in said FDB, flooding said data packet via one**

**and only one LAG port of said plurality of LAG
ports;**

Smith discloses claim 1[f]. EX1003 ¶¶114–116. Smith teaches the virtual network device sub-unit 122(1) looks up the destination address in a “lookup table” (*e.g.*, FDB). EX1004 ¶61. Claim 1[f] does not recite which forwarding database of the plurality of member line cards is checked.³ The only FDBs previously recited in the claim are associated with the member line cards and the LAG ports. Smith discloses sharing entries of the FDB with all its member line cards, and therefore, it does not matter which FDB is checked. *Id.* ¶63. A POSITA would have understood that any one of Smith’s FDBs could be checked, including the FDB associated with the member line cards. EX1003 ¶114. Another interpretation is the FDB that is checked is on the line card associated with the ingress port, which would be line card 304(1), continuing with the above example. *Id.* Although there is no line card or FDB associated with the ingress port recited in the claim, Smith would render

³ Petitioners intend to argue in the parallel litigation “said FDB” is indefinite for lacking a clear antecedent basis. However, the claims are invalid as obvious under any interpretation. Regardless, the question of indefiniteness is not before the Board and is no barrier to the Board’s inquiry. *Haas Automation, Inc. v. Olati LLC*, IPR2021-00146, Paper 11, 11–13 (P.T.A.B. May 18, 2021).

obvious this limitation under this interpretation as well. *Id.* Thus, under any possible interpretation of “said FDB,” Smith renders obvious this limitation. *Id.*

If there is no hit in the forwarding table, “the virtual network device sub-unit floods the packet to all egress ports and/or uplink interfaces in the incoming VLAN.” EX1004 ¶¶63,66. Smith expressly discloses or at least renders obvious that “flooding” is “via one and only one LAG port.” In particular, Smith teaches when the network device sends a packet “via the virtual link bundle,” it “selects one of the communication links on which to send the packet.” *Id.* ¶9; EX1003 ¶116.

h) Claim 1[g]: checking said MAC source address of the data packet against records in said FDB of said first line card; and

Smith discloses claim 1[g]. Smith teaches the virtual network device sub-unit learns the source identifier of the sending device (*e.g.*, MAC source address). EX1004 ¶65. These identifiers are stored in a lookup table (*e.g.*, FDB) on a virtual network device. *Id.* ¶61. It would have therefore been obvious that the MAC source address (*e.g.*, network device 120(1)) would have had to be checked in the records of the FDB of said first line card (*e.g.*, line card 304(1)) in order to determine whether it existed in the FDB. EX1003 ¶¶117–118.

i) Claim 1[h]: if said FDB of said first line card does not contain a record of an association of said MAC source address with said ingress port, creating a new record of said association, adding said new record to the FDB of said first line card, and sending a message of the

association to each member line card and said plurality of member line cards.

Smith discloses claim 1[h]. EX1003 ¶¶119–120. Smith teaches the virtual network device sub-unit learns the source identifier of the sending device (*e.g.*, MAC source address). EX1004 ¶65. Smith then teaches the network device sub-unit 122(2) sends a MAC notification (*e.g.*, a message) to update the forwarding engines (*e.g.*, sending a message of the association to each member line card) when it learns of a new association. *Id.* ¶63. “After being updated based on the MAC notification,” the forwarding engines “now know the location of the device identified by the destination address.” *Id.* A POSITA would have understood the MAC notification (*e.g.*, message) is therefore sent to each member line card. EX1003 ¶120. A POSITA would have also understood knowing the association between the MAC source address and ingress port requires adding an entry in the lookup table of, for example, line card 304(1) (*e.g.*, said first line card). *Id.*

- j) Claim 3: The method according to claim 1, and comprising receiving the message at the second line card, and responsively to the message, adding the record of the association to the FDB of the second line card if the record does not already exist in the FDB of the second line card.**

Smith discloses claim 3. Smith discloses the network device sub-unit 122(2) sends a MAC notification (*e.g.*, the message) to update the forwarding engines. EX1004 ¶63. Smith specifies “[a]fter being updated based on the MAC notification,

the forwarding engines in virtual network device sub-unit 122(1) now know the location of the device identified by the destination address.” *Id.* A POSITA would understand this includes the other line cards in virtual network device 202, such as line card 304(4). EX1003 ¶122. For example, line card 304(4) (*e.g.*, the second line card) would receive the MAC notification (*e.g.*, the message) and update its lookup tables. EX1004 ¶57. Thus, Smith discloses in response to the message, the record of association is added to the FDB of the second line card if the record does not already exist in the FDB of the second line card. EX1003 ¶¶121–122.




FIG. 3

- k) **Claim 6: The method according to claim 1, wherein sending the message comprises transmitting a synchronization packet from the first line card via**

switching core of the network node to at least the second line card.

Smith discloses claim 6. The MAC notification frames are used to keep the content of the L2 tables synchronized. *Id.* ¶62. Thus, it would have been obvious to a POSITA to implement the MAC notification frame as a synchronization packet and transmit it to at least the second line card via the switching core. EX1003 ¶¶123–125.

l) Claim 11[pre] 11: A node for network communication, comprising:

To the extent the preamble is limiting, Smith discloses claim 11[pre]. Smith’s system contains communications links that are coupled to network devices, which is a node. EX1004 ¶9; EX1003 ¶126.

m) Claim 11[a]: a switching core;

Smith discloses claim 11[a]. A POSITA would have understood that line cards need a way to communicate with each other. EX1003 ¶127. It would have been obvious to POSITA to do so via a switching core. *Id.*

n) Claim 11[b]: a plurality of ports;

Smith discloses claim 11[b]. *Id.* ¶128. The network devices in Smith contain several ports (*e.g.*, a plurality of ports). EX1004 ¶6 (“EtherChannel (TM) port bundle can be formed *from several ports on a switch....*”).

o) Claim 11[c]: a plurality of member line cards conjoined in a link aggregation (LAG) group of

parallel physical links between two endpoints in a Layer 2 data network joined together into a single logical link, having a plurality of LAG ports to forward packets through said switching core so that the node operates as a virtual media access control (MAC) bridge in said Layer 2 data network, said plurality of member line cards including at least first and second line cards, each line card having respective ports and having a respective forwarding database (FDB) to hold records associating MAC addresses with said respective ports of said line cards;

Smith discloses claim 11[c].⁴ §§VII.A.2.b)–d), VII.A.2.m); EX1003 ¶129.

- p) **Claim 11[d]: wherein said line cards are arranged so that upon receiving a data packet on an ingress line card from a MAC source address, said data packet specifying a MAC destination address, said ingress line card conveys said data packet via said switching core to at least said first line card for transmission to said MAC destination address, whereupon said first line card checks said MAC source address of said data packet against records in said FDB of said first line card, and if said FDB database of said first line card does not contain a record of an association of said MAC source address with said ingress port, adds said record to the FDB of said first line card and sends a message to at least said second line card informing said second line card of said association, and arranged, when said MAC destination address**

⁴ Petitioners also intend to argue “virtual MAC bridge” is indefinite. However, the claims are invalid as obvious under any interpretation. The question of indefiniteness is not before the Board and is no barrier to the Board’s inquiry. *Haas*, Paper 11, 11–13.

does not appear in said FDB, to flood said data packet via one and only one of said LAG ports.

Smith discloses claim 11[d]. §VII.A.2.e)–i); EX1003 ¶130.

- q) Claim 13: The node according to claim 11, wherein responsively to the message, the second line card adds the record of the association to the MAC database of the second line card if the record does not already exist in the FDB of the second line card.**

Smith discloses claim 13. §VII.A.2.j); EX1003 ¶131.

- r) Claim 16: The node according to claim 11, wherein the message comprises a synchronization packet, which is transmitted from the first line card via the switching core to at least the second line card.**

Smith discloses claim 16. §VII.A.2.k); EX1003 ¶132.

B. Ground 2: Claims 1–3 and 11–13 would have been obvious over Smith in view of Ishimori.


Although Smith renders obvious claims 1, 3, 11, and 13 alone (as described above), to the extent any elements are not disclosed or rendered obvious by Smith, they would have been obvious in combination with Ishimori. Petitioners incorporate by reference the analysis of Smith’s disclosures from above. §VII.A.2. Additionally, the combination of Smith and Ishimori renders obvious claims 2 and 12.

1. Overview of Ground 2

a) Ishimori

Ishimori is directed to a packet forwarding method that uses “path learning” and “link aggregation” to “prevent repeated flooding.” EX1005, Abstract.

Ishimori's system has node groups 60, 70 in between two terminals A10 and B20. *Id.*, Fig. 9. The node groups 60, 70 have communication cards #0–2 and #3–5, respectively. *Id.* ¶15. Ishimori teaches a method where the source MAC address of a received packet is learned by storing the source MAC address “in a buffer [*e.g.*, MAC table] had by each node.” *Id.* ¶2. Thus, the destination MAC address for a packet is searched for in the MAC table of the local device. *Id.* ¶4. If the result is found in the buffer (*e.g.*, MAC table), then the packet is transmitted according to the forwarding path that was learned. *Id.* Thus, “as long as a learning result relating to the destination MAC address (DA)” exists in the buffer, “the forwarding path to apply to the corresponding packet can be decided unambiguously by using this information,” and “there is no need for flooding.” *Id.*



However, if the MAC address is not found, Ishimori acknowledges the inefficiency of flooding to all nodes, because the learning results on the communication cards in the same trunk are not leveraged and “flooding is constantly performed.” *Id.* ¶2,19. Ishimori’s solution to this flooding problem is the same solution the ’400 patent later claimed. Specifically, Ishimori teaches, by leveraging “link aggregation,” which bundles a plurality of ports to function as one virtual port, that “one representative port is selected from among the large number of ports.” *Id.* ¶13. And the packet is therefore only flooded to “this **representative port.**” *Id.* “[A] packet received from the terminal 10 and addressed to terminal 20 is received at card #0” of node 60. *Id.* ¶15. “One representative port is selected” and “this packet is further transmitted.” *Id.* Additionally, “each communication card #0 to #5 of the node groups 60, 70” learns “information relating to the path leading to the location device” in “association with the source address.” *Id.* ¶16. Ishimori teaches a learning method that “each packet forwarding device is configured to generate a learn packet at a predetermined timing under predetermined conditions.” *Id.* ¶25. The “learn packet,” which informs the plurality of line cards regarding new associations, “is transmitted to all nodes having the corresponding trunk.” *Id.* ¶33.

Ishimori also teaches an aging process. When the buffers of the cards of each node learns a MAC address, the “hit bit” corresponding to the learned MAC address is set to “1.” *Id.* ¶9. Then, after the first aging cycle, the hit bit is set to “0,” which

indicates the MAC address should be deleted from the buffer on the next cycle. If the MAC address is learned again at the next aging process, the hit bit is reset to “1.” *Id.* This prevents “performing flooding repeatedly and thus inviting increased line traffic.” *Id.* ¶25.

b) Motivation to combine Smith and Ishimori

A POSITA would have been motivated to combine Smith and Ishimori, because both references disclose link bundling and path learning techniques in communication networks. As explained above, Smith discloses sending a packet to only “one of the communication links.” EX1004 ¶9. Smith further teaches for interfaces included in interface bundles, “the virtual network device sub-unit selects one egress interface per interface bundle via which to send the packet.” *Id.* ¶66. The purpose of flooding over only one link in a bundle is to prevent sending packets over unnecessary links. EX1003 ¶¶85–87. For example, a POSITA would have understood flooding would not occur over the link from which the packet was received. *Id.* ¶87.

Similarly, to the extent Smith does not explicitly disclose flooding via only one port in a LAG, Ishimori discloses another way of decreasing the number of links over which flooding occurs by explaining that flooding in a LAG need only happen over one representative port. EX1005 ¶13. Ishimori recognized the shortcomings in the approach of flooding all ports because the learning results on the

communication cards in the same trunk are not leveraged and “flooding is constantly performed.” *Id.* ¶19. Ishimori explains using “one representative port” prevents the application of “a different path in the same trunk to the same packet.” *Id.* ¶¶13–14. A POSITA looking to solve the shortcomings of flooding all the links would have therefore looked to Smith and Ishimori. EX1003 ¶¶85–87. Smith discloses the concept of selecting one egress interface on a virtual network bundle, and Ishimori provides using “one representative port” prevents constant flooding.

Additionally, Ishimori teaches a “learn packet” is generated at a “predetermined timing.” EX1005 ¶25. A POSITA looking to limit flooding over unnecessary links would have therefore been motivated also to implement Ishimori’s learn packet generated at a predetermined timing into Smith’s MAC notification method. EX1003 ¶88. This would have been within the knowledge of a POSA implementing such a method to make this modification and combine the teachings of Smith and Ishimori and have a reasonable expectation of success. EX1003 ¶¶85–88.

2. Analysis of Ground 2

a) Claim 1[c]

The combination of Smith and Ishimori renders obvious claim 1[c]. Smith discloses this limitation. §VII.A.2.d). To the extent this limitation is not disclosed by Smith, it is disclosed by the combination of Smith and Ishimori. Each of

Ishimori's communication cards has a "MAC table." EX1005 ¶2. Ishimori explains when a packet from terminal A is received at port #0 of card #0, the path information for address "A" is learned by storing the information in the MAC table "had by this card" (*e.g.*, card #0). *Id.* ¶3. Ishimori's line cards therefore maintain a MAC table (*e.g.*, FDB) to store path information between MAC address and the receiving card and port (*e.g.*, hold records associating MAC addresses with ports of said plurality of ports of said network node) *Id.* ¶3; EX1003 ¶¶134–136. It would have been obvious to substitute the line cards from Ishimori into Smith, and a POSITA would have had a reasonable expectation of success in making this simple substitution. *Id.* ¶136.

b) Claim 1[f]

The combination of Smith and Ishimori render obvious claim 1[f]. Smith discloses this limitation. §VII.A.2.g). To the extent this limitation is not disclosed by Smith, it is disclosed by the combination of Smith and Ishimori. Ishimori expressly discloses flooding is performed "when an address is not learned" on "one representative port" (*e.g.*, via one and only one LAG port). EX1005 ¶13. A POSITA would have been motivated to combine Smith with Ishimori because Ishimori explains using "one representative port" prevents "the application of a different path in the same trunk to the same packet." *Id.* ¶¶13,14; EX1003 ¶¶137–138. Thus, a POSITA would have looked to both Smith and Ishimori, which explain why it would

be beneficial to flood only one representative port, to solve the shortcomings of flooding all ports. EX1003 ¶¶87,138. Furthermore, a POSITA would have had a reasonable expectation of success in combining Smith with Ishimori because it would have been a simple application of Ishimori's methods on Smith's virtual network device. *Id.* ¶88,138.

c) Claim 1[g]

The combination of Smith and Ishimori renders obvious claim 1[g]. Smith discloses this limitation. §VII.A.2.h). To the extent this limitation is not disclosed by Smith, it is disclosed by the combination of Smith and Ishimori. Ishimori teaches the source MAC address is stored with the reception path information including the card number and port number. EX1005 ¶3. It would have been obvious to a POSITA to check the MAC source address against the records of the MAC table to determine whether to add a new record, or as Ishimori discloses, to overwrite the record if the MAC address already exists in the MAC table. *Id.*; EX1003 ¶¶139–140.

d) Claim 1[h]

The combination of Smith and Ishimori render obvious claim 1[h]. Smith discloses this limitation. §VII.A.2.i). To the extent this limitation is not disclosed by Smith, it is disclosed by the combination of Smith and Ishimori. Ishimori discloses the source MAC address is stored with the reception path information

including the card number and port number. EX1005 ¶3. It would have been obvious to a POSITA to check the MAC source address against the records of the MAC table to determine whether to add a new record, or as Ishimori discloses, to overwrite the record if the MAC address already exists in the MAC table. *Id.*

Ishimori further teaches a learning process where “each packet forwarding device is configured to generate a learn packet at a predetermined timing under predetermined conditions.” *Id.* ¶25. Ishimori teaches the “learn packet” that informs the plurality of line cards regarding new associations “is transmitted to all nodes having the corresponding trunk.” *Id.* ¶33. Ishimori therefore discloses sending a message of the association to each member line card for said plurality of member line cards. EX1003 ¶¶141–142. It would have been obvious to a POSITA to implement the MAC notification in Smith to perform the method taught in Ishimori to first check whether the MAC address is found in the MAC table, and if not, create a new record of the association. *Id.* ¶141. Both Smith and Ishimori teach a message of this association is sent to the plurality of member line cards. EX1004 ¶63; EX1005 ¶16.

- e) **Claim 2: The method according to claim 1, wherein sending the message comprises sending messages periodically at predefined times to inform at least the second line card of new associations between the MAC address and the respective ports.**

The combination of Smith and Ishimori discloses claim 2. Smith teaches the network device sub-unit 122(2) sends a MAC notification (*e.g.*, a message) to update the forwarding engines, but it does not disclose doing so periodically at predefined times. EX1004 ¶63. However, Ishimori discloses this. Ishimori teaches a “learn packet” (*e.g.*, message) is generated at a “predetermined timing” (*e.g.*, periodically at predefined times). EX1005 ¶25, Fig. 15. Ishimori teaches the flooding operation informs the plurality of line cards regarding new associations “is transmitted to all nodes in the corresponding trunk.” *Id.* ¶33. Thus, it would have been obvious to a POSITA that this must include at least the second line card. *Id.*; EX1003 ¶¶143–144.

f) Claim 3

The combination of Smith and Ishimori discloses claim 3. Smith discloses this limitation. §VII.A.2.j). To the extent this limitation is not disclosed by Smith, it is disclosed by the combination of Smith and Ishimori. Ishimori teaches the reception path information is associated with each other and learned, which is done by storing the record of association to the MAC table (*e.g.*, adding the record of the association to the FDB). EX1005 ¶2. The routes are transmitted via a “learn packet” that can “perform packet reception in both directions as appropriate, and path learning in both directions is performed reliably.” *Id.* ¶25. Ishimori discloses “a learn packet is transmitted to all nodes having the corresponding trunk,” which

necessarily includes at least the second line card. *Id.* ¶33; EX1003 ¶¶145–147. In Ishimori, each of the communication cards, which includes the second line card, would add a record of the association to the FDB in response to the learn packet. EX1005 ¶25. A POSITA could use Ishimori’s learn packet in place of Smith’s MAC notification frames in order to send a message to the second line card to add an entry in its MAC table. EX1003 ¶147.

g) Claim 11[c]

The combination of Smith and Ishimori discloses claim 11[c]. §§VII.A.2.b)–A.2.d)d), VII.A.2.m), VII.B.2.a); EX1003 ¶148.

h) Claim 11[d]

The combination of Smith and Ishimori discloses claim 11[d]. §§VII.A.2.e)–i), VII.B.2.b)–d); EX1003 ¶149.

i) Claim 12: The node according to claim 11, wherein at least the first line card is adapted to send messages periodically at predefined times to inform at least the second line card of new associations between the MAC addresses and the respective ports.

The combination of Smith and Ishimori discloses claim 12. §VII.B.2.e); EX1003 ¶150.

j) Claim 13

The combination of Smith and Ishimori discloses claim 12. §VII.B.2.f); EX1003 ¶151.

C. Ground 3: Claims 1, 4–7, 10–11, 14–17, and 20 would have been obvious over Smith in view of Ishimori in further view of Edsall.

Although Smith alone or in combination with Ishimori renders obvious claims 1–3, 6, 11–13, and 16, to the extent any elements are not rendered obvious by Smith and Ishimori, they would have been obvious in further view of Edsall. Petitioners incorporate by reference the analysis of Smith and Ishimori from above. §VII.B.2. Additionally, the combination of Smith, Ishimori, and Edsall renders obvious claims 4–5, 7, 10, 14–15, 17, and 20.

1. Overview of Ground 3

a) Edsall

Edsall is generally directed to techniques for updating and synchronizing “forwarding tables contained on line cards that are interconnected by a switch fabric of a distributed network switch.” EX1006, Abstract. The “forwarding table” has an L2 portion that is “used to execute forwarding decision operations for frames forwarded among ports of the line cards.” *Id.*, 5:66–6:4. Similar to Smith and Ishimori, “[i]f the frame is received at the ingress card for the first time,” the ingress forwarding engine learns the source MAC address of this frame. *Id.*, 6:26–31. This involves “creating/updating an entry of the L2 forwarding table with the source MAC address and its location (index) within the switch.” *Id.*, 6:31–34. Then, the ingress forwarding engine floods “copies of the fabric frame through its port

interfaces to all (egress) line cards of the network switch,” called the “flood-to-fabric (FF)” operation. *Id.*, 6:34–39, 18:44–47.

A “novel MN [MAC notification] frame” complements the flooding operation. *Id.*, 6:46–50. The MN frame “involves use of a primary input (PI) indicator,” which “denotes a primary input MAC address that is directly attached to a port of the line card associated with the forwarding table containing this entry.” *Id.*, 6:50–56. As Edsall explains, the PI indicator is “asserted” for “a MAC address that is learned from a frame sourced through one of the ports of the line card, as opposed to being learned through the switch fabric.” *Id.*, 6:56–60, 18:56–19:5. The frame additionally includes a POE (port-of-exit) field that “includes a plurality of bits, one for each port interface of the switch fabric.” *Id.*, 6:24–25. The POE bit instructs the switch which port interfaces on which line cards should receive the MN frame. *Id.*, 9:47–50.

b) Motivation to Combine Smith and Ishimori with Edsall

A POSITA would have been motivated to combine Smith, Ishimori, and Edsall. EX1003 ¶¶89–93. All three references disclose methods of using MAC forwarding tables in a distributed network switch. EX1004 ¶54; EX1005 ¶9; EX1006, Abstract. It would have been obvious to a POSITA to modify the MAC forwarding tables of Smith to implement the various features found in Edsall’s forwarding engine. *Id.* ¶89. Specifically, a POSITA could alter the MAC tables of

Smith to implement how Edsall's forwarding engine learns the source MAC address by "creating/updating an entry of the L2 forwarding table with the source MAC address and its location (index) within the switch." EX1006, 18:39–44; EX1003 ¶89.

A POSITA could further implement Edsall's "flood-to-fabric (FF) operations" by modifying the MAC tables to include POE bits that determine which port interfaces on which line cards should receive the MN frames. *Id.*, 18:47–50. Finally, it would have been obvious to a POSITA to modify the MAC tables to include the PI indicator as taught in Edsall in order to keep track of whether a MAC address was learned from a local line card or a different line card. EX1003 ¶89. It would have also been obvious to a POSITA to further modify the MAC forwarding tables of Smith to implement the POE field in order to determine which ports should receive the MN frame. *Id.*

Additionally, Smith and Edsall are both assigned to the same assignee, and so a POSITA would have had a motivation to combine these references for this additional reason. *Abbot Vascular, Inc. v. Flexstent*, IPR No. 2019-882, Paper 48, 28–29 (Oct. 2, 2020); *Laird Techs., Inc. v. Garftech Int'l Holdings, Inc.*, IPR No. 2014-24, Paper 46, 30–31 (Mar. 25, 2015). Furthermore, creating entries in a MAC forwarding table and adding a PI indicator field and POE field to the MAC tables

would have been a modification a POSITA would have known how to make with a reasonable expectation of success. EX1003 ¶90.

Ishimori additionally teaches an aging process using a “hit bit.” EX1005 ¶9. The “hit bit” keeps track of how long a MAC address has been in the MAC address table without having been refreshed. *Id.* If the MAC address has not been refreshed within a predefined time, then the MAC address is deleted from the buffer on the next cycle. *Id.* “As a result, performing flooding repeatedly and thus inviting increased line traffic can be prevented.” *Id.* ¶25. A POSITA would have understood that MAC addresses may become stale due to changes in the network and would have been motivated to look at solutions such as Ishimori’s aging parameter for making sure the MAC table is up to date. EX1003 ¶92. Moreover, because the size of a MAC table is limited, a POSITA would have known older entries have to be aged to make space for newer entries. *Id.* A POSITA looking to implement this solution would have therefore been motivated to combine Smith and Ishimori and it would have been an easy application of Ishimori’s methods to Smith’s system. *Id.* For example, aging information could be added to the MAC address table described in Smith using known techniques of adding information to a table. *Id.*

Furthermore, one of skill in the art would have had a reasonable expectation of success in implementing this combination because it would have required a simple addition of Ishimori’s aging process to Smith-Edsall. *Id.* ¶93. Smith, Ishimori, and

Edsall already leverage link aggregation techniques in the packet forwarding methods they disclose. Thus, it would have been within the knowledge of a POSITA to simply implement the additional aging process taught by Ishimori to Smith-Edsall, and a POSITA would have been able to do so with a reasonable expectation of success. *Id.*

2. Analysis of Ground 3

a) Claim 1[g]

To the extent this is not disclosed by the combination of Smith and Ishimori, Edsall discloses this. Edsall teaches the forwarding engine learns the *source* MAC address of a frame received at the ingress card for the first time. EX1006, 18:39–41. It does so by “creating/updating an entry of the L2 forwarding table with the source MAC address and its location (index) within the switch.” *Id.*, 18:42–44. “[I]f there is not a current entry,” the forwarding engine “learn[s] the source address/index of the frame.” *Id.*, 18:54–55. It would have been obvious to use Edsall’s learning methods with the Smith-Ishimori MAC tables, and a POSITA would have had a reasonable expectation of success in doing so. EX1003 ¶¶153–154.

b) Claim 1[h]

To the extent the combination of Smith and Ishimori does not render obvious this limitation, it would have been obvious to combine Smith and Ishimori with Edsall. Edsall teaches the forwarding engine learns the *source* MAC address of a

frame received at the ingress card for the first time. EX1006, 18:39–41. Edsall further teaches “if there is not a current entry,” the forwarding engine “learn[s] the source address/index of the frame.” *Id.*, 18:54–55. It does so by “creating/creating an entry of the L2 forwarding table with the source MAC address and its location (index) within the switch.” *Id.*, 18:42–44. It then floods copies of the fabric frame to all the egress line cards of the network switch, which Edsall calls the “flood-to-fabric (FF) operation.” *Id.*, 18:47–50. This “forces each forwarding engine associated with each egress card to either (i) update its current L2 forwarding table entry with the newly-learned source MAC address and index of the frame or, if there is not a current entry, (ii) learn the source address/index of the frame.” *Id.*, 18:50–55.

It would have been obvious to a POSITA to add Edsall’s learning method and flood-to-fabric operation to the Smith-Ishimori path learning operations. EX1003 ¶¶155–156. A POSITA would have been able to implement this teaching with a reasonable expectation of success, because the combination of Smith and Ishimori already teach a message of new associations are sent to the plurality of member line cards. EX1004 ¶63; EX1005 ¶16. It would have been within the knowledge of a POSITA to create new entries in a FDB when an association of a MAC address and port does not yet exist. EX1003 ¶156.

- c) Claim 4: The method according to claim 3 and comprising marking the records in the respective**

FDB of each line card to distinguish a first type of the records, which are added in response to data packets transmitted via a port of the line card, from a second type of the records, which are added in response to messages received from another of the line cards.

The combination of Smith, Ishimori, and Edsall discloses claim 4. Edsall discloses the “PI indicator” (*e.g.*, marking the records) is asserted when a forwarding table entry for the MAC address is learned through one of the ports (*e.g.*, data packets transmitted via a port of the line card) “as opposed to through the switch fabric” (*e.g.*, received from another of the line cards). EX1006, 6:46–64. Additionally, during prosecution, the Examiner found Edsall disclosed this limitation, which the applicant did not dispute. EX1002, 82,120–122. Specifically, Edsall discloses the “PI indicator is asserted for a destination MAC address entry of the forwarding table on the egress card and the DI contained in the switched fabric frame (*i.e.*, the ingress DI) is **different** from the DI stored in this egress forwarding table (*i.e.*, the egress DI).” EX1006, 18:56–19:5; EX1002, 82. Thus, the PI indicator is different for the MAC address entry learned through one of the ports as opposed to through the switch fabric. EX1006, 6:46–64, 18:56–19:5; EX1003 ¶¶157–158.

d) Claim 5[a]: The method according to claim 4, and comprising: associating a respective aging time with each of the records;

The combination of Smith, Ishimori, and Edsall discloses claim 5[a]. Ishimori teaches each node has a hit bit that is set to “1” at the time of learning (*e.g.*,

associating a respective aging time with each of the records). EX1005 ¶9. The hit bit is then set to “0,” so on the next aging process, “the corresponding MAC address is deleted from the buffer.” *Id.* Thus, the hit bit is how Ishimori associates a respective aging time with each of the records. EX1003 ¶¶159–161. It would have been obvious to a POSITA to implement a “hit bit” on the Smith’s lookup tables with Edsall’s PI indicator in order to associate a respective aging time with each of the records. *Id.* ¶161.

**e) Claim 5[b]: refreshing the records in the FDB
responsively to further packets transmitted by the line
cards; and**

The combination of Smith, Ishimori, and Edsall discloses claim 5[b]. As described above, Ishimori teaches when the hit bit is “0,” “the corresponding MAC address is deleted from the buffer.” EX1005 ¶¶9, 11. Ishimori then teaches when the destination MAC address is received at card #0, the source MAC address is learned by “overwriting,” and the hit bit is returned to “1” (*e.g.*, refreshing the records in the FDB responsively to further packets transmitted by the line cards). *Id.* ¶11. It would have been obvious to a POSITA implement Ishimori’s “hit bit” on Smith’s lookup tables with Edsall’s PI indicator and to refresh the hit bit by learning the source MAC address of the data packets transmitted by the line cards. EX1003 ¶¶162–163.

f) Claim 5[c]: removing the records from the respective FDB if the records are not refreshed within the respective aging time.

The combination of Smith, Ishimori, and Edsall discloses claim 5[c]. Ishimori teaches if the hit bit remains “0” in the next aging process, then the corresponding MAC address is deleted (*e.g.*, removing the records from the FDB if the records are not refreshed within the respective aging time). EX1005 ¶¶9,12. It would have been obvious to a POSITA implement Ishimori’s “hit bit” on Smith’s lookup tables with Edsall’s PI indicator and to remove the records from the respective lookup tables if the record has not been refreshed by the next aging process. EX1003 ¶¶164–165.

g) Claim 6

The combination of Smith, Ishimori, and Edsall discloses claim 6. Smith teaches sending a message that is a synchronization packet . §VII.A.2.k). Edsall teaches the “plurality of line cards” are “interconnected by a switch fabric 550,” which could be the “switching core.” EX1006, 8:20–27. A POSITA would have understood that the packet is sent from the first line card via Edsall’s “switching fabric” to at least the second line card. EX1003 ¶166.

h) Claim 7: The method according to claim 6, wherein sending the synchronization packet comprises, if the record in the FDB associates the MAC source address with a port different from the one of the ports on which the data packet was received, changing the record in the FDB of the first line card and sending a synchronization update packet to at least the second line card to indicate that the record has been changed

The combination of Smith, Ishimori, and Edsall discloses claim 7. EX1003 ¶¶167–169. Smith teaches its virtual network device includes several line cards, which include several interfaces. EX1004 ¶46. “[W]hen updating control protocol behavior of virtual link bundle 250(1), a user can simply access virtual network device sub-unit 122(1) (instead of accessing both virtual network device sub-units 122(1) and 122(2)).” *Id.* ¶59. “Virtual network device sub-unit 122(1) can then automatically propagate to network device 122(2) any changes made by the user to the control protocols.” *Id.* Smith teaches “MAC notification frames are used to keep the content of the L2 tables in virtual network device sub-unit 122(1) synchronized with the content of the L2 tables in virtual network device sub-unit 122(2) and vice versa.” *Id.* ¶62.




FIG. 3

Thus, if the record in the FDB of the first line card (*e.g.*, line card 304(1)) associates the MAC source address (*e.g.*, virtual device 120(1)) with a port different from the one of the ports on which the data packet was received (*e.g.*, a port other than 320(9)), the MAC notification frames will notify and update the L2 tables in the virtual network device 202, which would include at least the second line card (*e.g.*, line card 304(4)) to indicate the record has been changed. Similarly, Ishimori teaches the routes are transmitted via a “learn packet” (*e.g.*, synchronization packet). EX1005 ¶25. During the packet transfer, each of the communication cards learns information about the route (*e.g.*, the second line card). *Id.* ¶¶28,33; EX1003 ¶168.

Edsall further discloses the forwarding engine generates an MN frame (*e.g.*, synchronization packet) that may get sent to the SMC (switch management card) to ensure that FwdT0 (*e.g.*, the forwarding table) is synchronized. EX1006, 17:26–38. The forwarding engine also asserts an appropriate bit of the POE field (port-of-exit field) when generating the MN frame, which is a port different from one of the ports on which data was received. *Id.* It would have been obvious to a POSITA to modify the MN frame from Smith to include the POE field in order to note the port interface of the switch fabric. *Id.*, 6:24–25; EX1003 ¶169. This disclosure was cited by the examiner during prosecution and was not disputed by the applicant. EX1002, 83,120–122.

- i) **Claim 10[a]: The method according to claim 1, and comprising: conveying a further data packet, received**

from a further MAC source address, to the second line card for transmission over the network;

The combination of Smith, Ishimori, and Edsall discloses claim 10[a]. EX1003 ¶¶170–171. Smith teaches the uplink interface receives a data packet with the “sending device’s MAC address” (*e.g.*, a MAC source address). EX1004 ¶54; *see also* §VII.A.2.e). Based on Figure 3, a further data packet received from a further MAC source address (*e.g.*, network device 120(3)) would be conveyed to the second line card (*e.g.*, line card 304(4)) for transmission over the network.




FIG. 3

To the extent Smith does not teach this claim, Edsall does. Edsall teaches in each subsequent frame, the encoded address recognition logic (EARL) circuit looks up the MAC address and “sends the corresponding rewrite information over the local bus after the frame” (*e.g.*, conveying a further data packet to the second line card for transmission over the network). EX1006, 14:22–34. This was cited by the examiner during prosecution and was not disputed by the applicant. EX1002, 83–84, 120–122.

j) Claim 10[b]: checking the further MAC source address against the records in the FDB of the second line card; and

The combination of Smith, Ishimori, and Edsall discloses claim 10[b]. EX1003 ¶¶172–173. Smith teaches the virtual network device sub-unit learns the source identifier of the sending device (*e.g.*, MAC source address). EX1004 ¶65; *see also* §VII.A.2.h). These identifiers are stored in a lookup table (*e.g.*, FDB) on virtual network device sub-unit. *Id.* ¶61. The MAC source address (*e.g.*, network device 120(3)) would be checked in the records of the FDB of said second line card (*e.g.*, line card 304(4)). EX1003 ¶172.

To the extent Smith does not disclose this limitation, Edsall does. Edsall teaches checking to see if the rewrite information matches (*e.g.*, checking the MAC source address against the records in the FDB of the second line card). EX1006, 14:22–34 (“The destination port circuitry (or, alternatively, a UDlink or central rewrite engine) matches the frame with the rewrite information and modifies the

frame as needed by replacing, *inter alia*, the destination and source MAC addresses.”). This was cited by the examiner during prosecution and was not disputed by the applicant. EX1002, 83–84, 120–122.

- k) Claim 10[c] responsively to the further data packet, adding a further record with respect to the MAC source address to the FDB of the second line card and sending a further message to inform at least the first line card of the further record.**

The combination of Smith, Ishimori, and Edsall discloses claim 10[c]. EX1003 ¶¶174–175. Smith teaches the network device sub-unit 122(2) sends a MAC notification (*e.g.*, a message) to update the forwarding engines (*e.g.*, sending a further message to inform at least the first line card) when it learns of a new association (*e.g.*, the further record). EX1004 ¶63; *see also* §VII.A.2.h). “After being updated based on the MAC notification, the forwarding engines in the virtual network device sub-unit 122(1) now know the location of the device identified by the destination address.” *Id.* A POSITA would understand that the MAC notification (*e.g.*, message) is therefore sent to inform at least the first member line card (*e.g.*, line card 304(1)) of the new association. EX1003 ¶¶174–175.

To the extent Smith does not disclose this limitation, Edsall does. Edsall teaches the forwarding engine “modifies the frame as needed by replacing” the “destination and source MAC addresses” (*e.g.*, adding a further record with respect to the MAC source address to the FDB). EX1006, 14:22–34. This was cited by the

examiner during prosecution and was not disputed by the applicant. EX1002, 83–84,120–122. It would have been obvious to a POSITA to use any of the learning methods from Smith, Ishimori, or Edsall, which all disclose updating a second line card with a further record and sending a message to inform at least the first line card of the further record. EX1003 ¶175.

l) Claim 11[a]

To the extent this is not disclosed by the combination of Smith and Ishimori, Edsall discloses this. The '400 patent describes the “switching core” as linking the multiple line cards. EX1001, 6:8–10. Edsall teaches the “plurality of line cards” are “interconnected by a switch fabric 550.” EX1006, 8:20–27. Thus, a POSITA would have understood that the “switching fabric” in Edsall could be a “switching core.” EX1003 ¶176.

m) Claim 11[d]

The combination of Smith, Ishimori, and Edsall renders obvious claim 11[d]. §§VII.A.2.p),VII.B.2.h),VII.C.2.a)–b); EX1003 ¶177.

n) Claim 14: The node according to claim 13, wherein the records in the respective FDB of each line card are marked to distinguish a first type of the records, which are added in response to data packets transmitted via a port of the line card, from a second type of the records, which are added in response to messages received from another of the line cards.

The combination of Smith, Ishimori, and Edsall discloses claim 14.
§VII.C.2.c); EX1003 ¶178.

- o) Claim 15: The node according to claim 14, wherein a respective aging time is associated with each of the records, and wherein the line cards are operative to refresh the records in the FDB responsively to further packets transmitted by the line cards, and to remove the records from the respective FDB if the records are not refreshed within the respective aging time.**

The combination of Smith, Ishimori, and Edsall discloses claim 15.
§§VII.C.2.d)–f); EX1003 ¶179.

- p) Claim 16**

The combination of Smith, Ishimori, and Edsall discloses claim 16.
§VII.C.2.g); EX1003 ¶180.

- q) Claim 17: The node according to claim 16, wherein the line cards are operative so that if the record in the FDB associates the MAC source address with a port different from the one of the ports on which the data packet was received, the record in the FDB of the first line card is changed, and the synchronization packet comprises a synchronization update packet, which indicates to at least the second line card to indicate that the record has been changed.**

The combination of Smith, Ishimori, and Edsall discloses claim 17.
§VII.C.2.h); EX1003 ¶181.

- r) Claim 20: The node according to claim 11, wherein the line cards are adapted to forward a further data packet, received from a further MAC source address, to the second line card for transmission over the**

network, whereupon the second line card checks the further MAC source address against the records in the FDB of the second line card, and responsively to the further data packet, adds a further record with respect to the MAC source address to the FDB of the second line card and sends a further message to inform at least the first line card of the further record.

The combination of Smith, Ishimori, and Edsall discloses claim 20. §§VII.C.2.i)–k); EX1003 ¶182.

D. Ground 4: Claims 8–9 and 18–19 would have been obvious over Smith in view of Ishimori in further view of Zelig

The combination of Smith, Ishimori, and Zelig renders obvious claims 8–9 and 18–19.

1. Overview of Ground 4

a) Zelig

Zelig is directed to a data communication network that “includes a plurality of primary virtual bridges, interconnected by primary virtual connections.” EX1007, Abstract. Zelig discloses “MAC bridges that implement the 802.1D standard allow MAC devices attached to physically separated LANs to appear to each other as if they were attached to a single LAN.” *Id.* ¶3. A MAC bridge “includes two or more MAC devices that interconnect the bridge ports to respective LANs.” *Id.* The MAC bridges “maintain a database to map destination MAC address of the packets they receive to bridge ports.” *Id.* ¶4. This database is built “by means of a learning

process, in which it associates the source MAC address of each incoming packet with the port on which the packet was received.” *Id.* The purported invention in Zelig is directed to providing “improved mechanisms for protection from failure in virtual private networks (VPNs).” *Id.* ¶18. Zelig thus discloses a data communication network including “a plurality of primary virtual bridges” and “one or more backup virtual bridges” that are “arranged to transmit the packets using a virtual private LAN service (VPLS).” *Id.* ¶¶21–24. “[E]ach of the primary virtual bridges is adapted to maintain a respective media access control (MAC) table, and to forward the data packets in accordance with entries in the MAC table.” *Id.* ¶27. The VPN 20 includes multiple primary virtual bridges. *Id.* ¶42.




FIG. 1

b) Motivation to combine Smith and Ishimori with Zelig

A POSITA would have been motivated to combine Smith-Ishimori with Zelig. EX1003 ¶¶94–95. All three references disclose methods of using MAC bridges in a data communication network. EX1004 ¶54; EX1005 ¶9; EX1007, Abstract. Zelig teaches the MAC bridges “maintain a database to map destination MAC address of the packets they receive to bridge ports.” EX1007 ¶4. This database is built “by means of a learning process, in which it associates the source MAC address of each incoming packet with the port on which the packet was received.” *Id.* In order to “improve[] mechanisms for protection from failure in virtual private networks (VPNs),” Zelig teaches a data communication network including “a plurality of primary virtual bridges” and “one or more backup virtual bridges” can be “arranged to transmit the packets using a virtual private LAN service (VPLS).” *Id.* ¶¶18,21–24.

A POSITA looking to improve against failure in VPNs would look to Zelig and implement Zelig’s MAC bridges into the Smith-Ishimori system to create a separate MAC table for each MAC bridge as taught in Zelig. EX1003 ¶94. It would have also been obvious to a POSITA to configure each virtual MAC bridge to serve a respective VPN, as shown above in Figure 1 of Zelig. *Id.* ¶95. Furthermore, implementing the MAC bridges as taught in Zelig to the Smith-Ishimori data communication network would have been a modification that a POSITA would have known how to make with a reasonable expectation of success. *Id.*

2. Analysis of Ground 4

- a) Claim 8: The method according to claim 1, wherein the network node is configured to operate as multiple virtual MAC bridges in a Layer 2 virtual private network (VPN), wherein each virtual MAC bridge is configured to serve a respective VPN instance, and wherein the records associating the MAC addresses with the respective ports are maintained independently for each of the VPN instances.**

The combination of Smith, Ishimori, and Zelig discloses claim 8. Zelig discloses its VPN 20 contains multiple primary virtual bridges 30, 32, and 34 (*e.g.*, multiple virtual MAC bridges in a Layer 2 VPN). EX1007 ¶¶42–43. It would have been obvious to a POSITA that the virtual bridges servicing VPN 20 are configured to serve that VPN. EX1003 ¶¶184–185. Zelig further teaches “each of the primary virtual bridges is adapted to maintain a respective media access control (MAC) table, and to forward the data packets in accordance with entries in the MAC table” EX1007 ¶31. Thus, Zelig discloses each MAC bridge maintains its own MAC table, and because the MAC bridge serves only that VPN instance, the records associating the MAC addresses with the respective ports are maintained independently for each VPN instance. EX1003 ¶185.

- b) Claim 9: The method according to claim 8, wherein the VPN instance is a VPLS instance among multiple VPLS instances served by the network node, and wherein sending the message comprises identifying the VPLS instance in the message so as to inform all the line cards that serve the VPLS instance.**

The combination of Smith, Ishimori, and Zelig discloses claim 9. Smith teaches the MAC notification (*e.g.*, message) is “distributed to any other forwarding engines within virtual network device sub-unit 122(2).” EX1004 ¶63. Zelig provides “VPN 20 is built around a virtual private LAN service (VPLS), operating within a network 22.” EX1007 ¶42. It would have been obvious to a POSITA to adopt the notification scheme in Smith to the architecture in Zelig so the MAC notification (*e.g.*, message) is sent to all the line cards not just in the virtual network device sub-unit, but to all the line cards that serve the VPLS instances. EX1003 ¶186–187. Furthermore, the IEEE 802.1Q standards disclose a VLAN identifier (VID) as a twelve-bit field that “uniquely identif[ies] the VLAN to which the frame belongs.” EX1008 §9.3.2.3. Thus, it would have been within the knowledge of a POSITA to identify the VPLS instance in the message, such as using the VID, in order to inform all the line cards in the VPLS. EX1003 ¶187.

- c) **Claim 18: The node according to claim 11, wherein at least some of the line cards are configured so that the node operates as multiple virtual MAC bridges in a Layer 2 virtual private network (VPN), wherein each virtual MAC bridge is configured to serve a respective VPN instance, and wherein the records associating the MAC addresses with the respective ports are maintained independently for each of the VPN instances.**

The combination of Smith, Ishimori, and Zelig discloses claim 18. §VII.D.2.a); EX1003 ¶188.

- d) Claim 19: The node according to claim 18, wherein the VPN instance is a VPLS instance among multiple VPLS instances served by the network node, and wherein the VPLS instance is identified in the message so as to inform all the line cards that serve the VPLS instance of the association.**

The combination of Smith, Ishimori, and Zelig discloses claim 19. §VII.D.2.b); EX1003 ¶189.

E. Ground 5: Claims 9 and 19 would have been obvious over Smith in view of Ishimori, Zelig, and 802.1Q

The combination of Smith, Ishimori, Zelig, and 802.1Q renders obvious claims 9 and 19.

1. Overview of Ground 5

a) 802.1Q

The 802.1Q standard defines the architecture for Virtual Bridged LANs, its services, protocols, and algorithms. EX1008, Abstract. It was part of the effort of the IEEE to standardize virtual LAN services in bridged LANs. EX1003 ¶84. The standard discloses a twelve-bit VLAN identifier field is used to “uniquely identify the VLAN to which the frame belongs.” EX1008 §9.3.2.3. Such a standard would have been known and within the knowledge of a POSITA. EX1003 ¶84.

b) Motivation to combine Smith, Ishimori, and Zelig with 802.1Q

A POSITA would have been motivated to combine Smith-Ishimori-Zelig with 802.1Q-1998. *Id.* ¶96. All four references pertain to MAC bridges in a data

communication network. EX1004 ¶54; EX1005 ¶9; EX1007, Abstract; EX1008, Abstract. Moreover, Zelig expressly cites to the 802.1Q standard. EX1007 ¶12. A POSITA would have therefore been motivated to look to the teachings of 802.1Q in combination with Smith-Ishimori-Zelig to modify the messages from Smith-Ishimori-Zelig with a VLAN identifier as disclosed in 802.1Q in order to identify the VPLS instance. EX1003 ¶96. Furthermore, modifying the Smith-Ishimori-Zelig message to add the VLAN identifier would have been a modification that a POSITA would have known how to make with a reasonable expectation of success. *Id.*

2. Analysis of Ground 5

a) Claim 9

The combination of Smith, Ishimori, Zelig, and 802.1Q discloses claim 9. §VII.D.2.b); EX1003 ¶191.

b) Claim 19

The combination of Smith, Ishimori, Zelig, and 802.1Q discloses claim 19. §VII.D.2.b), VII.D.2.d); EX1003 ¶192.

VIII. PTAB DISCRETION SHOULD NOT PRECLUDE INSTITUTION

A. PTAB should not exercise its discretion to deny institution under *Fintiv*.

1. Factor 1: Institution will increase the likelihood of stay

The District Court has made clear a motion to stay prior to institution of an IPR would be premature, so Petitioners have not yet requested a stay. The Board

has repeatedly stated it will not speculate on how a district court will treat a motion to stay. *See, e.g., VMware Inc. v. Intellectual Ventures II LLC*, IPR2020-00859, Paper No. 13, at 12 (PTAB Nov. 5, 2020) (informative). Thus, this factor is, at worst, neutral.

2. Factor 2: District Court schedule

In the Cisco litigation, January 3, 2023 is the deadline to submit a Joint Claim Construction Statement. EX1013. The *Markman* hearing is not until February 15, 2023. EX1014. No date has been set for trial. A Final Written Decision (“FWD”) would be expected in the second quarter of 2024. The Dell litigation does have a scheduling order yet. The parties worked expeditiously to prepare and file this petition approximately four months after Cisco received Patent Owner’s infringement contentions on August 3, 2022.

Furthermore, the current average time-to-trial for the W.D. Texas district court (for the Cisco litigation) is over two years. EX1023 (showing W.D. Texas median time-to-trial of 28.8 months); *In re Apple Inc.*, 979 F.3d 1332, 1350 (Fed. Cir. 2020). Petitioners note “a court’s general ability to set a fast-paced schedule is not particularly relevant...where, like here, the forum itself has not historically resolved cases so quickly.” *Id.* at 1344. The current average time-to-trial for the Delaware district court (for the Dell litigation) is much longer. EX1024 (median time-to-trial of 39 months). Accordingly, this factor weighs against discretionary

denial. *See Sand Revolution II, LLC v. Continental Intermodal Grp.*, IPR2020-01393, Paper 24 (June 16, 2020).

3. Factor 3: Petitioners' investment in IPR outweighs forced investment in litigation to date

The co-pending litigations are in their early stages, and the investment in them has been minimal. The *Markman* hearing is not for another couple months in the Cisco litigation and discovery has not begun. EX1014; *see PEAG LLC v. Varta Microbattery GMBH*, IPR2020-01214, Paper 8, 17 (Jan. 6, 2021). The Dell litigation does not yet have a scheduling order.

4. Factor 4: The Petition raises unique issues

Preliminary invalidity contentions in the Cisco litigation were served on Patent Owner on September 28, 2022. Instituting a proceeding will allow the Board to address the asserted prior art, and the issues will be narrowed in the co-pending litigations due to the estoppel provisions of 35 U.S.C. § 315(e)(2).

Moreover, there will be no overlap of prior art issues. If the Board institutes trial, Petitioners will cease asserting in the co-pending litigations the combination of references on which trial is instituted for the claims on which trial is instituted, to the extent Petitioners even assert the same combination in the litigations against them. This factor weighs against discretionary denial. *See Verizon v. Huawei*, IPR2020-01079, Paper 10, 38 (Jan. 14, 2021).

5. Factor 5: Whether the Petitioners and Defendants in the parallel litigation are the same party

Petitioners are defendants in co-pending litigations. While the *Fintiv* case indicates that a difference between the district court defendant and the petitioner may weigh against discretionary denial, nothing within the *Fintiv* case suggests the same party between the proceedings weighs in favor of discretionary denial. *See HP Inc. v. Slingshot Printing LLC*, IPR2020-01084, Paper 13, 9 (Jan. 14, 2021). Accordingly, this factor is neutral and should not be a basis for denying institution.

6. Factor 6: Other circumstances support institution

The prior art presented in this Petition renders the Challenged Claims unpatentable as obvious. The merits of Petitioners' arguments are compelling because they lead to a strong conclusion that one or more claims are unpatentable by a preponderance of the evidence. This factor therefore weighs against discretionary denial. As such, the *Fintiv* factors are either neutral or weigh against discretionary denial. Because this Petition was filed approximately four months before the statutory bar date, institution should not be denied on discretionary factors.

B. PTAB should not exercise its discretion to deny institution under *Becton and Advanced Bionics*

Furthermore, institution should not be denied under the framework set forth in *Becton, Dickinson & Co. v. B. Braun Melsungen AG*, IPR2017-01586, Paper 8

(Dec. 15, 2017) (§III.C.5, ¶1, precedential) and *Advanced Bionics LLC v. MED-EL Elektromedizinische Gerate GmbH*, IPR2019-01469, Paper 6 (PTAB Feb. 13, 2020) (precedential). While Edsall was before the Examiner during prosecution, Edsall is only relied upon in this petition for disclosures of limitations that were not disputed by the applicant. Moreover, Smith alone and Smith in combination with Ishimori, which were not before the Examiner during prosecution form the basis of the disclosures for the independent claims here, to which Edsall serves as only a supplemental reference as disclosure for certain claim limitations. Therefore, the arguments presented in this Petition are not the same as, or substantially similar to, arguments previously presented to the Office. Thus, the first prong of the *Advanced Bionics* framework is not met, and there is no need to reach the second prong to resolve against discretionary denial under §325(d). *Oticon Med. AB v. Cochlear Ltd.*, IPR2019-00975 Paper 15, 20 (PTAB Oct. 16, 2019) (precedential).

C. Discretionary denial under *General Plastic* is not appropriate

The '400 patent has not been challenged in any prior IPR petition, so none of *General Plastic* discretionary institution factors apply here. See *Gen. Plastic Indus. Co., Ltd. v. Canon Kabushiki Kaisha*, IPR2016-01357, Paper 19, 16 (Sept. 6, 2016) (Section II.B.4.i. precedential).

IX. CONCLUSION

Petitioners have established a reasonable likelihood that the Challenged Claims are unpatentable. Petitioners therefore respectfully request that IPR be instituted.

DATED: December 23, 2022

Respectfully Submitted,

/s/ Stuart Rosenberg

Stuart Rosenberg

(Reg. No. 60,772)

Gibson, Dunn & Crutcher LLP

1881 Page Mill Road

Palo Alto, CA 94304-1211

Tel: 650-849-5389

SRosenberg@gibsondunn.com

*Attorney for Petitioners Dell Inc., Dell
Technologies Inc., and Cisco Systems,
Inc.*

CERTIFICATION UNDER 37 C.F.R. § 42.24(d)

Under the provisions of 37 C.F.R. § 42.24(d), the undersigned attorney hereby certifies that the word count for Sections I and III-VIII of the foregoing Petition for *Inter Partes* Review is 13,969, according to the word count tool in Microsoft Word.

DATED: December 23, 2022

Respectfully Submitted,

/s/ Stuart Rosenberg

Stuart Rosenberg
(Reg. No. 60,772)
Gibson, Dunn & Crutcher LLP
1881 Page Mill Road
Palo Alto, CA 94304-1211
Tel: 650-849-5389
SRosenberg@gibsondunn.com

*Attorney for Petitioners Dell
Technologies Inc., Dell Inc., and
Cisco Systems, Inc.*

CERTIFICATE OF SERVICE

The undersigned certifies service pursuant to 37 C.F.R. §§ 42.6(e) and 42.105(a), (b) on the Patent Owner via US Express Mail and FedEx Priority Overnight of a copy of this Petition for *Inter Partes* Review and supporting materials at the correspondence address of record for the '400 patent shown in USPTO PAIR:

May Patents Ltd. (attorney/agent of record)
c/o Dorit Shem-Tov
P. O. Box 7230
Ramat-Gan, - 5217102
Israel

Corrigent Corporation (assignee of record)
291 Main St.
C/O New England Intellectual Property
West Newbury, Massachusetts 01985

DATED: December 23, 2022

Respectfully Submitted,

/s/ Stuart Rosenberg

Stuart Rosenberg (Reg. No. 60,772)
Gibson, Dunn & Crutcher LLP
1881 Page Mill Road
Palo Alto, CA 94304-1211
Tel: 650-849-5389
SRosenberg@gibsondunn.com

*Attorney for Petitioners Dell
Technologies Inc., Dell Inc., and
Cisco Systems, Inc.*

APPENDIX: CHALLENGED CLAIM LISTING

No.	Limitation
1[pre]	A method for communication, comprising:
1[a]	configuring a network node having a plurality of ports, and at least first and second line cards with respective first and second ports, to operate as a distributed media access control (MAC) bridge in a Layer 2 data network;
1[b]	configuring a link aggregation (LAG) group of parallel physical links between two endpoints in said Layer 2 data network joined together into a single logical link, said LAG group having a plurality of LAG ports and a plurality of conjoined member line cards;
1[c]	providing for each of said member line cards a respective forwarding database (FDB) to hold records associating MAC addresses with ports of said plurality of ports of said network node;
1[d]	receiving a data packet on an ingress port of said network node from a MAC source address, said data packet specifying a MAC destination address on said Layer 2 data network;
1[e]	conveying, by transmitting said data packet to said MAC destination address via said first port, said received data packet in said network node to at least said first line card for transmission to said MAC destination address;
1[f]	if said MAC destination address does not appear in said FDB, flooding said data packet via one and only one LAG port of said plurality of LAG ports;
1[g]	checking said MAC source address of the data packet against records in said FDB of said first line card; and
1[h]	if said FDB of said first line card does not contain a record of an association of said MAC source address with said ingress port, creating a new record of said association, adding said new record to the FDB of said first line card, and sending a message of the association to each member line card of said plurality of member line cards.

No.	Limitation
2	The method according to claim 1, wherein sending the message comprises sending messages periodically at predefined times to inform at least the second line card of new associations between the MAC addresses and the respective ports.
3	The method according to claim 1, and comprising receiving the message at the second line card, and responsively to the message, adding the record of the association to the FDB of the second line card if the record does not already exist in the FDB of the second line card.
4	The method according to claim 3, and comprising marking the records in the respective FDB of each line card to distinguish a first type of the records, which are added in response to data packets transmitted via a port of the line card, from a second type of the records, which are added in response to messages received from another of the line cards.
5[pre]	The method according to claim 4, and comprising
5[a]	associating a respective aging time with each of the records;
5[b]	refreshing the records in the FDB responsively to further packets transmitted by the line cards; and
5[c]	removing the records from the respective FDB if the records are not refreshed within the respective aging time.
6	The method according to claim 1, wherein sending the message comprises transmitting a synchronization packet from the first line card via a switching core of the network node to at least the second line card
7	The method according to claim 6, wherein sending the synchronization packet comprises, if the record in the FDB associates the MAC source address with a port different from the one of the ports on which the data packet was received, changing the record in the FDB of the first line card and sending a synchronization update packet to at least the second line card to indicate that the record has been changed.

No.	Limitation
8	The method according to claim 1, wherein the network node is configured to operate as multiple virtual MAC bridges in a Layer 2 virtual private network (VPN), wherein each virtual MAC bridge is configured to serve a respective VPN instance, and wherein the records associating the MAC addresses with the respective ports are maintained independently for each of the VPN instances
9	The method according to claim 8, wherein the VPN instance is a VPLS instance among multiple VPLS instances served by the network node, and wherein sending the message comprises identifying the VPLS instance in the message so as to inform all the line cards that serve the VPLS instance.
10[pre]	The method according to claim 1, and comprising:
10[a]	conveying a further data packet, received from a further MAC source address, to the second line card for transmission over the network;
10[b]	checking the further MAC source address against the records in the FDB of the second line card; and
10[c]	responsively to the further data packet, adding a further record with respect to the MAC source address to the FDB of the second line card and sending a further message to inform at least the first line card of the further record.
11[pre]	A node for network communication, comprising:
11[a]	a switching core;
11[b]	a plurality of ports;
11[c]	a plurality of member line cards conjoined in a link aggregation (LAG) group of parallel physical links between two endpoints in a Layer 2 data network joined together into a single logical link, having a plurality of LAG ports to forward packets through said switching core so that the node operates as a virtual media access control (MAC) bridge in said Layer 2 data network, said plurality of member line cards including at least first and second line cards, each line card having respective ports and having a respective forwarding database (FDB) to hold records associating MAC addresses with said respective ports of said line cards,

No.	Limitation
11[d]	wherein said line cards are arranged so that upon receiving a data packet on an ingress line card from a MAC source address, said data packet specifying a MAC destination address, said ingress line card conveys said data packet via said switching core to at least said first line card for transmission to said MAC destination address, whereupon said first line card checks said MAC source address of said data packet against records in said FDB of said first line card, and if said FDB database of said first line card does not contain a record of an association of said MAC source address with said ingress port, adds said record to the FDB of said first line card and sends a message to at least said second line card informing said second line card of said association, and arranged, when said MAC destination address does not appear in said FDB, to flood said data packet via one and only one of said LAG ports.
12	The node according to claim 11, wherein at least the first line card is adapted to send messages periodically at predefined times to inform at least the second line card of new associations between the MAC addresses and the respective ports.
13	The node according to claim 11, wherein responsively to the message, the second line card adds the record of the association to the MAC database of the second line card if the record does not already exist in the FDB of the second line card.
14	The node according to claim 13, wherein the records in the respective FDB of each line card are marked to distinguish a first type of the records, which are added in response to data packets transmitted via a port of the line card, from a second type of the records, which are added in response to messages received from another of the line cards.
15	The node according to claim 14, wherein a respective aging time is associated with each of the records, and wherein the line cards are operative to refresh the records in the FDB responsively to further packets transmitted by the line cards, and to remove the records from the respective FDB if the records are not refreshed within the respective aging time.
16	The node according to claim 11, wherein the message comprises a synchronization packet, which is transmitted from the first line card via the switching core to at least the second line card.

No.	Limitation
17	The node according to claim 16, wherein the line cards are operative so that if the record in the FDB associates the MAC source address with a port different from the one of the ports on which the data packet was received, the record in the FDB of the first line card is changed, and the synchronization packet comprises a synchronization update packet, which indicates to at least the second line card to indicate that the record has been changed.
18	The node according to claim 11, wherein at least some of the line cards are configured so that the node operates as multiple virtual MAC bridges in a Layer 2 virtual private network (VPN), wherein each virtual MAC bridge is configured to serve a respective VPN instance, and wherein the records associating the MAC addresses with the respective ports are maintained independently for each of the VPN instances.
19	The node according to claim 18, wherein the VPN instance is a VPLS instance among multiple VPLS instances served by the network node, and wherein the VPLS instance is identified in the message so as to inform all the line cards that serve the VPLS instance of the association.
20	The node according to claim 11, wherein the line cards are adapted to forward a further data packet, received from a further MAC source address, to the second line card for transmission over the network, whereupon the second line card checks the further MAC source address against the records in the FDB of the second line card, and responsively to the further data packet, adds a further record with respect to the MAC source address to the FDB of the second line card and sends a further message to inform at least the first line card of the further record.