

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CISCO SYSTEMS, INC., MICROSOFT CORPORATION, AMAZON.COM,
INC., AMAZON WEB SERVICES, INC., AND AMAZON.COM SERVICES
LLC
Petitioners

v.

LS CLOUD STORAGE TECHNOLOGIES, LLC
Patent Owner

Case IPR2023-00733
U.S. Patent No. 10,154,092

**PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NO. 10,154,092**

TABLE OF CONTENTS

	<u>Page</u>
I. INTRODUCTION	1
II. MANDATORY NOTICES	2
A. Real Parties-in-Interest.....	2
B. Related Matters	2
C. Counsel and Service	3
D. Power of Attorney	5
E. Fees.....	6
III. REQUIREMENTS FOR <i>INTER PARTES</i> REVIEW.....	6
A. Grounds for Standing	6
B. Identification of Challenge and Relief Requested.....	6
IV. THE '092 PATENT	8
V. PROSECUTION HISTORY	10
VI. THE CHALLENGED CLAIMS ARE UNPATENTABLE.....	13
A. State of the Art.....	13
B. Overview of the Prior Art.....	14
a. Heil.....	14
b. Nakayama	16
c. Gulick	18
d. Berman	18
C. Level of Ordinary Skill in the Art	18
D. Claim Construction	19

E.	Ground 1: Claims 1-3, 7-12, and 19-23 are anticipated by Heil under 35 U.S.C. § 102(e)	22
a.	Claim 1	22
i.	Claim 1.pre	22
ii.	Claim 1.a: “a first interface configured to receive input/output (I/O) traffic from a first host device via a dedicated I/O channel, the I/O traffic comprising a read command”	23
iii.	Claim 1.b: “a second interface configured to receive first data via a network”	25
iv.	Claim 1.c: “a cache memory configured to store second data”	27
v.	Claim 1.d: “a storage device configured to store third data”	27
vi.	Claim 1.e: “a processor coupled to the cache memory, the processor coupled to the storage device via a communication path that is distinct from the dedicated I/O channel, the processor configured to access the cache memory during processing of the I/O traffic”	28
vii.	Claim 1.f: “the processor further configured to perform an access operation at the storage device based on the I/O traffic”	31
b.	Claim 2	32
c.	Claim 3	32
d.	Claim 7	34
e.	Claim 8	35
f.	Claim 9	36

i.	Claim 9.a: “the storage device and the dedicated I/O channel are independently accessible”	36
ii.	Claim 9.b: “wherein the processor is configured to read at least a portion of the second data from the cache memory based on the read command”	37
g.	Claim 10	38
h.	Claim 11	39
i.	Claim 11.a: “configuration manager circuitry configured to route an I/O request included in the I/O traffic to the cache memory, to route the I/O request to the storage device, or to deny the I/O request”	39
ii.	Claim 11.b: “front-end circuitry configured to process the I/O request”	40
iii.	Claim 11.c: “back-end circuitry configured to perform a read operation or a write operation at the storage device based on the I/O request”	41
i.	Claim 12	41
j.	Claim 19	42
i.	Claim 19.pre.....	42
ii.	Claim 19.a: “a first interface configured to receive input/output (I/O) traffic from a host computer via a dedicated I/O channel, the I/O traffic comprising one or more read commands, one or more write commands, or a combination thereof”	42
iii.	Claim 19.b: “a second interface configured to receive data via a network”	42
iv.	Claim 19.c: “a cache memory”	42
v.	Claim 19.d: “a storage device”	42

vi.	Claim 19.e: “a processor coupled via a communication path to the storage device, the processor configured to access the cache memory during processing of the I/O traffic”	42
vii.	Claim 19.f: “the processor further configured to perform one or more access operations at the storage device based on the I/O traffic, wherein the communication path is distinct from the dedicated I/O channel”	43
k.	Claim 20	43
l.	Claim 21	43
m.	Claim 22	43
n.	Claim 23	44
F.	Ground 2: Claims 10 and 11 are anticipated by Heil under 35 U.S.C. § 102(b)	44
G.	Ground 3: Claims 1-3 and 6-24 are obvious over Heil and Nakayama under 35 U.S.C. § 103(a)	45
a.	A POSITA Would Have Combined Heil and Nakayama	45
b.	Claim 1	46
i.	Claim 1.pre.....	46
ii.	Claim 1.a	46
iii.	Claim 1.b	48
iv.	Claim 1.c.....	48
v.	Claim 1.d	48
vi.	Claim 1.e.....	48
vii.	Claim 1.f.....	48

c.	Claim 2	49
d.	Claim 3	49
e.	Claim 6	49
f.	Claim 7	49
g.	Claim 8	49
h.	Claim 9	50
i.	Claim 10	50
j.	Claim 11	51
k.	Claim 12	52
l.	Claim 13	52
i.	Claim 13.pre.....	52
ii.	Claim 13.a	52
iii.	Claim 13.b	54
iv.	Claim 13.c.....	54
v.	Claim 13.d	54
vi.	Claim 13.e.....	54
vii.	Claim 13.f.....	54
viii.	Claim 14	56
ix.	Claim 15	56
m.	Claim 16	57
n.	Claim 17	58
i.	Claim 17.a	58
ii.	Claim 17.b	59

o.	Claim 18	60
p.	Claim 19	62
i.	Claim 19.pre.....	62
ii.	Claim 19.a	62
iii.	Claim 19.b	63
iv.	Claim 19.c.....	63
v.	Claim 19.d	63
vi.	Claim 19.e.....	63
vii.	Claim 19.f	63
q.	Claim 20	63
r.	Claim 21	63
s.	Claim 22	63
t.	Claim 23	64
u.	Claim 24	64
i.	Claim 24.pre.....	64
ii.	Claim 24.a	64
iii.	Claim 24.b	67
iv.	Claim 24.c.....	68
v.	Claim 24.d	69
vi.	Claim 24.e.....	70
H.	Ground 4: Claim 4 is obvious over Heil, Nakayama, and Gulick under 35 U.S.C. § 103(a)	73
a.	A POSITA Would Have Combined Heil, Nakayama, and Gulick.....	73

b.	Claim 4 is obvious	74
I.	Ground 5: Claim 5 is obvious over Heil, Nakayama, and Berman under 35 U.S.C. § 103(a)	75
a.	A POSITA Would Have Combined Heil, Nakayama, and Berman	75
b.	Claim 5 is obvious	76
VII.	SECONDARY CONSIDERATIONS	76
VIII.	DISCRETION	77
A.	The <i>Fintiv</i> Factors Strongly Favor Institution	77
B.	Discretionary Denial Under <i>General Plastic</i> Is Unwarranted.....	80
C.	35 U.S.C. § 325(d) Favors Institution	80
IX.	CONCLUSION	80

TABLE OF AUTHORITIES

Page(s)

Cases

<i>Brown v. 3M</i> , 265 F.3d 1349 (Fed. Cir. 2001)	34
<i>Samsung Elecs. Am. Inc. v. Prisia Eng'g Corp.</i> , 948 F.3d 1342 (Fed. Cir. 2020)	20
<i>Thryv, Inc. v. Click-To-Call Techs., LP</i> , 140 S. Ct. 1367 (2020)	79
<i>Williamson v. Citrix Online, LLC</i> , 792 F.3d 1339 (Fed. Cir. 2015)	19

Agency Decisions

<i>Apple Inc. v. Fintiv, Inc.</i> , IPR2020-00019, Paper 11 (Mar. 20, 2020)	77, 79
<i>Samsung Elecs. Co. Ltd. v. Dynamics Inc.</i> , IPR2020-00505, Paper 11 (Aug. 12, 2020)	79
<i>General Plastic Industrial Co., Ltd. v. Canon Kabushiki Kaisha</i> , IPR2016-01357, Paper No. 19 (Sept. 6, 2017)	80

Statutes

35 U.S.C. § 102	<i>passim</i>
35 U.S.C. § 103	<i>passim</i>
35 U.S.C. § 112	19, 20
35 U.S.C. § 132	7, 11
35 U.S.C. § 325	80

Regulations

37 C.F.R. § 42.10	5
-------------------------	---

37 C.F.R. § 42.100	19
--------------------------	----

37 C.F.R. § 42.104	20
--------------------------	----

Patent Procedures

MPEP § 211	12
------------------	----

PETITIONER’S EXHIBIT LIST

<u>Exhibit Number</u>	<u>Description</u>
1001	U.S. Patent No. 10,154,092
1002	File History of U.S. Patent No. 10,154,092 (Application No. 14/997,327)
1003	File History of U.S. Patent No. 6,549,988 (Application No. 09/236,409)
1004	Declaration of Dr. Paul Franzon, Ph.D.
1005	CV of Dr. Paul Franzon, Ph.D.
1006	U.S. Patent No. 6,173,374 (“ <i>Heil</i> ”)
1007	U.S. Patent No. 5,920,893 (“ <i>Nakayama</i> ”)
1008	U.S. Patent No. 5,692,211 (“ <i>Gulick</i> ”)
1009	U.S. Patent No. 6,118,776 (“ <i>Berman</i> ”)

1010	United States District Courts – National Judicial Caseload Profile, https://www.uscourts.gov/sites/default/files/fcms_na_distprofile0331.2022.pdf (accessed September 21, 2022)
1011	Order Staying Case in <i>LS Cloud Storage Technologies, LLC v. Google LLC</i> , 1:22-cv-00853 (W.D. Tex. 2022)
1012	USPTO Memo: Interim Procedure for Discretionary Denials in AIA Post-Grant Proceedings with Parallel District Court Litigation, June 21, 2022
1013	Order Setting Initial Pretrial Conference for 11/3/2022 in <i>LS Cloud Storage Technologies, LLC v. Google LLC</i> , 1:22-cv-00853 (W.D. Tex. 2022)
1014	D. Patterson, G. Gibson, R. Katz, entitled “A Case for Redundant Arrays of Inexpensive Disks (RAID)”. SIGMOD88: International Conference On Management of Data Chicago Illinois USA June 1 - 3, 1988
1015	R. H. Katz, “Network-attached storage systems,” <i>Proceedings Scalable High Performance Computing Conference SHPCC-92.</i> , Williamsburg, VA, USA, 1992, pp. 68-75

1016	Gang Ma and A. L. Narasimha Reddy, “An evaluation of storage systems based on network-attached disks,” Proceedings. 1998 International Conference on Parallel Processing (Cat. No.98EX205), Minneapolis, MN, USA, 1998, pp. 278-285.
1017	Average Time to Trial for Patent Cases before Judge Yeakel, retrieved from www.docketnavigator.com
1018	Infringement Contentions served in <i>LS Cloud Storage Technologies, LLC v. Google LLC</i> , 1:22-cv-00853 (W.D. Tex. 2022)

LISTING OF CHALLENGED CLAIMS

<u>Reference</u>	<u>Claim Limitation</u>
Claim 1	
[1.pre]	An apparatus comprising:
[1.a]	a first interface configured to receive input/output (I/O) traffic from a first host device via a dedicated I/O channel, the I/O traffic comprising a read command;
[1.b]	a second interface configured to receive first data via a network;
[1.c]	a cache memory configured to store second data;
[1.d]	a storage device configured to store third data; and
[1.e]	a processor coupled to the cache memory, the processor coupled to the storage device via a communication path that is distinct from the dedicated I/O channel, the processor configured to access the cache memory during processing of the I/O traffic,

[1.f]	the processor further configured to perform an access operation at the storage device based on the I/O traffic.
-------	---

Claim 2

[2.pre]	The apparatus of claim 1,
[2.a]	wherein the dedicated I/O channel is associated with dedicated throughput that corresponds to the I/O traffic.

Claim 3

[3.pre]	The apparatus of claim 1,
[3.a]	wherein the I/O traffic is distinct from the first data.

Claim 4

[4.pre]	The apparatus of claim 1,
[4.a]	wherein the dedicated I/O channel comprises a small computer system interface (SCSI) channel.

Claim 5

[5.pre]	The apparatus of claim 1,
---------	---------------------------

[5.a]	wherein the second interface comprises an Ethernet interface or an asynchronous transfer mode (ATM) interface.
-------	--

Claim 6	
[6.pre]	The apparatus of claim 1,
[6.a]	wherein the I/O traffic further comprises a write command.

Claim 7	
[7.pre]	The apparatus of claim 1,
[7.a]	wherein the first data comprises a second read request, a write request, or a combination thereof.

Claim 8	
[8.pre]	The apparatus of claim 1,
[8.a]	wherein the network comprises a plurality of interconnected computing devices.

Claim 9	
[9.pre]	The apparatus of claim 1,
[9.a]	wherein the storage device and the dedicated I/O channel are independently accessible,
[9.b]	and wherein the processor is configured to read at least a portion of the second data from the cache memory based on the read command.

Claim 10	
[10.pre]	The apparatus of claim 1,
[10.a]	wherein the processor is further configured to route the read command to the cache memory or to the storage device.

Claim 11	
[11.pre]	The apparatus of claim 1,
[11.a]	configuration manager circuitry configured to route an I/O request included in the I/O traffic to the cache memory, to route the I/O request to the

	storage device, or to deny the I/O request;
[11.b]	front-end circuitry configured to process the I/O request; and
[11.c]	back-end circuitry configured to perform a read operation or a write operation at the storage device based on the I/O request.

Claim 12

[12.pre]	The apparatus of claim 1,
[12.a]	wherein the host device comprises a server.

Claim 13

[13.pre]	A method comprising:
[13.a]	receiving input/output (I/O) traffic from a host device via a dedicated I/O channel at a first interface, the I/O traffic comprising a write command;
[13.b]	receiving first data via a network at a second interface;
[13.c]	storing second data at a cache memory;

[13.d]	storing third data at a storage device;
[13.e]	accessing the cache memory during processing of the I/O traffic; and
[13.f]	performing one or more access operations at the storage device based on the I/O traffic, the one or more access operations utilizing a communication path between a processor and the storage device, the communication path distinct from the dedicated I/O channel.

Claim 14

[14.pre]	The apparatus of claim 13,
[14.a]	wherein the I/O traffic is distinct from the first data.

Claim 15

[15.pre]	The apparatus of claim 13,
[15.a]	further comprising performing a first access operation at the storage device independently of the dedicated I/O channel.

Claim 16	
[16.pre]	The apparatus of claim 13,
[16.a]	wherein the second data is written at the cache memory in response to the write command.

Claim 17	
[17.pre]	The apparatus of claim 13,
[17.a]	further comprising: receiving a read request via the network; and
[17.b]	transmitting at least a portion of the second data from the cache memory via the network responsive to the read request being associated with a cache hit.

Claim 18	
[18.pre]	The apparatus of claim 13,
[18.a]	further comprising storing fourth data at the cache memory, wherein the second data is indicated by the write command, and wherein the fourth data is indicated by a second write

	command received at the second interface.
--	---

Claim 19	
[19.pre]	An apparatus comprising:
[19.a]	a first interface configured to receive input/output (I/O) traffic from a host computer via a dedicated I/O channel, the I/O traffic comprising one or more read commands, one or more write commands, or a combination thereof;
[19.b]	a second interface configured to receive data via a network;
[19.c]	a cache memory;
[19.d]	a storage device; and
[19.e]	a processor coupled via a communication path to the storage device, the processor configured to access the cache memory during processing of the I/O traffic,
[19.f]	the processor further configured to perform one or more access operations at the storage device based on the I/O traffic, wherein the communication

	path is distinct from the dedicated I/O channel.
--	--

Claim 20	
[20.pre]	The apparatus of claim 19,
[20.a]	wherein the dedicated I/O channel is associated with dedicated throughput that corresponds to the I/O traffic.

Claim 21	
[21.pre]	The apparatus of claim 19,
[21.a]	wherein the I/O traffic is distinct from the data.

Claim 22	
[22.pre]	The apparatus of claim 19,
[22.a]	wherein the data corresponds to one or more read requests, one or more write requests, or a combination thereof.

Claim 23	
[23.pre]	The apparatus of claim 19,

[23.a]	wherein the storage device and the dedicated I/O channel are independently accessible.
--------	--

Claim 24	
[24.pre]	An apparatus comprising:
[24.a]	means for receiving input/output (I/O) traffic from a first host device via a dedicated I/O channel, the I/O traffic comprising a write command;
[24.b]	means for receiving data via a network;
[24.c]	means for short-term data storage;
[24.d]	means for long-term data storage; and
[24.e]	means for performing one or more access operations at the means for short-term data storage during processing of the I/O traffic and for performing one or more access operations at the means for long-term data storage based on the I/O traffic, the means for performing distinct from the dedicated I/O channel.

I. INTRODUCTION

Cisco Systems, Inc. (“Cisco”), Microsoft Corporation (“Microsoft”), Amazon.com, Inc., Amazon Web Services, Inc., and Amazon.com Services LLC (“Amazon”) (collectively, “Petitioner”) request *inter partes* review (IPR) of claims 1-24 (the “Challenged Claims”) of U.S. Patent No. 10,154,092 (the “’092 Patent”) assigned to LS Cloud Storage Technologies LLC (“PO”).

The ’092 Patent relates to cached data storage systems and describes an apparatus and method in which host computers provide input to a cache memory or a disk storage device. The host may be a local host that accesses the cache or disk storage via a dedicated I/O channel, or a remote host that accesses the cache or disk storage via a network interface.

As the ’092 Patent admits, such data storage and caching architectures were already well known in the art. “Numerous prior art systems and methods exist for managing cache memory in a data storage system,” and a “typical data processing system generally involves a cached data storage system that connects to local host computers via I/O channels or remote host computers via network links.” Ex.1001, 1:56-58, 1:30-34. Indeed, U.S. Patent 6,173,374 to Heil discloses a data storage and caching architecture identical to that of the ’092 Patent that similarly provides for host access to cache memory and disk storage via either a dedicated I/O channel or a network interface.

Petitioner respectfully requests review and cancellation of the Challenged Claims based on the grounds as set forth below.

II. MANDATORY NOTICES

A. Real Parties-in-Interest

The following are the real parties-in-interest for this petition: Cisco, Microsoft, Amazon.com, Inc., Amazon Web Services, Inc., and Amazon.com Services LLC.

B. Related Matters

PO asserted the '092 Patent against Petitioner and other parties in the following cases in the Waco Division of the Western District of Texas:

- *LS Cloud Storage Technologies, LLC v. Cisco Systems, Inc.*, 6:22-cv-00319 (W.D. Tex. 2022)
 - Dismissed and refiled as *LS Cloud Storage Technologies, LLC v. Cisco Systems, Inc.*, 6:22-cv-00845 (W.D. Tex. 2022) (joint motion to transfer to the Austin Division granted November 22, 2022)
- *LS Cloud Storage Technologies, LLC v. Google LLC*, 6:22-cv-00318 (W.D. Tex. 2022)
 - Transferred to the Austin Division as *LS Cloud Storage Technologies, LLC v. Google LLC*, 1:22-cv-00853 (W.D. Tex. 2022) by agreement of the parties on August 15, 2022

- *LS Cloud Storage Technologies, LLC v. Amazon.com, Inc.*, 6:22-cv-00316 (W.D. Tex. 2022)
 - Transferred to the Austin Division as *LS Cloud Storage Technologies, LLC v. Amazon.com, Inc.*, 1:22-cv-01167 (W.D. Tex. 2022) by agreement of the parties on October 31, 2022
- *LS Cloud Storage Technologies, LLC v. Microsoft Corporation*, 6:22-cv-00321 (W.D. Tex. 2022)
 - Transferred to the Austin Division as *LS Cloud Storage Technologies, LLC v. Microsoft Corporation*, 1:22-cv-00974 (W.D. Tex. 2022)

The '092 Patent is also the subject of the following proceeding:

- *Google LLC v. LS Cloud Storage Technologies, LLC*, IPR2023-00120 (U.S. Patent No. 10,154,092)

Cisco, Microsoft, and Amazon were each first served with a complaint alleging infringement of the '092 Patent on March 28, 2022.

C. Counsel and Service

PETITIONER'S LEAD AND BACK-UP COUNSEL	
Lead Counsel	Brian E. Ferguson Reg. No. 36,801 WINSTON & STRAWN LLP 1901 L Street, N.W. Washington, D.C. 20036

	Telephone: 202-282-5276 Email: beferguson@winston.com <i>Counsel for Petitioners Cisco Systems, Inc. and Microsoft Corporation</i>
Back-Up Counsel	Juan C. Yaquian Reg. No. 70,755 WINSTON & STRAWN LLP 800 Capital Street, Suite 2400 Houston, TX 77002 Telephone: (713) 651-2600 Fax: (713) 651-2700 Email: jyaquian@winston.com <i>Counsel for Petitioners Cisco Systems, Inc. and Microsoft Corporation</i>
Back-Up Counsel	Krishnan Padmanabhan (to seek <i>pro hac vice</i> admission) WINSTON & STRAWN LLP 200 Park Avenue New York City, NY 10166 Telephone: (212) 294-6700 Fax: (212) 294-4700 Email: kpadmanabhan@winston.com <i>Counsel for Petitioner Cisco Systems, Inc.</i>
Back-Up Counsel	Michael Rueckheim (to seek <i>pro hac vice</i> admission) WINSTON & STRAWN LLP 255 Shoreline Drive, Suite 520 Redwood City, CA 94065 Telephone: (650) 858-6500 Fax: (650) 858-6550 Email: mrueckheim@winston.com <i>Counsel for Petitioner Microsoft</i>

	<i>Corporation</i>
Back-Up Counsel	Brian C. Nash Reg. No. 58,105 MORRISON & FOERSTER LLP 701 Brazos Street, Suite 1100 Austin, TX 78701 Telephone: (737) 309-0700 Email: BNash@mofo.com <i>Counsel for Petitioners Amazon.com, Inc., Amazon Web Services, Inc., and Amazon.com Services LLC</i>

Petitioner consents to service by electronic mail at the following addresses:

beferguson@winston.com;

jyaquian@winston.com;

kpadmanabhan@winston.com;

mrueckheim@winston.com;

WS-Cisco-LSCloud@winston.com;

Winston-Microsoft-LSCloud@winston.com;

BNash@mofo.com; and

Amazon-LSCloud@mofo.com.

D. Power of Attorney

Power of attorneys are filed herewith in accordance with 37 C.F.R. § 42.10(b).

E. Fees

Petitioner concurrently electronically submits the required fees for this Petition.

III. REQUIREMENTS FOR *INTER PARTES* REVIEW

A. Grounds for Standing

Petitioner certifies that the '092 Patent is available for IPR, and that Petitioner is not barred or estopped from requesting this IPR.

B. Identification of Challenge and Relief Requested

Petitioner requests cancellation of the Challenged Claims on the following grounds:

Ground 1: Anticipation of claims 1-3, 7-12, and 19-23 by U.S. Patent 6,173,374 ("*Heil*") issued January 9, 2001, under 35 U.S.C. § 102(e);

Ground 2: Anticipation of claims 10 and 11 by Heil under 35 U.S.C. § 102(b);

Ground 3: Obviousness of claims 1-3 and 6-24 based on Heil and U.S. Patent 5,920,893 ("*Nakayama*") issued July 6, 1999, under 35 U.S.C. § 103(a);

Ground 4: Obviousness of claim 4 based on Heil, Nakayama, and U.S. Patent 5,692,211 ("*Gulick*") issued November 25, 1997, under 35 U.S.C. § 103(a);
and

Ground 5: Obviousness of claim 5 based on Heil, Nakayama, and U.S. Patent 6,118,776 (“*Berman*”) issued September 12, 2000, under 35 U.S.C. § 103(a).

Each ground is supported by the Declaration of Dr. Paul Franzon, Ph.D., an expert in the field of computer networking and cache management. None of the prior art used in any of the grounds was cited during prosecution.

As explained in Section V, the ’092 Patent claims priority through a chain of applications, the earliest of which (U.S. Application No. 09/236,409) has a filing date of January 22, 1999. The written description support for claims 10 and 11 of the ’092 Patent, however, was not included in the original specification of the ’409 Application. Rather, it was introduced as new matter during the prosecution of the ’409 Application. Because “[n]o amendment shall introduce new matter into the disclosure of the invention,” claims 10 and 11 are not entitled to claim the January 22, 1999 priority date of the ’409 Application. 35 U.S.C. § 132. Instead, claims 10 and 11 have a priority date no earlier than the March 5, 2003 filing date of U.S. Application No. 10/382,016, the first child application of the ’409 Application in the ’092 Patent’s priority chain having the ’409 Application’s new matter in its original specification.

For all claims, Heil, Nakayama, and Berman are prior art pursuant to at least (pre-AIA) 35 U.S.C. § 102(e). Gulick is prior art pursuant to (pre-AIA) 35 U.S.C.

§ 102(b). To the extent that the '409 Application's new matter does not entitle claims 10 and 11 to the '409 Application's January 22, 1999 priority date, Heil and Nakayama further qualify as prior art under (pre-AIA) 35 U.S.C. § 102(b).

IV. THE '092 PATENT

The '092 Patent relates to cached data storage systems and describes an apparatus and method in which host computers provide input to a cache memory or a disk storage device. Ex.1001, abstract. Figure 1 shows the configuration of the preferred embodiment in which hosts 111 and 112 are connected via I/O channels 121 and 122 respectively to a PC data storage system 131. Ex.1001, 5:9-16. PC data storage system 131 includes a cache memory managed by a cache manager as well as disks 161 and 162. Ex.1001, 5:25-34. When the host accesses the cache or storage disk of a data storage system via an I/O channel, the '092 Patent considers the host to be a local host. Ex.1001, 1:30-36.





FIG. 1

Ex.1001, Fig. 1.

Alternatively, hosts 111 and 112 may access a remote data storage system 133 via a network connection 153. Ex.1001, 5:9-14, 25-34. Remote PC storage system 133 also includes a cache memory managed by a cache manager as well as disks 181 and 182. Ex.1001, 5:25-34. When the host accesses the cache or storage disk of a data storage system via a network connection, the '092 Patent considers the host to be a remote host. Ex.1001, 1:30-36.

As shown in Figure 3, the cache manager 320 may comprise a cache memory 322. Ex.1001, 5:48-61. Cache manager 320 is connected to hosts 111 and

112 via a front-end module 310 and is connected to disks 161 and 162 via a back-end module 330. Ex.1001, 5:48-61. Front-end module 310 is an interface for the I/O channel and may be a small computer system interface (“SCSI”) for a SCSI channel, which the ’092 Patent discloses is a “well-established” technology. Ex.1001, 5:62-67, 5:35-37. Figure 3 shows the front-end module 310 host connection as a distinct communication path from back-end module 330 disk connection:



Ex.1001, Fig. 3.

V. PROSECUTION HISTORY

The ’092 Patent issued on December 11, 2018, from U.S. Application No. 14/997,327, filed January 15, 2016, and claims priority through a chain of continuation and divisional applications to U.S. Application No. 09/236,409, filed January 22, 1999, and issued as U.S. Pat. No. 6,549,988. The specification of the ’092 Patent, however, differs from the originally-filed specification of the parent

'409 Application because on April 25, 2002, during prosecution of the '409 parent, the applicant added by amendment new disclosures corresponding to col. 6 ln. 25-49 of the '092 Patent describing the use of a volume access table to instruct the configuration manager to route I/O requests. Ex.1003, 132. This amended disclosure is unsupported by the '409 Application's original specification and constitutes new matter added during prosecution, in contravention of 35 U.S.C. § 132's prohibition against new matter. 35 U.S.C. § 132 ("No amendment shall introduce new matter into the disclosure of the invention."). Critically, this amended disclosure provides the only potential support in the '092 Patent for routing the read command to a cache or other storage as recited in claims 10 and 11.

New Matter Introduced During Prosecution of the '409 Application	Volume access table 450 (see FIG. 4) contains a mapping between hosts and volumes specifying an access mode value. If the access mode is set neither to shared nor exclusive, the configuration manager forwards I/O requests directly to disk. In addition to the access mode, the volume access table may contain other values that help the configuration manager and improve performance of the data storage system. In another embodiment of this application, shown in FIG. 5, the volume access table includes a translation module for a given host to facilitate volume mapping. The
---	--

	translation module is a dynamically loadable library that can be changed, compiled and linked at run-time.
'092 Patent, Claim 10	10. The apparatus of claim 1, wherein the processor is further configured to route the read command to the cache memory or to the storage device.
'092 Patent, Claim 11	11. The apparatus of claim 1, further comprising: configuration manager circuitry configured to route an I/O request included in the I/O traffic to the cache memory, to route the I/O request to the storage device, or to deny the I/O request; front-end circuitry configured to process the I/O request; and back-end circuitry configured to perform a read operation or a write operation at the storage device based on the I/O request.

Claims 10 and 11 are not entitled to the January 22, 1999, priority date of the '409 Application but are instead entitled to at most the March 5, 2003 filing date of U.S. Application No. 10/382,016, the first child application of the '409 Application in the '092 Patent's priority chain having the '409 Application's new matter in its original specification. MPEP § 211.

The Office issued one Non-Final Rejection during the course of prosecuting the '327 Application. Ex.1002, 97–116. On February 7, 2018, the Office rejected

all of the then-pending claims of the '327 Application after finding all elements present in the prior art. *Id.* The applicant did not present any arguments distinguishing the pending claims from the prior art but instead amended the pending claims with additional limitations requiring that the processor be coupled to the storage device via a communication path that is distinct from the I/O channel. Ex.1002, 78-82. A Notice of Allowance followed on May 16, 2018, and the '327 Application issued as the '092 Patent. Ex.1002, 7. None of the prior art references used in this Petition, which as explained disclose the amended and other limitations, were discussed or of record during the '327 Application's prosecution.

VI. THE CHALLENGED CLAIMS ARE UNPATENTABLE

A. State of the Art

As the '092 Patent's prosecution history and specification confirm, data storage and caching architectures like those disclosed by the '092 Patent were already well known in the art. "Numerous prior art systems and methods exist for managing cache memory in a data storage system," and a "typical data processing system generally involves a cached data storage system that connects to local host computers via I/O channels or remote host computers via network links." Ex.1001, 1:56-58, 1:30-34. Even in the early 90s, storage systems in which hosts are connected to peripheral devices through dedicated I/O channels such as SCSI interfaces, one of the main claimed features of the '092 Patent, were well-known.

Ex.1004, ¶¶37-50. During this time, storage systems transitioned from node-based storage architectures that mainly utilized dedicated I/O channels to network-based architectures that also connected hosts to multiple storage servers over a network.

Id. As a result of this paradigm shift, hosts and storage devices were frequently connected together through a network, another claimed feature of the '092 Patent.

Id. Similarly, providing caching for network-based storage was also well-known in the art. Ex.1004, ¶¶48-50.


Indeed, the Office determined during prosecution that all elements of the Challenged Claims, barring a dedicated I/O channel distinct from the communication path from the processor to the storage device, were known in the art. Ex.1002, 105-116. PO did not contest any of these findings. Ex.1002, 80-82. Nonetheless, even the dedicated I/O channel being distinct from the communication path from the processor to the storage device was known in the art, as shown by Heil. Ex.1006; Ex.1004, ¶¶68-73; Ex.1007, 4:40-5:9; Ex.1008, 5:37-51; Ex.1009, 12:9-11.

B. Overview of the Prior Art

a. Heil

U.S. Patent 6,173,374 to Heil et al. was filed February 11, 1998 and issued January 9, 2001. Like the '092 Patent, Heil is directed to data storage systems and describes an apparatus and method in which host computers may store files in a

cache or disk storage. Ex.1006, abstract, 1:9-14, 6:34-47. Heil's Figure 1 shows a configuration remarkably like that shown by Figure 1 of the '092 Patent. In Heil's Figure 1, a host system having a CPU 100 is connected via a PCI bus 116.5 to a host bus adapter ("HBA") 117. Ex.1006, 6:34-47. Host bus adapter 117 includes a cache memory 117.2 and local memory 116 and is further connected to local drives 118 via SCSI or fiber channel arbitrated loop ("FCAL") interfaces 117.8 and 117.9. Ex.1006, 7:9-15, 7:51-56. The host may process I/O requests locally by accessing the host bus adapter 117 via PCI bus 116.5. Ex.1006, 4:51-57, 6:34-47.



Ex.1006, Figure 1.

Alternatively, the host system (shown in node 1) may access another, remote data storage system (shown in node n) via a network connection 121. Ex.1006, 8:8-22. The remote data storage system of node n also includes a cache memory 126.2 and memory 127 and is further connected to drives 123 via SCSI or FCAL interfaces 122.3 and 122.4. Ex.1006, 8:61-9:10.


The host bus adapter 117 has a CPU 117.1 which may store data either to cache memory 117.2 or disk drives 118. Ex.1006, 7:21-34. The host bus adapter CPU 117.1 is connected to host CPU 100 via the PCI bus 116.5 and is connected to drives 118 via SCSI or FCAL interfaces 117.8 and 117.9. Ex.1006, 9:34-47. Heil's Figure 1 shows the PCI bus 116.5 is a distinct communication path from SCSI or FCAL interfaces 117.8 and 117.9. Ex.1006, Fig. 1.

b. Nakayama

U.S. Patent 5,920,893 to Nakayama et al. was filed June 2, 1997 and issued July 6, 1999. Like the '092 Patent, Nakayama is directed to cached data storage systems and describes an apparatus and method in which host computers provide input to a cache memory or a storage drive. Ex.1007, abstract, 4:52-65.

Nakayama's Figure 3 shows a host computer 300 connected via channel 301 to a disk control 302 having a cache memory 309. Ex.1007, 4:40-5:9. A host computer 303 may be connected to channel control 302 by a SCSI bus 304. Ex.1007, 4:56-60.

FIG. 3



Disk control 302 is connected to drives 315 and 316 via a drive interface 313.

Ex.1007, 4:56-60. Nakayama's Figure 3 shows that channel control 301 or SCSI bus 304 is a distinct communication path from drive interfaces 313 and 314.

c. Gulick

U.S. Patent 5,692,211 to Gulick et al. was filed September 11, 1995 and issued November 25, 1997. Gulick discloses that it is well known that PCI buses may be connected to SCSI adapters. Ex.1008, 1:17-25. Gulick describes a “host/PCI/cache bridge [that] interfaces to a peripheral expansion bus, preferably a peripheral component interconnect (PCI) bus.” Ex.1008, 5:37-51. A SCSI adapter may be coupled to the PCI bus to interface the bridge with various SCSI devices, such as a CD-ROM drive and a tape drive, as desired and well known in the art. Ex.1008, 5:37-51.

d. Berman

U.S. Patent 6,118,776 to Berman was filed August 7, 1997 and issued September 12, 2000. Berman is directed to networking systems utilizing fiber optic technology, specifically Fiber Channels. Ex.1009, abstract, 1:34-47. Berman suggests that it would be advantageous to interface Fiber Channels with “legacy networks” such as ethernet and discloses systems such as routers that connect Fiber Channel to ethernet. Ex.1009, 12:9-11, claims 13-15, 3:23-30.

C. Level of Ordinary Skill in the Art

A person of ordinary skill in the art (“POSITA”) as of the earliest potential priority date of the '092 Patent would have had a bachelor’s degree in Computer Science, Computer Engineering, or Electrical Engineering, or an equivalent course of study, as well as 2 or more years of academic or industry experience in the field

of computer networking, computer architecture and storage management. Ex.1004, ¶¶34-36. More education can supplement practical experience, and vice versa. *Id.*

D. Claim Construction

Claim terms are construed in accordance with their ordinary and customary meaning as understood by a POSITA and the prosecution history pertaining to the patent. 37 C.F.R. § 42.100(b). Petitioner is unaware of any “prior claim construction determination” related to the ’092 patent.

For the purposes of this Petition, Petitioner interprets all terms of claims 1-23 in accordance with their ordinary and customary meaning. Petitioner submits that elements of claim 24, including “means for receiving input/output (I/O) traffic from a first host device via a dedicated I/O channel, the I/O traffic comprising a write command”; “means for receiving data via a network”; “means for short-term data storage”; “means for long-term data storage”; and “means for performing one or more access operations at the means for short-term data storage during processing of the I/O traffic and for performing one or more access operations at the means for long-term data storage based on the I/O traffic” are means-plus-function limitations governed by 35 U.S.C. § 112. Use of the word “means” as a claim element creates a presumption that § 112, para. 6 applies. *Williamson v. Citrix Online, LLC*, 792 F.3d 1339, 1348 (Fed. Cir. 2015).

For the purposes of this proceeding, Petitioner interprets these terms as follows¹:

- “means for receiving input/output (I/O) traffic from a first host device via a dedicated I/O channel, the I/O traffic comprising a write command”:
 - Recited function: “receiving input/output (I/O) traffic from a first host device via a dedicated I/O channel”
 - Disclosed structure: A front-end module that receives data from hosts via I/O channels, and equivalents. *See* Ex.1001, col. 5:47-67; 6:13-24, Fig. 3 label 310.
- “means for receiving data via a network”:

¹ Because challenges under 35 U.S.C. § 112 are not available in IPR proceedings, Petitioner does not raise any section 112 challenge or concede that the claim limitations satisfy section 112. Petitioner provides the identification of disclosures in the specification that appear to be most related to performance of the claimed function only to satisfy 37 C.F.R. § 42.104(b) and expressly reserves the right to raise other challenges as permitted in the parallel district court litigation. The existence of section 112 issues does not preclude institution of this IPR. *Samsung Elecs. Am. Inc. v. Prisia Eng'g Corp.*, 948 F.3d 1342, 1352-53, 55 (Fed. Cir. 2020).

- Recited function: “receiving data via a network”
- Disclosed structure: A network adapter, and equivalents. See Ex.1001, abstract, col. 3:44-47, 4:1-2.
- “means for short-term data storage”:
 - Recited function: “short-term data storage”
 - Disclosed structure: A cache memory, and equivalents. See Ex.1001, col. 1:36-40, 6:1-8.
- “means for long-term data storage”:
 - Recited function: “long-term data storage”
 - Disclosed structure: A back-end storage device or local storage device. See Ex.1001, 5:9-34, 6:1-4.
- “means for performing one or more access operations at the means for short-term data storage during processing of the I/O traffic and for performing one or more access operations at the means for long-term data storage based on the I/O traffic”:
 - Recited function: “performing one or more access operations at the means for short-term data storage during processing of the I/O traffic and for performing one or more access operations at the means for long-term data storage based on the I/O traffic”

- Disclosed structure: A cache manager and a configuration manager, and equivalents. *See* Ex.1001, 6:9-32 (configuration manager for operations to local disks), 6:50-8:46 (cache manager for operations to cache, Figs. 4-4F).

To the extent PO makes any arguments that any of the foregoing are not means-plus-function terms or have a different function or disclosed structure, Petitioner reserves the right to further address them at that time.

E. Ground 1: Claims 1-3, 7-12, and 19-23 are anticipated by Heil under 35 U.S.C. § 102(e)

a. Claim 1

The claimed apparatus is anticipated by Heil. Ex.1004, ¶¶66-95.

i. Claim 1.pre

If the Board determines that the preamble is limiting, Heil discloses a data storage and communication apparatus and teaches the preamble. Ex.1006, 1:9-14 (“This invention relates to distribution of storage data within computer nodes of a clustered computing environment, and more particularly, to *an apparatus and method for I/O request shipping operable using peer-to-peer communications within host bus adapters of the clustered systems.*”); Ex.1004, ¶67.


ii. Claim 1.a: “a first interface configured to receive input/output (I/O) traffic from a first host device via a dedicated I/O channel, the I/O traffic comprising a read command”

Heil discloses this element. Ex.1004, ¶¶68-73. Heil describes a first interface that receives I/O traffic from a host device in the form of an interface for a PCI bus bridge 115, PCI bus 116.5, or Front End I/F 102 for receiving input from a host system (CPU 100 of Node 1):

Node 1 150 includes a host system connected to a HBA 117 via a Host-to-Peripheral Component Interconnect (hereinafter referred to as “PCI”) bus bridge 115 and PCI bus 116.5.

Ex.1006, 6:35-40; Ex.1004, ¶69.

Heil’s Figure 1 shows that the first interface (either PCI bus bridge 115, PCI bus 116.5, or Front End I/F 102) is between HBA 117 and Node 1’s host system (a first host device):



Ex.1006, Fig. 1; Ex.1004, ¶70.

Heil further discloses that communications between the Node 1 host and HBA 117 comprise I/O requests. Ex.1006, 7:36-37 (“As shown in FIG. 1, one HBA services the I/O requests from CPU 1 and CPU 2 100.”); Ex.1004, ¶71. The I/O requests are instructions to retrieve files from caches or local storage devices and are read commands. Ex.1006, 2:59-65 (“To retrieve files from a remote resource, software in the host intercepts file or block level I/O requests and determines whether the particular file or block is stored locally or remotely. If local, the software passes the request down to the local file system or block I/O driver.”), 4:63-69 (“Data returned from the HBA performing the remote processing (i.e., a read request) is passed by the initiating HBA to the requesting host.”), 12:46

(“The remote I/O blocks are sent from the HBA performing the remote processing (e.g., a read request) to the initiating HBA 460.”); Ex.1004, ¶71 (“[A] POSITA would understand Heil’s I/O requests to be read commands” because they instruct the processor to retrieve data from cache or drive storage).


The channel used by the PCI bus bridge 115 and PCI bus 116.5 is a dedicated I/O channel because it is a direct, dedicated connection between the host and the HBA that carries I/O traffic between them. Ex.1006, 6:35-40, 7:36-37; Ex.1004, ¶72.

iii. Claim 1.b: “a second interface configured to receive first data via a network”

Heil discloses this element. Ex.1004, ¶¶74-78. Heil describes a second interface that receives first data via a network in the form of an interface between fibre channel chip 120 and network fibre channel backbone 121 that receives data from the fibre network (i.e., fibre channel backbone 121):

Fibre Channel Chip 120 connects Node 1 150 to the rest of the clustered environment via the network Fibre Channel backbone 121. The Fibre Channel backbone is a high-speed communication medium and provides numerous advantages[] Fibre Channel makes use of unique addresses to connect to other processors or peripherals. The Fibre Channel port only needs to manage the point-to-point connection between the device and the backbone fabric.

Ex.1006, 8:24-38; 4:26-35 (“Similarly, HBAs in different nodes can communicate as peers via[]a Fibre Channel backbone that interconnects the HBAs.”). The data transmitted over the Fibre Channel backbone and received by Fibre Channel Chip 120 comprises a first data. Ex.1004, ¶75. Heil’s Figure 1 shows, in addition to the interface at the PCI bus bridge 115, the PCI bus 116.5, or the Front End I/F 102(a first interface), that HBA 117 has an interface at fibre channel chip 120 (a second interface) that connects to fibre network 121:




Ex.1006, Fig. 1; Ex.1004, ¶¶76-78.

The data transmitted over the fibre channel backbone and received by fibre channel chip 120 corresponds to the first data. Ex.1004, ¶77.

iv. **Claim 1.c: “a cache memory configured to store second data”**

Heil discloses a cache memory for storing second data in the form of cache


117.2:



Ex.1006, Fig. 1, 7:25-26 (“To improve I/O performance, CPU 117.1 temporarily caches disk data in cache memory 117.2 that is frequently accessed.”); Ex.1004, ¶¶79-82. The data cached within cache memory comprises a second data. Ex.1004, ¶81.

v. **Claim 1.d: “a storage device configured to store third data”**


Heil discloses a storage device memory for storing third data in the form of local drives 118 and 123:



Ex.1006, Fig. 1, 10:14-20 (“The front end interface 195 is a PCI-to-PCI Bus bridge that connects, via PCI bus 125.2, various peripheral I/O devices, such as the local drives 123, to centralized PCI bus 127.5. Front-end interface 195 is connected to SCSI or FCAL interfaces 122.3, 122.4 via PCI bus 125.2. These interfaces can be either SCSI or FCAL depending upon the bus utilized to connect the local drives 123.”), 7:26-29 (“Later, the CPU 117.1 stores the data permanently in the slower disk drive devices 118.”); Ex.1004, ¶84. The data stored within the local drives comprises a third data. Ex.1004, ¶85.


- vi. **Claim 1.e: “a processor coupled to the cache memory, the processor coupled to the storage device via a communication path that is distinct from the dedicated I/O channel, the processor configured to access the cache memory during processing of the I/O traffic”**

Heil discloses a processor (CPU 117.1) coupled to the cache memory (cache 117.2) and to a storage device (local drives 118):



Ex.1006, Fig. 1, 7:9-11 (“HBA 117 also includes a CPU 117.1, local memory 116, and cache memory 117.2, which are connected by internal processor bus 117.3.”); 7:25-28 (“To improve I/O performance, CPU 117.1 temporarily caches disk data in cache memory 117.2 that is frequently accessed. Later, the CPU 117.1 stores the data permanently in the slower disk drive devices 118.”); Ex.1004, ¶¶87-92.

Heil’s Figure 1 shows that the communication path from the processor to the storage device—i.e., from the CPU 117.1 to the local drives 118 via a SCSI interface, annotated below in **red**—is distinct from the dedicated I/O channel—i.e., from the Node 1 host system to HBA 117 via a PCI bus, annotated in **green**:



Ex.1006, Fig. 1; Ex.1004, ¶¶88-89.

Heil further discloses that the processor may access cache memory while processing I/O traffic, including as a result of processing the I/O command transmitted in the I/O traffic. Ex.1006, 5:25-28 (“Performance is further enhanced in that each intelligent HBA can provide disk caching of the local disks under its control, and can service hits out of this cache regardless of whether the requesting host is local or remote.”), 7:24-28 (“To improve I/O performance, CPU 117.1 temporarily caches disk data in cache memory 117.2 that is frequently accessed.”); Ex.1004, ¶90 (“[A] POSITA would understand from these disclosures that the cache memory would be accessed to retrieve data responsive to the I/O request. The processor (CPU 117.1 of the HBA) may access cache memory while

processing I/O traffic, including after processing an I/O command transmitted in the I/O traffic.”). This is further illustrated by Heil’s Figure 3. Once an I/O request is made by a CPU 100 (Ex.1006, 7:36-40), the HBA may determine that the request may be serviced locally (Ex.1006, 11:47-54). While Fig. 3 suggests responsive data is retrieved from the local disks, Heil discloses that “frequently accessed” data is cached in cache memory and may be provided in response to “improve I/O performance.” Ex.1006, 7:21-34, 12:1-4; Ex.1004, ¶¶91.

vii. Claim 1.f: “the processor further configured to perform an access operation at the storage device based on the I/O traffic”


Heil discloses that the processor (of the HBA) may access the storage device (the local drive) based on the I/O traffic. In response to an I/O request by CPU 100, the HBA may determine that the request may be serviced locally and retrieve the response data from the local disks. Ex.1006, 7:36-37 (“As shown in FIG. 1, one HBA services the I/O requests from CPU 1 and CPU 2 100.”), 11:45-53 (“After receiving a block I/O request 400, the 110 redirector has the means to search the directory and determining means to locate the local or remote disk drives that are storing the I/O requested blocks 410. If the location of the blocks is local 420, then the blocks are retrieved from the respective local disk drive.”); Ex.1004, ¶¶93-95.

b. Claim 2

Heil discloses the dedicated I/O channel is associated with dedicated throughput that corresponds to the I/O traffic. Ex.1004, ¶¶96-99. Heil discloses that an HBA “in general is a device which adapts (connects) a host computer system to an I/O device. Signals associated with the bus of the host computer system are adapted for exchange with a bus specific to the I/O device.” Ex.1006, 4:3-8. HBA 117 receives a request from a host via the dedicated I/O channel. Ex.1006, 4:50-51 (“In the preferred embodiment, an HBA receives a block I/O request from the host and searches through the directory in HBA memory to determine whether the block can be found locally or remotely.”). The dedicated I/O channel is associated with dedicated throughput that corresponds to the I/O traffic because the data passing through the I/O channel corresponds to the I/O traffic. Ex.1004, ¶99.

c. Claim 3

Heil discloses “the I/O traffic is distinct from the first data.” Ex.1004, ¶¶100-107. As shown in Heil’s Figure 1, the I/O traffic is received by HBA 117 via the PCI bus from the Node 1 host:



Ex.1004, ¶¶104-106 (“**The path for I/O traffic from Heil’s host server (including CPU 100) to its HBA is annotated in red. The network interface comprising fibre channel chip 120 receiving first data from the fibre channel network is annotated in green.**”); *see also* Ex.1006, 6:34-47, 7:36-50, 9:11-32, 9:49-59. Fig. 1 shows that the path for I/O is distinct from the first data which is network traffic received by fiber channel chip 120 from fiber channel network 121.

Ex.1004, ¶107 (“Figure 1 shows that the I/O channel is transmitted through a different communication path than the fibre channel network traffic.”); *see also* Ex.1006, 6:65-7:8, 8:23-38, 9:11-25, 9:49-59.

d. Claim 7

Heil discloses the first data received via a network may comprise a second read request. Ex.1004, ¶¶108-113. *Brown v. 3M*, 265 F.3d 1349, 1351 (Fed. Cir. 2001) (“When a claim covers several structures [as] alternatives, the claim is deemed anticipated if any of the structures[] within the scope of the claim is known in the prior art.”). Heil discloses that a requesting HBA may send read requests via the fibre network to a remote HBA having data stored in its disk. Ex.1006, 11:5456 (“If the requested I/O block is stored in a remote disk, the initiating HBA ships the block I/O request to the peer HBA that the initiating HBA has determined can perform the requisite processing 450.”), 12:4-6 (“The remote I/O blocks are sent from the HBA performing the remote processing (e.g., a read request) to the initiating HBA 460.”); Ex.1004, ¶111 (“[I]f the requested data is not found at the local HBA, then the local HBA transmits the I/O request via the fibre channel network to remote HBAs to service the request. An I/O request transmitted between two HBAs is transmitted via the fibre channel network and is distinct from an I/O request transmitted from a host system to a local HBA via the PCI bus.”). Because these read requests are different than the requests sent by local HBAs via the dedicated I/O channel, the read requests of remote HBAs constitute a second read request. Ex.1004, ¶112 (“A POSITA would understand that the second I/O request corresponds to the claimed second read request. If data responsive to

the second I/O request is not found at the local HBA, and the local HBA ships that request to a remote HBA, then the remote HBA is the apparatus that has received the second read request via the network.”). Alternatively, Heil discloses “future requests” which may also constitute a second read request. Ex.1006, 12:1-4, Ex.1004, ¶112 (“Heil discloses that an HBA “can cache blocks of data retrieved from its local disk, to service future local or remotely generated requests for these blocks.” (Ex.1006, 12:1-4.) Based on this disclosure, a POSITA would understand that more than one I/O request may be issued by a host system—Heil explicitly contemplates that the host system may issue a second I/O request to the local HBA even after the host system issues a first I/O request to the local HBA.”).

e. Claim 8

Heil discloses “the network comprises a plurality of interconnected computing devices.” Ex.1004, ¶¶114-119. Heil discloses a network comprising a plurality of interconnected HBA nodes, which are computer devices:

Node 150, Nodes 152, and Node N 151 are connected via a network communications medium. In the preferred embodiment, a Fibre Channel backbone 121 allows each CPU to be connected to peer CPUs. One skilled in the art will recognize all node-to-node inter-processor communication can occur on this Fibre Channel backbone 121.

Ex.1006, 8:12-24; see also Ex.1006, 6:35-44, 7:9-12, 4:3-5; Ex.1004, ¶¶117-119 (“Heil describes a network of multiple nodes interconnected via a fibre channel...


each node comprises a host system... Heil describes the host system and the HBA as computers or devices.”).

f. Claim 9

Heil discloses claim 9. Ex.1004, ¶¶120-127.

i. Claim 9.a: “the storage device and the dedicated I/O channel are independently accessible”

Heil discloses the storage device and dedicated I/O channel are independently accessible. Ex.1004, ¶¶121-124. Heil’s Fig. 1 shows two independent pathways (1) from the Node 1 host via the PCI bus to the processor (the dedicated I/O channel) and (2) from a remote Node N host via the fiber network and fibre channel chip 120 to the local drives via a SCSI or FCAL interface:



Ex.1004, ¶122.

In addition, Heil discloses that the storage device may be accessed based on a network request from a requestor HBA on the network, rather than the local HBA. Ex.1006, 11:54-12:6; Ex.1004, ¶123. Because requestor HBAs on the network use the fibre channel pathway rather than the PCI bus pathway, access operations by a requestor HBA on the network to the storage device is independent of I/O operations from the local HBA to the storage device. *Id.*

ii. Claim 9.b: “wherein the processor is configured to read at least a portion of the second data from the cache memory based on the read command”

Heil discloses that the processor may retrieve data responsive to a read request from the cache memory based on the read command and discloses this element. Ex.1006, 5:25-33 (“Performance is further enhanced in that each intelligent HBA can provide disk caching of the local disks under its control, and can service hits out of this cache regardless of whether the requesting host is local or remote.”), 7:24-30 (“To improve I/O performance, CPU 117.1 temporarily caches disk data in cache memory 117.2 that is frequently accessed.”), 11:66-12.1 (“The remote HBA retrieves the requested blocks of data from its HBA cache, or if the blocks of data are not within cache, then directly from the remote HBA’s local disk.”), *see also id.* at 12:1-4; Ex.1004, ¶125-127.

g. Claim 10


Heil discloses this element. Ex.1004, ¶¶128-131. Heil describes that the processor reads the cache in response to a read command, and if the data is not found in the cache, the processor then reads from the storage device. Ex.1006, 5:25-33 (“Performance is further enhanced in that each intelligent HBA can provide disk caching of the local disks under its control, and can service hits out of this cache regardless of whether the requesting host is local or remote.”), 7:24-30 (“To improve I/O performance, CPU 117.1 temporarily caches disk data in cache memory 117.2 that is frequently accessed.”), 11:66-12.1 (“The remote HBA retrieves the requested blocks of data from its HBA cache, or if the blocks of data are not within cache, then directly from the remote HBA’s local disk.”). This routes the read command to the cache or the storage device, depending on where the requested data may be found, because reading the cache or disk necessarily requires the cache or disk to receive the read command. Ex.1004, ¶131 (“By using the cache to store frequently requested data, Heil recognizes that I/O performance may be improved by initially attempting to retrieve requested data from the cache, and if the data is not found in the cache, then retrieving the requested data from disk storage...In my opinion, a POSITA would understand that this process routes the read command to the cache or the storage device, depending on where the

requested data may be found, because reading the cache or disk necessarily requires the cache or disk to receive the read command.”).

h. Claim 11

i. Claim 11.a: “configuration manager circuitry configured to route an I/O request included in the I/O traffic to the cache memory, to route the I/O request to the storage device, or to deny the I/O request”

Heil discloses this element. Ex.1004, ¶¶133-137. Heil discloses that an I/O redirector 240 as shown in Figure 2 facilitates retrieval of data related to an I/O request based on the mapping of the data within every storage subsystem of the server clusters having HBAs. Ex.1006, 12:19-34, 5:1-6; Ex.1004, ¶134.



Ex.1006, Fig. 2. Heil further discloses that caching enhances I/O performance and caches are part of the storage subsystem of the HBAs. Ex.1006, 5:25-33

(“Performance is further enhanced in that each intelligent HBA can provide disk caching of the local disks under its control, and can service hits out of this cache regardless of whether the requesting host is local or remote.”), 7:24-30 (“To improve I/O performance, CPU 117.1 temporarily caches disk data in cache memory 117.2 that is frequently accessed.”); Ex.1004, ¶135. When servicing an I/O request, the HBA looks to its cache and, if the requested data is not found in the HBA’s cache, the HBA then accesses the storage device. Ex.1006, 11:66-12:1 (“The remote HBA retrieves the requested blocks of data from its HBA cache, or if the blocks of data are not within cache, then directly from the remote HBA’s local disk.”); Ex.1004, ¶136-137. Heil discloses element 11.a. *Id.*, (“[T]his operation routes the I/O request to cache memory or to the storage device, depending on where the requested data is found, because reading the cache or disk necessarily requires the cache or disk to receive the read command.”).

ii. Claim 11.b: “front-end circuitry configured to process the I/O request”

Heil discloses this element. Ex.1004, ¶¶138-139. Heil discloses that the PCI bus 117 includes front-end interface 171 to process the I/O request. Ex.1006, 9:43-48 (“The front-end interface 171 is a PCI-to-PCI Bus bridge that connects, via PCI bus 117.4, various peripheral I/O devices, such as the local drives 118, to centralized PCI bus 116.5. The HBA interface to the local drives 117.8 and 117.9

can be SCSI or FCAL depending upon the bus supporting the local drives.”);

Ex.1004, ¶¶138-139.

iii. Claim 11.c: “back-end circuitry configured to perform a read operation or a write operation at the storage device based on the I/O request”

Heil discloses this element. Ex.1004, ¶¶140-141. Heil discloses HBA interfaces 117.8 and 117.9 which are back-end circuitry that directs the local drives 118 to perform a read operation at the storage device based on the I/O request. Ex.1006, 7:51-56 (“The HBA interfaces and controls the local drives 118 via HBA interface 117.8 and 117.9. These interfaces are[]small computer system interface (SCSI) or Fibre Channel Arbitrated Loop (FCAL). A person skilled in the art will recognize the interface utilized depends upon the bus used to connect the plurality of local disks 118.”); Ex.1004, ¶¶140-141.

i. Claim 12

Heil discloses “the host device comprises a server.” Ex.1004, ¶¶142-146. Specifically, Heil discloses that its CPUs 100 are server clusters (i.e., a server) which provide I/O requests to and are serviced by HBAs:

The present invention retrieves data across independent computer nodes of a server cluster by providing for I/O shipping of block level requests to peer intelligent host-bus adapters (hereinafter referred to as HBA).

Ex.1006, abstract; Ex.1004, ¶¶142-146.

j. Claim 19

The claimed apparatus is anticipated by Heil. Ex.1004, ¶¶147-153.

i. Claim 19.pre

See element 1.pre (Section VI.E.a.i).

ii. Claim 19.a: “a first interface configured to receive input/output (I/O) traffic from a host computer via a dedicated I/O channel, the I/O traffic comprising one or more read commands, one or more write commands, or a combination thereof”

See element 1.a (Section VI.E.a.ii).

iii. Claim 19.b: “a second interface configured to receive data via a network”

See element 1.b (Section VI.E.a.iii).

iv. Claim 19.c: “a cache memory”

See element 1.c (Section VI.E.a.iv).

v. Claim 19.d: “a storage device”

See element 1.d (Section VI.E.a.v).

vi. Claim 19.e: “a processor coupled via a communication path to the storage device, the processor configured to access the cache memory during processing of the I/O traffic”

See element 1.e (Section VI.E.a.vi) (Heil’s processor is coupled to the storage device via a communication path and configured to access the cache memory during processing of the I/O traffic).

vii. Claim 19.f: “the processor further configured to perform one or more access operations at the storage device based on the I/O traffic, wherein the communication path is distinct from the dedicated I/O channel”

See element 1.f (Section VI.E.a.vii) (Heil’s processor is configured to perform an access operation at the storage device based on the I/O traffic) *and* element 1.e (Section VI.E.a.vi) (the communication path is distinct from the I/O channel).

k. Claim 20

See claim 2 (Section VI.E.b). Heil discloses that the dedicated I/O channel is associated with dedicated throughput that corresponds to the I/O traffic because the data passing through the I/O channel corresponds to the I/O traffic. Ex.1004, ¶154.

l. Claim 21

See claim 3 (Section VI.E.c). Heil discloses the I/O traffic is received by HBA 117 via the PCI bus from the Node 1 host. Ex.1004, ¶155. This is distinct from the first data which is network traffic received by fibre channel chip 120 from fiber channel network 121. *Id.*


m. Claim 22

See claim 7 (Section VI.E.d). Heil discloses that a requesting HBA may send network traffic including one or more read requests to another HBA on the network. Ex.1006, 11:54-56, 12:4-6; Ex.1004, ¶156.

n. Claim 23

See claim 9 (Section VI.E.f).

Heil discloses “the storage device and the dedicated I/O channel are independently accessible.” Ex.1004, ¶157. Heil’s Fig. 1 shows two independent pathways (1) from the Node 1 host via the PCI bus to the processor (the dedicated I/O channel) and (2) from a remote Node N host via the fiber network and fibre channel chip 120 to the local drives via a SCSI or FCAL interface:



These paths are independent of each other as shown in Fig. 1. Ex.1004, ¶¶121-124.

F. Ground 2: Claims 10 and 11 are anticipated by Heil under 35 U.S.C. § 102(b)

Heil anticipates claims 10 and 11. See Sections VI.E.g and VI.E.h; Ex.1004, ¶¶128-141. As explained in Section V, claims 10 and 11 are not entitled to the January 22, 1999, priority date of the '409 application and have a priority date no earlier than March 5, 2003. In addition to qualifying as prior art under (pre-AIA)

35 U.S.C. § 102(e), Heil qualifies as prior art under (pre-AIA) 35 U.S.C. § 102(b) for these claims.

G. Ground 3: Claims 1-3 and 6-24 are obvious over Heil and Nakayama under 35 U.S.C. § 103(a)

a. A POSITA Would Have Combined Heil and Nakayama

It would have been obvious for a POSITA to combine Heil and Nakayama, with a reasonable chance of success. Ex.1004, ¶¶158-163. Heil and Nakayama are in the same field of invention and disclose similar cache and storage architectures. *Id.*, ¶160 (“Heil ‘relates to distribution of storage data within computer nodes’ [and] Nakayama relates to a ‘storage control system’ ‘for storing input/output data to/from host computers.’”). Both Heil and Nakayama recognize the importance of cache storage in addition to disk storage, and both disclose methods of reading and writing to cache and disk storage. *Id.* Heil discloses I/O requests as discussed in Section VI.E.a, and a POSITA would understand that read and write requests are types of I/O requests. *Id.* A POSITA would have recognized the complementary nature of read and write requests, and that enabling both read and write requests in any storage architecture would be critical to its success and usefulness. *Id.*, ¶160 (“Heil describes how ‘a large number of processors may be added to a computing environment all of which share access to common storage devices’ [and] discusses ‘file level access’ and ‘block level access.’ Nakayama discloses that ‘the data access command is a write command or a read command,’ and then describes how

write commands and read commands are executed. A POSITA would implement Heil in view of Nakayama by implementing the details of the read and write commands of the latter on the former.”). A POSITA seeking to improve storage access for a host computer in a storage system having both cache and disk storage would look to Nakayama to do so, and would use read and write operations such as those disclosed by Nakayama as I/O requests to access cache and disk storage as disclosed by Heil. *Id.* Read and write operations such as those disclosed by Nakayama are applicable to similar systems having both cache and disk storage and yield predictable results. *Id.*, ¶¶161-162 (“A practitioner would implement this combination with simple program modifications to the operating system and would have expected success.”).

b. Claim 1


The claimed apparatus is obvious in view of Heil and Nakayama. Ex.1004, ¶¶164-173.

i. Claim 1.pre

Heil discloses the preamble. *See* Section VI.E.a.i.

ii. Claim 1.a

Heil discloses element 1.a. *See* Section VI.E.a.ii. In addition, Nakayama discloses a similar data storage architecture as Heil in which host computer 300 is connected via channel 301 (which may be a SCSI bus control, as seen at 304) to disk control 302 having cache memory 309 as shown in Nakayama’s Figure 3:



Ex.1007, Fig. 3, 4:40-5:9, 4:56-60; Ex.1004, ¶¶166-170. To the extent Heil describes I/O requests between its host computers and HBAs as “access requests” and does not explicitly describe the access requests using the term “read command,” a POSITA would have found it obvious that Heil’s description of I/O requests includes write requests. Ex.1004, ¶¶166-170. Additionally, Nakayama explicitly teaches that the host computers may provide access requests in the form of read and write requests to storage, including the cache memory and disk drive. Ex.1007, 3:66-4:25, 4:48-51, 4:61-65. A POSITA would understand that Heil’s I/O traffic may comprise a read command as disclosed by Nakayama. Ex.1004, ¶170.

In addition, a POSITA would have implemented Nakayama's read and write commands with Heil's storage system for the reasons in Section VI.G.a.

iii. Claim 1.b

iv. Claim 1.c

v. Claim 1.d

vi. Claim 1.e

Heil discloses elements 1.b-1.e. *See* Sections VI.E.a.iii-VI.E.a.vi.

vii. Claim 1.f

Heil discloses element 1.f. *See* Section VI.E.a.vii. In addition, Nakayama discloses that the host computers may provide access requests in the form of read and write requests to storage, including the cache memory and disk drive. Ex.1007, 3:66-4:25, 4:48-51, 4:61-65; Ex.1004, ¶¶171-173. Nakayama further discloses reading or writing to the storage device based on whether a read request or a write request is issued. Ex.1007, 6:20-54. A POSITA would understand that Heil may read or write at the storage device based on whether the I/O request is a read request or a write request. Ex.1004, ¶¶171-173 ("reading and writing to a cache memory or a disk drive based on a read command or a write command provided by a host computer" is the processor accessing cache memory or the storage device as claimed). A POSITA would use Nakayama's read and write requests with Heil's storage system for the reasons in Section VI.G.a.

c. Claim 2

Heil discloses claim 2. *See* Sections VI.E.b and VI.G.b (Heil in view of Nakayama renders obvious claim 1).

d. Claim 3

Heil discloses claim 3. *See* Sections VI.E.c and VI.G.b.

e. Claim 6

Heil discloses I/O traffic that directs an HBA to access a cache memory or storage device. *See* claim 1 (Sections VI.E.a and VI.G.b). To the extent Heil describes I/O requests between its host computers and HBAs as “access requests” and does not explicitly describe the access requests using the term “write command,” Nakayama explicitly teaches that the host computers may provide access requests in the form of read and write requests to storage, including the cache memory and disk drive. Ex.1007, 3:66-4:25, 4:48-51, 4:61-65, 6:36-38, 7:3-10; Ex.1004, ¶¶178-184. A POSITA would understand that Heil’s I/O traffic may comprise a write command as disclosed by Nakayama. Ex.1004, ¶¶178-184. A POSITA would use Nakayama’s read and write commands with Heil’s storage system for the reasons in Section VI.G.a.

f. Claim 7

Heil discloses claim 7. *See* Sections VI.E.d and VI.G.b.

g. Claim 8

Heil discloses claim 8. *See* Sections VI.E.e and VI.G.b.

h. Claim 9

Heil discloses claim 9. *See* Sections VI.E.f and VI.G.b. To the extent Heil describes I/O requests between its host computers and HBAs as “access requests” and does not explicitly describe the access requests using the term “read command,” Nakayama explicitly teaches that the host computers may provide access requests in the form of read and write requests to storage, including the cache memory and disk drive. Ex.1007, 3:66-4:25, 4:48-51, 4:61-65; Ex.1004, ¶¶189-193. If a read command is issued, Nakayama further discloses that data is read from the cache memory. Ex.1007, 6:39-54. A POSITA would understand that Heil’s processor may read data from the cache memory based on a read command from the host. Ex.1004, ¶¶189-193. A POSITA would use Nakayama’s read and write commands with Heil’s storage system for the reasons in Section VI.G.a.

i. Claim 10

Heil discloses claim 10. *See* Sections VI.E.g and VI.G.b. In addition, Nakayama also discloses claim 10. Ex.1004, ¶¶194-197. Nakayama discloses that if the received command is a read command, the processor determines whether the read data exists in the cache memory. Ex.1007, 6:39-7:2. If the data is found within the cache, the data is read from the cache—otherwise, the processor reads the data from the drives. *Id.* The processor is configured to route the read command to the cache memory or the storage device, depending on where the data may be found.

Ex.1004, ¶196. A POSITA would use Nakayama's read and write requests with Heil's storage system for the reasons in Section VI.G.a.

j. Claim 11

Heil discloses claim 11. *See* Sections VI.E.h and VI.G.b. Heil retrieves data from the cache if the data is within the cache and from the local disk if the data is not within the cache for a remote HBA (Ex.1006, 5:25-28, 11:66-12:7), which a POSITA would perform for a local HBA "to improve I/O performance." Ex.1006, 11:27-34; Ex.1004, ¶¶198-200.

To the extent Heil describes I/O requests between its host computers and HBAs as "access requests" and does not explicitly describe the access requests using the term "read operation" or "write operation," Nakayama explicitly teaches that the host computers may provide access requests in the form of read and write requests to storage, including the cache memory and disk drive. Ex.1007, 3:664:25, 4:48-51, 4:61-65; Ex.1004, ¶¶201-204. Nakayama discloses reading or writing to the storage device based on whether a read request or a write request is issued. Ex.1007, 6:20-54. In addition, Nakayama discloses denying an I/O request if the processor fails to obtain access to a drive. Ex.1004, ¶203. A POSITA would understand that Heil in view of Nakayama discloses reading or writing at the storage device based on whether the I/O request is a read request or a write request, or denying the I/O request. Ex.1004, ¶¶198-204. A POSITA would use

Nakayama's read and write requests with Heil's storage system for the reasons in Section VI.G.a.

k. Claim 12

Heil discloses claim 12. *See* Sections VI.E.i and VI.G.b.

l. Claim 13


i. Claim 13.pre

If the Board determines that the preamble is limiting, Heil discloses a method of data storage and teaches the preamble. Ex.1006, 1:9-14 (“This invention relates to distribution of storage data within computer nodes of a clustered computing environment, and more particularly, to *an apparatus and method for I/O request shipping operable using peer-to-peer communications within host bus adapters of the clustered systems.*”); Ex.1004, ¶208.

ii. Claim 13.a

Heil discloses receiving input/output (I/O) traffic from a host device via a dedicated I/O channel at a first interface. *See* element 1.a (Section VI.E.a.ii).

In addition, Nakayama discloses a similar data storage architecture as Heil in which a host computer 300 is connected via channel 301 (which may be a SCSI bus control, as seen at 304) to a disk control 302 having a cache memory 309 as shown in Nakayama's Figure 3:



Ex.1007, Fig. 3, 4:40-5:9, 4:56-60; Ex.1004, ¶¶210-214. To the extent Heil describes I/O requests between its host computers and HBAs as “access requests” and does not explicitly describe the access requests using the term “write command,” Nakayama explicitly teaches that the host computers may provide access requests in the form of read and write requests to storage, including the cache memory and disk drive. Ex.1007, 3:66-4:25, 4:48-51, 4:61-65; Ex.1004, ¶213. A POSITA would understand that Heil’s I/O traffic may comprise a write command as disclosed by Nakayama. Ex.1004, ¶214. A POSITA would use

Nakayama's read and write commands with Heil's storage system for the reasons in Section VI.G.a.

iii. Claim 13.b

Heil discloses element 13.b. *See* element 1.b (Section VI.E.a.iii).

iv. Claim 13.c

Heil discloses element 13.c. *See* element 1.c (Section VI.E.a.iv). In addition, Heil in view of Nakayama renders obvious storing data (i.e., writing data) to a cache. *See* claim 6 (Section VI.G.e); Ex.1007, 6:30-38.

v. Claim 13.d

Heil discloses element 13.d. *See* element 1.d (Section VI.E.a.v). In addition, Heil in view of Nakayama renders obvious storing data (i.e., writing data) to disk storage. *See* claim 6 (Section VI.G.e); Ex.1007, 7:3-10.

vi. Claim 13.e


Heil discloses element 13.e. *See* element 1.e (Section VI.E.a.vi).

vii. Claim 13.f

Heil discloses element 13.f. *See* element 1.e (Section VI.E.a.vi) and element 1.f (Section VI.E.a.vii). Heil discloses that the processor (of the HBA) may access the storage device (the local drive) based on the I/O traffic. In response to an I/O request by CPU 100, the HBA may determine that the request may be serviced locally and retrieve the response data from the local disks. Ex.1006, 7:36-37 ("As shown in FIG. 1, one HBA services the I/O requests from CPU 1 and CPU 2

100.”), 11:45-53 (“After receiving a block I/O request 400, the 110 redirector has the means to search the directory and determining means to locate the local or remote disk drives that are storing the I/O requested blocks 410. If the location of the blocks is local 420, then the blocks are retrieved from the respective local disk drive.”); Ex.1004, ¶¶215-217.

Heil’s Figure 1 shows that the access operations at the storage device utilize a communication path from the processor to the storage device—i.e., from the CPU 117.1 to the local drives 118 via a SCSI interface, annotated below in **red**—that is distinct from the dedicated I/O channel—i.e., from the Node 1 host system to HBA 117 via a PCI bus, annotated in **green**:




Ex.1006, Fig. 1; Ex.1004, ¶216.

viii. Claim 14

See claim 3 (Sections VI.E.c and VI.G.d). Heil discloses the I/O traffic is received by HBA 117 via the PCI bus from the Node 1 host. Ex.1004, ¶218. This is distinct from the first data which is network traffic received by fiber channel chip 120 from fiber channel network 121. *Id.*

ix. Claim 15

Heil discloses “performing a first access operation at the storage device independently of the dedicated I/O channel.” Ex.1004, ¶¶219-223. Heil’s Fig. 1 shows two independent pathways (1) from the Node 1 host via the PCI bus to the processor (the dedicated I/O channel) and (2) from a remote Node N host via the fiber network and fibre channel chip 120 to the local drives via a SCSI or FCAL interface:



Ex.1004, ¶222.

In addition, Heil discloses that the storage device may be accessed based on a network request from a requestor HBA on the network, rather than the local HBA. Ex.1006, 11:54-12:6; Ex.1004, ¶223. Because requestor HBAs on the network use the fibre channel pathway rather than the PCI bus pathway, access operations by a requestor HBA on the network to the storage device are independent of I/O operations from the local HBA to the storage device. Ex.1004, ¶223 (“Heil discloses that the storage device is accessible based on a network request from a requestor HBA on the network, rather than the local HBA...Figure 1 shows that a requestor HBA servicing an I/O request may “ship” the request to another HBA having data responsive to the I/O request using the fibre channel pathway, rather than the PCI bus pathway for I/O requests from the host system. Access operations by a requestor HBA on the network to the storage device is independent of I/O operations from the host system to the local HBA using the dedicated I/O channel.”).

m. Claim 16

Heil discloses I/O requests from the host system to the HBA that directs accesses of either a cache memory or local drives. Ex.1006, 12:1-4 (“The remote HBA can cache blocks of data retrieved from its local disk, to service future local or remotely generated requests for these blocks.”), 6:36-38, 7:3-10, 3:34-40.

To the extent Heil describes I/O requests between its host computers and HBAs as “access requests” and does not explicitly describe the access requests using the term “write command,” Nakayama explicitly teaches that the host computers may provide access requests in the form of read and write requests to storage, including the cache memory and disk drive. Ex.1007, abstract, 3:66-4:25, 4:48-51, 4:61-65, 6:20-35; Ex.1004, ¶¶224-230. Nakayama further discloses reading or writing to the storage device based on whether a read request or a write request is issued. Ex.1007, 6:20-54. A POSITA would understand that Heil may read or write at the storage device based on whether the I/O request is a read request or a write request. Ex.1004, ¶¶224-230. A POSITA would use Nakayama’s read and write requests with the cache memory of Heil’s storage system for the reasons in Section VI.G.a.

n. Claim 17

i. Claim 17.a

Heil discloses that read requests from a remote host may be sent via the fiber network to HBAs connected to the network and discloses this element. Ex.1006, 11:56-12:1 (“If the requested I/O block is stored in a remote disk, the initiating HBA ships the block I/O request to the peer HBA that the initiating HBA has determined can perform the requisite processing 450...The remote HBA retrieves the requested blocks of data from its HBA cache, or if the blocks of data are not

within cache, then directly from the remote HBA's local disk."); Ex.1004, ¶¶233-237.

To the extent Heil describes I/O requests between its host computers and HBAs as "access requests" and does not explicitly describe the access requests using the term "read request," Nakayama explicitly teaches that the host computers may provide access requests in the form of read to storage, including the cache memory and disk drive. Ex.1007, 3:66-4:25, 4:48-51, 4:61-65; Ex.1004, ¶¶233-237. A POSITA would understand that Heil's network data may comprise read requests as disclosed by Nakayama. *Id.* A POSITA would use Nakayama's read and write requests with Heil's storage system for the reasons in Section VI.G.a.

ii. Claim 17.b

In response to a read request from a remote host, Heil discloses that data responsive to the read request is retrieved from cache memory and transmitted back to the remote host, thereby disclosing this element. Ex.1006, 12:1-10 ("The remote HBA retrieves the requested blocks of data from its HBA cache, or if the blocks of data are not within cache, then directly from the remote HBA's local disk. The remote HBA can cache blocks of data retrieved from its local disk, to service future local or remotely generated requests for these blocks. The remote I/O blocks are sent from the HBA performing the remote processing (e.g., a read request) to the initiating HBA 460."); Ex.1004, ¶¶238-239.

o. Claim 18

Heil in view of Nakayama renders claim 18 obvious. Ex.1004, ¶¶240-248. In addition to storing data in the cache (the second data), Heil discloses that the HBAs may also store a directory of block locations in cache. Ex.1006, 12:19-25 (“The HBA’s via I/O redirector 240, 331 has the means to create the directory storing the mapping of the local and/or remote location of blocks within the clustered computing environment. Specifically, the data within each respective local disk is mapped and coalesced into a directory.”), 11:8-9; Ex.1004, ¶243. Additionally, commands may be transmitted through the fiber network to update the directory (i.e., store directory updates, the fourth data). Ex.1006, 12:25-33 (“After the host(s) and HBA(s) initialize 500, the HBAs map the location of blocks within their respective local storage subsystem 502. The HBAs build a local directory containing the mapped location information for their respective local storage subsystem 504. The HBAs use their peer-to-peer capabilities to broadcast the contents of their directory to peer HBAs 506 and each HBA updates its respective directory with the broadcast information 508.”); Ex.1004, ¶243. Heil discloses storing fourth data indicated by a second I/O request received at the second interface. Ex.1004, ¶244.

Alternatively, Heil discloses “future requests” which may also constitute a second I/O request. Ex.1006, 12:1-4, Ex.1004, ¶245 (“Heil discloses that an HBA

‘can cache blocks of data retrieved from its local disk, to service future local or remotely generated requests for these blocks’...more than one I/O request may be issued by a host system—Heil explicitly contemplates that the host system may issue a second I/O request to the local HBA even after the host system issues a first I/O request to the local HBA...the second I/O request corresponds to the claimed second read request. If data responsive to the second I/O request is not found at the local HBA, and the local HBA ships that request to a remote HBA, then the remote HBA is the apparatus that has received the second read request via the network.”). Data written to cache memory in response to the future request constitutes the fourth data. *Id.*

To the extent Heil describes I/O requests between its host computers and HBAs as “access requests” and does not explicitly describe the access requests using the term “write command,” Nakayama explicitly teaches that the host computers may provide access requests in the form of read and write requests to storage, including the cache memory and disk drive. Ex.1007, 3:66-4:25, 4:48-51, 4:61-65; Ex.1004, ¶¶246-248. Nakayama further discloses reading or writing to the storage device based on whether a read request or a write request is issued. Ex.1007, 6:20-54. A POSITA would understand that Heil may read or write at the storage device based on whether the I/O request is a read request or a write request.

Ex.1004, ¶¶246-248. A POSITA would use Nakayama’s read and write requests with Heil’s storage system for the reasons in Section VI.G.a.

p. Claim 19

i. Claim 19.pre

Heil discloses the preamble. *See* Section VI.E.j.i.

ii. Claim 19.a

Heil discloses element 19.a. *See* Section VI.E.j.ii. To the extent Heil describes I/O requests between its host computers and HBAs as “access requests” and does not explicitly describe the access requests using the term “read command” or “write command,” Nakayama explicitly teaches that the host computers may provide access requests in the form of read and write requests to storage, including the cache memory and disk drive. Ex.1007, 3:66-4:25, 4:48-51, 4:61-65; Ex.1004, ¶¶251-254. Nakayama further discloses reading or writing to the storage device based on whether a read request or a write request is issued. Ex.1007, 6:20-54. A POSITA would understand that Heil may read or write at the storage device based on whether the I/O request is a read request or a write request. Ex.1004, ¶¶251-254. A POSITA would use Nakayama’s read and write requests with Heil’s storage system for the reasons in Section VI.G.a.

iii. Claim 19.b

iv. Claim 19.c

v. Claim 19.d

vi. Claim 19.e

vii. Claim 19.f

Heil discloses elements 19.b-19.f. *See* Sections VI.E.j.iii-VI.E.j.vii.

q. Claim 20

Heil discloses claim 20. *See* Sections VI.E.k and VI.G.p.

r. Claim 21

Heil discloses claim 21. *See* Sections VI.E.l and VI.G.p.

s. Claim 22

Heil discloses claim 22. *See* Section VI.E.m. To the extent Heil describes I/O requests between its host computers and HBAs as “access requests” and does not explicitly describe the access requests using the term “read request” or “write request,” Nakayama explicitly teaches that the host computers may provide access requests in the form of read and write requests to storage, including the cache memory and disk drive. Ex.1007, 3:66-4:25, 4:48-51, 4:61-65; Ex.1004, ¶¶259-263. A POSITA would understand that Heil’s network data may comprise read or write requests as disclosed by Nakayama. Ex.1004, ¶263. A POSITA would use Nakayama’s read and write requests with Heil’s storage system for the reasons in Section VI.G.a.

t. Claim 23

Heil discloses claim 23. *See* Sections VI.E.n and VI.G.p.


u. Claim 24

i. Claim 24.pre

If the Board determines that the preamble is limiting, Heil discloses a data storage and communication apparatus and teaches the preamble. Ex.1006, 1:9-14 (“This invention relates to distribution of storage data within computer nodes of a clustered computing environment, and more particularly, to *an apparatus and method for I/O request shipping operable using peer-to-peer communications within host bus adapters of the clustered systems.*”); Ex.1004, ¶267.

ii. Claim 24.a


As discussed in Section VI.D, Petitioner interprets this element as a means-plus-function limitation with a corresponding structure (as best understood for the purposes of this proceeding) of a front-end module that receives data from hosts via I/O channels, and equivalents. Heil discloses the function of receiving input/output (I/O) traffic from a first host device via a dedicated I/O channel. *See* element 1.a (Section VI.E.a.ii), element 13.a (Section VI.G.1.ii); Ex.1004, ¶¶268-276. Heil further discloses that this function is performed by the same corresponding structure, such as a front-end module (PCI bus bridge 115, PCI bus 116.5, or Front End I/F 102) that receives I/O traffic from a host device:



Ex.1006, Fig. 1, 6:35-40 (“Node 1 150 includes a host system connected to a HBA 117 via a Host-to-Peripheral Component Interconnect (hereinafter referred to as “PCI”) bus bridge 115 and PCI bus 116.5.”), 6:67-7:1 (“The front-end interface 102 is a PCI-to-PCI Bus bridge that connects, via PCI bus 117.6, various peripheral I/O devices, such as the local drives 118 and Fibre Channel Chip 120, to a centralized system PCI bus 116.5. Front-end interface 102 is connected to embedded Host-to-PCI Bus bridge 101 via PCI bus 117.6.”); Ex.1004, ¶¶268-276. The communication path between CPU 100 and the HBA, including PCI Bus Bridge 115, PCI Bus 116.5, and Front End I/F 102, is a dedicated I/O channel because it is dedicated to I/O between HBA 117 and the Node 1 host system. Ex.1004, ¶272.

Additional corresponding structure may include the software layers and protocols of the I/O bus that operate to allow the HBA to receive I/O traffic from CPU 100. Ex.1006, 6:8-11, 10:27-35, Fig. 2l; Ex.1004, ¶271.

In addition, Nakayama discloses a similar data storage architecture as Heil in which a host computer 300 is connected via channel 301 (which may be a SCSI bus control, as seen at 304) to a disk control 302 having a cache memory 309 as shown in Nakayama's Figure 3:




Ex.1007, Fig. 3, 4:40-5:9, 4:56-60; Ex.1004, ¶¶273-276. To the extent Heil describes I/O requests between its host computers and HBAs as “access requests”

and does not explicitly describe the access requests using the term “write command,” Nakayama explicitly teaches that the host computers may provide access requests in the form of read and write requests to storage, including the cache memory and disk drive. Ex.1007, 3:66-4:25, 4:48-51, 4:61-65; Ex.1004, ¶¶273-276. A POSITA would understand that Heil’s I/O traffic may comprise a write command as disclosed by Nakayama. A POSITA would use Nakayama’s read and write commands with Heil’s storage system for the reasons in Section VI.G.a.

iii. Claim 24.b

As discussed in Section VI.D, Petitioner interprets this element as a means-plus-function limitation with (as best understood for the purposes of this proceeding) a corresponding structure of a network link, and equivalents. Heil discloses the function of receiving data via and network. *See* element 1.b (Section VI.E.a.iii), element 13.b (Section VI.G.1.iii); Ex.1004, ¶¶277-280. Heil further discloses that this function is performed by the same corresponding structure of a network link in the form of a fibre channel chip 120 for receiving data from a fiber network 121:




Ex.1006, Fig. 1, 8:23-26 (“Fibre Channel Chip 120 connects Node 1 150 to the rest of the clustered environment via the network Fibre Channel backbone 121.”);

Ex.1004, ¶279.

iv. Claim 24.c


As discussed in Section VI.D, Petitioner interprets this element as a means-plus-function limitation with (as best understood for the purposes of this proceeding) a corresponding structure of a cache memory, and equivalents. Heil discloses the function of short-term data storage. *See* element 1.c (Section VI.E.a.iv), element 13.c (Section VI.G.1.iv); Ex.1004, ¶¶281-284. Heil further discloses that this function is performed by the same corresponding structure of a cache memory in the form of cache 117.2:



Ex.1006, Fig. 1, 7:25-26 (“To improve I/O performance, CPU 117.1 temporarily caches disk data in cache memory 117.2 that is frequently accessed.”); Ex.1004, ¶283.

v. Claim 24.d

As discussed in Section VI.D, Petitioner interprets this element as a means-plus-function limitation with (as best understood for the purposes of this proceeding) a corresponding structure of a back-end storage device or local storage device. Heil discloses the function of long-term data storage. *See* element 1.d (Section VI.E.a.v), element 13.d (Section VI.G.1.v); Ex.1004, ¶¶285-288. Heil further discloses that this function is performed by the same corresponding structure of a local storage device in the form of local drives 118 and 123:




Ex.1006, Fig. 1, 10:14-20 (“The front end interface 195 is a PCI-to-PCI Bus bridge that connects, via PCI bus 125.2, various peripheral I/O devices, such as the local drives 123, to centralized PCI bus 127.5.”); Ex.1004, ¶287.

vi. Claim 24.e

As discussed in Section VI.D, Petitioner interprets this element as a means-plus-function limitation with (as best understood for the purposes of this proceeding) a corresponding structure of a cache manager and a configuration manager, and equivalents. Heil discloses the function of performing one or more access operations at the means for short-term data storage during processing of the I/O traffic and for performing one or more access operations at the means for long-term data storage based on the I/O traffic. *See* element 1.e (Section VI.E.a.vi), 1.f (Section VI.E.a.vii), element 13.e (Section VI.G.1.vi), element 13.f (Section VI.G.1.vii); Ex.1004, ¶¶289-295. Heil further discloses that this function is performed by the same corresponding structure of a processor 117.1 that manages

access operations to cache memory (cache 117.2) and to a storage device (local drives 118), as well as software drivers referred to as ISMs for accessing cache and local storage:



Ex.1006, Fig. 1, 7:9-11 (“HBA 117 also includes a CPU 117.1, local memory 116, and cache memory 117.2, which are connected by internal processor bus 117.3.”), 11:11-19, Fig. 2; Ex.1004, ¶291.

Heil describes that its processor performs access operations to the cache during processing of I/O traffic and may access cache memory while processing I/O traffic. Ex.1006, 7:25-28 (“To improve I/O performance, CPU 117.1 temporarily caches disk data in cache memory 117.2 that is frequently accessed. Later, the CPU 117.1 stores the data permanently in the slower disk drive devices


118.”), 12:1-4 (“The remote HBA can cache blocks of data retrieved from its local disk, to service future local or remotely generated requests for these blocks.”); Ex.1004, ¶292.

Heil states that the processor performs access operations to the local storage device based on I/O traffic, for example if the I/O traffic directs retrieval of data from the cache, and from the local disk if the requested data is not found in cache:

The remote HBA retrieves the requested blocks of data from its HBA cache, or if the blocks of data are not within cache, then directly from the remote HBA’s local disk.

Ex.1006, 11:66-12:1; Ex.1004, ¶293.

Furthermore, Figure 1 shows that the processor (and its communication paths to the cache and the local drives) is distinct from the dedicated I/O channel (from the Node 1 host system to HBA 117 via a PCI bus):



Ex.1006, Fig. 1; Ex.1004, ¶294.

H. Ground 4: Claim 4 is obvious over Heil, Nakayama, and Gulick under 35 U.S.C. § 103(a)

a. A POSITA Would Have Combined Heil, Nakayama, and Gulick

It would have been obvious for a POSITA to combine Heil, Nakayama, and Gulick with a reasonable chance of success. Ex.1004, ¶¶298-299. A POSITA would have combined Heil and Nakayama as explained in Section VI.G.a. A POSITA at the time of invention would understand that PCI is a standardized format of local computer buses for attaching hardware devices to a computer and further understand that SCSI is a standard for connecting, interfacing, and transferring data between computers and peripheral devices. Ex.1004, ¶299. A

POSITA would know that PCI buses are interfaceable with SCSI adapters.

Ex.1004, ¶299. A POSITA would look to Gulick to interface SCSI adapters with PCI buses, as Gulick explicitly discloses that it is “desirable” to apply SCSI adapters to PCI buses to interface with SCSI devices, including drive storage, and would have been motivated to do so. Ex.1008, 5:37-51; Ex.1004, ¶299 (“SCSI is the most frequently encountered channel for small form disk drives in the mid to late 90s and offered advantages such as treating peripheral devices in a largely independent fashion. A POSITA wishing to interface a device having a PCI bus to peripheral small form disk drives would be interested in interfacing a PCI bus with a SCSI adapter to a SCSI compatible drive.”). Gulick would inform the POSITA to do so.

b. Claim 4 is obvious

As discussed in Section VI.E.a.ii regarding element 1.a, Heil describes a PCI bus bridge and PCI bus between the Node 1 host and the processor. Heil discloses a SCSI interface, but the SCSI interface is located between the processor and storage device—not between the processor and the host. Ex.1006, 10:12-20 (“Front-end interface 195 is connected to SCSI or FCAL interfaces 122.3, 122.4 via PCI bus 125.2. These interfaces can be either SCSI or FCAL depending upon the bus utilized to connect the local drives 123.”). Nonetheless, it would have been obvious to a POSITA in view of Gulick to use a SCSI channel instead of a PCI

channel. A POSITA would understand from Gulick that SCSI is compatible with PCI and is a known alternative to PCI, and a POSITA would incorporate Gulick's SCSI interface and channel in Heil to interface with a host system with predictable results. Ex.1008, 5:37-51; Ex.1004, ¶¶300-306 ("Gulick explicitly teaches the POSITA that SCSI channels may be used to interface computer systems, like Heil's host server, with other devices...a POSITA would have been aware that use of SCSI channels provides benefits over use of PCI, and would have been motivated to replace Heil's PCI channel between the host server and HBA by a SCSI channel as the dedicated I/O channel as taught by Gulick.").

I. Ground 5: Claim 5 is obvious over Heil, Nakayama, and Berman under 35 U.S.C. § 103(a)

a. A POSITA Would Have Combined Heil, Nakayama, and Berman

It would have been obvious for a POSITA to combine Heil, Nakayama, and Berman with a reasonable chance of success. Ex.1004, ¶¶309-310. A POSITA would have combined Heil and Nakayama as explained in Section VI.G.a. A POSITA would have known that it is desirable to interface fibre channels such as Heil's with Ethernet. Ex.1004, ¶¶309-310 ("A POSITA in 1999 would understand that fibre channel networks have improved performance over Ethernet networks. [T]he POSITA in 1999 would know that, for a variety of reasons including cost and existing infrastructure, Ethernet networks are more widely deployed than fibre

channel networks. [T]he POSITA would be motivated to use fibre channel networks where they are available and interface fibre channel networks with existing ('legacy') Ethernet networks where fibre channel networks have not yet been deployed.”). A POSITA would look to Berman to interface ethernet networks with Heil’s Fibre Channel network. Berman explicitly teaches that it is advantageous to do so, and a POSITA would understand that interfacing a Fibre Channel with a legacy network such as ethernet would yield predictable results. Ex.1009, 12:9-11, claims 13-15, 3:23-30; Ex.1004, ¶¶309-310.

b. Claim 5 is obvious

As discussed in Section VI.E.a.iii regarding element 1.b, Heil describes a second interface that receives data from a Fibre Channel network. While Heil does not explicitly teach an Ethernet interface to interface with an Ethernet network, it would have been obvious to a POSITA in view of Berman to incorporate an Ethernet interface to an Ethernet network. Ex.1009, 12:9-11; Ex.1004, ¶¶311-317 (“[A] POSITA would have been motivated...to interface Heil’s Fibre Channels with Ethernet networks [because] Berman explicitly teaches the POSITA to do so.”).

VII. SECONDARY CONSIDERATIONS

Petitioner is unaware of any evidence related to secondary considerations, whether identified by the prosecution history or by PO in the parallel litigation,

known to the industry, or identified by Dr. Franzon. Ex.1004, ¶¶318-320. To the extent PO makes any arguments regarding secondary considerations, Petitioner reserves the right to further address them.

VIII. DISCRETION

A. The *Fintiv* Factors Strongly Favor Institution

An evaluation of the factors under *Apple Inc. v. Fintiv, Inc.*, IPR2020-00019, Paper 11 (Mar. 20, 2020) (precedential), weighs against discretion to deny institution. This Petition presents compelling evidence of unpatentability and, per the PTO interim guidance dated 6/21/2022 (Ex.1012), demonstrates that the PTAB should not discretionarily deny institution.

Factor 1 is whether the court granted a stay or whether evidence exists that one may be granted if an IPR is instituted. This factor is neutral or inapplicable because Petitioner has not moved to stay the co-pending district court litigation pending IPRs.

Factor 2 is the proximity of the court's trial date to the projected statutory deadline for the PTAB's final written decision. The district court has not yet issued a scheduling order or set a trial date. The median time to trial in the Western District of Texas as of June 30, 2022 (the last time the statistics were published) was 28.3 months (*see* <https://www.uscourts.gov/statistics-reports/federal-court-management-statistics-june-2022>), meaning that any trial will likely not occur until

well after a final written decision here. Therefore, this factor weighs strongly against discretionary denial.

Factor 3 is the investment in the parallel proceeding by the court and the parties. The district court and the parties have invested *de minimis* resources for this litigation. The parallel district court litigation for Cisco, Microsoft, and Amazon was transferred from the Waco Division of the Western District of Texas to the Austin Division of the Western District of Texas during the early stages of the litigation. There is currently no schedule entered in the litigation, no claim construction proceedings have begun, and no discovery has taken place. Therefore, this factor weighs strongly against discretionary denial.

Factor 4 is the overlap between issues raised in the petition and in the parallel proceeding. This Petition challenges all claims (1-24) of the '092 Patent, which therefore will inherently overlap with the claims of the '092 Patent asserted against Cisco, Microsoft, and Amazon in the co-pending parallel district court litigation. But, PO has made clear that they are not asserting certain of the challenged claims, including claims 11 and 18, and therefore this Petition addresses a broader scope of claims than will be addressed in the parallel proceeding. Therefore, this factor weighs strongly against discretionary denial.

Factor 5 is whether the petitioner and the defendant in the parallel proceeding are the same party. The fact that Petitioner and PO are parties to the

WDTX litigation does not outweigh the other factors that strongly weigh against discretionary denial. Indeed, nothing in *Fintiv* suggests that the same party between the proceedings weighs in favor of discretionary denial. Therefore, this factor is neutral or inapplicable.

Finally, **Factor 6** are other circumstances that impact the Board's exercise of discretion, including the merits. Factor 6 weighs strongly against discretionary denial. Petitioner respectfully submits that the unpatentability challenge presented in this Petition is compelling. If the Board agrees, "that determination alone demonstrates that the PTAB should not discretionarily deny institution under *Fintiv*." Ex.1012 at 4-5. And as *Fintiv* noted, "if the merits of a ground raised in the petition seem particularly strong...institution of a trial may serve...overall system efficiency and integrity." *Fintiv*, 14-15.

Further, as discussed in Section V, the '092 Patent issued after a single non-final rejection that did not consider any of the art presented in this Petition. Institution is consistent with the significant public interest against "leaving bad patents enforceable." *Thryv, Inc. v. Click-To-Call Techs., LP*, 140 S. Ct. 1367, 1374 (2020).

Accordingly, based on a "holistic view of whether efficiency and integrity of the system are best served," the facts here weigh against exercising discretionary denial. *Samsung Elecs. Co. Ltd. v. Dynamics Inc.*, IPR2020-00505, Paper 11 at 15

(Aug. 12, 2020). At a minimum, Factors 2, 3, 4, and 6 (or combinations thereof) outweigh Factors 1 and 5, and favor institution.

B. Discretionary Denial Under *General Plastic* Is Unwarranted

General Plastic does not support denying institution because Petitioner filed this petition and motion for joinder timely, copying the same grounds and relying upon the same evidence and expert declaration as in IPR2023-00120 (the IPR Petitioner seeks to join). Additionally, because Petitioner agrees to be subject to the schedule in IPR2023-00120 and take an understudy role only, this petition will not tax the resources of the Board or impair the Board's ability to meet the statutory deadline for a final determination.

C. 35 U.S.C. § 325(d) Favors Institution

As discussed above, the '092 Patent was issued after a single non-final rejection without considering any of the art presented in this petition. The grounds presented in this petition are not cumulative or duplicative of any prior art or arguments previously presented to the Patent Office during prosecution.

To the extent PO makes any arguments regarding discretionary denial, Petitioner reserves the right to further address them.

IX. CONCLUSION

For the foregoing reasons, Petitioner requests institution of IPR for claims 1-24 of the '092 Patent based on the above-specified grounds.

Dated: March 20, 2023

Respectfully submitted,

/ Brian E. Ferguson /

Brian E. Ferguson
LEAD COUNSEL
Reg. No. 36,801
WINSTON & STRAWN LLP
1901 L Street, N.W.
Washington, D.C. 20036
Telephone: 202-282-5276
Email: beferguson@winston.com

*Counsel for Petitioners Cisco Systems, Inc. and
Microsoft Corporation*

Juan C. Yaquian
BACK-UP COUNSEL
Reg. No. 70,755
WINSTON & STRAWN LLP
800 Capital Street, Suite 2400
Houston, TX 77002
Telephone: (713) 651-2600
Fax: (713) 651-2700
Email: jyaquian@winston.com

*Counsel for Petitioners Cisco Systems, Inc. and
Microsoft Corporation*

Krishnan Padmanabhan
BACK-UP COUNSEL
(to seek *pro hac vice* admission)
WINSTON & STRAWN LLP
200 Park Avenue
New York City, NY 10166
Telephone: (212) 294-6700
Fax: (212) 294-4700
Email: kpadmanabhan@winston.com

Counsel for Petitioner Cisco Systems, Inc.

Michael Rueckheim

BACK-UP COUNSEL

(to seek *pro hac vice* admission)

WINSTON & STRAWN LLP

255 Shoreline Drive, Suite 520

Redwood City, CA 94065

Telephone: (650) 858-6500

Fax: (650) 858-6550

Email: mrueckheim@winston.com

Counsel for Petitioner Microsoft Corporation

Brian C. Nash

BACK-UP COUNSEL

Reg. No. 58,105

MORRISON & FOERSTER LLP

701 Brazos Street, Suite 1100

Austin, TX 78701

Telephone: (737) 309-0700

Email: BNash@mofo.com

*Counsel for Petitioners Amazon.com, Inc.,
Amazon Web Services, Inc., and Amazon.com
Services LLC*

CERTIFICATE OF COMPLIANCE

This Petition complies with the word count limits set forth in 37 C.F.R. § 42.24(a)(1)(i), because this Petition contains 13,617 words, excluding the parts of the Petition exempted by 37 C.F.R. § 42.24(a)(1), determined using the word count provided by Microsoft Word, which was used to prepare this Petition.

Dated: March 20, 2023

Respectfully submitted,

/ Brian E. Ferguson /

Brian E. Ferguson
Reg. No. 36,801
WINSTON & STRAWN LLP
1901 L Street, N.W.
Washington, D.C. 20036
Telephone: 202-282-5276
Email: beferguson@winston.com

Lead Counsel for Petitioners

CERTIFICATE OF SERVICE

The undersigned hereby confirms that the foregoing Petition for *Inter Partes* Review, associated exhibits, and Power of Attorney were caused to be served on March 20, 2023 via overnight courier upon the following counsel of record for Patent Owner in the related district court litigation:

William P. Ramey, III
Ramey LLP
5020 Montrose Blvd., Suite 800
Houston, Texas 77006
(713) 426-3923
wramey@rameyfirm.com

The foregoing were caused to be served on March 20, 2023 via overnight courier upon the following counsel of record for Patent Owner:

Toler Law Group
8500 Bluffstone Cove
Suite A201
Austin, Texas 78759

Respectfully submitted,

/ Brian E. Ferguson /

Brian E. Ferguson
Reg. No. 36,801
WINSTON & STRAWN LLP
1901 L Street, N.W.
Washington, D.C. 20036
Telephone: 202-282-5276
Email: beferguson@winston.com

Lead Counsel for Petitioners