

*Petition for Inter Partes Review
of U.S. Patent No. 8,275,827*

FUNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

DATABRICKS, INC.
Petitioner

v.

BYTEWEAVR, LLC
Patent Owner.

Inter Partes Review No. IPR2025-00716
Patent 8,275,827 B2

**PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT 8,275,827**

TABLE OF CONTENTS

	Page
MANDATORY NOTICES	1
A. Real Party-in-Interest [37 C.F.R. § 42.8(b)(1)]	1
B. Related Matters (37 C.F.R. § 42.8(b)(2)).....	1
C. Designation of Lead and Backup Counsel (37 C.F.R. § 42.8(b)(3)) and Service of Information (37 C.F.R. § 42.8(b)(4)).....	2
I. INTRODUCTION.....	4
II. CERTIFICATION OF GROUNDS FOR STANDING.....	4
III. OVERVIEW OF CHALLENGE AND RELIEF REQUESTED.....	4
A. Prior Art Patents and Printed Publications.....	4
B. Grounds for Challenge.....	5
IV. OVERVIEW OF THE '827 PATENT	5
A. The '827 patent.....	5
B. Summary of the Prosecution History.....	8
V. LEVEL OF ORDINARY SKILL IN THE ART.....	9
VI. CLAIM CONSTRUCTION (37 C.F.R. § 42.104(B)(3)).....	10
A. “NAS device” and “dedicated NAS device” (Claims 2 & 14)	11
B. “stand-alone dedicated NAS device” (Claims 2 & 14).....	11
VII. OVERVIEW OF THE PRIOR ART	14
A. Carter.....	14
B. OceanStore.....	18
C. Gibson.....	20

**TABLE OF CONTENTS
(Continued)**

	Page
D. Motivation to Combine Carter with Gibson	21
VIII. GROUNDS FOR UNPATENTABILITY	23
A. Collateral Estoppel Applies to Claims 2 and 14	23
B. Ground 1: Claims 2, 10-12, 14, and 22-24 are Obvious over Carter.....	28
1. Claim 1	28
2. Claim 13	38
3. Claims 2 and 14	40
4. Claims 9 and 21	44
5. Claims 10-12 and 22-24	46
a. Claim 10	46
b. Claim 11	47
c. Claim 12	49
d. Claim 22	51
e. Claims 23-24	52
C. Ground 2: Claims 2 and 14 are obvious over OceanStore.....	52
1. Claim 1	52
2. Claim 13	58
3. Claims 2 and 14	60
D. Ground 3: Claims 2 and 14 are obvious over Carter in view of Gibson.....	62

**TABLE OF CONTENTS
(Continued)**

	Page
IX. PTAB DISCRETION SHOULD NOT PRECLUDE INSTITUTION.....	65
A. Fintiv Analysis.....	65
B. <i>General Plastics</i> Analysis.....	67
C. §325(d): Advanced Bionics	69
X. CONCLUSION	69
APPENDIX: CHALLENGED CLAIM LISTING	73

TABLE OF AUTHORITIES

Cases

<i>Apple Inc. v. Fintiv, Inc.</i> , IPR2020-00019, Paper 11 (PTAB Mar. 20, 2020)	70
<i>Aylus Networks, Inc. v. Apple Inc.</i> , 856 F.3d 1353, 1360-61 (Fed. Cir. 2017)	29
<i>Byteweavr, LLC v. Databricks, Inc.</i> , No. 2:24-cv-00162-JRG-RSP (E.D. Tex., filed March 8, 2024)	4
<i>Databricks, Inc. v. Weisfield</i> , No. 24-cv-01417-RAJ (W.D. Wash., filed Sep. 6, 2024)	4
<i>E.I. du Pont Nemours & Co. v. Phillips Petroleum Co.</i> , 849 F.2d (Fed. Cir. 1988)	29
<i>EMC Corp. et al. v. Intellectual Ventures I LLC</i> , IPR2017-00374, Paper 53 (PTAB June 20, 2018)	4
<i>EMC Corp. et al. v. Intellectual Ventures I LLC</i> , IPR2017-00439, Paper 50 (PTAB June 20, 2018)	5
<i>General Plastic Indus. Co., Ltd. v. Kaisha</i> , IPR2016-01357, Paper 19 (PTAB Sept. 6, 2017)	
<i>Google LLC v. Parus Holdings, Inc.</i> , IPR2020-00847, Paper 9 (PTAB Oct. 21, 2020)	
<i>Intellectual Ventures I LLC et al. v. Lenovo Group Ltd.</i> , No. 1:16-cv-10860-IT (D.Mass., filed May 10, 2016)	4
<i>Intellectual Ventures I, LLC v. NetApp, Inc.</i> , No. 1:16-cv-10868-IT (D.Mass, filed May 11, 2016)	2
<i>Nikken USA, Inc. v. Robinsons-May, Inc.</i> , 51 F’Appx 874, 884 (Fed. Cir. 2002) ..	26
<i>Ohio Willow Wood Co. v. Alps S., LLC</i> , 735 F.3d (Fed. Cir. 2013)	
<i>Papst Licensing GMBH & Co. KG v. Samsung Elecs. Am., Inc.</i> , 924 F.3d (Fed. Cir. 2019)	
<i>Sand Revolution II, LLC v. Continental Intermodal Group-Trucking LLC</i> , IPR2019-01393, Paper 24 (PTAB June 16, 2020)	
<i>Soverain Software LLC v. Victoria’s Secret Direct Brand Mgmt., LLC</i> , 778 F.3d 1311,1319 (Fed. Cir. 2015)	
<i>Valve Corp. v. Elec. Scripting Prods., Inc.</i> , IPR2019-00062, Paper 11 at 10 (PTAB Apr. 2, 2019)	71

TABLE OF AUTHORITIES
(Continued)

	Page(s)
<i>Videndum Prod. Sol., Inc. v. Rotolight Ltd.</i> , IPR2023-01218, Paper 12 at 6 (PTO-Ofc. of Dir., Apr. 19, 2024)	71
<i>Weatherford U.S. v. Enventure Global Tech., Inc.</i> , IPR2020-01684, Paper 16 (PTAB Apr. 14, 2021)	69
 Statutes	
35 U.S.C. § 102(a)	17, 21, 23
35 U.S.C. § 102(b)	70
35 U.S.C. § 103	8, 70
35 U.S.C. § 311	72
35 U.S.C. §325(d)	72
35 U.S.C. §112	11
 Regulations	
37 C.F.R. § 42.10(b)	6
37 C.F.R. § 42.101	72
37 C.F.R. § 42.104(a)	7
37 C.F.R. § 42.104(b)(3)	13
37 C.F.R. § 42.8(b)(1)	4
37 C.F.R. § 42.8(b)(2)	4
37 C.F.R. § 42.8(b)(4)	6
37 C.F.R. §§ 42.22(a)(1)	7
37 C.F.R. §§ 42.6(e)	75
37 C.F.R. §42.24	74
37 C.F.R. §42.8	74

EXHIBIT LIST (37 C.F.R. § 42.63(e))

Exhibit	Description
1001	U.S. Patent No. 8,275,827
1002	File History of U.S. Patent No. 8,275,827
1003	Declaration of Dr. John D. Kubiadowicz, Ph.D. (“Kubiadowicz Decl.”)
1004	Curriculum Vitae of Dr. John D. Kubiadowicz, Ph.D
1005	International Application WO 98/22881 to Carter et al., titled Remote Access in a Globally Addressable Storage Environment,” published on May 28, 1998 (“Carter”)
1006	Kubiadowicz, John et al., “OceanStore: An Architecture for Global-Scale Persistent Storage,” stamped November 20, 2000 by the University of Michigan, and appearing in ASPLOS-IX proceedings: Ninth International Conference on Architectural Support for Programming Languages and Operating Systems, Cambridge, Mass., Nov. 12-15, 2000 (pp. 190-201) (“OceanStore”)
1007	Garth A. Gibson et al., File Server Scaling with Network-Attached Secure Disks (“Gibson”)
1008	Morris, et al., “Andrew, a Distributed Personal Computing Environment”
1009	Satyanarayanan et al., “The ITC Distributed Filed System: Principles and Design”
1010	Litzkow, et al., “Condor – A Hunter of Idle Workstations”
1011	Chien, “The Grid, Blueprint for a New Computing Infrastructure”
1012	Plank, et al., “The Internet Backplane Protocol: Storage in the Network”
1013	Bloom, “Space/Time Trade-offs in Hash Coding with Allowable Errors”
1014	Champine, et al., “Project Athena as a Distributed Computer System”

Exhibit	Description
1015	Joint Claim Construction and Prehearing Statement in <i>Byteweavr, LLC v. Databricks, Inc.</i> , No. 2:24-cv-00162-JRG-RSP (E.D. Tex.) (dkt. entry 55)
1016	Exhibit A to Joint Claim Construction and Prehearing Statement in <i>Byteweavr, LLC v. Databricks, Inc.</i> , No. 2:24-cv-00162-JRG-RSP (E.D. Tex.) (dkt. entry 55-1)
1017	Exhibit B to Joint Claim Construction and Prehearing Statement in <i>Byteweavr, LLC v. Databricks, Inc.</i> , No. 2:24-cv-00162-JRG-RSP (E.D. Tex.) (dkt. entry 55-2)
1018	Plaxton et al., “Accessing Nearby Copies of Replicated Objects in a Distributed Environment,” appearing in <i>Proceedings of the ACM Symposium on Parallelism in Algorithms and Architecture</i> (SPAA), June 1997
1019	IPR2017-00374, Paper 53 (Final Written Decision)
1020	IPR2017-00439, Paper 50 (Final Written Decision)
1021	Declaration of Rachel J. Watters
1022	Docket Control Order in <i>Byteweavr, LLC v. Databricks, Inc.</i> , No. 2:24-cv-00162-JRG-RSP (E.D. Tex.)
1023	Docket in <i>Byteweavr, LLC v. Databricks, Inc.</i> , No. 2:24-cv-00162-JRG-RSP (E.D. Tex.) (retrieved March 13, 2025)

MANDATORY NOTICES

A. Real Party-in-Interest [37 C.F.R. § 42.8(b)(1)]

The real party-in-interest is Petitioner Databricks, Inc. (“Petitioner” or “Databricks”). No entity not named above is funding, controlling, or directing this Petition, or otherwise has had an opportunity to control or direct this Petition or Petitioner’s participation in any resulting IPR.

B. Related Matters (37 C.F.R. § 42.8(b)(2))

Petitioner is aware of the following related matters. Patent Owner asserts U.S. Patent No. 8,275,827 against Petitioner in *Byteweavr, LLC v. Databricks, Inc.*, No. 2:24-cv-00162-JRG-RSP (E.D. Tex., filed March 8, 2024). On September 6, 2024, Petitioner asserted unfair competition claims against Patent Owner in *Databricks, Inc. v. Weisfield*, No. 24-cv-01417-RAJ (W.D. Wash., filed Sep. 6, 2024).

The ’827 patent has been asserted in two previous cases, including *Intellectual Ventures I LLC et al. v. Lenovo Group Ltd.*, No. 1:16-cv-10860-IT (D.Mass., filed May 10, 2016) in the District of Massachusetts and *Intellectual Ventures I, LLC v. NetApp, Inc.*, No. 1:16-cv-10868-IT (D.Mass, filed May 11, 2016).

The ’827 patent has been the subject of two prior IPRs. In IPR2017-00374, claims 1, 3-9, 13, 15, 16, and 18-21 were found invalid as obvious. *EMC Corp. et al. v. Intellectual Ventures I LLC*, IPR2017-00374, Paper 53 at 59 (PTAB June 20, 2018), *aff’d* 786 F.Appx. 1021 (Fed. Cir. Sep. 25, 2019). In IPR2017-00439, claims

1, 3-6, 8, 9, 13, 15-18, 20, and 21 were found invalid as anticipated or obvious. *EMC Corp. et al. v. Intellectual Ventures I LLC*, IPR2017-00439, Paper 50 at 61 (PTAB June 20, 2018) *aff'd* 786 F.Appx. 1021 (Fed. Cir. Sep. 25, 2019).

**C. Designation of Lead and Backup Counsel (37 C.F.R. § 42.8(b)(3))
and Service of Information (37 C.F.R. § 42.8(b)(4))**

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

Lead Counsel for Petitioner	Back-Up Counsel for Petitioner
Robin Snader (Reg. No. 66,085) Erise IP, P.A. 7015 College Blvd., Ste. 700 Overland Park, KS 66211 Phone: 913-777-5600 Fax: Fax: 913-777-5601 Email: robin.snader@eriseip.com PTAB@eriseip.com	Jennifer Bailey (Reg. No. 52,583) Erise IP, P.A. 7015 College Blvd., Ste. 700 Overland Park, KS 66211 Phone: 913-777-5600 Fax: Fax: 913-777-5601 Email: jennifer.bailey@eriseip.com Eric Buresh (Reg No. 50,394) Erise IP, P.A. 7015 College Blvd., Ste. 700 Overland Park, KS 66211 Phone: 913-777-5600 Fax: Fax: 913-777-5601 Email: eric.buresh@eriseip.com Kathleen Fitterling (Reg. No 62,950) Erise IP, P.A. 7015 College Blvd., Ste. 700 Overland Park, KS 66211 Phone: 913-777-5600

Lead Counsel for Petitioner	Back-Up Counsel for Petitioner
	Fax: Fax: 913-777-5601 Email: kathleen.fitterling@eriseip.com

A Power of Attorney accompanies this Petition in accordance with 37 C.F.R. § 42.10(b).

A. Service Information [37 C.F.R. § 42.8(b)(4)]

Please address all correspondence to the lead and back-up counsel as shown above. Petitioners consent to electronic service by e-mail at the e-mail addresses provided above.

I. INTRODUCTION

Petitioner Databricks, Inc. (“Petitioner”) requests *inter partes* review of claims 2, 10-12, 14, and 22-24 of US Patent No. 8,275,827 (“the ’827 patent”). The ’827 patent is assigned to Byteweavr, LLC (“Patent Owner”).

II. CERTIFICATION OF GROUNDS FOR STANDING

Petitioner certifies pursuant to 37 C.F.R. § 42.104(a) that the ’827 patent is available for *inter partes* review and that Petitioner is not barred or estopped from requesting an *inter partes* review challenging the patent claims on the grounds identified in this Petition.

III. OVERVIEW OF CHALLENGE AND RELIEF REQUESTED

Pursuant to 37 C.F.R. §§ 42.22(a)(1) and 42.104(1)-(2), Petitioner challenges claims 2, 10-12, 14, and 22-24 of the ’827 patent.

A. Prior Art Patents and Printed Publications

The following references are pertinent to the ground of unpatentability:

1. International Application WO 98/22881 to Carter et al., titled “Remote Access in a Globally Addressable Storage Environment,” published May 28, 1998 (“Carter”; EX1005).
2. Kubiawicz, John et al., “OceanStore: An Architecture for Global-Scale Persistent Storage,” published November 20, 2000 (“OceanStore”; EX1006).
3. Garth A. Gibson et al., File Server Scaling with Network-Attached Secure Disks, published 1998 (“Gibson”; EX1007)

B. Grounds for Challenge

This petition, supported by the Declaration of Dr. John D. Kubiawicz, Ph.D. (EX1003), requests cancellation of claims 2, 10-12, 14, and 22-24 as unpatentable under 35 U.S.C. § 103. 35 U.S.C. § 314(a). The grounds include:

Ground 1	Claims 2, 10-12, 14, 22-24	Obvious under 35 U.S.C. § 103 in view of Carter
Ground 2	Claims 2 and 14	Obvious under 35 U.S.C. § 103 in view of OceanStore
Ground 3	Claims 2 and 14	Obvious under 35 U.S.C. § 103 in view of Carter and Gibson

IV. OVERVIEW OF THE '827 PATENT


A. The '827 patent

The '827 patent is directed to software-based network-attached storage (NAS) services hosted on a massively-distributed network of devices. EX1001, Title, Abstract; EX1003, ¶21.

The '827 patent explains that, while ordinary desktop computers often sit idle for long periods of time, these computers have the same specifications and often greater functional capabilities than dedicated NAS devices. EX1001, 2:46-55; EX1003, ¶22. The '827 patent proposes making these otherwise unused storage resources available to other devices connected to the network, effectively turning

ordinary desktop computers into NAS devices for the network. EX1001, 2:66-3:1, 42:36-41; EX1003, ¶23.

The '827 patent claims focus on new subject matter added to the specification when the '827 patent application was filed. EX1003, ¶24. The applicant added new Figure 21, which provides “a distributed processing system that provides network attached storage (NAS) services, according to the present invention.” EX1001, 42:29-32; EX1003, ¶24.



EX1001, Fig. 21 (annotated)

This new disclosure is found at 42:29-46:10. EX1003, ¶24. The specification states that in Figure 21, “NAS embodiment 2100” includes a network 102 that connects user devices (red) and NAS client devices (green). EX1001, 43:6-11; EX1003, ¶25. Each of the distributed NAS client devices runs a client agent to

provide NAS functionality. EX1001, 43:26-30; EX1003, ¶25. Each of the user devices uses the “software-based NAS resources” provided by the distributed computing system. EX1001, 44:60-65; EX1003, ¶25.

The client devices are configured to “mimic” dedicated NAS devices. EX1001, 43:47-56, 44:30-32, 5:64-6:2. These devices can provide “both storage and data location functionality, thereby acting both to provide data storage services and distributed data location services.” EX1001, 44:9-12; EX1003, ¶27. For example, the data location libraries (blue) may provide links or directories to where desired data is stored. EX1001, 43:59-62, 44:1-3; EX1003, ¶27. The NAS client devices can also be combined with NAS server systems, which may include a file storage database system and a centralized data location library, which can provide services for the “files and other data being stored on the multitude of software-based NAS devices.” EX1001, 44:27-37; EX1003, ¶28.

These client devices include client agents that have corresponding software-based NAS components. EX1001, 31:43-32:9; EX1003, ¶¶29-33. For example, a “client agent provides a protected data storage area accessible to outside devices,” “space for a cache of other device addresses attached to the network,” “a simple indexing system that is automatically re-indexed when content is added or removed from the storage area,” among other functionalities. EX1001, 31:46-62, Fig. 14B; EX1003, ¶32.

B. Summary of the Prosecution History

The '827 patent issued from U.S. Patent Appl. No. 09/834,785, which was filed on April 13, 2001. It is a child of fourteen prior continuations-in-part. EX1001; EX1003, ¶35. The Board in the prior IPRs treated April 13, 2001 as the critical date because the '827 patent claims are directed to new subject matter added to the application filed on April 13, 2001. EX1019 (IPR2017-00374, Paper 53) at 4; EX1020 (IPR2017-00439, Paper 50) at 4.

The Board was correct for two reasons. First, during prosecution, independent Claims 1 and 13 (29 and 42 during prosecution) were rejected until the applicant narrowed them to require a “location distributed device to store location information associated with data stored by the selected distributed devices through use of the respective client agents for the particular distributed device.” EX1002, 904. This concept of a “location distributed device [] stor[ing] location information” was new subject-matter added to the '827 patent application on April 13, 2001. EX1003, ¶37.

Second, during prosecution, independent Claims 1 and 13 (29 and 42 during prosecution, respectively) were rejected under 35 U.S.C. §112 because the patent failed to provide written description for “configuring the client agent program to have a software-based NAS component to assess unused or under-utilized storage resources in selected distributed.” The Applicant overcame the rejection by relying on the new subject-matter of Figure 21 and the corresponding disclosure (EX1001,

61:13-21) to argue that “[s]ince the client agent programs configure the devices to mimic NAS devices by relying on *spare and/or unused capabilities* of the existing devices, then it follows that the client agent programs must be capable of *assessing* what capabilities are *spare and/or unused* (unused or under utilized).” EX1002, 365 (emphasis in original); *id.* (arguing “it would have been obvious to one of ordinary skill in the art” and that “assess” “is a descriptive term that describes the function of the client agent program used to configure a software-based NAS device using devices with spare and/or unused capabilities”); *id.* at 368 (overcoming rejection). EX1003, ¶¶39-40. Thus, the priority date of the claims is no earlier than April 13, 2001. EX1003, ¶¶35-42.

V. LEVEL OF ORDINARY SKILL IN THE ART

Petitioner submits that a POSITA would have the following education and experience, as the Board applied in the prior IPRs: “a bachelor’s degree in computer science, computer engineering, electrical engineering, or equivalent, with 1–2 years of work experience in the field of computer networking, storage systems, or a related field” and “that such a person would have been familiar with computer systems architecture and design, and would have been aware of design trends relating to distributed processing systems.” EX1019 (IPR2017-00374, Paper 53) at 18; EX1020 (IPR2017-00439, Paper 50) at 18; EX1003, ¶20.

VI. CLAIM CONSTRUCTION (37 C.F.R. § 42.104(B)(3))

For purposes of this Petition, the terms not specifically discussed in this section should have their plain and ordinary meaning in light of the specification, as commonly understood by a POSITA. EX1003, ¶43.

During the prior IPRs, the Board rejected the Patent Owner’s proposed construction of the term “[r]epresenting with the corresponding software-based NAS component that the selected distributed devices respectively comprise NAS devices” in claims 1 and 13. EX1019 (IPR2017-00374, Paper 53) at 8-14; EX1020 (IPR2017-00439, Paper 50) at 9-14. Patent Owner argued this term required “each selected distributed device *to separately appear* to network-connected user devices as a dedicated NAS device.”¹ EX1019 at 10; EX1020 at 9-10. The Board declined to adopt this proposal because it improperly limited the claim to one of the three disclosed embodiments—specifically a standalone dedicated NAS device. EX1019 at 12-13; EX1020 at 11-13. The Board also explained that Patent Owner’s proposed construction would not alter the outcome of the IPR, because Carter and OceanStore each disclosed the claim under Patent Owner’s proposal. EX1019 at 14; EX1020 at 14. While the prior IPRs were decided when the claim construction standard in effect was broadest reasonable interpretation, in each IPR the Board explained that

¹ Unless otherwise noted, all emphasis is added.

it was “not aware of any evidence of record in this proceeding that would lead to a different outcome under the *Phillip* [sic] claim construction standard.” EX1019 at 8 n3; EX1020 at 8 n4. Databricks agrees with the Board’s conclusions, and thus, analyzes invalidity of the challenged claims consistent with the Board’s findings in these IPRs.

A. “NAS device” and “dedicated NAS device” (Claims 2 & 14)

While the Board did not construe the terms “NAS device” and “dedicated NAS device” in either of the previous ’827 patent IPRs, it explained in each of its Final Written Decisions that a “‘NAS device’ as defined in the ’827 patent ‘broadly refers to a device that makes data storage resources available to network-connected user devices’ and a “‘dedicated NAS device’ is a device whose primary operational purpose is for providing NAS services.” EX1020 at 36; EX1019 at 32-33. EX1003, ¶44. Databricks applies the Board’s interpretations of these terms in this petition.

B. “stand-alone dedicated NAS device” (Claims 2 & 14)

In the prior IPRs, the Patent Owner attempted to limit the “representing...” term of claim 1 to the ’827 patent’s “stand-alone dedicated NAS device” implementation by arguing the “representing...” term should be construed, in part, as “...each selected distributed device [] separately appear to network-connected user devices as a dedicated NAS device.” EX1019 at 9-14; EX1020 at 9-14. Patent Owner argued that its proposed construction was correct based on the patent’s stand-

alone dedicated NAS device implementation in which “each individual NAS device ‘brings its resources to the network with the *appearance of a dedicated NAS device*.” EX1019 at 13; IPR2017-00374, Paper 24 at 10; EX1020 at 13; IPR2017-00439, Paper 21 at 11. Moreover, the Board understood Patent Owner’s construction as referring to a stand-alone dedicated NAS device. *Id.* Thus, Petitioner applies Patent Owner’s interpretation of “stand-alone dedicated NAS device” as a “distributed device that separately appears to network-connected devices as a dedicated NAS device.” EX1003, ¶¶45-47.

This proposed construction is also consistent with the usage of this term in the ’827 patent. The ’827 patent specification discloses “three different NAS architectures....: a stand-alone only NAS device implementation, a NAS device fabric implementation, and a server assisted implementation. EX 1001, 43:37-40. The ’827 patent discloses that in the stand-alone dedicated NAS device implementation, distributed devices *separately appear* as dedicated NAS devices. EX 1001, 43:44-50. Specifically, the ’827 patent discloses that a “large number of desktop PCs in the enterprise (Intranet) or on the Internet can each act as a NAS device by...bring[ing] its resources to the network with the *appearance of a dedicated NAS device*,” and thus, “mimic the functionality of stand-alone NAS devices.” EX 1001, 42:49-43:1; *id.* 43:44-50 (discussing “NAS devices...mimicking dedicated stand-alone NAS devices). Because these NAS

devices “each act as” or *separately* “*appear* as a single or dedicated NAS device” to network-connected devices “user device[s]...can directly access them as such.” Ex 1001, 42:59-65; 43:44-47.

The patent further explains why a “stand-alone dedicated NAS device” *separately appears* as a dedicated NAS device. While these devices “operate[] like a standard dedicated NAS device” they are actually “software-based NAS devices that rely on spare and/or unused capabilities of the existing devices to provide [the stand-alone dedicated] NAS services.” *Id.* at 43:50-58. Moreover, a “standalone dedicated NAS device” “separately appears” as a dedicated NAS device because it may not include the “data location libraries” that would otherwise provide “links or directions to data” on the NAS device. *Id.* at 43:56-58. Thus, a standalone dedicated NAS device *separately appears* as a dedicated NAS device so that user devices can directly access its stored data, similar to the way a user device would access data on a self-contained device with its own resources. EX1001, 42:49-43:1; 43:44-58; EX1003, ¶¶45-47.

In co-pending litigation against Petitioner, Patent Owner proposes that “stand-alone dedicated NAS device” be construed as “a dedicated NAS device that contains its own storage, processing, connectivity, and management resources” while Petitioner proposes it be construed as “a distributed device that separately appears to network connected user devices as a dedicated NAS device.” EX1016 at 8;

EX1017 at 4. This Petition applies both constructions and the prior art invalidates under either construction.

VII. OVERVIEW OF THE PRIOR ART


A. Carter

International Application WO 98/22881 to Carter et al., titled “Remote Access in a Globally Addressable Storage Environment,” published on May 28, 1998 (“Carter,” EX1005), is a continuation-in-part of two U.S. Applications (U.S. App. Nos. 08/754,481 and 08/827,534), and incorporates both applications by reference. EX1005, 1. As such, Carter is prior art to the ’827 patent under 35 U.S.C. §§ 102(a) and (b).

Carter describes a distributed storage system that allows a group of network-connected computers to share their unused storage resources. EX1003, ¶70. Each computer runs software that allocates a portion of the computer’s unused local storage resources to the aggregated storage. *Id.* Each of Carter’s nodes includes a “software module” with “a data control program 32a, and a shared memory subsystem 34a.” EX1005, 10; *id.*, Abstract, 28, Fig. 1. This software manages a portion of each node’s available storage space (e.g., unused hard drive space) and pools the unused storage on each node so that “nodes [can] transparently access portions of the structured store using standard memory access commands.” EX1005, 4, 9, 24; EX1003, ¶¶71-72.


The shared memory space is logically divided into sections of fixed size (e.g., 4kB, of the virtual memory space), each of which is associated with a globally unique identifier, which is also known as a global address signal and is used to access that block. EX1005, 4; 31-32. This global address signal is a virtual address because it is “decoupled from (i.e., unrelated to) the physical and virtual address spaces of the underlying computer hardware.” EX1005, 4; EX1003, ¶74. The system translates a global address signal into the address of the physical storage (e.g., RAM or a hard disk) using a hierarchical structure of directory pages stored in the shared memory system that maps a global address to the node providing the storage resources for that address. *Id.*, 5, 38-40; EX1003, ¶¶74, 76.

Figure 10 illustrates an example of a hierarchical directory structure 340, where each level of directory pages (shaded blue) subdivides the shared memory space into smaller sections. EX1003, ¶77



EX1005, Fig. 10 (annotated)

The “responsible node stores information regarding the location and attributes of data associated with a respective global address,” “optionally stores a copy of that page’s data,” and “tracks ownership and responsibility” for the page’s data. EX1005, 32, 40. Figure 11 illustrates a system 400 that implements the hierarchical directory page structure 340 similar to Figure 10.



EX1005, Fig. 11 (annotated)

Figure 11 shows how global addresses are translated by the system into physical memory locations. For example, if address consumer 408b requests the address corresponding to virtual memory page “5,” the directory manager accesses the root directory page and traverses the pointers in the hierarchical directory until it reaches a directory page that identifies the physical location of virtual page “5.” Ex 1005, 39-40, Figs. 10, 11; EX1003, ¶85. As shown in Figure 11, virtual page 5 is stored on node 406a, which represents itself as a “NAS device.” EX1003, ¶86-87. The shared memory subsystem makes it possible for any one of the computers in the networked group to store/access a particular page of memory in the shared memory space using its global address. Carter accomplishes this by distributing location

information among the computers in the group, independent of the location of the underlying data. EX1003, ¶88.

B. OceanStore

“OceanStore: An Architecture for Global-Scale Persistent Storage” by Dr. John Kubiatawicz *et al.* (“OceanStore,” EX1006) was published in the Proceedings of the Ninth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS IX) in November 2000. EX1003, ¶¶190-194. OceanStore was not disclosed to the Patent Office during prosecution of the ’827 patent. EX1001, Cover. As the Board found in IPR2017-00374, OceanStore is prior art to the ’827 patent under 35 U.S.C. § 102(a) for the ’827 patent’s priority date of April 13, 2001. EX1019 at 27.

OceanStore concerns “an architecture for global-scale persistent storage.” EX1006, Title. OceanStore sought to take advantage of unused storage resources on computers. EX1006, 191; EX1003, ¶194. OceanStore was designed to span the globe and offer “uniform and highly available access to information” through widely dispersed servers. EX1006, Abstract, 190; EX1003, ¶198.

Figure 1 below shows that user devices could come in many forms. EX1006, 191; EX1003, ¶199. As shown, an “ocean” of storage resources is provided by an aggregation of “pools” of storage, each of which can cache or serve data to the user devices. *Id.*




Figure 1: The OceanStore system. The core of the system is composed of a multitude of highly connected “pools”, among which data is allowed to “flow” freely. Clients connect to one or more pools, perhaps intermittently.

OceanStore (EX1006), Fig. 1 (annotated)

OceanStore’s fundamental unit of information was a “persistent object,” or “floating replica,” which contained “pointers” to other objects, representing directories or local namespaces. EX1006, 191-192. Each floating replica was assigned a “globally unique identifier (GUID),” which identified the object and could be used to access the object. EX1006, 191-193; EX1003, ¶200-201. In short, a user would connect to a nearby OceanStore server, as depicted in Figure 1 (dashed lines). EX1003, ¶202. The user could then read or update objects using a GUID to identify the relevant object being requested or updated. *Id.* But each object could be any place in the network and need not be on the server

or pool device that the user connected to. EX1006, 193-97, Figs. 2, 3; EX1003, ¶202.

C. Gibson

“File Server Scaling with Network-Attached Secure Disks” by Garth A. Gibson et al. (“Gibson,” EX1007) was published (and publicly available) in June of 1997. EX1007; EX1021. Gibson was not disclosed to the Patent Office during prosecution of the ’827 patent. EX1001, Cover. Gibson’s publication date predates the ’827 patent, and Gibson is prior art under 35 U.S.C. §§ 102 (a) and (b).

Gibson discloses that “[n]etwork-attached storage provides high bandwidth by directly attaching storage to the network” as it “avoid[s] file server store-and-forward operations and allow[s] data transfers to be striped over storage and switched-network links.” EX1007, 272. Gibson discloses network-attached secure disks (“NASD”), which are stand-alone dedicated NAS devices, as shown below in Figure 3. EX1007, 275-79; EX1003, ¶¶278-283.




Figure 3: Network-attached secure disks (NASD) are designed to offload more of the file system's simple and performance-critical operations. For example, in one potential protocol a client, prior to reading a file, requests access to that file from the file manager (1), which delivers a capability to the authorized client (2). So equipped, the client may make repeated accesses to different regions of the file (3, 4) without contacting the file manager again unless the file manager chooses to force reauthorization by revoking the capability (5).

EX1007, 276 (Fig. 3). Gibson explains that NASD drives offload “[c]ommon, data intensive operations, such as reads and writes” from the file manager. EX1007, 275. For example, “[b]efore issuing its first read of a file, the client authenticates itself with the file manager and requests access to the file.” EX1007, 276. After access is granted, the client may directly access the NASD drive containing the file until the client’s access is revoked. EX1007, 276; EX1003, ¶¶282-283.

D. Motivation to Combine Carter with Gibson

A POSITA would have been motivated to combine Carter with Gibson. First, both Carter and Gibson relate to the same field of endeavor, *i.e.*, a network system that shares resources available at distributed devices on the network. EX1005, Abstract (“The globally addressable storage environment allows data to be accessed and shared by and among the various computers on the plurality of networks.”); EX1007 (“Users are increasingly using distributed file systems to access data across local area networks” and “[n]etwork-attached storage provides high bandwidth by directly attaching storage to the network”); EX1003, ¶274.

Second, both Carter and Gibson are directed to solving problems related to the efficiency, scalability, and data-sharing capacity of distributed networks. EX1005, 4:2-11 (disclosing improving distributed “network performance and productivity,” “fault tolerance,” and capacity “to communicate and share data in a transparent and dynamic manner”); EX1007, 272 (disclosing providing users of

distributed file systems “with efficient, scalable, high-bandwidth access to stored data”). A POSITA would therefore have sought out various available solutions, including those described in Carter and Gibson, for improving the efficiency, scalability, and data-sharing capacity of distributed networks. EX1003, ¶275.

Carter solves the problem in part by creating and managing a virtual storage space that aggregates the unused storage resources of many storage devices connected to the network, including traditionally non-addressable data storage devices (*e.g.*, desktop computers). EX1005, Abstract, 4:15-17; EX1003, ¶276. Carter also discloses additional NAS devices beyond the aggregated virtual storage space—such as network memory device 226 illustrated in Fig. 6—that can be used by client devices to store or access data. EX1005, 29:21-30:2; Fig. 6. Gibson expands on the use of NAS devices by describing various examples of standalone dedicated NAS devices such as “Network-attached Secure Disks (NASD)” that “provide[] high bandwidth by directly attaching storage to the network” and “reduce[] the number of client-storage interactions that must be relayed through the file manager.” EX1007, 272, Fig. 3. For example, Gibson’s NASDs allow “common, data-intensive operations, such as reads and writes, [to] go straight to the disk...while policy decisions are made in the file manager.” *Id.*, 275. This improves the scalability of networks by providing more efficient access to a greater volume of

storage while significantly reducing file manager costs per client. *Id.*, 281 (Section 6.2), 282 (Section 7); EX1003, ¶276.

Third, a POSITA would have appreciated that Gibson's NASDs could have been easily incorporated into Carter by directly attaching them to Carter's network, similar to Carter's memory device 226. EX1003, ¶277. This combination would have yielded the predictable result of expanding the scalability of Carter's networks by providing more efficient access to a greater volume of storage through Gibson's NASDs operating in parallel with Carter's aggregated shared memory space. *Id.* Thus, a POSITA would have had a reasonable expectation of success as this is a predictable combination of known elements that would have been obvious to try, required nothing more than ordinary skill, and would have worked in an expected way. *Id.*

Motivation specific to claims 2 and 14 is discussed further in conjunction with those claims.

VIII. GROUNDS FOR UNPATENTABILITY

A. Collateral Estoppel Applies to Claims 2 and 14

A party seeking to invoke collateral estoppel must show: (1) the issue is identical to one decided in the first action; (2) the issue was actually litigated in the first action; (3) resolution of the issue was essential to a final judgment in the first action; and (4) the party against whom collateral estoppel is being asserted had a full

and fair opportunity to litigate the issue in the first action. *In re Freeman*, 30 F.3d 1459, 1465 (Fed. Cir. 1994); *Google*, 54 F. 4th 1377 at 1381.

As an initial matter, collateral estoppel applies to IPR proceedings. *Papst Licensing GMBH & Co. KG v. Samsung Elecs. Am., Inc.*, 924 F.3d 1243, 1251 (Fed. Cir. 2019). Regarding *Freeman*'s first prong, it is well established that patent claims need not be identical for collateral estoppel to apply. *Soverain Software LLC v. Victoria's Secret Direct Brand Mgmt., LLC*, 778 F.3d 1311,1319 (Fed. Cir. 2015). Only the issues of patentability must be identical. *Ohio Willow Wood Co. v. Alps S., LLC*, 735 F.3d 1333, 1342 (Fed. Cir. 2013). Thus, collateral estoppel may apply even if the patent claims "use slightly different language to describe substantially the same invention," so long as "the differences between the unadjudicated patent claims and adjudicated patent claims do not materially alter the question of invalidity." *Id.*

Collateral estoppel applies to claims 2 and 14. The limitation that claims 2 and 14 add to invalidated claims 1 and 13 do not materially alter the question of invalidity based on (1) the Patent Owner's own arguments about the scope of claims 1 and 13 in the prior IPRs, and (2) the Board's final written decision confirming that Carter and OceanStore disclosed the asserted claims even under Patent Owner's interpretation of the claims. *Aylus Networks, Inc. v. Apple Inc.*, 856 F.3d 1353, 1360-61 (Fed. Cir. 2017) (IPR proceedings are part of the intrinsic evidence); *accord Nikken USA, Inc. v. Robinsons-May, Inc.*, 51 F'Appx 874, 884 (Fed. Cir. 2002); *E.I.*

du Pont Nemours & Co. v. Phillips Petroleum Co., 849 F.2d 1430, 1439 (Fed. Cir. 1988)

As discussed above, the claim term “[r]epresenting with the corresponding software-based NAS component that the selected distributed devices respectively comprise NAS devices” was previously construed. EX1019 (IPR2017-00374, Paper 53) at 8-14; EX1020 (IPR2017-00439, Paper 50) at 9-14. In those IPRs, Patent Owner argued that the proper construction of this term was “using the corresponding software-based NAS components of the selected distributed devices to cause each selected distributed device to separately appear to network-connected user devices as a dedicated NAS device.” EX1019 at 10; EX1020 at 9-10. Claims 2 and 14 are substantively identical in scope to Patent Owner’s proposed construction, as shown in the table below. Additionally, as explained in Section VII above, the ’827 patent makes clear that a “stand-alone dedicated NAS device” is a “distributed device that separately appears to network-connected user devices as a dedicated NAS device.”

Disputed Term in Prior IPRs	PO’s Construction	Claims 2 and 14
...representing with the corresponding software-based NAS component that the selected distributed devices respectively comprise NAS devices...	“using the corresponding software-based NAS components of the selected distributed devices to cause each selected distributed device to separately appear to network-connected user devices as a dedicated NAS device.”	2. The method of claim 1, further comprising enabling at least one of the selected distributed devices to function as a stand-alone dedicated NAS device through use of the respective client

'827 patent, claims 1, 13	IPR2017-00374, Paper 24 (POR) at 8; IPR2017-00439, Paper 21 (POR) at 7.	agents for the particular distributed device. '827 patent, claim 2; <i>id.</i> claim 14.
---------------------------	---	---

EX1019 at 10; EX1020 at 9-10.

The Board found that Patent Owner's construction improperly attempted to limit claim 1 to "stand-alone dedicated NAS devices," which is one of three disclosed embodiments and is expressly recited by claims 2 and 14. EX1019 at 12-13; EX1020 at 11-13 (citing EX1001, 42:59-65, 43:44-47). As the Board explained, the '827 patent discloses three embodiments, including "stand-alone" NAS devices, a "NAS client device fabric," and a "server assisted [implementation]," but "Patent Owner's construction intentionally read[] out the embodiments covering a NAS fabric and the 'server assisted implementation.'" EX1019 at 12-13; EX1020 at 12-13. Indeed, Patent Owner argued its proposed construction is correct because claim 1 is limited to the stand-alone dedicated NAS device embodiment in which "each individual NAS device 'brings its resources to the network with the appearance of a dedicated NAS device.'" EX1019 at 13; EX1020 at 13 (emphasis in original in each). While the Board did not adopt Patent Owner's construction, it found that "even if we accepted Patent Owner's proposed construction, it would not alter our determination regarding the patentability of the challenged claims" because each of

Carter and OceanStore also disclosed Patent Owner’s proposed construction. EX1019 at 30-33; EX1020 at 32-37.

The Board explained that each of Carter and OceanStore discloses the “representing” limitation of claims 1 and 13 “based on the entirety of the record before us and the claim constructions for ‘representing’ proposed by either party.” EX1019 at 33; EX1020 at 37. Specifically, Carter discloses “a shared memory subsystem on *each node* that allows the node to mimic a dedicated NAS device by making its storage available to other nodes.” EX1020 at 37. OceanStore discloses “multiple storage resources are provided by an aggregation of ‘pools’ of storage, *each of which* can cache or serve data to the user devices,” “*each pool device* makes its storage available for use by other as depicted in Figures 2 and 3,” and OceanStore clients see more than a single NAS device.” EX1019 at 33. Thus, in addressing the validity of claims 1 and 13, the Board ruled that Patent Owner’s proposed construction—which captures the stand-alone dedicated NAS device implementation and is substantively identical in scope to claims 2 and 14—is disclosed and rendered obvious by each of Carter and OceanStore. EX1019 at 30-33; EX1020 at 32-37. Because the issue here is “identical to one decided in the first action,” “the differences between the unadjudicated patent claims and adjudicated patent claims [in view of Patent Owner’s proposed construction] do not materially alter the question of invalidity.” *Ohio*, 735 F.3d at 1342; *Freeman*, 30 F.3d at 1465.

Moreover, the issue was actually litigated in the first action as part of the prior IPRs (second factor), resolution of the issue was essential to a final judgment in the first action as explained in the respective IPR Final Written Decisions (third factor), and Patent Owner had a full and fair opportunity to litigate the issue in the IPRs and chose not to appeal the Board's decisions (fourth factor). *Id.* Thus, Patent Owner is estopped from arguing the similar-scope limitations of claims 2 and 14 are valid. *Id.*

B. Ground 1: Claims 2, 10-12, 14, and 22-24 are Obvious over Carter

Claims 1 and 13, which claims 2 and 14 depend upon, were invalidated in view of Carter. EX1019 (IPR2017-00439, Paper 50). Claims 2 and 14 were not challenged in that proceeding. The analysis of claim 1 below reflects the ground that was successful in that IPR.

1. Claim 1

1(pre) "A computer-implemented method comprising ..."

This preamble does not limit the claim. Regardless, the Board found that Carter discloses it. EX1020 (IPR2017-00439, Paper 50) at 24. Carter "relates in general to distributed computer workgroups and, more specifically, to a globally addressable storage environment that accommodates remote access and two or more interconnected computer networks." EX1005, Abstract; *id.*, 27 (disclosing "methods for providing a computer system"); EX1003, ¶¶99-101.

1(a) "configuring a distributed processing system of a plurality of distributed devices coupled to a network, wherein the plurality of

distributed devices include respective client agents configured to process respective portions of a workload for the distributed processing system”

The Board found that Carter discloses this limitation. IPR2017-00439, Paper 50 at 24-26 (citing EX1005, Abstract, 4, 9, 24, 29, 34); EX1003 ¶¶102-110. Carter teaches “computer network systems that have adaptable system configurations for dynamically exploiting *distributed network resources*” to “increas[e] network performance and productivity.” EX1005, 4. Carter’s distributed systems include “a *plurality of network nodes* [(i.e., “a plurality of distributed devices coupled to the network”)] that access a memory space storing a structured store of data, such as a structured file system or a database.” *Id.*, 9. EX1003, ¶102.

Carter discloses a software package of components that run on each of the distributed devices in the system, including the “shared memory subsystem” and the “data control program,” a “policy controller” or “policy element 290,” “distributor” component, “disk directory manager,” “migration controller,” “flow scheduler,” and “data control program.” EX1005, 10, 24-25, Figs. 2, 8. Collectively, this software is a “client agent” for processing portions of a distributed workload. EX1003, ¶¶102-104. As Dr. Kubiawicz explains, Carter’s software agent is equally a “client agent” under either of the proposed constructions in the district court litigation. *Id.*, ¶109. Carter explains that the shared memory subsystem is responsible for transfer and storage workloads for the various physical storage devices that make up the distributed system. EX1003, ¶106. For example, Carter discloses that “[e]ach node

212a-212c *through the shared memory subsystem* 220 can access each page 224a-224d stored in the virtual shared memory 222. EX1005, 29. Each page 224a-224d represents a unique entry of computer data stored within the virtual shared memory 222, and thus, using the shared memory subsystem, each page 224a- 224d can be accessed or modified by each one of the nodes 212a-212c. EX1005, 35, Fig. 8 (illustrating shared memory subsystem); EX1003, ¶¶106-108.

1(b) “wherein the respective client agents for particular distributed devices of the plurality of distributed devices have corresponding software-based network attached storage (NAS) components configured to assess unused or under- utilized storage resources in selected distributed devices of the plurality of distributed devices”

The Board found that Carter discloses this limitation. IPR2017-00439, Paper 50 at 26-32 (citing EX1005, 19, 28-31, 36-42, 142-143, Figs. 7, 10-11); EX1003 ¶¶111-118.

Carter’s shared memory subsystem is a “software-based NAS component.” As discussed above, Carter’s “shared memory subsystem 220 is a *software module*” that is part of the “client agent” that provides access to a shared addressable memory. IX.B.1.(b); EX1001, 28; EX1003, ¶¶111-113. Additionally, Carter’s shared memory subsystem communicates over the network with other nodes and makes data storage resources available to network-connected user devices. EX1005, 39-42, Figs. 10, 11; *id.*, Fig. 7. Thus, Carter’s shared memory subsystem is a “software-based NAS component.” EX1003, ¶¶111-113.

Carter’s shared memory subsystem also assesses unused or under-utilized storage resources in selected distributed devices. EX1003, ¶¶114-117. Carter’s policy element 290 is software that is part of the “client agent” (IX.B.1.b., *supra*) and is configured to assess unused or underutilized storage resources by “*determin[ing] the availability of resources such as printer capabilities, hard-disk, available RAM and other such resources.*” EX1005, 36:27-29; *id.*, 19:14-17, 142:14-17, 143:16-17; EX1003, ¶116. And as the Board found in the prior IPR—based on the testimony of Dr. Kubiawicz and Patent Owner’s declarant, Dr. Shenoy—Carter’s policy controller determines the free volume space available for allocation by evaluating the availability of hard disk space on the local node, as well as the availability of storage throughout the system by communicating with other nodes. EX1020 at 31; EX1003, ¶116.

Moreover, the policy element determines the availability of resources by assessing whether hard disk space is available, determining what volume of free space is available for allocation, and using that information to determine whether to proceed with allocation processing. EX1005, 30; EX1003, ¶116. Additionally, Carter discloses that the policy element assesses the extent to which storage resources are utilized and uses load balancing to prevent underutilization. *Id.* Indeed, the policy element “can employ any of the suitable heuristics to direct the elements, such as the paging controller, disk directory manager, and other elements to dynamically distribute

the available resources” and is configured to “allow portions of the distributed shared memory to move between persistent and electronic memory depending on predetermined conditions, such as recent usage.” EX1005, 36:29-37:2; 30:26-28; EX1003, ¶116.

Carter also distinguishes between “unused” storage space (*i.e.*, the pooled and shared virtual address storage) and “used” storage space (*i.e.*, “local storage that can be kept inaccessible to other nodes on the network” such as “private storage memory device 218” reserved for a node). EX1005, 28:19-26; EX1003, ¶117; *Cf.* EX1001, 12:26-31.

1(c) “representing with the corresponding software-based NAS component that the selected distributed devices respectively comprise NAS devices having an available amount of storage resources related to the unused or under-utilized storage resources for the selected distributed devices”


In the prior IPR, the Board found that Carter discloses this limitation. EX1020 at 32-37 (citing EX1005, 5, 9-11, 16, 41-42, Figs. 1, 11); EX1003 ¶¶119-130. As the Board found, Carter’s “shared memory subsystem” is a software-based NAS component on each node (IX.B.1.b., *supra*) that cooperates with the “shared memory subsystems” on other nodes to make unused or under-utilized physical data storage resources available to network-connected user devices. EX1020 at 37; EX1005, Figs. 10, 1; EX1003, ¶¶103-109, 112-116, 119. These “shared memory subsystems” allow each node to represent, or mimic, a dedicated NAS device by making its

unused or under-utilized storage available to other nodes. EX1020 at 36 (citing EX1005, 9:18-19, 9:26-27, 10:6-26, 10:30-11:2, Figs. 1, 8); EX1003, ¶¶119-120.

As the Board explained, “‘NAS device’ as defined in the ’827 patent ‘broadly refers to a device that makes data storage resources available to network-connected user devices’” and the ’827 patent further discloses “an integrated system referred to as a ‘NAS device fabric.’” EX1020 at 36 (citing EX1001, 2:31-34, 42:36-43:5). Like the ’827 patent, Carter discloses NAS devices and a similarly integrated NAS device system. For example, Carter’s “shared memory subsystem” causes individual nodes (e.g., commercially available PCs) to represent dedicated NAS devices, as their primary operational purpose is for providing NAS services, by making their unused or under-utilized storage resources available. EX1005, 10:6-26, Fig. 1; EX1003, ¶¶120-122. With reference to Fig. 1, Carter also discloses that “[e]ach of the nodes 12a-12d can access the addressable shared memory space 20 through the shared memory subsystems 34a-34d.” EX1005, 10:6-26, 10:30–11:2, Fig. 1. Carter explains that “[t]his allows the interconnected nodes 12a-12d to employ the addressable shared memory space 20 as a space for storing and retrieving data.” *Id.*, 10:5–6, Fig. 1; EX1003, ¶122.

Additionally, the shared memory subsystem (*i.e.*, claimed “corresponding software-based NAS component”) “represent[s]...that the selected distributed devices respectively comprise NAS devices having an available amount of storage”

through shared memory space 20. EX1020 at 36 (citing EX1005, 10:30-11:2, Fig. 1); EX1003, ¶123. This shared memory space transparently provides a disk-like interface that appears and is represented to external clients and user applications as a single large NAS storage device with a total size that is equal to the “current aggregate storage capacity in the addressable shared memory space” while “encapsulating the memory management operations” to conceal the underlying details of the participating computers that make up the shared storage space behind a layer of abstraction. EX1005, 31, 43; *id.*, 39-42, Figs. 10, 11; EX1003, ¶¶123-124. The translation of global addresses into physical locations on the participating distributed devices is hidden from the clients and applications requesting to use the available or unused resources of the shared memory space. EX1005, 41:25-29; EX1003, ¶125.



EX1005, Fig. 2 (annotated)

As another example, Figure 2 illustrates that user applications 92 and 96 (shaded red) see the aggregated unused storage resources on filesets 66 and 68 of the global file system as the separate “F:” and “N:” drives, respectively, each of which are represented as and appear to clients and application programs as separate NAS devices whose primary operational purpose is for providing NAS services. EX1005, Fig. 2; EX1003, ¶¶89-93, 126-127.

1(d) “processing one or more data storage or access workloads for the distributed processing system by accessing data from or storing data to at least a portion of the available amount of storage resources to provide NAS service to a client device coupled to a network”

In the previous IPR, the Patent Owner did not dispute that Carter discloses this limitation and the Board found that this limitation is disclosed. IPR2017-00439, Papers 9, 21 (Patent Owner failing to address Petitioner’s argument); EX1020 at 41 (Board finding Carter discloses claim 1); *see also* EX1003, ¶¶131-133, 118.

Carter discloses that storage or access workloads are processed by the “shared memory subsystem,” which is a component of the client agent. EX1005, Fig. 2. Specifically, Carter discloses that “[e]ach node 212a-212c through the shared memory subsystem 220 can access each page 224a-224d stored in the virtual shared memory 222.” EX1005, 29. Further, Carter discloses that “[e]ach page 224a-224d represents a unique entry of computer data stored within the virtual shared memory 222,” and thus, using the shared memory subsystem, each page 224a-224d can be

accessed or modified by each one of the nodes 212a-212c. *Id.* Carter discloses that each node can also “store additional pages of data within the virtual shared memory 222,” and “[e]ach newly stored page of data can be accessible to each of the other nodes 212a-212c.” *Id.*; EX1003, ¶132. Thus, the shared memory subsystem “process[es] one or more data storage or access workloads for the distributed processing system by accessing data from or storing data to at least a portion of the available amount of storage resources” as the claim requires. IX.B.1.(b), *supra*; EX1005, 29; EX1003, ¶132.

Carter’s “data control program”—another component of the client agent—uses the storage and access capabilities of the shared memory subsystem to “provide NAS service to a client device coupled to a network.” Specifically, the data control program uses the low-level storage services provided by the shared memory subsystem to maintain a “structured storage system, such as a file system, that is adapted to maintain a structured store of data,” which is provided to “external clients.” EX1005, 10, 35; EX1003, ¶133.

1(e) “enabling at least one of the selected distributed devices to function as a location distributed device to store location information associated with data stored by the selected distributed devices through use of the respective client agents for the particular distributed device”

In the previous IPR, the Patent Owner did not dispute that Carter discloses this limitation and the Board found this limitation is disclosed. IPR2017-00439,

Papers 9, 21 (Patent Owner failing to address Petitioner's argument); EX1020 at 41 (Board finding Carter discloses claim 1).

Carter's shared memory subsystem enables every node in the network to "function as a location distributed device" that stores "location information associated with data stored...through use of the respective client agents" (e.g., as illustrated by the hierarchical page directory structure 340 shown in Figs. 10 and 11). EX1005, 39-41, Figs. 10-11; EX1003, ¶135. For example, Carter explains that the shared memory space is logically divided into sections of fixed size (e.g., "4k bytes, of the virtual memory space"), each of which is associated with a globally unique identifier, also known as a global address signal. EX1005, 4; EX1003, ¶135-136. Carter further explains that any node (*i.e.*, a requesting node) connected to the system can use the global address signal corresponding to a particular block (what Carter refers to as a "4k page") to access the data stored in that block or write data to that block for later use. EX1005, 31-32; EX1003, ¶135-136. This global address signal is a virtual address because it is "decoupled from (*i.e.*, unrelated to) the physical and virtual address spaces of the underlying computer hardware." EX1005, 4; EX1003, ¶135. Carter discloses that the memory subsystem is able to translate a global address signal into the address of the physical storage (e.g. RAM or a hard disk) using a hierarchical structure of directory pages that is stored in the shared memory system. *Id.*, 5, 38-40, Fig. 10. EX1003, ¶135. These directory pages are

“location information” associated with the data stored in the shared memory by the various distributed devices. EX1005, 41, Figs. 10-11; EX1003, ¶135.

Further, Carter teaches that “stor[ing] location information associated with data” is provided “through use of the respective client agent[]” as the claim requires. The “shared memory subsystem” and “data control program”—which are components of the larger “client agent”—are “cooperating elements” that run on and enable each distributed device to provide this functionality. Specifically, the “data control program” on each node interfaces with the “shared memory subsystem” to “access and store data within the addressable shared memory space 20” using the directory structure discussed above. EX1005, 10, 39, 41, Figs. 8, 10-11; EX1003, ¶¶136-138

2. Claim 13

Claim 13 is a system claim that parallels the limitations of method claim 1. Carter renders obvious claim 13 for similar reasons as claim 1. EX1003, ¶¶98-138.

13(pre) “A system comprising ...”

Carter teaches systems. EX1005, 4 (disclosing “computer network systems”); Abstract, 27; EX1003, ¶140. The Board in the prior IPR agreed.

EX1020 at 24-41.

13(a) “a plurality of distributed devices configured to be coupled to a network, wherein the distributed devices include respective client agents configured to process respective portions of a workload for the distributed processing system”

Element 13[a] is satisfied for the same reasons as element 1[a].

13(b) “the respective client agents including respective instances of a software-based network attached storage (NAS) component wherein the NAS component is configured to: assess unused storage resources of the distributed devices”

Element 13[b] is satisfied for the same reasons as element 1[b].

13(c) “wherein the NAS component is configured to: ... allocate respective available amount of unused storage resources in selected distributed devices of the plurality of distributed devices;

Element 13[c] is satisfied for the same reasons as element 1[c]. EX1003, ¶¶143-146.

Additionally, as the Board found, “allocate” in Carter has the same meaning as in the ’827 patent, and Carter discloses this limitation. EX1020 at 39-41. Indeed, Carter discloses a policy element 290 that assesses used, unused, and allocated storage space “to determine the availability of resources, such as...hard-disk space, available RAM and other such resources” and “dynamically distribut[es] [these] available resources” to other devices on the network. EX1020 at 39; EX1005, 36:27-37:2; 9:18-21; EX1103, ¶145.

Carter discloses several other instances of allocating. For example, Carter discloses “reserv[ing] (sometimes much) larger than requested memory space for a file, and substantiates the storage by allocating backing storage page by page.” EX1005, 19:14-17. Carter further discloses a “distributed directory manager...for allocating a portion of the addressable memory space” and that “data streams are

mapped onto pages allocated from the addressable shared memory space 20 for storage.” EX1005, 26:18-19; *id.*, 15:9-11, 16:21-23, 17:18-19; EX1003, ¶146.

13(d) “wherein the NAS component is configured to: ... represent that the selected distributed devices comprise NAS devices having the respective available amounts of storage resources;”

Element 13[d] is satisfied for the same reasons as element 1[c].

13(e) “wherein the NAS component is configured to: ... process workloads associated with data storage and access by accessing data from and storing data into portions of the storage resources in the selected distributed devices to provide NAS service to a client device”

Element 13[e] is satisfied for the same reasons as element 1[d].

13(f) “wherein the respective client agents are configured to enable at least one of the selected distributed devices to function as a location distributed device to store location information for data stored by the selected distributed devices”

Element 13[f] is satisfied for the same reasons as element 1[e].

3. Claims 2 and 14

Claim 2 recites the method of claim 1, “further comprising enabling at least one of the selected distributed devices to function as a stand-alone dedicated NAS device through use of the respective client agents for the particular distributed device.” Claim 14 is of similar scope but is a system claim. Carter discloses a “NAS device” and “respective client agents for the particular distributed device” for the reasons discussed above regarding limitations 1(a) and 1(c).

For the reasons discussed above, the limitation claims 2 and 14 add to claims 1 and 13 do not materially alter the question of invalidity, and thus, Patent Owner is

estopped from arguing the validity of the similar-scope limitations of claims 2 and 14.

As explained above, Petitioner applies Patent Owner’s interpretation of “stand-alone dedicated NAS device” from the previous IPR, which is a “device that separately appears to network-connected devices as a dedicated NAS device.” The Board in the prior IPR found that Carter discloses this interpretation as Carter “uses software that causes *individual nodes* (e.g., commercially available PCs) *to represent, or mimic, dedicated NAS devices*” and “each node may be responsible for storing a particular element or elements of the structured store of data.” EX1020 at 36, citing EX1005, 9:18-19; 10:6-26, Fig. 1. The Board also found that “Carter’s Figure 11 shows multiple commercially available PC nodes (406a and 406b) *each mimicking NAS devices* by making storage available.” *Id.*, citing EX1005, 41:12–21.

Moreover, Carter discloses providing the recited “stand-alone dedicated NAS device” functionality “through use of the respective client agents for the particular distributed device” as the claim requires. As the Board found regarding claim 1, the shared memory subsystem is part of the software package on each node that is the “client agent” that provides the required NAS device functionality. EX1005, 39–42, Figs. 7, 10, 11. Thus, Carter discloses a client agent that enables a node to function

as a stand-alone dedicated NAS device, as required by claims 2 and 14. EX1003, ¶150-151.

Additional aspects of Carter also disclose this limitation. For example, Figure 2 shows “[t]wo network nodes 84 and 86 [] accessing the addressable shared memory space 20 . . . via the logical drives 90 and 94,” “F:” and “N:” respectively. EX1005, 13:8-14, Fig. 2; EX1003, ¶152-153. It would have been obvious to a POSITA that the “F:” and “N:” drives are wholly distinct logical drives, and that data on each of these drives could be hosted on separate physical devices that could operate completely independent of one another without knowledge of the other drive or the set of memory pages associated with the other drive. EX1003, ¶152. For example, a POSITA would have understood that the policy controller (EX1005, 36:29-37:2, Fig. 8 (290)) could be used to implement the F: and N: drives as physically separate devices so that the storage access utilized by one client node could be isolated and to prevent any impact on the performance of storage access by another client node. EX1003, ¶152. Additionally, these N: and F: drives would separately appear to the application programs as dedicated NAS devices whose primary operational purpose is to provide NAS service. EX1003, ¶153. Indeed, each drive is separately addressed by application programs by reference to the respective drive name “F:” or “N:” as each drive appears to application programs as a device whose primary operation purpose is providing NAS services. EX1003, ¶153; EX1005, 40-41 (“during a

memory access operation, the memory subsystem [] passes to the responsible node . . . the address of the page being accessed” from “a node that stores a copy of the page being accessed”); Fig 11.

The '827 patent also discloses potential additional features of a standalone, dedicated NAS device such as having “[its] own storage, processing, connectivity and management resources” (EX1001, 2:10-14), may act without the knowledge of other NAS devices connected to the network (*id.*, 3:17-21, 43:47-52), and may lack data location libraries 2106A and 2106B (*id.* 43:30-33, 4:56-58). EX1003, ¶156.

Carter discloses NAS devices with these additional features. For example, network memory device 226 is separately connected to the network and contains its own storage “network disk,” which are resources that are independently managed by network memory device 226. *Id.*, Fig. 6; EX1003, ¶157. Second, network memory device 226 does not include the memory subsystems 232a-232c illustrated within nodes 212a-212c and thus acts without knowledge of other NAS devices connected to the network. EX1003, ¶157. Third, because network memory device 226 does not include the memory subsystem, it does not have access to the hierarchical structure of directory pages (i.e., “location information”) thus lacks data location libraries. EX1005, 5, 38-41, Figs. 10, 11; EX1003, ¶¶158-160.

Carter also discloses that not all nodes would have to store directory information, and thus these nodes function as stand-alone dedicated NAS devices

without data location libraries. Specifically, Carter discloses that “a node selector can select a responsible node for each page...so the leaf entry 356 provide information of the address and responsible node of the page being accessed.” EX1005, 40; Fig. 10. “[T]his directory tracks ownership and responsibility for data to provide a level of indirection between the directory and the physical location of the data.” EX1005, 40. Thus, a POSITA would have understood that not all of Carter’s nodes need to contain directory information, and the directory information could be stored anywhere in the system, including on only a subset of the nodes. EX1003, ¶¶159-169.

4. Claims 9 and 21

Claims 9 and 21, which claims 10-12 and 22-24 depend upon, were invalidated over Carter. EX1020. Claims 10-12 and 22-24 were not asserted in that proceeding. The analysis of claim 9 reflects the ground that was successful in that IPR.

Claim 9 depends from claim 1 and further requires “managing storage resources for the selected distributed devices with a storage priority control that facilitates use of the available amount of storage resources for the selected distributed devices.” As the Board found, Carter’s shared memory subsystems discloses this limitation as it includes logic to assess the “*availability of resources*, such as printer capabilities, *hard-disk space*, available RAM and other such

resources.” EX1020, citing EX1005, 36-37; EX1003, ¶¶171-173. The shared memory subsystem also includes policy element 290 that can act as a “policy controller [that] can employ any of the suitable heuristics to direct the elements, such as the paging controller, disk directory manager, and other elements to dynamically *distribute the available resources.*” EX1005, 36-37. EX1003, ¶171.

Carter also discloses that the shared memory subsystem “includes persistent storage for portions of the distributed memory” and “can provide integrated control of persistent storage devices and electronic memory.” EX1005, 30:15-25, Fig. 6. As one example of the “integrated control”—i.e., the recited “storage priority control that facilitates use of the available amount of storage”—the shared memory subsystem controls what “portions of the distributed shared memory to move between persistent and electronic memory *depending on predetermined conditions, such as recent usage.*” EX1005, 30:26-28; EX1003, ¶172.

Carter further explains that “the memory subsystems 232a-232c can include a hierarchy manager that provides hierarchical control for the distribution of data” by “controlling the migration controller, and policy controller...to perform hierarchical data migration and load balancing, such that data migrates primarily between computers of the same group, and passes to other groups in hierarchical order.” Carter explains that “[r]esource distribution is similarly managed.” EX1005, 31; EX1003, ¶¶173.

Carter discloses claim 21 for the same reasons. EX1020 at 48-51.

5. Claims 10-12 and 22-24

Claims 10 and 11 depend from claim 9, and claim 12 depends from claim 10.

Claim 22 depends from claim 21, and claims 23 and 24 depend from claim 22.

a. Claim 10

Claim 10 recites “further comprising accepting selection of the storage priority control, by the client device, that comprises a parameter.” As discussed above regarding Claim 9, Carter discloses that the shared memory subsystem provides integrated control of persistent storage devices—i.e., the recited “storage priority control”—by (1) controlling what “portions of the distributed shared memory to move between persistent and electronic memory depending on predetermined conditions, such as recent usage” and (2) “controlling...hierarchical data migration and load balancing, such that data migrates primarily between computers of the same group, and passes to other groups in hierarchical order.” EX1005, 30-31; EX1003, ¶¶147-150, 152. The identification of which “portions of the distributed shared memory to move,” whether data should be migrated “between computers of the same group,” whether data should be “pass[ed] to other groups in hierarchical order,” and how to “load balanc[e],” are each examples of storage priority control parameters, as they identify information that is needed to execute the relevant storage priority control. EX1003, ¶153.

Moreover, a POSITA would have understood, and found it obvious, that the client device, which may have a portion of its storage resources allocated to provide NAS services and is the target of the “storage priority control,” would have to accept the parameter defining the relevant storage priority control when executing the storage priority control. EX1001, 2:66-3:1, 42:36-41; 30-31; EX1003, ¶¶174-177.

b. Claim 11

Claim 11 recites “further comprising prioritizing one or more of storage or deletion of data using the storage priority control that comprises storage priority level schemes.” The ’827 patent clarifies that “storage priority level scheme [] prioritize[s] data storage and deletion.” EX1001, 3:36-37. Carter discloses claim 11. EX1003, ¶¶178-183.

As discussed above regarding claims 9 and 10, Carter discloses “storage priority control” that controls “what portions of the distributed shared memory to move” and controls “hierarchical data migration and load balancing.” EX1005, 30-31; EX1003, ¶¶170-177, 179. Each of “mov[ing],” “migrat[ing],” and “load balancing” data involves removing data from its original physical location and storing data at a new physical location. For example, a POSITA would have understood that “electronic memory” (which is used interchangeably with “volatile” or “RAM” memory in Carter) would typically be more limited in quantity but faster than “persistent memory” (e.g. disk). Further, since electronic memory is not

persistent, it would typically be used to hold temporary (cached) copies of data as a way of enhancing performance. A POSITA would have understood that a storage priority level scheme was used to prioritize which portions of the distributed shared memory should be used as a temporary cache and which portions of the shared memory to use as persistent memory. Further, choosing to place a copy of one application's data into the electronic memory of a node would clearly involve the deletion of data from a different application to make room for this new data. Thus, Carter discloses "prioritizing one or more of storage or deletion of data using the storage priority control that comprises storage priority level schemes" as required by the claim. EX1003, ¶180-181.

As another example, Carter further discloses utilizing the electronic memory and/or excess persistent memory for a distributed cache: "commonly used pages are copied and distributed, and rarely used pages are shuffled off to disk" and "directory pages will migrate to those nodes that access them most, providing a degree of self-organization that reduces network traffic." EX1005, 41:8-11; *id.*, 13:3-4 (disclosing "automatically and dynamically migrat[ing] data to account for varying network usage and traffic patterns"). Here, a POSITA would have understood that the "rarely used pages are shuffled off to disk" because disk is slower, allowing less used data to be deleted from electronic memory (i.e., deprioritizing the data) and freeing this precious resource to be utilized for more popular data (i.e., prioritizing the data).

This migration of data, including, copying, distributing, and deleting data, according to Carter's storage priority control and storage priority level schemes meets the requirement of claim 11. EX1003, ¶181-182.

c. Claim 12

Claim 12 recites "wherein the managing storage resources further comprises marking directly data or files." Carter discloses "a global address index for accessing the cached pages of the distributed memory and the attributes based on that page." EX1005, 37:16-18. Specifically, Carter explains that "[o]ptionally, *with each listed global address, the index can further provide attribute information.*" This attribute information may include (1) "a version tag in the index that is representative of the version of the data," (2) "a dirty bit representative of whether the RAM cached data is a copy of the data held on disk, or whether the RAM cached data has been modified but not yet flushed to disk," (3) "a volatile bit to indicate if the page is backed by backing store in persistent memory," and "(4) *other such attribute information useful for managing the coherency of the stored data.*" EX1005, 37:16-26; EX1003, ¶185. A POSITA would have understood from this disclosure at the time of invention that the "storage priority control" and "storage priority schemes" discussed regarding claims 9-11 above could use this attribute information to mark data or files that should be moved, migrated, load balanced, or distributed. EX1003, ¶185.

For example, a POSITA would have understood from these disclosures in Carter that it “would be useful for managing the coherency of the stored data” to mark each “commonly used page[]” with Carter’s disclosed attribute information such as a “dirty bit,” to indicate that these pages have been “copied and distributed” because they are so frequently accessed. EX1005, 41:8-9; 37:16-26; EX1003, ¶186. For the same reason, a POSITA would have understood that attribute information could be used to mark “rarely used pages” that should be “shuffled off to disk.” EX1005, 41:9; EX1003, ¶186.

As another example, a POSITA would have understood that Carter’s attribute information would be used to mark each “directory page” by identifying the “nodes that access them most” so the “directory pages will migrate to those nodes.” EX1005, 41:9-10; EX1003, ¶186. As another example, a POSITA would have understood that Carter’s attribute information could be used to mark each data file that is sent to persistent storage based on “recent usage.” EX1005, 30:27-28; EX1003, ¶186.

In each of these examples, a POSITA would have understood that Carter’s attribute information would be used together with the storage priority control to manage the coherency of the stored data as taught by Carter. EX1005, 37:21-26; EX1003, ¶186.

d. Claim 22

Claim 22 recites “[t]he system of claim 21, wherein the storage priority control comprises a parameter selectable by the client device.” Carter renders obvious claim 22 for similar reasons as claims 9-12 and 21. EX1003, ¶¶187-188.

As discussed regarding those claims, Carter discloses “storage priority control” that controls “what portions of the distributed shared memory to move” and controls “hierarchical data migration and load balancing.” EX1005, 30-31; EX1003, ¶¶179-182. Where a portion of a client device is providing NAS service, and the client devices’ storage memory subsystem is deciding whether and how to “move,” “migrat[e],” or “load balan[ce]” data, the client device “selects” the appropriate storage priority control parameter. EX1005, 30-31; EX1003, ¶¶179-182. Indeed, the identification of which “portions of the distributed shared memory to move,” whether data should be migrated “between computers of the same group,” whether data should be “pass[ed] to other groups in hierarchical order,” and how to “load balanc[e],” are each examples of storage priority control parameters that are selectable by the client device that is setting the storage priority control. EX1005, 30-31; EX1003, ¶¶179-182.

As another example, when a client device’s storage memory subsystem determines what portions of the distributed shared memory to move from electronic storage to persistent storage, it must “select” whether to make the decision based on

“recent usage” or some other parameter—e.g., whether the portion contains the most recent version of the data, whether the data is currently backed up in the persistent storage, or “whether the RAM cached data has been modified but not yet flushed to disk.” EX1005, 37:16-26; EX1003, ¶188.

e. Claims 23-24

Claims 23 and 24 recite substantively identical limitations to claims 11 and 12, except in the form of system claims. Carter discloses each of claims 23 and 24 for the same reasons as 11 and 12. EX1003, ¶189.

C. Ground 2: Claims 2 and 14 are obvious over OceanStore

Claims 1 and 13, which claims 2 and 14 depend upon, were invalidated over OceanStore. EX1019 (IPR2017-00374, Paper 53). Claims 2 and 14 were not asserted in that proceeding. The analysis of claim 1 below reflects the obviousness ground that was successful in that IPR.

1. Claim 1

1(pre)

This preamble does not limit the claim. Regardless, the Board in the prior IPR found that OceanStore discloses a “computer-implemented method.” EX1019 at 21-22, 27.

As the Board explained, OceanStore is designed for computers that are geographically dispersed and allows for continuous access to persistent information stored on remotely located storage devices. EX1019 at 21 (citing EX1003, Abstract.)

OceanStore also discloses support of legacy computer coding languages by providing familiar computer interfaces such as the Unix interface. *Id.* (citing EX1003, 191); EX1003, ¶¶220-225.

1(a)

The Board in the prior IPR found that OceanStore discloses this limitation. EX1019 at 22-23, 28.

OceanStore describes a widely-distributed network of devices including desktop workstations and servers with excess storage space to provide a “utility infrastructure designed to span the globe.” EX1006, Abstract, 190-191, Fig. 1; EX1003, ¶226. For example, computers that store data are referred to as “nodes” or “servers” and OceanStore clients access data by connecting to one or more “pools” of nodes. *Id.* OceanStore’s pools are also connected to each other, which allows data to “flow freely” between the pools. *Id.*; EX1003, ¶¶199, 226.

As Dr. Kubiawicz explains, a POSITA would have understood that the OceanStore’s various server functionalities (*e.g.*, assessing excess storage capacity, making that excess capacity available for use, handling of read and update requests for objects, routing of requests if needed) would be implemented by software, which is an “agent” running on the networked distributed devices consistent with the way “client agents” are described in the ’827 patent. EX1006, 197-198; EX1003, ¶¶227-230; EX1001, 33:19-21, 41:54-57, 43:26-30; Figs. 15, 20 (illustrating what a

client agent *should* contain, including “communication interface” 1508, “core agent” 1502, “project component” 1504, “task module” 1512,” and “user interface component”).

Moreover, OceanStore’s agent functionality includes the ability to process respective portions of a workload as recited by the claim. EX1003, 197-198 (“The base [OceanStore] API provides full access to OceanStore functionality in terms of sessions” where each “sessions is a sequence of reads and write to potentially different objects”); EX1003, ¶230.

1(b)

The Board in the prior IPR found that OceanStore discloses this limitation based on Dr. Kubiawicz’s testimony. EX1019 at 28-30.

As Dr. Kubiawicz explains, a POSITA would appreciate that the various server functionalities of OceanStore would be implemented by software. EX1003, ¶231. Further, a POSITA would appreciate that an OceanStore software-based component would assess the unused storage resources of the device (e.g., desktop workstation with excess storage space) when making this excess storage available for use by other devices on the network, as depicted by Figure 1. EX1003, ¶231; EX1006, 190,191 (“OceanStore hopes to take advantage of this excess of storage space to make the management of data seamless and carefree” and “[e]ach user...could consume storage and bandwidth resources form many different providers,” each configured with OceanStore software).

Additionally, as Dr. Kubiadowicz explains, the '827 patent never mentions “assess” or describes any particular way in which the assessment is performed. EX1003, ¶232. As the applicant explained during prosecution, “it would have been obvious to a POSITA that the client agent programs must assess unused or underutilized capabilities for it to be able to configure the devices...using their spare and/or unused capabilities.” EX1002, 358. EX1003, ¶232. OceanStore provides at least similar or more disclosure than the '827 patent for this operation. EX1003, ¶¶231-233.

I(c)

The Board in the prior IPR found that OceanStore discloses this limitation. EX1019 at 30-33.

As discussed in Ground 1, the Board explained “‘NAS device’ as defined in the '827 patent ‘broadly refers to a device that makes data storage resources available to network-connected user devices.’ EX1019 at 32-33 (citing EX1001, 2:31-34, 42:36-43:5). The OceanStore software causes the pool devices (e.g., desktop workstation with excess storage space, servers, and/or other pool devices) to represent, or mimic, NAS devices by making storage resources available to network-connected user devices. EX1006, 191-199. EX1003, ¶234. As the Board found, “[i]n OceanStore, multiple storage resources are provided by an aggregation of “pools” of storage, *each of which* can cache or serve data to the user devices” and “*each pool*

device makes its storage available for use by others as depicted in Figures 2 and 3.” EX1019 at 33 (citing EX1006, 191-193). Thus, a POSITA would appreciate the OceanStore software causes each pool device to appear, function, and act as a NAS device and “ha[s] an available amount of storage resources related to the unused or under-utilized storage resources for the selected distributed devices.” EX1003, ¶¶234-239.

1(d)

Patent Owner did not dispute that OceanStore discloses this limitation and the Board found that OceanStore renders claim 1 obvious. IPR2017-00374, Paper 9, 24; EX1019 at 37-38.

OceanStore’s software provides the ability to process one or more data storage or access workloads when processing “sessions,” where “[a] session is a sequence of reads and writes” to a portion of the available amount of storage resources, thereby providing NAS service to a client device coupled to a network. EX1006, 197-198; EX1003, ¶241. A POSITA would readily appreciate that the read and write requests of these sessions are directed to the excess storage space of, for example, a desktop workstation that is part of OceanStore’s “pool” of storage—i.e., the session read and write requests are directed “to at least a portion of the available storage resources.” EX1003, ¶242. Moreover, these OceanStore sessions “provides NAS service to a client device” coupled to the OceanStore network as depicted in Figure 1. EX1006,

190 (Abstract); *id.*, 197-198 (API provides full access to functionality of sessions). EX1003, ¶242.

I(e)

Patent Owner did not dispute that OceanStore discloses this limitation and the Board found that OceanStore renders claim 1 obvious. IPR2017-00374, Paper 9, 24; EX1019 at 37-38.

OceanStore satisfies this limitation in multiple ways. EX1003, ¶¶243-246. OceanStore's local, probabilistic algorithm uses an "attenuated Bloom filter" and is a "fully distributed" routing and location algorithm. EX1006, 193-94. Each node determines if it has the requested data, and if it does not, it routes the request to another node using locally stored routing information. EX1006, 193. Thus, in OceanStore, the distributed devices acting as NAS devices can both hold data in available storage, and hold distributed location information to forward a request if the query cannot be satisfied. EX1006, 193. This OceanStore functionality mirrors the '827 patent which discloses that some of the distributed devices providing software-based NAS services can provide "both storage and data location functionality, thereby acting both to provide data storage services and distributed data location services." EX1001, 44:9-12. EX1003, ¶244.

OceanStore's global "wide-scale distributed data location" algorithm teaches this limitation for similar reasons as it causes a node to use locally-stored routing

information to forward the request to another node until eventually a root node for the object is reached, which then points to the node storing the object. EX1006, 194; EX1003, ¶245.

OceanStore further discloses that stored objects themselves can include pointers to other versions of the object and thus are “location information associated with the data.” EX1006, 192 (“Certain OceanStore objects act as directories...we allow directories to contain pointers to other directories.”).

OceanStore also discloses that “[e]very addressable entity in the OceanStore (e.g. floating replica, archival fragment, or client) is identified by one or more GUIDs” and that “the role of the OceanStore routing layer is to route messages directly to the closest node that matches the predicate and has the desired GUID.” EX1006, p. 193. A POSITA would appreciate the pointers include these “global unique identifiers” (GUIDs) to objects and that GUIDs are information used to locate other versions of the object. EX1006, 193; EX1003, ¶246.

2. Claim 13

Claim 13 is a system claim parallel to method claim 1. As the Board found, OceanStore discloses every limitation of claim 13 for similar reasons as claim 1. EX1019 at 27-37.

13(pre)

Ocean Store is a system with the servers (“pools”) depicted in Figure 1. EX1006, 191, Fig. 1. EX1003, ¶¶222-225, 249.

13(a)

Element 13[a] is satisfied for the same reasons as element 1[a].

13(b)

Element 13[b] is satisfied for the same reasons as element 1[b].

13(c)

Although the term “allocating” is recited in claim 13, there is no specific method of allocating required by the claim or disclosed in the specification. *Id.* Thus, OceanStore discloses this limitation for the same reasons it discloses claim 1[c]. *Id.*; EX1003, ¶¶234-239, 252-255.

Specifically, OceanStore discloses that desktop workstations running OceanStore with excess storage would allocate the excess storage for use by other OceanStore users (i.e., making it available for other devices to use) but would also reserve other workstation storage for purely local use, e.g., to store computer programs, local operating system and the like. EX1003, ¶¶252-255. The Board found that the purpose of OceanStore is to provide storage space in servers around the world, allowing data to be “cached anywhere” (EX1006, 190 (Abstract)), and credited Dr. Kubiawicz’s opinion that a POSITA “would clearly understand that you need to allocate data in order to represent that you have an available amount of unused or under-utilized storage.” EX1019 at 34-37; EX1003, ¶255.

13(d)

Element 13[d] is satisfied for the same reasons as element 1[c].

13(e)

Element 13[e] is satisfied for the same reasons as element 1[d].

13(f)

Element 13[f] is satisfied for the same reasons as element 1[e].

3. Claims 2 and 14

OceanStore discloses a “NAS device” and “respective client agents for the particular distributed device” for the reasons discussed above regarding limitations 1(a) and 1(c).

As explained above, Petitioner applies Patent Owner’s interpretation of “stand-alone dedicated NAS device” from the previous IPR, which is a “device that separately appears to network-connected devices as a dedicated NAS device.” And the Board in the prior IPR found that OceanStore discloses this interpretation as “[i]n OceanStore, multiple storage resources are provided by an aggregation of “pools” of storage, *each of which* can cache or serve data to the user devices.” EX1019 at 33 (citing EX1006, 191) (emphasis in original). The Board also found that “*each pool device* makes its storage available for use by others as depicted in Figures 2 and 3.” *Id.* OceanStore further allows a user’s email to migrate closer to his client and use locally cached email, thereby indicating that the OceanStore clients see more than a single NAS device.” *Id.* (citing EX1006, 192, 193, 199); EX1003, ¶260.

Moreover, OceanStore discloses providing the recited “stand-alone dedicated NAS device” functionality “through use of the respective client agents for the particular distributed device” as the claim requires. As the Board found, the OceanStore software running on each node would be considered the “client agent” that provides the required NAS device functionality. *Id.*, 22-23 (citing EX1006, 191). Thus, OceanStore discloses a client agent that enables a node to function as a stand-alone dedicated NAS device, as required by claims 2 and 14. EX1003, ¶261. As Dr. Kubiawicz explains, the OceanState application is equally a “client agent” under either of the constructions proposed in the district court litigation. *Id.*, ¶¶227-230, 261.


As discussed above, the ’827 patent discloses additional possible features of a “standalone dedicated NAS device” including that it may “contain [its] own storage, processing, connectivity and management resources,” (EX1001, 2:10-14) and may lack data location libraries (*id.* 43:30-33, 4:56-58). EX1003, ¶¶262-273. OceanStore also discloses these features. OceanStore is partitioned into a variety of components, including “Pool servers” for storage, and networking and routing layers for data location and routing. EX1006, 193 (Section 4.3). EX1003, ¶264. Each “pool server” contains its own “storage, processing, connectivity and management resources.” *Id.* And it is the networking and routing layers—and not the individual pools of storage resources—which has the knowledge and ability to locate data and

act as a data location service. EX1006, 193 (“the network takes an active role in routing messages” such that the “networking layer consults a distributed, fault-tolerant data structure that explicitly tracks the location of all objects” and “the “routing layer...routes messages directly to the closest node that matches the predicate and has the desired GUID” for the requested data.); EX1003, ¶¶264-273 (discussing OceanStore’s Attenuated Bloom Filters and Global Algorithm for data location and routing).

D. Ground 3: Claims 2 and 14 are obvious over Carter in view of Gibson

Carter in view of Gibson alternatively renders obvious claims 2 and 14. EX1003, ¶¶285-293. Gibson states that its “principal contribution” is to “demonstrate the potential of network-attached storage devices,” which “provide[] high bandwidth by directly attaching storage to the network,” “avoid[] file server store-and-forward operations,” and “allow[] data transfers to be striped over storage and switched-network links.” EX1007, 272; EX1003, ¶286. Gibson describes several types of NAS devices including “Network-Attached Secure Disks (NASD)” which “are drives that support independent client access to drive object services.” EX1007, 272; EX1003, ¶287. As shown in Figure 3, each of Gibson’s NASDs is a “standalone dedicated NAS device” as each NASD makes data storage resources available to client devices over the local area network, the primary purpose of each NASD is to provide NAS service (as required by the construction of “dedicated NAS

device”), and each NASD is “separately appear[s] to network-connected user devices as a dedicated NAS device” EX1007, 275-276; Fig. 3. EX1003, ¶1288.



Moreover, Gibson discloses that each NASD functions as a stand-alone dedicated NAS device accessed directly by the client (i.e., recited “client agent”) by first consulting the “file manager”—similar to the use of the “responsible node” of Carter. For example, Gibson discloses that “[b]efore issuing its first read of a file, the client authenticates itself with the file manager and requests access to the file.” EX1007, 276, 269. After access is granted, the client may directly access the NASD drive containing the file until the client’s access is revoked. *Id.* Thus, Claims 2 and 14 are obvious in view of the combination of Carter and Gibson.

As discussed above, a POSITA would have been motivated to combine Gibson’s NASDs with Carter as both relate to the same field of endeavor and both disclose the implementation of NAS devices to improve efficiency, scalability and data sharing capacity of distributed networks. Moreover, a POSITA would have

been motivated to expand Carter's network-attached storage (NAS) devices (e.g., network memory device 226 illustrated in Fig. 6) by implementing Gibson's NASDs to work in parallel with Carter's aggregated virtual storage space. EX1003, ¶¶274-277. Gibsons' NASDs would attach directly to Carter's network similar to memory device 226 shown in Fig. 6 and the NASDs would be managed by Carter's "client agent" software package, which would be expanded to include Gibson's file manager functionality. EX005, Fig. 6; EX1007, Fig. 3 (illustrating NASD access over LAN and describing file manager). A POSITA would have recognized that this combined system would realize the benefits of improving network efficiency, scalability, and data sharing capacity as taught by each of Carter and Gibson in a single combined system. Moreover, combining the teachings of these references would have involved nothing more than the use of known features to yield predictable results. EX1003, ¶¶274-277. The '827 patent itself states that it is a "relatively simple implementation [] to have the devices simply mimic the functionality of stand-alone NAS devices." EX1001, 42:66–43:1. A POSITA thus would have been motivated to combine Gibson's NASDs with Carter's disclosed system, and would have had a reasonable expectation of success. EX1003, ¶¶274-277.

IX. PTAB DISCRETION SHOULD NOT PRECLUDE INSTITUTION

A. Fintiv Analysis

Discretionary denial under *Fintiv* is not warranted. Although there is currently a district court case in which the '827 patent is asserted, the six-factor test set forth in *Apple Inc. v. Fintiv, Inc.*, IPR2020-00019, Paper 11 at 5–6 (PTAB Mar. 20, 2020) (precedential), favors institution. Ultimately, given the evidence of unpatentability described above, the *Fintiv* factors should not lead the Board to discretionarily deny institution.

Factor 1. Petitioner has filed a Motion to Transfer Venue to the Northern District of California. Ex. 1022. “Given NDCA’s high stay rates, NDCA may well stay [Petitioner’s] case ... pending IPR.” *Google LLC v. Parus Holdings, Inc.*, IPR2020-00847, Paper 9 at 12 (PTAB Oct. 21, 2020) (holding Factor 1 neutral where a similar motion was pending). Regardless of the venue, Petitioner intends to file a motion to stay the Current Litigation pending resolution of this IPR should this Petition be instituted. As such, this factor is neutral.

Factor 2. Trial in the Current Litigation is currently scheduled to begin on November 17, 2025. Ex. 1007, dkt. entry 28 (setting Jury Selection for 11/17/2025). However, as described above, Petitioner have moved to transfer the venue to the Northern District of California, which would, in all likelihood, result in trial being delayed past the date of any Final Written Decision in an IPR instituted based on the

current Petition. Accordingly, any currently scheduled trial date is highly speculative. Thus, this factor remains neutral.

Factor 3. The Current Litigation is still in its early stages: the parties only recently filed their joint claim construction statement, only limited discovery requests have been served, and expert discovery remains far on the horizon. Ex. 1023, 3. Although the parties have exchanged initial infringement and invalidity contentions, “this investment is not significant.” *Vudu*, IPR2020-01688, Paper 16 at 12. This factor thus weighs against the Board exercising its discretion to deny the Petition.

Factor 4. There is minimal risk of overlap with arguments made in the district court because Petitioner hereby stipulates that, if IPR is instituted, Petitioner will not pursue in the district court litigation any ground for invalidity that was raised in the instituted IPR. PTAB case law holds that such a stipulation “mitigates to some degree the concerns of duplicative efforts between the district court and the Board, as well as concerns of potentially conflicting decisions.” *Sand Revolution II, LLC v. Continental Intermodal Group-Trucking LLC*, IPR2019-01393, Paper 24 (PTAB June 16, 2020) (informative). This factor thus weighs against denying the Petition.

Factor 5. This factor has no weight because the Petitioner and Patent Owner are the same parties as in the district court. *Weatherford U.S. v. Enventure Global Tech., Inc.*, IPR2020-01684, Paper 16 at 11–13 (PTAB Apr. 14, 2021) (holding only

that “this factor does not weigh against exercising discretion to deny institution” where the Petitioner and Patent Owner are, respectively, the defendant and complainant in the parallel district court proceeding).

Factor 6. The prior art presented in this Petition renders the Challenged Claims unpatentable as obvious. The merits of Petitioner’s arguments are extremely compelling because they lead to a strong conclusion that one or more claims are unpatentable by a preponderance of the evidence. In the previous IPRs challenging claims of the ’827 Patent on the basis of the same prior art, the Board held “Petitioner has shown by a preponderance of the evidence that claims 1, 3–9, 13, 15, 16, and 18–21 of the ’827 patent are unpatentable under 35 U.S.C. § 103 as obvious in view of OceanStore” and “Petitioner has shown by a preponderance of the evidence that (1) claims 1, 3, 4, 6, 8, 9, 13, 15, 16, 18, 20, and 21 under 35 U.S.C. § 102(b) as unpatentable over Carter.” *EMC Corp. v. Intellectual Ventures I*, IPR2017-00374, Paper 53 at 59; *EMC Corp. v. Intellectual Ventures I*, IPR2017-439, Paper 50 at 61. This factor therefore weighs strongly against discretionary denial. In totality, the *Fintiv* factors weigh strongly in favor of institution.

B. General Plastics Analysis

The factors set forth in General Plastic do not support discretionary denial here because Petitioner has not previously submitted any petitions for IPR on the ’827 patent—i.e., this is not a “follow-on petition.” *General Plastic Indus. Co., Ltd.*

v. Kaisha, IPR2016-01357, Paper 19 at 8–9 (PTAB Sept. 6, 2017) (precedential) (summarizing factors for exercising discretion to deny “follow-on petitions”). While it is true that there have been petitions with these references before, those petitions were filed by different petitioners (EMC Corporation, Lenovo (United States), and NetApp. Inc.) against a different Patent Owner (Intellectual Ventures I), none of which are party to the current litigation between Byteweavr and Databricks. *See Ex. 1023*. As such, there is no “significant relationship” between Petitioner and the petitioners in the previously filed petition. *See Valve Corp. v. Elec. Scripting Prods., Inc.*, IPR2019-00062, Paper 11 at 10 (PTAB Apr. 2, 2019) (precedential) (finding a “significant relationship” only because the two petitions were filed by co-defendants in the same litigation related to the same technology).

“Under USPTO policy and precedent, General Plastic has not been extended to any cases in which the first and second petitioners do not have a significant relationship.” *Videndum Prod. Sol., Inc. v. Rotolight Ltd.*, IPR2023-01218, Paper 12 at 6 (PTO-Ofc. of Dir., Apr. 19, 2024). Accordingly, “where ... the first and second petitioners are neither the same party, nor possess a significant relationship under Valve, General Plastic factor one necessarily outweighs the other General Plastic factors.” *Id.* Accordingly, the Board should therefore not exercise its discretion to deny institution under *General Plastics*.

C. §325(d): Advanced Bionics

Under *Advanced Bionics* and applying the *Becton Dickinson* factors, the Board should not deny institution under §325(d). Here, *Advanced Bionics*, Part 1, is not satisfied because the same or substantially the same art or arguments have not been “previously...presented to the Office[]” under 35 U.S.C. §325(d). During the original prosecution of the ‘827 Patent, none of the references relied upon were cited or relied upon by the examiner. *See Ex. 1002*. Although the ‘827 has previously been subject to IPR based on the art presented herein, the currently challenged claims were not challenged in the prior IPRs. *EMC*, IPR2017-00374, Paper 53, 2-3; *EMC*, IPR2017-00439, Paper 50, 2-3.

X. CONCLUSION

For the reasons given above, Petitioner hereby requests *inter partes* review under 35 U.S.C. § 311 and 37 C.F.R. § 42.101 of the ‘827 patent.

*Petition for Inter Partes Review
of U.S. Patent No. 8,275,827*

DATED: March 13, 2025

Respectfully Submitted,

/Robin A. Snader/

Robin Snader (Reg. No 66,085)

Erise IP, P.A.

7015 College Blvd, Ste 700

Overland Park, KS 66214

robin.snader@eriseip.com

Attorney for Petitioner Databricks, Inc.

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

The undersigned certifies pursuant to 37 C.F.R. §42.24 that the foregoing Petition for *Inter Partes* Review, excluding any table of contents, mandatory notices under 37 C.F.R. §42.8, certificates of service or word count, or appendix of exhibits, contains 13891 words according to the word-processing program used to prepare this document (Microsoft Word).

DATED: March 13, 2025

/Robin A. Snader/

Robin Snader (Reg. No 66,085)
Erise IP, P.A.
7015 College Blvd, Ste 700
Overland Park, KS 66214
robin.snader@eriseip.com

*Attorney for Petitioner Databricks,
Inc.*

CERTIFICATE OF SERVICE

The undersigned certifies service pursuant to 37 C.F.R. §§ 42.6(e) and 42.105(a), (b) on the Patent Owner via FedEx overnight mail of a copy of this Petition for *Inter Partes* Review and supporting materials at the Patent Owner at the correspondence address of record for the '827 patent:

STERNE, KESSLER, GOLDSTEIN & FOX P.L.L.C.
1101 K Street, NW
10th Floor
WASHINGTON, DC 20005
UNITED STATES

Further, a courtesy copy of this Petition for *Inter Partes* Review was sent via email to Patent Owner's litigation counsel:

Brandon V. Zuniga (bzuniga@bosfirm.com)
Marcus Benavides (mbenavides@bosfirm.com)
Mark Douglass (mdouglass@bosfirm.com)
Terry Afif Saad (tsaad@bosfirm.com)
Vivek Ramachandran (vramachandran@bosfirm.com)
Jeffrey Ray Bragalone (jbragalone@bosfirm.com)

DATED: March 13, 2025

By: /Robin A. Snader/

Robin Snader (Reg. No 66,085)
Erise IP, P.A.
7015 College Blvd, Ste 700
Overland Park, KS 66214
robin.snader@eriseip.com

APPENDIX: CHALLENGED CLAIM LISTING

No.	Limitation
1(pre)	A computer-implemented method comprising:
1(a)	configuring a distributed processing system of a plurality of distributed devices coupled to a network, wherein the plurality of distributed devices include respective client agents configured to process respective portions of a workload for the distributed processing system,
1(b)	wherein the respective client agents for particular distributed devices of the plurality of distributed devices have corresponding software-based network attached storage (NAS) components configured to assess unused or under-utilized storage resources in selected distributed devices of the plurality of distributed devices;
1(c)	representing with the corresponding software-based NAS component that the selected distributed devices respectively comprise NAS devices having an available amount of storage resources related to the unused and under-utilized storage resources for the selected distributed devices;
1(d)	processing one or more of data storage or access workloads for the distributed processing system by accessing data from or storing data to at least a portion of the available amount of storage resources to provide NAS service to a client device coupled to the network; and
1(e)	enabling at least one of the selected distributed devices to function as a location distributed device to store location information associated with data stored by the selected distributed devices through use of the respective client agents for the particular distributed device.
2	The method of claim 1, further comprising enabling at least one of the selected distributed devices to function as a stand-alone dedicated NAS device through use of the respective client agents for the particular distributed device.
9	The method of claim 1, further comprising managing storage resources for the selected distributed devices with a storage priority control that facilitates use of the available amount of storage resources for the selected distributed devices.

10	The method of claim 9, further comprising accepting selection of the storage priority control, by the client device, that comprises a parameter.
11	The method of claim 9, further comprising prioritizing one or more of storage or deletion of data using the storage priority control that comprises storage priority level Schemes.
12	The method of claim 10, wherein the managing storage resources further comprises marking directly data or files.
13(pre)	A system comprising:
13(a)	a plurality of distributed devices configured to be coupled to a network, wherein the distributed devices include respective client agents configured to process respective portions of workloads for the distributed processing system,
13(b)	the respective client agents including respective instances of a software-based network attached storage (NAS) component, wherein the NAS component is configured to: assess unused storage resources of the distributed devices;
13(c)	wherein the NAS component is configured to: ... allocate respective available amount of unused storage resources in selected distributed devices of the plurality of distributed devices;
13(d)	wherein the NAS component is configured to: ... represent that the selected distributed devices comprise respective NAS devices having the respective available amounts of storage resources;
13(e)	wherein the NAS component is configured to: ... process workloads associated with data storage and access by accessing data from and storing data into portions of the storage resources in the selected distributed devices to provide NAS service to a client device: and
13(f)	wherein the respective client agents are configured to enable at least one of the selected distributed devices to function as a location distributed device to store location information for data stored by the selected distributed devices.
14	The system of claim 13, wherein at least one of the respective client agents are configured to enable at least one of the selected distributed devices to function as a stand-alone dedicated NAS device.
21	The system of claim 13, further comprising a storage priority control configured to facilitate use of the respective available

	amount of unused storage resources for the selected distributed devices.
22	The system of claim 21, wherein the storage priority control comprises a parameter selectable by the client device.
23	The system of claim 22, wherein the storage priority control comprises storage priority level Schemes that prioritize one or more of storage or deletion of data.
24	The system of claim 22, wherein the storage priority control comprises a priority marking directly given to data or files.