

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

EMC CORPORATION, LENOVO (UNITED STATES) INC.,
NETAPP, INC.,
Petitioner,

v.

INTELLECTUAL VENTURES I LLC,
Patent Owner.

Case IPR2017-00439
Patent 8,275,827

Before JEFFREY S. SMITH, GEORGIANNA W. BRADEN, and
DANIEL J. GALLIGAN, *Administrative Patent Judges*.

BRADEN, *Administrative Patent Judge*.

FINAL WRITTEN DECISION

35 U.S.C. § 318

I. INTRODUCTION

We have jurisdiction to hear this *inter partes* review under 35 U.S.C. § 6, and this Final Written Decision is issued pursuant to 35 U.S.C. § 318(a). For the reasons that follow, we determine that Petitioner has shown by a preponderance of the evidence that claims 1, 3–6, 8, 9, 13, and 15, 16–18, 20, and 21 of U.S. Patent No. 8,275,827 (Ex. 1101, “the ’827 patent”) are unpatentable.

A. Procedural History

EMC Corporation, Lenovo (United States) Inc., and NetApp, Inc. (“Petitioner”) filed a Petition (Paper 5, “Pet.”) to institute an *inter partes* review of the challenged claims of the ’827 patent. Intellectual Ventures I LLC (“Patent Owner”) timely filed a Preliminary Response (Paper 9, “Prelim. Resp.”). Pursuant to 35 U.S.C. § 314(a), an *inter partes* review was originally instituted for (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 Carter¹; and (2) claims 5 and 17 under 35 U.S.C. § 103(a) as unpatentable over Carter and Pitzel.² See Paper 11 (“Dec. to Inst.”), 42.

On April 24, 2018, the Supreme Court held that a decision to institute under 35 U.S.C. § 314 may not institute on fewer than all claims challenged in the petition. *SAS Inst., Inc. v. Iancu*, 138 S. Ct. 1348 (2018). Thereafter, we modified the institution decision to institute on all of the challenged claims and all of the grounds presented in the Petition. Accordingly, this

¹ PCT Pub No. WO 98/22881, App. No. PCT/US97/21460, May 28, 1998 (Ex. 1140, “Carter”).

² US Pat. No. 7,062,765 B1, filed May 25, 1999, issued June 13, 2006 (Ex. 1142, “Pitzel”).

inter partes review also includes review of (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; (2) claims 5 and 17 under 35 U.S.C. § 103(a) as unpatentable over Carter and Pitzel; and (3) claims 4, 6, 7, 16, 18, and 19 under 35 U.S.C. § 103(a) as unpatentable over Carter and Condor.³ See Paper 48, 2; *see also* Ex. 3001 (email “advis[ing] the Board that Patent Owner will waive its right to file a [Patent Owner Response] and will not request an oral hearing on the newly instituted grounds in these proceedings. Both parties have conferred and believe this closes the record and allows the Board to issue a final written decision based on the current record.”).

After institution of trial, Patent Owner filed a Patent Owner Response (Paper 21, “PO Resp.”), to which Petitioner filed a Reply (Paper 29, “Reply”). Patent Owner also filed Objections to Evidence (Papers 14, 17) and a Motion for Observation on Cross-Examination (Paper 32). Petitioner responded to Patent Owner’s Motion for Observation on Cross-Examination. Paper 36. In addition, and with our authorization, Patent Owner filed a list of alleged improper reply arguments (Paper 37), to which Petitioner filed a Response (Paper 38).

An oral argument was held on March 5, 2018, together with co-pending Case IPR2017-00374. A transcript of the oral argument is included in the record. Paper 49 (“Tr.”).

B. Related Proceedings

The parties inform us that the challenged patent is the subject of two district court proceedings in the District of Massachusetts, captioned

³ Litzkow, M.J. et al., Condor – A Hunter of Idle Workstations, Computer Sciences Technical Report #730, Dec. 1987 (Ex. 1117, “Condor”).

Intellectual Ventures I LLC et al. v. Lenovo Group Ltd. (1:16-cv-10860-IT) and *Intellectual Ventures I, LLC et al. v. NetApp, Inc.* (1:16-cv-10868-IT). Pet. 2; Paper 8, 1 (Patent Owner’s Mandatory Notices).

C. The ’827 Patent

The ’827 patent was filed on April 13, 2001, and is titled “Software-Based Network Attached Storage Services Hosted on Massively Distributed Parallel Computing Networks.” Ex. 1101, Title. It describes systems and methods for software-based network attached storage (NAS) services hosted on a massively-distributed network of devices. *Id.* at Abstract. The ’827 patent discloses that the term “NAS device” generally refers to “a specialized file server that connects to the network to provide storage capacity to network-connected users,” (*id.* at 2:18–20), and, in the context of the ’827 patent, “broadly refers to a device that makes data storage resources available to network-connected user devices,” (*id.* at 2:31–34). The ’827 patent further discloses that a NAS device typically takes the form of a standalone dedicated hardware device that contains its “own storage, processing, connectivity, and management resources.” *Id.* at 2:12–14. According to the ’827 patent, ordinary desktop computers often sit idle for long periods of time but have the same specifications (e.g., a Pentium processor, a hard drive, and an Ethernet connection) and often greater functional capabilities than dedicated NAS devices. *Id.* at 2:46–55. The ’827 patent proposes making these otherwise unused storage resources available to other devices connected to the network. *Id.* at 2:66–3:1, 42:36–41.

In one embodiment, the '827 patent describes a distributed NAS system, shown in Figure 21 reproduced below, that includes user devices 152 and NAS client devices 108. *Id.* at 43:6–15.




FIG. 21

As illustrated above in Figure 21 of the '827 patent, the server systems may include user server systems 154D, which have attached database systems 2112, and NAS server systems 104 with attached NAS database systems 2108 and centralized data location library 2110. Ex. 1101, 43:22–26. To provide NAS functionality, the '827 patent describes client devices 108A, 108B, . . . 108C as running a client agent that provides NAS functionality by running a NAS component as one of the projects for the distributed device. *Id.* at 43:26–30. According to the '827 patent, some of these NAS devices, such as NAS client devices 108A and 108B, may store distributed data location information, such as data location libraries 2106A and 2106B. *Id.* at 43:30–33.

The '827 patent discloses that the NAS client devices within client device fabric 2104 can act as stand-alone NAS devices and user devices within user device fabric 2102 can directly access them as such. *Id.* at 43:44–46. Additionally, at least some of the NAS client devices within

client device fabric 2104 include data location information or libraries that are used to provide links or directions to data distributed on various NAS client devices 108A, 108B, . . . 108C. *Id.* at 43:59–63. These links help provide a NAS device fabric from which user server systems 154D and user devices 152A, 152B, . . . 152C within user device fabric 2102 can draw storage resources. *Id.* at 43:63–66.

The '827 patent also contemplates that server systems 104 facilitate the NAS functionality, for example, by providing centralized management, control, and support for the NAS architecture. *Id.* at 44:24–27. According to the '827 patent, NAS servers systems 104 may include, for example, file storage database systems 2108 and centralized data location library 2110. *Id.* at 44:27–29. Additionally, the '827 patent discloses that the user devices, which can be any type of computer devices, are configured to access data stored on the NAS client devices. *Id.* at 44:60–64.

D. Illustrative Claims

As noted above, an *inter partes* review was instituted as to claims 1, 3–9, 13, and 15–21 with claims 1 and 13 being independent. Claims 1 and 13 are illustrative of the challenged claims, and are reproduced below:

1. A computer-implemented method comprising:
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,
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;

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;

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

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.

Ex. 1101, 46:28–56.

13. A system comprising:

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, 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;

allocate respective available amount of unused storage resources in selected distributed devices of the plurality of distributed devices;

represent that the selected distributed devices comprise respective NAS devices having the respective available amounts of storage resources;

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

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.

Id. at 47:28–48:7.

II. DISCUSSION

A. Claim Construction

In an *inter partes* review, claim terms in an unexpired patent are interpreted according to their broadest reasonable construction in light of the specification of the patent in which they appear.⁴ 37 C.F.R. § 42.100(b); *see Cuozzo Speed Techs., LLC v. Lee*, 136 S. Ct. 2131, 2144–46 (2016) (“We conclude that the regulation represents a reasonable exercise of the rulemaking authority that Congress delegated to the Patent Office.”). Under that standard, and absent any special definitions, we give claim terms their ordinary and customary meaning, as would be understood by one of ordinary skill in the art at the time of the invention. *In re Translogic Tech., Inc.*, 504 F.3d 1249, 1257 (Fed. Cir. 2007). An inventor, however, may provide a meaning for a term that is different from its ordinary meaning by defining the term in the specification with “reasonable clarity, deliberateness, and precision.” *In re Paulsen*, 30 F.3d 1475, 1480 (Fed. Cir. 1994). Limitations, however, are not to be read from the specification into the claims. *In re Van Geuns*, 988 F.2d 1181, 1184 (Fed. Cir. 1993). In addition,

⁴ We are not aware of any evidence of record in this proceeding that would lead to a different outcome under the *Phillip* claim construction standard. *See Phillips v. AWH Corp.*, 415 F.3d 1303 (Fed. Cir. 2005).

the Board may not “construe claims during [an *inter partes* review] so broadly that its constructions are *unreasonable* under general claim construction principles.” *Microsoft Corp. v. Proxyconn, Inc.*, 789 F.3d 1292, 1298 (Fed. Cir. 2015), *overruled on other grounds by Aqua Prods., Inc. v. Matal*, 872 F.3d 1290 (Fed. Cir. 2017).

Thus, when construing the claims at issue (1) we look at the language of the claims themselves, (2) we consult the patent’s specification to help clarify the meaning of claim terms, because the claims “must be read in view of the specification, of which they are a part,” *Trading Techs. Int’l, v. eSpeed, Inc.*, 595 F.3d 1340, 1352 (Fed. Cir. 2010) (quoting *Markman v. Westview Instruments, Inc.*, 52 F.3d 967, 979 (Fed. Cir. 1995) (en banc), *aff’d*, 517 U.S. 370 (1996)), and (3) we review “the patent’s prosecution history in proceedings in which the patent has been brought back to the [US Patent and Trademark Office] for a second review” (*Microsoft Corp.*, 789 F.3d at 1298).

The parties dispute the proper construction of the limitation “[r]epresenting with the corresponding software-based NAS component that the selected distributed devices respectively comprise NAS devices” as recited in claims 1 and 13. *See* Pet. 10–11; PO Resp. 7–11; Reply 8–14. Specifically, the parties proposed the following constructions:

Petitioner	“using the corresponding software-based NAS components of the selected distributed devices to mimic dedicated NAS devices”
Patent Owner	“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”
--	--

Patent Owner argues the explicit language of the claims requires that each distributed device includes its own NAS device. PO Resp. 8. Patent Owner notes that the term “respectively” has “a well-settled meaning of ‘[s]eparately or individually and in the order already mentioned.’” *Id.* (citing Exhibit 2001, 1219). According to Patent Owner, the word “respectively” as used in claims 1 and 13 modifies the clause “selected distributed devices.” *Id.* Thus, Patent Owner argues that the explicit and plain language of the claims requires a representation that each of the “selected distributed devices,” separately or individually, appears to the network-connected user devices as a “NAS device.” *Id.* Patent Owner further argues that its construction is consistent with the specification of the ’827 patent because “[t]he ’827 patent defines a ‘NAS device’ as ‘a device that makes data storage resources available **to network-connected user devices.**’” *Id.* (quoting Ex. 1101, 2:31–34). Patent Owner contends that “[w]hen this definition of NAS device is applied, the ‘*representing*’ phrase requires a representation that each individual ‘selected distributed device’ comprises ‘a device that makes data storage resources available to network-connected user devices.’” *Id.*

Patent Owner contends that an integrated system that appears as a single device cannot satisfy this requirement, as the individual distributed devices in this integrated system would not separately make or appear to make their resources available to network-connected user devices. Patent Owner argues that “the explicit language of the claims excludes this

embodiment by requiring the representation that the ‘selected distributed devices **respectively** comprise NAS devices.’” *Id.* at 10 (citing Ex. 1101, 46:41–46 (emphasis added).) Instead, according to Patent Owner, the claim language recites the embodiment described first in this passage—that each individual NAS device “brings **its** resources to the network **with the appearance of a dedicated NAS device.**” *Id.* (citing Ex. 1101, 42:59–65 (emphasis added by Patent Owner)).

We agree with Patent Owner regarding the ordinary meaning of the term “respectively.” Specifically, we understand that the word “respectively” can be an adverb that means “for each separately or in turn, and in the order mentioned.” The use of “respectively” as an adverb in the English language requires parallelism. *See* Ex. 3002. In other words, there must be a direct one-to-one correspondence and specific referents must be *individually* identified before there can be a “parallel list” of attribute values applicable to *each, in that order*. The term “respectively” also can be an attributive adjective that means “concerning or appropriate to each of several individually.” *Id.* In the present case, the term “respectively” appears to modify “the selected distributed devices.” Because specific referents are not individually identified, we understand the term to act as an attributive adjective modifying each of several selected distributed devices.

Although we agree with Patent Owner regarding the use of the term “respectively,” we do not understand its use to require “the selected distributed devices to cause each selected distributed device to separately appear to network-connected user devices as a dedicated NAS device.” We cannot construe individual terms in isolation. Rather, we must consider claims as a whole in view of the specification. Specifically, we must

consider the implications of other terms recited in claims 1 and 13 (e.g., “location information”) and how the claims are supported by the disclosed embodiments.

As a general rule, “there is a strong presumption against a claim construction that excludes a disclosed embodiment.” *See In re Katz Interactive Call Processing Patent Litig.*, 639 F.3d 1303, 1324 (Fed. Cir. 2011). Several exceptions to this presumption apply. For example, a claim may be interpreted to exclude embodiments “where those embodiments are clearly disclaimed in the specification . . . or prosecution history.” *Oatey Co. v. IPS Corp.*, 514 F.3d 1271, 1277 (Fed. Cir. 2008). Similarly, where the disputed term is not present in the other independent claims of the patent-in-suit, it is permissible to construe the term so as to exclude an embodiment. This is because the other claim terms “leave[] open the possibility that claims not at issue in [the claim construction] encompass omitted embodiments.” *Helmsderfer v. Bobrick Washroom Equip., Inc.*, 527 F.3d 1379, 1383 (Fed. Cir. 2008). Otherwise, “where claims can reasonably [be] interpreted to include a specific embodiment, it is incorrect to construe the claims to exclude that embodiment, absent probative evidence [to] the contrary.” *Oatey Co.*, 514 F.3d at 1277.

In this case, the disclosure in the ’827 patent teaches at least three embodiments using the selected distributed devices in a variety of ways, including as “stand-alone” devices (Ex. 1101, 43:44–47), as part of a “NAS device fabric” (*id.* at 43:59–63, 44:9–12), and in a “server assisted [implementation]” (*id.* at 44:24–27). 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 running a client agent program and

NAS component that brings its resources to the network with the appearance of a dedicated NAS device or **as part of an integrated system that appears as a single or dedicated NAS device.**” See Ex. 1101, 42:59–65 (emphasis added). The ’827 patent further discloses:

A relatively simple implementation is to have the devices simply mimic the functionality of stand-alone NAS devices, while a more complex implementation is to provide a storage fabric across the Intranet or Internet where hundreds or thousands of devices work cooperatively as NAS devices in conjunction with server systems that assist in controlling and managing the storage function of the NAS fabric.

See Ex. 1101, 42:66–43:5. Yet, Patent Owner’s construction intentionally reads out the embodiments covering a NAS fabric and the “server assisted implementation” of the devices. Patent Owner specifically argues:

[T]he explicit language of the claims excludes this embodiment by requiring the representation that the “selected distributed devices **respectively** comprise NAS devices.” (’827 patent, 46:41–46 (emphasis added).) Instead, the claim language recites the embodiment described first in this passage—that each individual NAS device “brings **its** resources to the network **with the appearance of a dedicated NAS device.**” (’827 patent, 42:59–65 (emphasis added).)

PO Resp. 10. Patent Owner’s position is contrary to the recited claims when considered as a whole and runs counter to prevailing case law. See *Oatey Co.*, 514 F.3d at 1277; *Helmsderfer, Inc.*, 527 F.3d at 1383.

Both independent claims 1 and 13 recite “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.” As Petitioner argues, the NAS fabric embodiment disclosed in the ’827 patent is “the **only** embodiment in which [Patent

Owner's declarant], Dr. Shenoy, could identify disclosure of the recited 'location information' that was required to secure allowance of the claims." Reply 10–11 (citing Ex. 1144, 43:16–44:16); see Ex. 1144, 57:7–21. According to Dr. Shenoy, one reason that a "stand-alone implementation does not require the data location library to be used" is "presumably because it's a stand-alone implementation, and clients connect to the stand-alone NAS devices directly." Ex. 1144, 53:18–54:8. Thus, Patent Owner's argument that each selected distributed device must appear separately to network-connected user devices as a dedicated NAS device improperly imports selected limitations from the specification into the claim and ignores additional required elements recited in the claims. See *In re Van Geuns*, 988 F.2d 1181, 1184 (Fed. Cir. 1993).

Accordingly, we do not agree with Patent Owner's position because Patent Owner's proposed construction narrows unnecessarily the scope of the challenged claims. Furthermore, as detailed below, even if we accepted Patent Owner's proposed construction, it would not alter our determination regarding the patentability of the challenged claims. No further construction is necessary for our analysis.

B. Principles of Law

A claim is unpatentable under 35 U.S.C. § 102 if a prior art reference discloses every limitation of the claimed invention, either explicitly or inherently. *Glaxo Inc. v. Novopharm Ltd.*, 52 F.3d 1043, 1047 (Fed. Cir. 1995). If the prior art reference does not expressly set forth a particular element of the claim, that reference still may anticipate if that element is "inherent" in its disclosure. *In re Robertson*, 169 F.3d 743, 745 (Fed. Cir. 1999). To establish inherency, extrinsic evidence, when relied upon, "must

make clear that the missing descriptive matter is necessarily present,” or inherent, in the single anticipating reference. *Continental Can Co. v. Monsanto Co.*, 948 F.2d 1264, 1268 (Fed. Cir. 1991). “Inherency, however, may not be established by *probabilities* or *possibilities*. The mere fact that a certain thing may result from a given set of circumstances is not sufficient.” *Id.* at 1269 (emphasis added) (quoting *In re Oelrich*, 666 F.2d 578, 581 (CCPA 1981)). Inherent anticipation, however, does not require recognition in the prior art. *Schering Corp. v. Geneva Pharmaceuticals*, 339 F.3d 1373, 1377 (Fed. Cir. 2003) (“*Continental Can* does not stand for the proposition that an inherent feature of a prior art reference must be perceived as such by a person of ordinary skill in the art before the critical date.”).

A claim is unpatentable under 35 U.S.C. § 103(a) if “the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains.” *KSR Int’l Co. v. Teleflex Inc.*, 550 U.S. 398, 406 (2007). The question of obviousness is resolved on the basis of underlying factual determinations, including: (1) the scope and content of the prior art; (2) any differences between the claimed subject matter and the prior art; (3) the level of skill in the art; and (4) objective evidence of nonobviousness, i.e., secondary considerations. *See Graham v. John Deere Co. of Kansas City*, 383 U.S. 1, 17 (1966).

“[I]t is error to reach a conclusion of obviousness until all [the *Graham*] factors are considered.” *Apple Inc. v. Samsung Elecs. Co., Ltd.*, 839 F.3d 1034, 1048 (Fed. Cir. 2016) (en banc) (citations omitted). “This

requirement is in recognition of the fact that each of the *Graham* factors helps inform the ultimate obviousness determination.” *Id.*

“In an [*inter partes* review], the petitioner has the burden from the onset to show with particularity why the patent it challenges is unpatentable.” *Harmonic Inc. v. Avid Tech., Inc.*, 815 F.3d 1356, 1363 (Fed. Cir. 2016) (citing 35 U.S.C. § 312(a)(3) (requiring *inter partes* review petitions to identify “with particularity . . . the evidence that supports the grounds for the challenge to each claim”)). This burden of persuasion never shifts to Patent Owner. *See Dynamic Drinkware, LLC v. Nat’l Graphics, Inc.*, 800 F.3d 1375, 1378 (Fed. Cir. 2015) (discussing the burden of proof in *inter partes* review). Furthermore, Petitioner cannot satisfy its burden of proving obviousness by employing “mere conclusory statements.” *In re Magnum Oil Tools Int’l, Ltd.*, 829 F.3d 1364, 1380 (Fed. Cir. 2016).

Thus, to prevail in an *inter partes* review, Petitioner must explain how the proposed combinations of prior art would have rendered the challenged claims unpatentable. At this final stage, we determine whether a preponderance of the evidence of record shows that the challenged claims are anticipated by or would have been obvious over the cited prior art.

We analyze the instituted grounds of unpatentability in accordance with the above-stated principles.

C. Level of Ordinary Skill in the Art

In determining whether an invention would have been obvious at the time it was made, we consider the level of ordinary skill in the pertinent art at the time of the invention. *Graham*, 383 U.S. at 17. “The importance of resolving the level of ordinary skill in the art lies in the necessity of maintaining objectivity in the obviousness inquiry.” *Ryko Mfg. Co. v.*

Nu-Star, Inc., 950 F.2d 714, 718 (Fed. Cir. 1991). The person of ordinary skill in the art is a hypothetical person who is presumed to have known the relevant art at the time of the invention. *In re GPAC, Inc.*, 57 F.3d 1573, 1579 (Fed. Cir. 1995). The level of ordinary skill in the art may be reflected by the prior art of record. *Okajima v. Bourdeau*, 261 F.3d 1350, 1355 (Fed. Cir. 2001). Factors that may be considered in determining the level of ordinary skill in the art include, but are not limited to, the types of problems encountered in the art, the sophistication of the technology, and educational level of active workers in the field. *GPAC*, 57 F.3d at 1579. In a given case, one or more factors may predominate. *Id.* Generally, it is easier to establish obviousness under a higher level of ordinary skill in the art. *Innovation Toys, LLC v. MGA Entm't, Inc.*, 637 F.3d 1314, 1323 (Fed. Cir. 2011) (“A less sophisticated level of skill generally favors a determination of nonobviousness . . . while a higher level of skill favors the reverse.”).

Petitioner argues that a person of ordinary skill in the art relevant to the '827 patent would have had “at least a bachelor’s degree in computer science or electrical engineering and 3–5 years of professional experience in computer systems design, or a master’s or doctorate and 1–2 years of professional experience in computer systems design, or equivalent academic experience.” Pet. 4–5. According to Petitioner, “[s]uch 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.” *Id.* at 5. Petitioner’s declarant, John D. Kubiawicz, Ph.D. (“Dr. Kubiawicz”), proffers the same level of skill as that argued by Petitioner. Ex. 1102 ¶ 20.

Patent Owner contests Petitioner's definition of a person of ordinary skill in the art, specifically arguing that "Petitioners' proposed requirements especially with regard to work experience, are higher than can be reasonably supported." PO Resp. 5. Patent Owner notes that Petitioners' declarant, Dr. Kubiatowicz, testified on cross-examination that "somebody coming directly out of college, likely, would need **a couple** of extra years to understand this," and that "a little bit of professional experience might help them -- would help them understand the details involved in this case." *Id.* (citing Ex. 2002, 25:13–14; 26:2–5). Patent Owner's declarant, Dr. Shenoy, testified that, "at the time of the '827 patent, a person of ordinary skill in the art should have had 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." Ex. 2003 ¶ 19. According to Dr. Shenoy, "[l]ack of experience can be remedied by additional education or vice versa." *Id.*

Based on our review of the '827 patent, the types of problems and solutions described in the '827 patent and cited prior art, and the testimony of Dr. Kubiatowicz and Dr. Shenoy, we adopt and apply Patent Owner's definition of a person of ordinary skill in the art, specifically that a person of ordinary skill in the art at the time of the claimed invention would have had "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." *See id.* We also determine 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.

D. Alleged Anticipation of Claims 1, 3, 4, 6, 8, 9, 13, 15, 16, 18, 20, and 21 of the '827 Patent by Carter

Petitioner contends claims 1, 3, 4, 6, 8, 9, 13, 15, 16, 18, 20, and 21 of the '827 patent are unpatentable under 35 U.S.C. § 102(b) as anticipated by Carter. Pet. 18–55. Patent Owner disputes Petitioner's contentions. PO. Resp. 11–50. For reasons that follow, we determine Petitioner has established by a preponderance of the evidence that these claims of the '827 patent are unpatentable under 35 U.S.C. § 102(b) as anticipated by Carter.

1. Overview of Carter

Carter is a PCT publication titled “Remote Access in a Globally Addressable Storage Environment.” Ex. 1140, Title. Carter is directed to “a globally addressable storage environment that allows a plurality of networked computers to access data by addressing even when the data is stored on a persistent storage device such as a computer hard disk and other traditionally non-addressable data storage devices.” *Id.* at Abstract. Carter discloses that “computers can be located on a single computer network or on a plurality of interconnected computer networks such as two local area networks (LANs) coupled by a wide area network (WAN).” *Id.* According to Carter, “[t]he globally addressable storage environment allows data to be accessed and shared by and among the various computers on the plurality of networks.” *Id.* Additionally, Carter discloses that at least one of the computers can be located remote from the network and coupled thereto by a communications adapter such as a modem. *Id.* at 5:26–28. Carter further contemplates that “[t]he globally addressable data storage system preferably replicates and/or migrates data stored in the persistent data storage device among two or more of the computers based on, for example, accesses by the

computers of the globally addressable data storage system to obtain data stored in the persistent data storage device.” *Id.* at 5:28–6:2.

In Carter, computer groups are referred to as “nodes” and each of the nodes runs a software program that accesses and manages a structured store of data that is stored in an addressable shared memory space. *Id.* at 9:14–17. “The structured store of data may be stored in an addressable shared memory or the structured store may be stored in a more traditional fashion. For example, each node may be responsible for storing a particular element or elements of the structured store of data.” Ex. 1140, 9:17–19. One embodiment disclosed in Carter is shown in Figure 1, reproduced below.





Figure 1, above, illustrates a block diagram of a distributed addressable shared memory structured data storage system. *Id.* at 8:5–6. Carter discloses that network system 10 includes a plurality of network nodes 12a–12d and addressable shared memory space 20 that has portion 22 for storing a structured store of data 28. *Id.* at 10:6–8. According to Carter, each of nodes 12a–12d can include several sub-elements, such as processor 30a, data control program 32a, and shared memory subsystem 34a. *Id.* at 10:8–9. For the embodiment shown above in Figure 1, “two of the nodes, 12a and 12c,

include monitors that provide displays 40 and 42 graphically depicting the structured store of data 28 within addressable shared memory space 20.” *Id.* at 10:10–12. Carter discloses that addressable shared memory space 20 interconnects each of network nodes 12a–12d and provides each node 12a–12d with access to the structured store of data 28 contained within addressable shared memory space 20. *Id.* at 10:12–14. Carter contemplates that system 10 can provide, among other things, each network node 12a–12d with shared control over the structured store of data 28 and, therefore, system 10 can distribute control of the data store across the nodes of the network. Ex. 1140, 10:15–17. To this end, Carter states that each node of system 10, such as node 12a, includes data control program 32a that interfaces to shared memory subsystem 34a. *Id.* at 10:17–18.


Carter discloses that its shared memory subsystem is a software module that can couple between an operating system (“OS”) and a virtual shared memory and can form an interface between the OS and the virtual shared memory. *Id.* at 28:27–30. This allows the OS to access the virtual shared memory. *Id.* One embodiment of Carter’s shared memory subsystem is shown in Figure 8, reproduced below.



As shown in Figure 8, above, shared memory subsystem 270 includes remote operations element 274, local random access memory (“RAM”) cache 276, global RAM directory 280, global disk directory 284, configuration manager 288, policy element 290, local disk cache 294, physical memory 300, OS 216, and memory access request 312. Ex. 1140, 34:28–35:4. Carter discloses that “policy element 290 . . . can be a software module that acts as a controller to determine the availability of resources, such as printer capabilities, hard-disk space, available RAM and other such resources.” *Id.* at 36:27–29. Carter also states that policy controller 290 “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.” *Id.* at 36:29–37:2.

Carter indicates its system can provide nodes with access to distributed memory space by using the coordinated operation of a directory manager. *Id.* at 37:27–29. According to Carter, a directory manager can provide a directory structure that indexes the shared address space. *Id.* at

38:3–4. Figure 10, reproduced below, illustrates one embodiment of a directory structure disclosed in Carter.



As shown in Figure 10, above, directory 340 is formed from directory pages 342, 350–354, and 360–366, and provides location information to the pages of the distributed shared memory space depicted as pages 370–384.

Ex. 1140, 39:12–15. Carter states that directory page 342 acts like a root directory page and can be located at a static address that is known to each node coupled to the distributed address space. *Id.* at 39:16–17.

2. Analysis of Cited Art as Applied to Independent Claims 1 and 13

a. Preamble to Claims 1 and 13

Preamble language that merely states the purpose or intended use of an invention generally is not treated as limiting the scope of a claim. See *Boehringer Ingelheim Vetmedica, Inc. v. Schering-Plough Corp.*, 320 F.3d 1339, 1345 (Fed. Cir. 2003); *Rowe v. Dror*, 112 F.3d 473, 478 (Fed. Cir. 1997). When the limitations in the body of the claim rely upon or derive essential structure from the preamble, however, then the preamble acts as a

necessary component of the claimed invention and is limiting. *See Eaton Corp. v. Rockwell Int’l Corp.*, 323 F.3d 1332, 1339 (Fed. Cir. 2003).

The preamble of claim 1 recites a “computer-implemented method comprising,” and the preamble of claim 13 recites a “system comprising.” Ex. 1101, 46:28, 47:28. Petitioner first contends the preamble of claim 1 does not limit the claim under the claim’s broadest reasonable interpretation. Petitioner then contends that, nonetheless, Carter discloses a “computer-implemented method” as required by claim 1 and a system as required by claim 13. Pet. 32 (citing Ex. 1102 ¶ 102); 53 (citing Ex. 1102 ¶ 154).

Patent Owner does not address the preambles of independent claims 1 and 13. *See generally* PO Resp. The burden remains, however, on Petitioner to demonstrate unpatentability. *See Dynamic Drinkware*, 800 F.3d at 1378.

We agree with Petitioner, and we find, that Carter discloses a “computer-implemented method” and a system. Ex. 1102 ¶¶ 102, 154.

- b. *“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”*

Petitioner contends Carter discloses this limitation as recited by claim 1 with similar language recited in claim 13. Pet. 33 (citing Ex. 1102 ¶ 103), 53–54 (citing Ex. 1102 ¶ 155). As noted by Petitioner, Carter is directed to “provid[ing] computer network systems that have adaptable system configurations for dynamically exploiting distributed network resources” with “a plurality of network nodes that access a memory space storing a structured store of data, such as a structured file system or a

database.” Ex. 1140, 4, 9, *cited in* Pet. 33; *see* Ex. 1102 ¶ 105. Petitioner argues that

the systems of Carter are implemented using a computer program that comprises various software components that run on each of the distributed devices in the system, including the “shared memory subsystem” and the “data control program” described above, as well as additional components including, among others, a “distributor” component, a “disk directory manager,” and a “migration controller.”

Pet. 33–34 (citing Ex. 1140, 24). Petitioner further argues that, collectively, Carter’s software is a “client agent” for processing portions of a distributed workload. *Id.* at 34 (citing Ex. 1102 ¶ 107).

Petitioner relies on the testimony of its declarant, Dr. Kubiadowicz, to support its position. Dr. Kubiadowicz testifies that storage or access workloads are processed in Carter by the “shared memory subsystem,” which is a component of the client agent. Ex. 1102 ¶ 108 (citing Ex. 1140, Fig. 2). According to Dr. Kubiadowicz, “[t]he shared memory subsystems running on the distributed nodes communicate with each other over the network, and cooperate to manage the transfer and storage of data among the various physical storage devices that make up the distributed system. They serve to export storage resources to client devices over the network and are ‘NAS components’” *Id.* Dr. Kubiadowicz further testifies that “[t]he shared memory subsystem is responsible for processing storage workloads generated by the distributed devices.” *Id.* ¶ 109. Dr. Kubiadowicz then opines that “[a]ll of these storage and access capabilities are provided by the shared memory subsystem, which is a component of the software package described by Carter.” *Id.* (citing Ex. 1140, 29).

Patent Owner does not address this limitation of independent claims 1 and 13, but the burden remains on Petitioner to demonstrate unpatentability. *See Dynamic Drinkware*, 800 F.3d at 1378.

Based on our review of the entirety of the record before us, we agree with Petitioner and find that Carter discloses “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” as recited in claim 1 and similarly recited in claim 13. *See* Ex. 1140, 4, 9, 24, 29, 34; *see* Ex. 1102 ¶¶ 103, 105, 107–109, 155.

- c. *“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”*

Petitioner contends Carter discloses this limitation as recited by claim 1 and similarly recited in claim 13. Pet. 37 (citing Ex. 1102 ¶ 114), 54 (citing Ex. 1102 ¶ 156). Petitioner first argues

the shared memory subsystem is a component of the larger “client agent” provided by Carter, which allows low-level access to a shared addressable memory. By communicating over the network with other nodes in the Carter system, the shared memory subsystem makes data storage resources available to network-connected user devices, and thus, it is a “NAS component.”

Id. at 37 (citing Ex. 1140, 39–42, Figs. 7, 10, 11; Ex. 1102 ¶ 115). Petitioner then argues that the shared memory subsystem is implemented in software because it is a “software module.” *Id.* at 38 (citing Ex. 1140, 28; Ex. 1102

¶ 116). Petitioner further argues that the shared memory subsystem “assesses unused or under-utilized storage resources” because Carter describes nodes as having an “optional private memory device 218” that

provides the node with local storage that can be kept inaccessible to the other nodes on the network. Typically the private memory device 218 includes a RAM, or a portion of a RAM memory, for temporarily storing data and application programs and for providing the processor 214 with memory storage for executing programs. The private memory device 18 can also include persistent memory storage, typically a hard disk unit or a portion of a hard disk unit, for the persistent storage of data.

Id. (quoting Ex. 1140, 28). Petitioner, thus, asserts that “Carter teaches systems that distinguish between ‘unused’ storage space (*i.e.*, the storage space that is pooled to form the shared virtual address space) and ‘used’ storage space (*i.e.*, private storage that is reserved for the use of a particular node).” Pet. 38 (citing Ex. 1102 ¶ 117).

In addition, Petitioner argues Carter discloses this limitation because “the shared memory subsystem assesses the extent to which storage resources are being utilized and uses load balancing and resource distribution techniques to prevent underutilization.” Pet. 38. In support, Petitioner relies on the declaration of Dr. Kubiadowicz, who testifies that memory subsystems 232a–232c in Carter can include a hierarchy manager that provides hierarchical control for the distribution of data, which controls load balancing and resource distribution. Ex. 1102 ¶ 118 (citing Ex. 1140, 31). Dr. Kubiadowicz further testifies that “[t]he shared memory subsystem may also include 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” 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.’” *Id.* (alteration in original) (emphases omitted) (quoting Ex. 1140, 30, 36–37). Therefore, Dr. Kubiawicz opines that “the shared memory subsystem assesses the extent to which storage resources are being utilized and uses load balancing and resource distribution techniques to prevent underutilization.” *Id.* (citing Ex. 1140, 30).

Patent Owner contests Petitioner’s position. PO Resp. 19–29. Patent Owner argues that Petitioner presents two theories with respect to the “assessing” limitation of claims 1 and 13, both of which fail. *Id.* at 19. According to Patent Owner, the theory that Carter’s “private” storage performs the claimed assessment fails because the cited portion of Carter does not expressly disclose assessing unused or under-used storage resources and because there is no evidence that Carter’s “shared memory system” must perform some assessment to distinguish between private storage and non-private storage. *Id.* at 20. Patent Owner further argues that any distinguishing between “private or non-private storage has no bearing on whether such storage is ‘unused’ or ‘under-utilized’” because “the fact that storage is private does not mean it is ‘used’” and the fact that storage is non-private does not mean such storage is “‘unused.’” *Id.* at 21 (citing Pet. 38). Patent Owner insists that “distinguishing between private and non-private storage does not **necessitate** ‘assessing unused or under-utilized storage.’” *Id.* at 22 (citing Ex. 2003 ¶ 55). Patent Owner contends that a wide variety of different methods could be used by Carter to distinguish between private and non-private storage without performing any assessment of whether

storage is unused or underused. *Id.* (citing Ex. 2003 ¶ 55). Patent Owner notes, for example, that “Carter could reserve a static location and fixed amount of storage as ‘private storage’ on every node. In such a case, no assessment would be required as the private storage is located in the same place on every device, and has the same size.” *Id.* (citing Ex. 2003 ¶ 56).

For Petitioner’s second theory regarding Carter’s resource distribution, Patent Owner argues that Petitioner fails to demonstrate that the Carter’s policy controller assesses unused or underused storage resources as required by the claims. *Id.* Patent Owner further argues that distribution of resources could be done using heuristics based on any factor, not necessarily that of unused or underused storage resources. *Id.* Patent Owner, thus, contends that Petitioner has failed to show assessment is either explicitly or inherently disclosed in its second theory. *Id.*

Patent Owner also argues that Carter “states that the policy controller ‘direct[s] the elements, such as the paging controller, disk directory manager, and other elements **to dynamically distribute the available resources.**’” *Id.* (citing Ex. 1140, 36–37 (emphasis added by Patent Owner)). Patent Owner further argues that “dynamically distributing available resources” and assessing unused storage resources are “completely different options (i.e., distribution vs. assessment). *Id.* (citing Ex. 2003 ¶ 62). Patent Owner then argues that “Carter itself contradicts Petitioners’ conclusion” because “[t]he cited sentence states that the ‘paging controller’ and ‘disk directory manager’ distribute available resources. A paging controller and disk directory manager have nothing to do with assessing storage resources.” *Id.* at 26 (citing Ex. 2003 ¶ 63). Rather, according to Patent Owner, “the page controller can remap portions of the addressable

memory space between the volatile memories and persistent memories and the disk directory manager to track portions of the addressable memory space mapped to the persistent memories.” *Id.* (citing Ex. 1140, 25, 36). Patent Owner asserts that “Carter’s mention of heuristics also cannot salvage Petitioners’ argument” because “Carter does not disclose what ‘heuristics’ are used, let alone disclose that such ‘heuristics’ would be based on unused or underutilized storage resources.” *Id.*

Patent Owner also contends that Petitioner’s reliance on Carter’s shared memory caching embodiment is misplaced. *Id.* at 29. Patent Owner specifically notes that the Petition relies only on Carter’s statement that “the shared memory subsystem 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.’” *Id.* (citing Pet. 39; Ex. 1140, 30 (emphasis added by Petitioner)). Patent Owner argues that neither Petitioner nor its declarant, Dr. Kubiawicz, explain how nor why Carter’s mention of “recent usage” relates to an assessment of unused storage resources. *Id.* (citing Pet. 39; Ex. 1102 ¶ 118). According to Patent Owner, “[e]ven a cursory reading of Carter illustrates that ‘recent usage’ has nothing to do with assessing ‘unused storage resources.’” *Id.* (citing Ex. 2003 ¶ 69). Patent Owner asserts that “Carter’s reference to ‘recent usage’ refers to the concept that more frequently accessed data (*i.e.*, recently used), should be stored in faster speed RAM and not on slower speed hard disks (*i.e.*, it should be cached).” *Id.* (citing Ex. 2003 ¶ 69). Patent Owner, thus, concludes that Carter’s discussion of shared memory caching does not expressly, or inherently, disclose “assess[ing] unused or underutilized storage resources.” *Id.* (citing Ex. 2003 ¶ 69).

Despite Patent Owner's arguments, we find that Carter's policy element 290 (a software module) is configured to assess unused or under-utilized storage resources in selected distributed devices of the plurality of distributed devices. We understand policy element 290 to satisfy this limitation because Carter specifically discloses that policy element 290 "acts as a controller *to determine the availability of resources, such as printer capabilities, hard-disk space, available RAM and other such resources.*" Ex. 1140, 36:27–29 (emphases added). We credit the testimony of Dr. Kubiawicz, who during his deposition stated, "I would say this policy controller certainly can evaluate the hard disk space on the local node, but could also potentially evaluate storage throughout the system by communicating with other nodes." Ex. 2002, 86:13–16; *see id.* at 88:5–7; *see also id.* at 88:12–16 ("What I said was that this policy element is determining the availability of resources. So it's assessing whether or not hard disk space of various sorts is available. So it's availability, not just available."); *see id.* at 88:18–89:18, 90:5–91:1. We further credit the testimony of Patent Owner's declarant, Dr. Shenoy, who testified that Carter discloses "determin[ing] the free volume space available for allocation." *See* Ex. 1144, 142:19–20 (referencing Ex. 1140, 19:14–17). Reading directly from Carter, Dr. Shenoy further testified that Carter discloses "request[ing] the total available space information and us[ing] this information to quickly determine whether to proceed with the allocation processing." *Id.* at 142:14–17, 143:16–17. In view of the explicit disclosure of Carter and the testimony of both declarants, we find that Carter's controller performs an "assessment" in order to "determine the availability of resources." Additionally, Carter states that policy controller 290 "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.” Ex. 1140, 36:29–37:2. Thus, we agree with Petitioner’s argument, as supported by the testimony of Dr. Kubiawicz, that policy element 290 assesses the unused or underused storage resources when it determines the availability of such resources and uses some heuristics when distributing those resources that were assessed. See Pet 39 (citing Ex. 1102 ¶ 118).

Accordingly, based on the entirety of the record before us, we find Petitioner has shown by a preponderance of the evidence that Carter discloses the “assessing” limitations of claims 1 and 13.

d. *“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”*

Petitioner contends Carter discloses this limitation as recited by claim 1 and similarly recited in claim 13. Pet. 39 (citing Ex. 1102 ¶ 119), 54 (citing Ex. 1102 ¶ 157). Petitioner argues that “the shared memory subsystem of Carter satisfies this requirement by making the aggregated, shared storage space *appear* to be ‘a device that makes data storage resources available to network-connected user devices.’” *Id.* (citing Ex. 1102 ¶ 119). Petitioner argues that, “by communicating over the network with other nodes in the Carter system, the shared memory subsystem makes data storage resources available to network-connected user devices, thus providing the functionality of a NAS device.” *Id.* at 39–40 (citing Ex. 1102 ¶ 120). According to Petitioner, “in providing access to data storage

resources, the shared memory subsystem *appears* to be (*i.e.*, ‘mimics’) a NAS device, transparently providing a disk-like interface to external clients and user applications, while concealing the underlying details of the shared storage space behind a layer of abstraction.” *Id.* at 40.

Petitioner relies on the declaration of Dr. Kubiatowicz, who testifies that the ’827 patent provides scant detail as to what is meant by the “representing” requirement of the claim and further testifies that because Carter’s “shared memory subsystem hides the details of (or ‘encapsulates’) memory management from applications running on the local node, the local node ‘views’ the shared storage space provided by Carter as a single large storage device, rather than the aggregated storage resources of an entire network of participating computers.” Ex. 1102 ¶ 119 (citing Ex. 1101, 2:32–35), ¶ 122.

Petitioner, therefore, argues that “even though the shared addressable memory of Carter is created by pooling together unused space from each of the distributed nodes in the system, this pooled storage appears to application programs as a dedicated network-attached storage (NAS) device.” Pet. 42 (citing Ex. 1102 ¶ 125). Petitioner cites to Carter’s disclosure that “the total volume size” that is provided to the user is equal to the “*current aggregate storage capacity* in the addressable shared memory space 20” as support for its position. *Id.* at 43 (citing Ex. 1140, 16⁵). That is, according to Petitioner, “using the shared memory subsystem, the shared storage space of Carter is represented as a NAS device that has a capacity equal to the total aggregate capacity.” *Id.*

⁵ Although Petitioner’s citation is to page 15 of Carter, the quoted material appears on page 16 of Carter.

Patent Owner contests Petitioner's position. PO Resp. 30–35. Patent Owner first contends that the NAS component must “represent” “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.” *Id.* at 30. According to Patent Owner, the Carter system appears as a single NAS device because “Carter only discloses a single shared memory space across all nodes, and thus cannot disclose separate NAS devices.” *Id.* at 31 (citing Ex. 2003 ¶ 73).

Patent Owner further contends that the term “NAS devices” indicates a plurality and that the claims required multiple selected distributed devices each including a NAS device. *Id.* at 31–32 (citing Ex. 2003 ¶ 74). Patent Owner argues (*id.* at 31) that Dr. Kubiawicz's testimony during cross-examination supports its position regarding Petitioner's claim construction:

Q. But you would agree, the single NAS device cannot satisfy the claims of the '827 patent, correct?

....

A. So, it says here that they comprise NAS devices more than one. So, if an analysis were to come out with only one NAS device be represented, it wouldn't satisfy this claim language. However, as I mentioned before, I said multiple NAS devices, both at the individual level, aggregates because there is more than one Carter instance, and the file system example I give here where there are two NAS devices.

....

Q. You would agree that a single NAS device could not meet the limitations of the claims of the '827 patent? Yes or no?

....

A. So I -- so the read on this claim language is potentially a little loose, in the sense that it's, you know, the selected devices comprise NAS devices. But, I will-- you know, I will stipulate that having more than one NAS device when you're done here, seems like it's the closest read on the claim language.

Q. I am just trying to understand what you're understanding is and you are kind of giving me these wishy-washy answers.

A. Fine. It requires more than one NAS device.

Ex. 2002, 80:18–81:25 (objections omitted). Patent Owner further argues that Carter only describes a single shared storage space across all nodes, not separate storage spaces. PO Resp. 32 (citing Ex. 1140, 5). Patent Owner notes that “[e]ven Carter’s title acknowledges that it only discloses a single coherent storage space: ‘**globally** addressable storage environment.’” *Id.* (citing Ex. 1140, Title; Ex. 2003 ¶ 75). And, according to Patent Owner, “Petitioners repeatedly confirm throughout the Petition that Carter discloses only as a single NAS device.” *Id.* (citing Pet. 39, 40, 41).

Patent Owner next contends that Dr. Kubiadowicz suggested, in his deposition, that “multiple distributed computer work groups of Carter can each include a single aggregated shared memory, thereby making up multiple shared memory spaces,” and Patent Owner contends this is procedurally improper and also technically flawed. *Id.* at 34 (citing Ex. 2002, 78:9–80:17). According to Patent Owner, the testimony is a new theory of operation not supported by Dr. Kubiadowicz’s declaration and even if Carter did disclose each computer workgroup having a shared memory space, Carter still does not disclose that each node represents its own

respective NAS device as required by the claims under the proper interpretation of the “representing” element. *Id.* at 35 (citing Ex. 2003 ¶ 78).

We do not agree with Patent Owner for several reasons. First, “NAS device” as defined in the ’827 patent “broadly refers to a device that makes data storage resources available to network-connected user devices. A ‘dedicated NAS device’ is a device whose primary operational purpose is for providing NAS services.” *See* Ex. 1101, 2:31–34. And as discussed previously, one embodiment disclosed in the ’827 patent includes an integrated system referred to as a “NAS device fabric.” *See* Ex. 1101, 42:65–43:5. Carter discloses a similarly integrated system. Specifically, Carter uses software that causes individual nodes (e.g., commercially available PCs) to represent, or mimic, dedicated NAS devices. Ex. 1140, 10:6–26, Fig. 1. Carter states that “each node may be responsible for storing a particular element or elements of the structured store of data.” *Id.* at 9:18–19. Carter discloses that the system can be any structured storage system that maintains an organized set of data. *Id.* at 9:26–27. Carter further discloses with reference to Fig. 1 that “[e]ach of the nodes 12a-12d can access the addressable shared memory space 20 through the shared memory subsystems 34a-34d.” *Id.* at 10:30–11:2, Fig. 1.

Additionally, Carter’s Figure 11 shows multiple commercially available PC nodes (406a and 406b) each mimicking NAS devices by making storage available. *See* Ex. 1140, 41:12–21. Carter states that, as shown in Figure 11, the “data associated with the directory pages are distributively stored across the two local memories and duplicate copies can exist.” *Id.* at 41:29–30, Fig. 11. Carter further discloses that “the data can move between different local memories and also move, or page, between

volatile and persistent storage.” *Id.* at 41:30–42:1. According to Carter, “[t]he data movement can be responsive to data requests made by memory users like application programs, or by operation of the migration controller.” *Id.* at 42:1–3.

Based on the disclosure of Carter, we agree with Petitioner, as supported by the testimony of Dr. Kubiatowicz, that Carter’s “shared memory subsystem” is the software-based NAS component on each node that cooperates with “shared memory subsystems” on other nodes to make physical storage available. *See* Pet. 22–23, 31–32; Ex. 1102 ¶¶ 79–80, 98–100. Thus, we do not agree with Patent Owner that Carter does not disclose a “‘shared memory subsystem’ . . . [that] appears as a **single** shared resource across all of Carter’s nodes.” *See* PO Resp. 30. Rather, we find that 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.

Accordingly, having reviewed the entirety of the record before us and the claim constructions for “representing” proposed by either party, we find that Carter discloses “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” as recited in claim 1 and similarly recited in claim 13.

e. “allocat[ing] respective available amount of unused storage resources in selected distributed devices of the plurality of distributed devices”

Petitioner contends Carter discloses the limitation “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; represent that the selected distributed devices comprise respective NAS devices having the respective available amounts of storage resources,” as recited in claim 13. Pet. 54. Petitioner specifically states that this “limitation is satisfied for the reasons given above in the context of claim 1.” *Id.* (citing Pet. 37–43; Ex. 1102 ¶ 157).

Patent Owner contests Petitioner’s position, arguing that “distributing available resources” does not meet the allocating requirements of the claims. PO Resp. 36–37. Patent Owner contends Petitioner’s position and Dr. Kubiawicz’s testimony are conclusory and offer no support or analysis to explain how the disputed limitation is disclosed in Carter. *Id.* at 37. According to Patent Owner, “merely distributing the available resources across the nodes has nothing to do with ‘allocate[ing] [sic] respective available amount of unused storage resources.’” *Id.* (alteration in original) (citing Ex. 2003 ¶ 109). Patent Owner asserts that to “allocate” means “[t]o reserve a resource, such as sufficient memory, for use by a program.” *Id.* (citing Ex. 2005, 20). Patent Owner cites to the deposition testimony of Dr. Kubiawicz to support its position (*id.* at 37–38) because Dr. Kubiawicz testified as follows:

- Q. So allocation means reserving storage on the node?
- A. Allocation means basically making the storage available exclusively to the NAS component so that you can represent that you have that amount of storage.
- Q. And so because it’s available exclusively to the NAS component, you would agree that when storage is allocated, it can’t be used for storage purposes other than for use by the NAS component; correct?

....

- A. So we were talking in the context of claim 13. And in the context of claim 13, allocate here means in the context of claim 13 locking down the storage resources so that when you represent that the NAS device has an available amount of storage resources, they are not going to be taken over by local node use.

Ex. 2002, 105:21–106:21 (objection omitted).

Patent Owner then contends “Carter does not explicitly disclose reserving or ‘locking down’ resources as part of its policy controller’s distribution of resource activities to make the[m] ‘exclusively available’ to Carter’s shared memory subsystem.” PO. Resp. 38 (citing Ex. 1102 ¶ 111). According to Patent Owner, the mere fact “[t]hat Carter discloses distributing resources does not necessarily mean that unused or under-utilized resources are allocated.” *Id.*

We do not agree with Patent Owner. Rather, we find Petitioner has shown that Carter explicitly discloses “allocation” as required by claim 13 for at least the following reasons. First, Petitioner cites to Carter’s policy controller because the policy element 290 “acts as a controller to determine the availability of resources, such as . . . hard-disk space, available RAM and other such resources.” See Pet. 39 (citing Ex. 1140, 36:27–29). In order to determine the availability of memory resources, we understand that the controller assesses used, unused, and allocated space. See Ex. 1144, 131:19–22 (“the available amount of unused storage in the allocate step is the amount of unused storage that you determine in the assessment step.”). Once data has been moved into previously unused memory space, that space is now allocated to the stored data. We do not agree with Patent Owner’s

overly narrow reading of “allocate” because the claim does not require pre-allocation of specific memory space, nor does it require “exclusive” allocation of memory space for any specific duration. See Tr. 32:7–9, 32:21–33:14, 35:1–37:3. As Patent Owner agreed during oral argument, there is no temporal limitation on claim 13’s “allocation” so the “allocation” need not be exclusive for a given time period. *Id.* at 32:3–33:5.

Second, we credit the testimony of Dr. Shenoy, who testifies that Carter discloses “allocation” of memory space. Ex. 1144, 133:20–134:3, 127:1–8, 128:1–7, 135:5–17. Additionally, Dr. Shenoy testified that “allocate” has the same meaning when used in Carter as it has in the ’827 patent. *Id.* at 124:3–6.

Third, we credit the testimony of Dr. Kubiawicz, who testifies that both assessing and allocating are old, well-known, and basic functions and would not require significant description. See Ex. 1145 ¶ 21. Dr. Kubiawicz’s testimony supports Petitioner’s argument that the ’827 patent specification lacks a description of both functions because these functions were well understood by a person of ordinary skill in the art. Reply 7 (citing *Microsoft Corp. v. Biscotti, Inc.*, 878 F.3d 1052, 1068 (Fed. Cir. 2017)); see also *Kennametal, Inc. v. Ingersoll Cutting Tool Co.*, 780 F.3d 1376, 1381 (Fed. Cir. 2015) (“a reference can anticipate a claim even if it ‘d[oes] not expressly spell out’ all the limitations arranged or combined as in the claim, if a person of skill in the art, reading the reference, would ‘at once envisage’ the claimed arrangement or combination.”).

Lastly, we find that Carter explicitly discloses the action of “allocation.” Carter specifically “allocates” when it “reserves (sometimes much) larger than requested memory space for a file, and substantiates the

storage by allocating backing storage page by page.” See Ex. 1140, 19:2–4. Carter further discloses that file system 60 requests the total available space information and uses this information to quickly determine whether to proceed with the allocation processing. *Id.* at 19:14–17. Carter then discloses that “[t]he distributed directory manager can include a directory page generator for allocating a portion of the addressable memory space.” *Id.* at 26:18–19. Carter also discloses that its “data streams are mapped onto pages allocated from the addressable shared memory space 20 for storage.” *Id.* at 17:18–19. Thus, given Petitioner’s argument, Dr. Kubiawicz’s and Dr. Shenoy’s testimony, and the explicit disclosures in Carter, we find that Carter’s assessment and allocation of memory space for files meets the challenged claim limitation.

f. Summary

Accordingly, having analyzed the entirety of the record, we determine Petitioner has established by a preponderance of the evidence that claims 1 and 13 are anticipated under 35 U.S.C. § 102(b) based on Carter.

3. Analysis of Cited Art as Applied to Dependent Claims 3, 4, 6, 8, 9, 15, 16, 18, 20, and 21

Petitioner contends dependent claims 3, 4, 6, 8, 9, 15, 16, 18, 20, and 21 of the ’827 patent are unpatentable under 35 U.S.C. § 102(b) in view of Carter and provides specific arguments for each challenged claim. Pet. 47–53, 55. Patent Owner specifically disputes Petitioner’s contentions for claims 4, 9, 16, and 21. PO Resp. 39–50. We address arguments regarding the dependent claims in detail below.

a. Analysis of Cited Art as Applied to Dependent Claims 3, 6, 8, 15, and 18–20

Petitioner contends that dependent claims 3, 6–8, 15, and 18–20 of the '827 patent are unpatentable under 35 U.S.C. § 12(b) in view of Carter. Pet. 47–48, 51–52, 55.

For claims 3 and 15, which require “receiving an access request from the client device and directing the client device to data requested on at least one of the selected distributed devices,” Petitioner argues Carter meets the claim because the shared memory subsystem includes a “flow scheduler” that is a semi-autonomous software agent that routes and services requests from the distributed nodes, as indicated by red arrows in an annotated version of Carter’s Figure 8 provided on page 35 of the Petition. *Id.* at 47–48. Petitioner relies on Carter’s disclosure that “flow scheduler 272 can be a state machine that monitors and responds to the requests 312 and remote requests through network 304 which can be instructions for memory operations and which can include signals representative of the global addresses being operated on. These memory operation requests 312 can act as op-codes for primitive operations on one or more global addresses. They can be read and write requests, or other memory operations.” Ex. 1140, 35:9–14, *cited in* Pet. 48; *see also* Ex. 1102 ¶ 136. In response to such requests, Petitioner cites to Carter’s explanation 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.” *Id.* (citing Ex. 1140, 29 (emphasis added by Petitioner)). Carter further discloses that “[e]ach page 224a-224d represents a unique entry of computer data stored within the virtual shared memory 222.” Ex. 1140, 29:14–15. Petitioner asserts that, using the shared

memory subsystem, each page 224a-224d can be accessed or modified by each one of the nodes 212a-212c. Pet. 48 (citing Ex. 1140, 29). It is in this way, according to Petitioner, that “the shared memory subsystem ‘receiv[es] an access request from the client device and directing the client device to data requested on at least one of the selected distributed devices,’ as required by the claim.” *Id.* at 48 (citing Ex. 1102 ¶ 137).

For claims 6 and 18, which require storing, with the centralized server, location information associated with the data stored in the selected distributed devices, Petitioner contends “Carter teaches that the node responsible for storing the root of the filesystem acts as a centralized server.” *Id.* at 51. According to Petitioner, the root of the file system is “location information associated with the data stored in the selected distributed devices,” as required by claim 6. *Id.* Similarly, Petitioner argues that the node responsible for storing root directory page 342 acts as a centralized server. *Id.* Petitioner notes that the root directory page contains “location information associated with the data stored in the selected distributed devices,” as required by claim 6. *Id.* (citing Ex. 1140, Fig. 10; Ex. 1102 ¶ 146).

For claim 8, which requires that the method be “at least partially performed through the use of the Internet,” Petitioner argues that Carter points to the need for strong authentication “because of the ability of the invention to run over physically insecure WANs *such as the Internet.*” *Id.* at 52 (citing Ex. 1140, 63; Ex. 1102 ¶ 149).

For claim 20, which depends from claim 1 and further requires “wherein the distributed devices are configured to be coupled to the

Internet.” Ex. 1101, 48:30–31. Petitioner contends claim 20 is satisfied for the same reasons as claim 8. Pet. 55 (citing Ex. 1102 ¶ 167).

Patent Owner disputes Petitioner’s contentions, arguing that dependent claims are not anticipated by Carter for the same reasons as claims 1 and 13. PO Resp. 11.

After considering the parties’ arguments and evidence, we agree with Petitioner’s position and adopt its arguments as our own. Accordingly, we determine Petitioner has shown by a preponderance of the evidence that dependent claims 3, 6, 8, 15, and 18–20 of the ’827 patent are unpatentable under 35 U.S.C. § 102(b) in view of Carter.

b. Analysis of the Prior Art as Applied to Dependent Claims 4 and 16

Claims 4 depends from claim 3, while claim 16 depends from claim 15, and each further require “manag[ing] the NAS service for the [at least one of the particular] distributed devices at least in part utilizing a centralized server.” Ex. 1101, 46:65–67, 48:15–17.

Petitioner argues that Carter discloses using a “centralized server,” as recited in claims 4 and 16, because Carter uses a root directory page that “can be located at a *static* address that is known to each node coupled to the distributed address space.” Pet. 49 (quoting Ex. 1140, 39), 55 (citing Ex. 1102 ¶ 163). According to Petitioner, “[w]hile the other structures stored in the addressable shared memory space can be shifted dynamically from one underlying device to another, the root node is located at a ‘static’ location – *i.e.*, a known physical address on one of the nodes attached to the system.” *Id.* (citing Ex. 1140, 14; Ex. 1102 ¶ 139). Petitioner relies on the testimony

of its declarant, Dr. Kubiadowicz, to support its position. Dr. Kubiadowicz testifies that in Carter

the node that stores the root directory page 342 is *centralized*, in the sense that it is known to all of the other nodes on the system as a starting point for translating a global address into a physical one. While the traversal of the directory page structure shown in Fig. 10 may involve many different nodes in the system, it always starts with the node that stores the root directory page. Unlike other data structures in the Carter system, which are distributed fluidly among many different computers, this root node is stored in a centralized location.

Ex. 1102 ¶ 142. Dr. Kubiadowicz further testifies that nodes that store Carter’s root directory page 342 and root entry point 64 are servers because they receive and respond to requests for information from other nodes in the system. *Id.* ¶¶ 143–144.

Patent Owner disputes Petitioner’s contentions, arguing that the claimed “centralized server” must “form[] a known, physical ‘center’ from the perspective of the distributed devices” and that it “must be physically separate from the client devices (the claimed ‘distributed devices’) that it serves.”⁶ PO Resp. 41–42 (citing Ex. 2003 ¶ 85). Patent Owner contends that the nodes storing Carter’s root directory page 342 and root 64 are neither centralized nor separate from Carter’s distributed nodes. *Id.* at 43 (citing Ex. 2003 ¶ 88). Patent Owner argues that “Petitioners’ reliance on Carter’s discussion of a ‘static address or location’ is misplaced because

⁶ We note this proposed construction of “centralized server” differs from Patent Owner’s proposed construction in related IPR2017-00374. *See* IPR2017-00374, Paper 24, 29–30 (“centralized server” requires “a single server or a single group of servers that are commonly controlled and managed.”)

Carter's reference to a 'static address or location' is with respect to memory addresses in the **virtual shared memory space**, which is a logical address decoupled from the actual physical nodes." *Id.* at 44 (citing Ex. 2003 ¶ 91). According to Patent Owner, "Carter, at best, discloses that the logical address of root directory page 342 or root 64 in the virtual shared memory space can be static and known to each node." *Id.* Patent Owner asserts that "Carter does not teach, or even suggest, that the actual physical node(s) in which root directory page 342 or root 64 is stored is known also to each node." *Id.* Furthermore, according to Patent Owner, "the logical address of root directory page 342 or root 64 can be potentially linked to any node in Carter's systems" such that "Carter contemplates that all its nodes could be linked to the logical address of root directory page 342 or root 64." *Id.* at 44–45 (citing Ex. 2003 ¶ 92).

Patent Owner also contends that applying Carter's distributed architecture to claims 4 and 16 reads the term "centralized" out of the claims. *Id.* at 45. According to Patent Owner, "the term 'centralized' becomes meaningless because it is nonsensical to say there is still a known 'center' in a system when every node in the system is part of the 'center.'" *Id.* (citing Ex. 2003 ¶ 92). Patent Owner argues that, although in one embodiment shown in Figure 11, the root page is stored in persistent memory 262a on one of the nodes, along with some other pages of directory structure 340, Carter neither discloses that the node storing the root page is known to other nodes nor discloses that there is only one, or a limited number of, "centralized" node(s) that store the root page. *Id.* (citing Ex. 2003 ¶ 93). Patent Owner further argues that the root directory page in Figure 11 cannot be the centralized server because it violates one of the key

requirements of a centralized server—that the server is separate from the clients that it serves. *Id.* at 46 (citing Ex. 2003 ¶ 95). Patent Owner asserts that, “[i]n Carter, the root directory is stored in the same node in which distributed data is stored (*i.e.*, a client device), as opposed to being stored in a centralized server, separate from the distributed devices.” *Id.* Thus, Patent Owner argues that, “[e]ven if the physical node storing the root directory page was known to other nodes, such a node is merely one of the ‘*distributed devices*,’ as opposed to a ‘*centralized server*.’” *Id.*

Patent Owner further contends that centralized servers are the antithesis of Carter’s architecture. *Id.* at 43–44. According to Patent Owner, “Carter recognized that centralized client-server approaches had limitations such as ‘fail[ing] [sic] to exploit available resources as network activity changes,’ not ‘improve[ing] system performance,’ being ‘subject to complete system failure each time the server system crashes,’ and having ‘complex’ functionality.” *Id.* at 43 (citing Ex. 1140, 2–3; Ex. 2003 ¶ 89). Therefore, Patent Owner argues, “Carter designed its architecture to be completely distributed, with no conventional ‘client-server’ systems having servers acting as ‘a central controller.’” *Id.* (citing Ex. 1140, 1, 2–4; Ex. 2003 ¶ 89). Patent Owner asserts that Petitioner’s position that Carter includes a centralized server contradicts Carter’s motivations and objectives. *Id.* at 44 (citing Ex. 2003 ¶ 89).

We do not agree with Patent Owner. First, Patent Owner’s position regarding an alleged required “single centralized server” is unduly narrow given the language of claim 4 reciting “managing the NAS service for the at least one of the particular distributed devices *at least in part utilizing a centralized server*.” Ex. 1101, 46:65–67 (emphasis added). The express

language of claim 4, therefore, contemplates the use of a “centralized server” “at least in part,” not exclusively (i.e., not “[only a single] centralized server”). Still further, multiple such “centralized server[s]” may be used within the scope of claims 4 and 16.

Furthermore, Carter’s description of root directories as providing a “starting point to locate the filesets” that are used “[w]hen a node is accessing a fileset for the first time” (Ex. 1140, 14:1–6) appears consistent with the ’827 patent’s description of “centralized”: “Thus, the NAS server systems storing data location information in the centralized data location library 2110, for example, would receive storage and access requests from user devices and would at least in part direct the user devices to the distributed devices storing the requested data.” Ex. 1101, 44:37–42. Dr. Kubiadowicz’s testimony that “the node that stores the root directory page 342 is *centralized*, in the sense that it is known to all of the other nodes on the system as a starting point for translating a global address into a physical one” (Ex. 1102 ¶ 142) also appears consistent with the notion of centralization in the ’827 patent.

Accordingly, after consideration of the entire record before us, we determine Petitioner has shown by a preponderance of the evidence that dependent claims 4 and 16 of the ’827 patent are unpatentable under 35 U.S.C. § 102(b) in view of Carter.

c. Analysis of the Prior Art as Applied to Dependent Claims 9 and 21

Claim 9 recites “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.”

Ex. 1101, 47:15–19. Similarly, claim 21 recites “a storage priority control configured to facilitate use of the respective available amount of unused storage resources for the selected distributed devices.” *Id.* at 48:32–35.

Petitioner contends that Carter describes “managing storage resources for the selected distributed devices with a storage priority control,” as recited in claims 9 and 21. Pet. 52–53 (citing Ex. 1102 ¶ 150), 55 (citing Ex. 1102 ¶ 169). Petitioner argues that a “storage priority control” is demonstrated by Carter’s disclosure that the shared memory subsystems can include logic to assess the “*availability of resources*, such as printer capabilities, *hard-disk space*, available RAM and other such resources.” *Id.* at 52 (citing Ex. 1140, 36–37). Petitioner further notes that Carter’s shared memory subsystem may include 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*.” *Id.* at 52–53 (alteration in original) (quoting Ex. 1140, 36–37; citing Ex. 1102 ¶ 151). Petitioner also explains that Carter’s “memory subsystems 232a–232c can include a hierarchy manager that provides hierarchical control for the distribution of data. This includes controlling the migration controller, and policy controller, which are discussed in detail below, 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.” *Id.* at 53 (quoting Ex. 1140, 31 (“*Resource distribution is similarly managed*”); Ex. 1102 ¶ 152).

Patent Owner contests Petitioner’s position, arguing that Petitioner fails to demonstrate how Carter describes the management of storage resources by any component, “let alone management that facilitates the use

of an available amount of **unused** storage resources of distributed devices.” PO Resp. 48. Specifically, Patent Owner argues that “Carter does not disclose that it determines availability based on priority (the ranks, positions, or privileges) assigned to different storage resources for the selected distributed devices.” *Id.* at 49 (citing Ex. 2003 ¶ 100). Patent Owner further argues that Petitioner does not explain how or why using “heuristics” teaches that the policy controller facilitates the usage of unused storage resources in any manner. *Id.* Patent Owner then argues that Carter’s heuristics would not necessarily be “a storage priority control” as required by the challenged claims. *Id.* at 50. According to Patent Owner, “Carter’s disclosure of ‘load balancing’ or ‘resource distribution’ does not necessitate ‘storage priority control.’” *Id.* (citing Ex. 2003 ¶ 102).

We do not agree with Patent Owner. Rather, we find that “storage priority control” is demonstrated by Carter’s disclosure that the shared memory subsystems can include logic to assess the “*availability of resources*, such as printer capabilities, *hard-disk space*, available RAM and other such resources.” Ex. 1140, 36:27–29 (emphases added). Additionally, we agree with Petitioner that Carter’s “policy controller,” which “dynamically distribute[s] the available resources” based on “suitable heuristics,” and Carter’s memory subsystems that prioritize storage resources using functionality like “hierarchical data migration,” “load balancing,” and “resource distribution” describe the disputed claim limitation. See Pet. 52–53; Ex. 1102 ¶ 151; Ex. 1140, 31, 36.

Accordingly, after consideration of the entire record before us, we determine Petitioner has shown by a preponderance of the evidence that

dependent claims 9 and 21 of the '827 patent are unpatentable under 35 U.S.C. § 102(b) in view of Carter.

d. Summary

On this record and for purposes of this Decision, we determine Petitioner has shown by a preponderance of the evidence that dependent claims 3, 4, 6, 8, 9, 15, 16, 18, 20, and 21 are anticipated under 35 U.S.C. § 102(b) based on Carter.

E. Alleged Obviousness of Dependent Claims 5 and 17 of the '827 Patent in View of Carter and Pitzel

Petitioner contends dependent claims 5 and 17 of the '827 patent are unpatentable under 35 U.S.C. § 103(a) in view of Carter and Pitzel. Pet. 55–59. Patent Owner disputes Petitioner's contentions, arguing Pitzel fails to remedy the alleged deficiencies of Carter with respect to independent claims 1 and 13. PO Resp. 51. For reasons that follow, we determine Petitioner has shown by a preponderance of the evidence that the challenged claims would have been obvious in view of Carter and Pitzel.

1. Overview of Carter

See supra Section II.D.1.

2. Overview of Pitzel

Pitzel is a US Patent titled “System and Method for Updating Information via a Network.” Ex. 1142, Title. Pitzel is directed to a “method for upgrading software for various users, each of the users having diverse hardware and software configurations.” *Id.* at Abstract. The invention of Pitzel “includes a method for automatically determining the hardware and software configuration of the client computer without user intervention.” *Id.* According to Pitzel, “[t]o begin the upgrade, the client computer reads a configuration file that identifies the components for downloading.” *Id.*

Pitzel discloses that “[u]pon receipt of the configuration file, the client computer determines the hardware and software configuration of the client computer.” *Id.* The client computer of Pitzel then “automatically communicates with a component server via a network and requests one or more components for downloading to the client computer based upon the configuration file and the hardware and software conditions of the client computer.” *Id.*

3. *Analysis of the Prior Art as Applied to Dependent Claims*

Claims 5 and 17 depend from dependent claims 4 and 16, respectively. As discussed previously (*see supra* Section II.D.3.b), we are persuaded by Petitioner’s arguments that Carter satisfies the limitations in claims 4 and 16 requiring a centralized server. For claims 5 and 17, which require downloading the software-based NAS component to the selected distributed devices (Ex. 1101, 47:1–3, 48:19–21), Petitioner contends Pitzel meets this limitation. Pet. 56. According to Petitioner, “Pitzel teaches ‘downloading’ of a component to a client, which is equivalent to the ‘selected distributed device[]’” and “teaches a network with multiple clients, each of which can download and install a component.” *Id.* (citing Ex. 1142, Abstract, 3:23–27). Petitioner contends that a person of ordinary skill in the art would have had a reason to combine the teachings of Carter with that of Pitzel because (1) the references both are in the same field of endeavor, *id.* at 57; (2) a person of ordinary skill in the art would have recognized the benefits of incorporating the teachings of Pitzel to the system of Carter, *id.* at 58; and (3) Pitzel could have been easily incorporated into the Carter system, and such combination would have yielded predictable results, *id.*

Petitioner argues that “the network system in Carter can advantageously incorporate the downloading mechanism in Pitzel to provide a software-based NAS component to a computer in the network. The computer can install the downloaded component to provide NAS service to other computers in the network.” *Id.* According to Petitioner, “[t]he downloading mechanism in Pitzel can also provide an update to an existing software-based NAS component on a computer in the network. Such an update can make a version change, fix software issues, and/or implement a feature change.” *Id.* (citing Ex. 1102 ¶ 177). Petitioner supports its position with the declaration of Dr. Kubiawicz, who testifies that “[a] [person of ordinary skill in the art] would further appreciate that any software of the level of complexity of Carter would greatly benefit from the ability to perform updates—especially over geographic distances such as illustrated by Carter’s discussion of Figs. 12A, 12B and 13.” Ex. 1102 ¶ 177 (citing Ex. 1140, 49–50, Figs. 12A, 12B, 13).

Patent Owner disputes Petitioner’s position, arguing that Pitzel fails to cure the alleged deficiencies of Carter with respect to independent claims 1 and 13. PO. Resp. 51.

After considering the entirety of the record and analyzing the parties’ arguments and evidence, we agree with Petitioner’s position and adopt its arguments as our own. *See, e.g.*, Ex. 1102 ¶ 177. In particular, we credit Dr. Kubiawicz’s testimony regarding the advantages of being able to perform software updates. *See id.* Accordingly, we determine Petitioner has shown by a preponderance of the evidence that dependent claims 5 and 17 would have been obvious under 35 U.S.C. § 103(a) in view of Carter and Pitzel.

*F. Alleged Obviousness of Dependent Claims 4, 6, 7, 16, 18, and 19
in view of Carter and Condor*

Petitioner contends dependent claims 4, 6, 7, 16, 18, and 19 of the '827 patent are unpatentable under 35 U.S.C. § 103(a) in view of Carter and Condor. Pet. 59–66. Prior to our Decision to Institute, Patent Owner disputed Petitioner's contentions. Prelim. Resp. 29–35. As discussed previously, following the issuance of *SAS Institute*, we modified our Decision to Institute to include all of the challenged claims and all of the grounds presented in the Petition. Paper 48.

1. Overview of Carter

See supra Section II.D.1.

2. Overview of Condor

Condor is a publication titled “Condor – A Hunter of Idle Workstations.” Ex. 1117, Title. Condor is directed to a scheduling system that operates in a workstation environment. *Id.* at Abstract. Condor discloses “maximiz[ing] the utilization of workstations with as little interference as possible between the jobs it schedules and the activities of the people who own workstations.” *Id.* It teaches identifying idle workstations and scheduling background jobs on them, so that when the owner of a workstation resumes activity at a station, Condor checkpoints the remote job running on the station and transfers it to another workstation. *Id.* According to Condor, the “system guarantees that the job will eventually complete, and that very little, if any, work will be performed more than once.” *Id.*

Condor contemplates a “wide spectrum of scheduling structures” that use scheduling coordinators that “follow[] the principle that workstations are autonomous computing resources and they should be managed by their own

users. *Id.* at 4, 6. Condor states that “[t]his also helps to keep the responsibilities of the coordinator simple.” Ex. 1117, 6. According to Condor, “[s]implicity is important so that a central site is not required to maintain a great amount of information about each workstation.” *Id.* Condor contemplates that either a centralized, static coordinator could assign background jobs to execute at available remote workstations (*id.* at 4–5) or a distributed approach could be used where assignment of available processors is accomplished by each workstation cooperating to conduct a scheduling policy (*id.*).

One embodiment disclosed in Condor is shown in Figure 1, reproduced below.




Figure 1, above, shows one version of the Condor background job scheduler that lies between a centralized, static approach and the fully distributed approach. Condor explains that “[t]his approach uses the efficiency of scheduling with a central node to avoid the overhead of messages to decide which workstations should be allocated available capacity.” *Id.* at 5. Condor discloses that its system can be extendible to a large number of workstations and eases the required recovery when a centralized coordinator fails. *Id.* at 6–7. According to the teachings of Condor, “[l]ocal schedulers

are not affected if a remote site discontinues service.” *Id.* at 7. Condor discloses that “[i]f the site on which the coordinator is executing fails, remotely executing jobs initiated and executing on other machines are not affected. Only the allocation of new capacity to requesting users is affected.” *Id.* Condor explains that because “the coordinator has few duties, its recovery at another site is simplified in relation to a fully centralized strategy.” *Id.* In order to balance the burden of coordination, Condor states that the central coordinator can be moved to other locations. *Id.*

3. Analysis of Cited Art as Applied to the Challenged Dependent Claims

Claim 4 depends from claim 3, and further requires “managing the NAS service for the at least one of the particular distributed devices at least in part utilizing a centralized server.” Ex. 1101, 46:65–67. Claim 6 depends from claim 4, and requires the method of claim 4 to further comprise “storing, with the centralized server, location information associated with the data stored in the selected distributed devices.” *Id.* at 47:4–6. Claim 7 depends from claim 6, and requires the method of claim 6 to further comprise “utilizing the centralized server to receive and route the data for storage to the selected distributed devices based upon individual capabilities of the selected distributed devices indicated in a capabilities database.” *Id.* at 47:7–12. Dependent claims 16, 18, and 19 recite substantially similar limitations as claims 4, 6, and 7 but depend from independent claim 13. *Id.* at 48:15–29.

Petitioner contends these dependent claims would have been obvious in view of the combined teachings of Carter and Condor. Pet. 60 (citing Ex. 1102 ¶ 184). In particular, for claim 4, Petitioner argues that “Condor teaches ‘managing’ the use of ‘resources of the workstation [that] are

underutilized’ at least in part using a ‘central coordinator.’” *Id.* (brackets in original) (citing Ex. 1117, 1, 5–6). According to Petitioner, “[t]he central coordinator is held on one of the workstations that also has its own local scheduler and background job queue.” *Id.* (citing Ex. 1117, 6). Petitioner argues that the central coordinator “allocates capacity from idle workstations to local schedulers on workstations that have background jobs waiting.” *Id.* Petitioner further argues the central coordinator knows “not only which workstation’s processor is available . . . but [also] how much disk space is available on each workstation.” *Id.* (citing Ex. 1117, 16; Ex. 1102 ¶ 184).

For dependent claim 6, Petitioner contends “Condor teaches ‘storing’ ‘location information associated with the data’ stored on the distributed devices with a ‘central coordinator.’” Pet. 62 (citing Ex. 1117, 1, 5–6). According to Petitioner, the central coordinator “allocates capacity from idle workstations to local schedulers on workstations that have background jobs waiting.” *Id.* (citing Ex. 1117, 6). Petitioner further argues that the central coordinator knows “not only which workstation’s processor is available . . . but [also] how much disk space is available on each workstation.” *Id.* (citing Ex. 1117, 16; Ex. 1102 ¶ 189). For claim 7, Petitioner argues that “Condor teaches ‘utilizing’ a central coordinator to ‘receive and route data for storage’ based upon a ‘capabilities database.’” *Id.* at 63. Petitioner relies on Condor’s disclosure of central coordinator that “allocates capacity from idle workstations to local schedulers on workstations that have background jobs waiting,” and which knows “not only which workstation’s processor is available . . . but [also] how much disk space is available on each workstation.” *Id.* at 63–64 (citing Ex. 1117, 6; Ex. 1102 ¶ 192). For claims

16, 18, and 19, Petitioner states that the arguments and rationales for claims 4, 6, and 7 apply. *Id.* at 64–65.

Petitioner supports its position with the declaration of Dr. Kubiadowicz, who testifies that the challenged dependent claims would have been obvious in light of the combination of Carter and Condor. Ex. 1102 ¶¶ 183–201. Dr. Kubiadowicz specifically testifies that “[t]he advantages and disadvantages of both centralized and distributed approaches to coordination and management were well[-]known even at the time Condor was published.” *Id.* ¶ 185 (citing Ex. 1117, 5 (“Both the centralized and the distributed approaches have well known advantages and disadvantages.”)). Dr. Kubiadowicz specifically testifies that “it would have been obvious to a person of ordinary skill in the art that both approaches have advantages that could have been exploited using well known techniques, as well as disadvantages that [could have been] avoided using well known techniques, just like Condor’s central coordinator.” Ex. 1102 ¶ 186. According to Dr. Kubiadowicz, “[i]ncorporating Condor’s explicit teaching of a ‘central coordinator’ into Carter would have been obvious to one of ordinary skill in the art, would have yielded predictable results, and would have involved tools and techniques that were well known at the time.” *Id.* ¶ 187. Dr. Kubiadowicz further testifies that a person of ordinary skill in the art “would readily recognize the benefits of incorporating the teachings of Condor to the system of Carter. For example, the network system in Carter can advantageously incorporate the central coordinator in Condor to provide a centralized mechanism to manage resource sharing.” *Id.* ¶ 199.

Patent Owner contests Petitioner’s position, arguing Petitioner fails to provide an adequate rationale as to why a person of skill in the art would

combine an entirely distributed system from Carter with a centralized system from Condor, especially in light of Carter criticizing centralized servers such as those required in claims 4 and 6. Prelim. Resp. 31–32. Patent Owner specifically argues that Petitioner provides no analysis regarding the combination of Carter and Condor but, rather, offers only conclusory statements that merely are reiterated by Dr. Kubiadowicz. *Id.* at 32–33.

We agree with Patent Owner that Petitioner fails to offer sufficient explanation as to why a person of ordinary skill in the art would have combined the teachings of Carter and Condor. Rather, Petitioner appears to argue that, based on the advantages described by Condor (Ex. 1017, 5), and using well-known techniques with unspecified predictable results, a person of ordinary skill in the art would have had sufficient reason to combine these disparate elements taught in the prior art references to render the challenged claims obvious. Pet. 61–66. Yet, other than Petitioner’s cursory assertion of “advantageously incorporat[ing] the central coordinator in Condor to provide a centralized mechanism to manage resource sharing” (Pet. 66), which cites equally cursory and unsupported testimony by Dr. Kubiadowicz (Ex. 1102 ¶ 199), Petitioner fails to identify what particular advantages or efficiencies would have prompted a person of ordinary skill in the art to combine Condor’s scheduler with Carter’s system. Furthermore, Dr. Kubiadowicz’s testimony regarding predictable results and obviousness to try are similarly lacking in elaboration and substantiation. See Ex. 1102 ¶¶ 200–201. Such conclusory statements and unspecific testimony without evidence on the record cannot satisfy the Petitioner’s burden of demonstrating obviousness. *In re Magnum Oil Tools Int’l, Ltd.*, 829 F.3d 1364, 1380 (Fed. Cir. 2016); *I/P Engine, Inc. v. AOL Inc.*, 576 F. App’x 982,

999 (Fed. Cir. 2014) (“The testimony of the Defendants’ expert amounts to a ‘mere conclusory statement’ that may not serve as a basis for finding the asserted claims obvious.”); *see also InTouch Techs., Inc. v. VGO Commc’ns, Inc.*, 751 F.3d 1327, 1352 (Fed. Cir. 2014) (holding that expert’s testimony could not support a finding of obviousness where “testimony primarily consisted of conclusory references to [the expert’s] belief that one of ordinary skill in the art *could* combine these references, not that they *would* have been motivated to do so”); *see also* 37 C.F.R. 42.65(a) (“Expert testimony that does not disclose the underlying facts or data on which the opinion is based is entitled to little or no weight.”).

Although Petitioner has not provided a sufficient rationale to combine in this asserted ground of unpatentability, Petitioner has otherwise demonstrated that claims 4, 6, 16, and 18 are unpatentable under 35 U.S.C. § 102(b) because they are anticipated by Carter. Petitioner, however, has not asserted claims 7 and 19 are unpatentable under any other ground in this proceeding. Therefore, we determine Petitioner has failed to establish by a preponderance of the evidence that claims 7 and 19 of the ’827 patent are unpatentable under 35 U.S.C. § 103(a) in view of Carter and Condor.

III. CHALLENGE TO PETITIONER’S REPLY ARGUMENTS

With Board authorization (Paper 37), Patent Owner filed a list of alleged improper reply arguments (Paper 37), to which Petitioner filed a Reply (Paper 38). Patent Owner lists several portions of Petitioner’s Reply and evidence allegedly beyond the scope of what can be considered appropriate for a reply. *See* Paper 37 (citing Reply).

We have considered Patent Owner’s listing, but we disagree that the cited portions of Petitioner’s Reply and reply evidence are beyond the scope

of what is appropriate for a reply. Replies are a vehicle for responding to arguments raised in a corresponding Patent Owner Response. Petitioner's arguments and evidence are not beyond the proper scope of a reply because we find that they fairly respond to Patent Owner's arguments raised in Patent Owner's Response. *See Idemitsu Kosan Co., LTD. v. SFC Co. Ltd.*, 870 F.3d 1376, 1381 (Fed. Cir. 2017) ("This back-and-forth shows that what [Patent Owner] characterizes as an argument raised 'too late' is simply the by-product of one party necessarily getting the last word. If anything, [Patent Owner] is the party that first raised this issue, by arguing—at least implicitly—that [a reference] teaches away from non-energy-gap combinations. [Petitioner] simply countered, as it was entitled to do."); *see also Anacor Pharms., Inc. v. Iancu*, 889 F.3d 1372, 1380 (Fed. Cir. 2018) ("In addition, the petitioner in an inter partes review proceeding may introduce new evidence after the petition stage if the evidence is a legitimate reply to evidence introduced by the patent owner . . .").

IV. CONCLUSION

For the foregoing reasons, we conclude 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, and (2) claims 5 and 17 under 35 U.S.C. § 103(a) as unpatentable over Carter and Pitzel. We find, however, Petitioner has failed to provide an adequate rationale to combine the teachings of Carter and Condor, and we determine Petitioner has failed to establish by a preponderance of the evidence that claims 7 and 19 of the '827 patent are unpatentable under 35 U.S.C. § 103(a) in view of Carter and Condor.

V. ORDER

For the reasons given, it is

ORDERED that, by a preponderance of the evidence, claims 1, 3–6, 8, 9, 13, 15, 16–18, 20, and 21 of the '827 patent are unpatentable; and

FURTHER ORDERED that because this is a Final Written Decision, parties to the proceeding seeking judicial review of the decision must comply with the notice and service requirements of 37 C.F.R. § 90.2.

IPR2017-00439
Patent 8,275,827

PETITIONER:

Peter Dichiara
peter.dichiara@wilmerhale.com

Michael Smith
michaelh.smith@wilmerhale.com

PATENT OWNER:

Byron Pickard
bpickard-ptab@sternekessler.com

Daniel Block
dblock-ptab@sternekessler.com

Tim Seeley
tims@intven.com

James Hietala
jhietala@intven.com

Steven Pappas
spappas-ptab@sternekessler.com