

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

GOOGLE LLC
Petitioner,

v.

LS CLOUD STORAGE TECHNOLOGIES, LLC,
Patent Owner.

IPR2023-00777
Patent 6,549,988

**PATENT OWNER'S PRELIMINARY RESPONSE TO
GOOGLE LLC's PETITION FOR INTER PARTES REVIEW**

TABLE OF CONTENTS

I. INTRODUCTION	1
II. OVERVIEW OF THE '988 PATENT	2
A. Brief Description	2
B. Prosecution History	4
C. Claim Construction.....	4
IV. THE ASSERTED PRIOR ART	8
A. Heil (Ex. 1006)	8
B. Nakayama (Ex. 1007).....	10
C. Srinivasan (Ex. 1018).....	12
D. Macklem (Ex. 1009).....	13
V. PETITIONER HAS NOT DEMONSTRATED A REASONABLE LIKELIHOOD OF SUCCESS FOR THE GROUNDS ADVANCED IN THE PETITION, AND THE PETITION SHOULD BE DENIED	13
A. Requirements for Showing Obviousness Under 35 U.S.C. § 103	13
B. Petitioner Fails to Establish in Ground 1 that Challenged Claims 1-4 and 6 are Obvious Over Heil and Nakayama.....	15
1. Claim 1	15
a. Heil and Nakayama do not teach the limitation [1.b].....	16
b. Heil in view of Nakayama Does Not Render Claim 1 Obvious	20
2. Claim 2	21
a. Heil and Nakayama do not teach the limitation [2.b].....	21
b. Heil in view of Nakayama Does Not Render Claim 2 Obvious	25
3. Claim 3	26
4. Claim 4	26
5. Claim 6	26
C. Petitioner Fails to Establish in Ground 2 that Challenged Claim 5 is Obvious Over Heil in view of Nakayama, Macklem and Srinivasan.....	27
a. The combination of Heil, Nakayama, Macklem and Srinivasan does not teach the limitation [5.a].....	28

b. The combination of Heil, Nakayama, Macklem and Srinivasan does not teach the limitations [5.b] and [5.c].....	30
--	----

TABLE OF AUTHORITIES

Cases

<i>Chef America, Inc. v. Lamb-Weston, Inc.</i> , 358 F.3d 1371, 1372, 69 USPQ2d 1857 (Fed. Cir.2004).	5
<i>Vitronics Corp. v. Conceptronic, Inc.</i> , 90 F.3d 1576, 1582 (Fed.Cir.1996).	6
<i>Telemac Cellular Corp. v. Topp Telecom, Inc.</i> , 247 F.3d 1316, 1326 (Fed.Cir.2001).	6
<i>Renishaw PLC v. Marposs Societa' per Azioni</i> , 158 F.3d 1243, 1250 (Fed.Cir.1998).	6
<i>Graham v. John Deere Co.</i> , 383 U.S. 1 (1966)	14
<i>Pers. Web Techs., LLC v. Apple, Inc.</i> , 848 F.3d 987 (Fed. Cir. 2017)	14
<i>Intelligent Bio-Sys., Inc. v. Illumina Cambridge Ltd.</i> , 821 F.3d 1359 (Fed. Cir. 2016)	14
<i>Los Angeles Biomedical Res. Inst. at Harbor-UCLA Med. Ctr. v. Eli Lilly Co.</i> , 849 F.3d 1049 (Fed. Cir. 2017)	14
<i>Broadcom Corp. v. Emulex Corp.</i> , 732 F.3d 1325 (Fed. Cir. 2013).	15
<i>In re Van Os</i> , 844 F.3d 1359 (Fed. Cir. 2017).	16
<i>Innogenetics, N.V. v. Abbott Labs.</i> , 512 F.3d 1363 (Fed. Cir. 2008).	16
<i>CFMT, Inc. v. Yieldup Intern. Corp.</i> , 349 F.3d 1333, 1342 (Fed. Cir. 2003).	22

PTAB Cases

Apple, Inc. v. Contentguard Holdings, Inc., IPR2015-00355, Paper 9 14

Statutes

35 U.S.C. §103 15

Rules

37 C.F.R. § 42.107 1

37 C.F.R. . §42.100(b) 4

I. INTRODUCTION

On March 24, 2023, Google LLC (“Petitioner”) submitted a Petition for *Inter Partes* Review (“Petition”) of U.S. Patent No. 6,549,988 (“the ’988 Patent,” “Ex. 1001”), challenging Claims 1-6 (“Challenged Claims”). Pursuant to 37 C.F.R. § 42.107, LS Cloud Storage Technologies, LLC, (“Patent Owner”) files this Preliminary Patent Owner Response to the Petition, setting forth reasons the Board should deny institution.

Petitioner challenges the patentability of Claims 1-6 of the ’988 Patent based on four (4) different alleged prior art references as follows.

U.S Patent No. 6,173,374 (“*Heil*”);

U.S Patent No. 5,920,893, (“*Nakayama*”);

Not Quite NFS, Soft Cache Consistency for NFS, Macklem, Proceedings of the Winter 1994 USENIX Conference, January 1994. 1982 (“Macklem”);

Experiments with Cache Consistency Protocols, Proceedings of the Twelfth ACM Symposium on Operating Systems Principles, December 1989 to Srinivasan *et al.* (“Srinivasan”).

Petitioner requests institution of the IPR proceeding on the following grounds.

Ground	Claims	Basis
1	1-4 and 6	§103– Heil in view of Nakayama
2	5	§103– Heil in view of Nakayama, Macklem and Srinivasan

Petitioner fails to establish a reasonable likelihood that the Challenged Claims are unpatentable. Patent Owner will specifically address each of Petitioners arguments in the following sections.

II. OVERVIEW OF THE '988 PATENT

A. Brief Description


A typical data processing System generally involves a cached data storage system that connects to local host computers via I/O channels or remote host computers via network links. The purpose of the data storage system is to improve the performance of applications running on the host computer by offloading I/O processing from the host to the data storage system. The purpose of the cache memory in a data storage system is to further improve the performance of the applications by temporarily storing data buffers in the cache so that the references to those buffers can be resolved efficiently as “cache hits.” Ex. 1001, 1:18-28.

Prior art data storage systems are implemented using proprietary hardware and very low-level software, resulting in expensive and not portable systems. In contrast to the prior art systems, the '988 Patent uses standard hardware and software components. Ex. 1001, 1:35-42.

A data Storage System of the '988 Patent comprises a network of PCs each of which includes a cache memory, an I/O channel adapter for transmitting data over the channel and a network adapter for transmitting data and control signals over the

network. A method for managing resources in a cache memory ensures consistency of data stored in the distributed cache. Ex. 1001, 3:57-64.

Figure 1 from the '988 Patent, shown below, illustrates “[t]he PC data storage system 131 services a plurality of channel attached host processors 111, 112 using channels 121,122, and a plurality of network attached host processors 106, 107 using network link 151, and a plurality of network attached data storage systems 132, 133 using network links 152, 153. PC storage system 132 services channel attached hosts 157, 158.” EX1001, 4:67-5:6.



'988 Patent – Figure 1

B. Prosecution History

The '988 Patent issued on April 15, 2003, from U.S. Patent Application No. 09/236,409 (“’409 Application”), filed on January 22, 1999. Ex. 1001.

References cited by the applicant and considered by the Examiner include US 6,044,438 (Olnowich), Ex.1002, 80-85. At the close of prosecution, the Examiner issued a Notice of Allowance. The Examiner provided reasons for allowance, including that the taught by the prior art did not teach back-end software responsive to the configuration manager software for handling reads and writes to disks corresponding to the I/O requests but without communication over the I/O channel adapter, thereby separating disk operations from network and I/O traffic. Ex.1002, 225-230.

C. Claim Construction

According to 37 C.F.R. §42.100(b), in an *inter partes* review, claim terms are construed using the same standard used in a civil action under 35 U.S.C. 282(b), including construing the claim in accordance with the ordinary and customary meaning of such claim as understood by one of ordinary skill in the art and the prosecution history pertaining to the patent.

The ordinary and customary meaning of a term may be evidenced by a variety of sources, including the words of the claims themselves, the specification, drawings, and prior art. However, the best source for determining the meaning of a

claim term is the specification – the greatest clarity is obtained when the specification serves as a glossary for the claim terms. The words of the claim must be given their plain meaning unless the plain meaning is inconsistent with the specification. *In re Zletz*, 893 F.2d 319, 321, 13 USPQ2d 1320, 1322 (Fed. Cir. 1989); *Chef America, Inc. v. Lamb-Weston, Inc.*, 358 F.3d 1371, 1372, 69 USPQ2d 1857 (Fed. Cir.2004).

In accordance with these principles, Patent Owner submits that the terms of the '988 Patent are clear on their face, except the following claim term which should be construed as follows:

a. Volume Access Table (Claims 2, 4, 5, and 6)

“Volume access table” was not a commonly known term when the '409 application was filed.

In construing a claim term, we look to the words of the claim itself. *Vitronics Corp. v. Conceptronic, Inc.*, 90 F.3d 1576, 1582 (Fed.Cir.1996). If the claim term has a plain and ordinary meaning, our inquiry ends. *Id.* If, however, the claim term does not have an ordinary meaning, and its meaning is not clear from a plain reading of the claim, “we turn to the remaining intrinsic evidence, including the written description, to aid in our construction of that term.” *Telemac Cellular Corp. v. Topp Telecom, Inc.*, 247 F.3d 1316, 1326 (Fed.Cir.2001). As we have said, “the specification is always highly relevant to the claim construction analysis. Usually, it

is dispositive; it is the single best guide to the meaning of a disputed term.” *Vitronics*, 90 F.3d at 1582. The construction that stays true to the claim language and most naturally aligns with the patent's description of the invention will be, in the end, the correct construction. *Renishaw PLC v. Marposs Societa' per Azioni*, 158 F.3d 1243, 1250 (Fed.Cir.1998).

With these principles in mind, since “volume access table” does not have an ordinary meaning, The construction of the term “volume access table” must align with the patent's description of the invention.

The specification of the '988 Patent describes the volume access table:

Volume access table (450) in FIG. 4 contains a mapping between hosts and volumes specifying an access mode value. If the access mode is set to neither shared nor exclusive configuration manager forwards I/O requests directly to disk. In addition to the access mode said volume access table may contain other values that help to manager and improve performance of said data storage system. Ex. 1001, 6: 15-21.




Figure 4 (Excerpt) – '988 Patent

The patent thus contemplates that the volume access table is a table used to manage a data storage system, comprising a mapping between hosts and volumes specifying an access mode value, wherein the access mode value is shared or exclusive during cache processing.

The intrinsic record thus clearly indicates the '988 Patent 's precise conception, which here must control the construction.

Patent Owner reserves all rights to raise claim construction arguments in district court and in this proceeding should the Board grant institution.

III. LEVEL OF ORDINARY SKILL IN THE ART


For purposes of this Response, PO adopt Petitioner's proffered level of ordinary skill in the art because it comports with the technology and claims of the '988 Patent as well as the asserted prior art.

IV. THE ASSERTED PRIOR ART

A. Heil (Ex. 1006)

Heil is titled as “System and method for peer-to-peer accelerated I/O shipping between host bus adapters in clustered computer network,” and is directed to “retrieves data across independent computer nodes of a server cluster by providing for I/O shipping of block level requests to peer intelligent host-bus adapters (HBA).” EX1006, Abstract.

The architecture of a host system is shown in Figure 1 (reproduced below).



Heil – Figure 1

The system includes a CPU 100 being connected via a PCI bus 116.5 to a host bus adapter (HBA) 117. Ex.1006, 6:34-47. Host bus adapter 117 includes cache memory 117.2 and local memory 116 and is further connected to local drives 118

via interfaces 117.8 and 117.9. Ex.1006, 7:9-15 The host may process I/O requests locally by accessing the host bus adapter 117 via PCI bus 116.5. Ex.1006, 6:34-47.

The process of HBAs retrieving data corresponding to an I/O request for stored data blocks either locally or remotely is described as follows:

After receiving a block I/O request 400, the 110 redirector has the means to search the directory and determining means to locate the local or remote disk drives that are storing the I/O requested blocks 410. If the location of the blocks is local 420, then the blocks are retrieved from the respective local disk drive. Ex.1006, 11:45-52.


If the requested I/O block is stored in a remote disk, the initiating HBA ships the block I/O request to the peer HBA that the initiating HBA has determined can perform the requisite processing 450. Ex.1006, 11: 54-56. The remote HBA retrieves the requested blocks of data from its HBA cache, or if the blocks of data are not within cache, then directly from the remote HBA's local disk. Ex.1006,11:66-12:1.

If the search through the directory does not find a particular data block, the HBA queries peer HBAs to determine which HBA can retrieve the particular data block in the block I/O request. The initiating HBA sends a broadcast message to all peer HBAs and awaits for a response from the peer HBA in the node where the data blocks reside in local storage. After the remote HBA responds, the initiating HBA "ships" the block I/O request to the remote HBA. Ex.1006,12:8-16.

B. Nakayama (Ex. 1007)

Nakayama, titled “Storage control and computer system using the same,” discloses a data storage system, including a storage control enabling data on various storage media to be shared among host computers having various different host computer input/output interfaces. EX1007, Abstract.

The architecture of the data storage system is shown in Figure 3 (reproduced below).



Nakayama – Figure 3

In the system, a disk control 302 is connected to a host 40 computer 300 via a channel control 301 on the host side and is also connected to a host computer 303 via a small computer system interface. EX1007, 4:39-43. The disk control 302 is connected to drives 315 and 316 to perform read/write of data on the drives 315 and 316 in response to the requests of the host computers 300 and 303. The control processors 305 and 306 mainly perform data transfer between the host computers 300 and 303 and the cache memory 309. The control processors 310 and 311 mainly perform data transfer between the cache memory 309 and the drives 315 and 316. EX1007, 4:49-65.

The cache memory 309 is accessible from all control processors 305,306, 310 and 311 and is used for temporarily storing data which is read from the drives 315 and 316. A cache slot 312 is the data management unit quantity in the cache memory 309. EX1007, 5:5-9.

Nakayama describes the operation of data access processing as follows:

When a control processor 305 receives a data access command from the host computer 300, determination as to whether reservation of DCB is succeeded or not is made. If failed, the data access processing is ended. If it is successful, following operation will be conducted. EX1007, 6:7-10.

Firstly, reservation of the cache slot is conducted (step 603). Subsequently, a determination is made as to whether or not the data access command is a write command or a read command. EX1007, 6:16-19

If it is a write command, the control processor 305 writes the converted write data into the cache slot 312; The write data on the cache slot 312 is written into the drives 315, 316 by the control processors 310, 311 asynchronously with the operation of the host computers 300, 303. If the received command is a read command, data on the cache slot 312 is read, Subsequently, the cache slot 312 is released (step 614) and DCB is released. EX1007, 6:20-7-9.

C. Srinivasan (Ex. 1018)

Srinivasan discloses “Spritely NFS: Experiments with Cache Consistency Protocols.” in an ACM periodical. EX1018, 45. Srinivasan describes cache management strategies. A summary of NFS consistency model and Sprite consistency model is provided. In the Sprite consistency model, by tracking open and close operations the Sprite file server knows not only which clients are currently using a file, but whether any are potentially writers. EX1018, 46. If a file is write-shared, none of the clients (writers or readers) are allowed to cache it. EX1018, 47. When multiple clients open a file, they do not all issue their opens simultaneously. When a file that is open read-only is subsequently opened for write, the Sprite file

server must notify not only the newly arrived writer, but also the existing readers, that the file is no longer cachable. EX1018, 47.

D. Macklem (Ex. 1009)

Macklem discloses “Not Quite NFS, Soft Cache Consistency for NFS” in a conference paper presented at the Proceedings of the Winter 1994 USENIX Conference in January, 1994. Macklem describes an NFS-like protocol named Not Quite NFS providing full cache consistency during normal operation. EX1009, 261.

In the approach used by NQNFS is a variant of the Leases mechanism. In this model, the server issues to a client a promise, referred to as a "lease," that the client may cache a specific object without fear of conflict. EX1009, 262.

Whenever a client wishes to access a file's data it must hold one of three types of lease: read-caching, write caching or non-caching. The latter type requires that all file operations be done synchronously with the server via the appropriate RPCs. EX1009, 264.

V. PETITIONER HAS NOT DEMONSTRATED A REASONABLE LIKELIHOOD OF SUCCESS FOR THE GROUNDS ADVANCED IN THE PETITION, AND THE PETITION SHOULD BE DENIED

A. Requirements for Showing Obviousness Under 35 U.S.C. § 103

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) where in evidence, so called secondary considerations. *Graham v. John Deere Co.*, 383 U.S. 1, 17–18 (1966). The Board has held that a failure to identify the differences between the claimed subject matter and the prior art is fatal to an obviousness challenge. See, *Apple, Inc. v. Contentguard Holdings, Inc.*, IPR2015-00355, Decision Denying Institution of *Inter Partes* Review, Paper 9 at 9-10 (P.T.A.B. June 26, 2015) (denying institution for failure to identify the differences between the claimed subject matter and the prior art).

In arriving at an obviousness determination, the Board must sufficiently explain and support the conclusions that the prior art references disclose all the elements recited in the Challenged Claims and a relevant skilled artisan not only could have made, but would have been motivated to combine all the prior-art references in the way the patent claims. *Pers. Web Techs., LLC v. Apple, Inc.*, 848 F.3d 987, 994 (Fed. Cir. 2017). That is, even if all the claim elements are found across a number of references, an obviousness determination must consider whether a person of ordinary skill in the art would have been motivated to combine those references. *Intelligent Bio-Sys., Inc. v. Illumina Cambridge Ltd.*, 821 F.3d 1359, 1368 (Fed. Cir. 2016); *Los Angeles Biomedical Res. Inst. at Harbor-UCLA Med. Ctr. v. Eli Lilly Co.*, 849 F.3d 1049, 1067 (Fed. Cir. 2017) (vacating and remanding an obviousness determination, in part, because the Board did not make factual findings as to whether there was an apparent reason to combine all three prior art references

to achieve the claimed invention and whether a POSITA would have had a reasonable expectation of success from such a combination). This combination determination, as supported by an articulated motivation to combine, requires a plausible rationale as to why those prior art references would have worked together. *Broadcom Corp. v. Emulex Corp.*, 732 F.3d 1325, 1335 (Fed. Cir. 2013). Absent some articulated rationale, a “common sense” finding is no different than the conclusory statement “would have been obvious.” *In re Van Os*, 844 F.3d 1359, 1361 (Fed. Cir. 2017). Of additional importance, “knowledge of a problem and motivation to solve it are entirely different from motivation to combine particular references.” *Innogenetics, N.V. v. Abbott Labs.*, 512 F.3d 1363, 1373 (Fed. Cir. 2008).

B. Petitioner Fails to Establish in Ground 1 that Challenged Claims 1-4 and 6 are Obvious Over Heil and Nakayama

1. Claim 1

Independent claim 1 has been divided into elements for consideration by the Petitioner as follows:

[1. pre] A computer suitable for use in a data storage system comprising a network interconnecting a plurality of such computers, the computer comprising:

[1.a] an I/O channel adapter for accepting an incoming I/O request from a host;

[1.b] configuration manager software for enabling said I/O channel adapter to decide whether (i) to route said request to cache, (ii) to route said request to disk, or (iii) to reject said request;

[1.c] a network adapter for handling network control traffic;

[1.d] a cache memory;

[1.e] front-end software for handling I/O requests arriving at the I/O channel adapter or the network adapter;

[1.f] cache manager software, responsive to said front-end software, for handling data stored in said cache memory; and

[1.g] back-end software, responsive to said configuration manager software, for handling reads and writes to disks corresponding to the I/O requests but without communication over the I/O channel adapter, thereby separating disk operations from network and I/O traffic.

a. Heil and Nakayama do not teach the limitation [1.b]

Regarding the limitation [1.b] “configuration manager software for enabling said I/O channel adapter to decide whether (i) to route said request to cache, (ii) to route said request to disk, or (iii) to reject said request,” the ’988 Patent describes “configuration manager software for enabling said I/O channel adapter to decide whether (i) to route said request to cache” as follows:

The front-end module 310 of FIG. 3 allocates a channel and waits for I/O requests from the initiating hosts 111 or 112. Remote operations begin in step 402. Depending upon the status of the value in a volume access table 450 the requests are routed either as shown in FIG. 4A for write exclusive mode, FIG. 4B for read exclusive, FIG. 4C for write shared or FIG. 4D for read shared. Ex.1001, 6: 4-10.

The '988 Patent describes “(ii) to route said request to disk” as follows:

Volume access table (450) in FIG. 4 contains a mapping between hosts and volumes specifying an access mode value. If the access mode is set to neither shared nor exclusive configuration manager forwards I/O requests directly to disk. Ex.1001, 15-19.

The language of the limitation [1.b] and corresponding description in the specification shows that (i) to route said request to cache and (ii) to route said request to disk are independent and distinct access processes, which means “to route said request to disk” does not comprise “to route said request to cache,” since “to route said request to disk” is to route said request directly to disk.

To argue that Heil discloses “configuration manager software for enabling said I/O channel adapter to decide whether (ii) to route said request to disk, Petitioner proffers that:

Heil shows that I/O requests are routed from front-end interface 102 to either cache 117.2 or local drives 118:

“To improve I/O performance, CPU 117.1 temporarily caches disk data in cache memory 117.2 that is frequently accessed. Later, the CPU 117.1 stores the data permanently in the slower disk drive devices 118.” (Citing Ex.1006, 7:24-30). “The remote HBA retrieves the requested blocks of data from its HBA cache, or if the blocks of data are not within cache, then directly from the remote HBA's local disk.” (Citing Ex.1006, 11:66-12:5). Pet. 23-24.


Patent Owner respectfully submits that although the HBA can route the request to disk, it does not route the request to disk directly, on the contrary, the HBA first route the request to its HBA cache, if the blocks of data are not within cache, then directly from the remote HBA's local disk. It means that “to route said request to disk” comprises “to route said request to cache,” which is contrary to the limitation [1.b].

Therefore, Heil does not disclose the limitation [1.b], particularly, Heil does not disclose the limitation “configuration manager software for enabling said I/O channel adapter to decide whether (ii) to route said request to disk.

Referring to Fig. 3 below, Nakayama discloses a disk subsystem having a cache memory, The control processors 305 and 306 are connected to the host computers 300 and 303 through the channel control 301 are connected to the drives 315 and 316 through the drive interfaces 313 and 314, respectively. The control processors 305 and 306 mainly perform data transfer between the host computers

300 and 303 and the cache memory 309. The and the SCSI bus control 304, respectively and the control processors 310 and 311 control processors 310 and 311 mainly perform data transfer between the cache memory 309 and the drives 315 and 316. Ex.1007, 4: 56-65.

The written description of Nakayama shows that the I/O channel adapter is connected with cache memory, the I/O channel adapter is connected with disk through cache memory. the I/O channel adapter does not connect to disk directly, thus configuration manager software of Nakayama can only enable the I/O channel adapter to route the request to cache, cannot enable the I/O channel adapter to route said request to disk directly. Nakayama does not disclose the limitation “configuration manager software for enabling said I/O channel adapter to decide whether (ii) to route said request to disk.



Nakayama – Figure 3

b. Heil in view of Nakayama Does Not Render Claim 1 Obvious

Petitioner fails to undertake a Graham analysis, because Petitioner's combination of Smith and Baumann lacks evidence of crucial limitations of claim 1, such as the elements [1.b]. *Graham v. John Deere Co.*, 383 U.S. 1, 17–18 (1966). Therefore, combination of Heil and Nakayama does not render claim 1 obvious because it has not established its prima facie case, since “obviousness requires a

suggestion of all limitations in a claim.” *CFMT, Inc. v. Yieldup Intern. Corp.*, 349 F.3d 1333, 1342 (Fed. Cir. 2003).

2. Claim 2

Dependent claim 2 has been divided into elements for consideration by the Petitioner as follows:

[2.pre] The system of claim 1 wherein the computers comprise off-the-shelf hardware and operating systems and further comprise:

[2.a] an adapter I/O software for accepting incoming I/O requests from a host;
and

[2.b] a volume access table employed by the configuration manager to ensure consistency of data stored on the network.

a. Heil and Nakayama do not teach the limitation [2.b]

To argue that Heil discloses the limitation [2.b], Petitioner proffers that:

Heil discloses that its I/O redirector software 240 (the configuration manager) uses a directory (a volume access table) to maintain consistency of the data stored on the network. For example, Heil’s I/O redirector software 240 decides whether to route an I/O request received by the HBA from the host by accessing a directory “mapping the location of blocks in the entire cluster”—including both local and/or remote locations of blocks—to determine the location of data on the network responsive to the I/O request. (Citing Ex.1006, 12:19-25, 12:34-42, 11:45-54). The

directory stores “location information regarding blocks of data stored” to determine where to route the I/O request. (citing Ex.1006, 4:8-12). The directory is updated as new blocks of data are stored to reduce latency on the network. Ex.1006, 12:15-19, 13:36-48. As Heil explicitly describes methods of using the directory to maintain the consistency of data in the storage system, a POSITA would have understood that using the directory with the I/O redirector would similarly maintain the consistency of data in the storage system. (Citing Ex.1004, ¶¶123-124). The POSITA would thus understand that Heil’s directory has the same functionality as and discloses the claimed volume access table of claim 2 of the ’988 Patent. Pet. 41-42.

Patent Owner respectfully submits that petitioner errs on equating Heil’s “directory” to “volume access table” of the ’988 Patent.

As set forth in II.C.a, “volume access table” is a table comprising a mapping between hosts and volumes specifying an access mode value. wherein the access mode value is shared or exclusive during cache processing (it is used by cache manager 320, Ex. 1001, 6:43-45).

Heil’s “directory” is used for storing location information regarding blocks of data stored within the plurality of storage devices. Ex.1006, 4:8-10. Heil does not specify that the “directory” is a kind of table. To the extent that the “directory” includes a table, the table will comprise a mapping between hosts and volumes specifying whether or not a volume is hold by a local host or a remote host, such that

“searching the directory will determine whether to locally or remotely retrieve blocks of data.” Ex.1006, 4:10-13. Since “directory” specifies that some data volumes can be accessed by particular hosts, and does not specify the cache mode by particular hosts “directory” is irrelevant to cache processing.

“volume access table” of the ’988 Patent is a table comprising a mapping between hosts and volumes specifying an access mode value. wherein the access mode value is shared or exclusive during cache processing, shared/exclusive access mode in the ’988 Patent specifies the cache mode by particular hosts, but does not specify some data volumes can be accessed by particular hosts. That is, “directory” in Heil and “volume access table” in the ’988 Patent have different structures.

The function of the “volume access table” in the ’988 Patent is to ensure consistency of data stored in the cache during cache processing, wherein “directory” in Heil is to maintain the consistency of data in the storage system. Therefore, “directory” in Heil and “volume access table” in the ’988 Patent have different structures and functions, Heil does not disclose “volume access table” as required by the limitation [2.b].

Nakayama does not disclose “volume access table” either.

Petitioner argues that the term “volume access table” should be construed broadly without reading any particular embodiment of the ’988 Patent into the construction, **to the extent a volume access table contains additional data**

specifying an access mode, and that this is rendered obvious by Nakayama. (Citing Ex. 1004, ¶¶125-126). For example, Nakayama discloses storing data indicating whether a host has exclusive access to a drive (i.e., the host reserves the drive, thereby controlling access to it and excluding other host computers from accessing the drive): “The **interhost exclusive information** 403 is used when the host computer 300 or 303 exclusively controls the drive access from the other host computer and sets the information 403 ‘on’ when the processor reserves access of the specified drive number and sets it ‘off’ when the host computer releases the access right.” Ex.1007, 5:43-47, 5:31-34. Nakayama further discloses that it is desirable to prevent any disk storage “from being predominantly used by” a host computer, which would reduce the availability of the disk storage for other hosts that may want to retrieve data from the disk storage or store data to the disk storage. citing Ex.1006, 7:54-56. To prevent a particular host from predominantly using a specific disk storage, the POSITA would include this permission information taught by Nakayama in Heil’s directory to control access to each disk storage using the exclusivity indicator as taught by Nakayama. Ex.1007, 5:43-47, 5:31-34. Moreover, Heil teaches controlling access to blocks of disk storage to maintain consistency (citing Ex.1006, 2:18-32), and a POSITA would understand that Nakayama’s access modes would improve control of access to disk drives and would help achieve Heil’s

objective of controlling access to disk to maintain consistency of the stored data. citing Ex.1004, ¶¶125-126. Pet.42-43.

However, the cited portion of Petition shows that Petitioner confused “interhost exclusive information” of Nakayama with “shared/exclusive” access mode in the ’988 Patent.

Not unlike in Heil, “interhost exclusive information” is used to determine specific host computer exclusively controls specific drive access, not the access mode used in the period of cache processing. The function of “interhost exclusive information” is to maintain the consistency of data in the storage system.

Therefore, “interhost exclusive information” in Nakayama and “volume access table” in the ’988 Patent have different structures and functions, Nakayama does not disclose “volume access table” as required by the limitation [2.b].

Thus, Petitioner has not met his burden to show how the combination of Heil and Nakayama teaches the limitation [2.3].

b. Heil in view of Nakayama Does Not Render Claim 2 Obvious

Petitioner fails to undertake a Graham analysis, because Petitioner’s combination of Heil and Nakayama lacks evidence of crucial limitations of claim 2, such as the element [2.b]. *Graham v. John Deere Co.*, 383 U.S. 1, 17–18 (1966). Therefore, combination of Bergsten and Blumenau does not render claim 2 obvious because it has not established its *prima facie* case, since “obviousness requires a

suggestion of all limitations in a claim.” *CFMT, Inc. v. Yieldup Intern. Corp.*, 349 F.3d 1333, 1342 (Fed. Cir. 2003).

3. Claim 3

Claim 3 depends from claim 1 and is not anticipated by Bergsten for at least the same reasons.

4. Claim 4

Claim 4 recites: The system of claim 3 further comprising a volume access table employed by the configuration manager to ensure consistency of data stored in the distributed cache.

Claim 4 comprises limitation “volume access table,” similar to claim 2, the combination of Heil and Nakayama does not disclose “volume access table.” Therefore, claim 4 is not obvious over Heil in view of Nakayama.

5. Claim 6

Independent claim 6 has been divided into elements for consideration by the Petitioner as follows:

[6.pre] A method of accessing a remote disk over a computer network without incurring network overhead, the method comprising the steps of:

[6.a] a. causing a local host to issue a request over an I/O channel to a local computer;

[6.b] b. providing a configuration manager on the local computer, the configuration manager routing the request to a remote computer via the computer network;

[6.c] c. causing the remote computer to check the request against a volume access table;

[6.d] d. causing the remote computer to perform an I/O operation on a disk located on the remote computer and to return data to the local computer;

[6.e] e. causing the local computer to provide the returned data to the local host via the I/O channel; and

[6.f] f. causing the local computer to check the data against the volume access table to ensure consistency of the data on the local and the remote computers.

The elements [6.c] and [6.f] comprise limitation “volume access table,” similar to claim 2, the combination of Heil and Nakayama does not disclose “volume access table.” Therefore, claim 6 is not obvious over Heil in view of Nakayama, the Board should deny Petitioner’s Petition because it has not established its *prima facie* case.

C. Petitioner Fails to Establish in Ground 2 that Challenged Claim 5 is Obvious Over Heil in view of Nakayama, Macklem and Srinivasan

Dependent claim 5 has been divided into elements for consideration by the Petitioner as follows:

[5.pre] The system of claim 4,

[5.a] wherein the configuration manager includes software that checks an access mode in the volume access table and

[5.b] (i) if the access mode is set to an exclusive mode, causes both reads and writes to be stored in the cache memory, and causes invalidate messages to be sent to remote storage systems;

[5.c] (ii) if the access mode is set to shared, causes only reads to be stored in the cache memory; and

[5.d] (iii) if the access mode is set to a value other than exclusive or shared, causes reads and writes to be performed directly to a disk without using the cache memory.

a. The combination of Heil, Nakayama, Macklem and Srinivasan does not teach the limitation [5.a]

The limitation [5.a] comprises element “volume access table,” similar to claim 2, the combination of Heil and Nakayama does not disclose “volume access table.”

Srinivasan also does not teach “volume access table.”

The volume access table of the '988 patent contains a mapping between hosts and volumes specifying an access mode value.

Srinivasan discloses a server state table, in the server state table, each file may be in one of several states. The states are:

CLOSED	File not open by any client.
CLOSED_DIRTY	File not open, but the last writer may still have dirty blocks.
ONE_READER	File open read-only by one client.
ONE_RDRDIRTY	File open read-only by one client, which may have dirty blocks cached from a previous open.
MULT_READERS	File open read-only by two or more clients.
ONE_WRITER	File open read-write by one client.
WRITE_SHARED	File open by two or more clients, including at least one writer.

EX1018, 50

The “server state table” is used to record the states of files used by clients, it does not specify access mode value for any file in relate to clients, “server state table” is distinct from “volume access table” of the '988 Patent.

Macklem does not teach “volume access table” either.

Macklem discloses in the NQNFS model, the server issues to a client a promise, referred to as a "lease," that the client may cache a specific object without fear of conflict. Ex. 1009, 262. Whenever a client wishes to access a file's data it must hold one of three types of lease: read-caching, write caching or non-caching. Ex. 1009, 264. The server must maintain state information about all current client leases. Ex. 1009, 264.

Although “current client lease” contains a mapping between clients and files specifying an access mode, it is not a mapping between clients and volumes of the storage device, files and volumes are distinct, volumes of the storage device can

store different files and files can be stored in different volumes. Further, the access modes are read-caching, write caching or non-caching, not “shared or exclusive” mode of the ’988 Patent. Therefore, “current client leases” in Macklem and “volume access table” in the ’988 Patent have different structures. Macklem does not disclose “volume access table” as required by the limitation [5.a].

Therefore, the combination of Heil, Nakayama, Srinivasan and Macklem does not disclose the limitation [5.a].

b. The combination of Heil, Nakayama, Macklem and Srinivasan does not teach the limitations [5.b] and [5.c]

Petitioner asserts that Heil in view of Nakayama discloses checking an access mode in a volume access table. Both Srinivasan and Macklem further disclose an exclusive mode that causes reads and writes to be stored in the cache memory. Pet., 58.

First, as set forth above, Heil in view of Nakayama does not disclose “volume access table,” thus, Heil in view of Nakayama does not disclose checking an access mode in a volume access table.

Second, Srinivasan does not disclose limitations [5.b] and [5.c].

Petitioner contends that Srinivasan describes at least three protocols that limit read and write operations on a cache : a “single-writer” “non-write-shared” mode, a “read-only” “non-write-shared” mode, and a “write-shared” mode. Pet., 59.

Srinivasan discloses the server tracking operation to get the state of access mode of files as follows:

“by tracking open and close operations the Sprite file server knows not only which clients are currently using a file, but whether any of them are potentially writers.” Ex. 1018, 46.

Although “single-writer” “non-write-shared” mode, a “read only” “non-write-shared” mode, and a “write-shared” mode in Srinivasan are access mode, it is **tracked** by server, it is the states of cache mode used by files, not access mode **set** by the server. Thus, Srinivasan does not disclose the access mode is set to an exclusive mode or shared mode as required by the limitations [5.b] and [5.c]. Further Srinivasan does not disclose “if the access mode is set to an exclusive mode, causes both reads and writes to be stored in the cache memory, and **causes invalidate messages to be sent to remote storage systems**” as required by the limitation [5.b].

Third, Macklem does not disclose limitations [5.b] and [5.c].

Although Macklem describes multiple modes of operation that restrict client access to cache, including “write-caching lease,” “read-caching lease,” and a “non-caching” mode. Macklem does not disclose exclusive mode or shared mode as required by the limitations [5.b] and [5.c]. Further Macklem does not disclose “if the access mode is set to an exclusive mode, causes both reads and writes to be stored

in the cache memory, and causes invalidate messages to be sent to remote storage systems” as required by the limitation [5.b].

Therefore, combination of Heil, Nakayama, Macklem and Srinivasan does not render claim 5 obvious because it has not established its *prima facie* case, since “obviousness requires a suggestion of all limitations in a claim.” *CFMT, Inc. v. Yieldup Intern. Corp.*, 349 F.3d 1333, 1342 (Fed. Cir. 2003).

It is wrong to use the patent in suit as a guide through the maze of prior art references, combining the right references in the right way so as to achieve [a desired result]. *Orthopedic Equip. Co. v. United States*, 702 F.2d 1005, 1012, 217 USPQ 193, 199 (Fed. Cir. 1983). Petitioner’s usage of the ’988 Patent’s teaching to explain the combination of prior art references is an exercise in impermissible hindsight.

VI. CONCLUSION

For the foregoing reasons, Patent Owner respectfully requests that the petition for *inter partes* review be denied.

Dated: August 25, 2023

Respectfully submitted,
RAMEY LLP

/Jacob B. Henry/
Jacob B. Henry
Registration No. 61,093
5020 Montrose Blvd., Suite 800
Houston, Texas 77006
Telephone: (832) 454-1456
Facsimile: (832) 900-4941
jhenry@rameyfirm.com

CERTIFICATION OF WORD NUMBER

Counsel for Patent Owner hereby certifies that the foregoing document contains no more than 8,000 words as calculated using the word count function of Microsoft Word, not including material excluded under 37 C.F.R. 42.24(a)(1).

Date: August 25, 2023

/Jacob B. Henry/
Jacob B. Henry

CERTIFICATE OF SERVICE

I hereby certify that all counsel for Petitioner who have appeared in this case and identified below are being served on this day of August 25, 2023, with a copy of the foregoing document via e-mail.

/s/ Jacob B. Henry
Jacob B. Henry

Lead Counsel	Backup Counsel
James T. Carmichael Reg. No. 45,306 jim@carmichaelip.com Carmichael IP, PLLC 8607 Westwood Center Drive, Suite 270 Tysons, VA 22182 Tel: 703-646-9255	Minghui Yang Reg. No. 71,989 mitch@carmichaelip.com Carmichael IP, PLLC 8607 Westwood Center Drive, Suite 270 Tysons, VA 22182 Tel: 703-646-9248