

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

GOOGLE LLC
Petitioner

v.

LS CLOUD STORAGE TECHNOLOGIES LLC
Patent Owner

U.S. Patent No. 10,154,092
IPR2023-00120

EX. 1004

DECLARATION OF DR. PAUL FRANZON, PH.D.

TABLE OF CONTENTS

	<u>Page</u>
I. INTRODUCTION	1
II. QUALIFICATIONS	2
III. MATERIALS CONSIDERED	5
IV. APPLICABLE LEGAL STANDARDS	7
A. Claim Construction	7
B. Anticipation (35 U.S.C. § 102)	8
C. Obviousness (35 U.S.C. § 103)	8
V. LEVEL OF ORDINARY SKILL IN THE ART	12
VI. THE STATE OF THE ART AT THE TIME OF INVENTION	13
VII. THE '092 PATENT	20
A. Claim Construction	22
VIII. UNPATENTABILITY ANALYSIS	25
A. Overview of the Prior Art	25
1. Heil	25
2. Nakayama	26
3. Gulick	27
4. Berman	28
B. Grounds 1 and 2: Anticipation of Claims 1-3, 7-12, and 19-23 by Heil	28

1.	Claim 1	29
a)	1.pre: “An apparatus comprising”	29
b)	1.a: “a first interface configured to receive input/output (I/O) traffic from a first host device via a dedicated I/O channel, the I/O traffic comprising a read command”	29
c)	1.b: “a second interface configured to receive first data via a network”	32
d)	1.c: “a cache memory configured to store second data”	34
e)	1.d: “a storage device configured to store third data”	35
f)	1.e: “a processor coupled to the cache memory, the processor coupled to the storage device via a communication path that is distinct from the dedicated I/O channel, the processor configured to access the cache memory during processing of the I/O traffic”	36
g)	1.f: “the processor further configured to perform an access operation at the storage device based on the I/O traffic”	40
2.	Claim 2	41
3.	Claim 3	42
4.	Claim 7	44
5.	Claim 8	46
6.	Claim 9	48

a)	9.a: “the storage device and the dedicated I/O channel are independently accessible”	48
b)	9.b: “wherein the processor is configured to read at least a portion of the second data from the cache memory based on the read command”	50
7.	Claim 10	50
8.	Claim 11	52
a)	11.a: “configuration manager circuitry configured to route an I/O request included in the I/O traffic to the cache memory, to route the I/O request to the storage device, or to deny the I/O request”	52
b)	11.b: “front-end circuitry configured to process the I/O request”	54
c)	11.c: “back-end circuitry configured to perform a read operation or a write operation at the storage device based on the I/O request”	55
9.	Claim 12	56
10.	Claim 19	57
a)	19.pre: “An apparatus comprising”	57
b)	19.a: “a first interface configured to receive input/output (I/O) traffic from a host computer via a dedicated I/O channel, the I/O traffic comprising one or more read commands, one or more write commands, or a combination thereof”	57
c)	19.b: “a second interface configured to receive data via a network”	57

d)	19.c: “a cache memory”	58
e)	19.d: “a storage device”	58
f)	19.e and 19.f: “a processor coupled via a communication path to the storage device, the processor configured to access the cache memory during processing of the I/O traffic, the processor further configured to perform one or more access operations at the storage device based on the I/O traffic, wherein the communication path is distinct from the dedicated”	58
11.	Claim 20	58
12.	Claim 21	59
13.	Claim 22	59
14.	Claim 23	59
C.	Ground 3: Obviousness of Claims 1-3 and 6-24 Over Heil and Nakayama.....	59
1.	Motivation to Combine.....	60
2.	Claim 1	62
a)	1.pre: “An apparatus comprising”	63
b)	1.a: “a first interface configured to receive input/output (I/O) traffic from a first host device via a dedicated I/O channel, the I/O traffic comprising a read command”	63
c)	1.f: “the processor further configured to perform an access operation at the storage device based on the I/O traffic”	66

3.	Claim 2	67
4.	Claim 3	68
5.	Claim 6	68
6.	Claim 7	70
7.	Claim 8	70
8.	Claim 9	70
9.	Claim 10	72
10.	Claim 11	73
11.	Claim 12	76
12.	Claim 13	77
	a) 13.pre: “A method comprising:”	77
	b) 13.a: “receiving input/output (I/O) traffic from a host device via a dedicated I/O channel at a first interface, the I/O traffic comprising a write command”	78
	c) 13.f: “performing one or more access operations at the storage device based on the I/O traffic, the one or more access operations utilizing a communication path between a processor and the storage device, the communication path distinct from the dedicated I/O channel.”	81
13.	Claim 14	82
14.	Claim 15	83
15.	Claim 16	85
16.	Claim 17	87

a)	17.a: “The method of claim 13, receiving a read request via the network”	87
b)	17.b: “transmitting at least a portion of the second data from the cache memory via the network responsive to the read request being associated with a cache hit”	89
17.	Claim 18	90
18.	Claim 19	93
a)	19.pre: “An apparatus comprising:”	94
b)	19.a: “a first interface configured to receive input/output (I/O) traffic from a host computer via a dedicated I/O channel, the I/O traffic comprising one or more read commands, one or more write commands or a combination thereof”	95
19.	Claim 20	96
20.	Claim 21	97
21.	Claim 22	97
22.	Claim 23	99
23.	Claim 24	99
a)	24.pre: “An apparatus comprising”	99
b)	24.a: “means for receiving input/output (I/O) traffic from a first host device via a dedicated I/O channel, the I/O traffic comprising a write command”	100
c)	24.b: “means for receiving data via a network”	104
d)	24.c: “means for short-term data storage”	106

e)	24.d: “means for long-term data storage”	108
f)	24.e: “means for performing one or more access operations at the means for short-term data storage during processing of the I/O traffic and for performing one or more access operations at the means for long-term data storage based on the I/O traffic, the means for performing distinct from the dedicated I/O channel”	109
D.	Ground 4: Obviousness of Claim 4 Over Heil, Nakayama, and Gulick	113
1.	Motivation to Combine	113
2.	Claim 4	114
E.	Ground 5: Obviousness of Claim 5 Over Heil, Nakayama, and Berman	116
1.	Motivation to Combine	117
2.	Claim 5	118
F.	Secondary Considerations of Non-Obviousness	120
IX.	CONCLUSION	121

I. INTRODUCTION

1. I have been asked to submit this declaration on behalf of Google, LLC (“Petitioner”) in connection with a petition for *inter partes* review of U.S. Patent No. 10,154,092 (“the ’092 Patent,” Ex. 1001). Specifically, I have been retained as an independent expert consultant by Petitioner to provide my opinions on the technology claimed in, and the patentability or unpatentability of, claims 1-24 of the ’092 Patent (“the challenged claims”). Although I am being compensated at my usual rate of \$500 per hour for the time I spend on this matter, no part of my compensation depends on the outcome of this proceeding, I have no financial interest in any of the parties, and I have no other interest in this proceeding.

2. Unless otherwise noted, the statements made herein are based on my personal knowledge, and if called to testify about this declaration, I could and would do so competently and truthfully.

3. I am not a legal expert and offer no opinions on the law. However, I have been informed by counsel of the various legal standards that apply, and I have applied those standards in arriving at my conclusions.

4. The opinions I set forth herein are my own and are based on my education, experience, training and skill I have accumulated.

II. QUALIFICATIONS

5. I understand that my curriculum vitae (“CV”) is attached to this Report as Appendix A. It provides an accurate description of my background and experience qualifying me to provide my opinions in this matter.

6. I completed my Ph.D. in Electrical and Electronic Engineering in 1989 from the University of Adelaide in Australia. I obtained two additional degrees from the University of Adelaide: a Bachelor of Engineering in Electrical and Electronic Engineering in 1984 and a Bachelor of Science in Physics and Mathematics in 1983.

7. I am currently Distinguished Professor and Director of Graduate Programs in the Department of Electrical and Computer Engineering at North Carolina State University (“NCSU”) in Raleigh, North Carolina. I have been affiliated with NCSU in various roles since 1989.

8. I am a Fellow of the IEEE.

9. I have over 30 years’ experience designing digital hardware including computers, computer interfaces and peripherals. I detail some of those experiences relevant to this report below.

10. In 1981/82 I was an intern at Australia Telecom where I reprogrammed a database system for logging system alarms in the state's communications backbone. The goal was to speed up the system. The system was implemented on a minicomputer and how to manage disk input/output was a major consideration.

11. In the late 1980's some colleagues and I started a company called Communica Pty. Ltd. The company's initial goal was to build a microwave link between two Ethernet based local area networks. I was involved in planning and design.

12. In the 1990's, I taught ECE 521 Computer Design and Technology and ECE 218 Computer Organization and Microprocessors. ECE 521 dealt with the design tradeoffs in pipelined and multiple-issue (i.e. superscalar) microprocessors. It included a section on input/output systems and standards, including disk drives. ECE 218 focused primarily on how to program a M68000 microprocessor, and included a section on Input/Output that emphasized interrupt and exception handling, as well as polling. Students had to get the Input/Output structures working as interfaced to a keyboard and display.

13. Around the same time I was doing research on computer architecture. This included a project designing a superscalar microprocessor disaggregated onto a multichip module.

14. I cofounded Lightspin Technologies in 2001 and filled the role of Vice President of Engineering for two years. The goal of the company was to commercialize a new high-speed modulated LED light source. The target application was Metropolitan Area Networks so we spent some time understanding their architecture. The company ran out of funds as the .com era of rapid growth came to a close.

15. In the early 2000's I was the NCSU Principal Investigator for an optical networking project. My team contributed to the logic design of the router. We filed two patents on the forwarding engine algorithms. These patents were applicable to a range of networks, not just optical ones.

16. In the mid-2000's I was the NCSU PI on a network security project. The idea was to develop ASICs to detect intruder attacks based on the behavior of the network traffic. We did partial implementations using FPGAs. We also developed hardware to build filters to detect and block intruder attacks. We developed an accelerator for a Snort packet filter.

17. In the 2010's I was the PI of a project building a 3D heterogeneous processor stack. The project covered CPU, cache and IO design. The CPU stack was fabricated at Global Foundries and Ziptronix. We also architected similar concepts in an x86 CPU stack.

18. In the late 2010's I led a series of projects intended to develop a low power SIMD CPU aimed at executing neuromorphic applications. Several versions of the CPU were designed and/or built including conventional versions and chipletized versions.

19. Since 2020 I have been leading a project implementing a FPGA system to intercept RansomWare attacks in real time. The system sits on the network of a computer or data system and monitors incoming network traffic. It can identify an attack in less than 1 ms.

III. MATERIALS CONSIDERED

20. In forming my opinions expressed in this declaration, I have considered, among other things, the following documents. I understand the documents have been given the following exhibit numbers in this proceeding:

<u>Exhibit Number</u>	<u>Description</u>
1001	U.S. Patent No. 10,154,092
1002	File History of U.S. Patent No. 10,154,092 (Application No. 14/997,327)
1003	File History of U.S. Patent No. 6,549,988 (Application No. 09/236,409)
1006	U.S. Patent No. 6,173,374 (“ <i>Heil</i> ”)
1007	U.S. Patent No. 5,920,893 (“ <i>Nakayama</i> ”)
1008	U.S. Patent No. 5,692,211 (“ <i>Gulick</i> ”)
1009	U.S. Patent No. 6,118,776 (“ <i>Berman</i> ”)
1014	D. Patterson, G. Gibson, R. Katz, entitled “A Case for Redundant Arrays of Inexpensive Disks (RAID)”. SIGMOD88: International Conference On Management of Data Chicago Illinois USA June 1 - 3, 1988
1015	R. H. Katz, "Network-attached storage systems," <i>Proceedings Scalable High Performance Computing Conference SHPCC-92.</i> , Williamsburg, VA, USA, 1992, pp. 68-75.

1016	Gang Ma and A. L. Narasimha Reddy, "An evaluation of storage systems based on network-attached disks," Proceedings. 1998 International Conference on Parallel Processing (Cat. No.98EX205), Minneapolis, MN, USA, 1998, pp. 278-285. doi: 10.1109/ICPP.1998.708497
------	---

IV. APPLICABLE LEGAL STANDARDS

A. Claim Construction

21. I understand that claims of a patent are generally interpreted according to their ordinary and customary meaning taking into consideration the so-called “intrinsic evidence” of the patent consisting of (1) the claim language; (2) the specification; and (3) the prosecution history.

22. I understand that claim terms may sometimes be defined by the specification. For example, a claim term may be explicitly defined in the patent specification, or it may be implicitly defined through consistent usage in the specification. I also understand that the scope of claim terms may be limited by

statements in the specification or prosecution history where the applicant clearly disavows or disclaims subject matter.

B. Anticipation (35 U.S.C. § 102)

23. I have been informed that to obtain a patent the claimed invention must have, as of the date of the invention, been novel over the prior art. I understand that an invention is not novel if it is anticipated by the prior art. I understand that an invention is anticipated if a single prior art reference discloses all the limitations or elements as set forth in the claim. I understand that in some cases an element or limitation, while not explicitly or expressly disclosed in a prior art reference, may still be disclosed by the prior art reference if that element or limitation is inherent in the reference. I understand that an element or limitation is inherent if the nature of the reference is such that the element or limitation is necessarily present, or put another way, that it would have to follow from the disclosed system or method that the element or limitation must be present.

C. Obviousness (35 U.S.C. § 103)

24. I have been informed that a patentable invention must, as of the date of the invention, have been non-obvious in view of the prior art in the field. I understand that an invention is obvious 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 of the invention to a person having ordinary skill in the art (POSITA).

25. I have been informed that the following factors are evaluated to determine whether a Petitioner has met its burden of proof that a claimed invention is obvious:

- a. The scope and content of the prior art relied upon;
- b. The difference or differences, if any, between each claim of the patent and the prior art; and
- c. The level of ordinary skill in the art at the time the invention of the patent was made.

26. These factual inquiries determine whether the claimed subject matter would have been obvious to one of ordinary skill in the art at the time the alleged invention was made.

27. I understand that various objective indicia can be evidence of non-obviousness, including, for example: (1) commercial success of the invention; (2) a long-felt need for the invention; (3) copying by others of the subject matter of the claim invention; (4) failure by others to find the solution provided by the

invention; (5) skepticism or expressions of surprise from experts and those skilled in the art; (6) unexpected results of the claimed invention; (7) acceptance of others and industry praise; and (8) licensing of the patents.

28. I also understand that “obviousness” is a legal conclusion based on the underlying factual issues of the scope and content of the prior art, the differences between the claimed invention and the prior art, the level of ordinary skill in the prior art, and any objective indicia of non-obviousness. For that reason, I am not rendering a legal opinion on the ultimate legal question of obviousness. Rather, my testimony addresses the underlying facts and factual analysis that would support a legal conclusion of obviousness or non-obviousness, and when I use the term obvious, I am referring to the perspective of a POSITA at the time of invention.

29. I understand that to prove that prior art or a combination of prior art renders a patent obvious, it is necessary to: (1) identify the prior art references, that singly or in combination, render the patent obvious; (2) specifically identify which limitations of a patent claim appear in the prior art references; and (3) explain how the POSITA would have been combined the prior art references in order to render the invention defined by the patent claim.

30. I have been informed that a finding of obviousness may be supported by showing that a person of ordinary skill would have had an apparent reason to combine or modify the prior art references as proposed. I understand that sometimes this is referred to as a motivation to combine. The apparent reason or motivation to combine need not be explicitly found in the prior art. On the contrary, a reason to combine may be based on a variety of factors, such as the nature of the problem to be solved, the existence of a need or problem known in the field, common knowledge and design skill in the field, and the ordinary creativity and common sense of a POSITA.

31. In some cases, an invention is presumptively obvious where prior art elements with known functions are combined to produce a predictable result. In other words, the invention is obvious because it is merely the combination of known elements producing an expected result. Under some circumstances, obviousness under §103 may be proved merely by showing that a combination of elements was obvious to try.

32. I understand that the obviousness inquiry should guard against hindsight bias or hindsight reconstruction where after-the-fact reasoning is applied to combine prior art elements using the claimed invention as a template, without

establishing that, as of the date of the invention, there exists a rationale to combine the prior art as proposed or that the combination is otherwise obvious.

33. In assessing obviousness, I have been instructed that a reference must qualify as analogous art in order to qualify as proper prior art for an obviousness analysis. I have been informed that a reference qualifies as analogous art with respect to the claims if it is either: (1) from the same field or endeavor as the patent; or (2) the reference is reasonably pertinent to the particular problem addressed by the invention.

V. LEVEL OF ORDINARY SKILL IN THE ART

34. My opinion on the level of ordinary skill in the art is based on the date of the invention, which I have been informed for my analysis to assume a date no earlier than January 22, 1999.

35. In deciding the level of ordinary skill as of the time of the invention, I considered:

- a. The types of problems encountered in the field and the solutions developed in the field;
- b. The sophistication of the technology;

c. The levels of education and experience of persons working in the field; and

d. My own experience in working with those of skill in the art.

36. Based on my experience in the field and the disclosure of the '092 Patent, a person of ordinary skill in the art ("POSITA" or "skilled artisan") in the field of the '092 Patent at the time of the invention would have at least a bachelor's degree in electrical engineering, computer engineering, computer science or equivalent coursework, and at least two years of experience in computer architecture, networking or storage. Less work experience may be compensated by a higher level of education, such as a master's degree, and vice versa.

VI. THE STATE OF THE ART AT THE TIME OF INVENTION

37. In the 1980s disk drives were considered to be expensive resources and were generally tied to individual computers. For example, when I interned at Australia Telecom in 1981-82, the removable disk drive was big – about 30 inches – across and disk drive space was considered a valuable resource. When I was a PhD student in the mid-1980s, we time shared a single Vax 11/780 with its own multi-cabinet scale disk drive storage system.

38. Meanwhile the Personal Computer (PC) revolution was occurring with attendant cost reduction in all PC hardware, from CPUs to disk drives. In 1988, David Patterson, Garth Gibson and Randy Kratz, all of the University of California, Berkeley, wrote a seminal paper entitled “A Case for Redundant Arrays of Inexpensive Disks (RAID)”. (Ex. 1014.) In this paper, they proposed replacing “Single Large Expensive Disks (SLED)” with “Redundant Arrays of Inexpensive Disks (RAID) based on magnetic disk technology developed for personal computers”. (*Id.* abstract.)

39. Table 1, reproduced below, justifies that proposal. A Connors CP3100 PC disk costs 2-3x less per MB than the IBM and Fujitsu SLEDs. It has lower throughput (measured as I/Os per second) and worse reliability when it is realized that multiple PC disk drives are needed to replace each SLED. (*Id.* p.110.)

<i>Characteristics</i>	<i>IBM 3380</i>	<i>Fujitsu M2361A</i>	<i>Conners CP3100</i>	<i>3380 v. 2361 v 3100 3100</i>	
				<i>(>1 means 3100 is better)</i>	
Disk diameter (inches)	14	10.5	3.5	4	3
Formatted Data Capacity (MB)	7500	600	100	.01	.2
Price/MB(controller incl.)	\$18-\$10	\$20-\$17	\$10-\$7	1-2.5	1.7-3
MTTF Rated (hours)	30,000	20,000	30,000	1	1.5
MTTF in practice (hours)	100,000	?	?	?	?
No. Actuators	4	1	1	.2	1
Maximum I/O's/second/Actuator	50	40	30	.6	.8
Typical I/O's/second/Actuator	30	24	20	.7	.8
Maximum I/O's/second/box	200	40	30	.2	.8
Typical I/O's/second/box	120	24	20	.2	.8
Transfer Rate (MB/sec)	3	2.5	1	.3	.4
Power/box (W)	6,600	640	10	660	64
Volume (cu. ft.)	24	3.4	.03	800	110

40. They proposed that a traditional disk system be replaced by commodity parts – figure 1, reproduced below.




Figure 1. *Comparison of organizations for typical mainframe and small computer disk interfaces. Single chip SCSI interfaces such as the Adaptec AIC-6250 allow the small computer to use a single chip to be the DMA interface as well as provide an embedded controller for each disk [Adeptec 87] . (The price per megabyte in Table I includes everything in the shaded boxes above.)*

41. Most of the rest of the paper proposes techniques to improve reliability (e.g., via redundancy and error correction) and techniques to improve throughput (via parallelism).

42. In the early 1990's it was realized that storage could be network based rather than node based. This heralded a transition of storage from using dedicated channels between hosts and storage to using network communications between hosts and storage. For example, Randy Katz, still at the University of California,

Berkeley wrote a paper entitled “Network-Attached Storage Systems” in 1992.
(Ex. 1015.)

43. Katz introduces his paper with the statement “[t]oday, we are experiencing a major paradigm shift away from centralized mainframes to a distributed model of computation based on workstations, computation servers, and storage servers connected by high performance networks”. (*Id.* p.68.)

44. Katz describes two channel architectures to connect “storage devices” to “I/O controllers”, including the “low end” SCSI (SCSI stands for “Small Computer System Interface”). “Channels provide the logical and physical pathways between I/O controllers and storage devices.” (*Id.* p.69.) “SCSI is the channel interface most frequently encountered in small form factor (5.25” diameter and smaller) disk drives. SCSI treats peripheral devices in a largely independent fashion. For example, a disk drive is viewed as a linear byte stream.” (*Id.* p.70.)

45. Figure 3, reproduced below, shows the internal organization of an I/O Controller. (*Id.* Figure 3.)




Figure 3: Internal Organization of an I/O Controller

46. Katz describes the controller as the element on the right of this figure, “the controller contains a microprocessor, a modest amount of memory dedicated to buffers and run-time data structures, a ROM to hold the controller firmware, a DME/peripheral bus interfaces, and I/O channel interface.” It’s clear from the discussion above that the I/O channel interface is connected to the disk drives, e.g., via a SCSI channel, and that the connection to the host is via a separate path. (*Id.* p.72.)

47. Thus, a POSITA as of 1999 would be familiar with storage systems using dedicated I/O channels, including SCSI, between the host and storage, as well as distributed network communications between storage servers.

48. Caching was also common in this time frame. Ma and Reddy of Texas A&M University published a paper in 1998 entitled, "An evaluation of storage systems based on network-attached disks" that included a discussion of cached network storage systems. (Ex. 1016.)

49. Figure 2 of the paper, replicated below, shows network-based storage with caches.




Figure 2: Net-disk:Storage system based on Network-attached disks.

50. They describe the operation of this system. "Net-disk and server disk organizations employ caching at different places in the system. A single large cache space in the file manager in server-disk enables efficient utilization of cache space across all the data sets in the system. The smaller caches at each disk in net-


disk organizations can lead to inefficiencies if disks are accessed non-uniformly. Our assumption of system-wide disk striping [10] however reduces this disadvantage.” (*Id.* p. 3.)

VII. THE '092 PATENT

51. U.S. Patent 10,154,092 is entitled “Data sharing using distributed cache in a network of heterogeneous computers” (Ex. 1001.) The inventor is Ilya Gertner. It was filed on Jan. 15, 2016 and granted on Dec. 11, 2018. I’ve been informed that the priority date is Jan. 22, 1999. I shall refer to this patent as the ‘092 patent.

52. The field of the alleged invention is described as “cached storage systems and more particularly to a data storage system that permits independent access from local hosts connected via I/O channels and independent access from remote hosts and remote storage systems connected via network links.” (*Id.*, abstract.) It describes a system built from a “network of PCs” “using off-the-shelf components at reduced cost.” (*Id.*, 1:15-26.)

53. Figure 1 from the patent illustrating “data storage configurations of the preferred embodiment” is reproduced below. (*Id.*, 5:9-10, Fig. 1.)



54. Figure 1 shows hosts 111, 112 attached to a PC storage system 131 via I/O channels 121, 122. SCSI channels can be used to carry I/O traffic. There are also networked attached hosts 106, 107 that can access the PC storage systems 131, 132, 133 using network links 151, 152, 153. The preferred embodiment uses “SCSI channels for I/O traffic and Ethernet link for network traffic.” (*Id.*, 5:35-

47.) An alternate embodiment uses “fiber channel technology for both I/O and traffic and network traffic.” (*Id.*, 5:9-47.)

55. The PC storage systems employ “fast access cache memory.” (*Id.*, 6:1.) “The front-end module 310 handles channel commands and places the results in the cache memory 322 in the form of new data or modification to data stored on the disk 161.” (*Id.*, 5:48-6:8.) The cache also facilitates reads. (*Id.*, 7:5-18.)

A. Claim Construction

56. For the purposes of this proceeding, I have applied the plain and ordinary meaning to each term in the patent claims – except for the means-plus-function terms in claim 24 – as they would be understood by me and by other POSITAs.

57. For the purposes of this proceeding, I have been informed that the following terms in claim 24 are means-plus-function terms where the claim recites a function and the POSITA looks to corresponding structure found in the patent specification that performs the recited function. I have been informed that the Petitioner has identified the following recited functions and corresponding

structures, and I agree that the identified corresponding structures appear to be the '092 Patent's disclosure most related to performance of the recited functions.

- “means for receiving input/output (I/O) traffic from a first host device via a dedicated I/O channel, the I/O traffic comprising a write command”:
 - Recited function: “receiving input/output (I/O) traffic from a first host device via a dedicated I/O channel”
 - Disclosed structure: A front-end module that receives data from hosts via I/O channels, and equivalents. (Ex. 1001, col. 5:47-67; 6:13-24, Fig. 3 label 310.)
- “means for receiving data via a network”:
 - Recited function: “receiving data via a network”
 - Disclosed structure: A network adapter, and equivalents. (Ex. 1001, abstract, col. 3:44-47, 4:1-2.)
- “means for short-term data storage”:
 - Recited function: “short-term data storage”


- Disclosed structure: A cache memory, and equivalents. (Ex. 1001, col. 1:36-40, 6:1-8.)
- “means for long-term data storage”:
 - Recited function: “long-term data storage”
 - Disclosed structure: A back-end storage device or local storage device. (Ex. 1001, 5:9-34, 6:1-4.)
- “means for performing one or more access operations at the means for short-term data storage during processing of the I/O traffic and for performing one or more access operations at the means for long-term data storage based on the I/O traffic”:
 - Recited function: “performing one or more access operations at the means for short-term data storage during processing of the I/O traffic and for performing one or more access operations at the means for long-term data storage based on the I/O traffic”
 - Disclosed structure: A cache manager and a configuration manager, and equivalents. (Ex. 1001, 6:9-32 (configuration manager for operations to local disks), 6:50-8:46 (cache manager for operations to cache), Figs. 4-4F.)

VIII. UNPATENTABILITY ANALYSIS

A. Overview of the Prior Art

1. Heil

58. Heil (Ex. 1006) is a U.S. Patent filed February 11, 1998 and issued January 9, 2001. Heil describes an apparatus and method for data storage utilizing host computers to store files in a cache or disk storage. (Ex. 1006, abstract, 1:9-14, 6:34-47.) Figure 1 of Heil shows one embodiment of the invention:




59. A host system comprising a CPU 100 is connected by a PCI bus 116.5 to a host bus adapter (HBA) 117. (Ex. 1006, 6:34-47.) Host bus adapter 117 includes a cache memory 117.2 and local memory 116 and is connected to local

drives 118 by SCSI or FCAL interfaces 117.8 and 117.9. (Ex. 1006, 7:9-15, 7:51-56.) The host processes I/O requests locally by accessing the host bus adapter 117 through PCI bus 116.5. (Ex. 1006, 4:51-57, 6:34-47.)

60. A first host and HBA node (referred to as node 1) can access other nodes (referred to as node n) comprising a host system and an HBA through network connection 121. (Ex. 1006, 8:8-22.) The host and HBA architecture of node 1 is shown as identical to that of node n. (Ex. 1006, 8:61-9:10.)

2. Nakayama

61. Nakayama (Ex. 1007) is a U.S. Patent filed June 2, 1997 and issued July 6, 1999. Nakayama describes an apparatus and method for a cached data storage system using host computers to provide input to a cache memory or a storage drive. (Ex. 1007, abstract, 4:52-65.) Nakayama Figure 3 shows an embodiment of the invention having a host computer 300 connected by channel 301 to a disk control 302 having a cache memory 309. (Ex. 1007, 4:40-5:9.) Host computer 303 can also be connected to channel control 302 by a SCSI bus 304. (Ex. 1007, 4:56-60.)



62. Disk control 302 is connected to drives 315 and 316 through a drive interface 313. (Ex. 1007, 4:56-60.) In addition, Nakayama Figure 3 shows that channel control 301 or SCSI bus 304 is a distinct communication path from drive interfaces 313 and 314.

3. Gulick

63. Gulick (Ex. 1008) is a U.S. Patent filed September 11, 1995 and issued November 25, 1997. Gulick describes connecting PCI buses to SCSI adapters. (Ex.

1008, 1:17-25.) For example, Gulick describes a “host/PCI/cache bridge [that] interfaces to a peripheral expansion bus, preferably a peripheral component interconnect (PCI) bus.” (Ex. 1008, 5:37-51.) A SCSI adapter is used to couple the PCI bus with various SCSI devices, such as a CD-ROM drive and a tape drive. (Ex. 1008, 5:37-51.)

4. Berman

64. Berman (Ex. 1009) is a U.S. Patent filed August 7, 1997 and issued September 12, 2000. Berman describes an improved networking system using “Fibre Channels”. (Ex. 1009, abstract, 1:34-47.) In particular, Berman suggests that it would be advantageous to interface Fibre Channels with “legacy networks” such as ethernet and discloses systems such as routers that connect Fibre Channel to ethernet. (Ex. 1009, 12:9-11, claims 13-15, 3:23-30.)

B. Grounds 1 and 2: Anticipation of Claims 1-3, 7-12, and 19-23 by Heil

65. In my opinion, when viewed and considered considering the knowledge possessed by one of ordinary skill in the art, Heil discloses all the limitations of, and anticipates, claims 1-3, 7-12, and 19-23 of the '092 Patent.

1. Claim 1

66. In my opinion, Heil discloses and anticipates all of the limitations of claim 1.

a) 1.pre: “An apparatus comprising”

67. The preamble of claim 1 recites “an apparatus.” In my opinion, Heil discloses the preamble. Heil explicitly states that: “*This invention relates to* distribution of storage data within computer nodes of a clustered computing environment, and more particularly, to *an apparatus* and method for I/O request shipping operable using peer-to-peer communications within host bus adapters of the clustered systems.” (Ex. 1006, 1:9-14.)

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


68. Claim 1 additionally requires “a first interface configured to receive input/output (I/O) traffic from a first host device via a dedicated I/O channel, the I/O traffic comprising a read command.” Heil discloses this limitation.

69. Specifically, Heil describes a first interface that receives I/O traffic from a host device in the form of an interface for a PCI bus bridge 115, PCI bus 116.5, or Front End I/F 102 for receiving input from a host system (CPU 100 of Node 1):

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

(Ex. 1006, 6:35-40.)

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



(Ex. 1006, Fig. 1 (annotated).)

71. In addition, Heil discloses that communications between the Node 1 host and HBA 117 comprise I/O requests. (Ex. 1006, 7:36-37 (“As shown in FIG. 1, one HBA services the I/O requests from CPU 1 and CPU 2 100.”).) Heil describes the I/O requests as instructions to retrieve files from caches or local storage devices. (Ex. 1006, 2:59-65 (“To retrieve files from a remote resource, software in the host intercepts file or block level I/O requests and determines whether the particular file or block is stored locally or remotely. If local, the software passes the request down to the local file system or block I/O driver.”), 4:63-69 (“Data returned from the HBA performing the remote processing (i.e., a read request) is passed by the initiating HBA to the requesting host.”), 12:4-6 (“The remote I/O blocks are sent from the HBA performing the remote processing (e.g., a read request) to the initiating HBA 460.”).) Based on these disclosures, my opinion is that a POSITA would understand Heil’s I/O requests to include read commands.

72. Finally, it is my opinion that a POSITA would understand that the channel used by the PCI bus bridge 115 and PCI bus 116.5 is a dedicated I/O channel because Heil discloses that I/O traffic between HBA 117 and the Node 1 host system utilizes this communication path and it is a direct, dedicated

connection between the host and the HBA that carries I/O traffic between them.

(Ex. 1006, 6:35-40, 7:36-37.)

73. Heil therefore anticipates element 1.a.

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

74. Claim 1 additionally requires “a second interface configured to receive first data via a network.” Heil discloses this limitation.


75. Specifically, Heil describes a second interface that receives first data via a network in the form of an interface between “fibre channel chip 120” and “network fibre channel backbone 121” that receives data from the fibre network (i.e., fibre channel backbone 121):

Fibre Channel Chip 120 connects Node 1 150 to the rest of the clustered environment via the network Fibre Channel backbone 121. The Fibre Channel backbone is a high-speed communication medium and provides numerous advantages. In particular, there are no complex routing algorithms to compute when sending data from a first Fibre Channel port to a second Fibre Channel port. Fibre Channel makes use of unique addresses to connect to other processors or peripherals. The Fibre Channel port only needs to manage the point-to-point connection between the

device and the backbone fabric. There are three classes of service available through the Fibre Channel backbone: dedicated service, switched service without a dedicated connection but providing guaranteed delivery with an acknowledgment of receipt, and connectionless service without acknowledgment of receipt.

(Ex. 1006, 8:24-38; 4:26-35 (“Similarly, HBAs in different nodes can communicate as peers via, for example, a Fibre Channel backbone that interconnects the HBAs.”).)

76. Heil’s Figure 1 shows, in addition to the interface at the PCI bus 116.5 or PCI bus bridge 115 (a first interface), that HBA 117 has an interface at fibre channel chip 120 (a second interface) that connects to fibre network 121:



(Ex. 1006, Fig. 1 (annotated).)


77. The data transmitted over the fibre channel backbone and received by fibre channel chip 120 would correspond to the claimed first data.

78. Heil therefore anticipates element 1.b.

d) 1.c: “a cache memory configured to store second data”

79. Claim 1 additionally requires “a cache memory configured to store second data.” Heil discloses this limitation.

80. Specifically, cache 117.2 as shown in Heil’s Figure 1 is a cache memory that stores second data:



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


81. The data cached within cache memory would correspond to the claimed second data.

82. Heil therefore anticipates element 1.c.

e) 1.d: “a storage device configured to store third data”

83. Claim 1 additionally requires “a storage device configured to store third data.” Heil discloses this limitation.

84. Specifically, local drives 118 and 123 as shown by Heil’s Fig. 1 are storage devices configured to store third data:



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


85. The data stored within the local drives would correspond to the claimed third data.

86. Heil therefore anticipates element 1.d.

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

87. Claim 1 additionally requires “a processor coupled to the cache memory, the processor coupled to the storage device via a communication path that is distinct from the dedicated I/O channel, the processor configured to access the cache memory during processing of the I/O traffic.” Heil discloses this limitation.


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



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

89. Heil’s Figure 1 shows that the communication path from the processor to the storage device—i.e., from the CPU 117.1 to the local drives 118 via a SCSI

interface, annotated below in **red**—is distinct from the dedicated I/O channel—
i.e., from the Node 1 host system to HBA 117 via a PCI bus, annotated below in **green**:



(Ex. 1006, Fig. 1 (annotated).)

90. Additionally, Heil states that, to improve performance, the cache memory stores data that is frequently responsive to I/O requests. (Ex. 1006, 5:25-28 (“Performance is further enhanced in that each intelligent HBA can provide disk caching of the local disks under its control, and can service hits out of this

cache regardless of whether the requesting host is local or remote.”), 7:24-28 (“To improve I/O performance, CPU 117.1 temporarily caches disk data in cache memory 117.2 that is frequently accessed.”).) In my opinion, a POSITA would understand from these disclosures that the cache memory would be accessed to retrieve data responsive to the I/O request. Therefore, the processor (CPU 117.1 of the HBA) may access cache memory while processing I/O traffic, including after processing an I/O command transmitted in the I/O traffic.

91. My opinion is further supported by Figure 3 of Heil. Figure 3 shows that, when an I/O request is made by a CPU 100 (Ex. 1006, 7:36-40), the HBA determines whether the request may be serviced locally. (Ex. 1006, 11:47-54.) While Fig. 3 suggests responsive data is retrieved from the local disks, Heil discloses that “frequently accessed” data is cached in cache memory and may be provided in response to “improve I/O performance.” (Ex. 1006, 7:21-34, 12:1-4.) This again confirms that the processor (CPU 117.1 of the HBA) may access cache memory while processing I/O traffic.

92. Heil therefore anticipates element 1.e.

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

93. Finally, claim 1 requires “the processor further configured to perform an access operation at the storage device based on the I/O traffic.” Heil discloses this limitation.

94. Heil discloses that the HBA having a processor CPU 117.1 accesses the local drives 118 based on the I/O traffic. For example, in response to an I/O request by CPU 100, the HBA may determine that the request may be serviced locally and retrieve the response data from the local disks. (Ex. 1006, 7:36-37 (“As shown in FIG. 1, one HBA services the I/O requests from CPU 1 and CPU 2 100.”), 11:45-53 (“After receiving a block I/O request 400, the 110 redirector has the means to search the directory and determining means to locate the local or remote disk drives that are storing the I/O requested blocks 410. If the location of the blocks is local 420, then the blocks are retrieved from the respective local disk drive.”).)

95. Heil therefore anticipates element 1.f.

2. Claim 2

96. In my opinion, Heil discloses and anticipates all of the limitations required by claim 2.

97. As described above in Section VIII.B.1, in my opinion Heil anticipates claim 1.

98. Claim 2 depends from claim 1 and further requires that “the dedicated I/O channel is associated with dedicated throughput that corresponds to the I/O.” Heil discloses this limitation.

99. Heil discloses that an HBA “in general is a device which adapts (connects) a host computer system to an I/O device. Signals associated with the bus of the host computer system (e.g., PCI, ISA, etc.) are adapted for exchange with a bus specific to the I/O device (e.g., SCSI, Fibre Channel, LAN, etc.).” (Ex. 1006, 4:3-8.) In addition, HBA 117 receives a request from a host via the dedicated I/O channel. (Ex. 1006, 4:50-51 (“In the preferred embodiment, an HBA receives a block I/O request from the host and searches through the directory in HBA memory to determine whether the block can be found locally or remotely.”).) Based on these disclosures, a POSITA would understand that the dedicated I/O channel is associated with dedicated throughput that corresponds to the I/O traffic

because the data passing through the I/O channel corresponds to the I/O traffic.

Heil therefore discloses claim 2.

3. Claim 3


100. In my opinion, Heil discloses and anticipates the limitations required by claim 3.

101. As described above in Section VIII.B.1, in my opinion Heil anticipates claim 1.

102. Claim 3 depends from claim 1 and further requires that “the I/O traffic is distinct from the first data.” Heil discloses this limitation.

103. According to claim 1, the first data is data received by the second interface via the network. In my opinion, Heil discloses that the I/O traffic from its host server to its HBA is distinct from the data received by its HBA from the network.

104. This is shown by the annotations to Heil’s Figure 1 below:



(Ex. 1006, Fig. 1 (annotated).)

105. **The path for I/O traffic from Heil's host server (including CPU 100) to its HBA is annotated in red. The network interface comprising fibre channel chip 102 receiving first data from the fibre channel network is annotated in green.**

106. Heil discloses that the I/O traffic is an I/O request from the host server to the HBA. (Ex. 1006, 7:36-50, 9:11-32.) By contrast, the fibre channel network traffic comprises communications between HBAs that may be responsive to an I/O

request. (Ex. 1006, 8:8-38, 9:49-67.) It is my opinion that a POSITA would understand that these are different.

107. Furthermore, Figure 1 shows that the I/O channel is transmitted through a different communication path than the fibre channel network traffic. Therefore, it is my opinion that a POSITA would understand that the I/O traffic in Heil is distinct from the network traffic received from the fibre channel. Thus, Heil anticipates claim 3.

4. Claim 7

108. In my opinion, Heil discloses and anticipates all of the limitations required by claim 7.

109. As described above in Section VIII.B.1, in my opinion Heil anticipates claim 1.

110. Claim 7 depends from claim 1 and further requires that “the first data comprises a second read request, a write request, or a combination thereof.” Heil discloses this limitation.

111. Heil discloses that, for each I/O request, if the requested data is not found at the local HBA, then the local HBA transmits the I/O request via the fibre channel network to remote HBAs to service the request. (Ex. 1006, 11:45-65

(“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. 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.”.) An I/O request transmitted between two HBAs is transmitted via the fibre channel network and is distinct from an I/O request transmitted from a host system to a local HBA via the PCI bus. Therefore, a POSITA would understand that an I/O request received by an HBA from another HBA via the fibre channel network is first data comprising a second read request received by the second interface via the network.

112. In addition, Heil discloses that an HBA “can cache blocks of data retrieved from its local disk, to service future local or remotely generated requests for these blocks.” (Ex. 1006, 12:1-4.) Based on this disclosure, a POSITA would understand that more than one I/O request may be issued by a host system—Heil explicitly contemplates that the host system may issue a second I/O request to the local HBA even after the host system issues a first I/O request to the local HBA. A

POSITA would understand that the second I/O request corresponds to the claimed second read request. If data responsive to the second I/O request is not found at the local HBA, and the local HBA ships that request to a remote HBA, then the remote HBA is the apparatus that has received the second read request via the network.

113. Thus, Heil anticipates claim 7.

5. Claim 8

114. In my opinion, Heil discloses and anticipates all of the limitations required by claim 8.

115. As described above in Section VIII.B.1, in my opinion Heil anticipates claim 1.

116. Claim 8 depends from claim 1 and further requires that “the network comprises a plurality of interconnected computing devices.” Heil discloses this limitation.

117. Specifically, Heil describes a network of multiple nodes interconnected via a fibre channel:

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

node inter-processor communication can occur on this Fibre Channel backbone 121. Alternatively, the Fibre Channel backbone 121 can be used exclusively for storage communications and inter-processor communications can occur over a separate network. One skilled in the art will recognize other connectivity standards besides Fibre Channel can be used.

(Ex. 1006, 8:12-24.)

118. Heil further discloses that each node comprises a host system having a CPU 100 connected to a host bus adapter (HBA) via a PCI bus bridge: “Node 1 150 includes a host system connected to a HBA 117 via a Host-to-Peripheral Component Interconnect (hereinafter referred to as “PCI”) bus bridge 115 and PCI bus 116.5. The host system includes one or more Central Processing Units (hereinafter referred to as “CPU”).” (Ex. 1006, 6:35-44.) “HBA 117 also includes a CPU 117.1, local memory 116, and cache memory 117.2, which are connected by internal processor bus 117.3.” (*Id.*, 7:9-12.)

119. Heil describes the host system and the HBA as computers or devices: “An HBA in general is a device which adapts (connects) a host computer system to an I/O device.” (*Id.*, 4:3-5.) In my opinion, a POSITA would therefore understand from these disclosures that both the host servers having CPUs 100 and

the HBAs having a CPU and memory are computing devices, and they are interconnected via a fibre channel network. Heil therefore anticipates claim 8.


6. Claim 9

120. In my opinion, Heil discloses and anticipates all of the limitations required by claim 9.

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

121. First, claim 9 requires “the storage device and the dedicated I/O channel are independently accessible.” Heil discloses this limitation.

122. Specifically, Heil’s Fig. 1 shows two independent pathways (1) from the Node 1 host via the PCI bus to the processor (the dedicated I/O channel) and (2) from a remote Node N host via the fiber network and fibre channel chip 120 to the local drives via a SCSI or FCAL interface:



(Ex. 1006, Fig. 1 (annotated).)

123. Heil further discloses that the storage device may be accessed based on a network request from a requestor HBA on the network, rather than the local HBA. (Ex. 1006, 11:54-12:6.) Because requestor HBAs on the network use the fibre channel pathway rather than the PCI bus pathway, a POSITA would understand that access operations by a requestor HBA on the network to the storage device is independent of I/O operations from the local HBA to the storage device.

124. Therefore, Heil discloses the storage device and dedicated I/O channel are independently accessible.

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

125. Second, claim 9 requires “the processor is configured to read at least a portion of the second data from the cache memory based on the read command.”

Heil also discloses this limitation.

126. Specifically, Heil discloses CPU 117.1 of the HBA reads data from the cache based on the I/O request. (Ex. 1006, 5:25-33 (“Performance is further enhanced in that each intelligent HBA can provide disk caching of the local disks under its control, and can service hits out of this cache regardless of whether the requesting host is local or remote.”), 7:24-30 (“To improve I/O performance, CPU 117.1 temporarily caches disk data in cache memory 117.2 that is frequently accessed.”), 11:66-12.1 (“The remote HBA retrieves the requested blocks of data from its HBA cache, or if the blocks of data are not within cache, then directly from the remote HBA's local disk.”).)

127. Therefore, Heil discloses element 9.b.

7. Claim 10

128. In my opinion, Heil discloses and anticipates all of the limitations required by claim 10.

129. As described above in Section VIII.B.1, in my opinion Heil anticipates claim 1.

130. Claim 10 depends from claim 1 and further requires that “the processor is further configured to route the read command to the cache memory or to the storage device.” Heil discloses this limitation.

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

or the storage device, depending on where the requested data may be found, because reading the cache or disk necessarily requires the cache or disk to receive the read command. Heil therefore discloses claim 10.

8. Claim 11


132. In my opinion, Heil discloses and anticipates all of the limitations required by claim 11.

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

133. The first element of claim 11 requires “configuration manager circuitry configured to route an I/O request included in the I/O traffic to the cache memory, to route the I/O request to the storage device, or to deny the I/O request.” Heil discloses this limitation.

134. Specifically, Heil discloses that an I/O redirector 240 as shown in Figure 2 facilitates retrieval of data related to an I/O request based on the mapping

of the data within every storage subsystem of the server clusters having HBAs:



(Ex. 1006, Fig. 2, 12:19-34, 5:1-6.)

135. Additionally, Heil discloses that caching enhances I/O performance and caches are part of the storage subsystem of the HBAs. (Ex. 1006, 5:25-33 (“Performance is further enhanced in that each intelligent HBA can provide disk caching of the local disks under its control, and can service hits out of this cache regardless of whether the requesting host is local or remote.”), 7:24-30 (“To

improve I/O performance, CPU 117.1 temporarily caches disk data in cache memory 117.2 that is frequently accessed.”.)

136. Heil discloses that the HBA looks to its cache when servicing an I/O request, and, if the requested data is not found in the HBA’s cache, the HBA then accesses the storage device. (Ex. 1006, 11:66-12:1 (“The remote HBA retrieves the requested blocks of data from its HBA cache, or if the blocks of data are not within cache, then directly from the remote HBA's local disk.”).)

137. A POSITA would understand that this operation routes the I/O request to cache memory or to the storage device, depending on where the requested data is found. Therefore, the circuitry of Heil’s I/O redirector corresponds to the claimed configuration manager circuitry and Heil discloses element 11.a.

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

138. The second element of claim 11 requires “front-end circuitry configured to process the I/O request.” Heil discloses this limitation.

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

The HBA interface to the local drives 117.8 and 117.9 can be SCSI or FCAL depending upon the bus supporting the local drives.”.) Therefore, the front-end interface 171 is “front-end circuitry configured to process the I/O request.”

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

140. The third element of claim 11 requires “back-end circuitry configured to perform a read operation or a write operation at the storage device based on the I/O request.” Heil discloses this limitation.

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

9. Claim 12

142. In my opinion, Heil discloses and anticipates all of the limitations required by claim 12.

143. As described above in Section VIII.B.1, in my opinion Heil anticipates claim 1.

144. Claim 12 depends from claim 1 and further requires that “the host device comprises a server.”

145. Heil explicitly describes its invention as retrieving “across independent computer nodes of a server cluster by providing for I/O shipping of block level requests to peer intelligent host-bus adapters (hereinafter referred to as HBA).” (Ex. 1006, abstract.) Heil further describes its host as a computer system connected to an HBA: “An HBA in general is a device which adapts (connects) a host computer system to an I/O device.” (*Id.*, 4:3-5.) In addition, the host system and HBA together comprise the node of a fibre channel network. (Ex. 1006, 6:35-44, 8:12-24.)

146. A POSITA would therefore understand that Heil’s host system is a server of a server cluster which facilitates retrieval across the cluster by providing I/O requests to and are serviced by HBAs. Heil therefore anticipates claim 12.

10. Claim 19

147. In my opinion, Heil discloses and anticipates all of the limitations of claim 19.

a) 19.pre: “An apparatus comprising”

148. In my opinion, Heil anticipates element 19.pre for the same reasons it anticipates element 1.pre as discussed in Section VIII.B.1.

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

149. In my opinion, Heil anticipates element 19.a for the same reasons it anticipates element 1.a. As discussed in Section VIII.B.1, Heil discloses “a first interface configured to receive input/output (I/O) traffic from a first host device via a dedicated I/O channel, the I/O traffic comprising a read command” and therefore discloses “the I/O traffic comprising one or more read commands, one or more write commands, or a combination thereof.”

c) 19.b: “a second interface configured to receive data via a network”

150. In my opinion, Heil anticipates element 19.b for the same reasons it anticipates element 1.b as discussed in Section VIII.B.1.

d) 19.c: “a cache memory”

151. In my opinion, Heil anticipates element 19.c for the same reasons it anticipates element 1.c as discussed in Section VIII.B.1.

e) 19.d: “a storage device”

152. In my opinion, Heil anticipates element 19.d for the same reasons it anticipates element 1.d as discussed in Section VIII.B.1.

f) 19.e and 19.f: “a processor coupled via a communication path to the storage device, the processor configured to access the cache memory during processing of the I/O traffic, the processor further configured to perform one or more access operations at the storage device based on the I/O traffic, wherein the communication path is distinct from the dedicated”

153. In my opinion, Heil anticipates elements 19.e and 19.f for the same reasons it anticipates elements 1.e and 1.f as discussed in Section VIII.B.1.

11. Claim 20

154. In my opinion, Heil anticipates claim 20 for the same reasons it anticipates claim 2. As discussed in Section VIII.B.2 above, Heil’s I/O requests are “dedicated throughput that corresponds to the I/O traffic” and “associated” with the I/O channel.

12. Claim 21

155. In my opinion, Heil anticipates claim 21 for the same reasons it anticipates claim 3. As discussed in Section VIII.B.3 above, Heil's I/O traffic is distinct from the data received by the second interface via the network.

13. Claim 22

156. In my opinion, Heil anticipates claim 22 for the same reasons it anticipates claim 7. As discussed in Section VIII.B.4 above, the data received by the second interface via the network may correspond to one or more read requests transmitted by another HBA on Heil's fibre channel network.

14. Claim 23

157. In my opinion, Heil anticipates claim 23 for the same reasons it anticipates claim 9. As discussed in Section VIII.B.6 above, Heil's disk storage (*i.e.*, the storage device) and the dedicated I/O channel are independently accessible.

C. Ground 3: Obviousness of Claims 1-3 and 6-24 Over Heil and Nakayama

158. In my opinion, when viewed and considered considering the knowledge possessed by one of ordinary skill in the art, Heil in view of Nakayama

discloses all the limitations of, and renders obvious, claims 1-3 and 6-24 of the '092 Patent.

159. As described in detail below, one of ordinary skill in the art would have been motivated to combine and collectively consider the teachings of these references to arrive at the subject matter claimed by the '092 Patent and would have had a reasonable expectation of success when doing so.

1. Motivation to Combine

160. It would have been obvious for a POSITA to combine Heil and Nakayama, with a reasonable chance of success in doing so. Heil and Nakayama are in the same field of invention and disclose similar cache and storage architectures. Heil “relates to distribution of storage data within computer nodes.” (Ex. 1006, 1:8-9.) Nakayama relates to a “storage control system” “for storing input/output data to/from host computers to be shared among said host computers.” (Ex. 1007, 1:5-10.)

161. Both Heil and Nakayama recognize the importance of cache storage in addition to disk storage, and both disclose methods of reading and writing to cache and disk storage from a processor. Heil explains that “[p]erformance is enhanced in that each intelligent HBA can provide disk caching of the local disks under its

control.” (Ex. 1006, 5:25-27.) Nakayama likewise states that 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.” (Ex. 1007, 5:5-9.) A POSITA would understand that this temporary storage in a disk cache allows data to be accessed more quickly.

162. Heil describes how “a large number of processors may be added to a computing environment all of which share access to common storage devices.” (Ex. 1006, 1:39-42.) Furthermore, Heil discusses “file level access” and “block level access.” (*Id.*, 1:50-52.) Nakayama discloses that “the data access command is a write command or a read command.” (Ex. 1007, 6:18-19), and then describes how write commands and read commands are executed. (*Id.*, 2:36-61, 6:20-7:10.) Thus, a POSITA would implement Heil in view of Nakayama by implementing the details of the read and write commands of the latter on the former. Note that a practitioner of Heil would be operating with a single data format and thus would skip the file format conversion steps of Nakayama. A POSITA seeking to improve storage access for a host computer in a storage system having both cache and disk storage would therefore look to Nakayama to do so, and would use read and write operations such as those disclosed by Nakayama to access cache and disk storage

in Heil. A practitioner would implement this combination with simple program modifications to the operating system within a POSITA's skill level and would have reasonably expected success.

163. Read and write operations such as those disclosed by Nakayama are applicable to similar systems having both cache and disk storage and yield predictable results.

2. Claim 1

164. In my opinion, Heil in view of Nakayama discloses and renders obvious all of the limitations of claim 1.

a) 1.pre: “An apparatus comprising”

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

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

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


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

165. In my opinion, Heil and Nakayama renders elements 1.pre, 1.b, 1.c, 1.d, and 1.e obvious for the same reasons Heil anticipates these elements as discussed in Section VIII.B.1 above.

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

166. In my opinion, Heil discloses element 1.a for the reasons discussed in Section VIII.B.1 above.

167. As I discussed in Section VIII.A, Nakayama is directed to a data storage architecture similar to Heil's in which a host computer 300 is connected by a channel 301 (which may be a SCSI bus control, as seen at 304) to a disk control 302 having a cache memory 309. (Ex. 1007, 4:40-5:9, 4:56-60.) This is shown by Fig. 3 of Nakayama:



(Ex. 1007, Fig. 3.)

168. Heil characterizes the communications between its host system and its HBA as “I/O requests” and describes responding to such I/O requests by searching through data in the cache memory and disk drives of HBAs local and remote to the host system issuing the I/O request. (Ex. 1006, abstract, 4:29-67, 7:36-50, 9:11-32, 11:45-12:7.)

169. In my opinion, a POSITA would understand that an “I/O request” includes both read and write commands. Nakayama confirms my opinion and explicitly states that its host computers make access requests in the form of read and write requests to storage including the cache memory and disk drive. (Ex. 1007, 3:66-4:25 (“An access request 204 (write), 205 (read), 206 (write) and 207 (read) which are generated by the host computers A200 and B201 to the storage media 203 are issued by the storage control 210. These access requests are executed by the data access unit 202.”), 4:48-51, 4:61-65, Fig. 6.)

170. To the extent Heil does not use the word “read command” or “write command” to describe its “I/O requests,” it is my opinion that a POSITA would understand from Nakayama’s explicit teaching of access requests in the form of read and write commands that I/O traffic may comprise a read command. In addition, as I discussed in Section VIII.C.1, a POSITA would use Nakayama’s

read and write commands with Heil's storage system. Therefore, Heil in view of Nakayama renders obvious "I/O traffic comprising a read command" as in element 1.a.

c) 1.f: "the processor further configured to perform an access operation at the storage device based on the I/O traffic"

171. In my opinion, Heil discloses element 1.f for the reasons discussed in Section VIII.B.1 above.

172. In my opinion, Nakayama also discloses this limitation. Nakayama discloses that its host computers provide access requests such as read and write commands to storage, including the cache memory and disk drive. (Ex. 1007, 3:66-4:25 ("An access request 204 (write), 205 (read), 206 (write) and 207 (read) which are generated by the host computers A200 and B201 to the storage media 203 are issued by the storage control 210. These access requests are executed by the data access unit 202."), 4:48-51 ("The disk control 302 performs read/write of data on the drives 315 and 316 in response to the requests of the host computers 300 and 303"), 4:61-65 ("The control processors 305 and 306 mainly perform data transfer between the host computers 300 and 303 and the cache memory 309").) Nakayama further discloses reading or writing to the storage device based on

whether a read command or a write command is issued. (Ex. 1007, 6:20-54.) In my opinion, reading and writing to a cache memory or a disk drive based on a read command or a write command provided by a host computer is “the processor configured to access the cache memory during processing of the I/O traffic” and “the processor further configured to perform an access operation at the storage device based on the I/O traffic” as claimed.

173. As I discussed in Section VIII.C.1, a POSITA would use Nakayama’s read and write commands with Heil’s storage system. Because both Heil and Nakayama separately disclose this limitation, Heil in view of Nakayama renders obvious element 1.f.

3. Claim 2

174. Claim 2 depends from claim 1 and further requires “the dedicated I/O channel is associated with dedicated throughput that corresponds to the I/O traffic.” As described above in Section VIII.C.2, in my opinion Heil in view of Nakayama renders obvious claim 1.

175. In my opinion, Heil and Nakayama renders claim 2 obvious for the same reasons Heil anticipates claim 2 as discussed in Section VIII.B.2 above.

4. Claim 3

176. Claim 3 depends from claim 1 and further requires “the I/O traffic is distinct from the first data.” As described above in Section VIII.C.2, in my opinion Heil in view of Nakayama renders obvious claim 1.

177. In my opinion, Heil and Nakayama renders claim 3 obvious for the same reasons Heil anticipates claim 3 as discussed in Section VIII.B.3 above.

5. Claim 6

178. Claim 6 depends from claim 1 and further requires “the I/O traffic further comprises a write command.”

179. In my opinion, Heil in view of Nakayama discloses and renders obvious all of the limitations required by claim 6.

180. As described above in Section VIII.C.2, in my opinion Heil in view of Nakayama renders obvious claim 1.

181. Heil characterizes the communications between its host system and its HBA as “I/O requests” and describes responding to such I/O requests by searching through data in the cache memory and disk drives of HBAs local and remote to the host system issuing the I/O request. (Ex. 1006, abstract, 4:29-67, 7:36-50, 9:11-32, 11:45-12:7.)

182. In my opinion, a POSITA would understand that an “I/O request” includes both read and write commands.

183. Nakayama confirms my opinion and explicitly states that its host computers make access requests in the form of read and write requests to storage including the cache memory and disk drive. (Ex. 1007, 3:66-4:25 (“An access request 204 (write), 205 (read), 206 (write) and 207 (read) which are generated by the host computers A200 and B201 to the storage media 203 are issued by the storage control 210. These access requests are executed by the data access unit 202.”), 4:48-51, 4:61-65.)

184. To the extent Heil does not use the word “write command” as in claim 6 or “read command” as in claim 1 to describe its “I/O requests,” it is my opinion that a POSITA would understand from Nakayama’s explicit teaching of access requests in the form of read and write commands that I/O traffic may comprise a write command. In addition, as I discussed in Section VIII.C.1, a POSITA would use Nakayama’s read and write commands with Heil’s storage system. Therefore, Heil in view of Nakayama discloses “I/O traffic further comprises a write command” as in claim 6.

6. Claim 7

185. Claim 7 depends from claim 1 and further requires “the first data comprises a second read request, a write request, or a combination thereof.” As described above in Section VIII.C.2, in my opinion Heil in view of Nakayama renders obvious claim 1.

186. In my opinion, Heil and Nakayama renders claim 7 obvious for the same reasons Heil anticipates claim 7 as discussed in Section VIII.B.4 above.

7. Claim 8

187. Claim 8 depends from claim 1 and further requires “the network comprises a plurality of interconnected computing devices.” As described above in Section VIII.C.2, in my opinion Heil in view of Nakayama renders obvious claim 1.

188. In my opinion, Heil and Nakayama renders claim 8 obvious for the same reasons Heil anticipates claim 8 as discussed in Section VIII.B.5 above.

8. Claim 9

189. Claim 9 depends from claim 1 and further requires “the storage device and the dedicated I/O channel are independently accessible, and wherein the processor is configured to read at least a portion of the second data from the cache

memory based on the read command.” As described above in Section VIII.C.2, in my opinion Heil in view of Nakayama renders obvious claim 1.

190. In my opinion, Heil and Nakayama renders claim 9 obvious for the same reasons Heil anticipates claim 9 as discussed in Section VIII.B.6 above.

191. In addition, to the extent Heil does not use the words “read command” or “write command,” Nakayama explicitly discloses read and write commands. In my opinion, a POSITA would apply Nakayama’s teaching of read and write commands to Heil.

192. I base my opinion on the fact that Heil characterizes the communications between its host system and its HBA as “I/O requests” and describes responding to such I/O requests by searching through data in the cache memory and disk drives of HBAs local and remote to the host system issuing the I/O request. (Ex. 1006, abstract, 4:29-67, 7:36-50, 9:11-32, 11:45-12:7.) A POSITA would understand that an “I/O request” includes both read and write commands. Nakayama confirms my opinion and explicitly states that its host computers make access requests in the form of read and write requests to storage including the cache memory and disk drive. (Ex. 1007, 3:66-4:25 (“An access request 204 (write), 205 (read), 206 (write) and 207 (read) which are generated by

the host computers A200 and B201 to the storage media 203 are issued by the storage control 210. These access requests are executed by the data access unit 202.”), 4:48-51, 4:61-65.)

193. In my opinion, a POSITA would understand from Nakayama’s explicit teaching of access requests in the form of read and write commands that I/O traffic may comprise a read command. In addition, as I discussed in Section VIII.C.1, a POSITA would use Nakayama’s read and write commands with Heil’s storage system. Therefore, Heil in view of Nakayama renders obvious claim 9.

9. Claim 10

194. Claim 10 depends from claim 1 and further requires “the processor is further configured to route the read command to the cache memory or to the storage device.” As described above in Section VIII.C.2, in my opinion Heil in view of Nakayama renders obvious claim 1.

195. In my opinion, Heil and Nakayama renders claim 10 obvious for the same reasons Heil anticipates claim 10 as discussed in Section VIII.B.7 above.

196. Furthermore, in my opinion Nakayama also discloses claim 10. Nakayama discloses that if the received command is a read command, the processor determines whether the read data exists in the cache memory. (Ex. 1007,

6:39-7:2.) If the data is found within the cache, the data is read from the cache.

(*Id.*) Otherwise, the processor reads the data from the drives. (*Id.*) In my opinion, a POSITA would understand that this process routes the read command to the cache or the storage device, depending on where the requested data may be found, because reading the cache or disk necessarily requires the cache or disk to receive the read command.

197. As I discussed in Section VIII.C.1, a POSITA would use Nakayama's read and write commands with Heil's storage system. Therefore, Heil in view of Nakayama renders obvious claim 10.

10. Claim 11

198. Claim 11 depends from claim 1 and further requires "configuration manager circuitry configured to route an I/O request included in the I/O traffic to the cache memory, to route the I/O request to the storage device, or to deny the I/O request; front-end circuitry configured to process the I/O request; and back-end circuitry configured to perform a read operation or a write operation at the storage device based on the I/O request." As described above in Section VIII.C.2, in my opinion Heil in view of Nakayama renders obvious claim 1.

199. In my opinion, Heil and Nakayama renders claim 11 obvious for the same reasons Heil anticipates claim 11 as discussed in Section VIII.B.8 above.

200. As discussed in Section VIII.B.8, I cited to col. 11 ln. 66 to col. 12 ln. 7 of Heil as disclosing routing the I/O request to cache memory or storage. This section describes a process for responding to an I/O request that retrieves data from the cache if the requested data is found within the cache, or from the disk drive if the requested data is not found within the cache, for a remote HBA, which a POSITA would also find obvious to perform for a local HBA. (*See also* Ex. 1006, 5:25-28.) Heil explicitly teaches caches for both local and remote HBAs as shown in Fig. 1, and further explicitly teaches that caches should store frequently requested data “to improve I/O performance.” (Ex. 1006, 11:27-34.) Accordingly, a POSITA in view of these disclosures would understand that the process of retrieving data responsive to an I/O request to a remote HBA is equally applicable to retrieving data responsive to an I/O request to a local HBA.

201. As I discussed in Section VIII.B.1 above, Heil discloses “perform[ing] a read operation or a write operation at the storage device based on the I/O request” because it accesses disk storage based on I/O requests from the host system. A POSITA would understand that Heil’s accessing of disk storage to be a

read operation even if Heil does not explicitly use the words “read operation” or “write operation” to describe the access. Furthermore, as I explain below, Nakayama explicitly discloses read and write operations. In my opinion, as discussed in Section VIII.C.1, a POSITA would apply Nakayama’s teaching of read and write operations to Heil.

202. I base my opinion on the fact that Heil characterizes the communications between its host system and its HBA as “I/O requests” and describes responding to such I/O requests by searching through data in the cache memory and disk drives of HBAs local and remote to the host system issuing the I/O request. (Ex. 1006, abstract, 4:29-67, 7:36-50, 9:11-32, 11:45-12:7.) A POSITA would understand that an “I/O request” includes both read and write operations. Nakayama confirms my opinion and explicitly states that its host computers make access requests in the form of read and write operations to storage including the cache memory and disk drive. (Ex. 1007, 3:66-4:25 (“An access request 204 (write), 205 (read), 206 (write) and 207 (read) which are generated by the host computers A200 and B201 to the storage media 203 are issued by the storage control 210. These access requests are executed by the data access unit 202.”), 4:48-51, 4:61-65.)

203. Furthermore, in addition to routing an I/O request to a cache memory or to a storage device, Nakayama discloses denying an I/O request. (Ex. 1007, 6:11-14.) When Nakayama's control processor 305 receives a data access command (which can include a read command or a write command), the processor attempts to obtain access rights to a drive. (*Id.*) If the processor fails to obtain access rights to the drive, data access ends and the I/O request is denied. (*Id.*) A POSITA would understand that this process is equally applicable to Heil should Heil's cache or disk storage be similarly inaccessible as part of an access request.

204. In my opinion, a POSITA would understand from Nakayama's explicit teaching of access requests in the form of read and write operations that I/O traffic may comprise a read request or a write operations. In addition, as I discussed in Section VIII.C.1, a POSITA would use Nakayama's read and write operations with Heil's storage system. Therefore, Heil in view of Nakayama renders obvious claim 11.

11. Claim 12

205. Claim 12 depends from claim 1 and further requires "the host device comprises a server." As described above in Section VIII.C.2, in my opinion Heil in view of Nakayama renders obvious claim 1.

206. In my opinion, Heil and Nakayama renders claim 12 obvious for the same reasons Heil anticipates claim 12 as discussed in Section VIII.B.9 above.

12. Claim 13

207. In my opinion, Heil in view of Nakayama discloses and renders obvious all of the limitations of claim 13.

a) 13.pre: “A method comprising:”

13.b: “receiving first data via a network at a second interface”

13.c: “storing second data at a cache memory”

13.d: “storing third data at a storage device”

13.e: “accessing the cache memory during processing of the I/O traffic”

208. In my opinion, Heil and Nakayama renders elements 1.pre, 1.b, 1.c, 1.d, and 1.e obvious for the same reasons Heil anticipates these elements as discussed in Section VIII.B.1 above. Heil discloses a method of data storage and thus teaches the preamble. (Ex. 1006, 1:9-14 (“This invention relates to distribution of storage data within computer nodes of a clustered computing environment, and more particularly, to *an apparatus and method for I/O request*


shipping operable using peer-to-peer communications within host bus adapters of the clustered systems.”.)

209. In addition, Heil in view of Nakayama renders obvious storing data (i.e., writing data) to a cache or disk storage for the reasons discussed in Section VIII.C.5 regarding claim 6 above.

b) 13.a: “receiving input/output (I/O) traffic from a host device via a dedicated I/O channel at a first interface, the I/O traffic comprising a write command”

210. In my opinion, Heil discloses receiving input/output (I/O) traffic from a host device via a dedicated I/O channel at a first interface for the reasons discussed in Section VIII.B.1 above regarding element 1.a.

211. As I discussed in Section VIII.A, Nakayama is directed to a data storage architecture similar to Heil’s in which a host computer 300 is connected by a channel 301 (which may be a SCSI bus control, as seen at 304) to a disk control 302 having a cache memory 309. (Ex. 1007, 4:40-5:9, 4:56-60.) This is shown by Fig. 3 of Nakayama:



(Ex. 1007, Fig. 3.)

212. Heil characterizes the communications between its host system and its HBA as “I/O requests” and describes responding to such I/O requests by searching through data in the cache memory and disk drives of HBAs local and remote to the host system issuing the I/O request. (Ex. 1006, abstract, 4:29-67, 7:36-50, 9:11-32, 11:45-12:7.)

213. In my opinion, a POSITA would understand that an “I/O request” includes both read and write commands. Nakayama confirms my opinion and explicitly states that its host computers make access requests in the form of read and write commands to storage including the cache memory and disk drive. (Ex. 1007, 3:66-4:25 (“An access request 204 (write), 205 (read), 206 (write) and 207 (read) which are generated by the host computers A200 and B201 to the storage media 203 are issued by the storage control 210. These access requests are executed by the data access unit 202.”), 4:48-51, 4:61-65.)


214. To the extent Heil does not use the words “read command” or “write command” to describe its “I/O requests,” it is my opinion that a POSITA would understand from Nakayama’s explicit teaching of access requests in the form of read and write commands that I/O traffic may comprise a write command. In addition, as I discussed in Section VIII.C.1, a POSITA would use Nakayama’s read and write commands with Heil’s storage system. Therefore, Heil in view of Nakayama renders obvious element 13.a.

- c) **13.f: “performing one or more access operations at the storage device based on the I/O traffic, the one or more access operations utilizing a communication path between a processor and the storage device, the communication path distinct from the dedicated I/O channel.”**

215. In my opinion, Heil discloses element 13.f for the reasons discussed in Section VIII.B.1 above regarding elements 1.e and 1.f. Heil discloses that the processor (of the HBA) may access the storage device (the local drive) based on the I/O traffic. For example, in response to an I/O request by CPU 100, the HBA may determine that the request may be serviced locally and retrieve the response data from the local disks. (Ex. 1006, 7:36-37 (“As shown in FIG. 1, one HBA services the I/O requests from CPU 1 and CPU 2 100.”), 11:45-53 (“After receiving a block I/O request 400, the 110 redirector has the means to search the directory and determining means to locate the local or remote disk drives that are storing the I/O requested blocks 410. If the location of the blocks is local 420, then the blocks are retrieved from the respective local disk drive.”).)

216. In addition, Figure 1 of Heil shows that the access operations at the storage device utilize a communication path from the processor to the storage device—i.e., from the CPU 117.1 to the local drives 118 via a SCSI interface,

annotated below in **red**—that is distinct from the dedicated I/O channel—i.e., from the Node 1 host system to HBA 117 via a PCI bus, annotated below in **green**:



(Ex. 1006, Fig. 1 (annotated).)

217. Thus, Heil in view of Nakayama renders obvious claim 13.

13. Claim 14

218. In my opinion, Heil and Nakayama renders obvious claim 14 for the same reasons Heil anticipates claim 3. As discussed in Section VIII.B.3 above,

Heil's I/O traffic is distinct from the first data received by the second interface via the network.


14. Claim 15

219. In my opinion, Heil in view of Nakayama discloses and renders obvious all of the limitations required by claim 15.

220. As described above in Section VIII.C.12, in my opinion Heil in view of Nakayama renders obvious claim 13.

221. Claim 15 depends from claim 13 and further requires that "performing a first access operation at the storage device independently of the dedicated I/O channel." Heil in view of Nakayama renders obvious this limitation.

222. As discussed in Section VIII.B.6 above with respect to claim 9, Heil discloses that the disk storage device and the PCI bus I/O channel are independently accessible. For example, Fig. 1 of Heil shows two independent pathways (1) from the Node 1 host via the PCI bus to the processor (the dedicated I/O channel) and (2) from a remote Node N host via the fiber network and fibre channel chip 120 to the local drives via a SCSI or FCAL interface:



(Ex. 1006, Fig. 1 (annotated).)

223. Additionally, Heil discloses that the storage device is accessible based on a network request from a requestor HBA on the network, rather than the local HBA. (Ex. 1006, 11:54-12:6.) Heil discloses and Figure 1 shows that a requestor HBA servicing an I/O request may “ship” the request to another HBA having data responsive to the I/O request using the fibre channel pathway, rather than the PCI bus pathway for I/O requests from the host system. (Ex. 1006, 4:29-39, 4:58-67, 9:11-25, 12:8-18.) Therefore, access operations by a requestor HBA on the network to the storage device is independent of I/O operations from the host

system to the local HBA using the dedicated I/O channel. Thus, Heil in view of Nakayama renders obvious claim 15.

15. Claim 16

224. In my opinion, Heil in view of Nakayama discloses and renders obvious all of the limitations required by claim 16.

225. As described above in Section VIII.C.12, in my opinion Heil in view of Nakayama renders obvious claim 13.

226. Claim 16 depends from claim 13 and further requires “the second data is written at the cache memory in response to the write command.” Heil in view of Nakayama renders obvious this limitation.

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

228. To the extent claim 16 requires a “write command” and Heil describes I/O requests between its host computers and HBAs as “access requests” but does not use the word “write command” to describe its access requests, Nakayama explicitly discloses read and write commands. Nakayama further discloses a write

command that writes data to the cache memory. (Ex. 1007, abstract (“The control processor writes write data in a cache slot of the cache memory.”), 6:20-35 (“If it is a write command, operation will be conducted as follows:...the control processor 305[] writes the converted write data into the cache slot.”).) In my opinion, a POSITA would apply Nakayama’s teaching of read and write commands to Heil.

229. I base my opinion on the fact that Heil characterizes the communications between its host system and its HBA as “I/O requests” and describes responding to such I/O requests by searching through data in the cache memory and disk drives of HBAs local and remote to the host system issuing the I/O request. (Ex. 1006, abstract, 4:29-67, 7:36-50, 9:11-32, 11:45-12:7.) A POSITA would understand that an “I/O request” includes both read and write commands. Nakayama confirms my opinion and explicitly states that its host computers make access requests in the form of read and write commands to storage including the cache memory and disk drive. (Ex. 1007, 3:66-4:25 (“An access request 204 (write), 205 (read), 206 (write) and 207 (read) which are generated by the host computers A200 and B201 to the storage media 203 are issued by the storage control 210. These access requests are executed by the data access unit 202.”), 4:48-51, 4:61-65.)

230. In my opinion, a POSITA would understand from Nakayama's explicit teaching of access requests in the form of read and write commands that I/O traffic may comprise a read request or a write command. In addition, as I discussed in Section VIII.C.1, a POSITA would use Nakayama's read and write commands with the cache memory of Heil's storage system. Therefore, Heil in view of Nakayama renders obvious claim 16.

16. Claim 17

231. In my opinion, Heil in view of Nakayama discloses and renders obvious all of the limitations required by claim 17.

232. As described above in Section VIII.C.12, in my opinion Heil in view of Nakayama renders obvious claim 17's parent claim, claim 13.

a) 17.a: "The method of claim 13, receiving a read request via the network"

233. Claim 17 depends from claim 13. The first element, 17.a, requires "receiving a read request via the network." In my opinion, Heil in view of Nakayama renders obvious this element.

234. Specifically, Heil discloses that read requests from a remote host may be sent via the fiber network to HBAs connected to the network. (Ex. 1006, 11:56-12:1 ("If the requested I/O block is stored in a remote disk, the initiating HBA

ships the block I/O request to the peer HBA that the initiating HBA has determined can perform the requisite processing 450...The remote HBA retrieves the requested blocks of data from its HBA cache, or if the blocks of data are not within cache, then directly from the remote HBA's local disk.”.)

235. In addition, to the extent Heil does not use the words “read request” or “write request,” Nakayama explicitly discloses read and write requests. In my opinion, a POSITA would apply Nakayama’s teaching of read and write requests to Heil.

236. I base my opinion on the fact that Heil characterizes the communications between its host system and its HBA as “I/O requests” and describes responding to such I/O requests by searching through data in the cache memory and disk drives of HBAs local and remote to the host system issuing the I/O request. (Ex. 1006, abstract, 4:29-67, 7:36-50, 9:11-32, 11:45-12:7.) A POSITA would understand that an “I/O request” includes both read and write requests. Nakayama confirms my opinion and explicitly states that its host computers make access requests in the form of read and write requests to storage including the cache memory and disk drive. (Ex. 1007, 3:66-4:25 (“An access request 204 (write), 205 (read), 206 (write) and 207 (read) which are generated by

the host computers A200 and B201 to the storage media 203 are issued by the storage control 210. These access requests are executed by the data access unit 202.”), 4:48-51, 4:61-65.)

237. In my opinion, a POSITA would understand from Nakayama’s explicit teaching of access requests in the form of read and write requests that I/O traffic may comprise a read request. In addition, as I discussed in Section VIII.C.1, a POSITA would use Nakayama’s read and write requests with Heil’s storage system. Therefore, Heil in view of Nakayama renders obvious element 17.a.

b) 17.b: “transmitting at least a portion of the second data from the cache memory via the network responsive to the read request being associated with a cache hit”

238. The first element, 17.b, requires “transmitting at least a portion of the second data from the cache memory via the network responsive to the read request being associated with a cache hit.” In my opinion, Heil in view of Nakayama renders obvious this element.

239. Specifically, Heil discloses that, in response to a read request from a remote host, data responsive to the read request is retrieved from cache memory and transmitted back to the remote host. (Ex. 1006, 12:1-10 (“The remote HBA retrieves the requested blocks of data from its HBA cache, or if the blocks of data

are not within cache, then directly from the remote HBA's local disk. The remote HBA can cache blocks of data retrieved from its local disk, to service future local or remotely generated requests for these blocks. The remote I/O blocks are sent from the HBA performing the remote processing (e.g., a read request) to the initiating HBA 460.”.) Thus, Heil discloses element 17.b.

17. Claim 18

240. In my opinion, Heil in view of Nakayama discloses and renders obvious all of the limitations required by claim 18.

241. As described above in Section VIII.C.12, in my opinion Heil in view of Nakayama renders obvious claim 13.

242. Claim 18 depends from claim 13 and further requires “storing fourth data at the cache memory, wherein the second data is indicated by the write command, and wherein the fourth data is indicated by a second write command received at the second interface.” Heil in view of Nakayama renders obvious this limitation.

243. Heil discloses that a HBA may store in its cache a directory of block locations in cache in addition to requested data. (Ex. 1006, 12:19-25 (“The HBA's via I/O redirector 240, 331 has the means to create the directory storing the

mapping of the local and/or remote location of blocks within the clustered computing environment. Specifically, the data within each respective local disk is mapped and coalesced into a directory. FIG. 4A shows the preferred method to build the directory.”), 11:8-9.) In addition, Heil discloses that I/O requests may be transmitted through the fiber network to update the directory (i.e., store directory updates, the fourth data). (Ex. 1006, 12:25-33 (“After the host(s) and HBA(s) initialize 500, the HBAs map the location of blocks within their respective local storage subsystem 502. The HBAs build a local directory containing the mapped location information for their respective local storage subsystem 504. The HBAs use their peer-to-peer capabilities to broadcast the contents of their directory to peer HBAs 506 and each HBA updates its respective directory with the broadcast information 508.”).)

244. Based on these disclosures, a POSITA would understand that the directory is fourth data stored by a second I/O request received at the second, network interface.

245. In addition, Heil discloses that an HBA “can cache blocks of data retrieved from its local disk, to service future local or remotely generated requests for these blocks.” (Ex. 1006, 12:1-4.) Based on this disclosure, a POSITA would

understand that more than one I/O request may be issued by a host system—Heil explicitly contemplates that the host system may issue a second I/O request to the local HBA even after the host system issues a first I/O request to the local HBA. A POSITA would understand that the second I/O request corresponds to the claimed second read request. If data responsive to the second I/O request is not found at the local HBA, and the local HBA ships that request to a remote HBA, then the remote HBA is the apparatus that has received the second read request via the network.

246. To the extent claim 18 requires a “write command” and Heil does not use the word “write command” to describe its access requests, Nakayama explicitly discloses read and write commands. In my opinion, a POSITA would apply Nakayama’s teaching of read and write commands to Heil.

247. I base my opinion on the fact that Heil characterizes the communications between its host system and its HBA as “I/O requests” and describes responding to such I/O requests by searching through data in the cache memory and disk drives of HBAs local and remote to the host system issuing the I/O request. (Ex. 1006, abstract, 4:29-67, 7:36-50, 9:11-32, 11:45-12:7.) A POSITA would understand that an “I/O request” includes both read and write commands. Nakayama confirms my opinion and explicitly states that its host

computers make access requests in the form of read and write commands to storage including the cache memory and disk drive. (Ex. 1007, 3:66-4:25 (“An access request 204 (write), 205 (read), 206 (write) and 207 (read) which are generated by the host computers A200 and B201 to the storage media 203 are issued by the storage control 210. These access requests are executed by the data access unit 202.”), 4:48-51, 4:61-65.)

248. In my opinion, a POSITA would understand from Nakayama’s explicit teaching of access requests in the form of read and write commands that I/O traffic may comprise a read request or a write command. In addition, as I discussed in Section VIII.C.1, a POSITA would use Nakayama’s read and write commands with Heil’s storage system. Therefore, Heil in view of Nakayama renders obvious claim 18.

18. Claim 19

249. In my opinion, Heil in view of Nakayama discloses and renders obvious all of the limitations of claim 19.

- a) **19.pre: “An apparatus comprising:”**
- 19.b: “a second interface configured to receive data via
a network”**
- 19.c: “a cache memory”**
- 19.d: “a storage device”**
- 19.e: “a processor coupled via a communication path to
the storage device, the processor configured to
access the cache memory during processing of the
I/O traffic”**
- 19.f: “the processor further configured to perform one
or more access operations at the storage device
based on the I/O traffic, wherein the
communication path is distinct from the
dedicated I/O channel.”**

250. In my opinion, Heil and Nakayama renders elements 19.pre, 19.b, 19.c, 19.d, 19.e, and 19.f obvious for the same reasons Heil anticipates these elements as discussed in Section VIII.B.10 above.

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

251. In my opinion, Heil and Nakayama renders element 19.a obvious for the same reasons Heil anticipates this element as discussed in Section VIII.B.10 above.

252. To the extent element 19.a requires “the I/O traffic comprising one or more read commands, one or more write commands, or a combination thereof” and Heil does not use the words “read command” or “write command” to describe its access requests, Nakayama explicitly discloses read and write commands. In my opinion, a POSITA would apply Nakayama’s teaching of read and write commands to Heil.

253. I base my opinion on the fact that Heil characterizes the communications between its host system and its HBA as “I/O requests” and describes responding to such I/O requests by searching through data in the cache memory and disk drives of HBAs local and remote to the host system issuing the I/O request. (Ex. 1006, abstract, 4:29-67, 7:36-50, 9:11-32, 11:45-12:7.) A POSITA would understand that an “I/O request” includes both read and write

commands. Nakayama confirms my opinion and explicitly states that its host computers make access requests in the form of read and write commands to storage including the cache memory and disk drive. (Ex. 1007, 3:66-4:25 (“An access request 204 (write), 205 (read), 206 (write) and 207 (read) which are generated by the host computers A200 and B201 to the storage media 203 are issued by the storage control 210. These access requests are executed by the data access unit 202.”), 4:48-51, 4:61-65.)

254. In my opinion, a POSITA would understand from Nakayama’s explicit teaching of access requests in the form of read and write commands that I/O traffic may comprise a read command or a write command. In addition, as I discussed in Section VIII.C.1, a POSITA would use Nakayama’s read and write commands with Heil’s storage system. Therefore, Heil in view of Nakayama renders obvious claim 19.

19. Claim 20

255. Claim 20 depends from claim 19 and further requires “the dedicated I/O channel is associated with dedicated throughput that corresponds to the I/O traffic.” As explained above, Heil in view of Nakayama renders obvious claim 19.

256. In my opinion, Heil and Nakayama renders claim 20 obvious for the same reasons Heil anticipates claim 20 as discussed in Section VIII.B.11 above.

20. Claim 21

257. Claim 21 depends from claim 19 and further requires “the I/O traffic is distinct from the data.” As explained above, Heil in view of Nakayama renders obvious claim 19.

258. In my opinion, Heil and Nakayama renders claim 21 obvious for the same reasons Heil anticipates claim 21 as discussed in Section VIII.B.12 above.

21. Claim 22

259. Claim 22 depends from claim 19 and further requires “the data corresponds to one or more read requests, one or more write requests, or a combination thereof.” As explained above, Heil in view of Nakayama renders obvious claim 19.

260. In my opinion, Heil and Nakayama renders claim 22 obvious for the same reasons Heil anticipates claim 22 as discussed in Section VIII.B.13 above.

261. In addition, to the extent Heil does not use the words “read request” or “write request” to describe its access requests, Nakayama explicitly discloses read

and write requests. In my opinion, a POSITA would apply Nakayama's teaching of read and write requests to Heil.

262. I base my opinion on the fact that Heil characterizes the communications between its host system and its HBA as "I/O requests" and describes responding to such I/O requests by searching through data in the cache memory and disk drives of HBAs local and remote to the host system issuing the I/O request. (Ex. 1006, abstract, 4:29-67, 7:36-50, 9:11-32, 11:45-12:7.) A POSITA would understand that an "I/O request" includes both read and write commands. Nakayama confirms my opinion and explicitly states that its host computers make access requests in the form of read and write requests to storage including the cache memory and disk drive. (Ex. 1007, 3:66-4:25 ("An access request 204 (write), 205 (read), 206 (write) and 207 (read) which are generated by the host computers A200 and B201 to the storage media 203 are issued by the storage control 210. These access requests are executed by the data access unit 202."), 4:48-51, 4:61-65.)

263. In my opinion, a POSITA would understand from Nakayama's explicit teaching of access requests in the form of read and write requests that I/O traffic may comprise a read request or a write request. In addition, as I discussed in

Section VIII.C.1, a POSITA would use Nakayama's read and write request with Heil's storage system. Therefore, Heil in view of Nakayama renders obvious claim 22.

22. Claim 23

264. Claim 23 depends from claim 19 and further requires "the storage device and the dedicated I/O channel are independently accessible." As explained above, Heil in view of Nakayama renders obvious claim 19.

265. In my opinion, Heil and Nakayama renders claim 23 obvious for the same reasons Heil anticipates claim 23 as discussed in Section VIII.B.14 above.

23. Claim 24

266. In my opinion, Heil in view of Nakayama renders obvious all of the limitations of claim 24.

a) 24.pre: "An apparatus comprising"

267. The preamble of claim 24 recites "an apparatus." In my opinion, Heil discloses the preamble. Heil explicitly states that: "*This invention relates to* distribution of storage data within computer nodes of a clustered computing environment, and more particularly, to *an apparatus* and method for I/O request

shipping operable using peer-to-peer communications within host bus adapters of the clustered systems.” (Ex. 1006, 1:9-14.)


b) 24.a: “means for receiving input/output (I/O) traffic from a first host device via a dedicated I/O channel, the I/O traffic comprising a write command”

268. Element 24.a requires “means for receiving input/output (I/O) traffic from a first host device via a dedicated I/O channel, the I/O traffic comprising a write command.” As discussed in Section VII.A, I understand that this is a means-plus-function limitation performing the function of receiving input/output (I/O) traffic from a first host device via a dedicated I/O channel with a corresponding structure of a front-end module that receives data from hosts via I/O channels, and equivalents. Heil discloses both the function and the same structure that performs the function as the ’092 Patent.

269. As discussed in Sections VIII.C.2.b and VIII.C.12.b for elements 1.a and 13.a, Heil discloses performing the function of “receiving input/output (I/O) traffic from a first host device via a dedicated I/O channel.”

270. Heil discloses that this function is performed by the same corresponding structure as the ’092 Patent. Specifically, Heil discloses that a front-

end module (PCI bus bridge 115, PCI bus 116.5, or Front End I/F 102) receives I/O traffic from a host device:




(Ex. 1006, Fig. 1, 6:35-40 (“Node 1 150 includes a host system connected to a HBA 117 via a Host-to-Peripheral Component Interconnect (hereinafter referred to as “PCI”) bus bridge 115 and PCI bus 116.5.”), 6:67-7:1 (“The front-end interface 102 is a PCI-to-PCI Bus bridge that connects, via PCI bus 117.6, various peripheral I/O devices, such as the local drives 118 and Fibre Channel Chip 120, to

a centralized system PCI bus 116.5. Front-end interface 102 is connected to embedded Host-to-PCI Bus bridge 101 via PCI bus 117.6.”.)

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

272. It is my opinion that a POSITA would understand that the communication path between CPU 100 and the HBA, including PCI Bus Bridge 115, PCI Bus 116.5, and Front End I/F 102, is a dedicated I/O channel because it is dedicated to I/O between HBA 117 and the Node 1 host system.

273. As I discussed in Section VIII.A, Nakayama is directed to a data storage architecture similar to Heil’s in which a host computer 300 is connected by a channel 301 (which may be a SCSI bus control, as seen at 304) to a disk control 302 having a cache memory 309. (Ex. 1007, 4:40-5:9, 4:56-60.) This is shown by Fig. 3 of Nakayama:



(Ex. 1007, Fig. 3.)

274. Heil characterizes the communications between its host system and its HBA as “I/O requests” and describes responding to such I/O requests by searching through data in the cache memory and disk drives of HBAs local and remote to the host system issuing the I/O request. (Ex. 1006, abstract, 4:29-67, 7:36-50, 9:11-32, 11:45-12:7.)

275. In my opinion, a POSITA would understand that an “I/O request” includes both read and write commands. Nakayama confirms my opinion and explicitly states that its host computers make access requests in the form of read and write requests to storage including the cache memory and disk drive. (Ex. 1007, 3:66-4:25 (“An access request 204 (write), 205 (read), 206 (write) and 207 (read) which are generated by the host computers A200 and B201 to the storage media 203 are issued by the storage control 210. These access requests are executed by the data access unit 202.”), 4:48-51, 4:61-65.)

276. To the extent Heil does not explicitly use the words “read command” or “write command” to describe its “I/O requests,” it is my opinion that a POSITA would understand from Nakayama’s explicit teaching of access requests in the form of read and write commands that I/O traffic may comprise a write command. In addition, as I discussed in Section VIII.C.1, a POSITA would use Nakayama’s read and write commands with Heil’s storage system. Therefore, Heil in view of Nakayama renders obvious element 24.a.


c) 24.b: “means for receiving data via a network”

277. Element 24.b requires “means for receiving data via a network.” As discussed in Section VII.A, I understand that this is a means-plus-function

limitation performing the function of receiving data via a network with a corresponding structure of a network adapter, and equivalents. A network adapter is commonly understood as a component of a computer system that communicates over a network with another computer system, including by receiving data via the network. Heil discloses both the function and the same structure that performs the function as the '092 Patent.

278. In my opinion, Heil discloses performing the function of “receiving data via a network” for the reasons discussed in Section VIII.B.1.c for element 1.b above.

279. Heil also discloses that this function is performed by the same corresponding structure: a network adapter. Specifically, Heil discloses a network adapter in the form of a fibre channel chip 120 that receives data from a fiber network 121:



(Ex. 1006, Fig. 1, 8:23-26 (“Fibre Channel Chip 120 connects Node 1 150 to the rest of the clustered environment via the network Fibre Channel backbone 121.”).) Fibre channel chip 120 is a component of the HBA that communicates with other HBAs over a network (disclosed by Heil as the fibre channel network 121). Therefore, a POSITA would understand that a fibre channel chip is a network adapter.

280. Thus, Heil in view of Nakayama renders obvious element 24.b.


d) 24.c: “means for short-term data storage”

281. Element 24.c requires “means for short-term data storage.” As discussed in Section VII.A, I understand that this is a means-plus-function limitation performing the function of short-term data storage with a corresponding

structure of a cache memory, and equivalents. Heil discloses both the function and the same structure that performs the function as the '092 Patent.

282. In my opinion, Heil discloses performing the function of “short-term data storage” for the reasons discussed in Section VIII.B.1.d for element 1.c above.

283. Heil also discloses that this function is performed by the same corresponding structure of a cache memory in the form of cache 117.2:



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


284. Thus, Heil in view of Nakayama renders obvious element 24.c.

e) **24.d: “means for long-term data storage”**

285. Element 24.d requires “means for long-term data storage.” As discussed in Section VII.A, I understand that this is a means-plus-function limitation performing the function of long-term data storage with a corresponding structure of a back-end storage device or local storage device, and equivalents. Heil discloses both the function and the same structure that performs the function as the '092 Patent.

286. In my opinion, Heil discloses performing the function of “long-term data storage” for the reasons discussed in Section VIII.B.1.e for element 1.d above.

287. Heil also discloses that this function is performed by the same corresponding structure of a local storage device in the form of local drives 118 and 123:



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


288. Thus, Heil in view of Nakayama renders obvious element 24.d.

f) **24.e: “means for performing one or more access operations at the means for short-term data storage during processing of the I/O traffic and for performing one or more access operations at the means for long-term data storage based on the I/O traffic, the means for performing distinct from the dedicated I/O channel”**

289. Element 24.e requires “means for performing one or more access operations at the means for short-term data storage during processing of the I/O traffic and for performing one or more access operations at the means for long-term data storage based on the I/O traffic, the means for performing distinct from the dedicated I/O channel.” As discussed in Section VII.A, I understand that this is a means-plus-function limitation performing the function of performing one or more access operations at the means for short-term data storage during processing of the I/O traffic and performing one or more access operations at the means for long-term data storage based on the I/O traffic with a corresponding structure of a cache manager and a configuration manager, and equivalents. Heil discloses both the function and the same structure that performs the function as the '092 Patent.

290. In my opinion, Heil discloses performing the function of performing one or more access operations at the means for short-term data storage during processing of the I/O traffic and performing one or more access operations at the means for long-term data storage based on the I/O traffic for the reasons discussed in Sections VIII.B.1.e and VIII.B.1.g for elements 1.d and 1.f above.

291. Heil also discloses that this function is performed by the same corresponding structure of a processor 117.1 that manages access operations to cache memory (cache 117.2) and to a storage device (local drives 118), as well as software drivers referred to as ISMs for accessing cache and local storage:



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


292. Heil describes that its processor supported by ISMs performs access operations to the cache during processing of I/O traffic and accesses cache memory while processing I/O traffic. (Ex. 1006, 7:25-28 (“To improve I/O performance, CPU 117.1 temporarily caches disk data in cache memory 117.2 that is frequently accessed. Later, the CPU 117.1 stores the data permanently in the slower disk drive devices 118.”), 12:1-4 (“The remote HBA can cache blocks of data retrieved from its local disk, to service future local or remotely generated requests for these blocks.”).)

293. In addition, the processor supported by ISMs performs access operations to the local storage device based on I/O traffic, for example if the I/O traffic directs retrieval of data from the cache, and from the local disk if the requested data is not found in cache:

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

(Ex. 1006, 11:66-12:1.)

294. As I explained in Section VIII.B.1.f regarding element 1.e and Section VIII.B.3 regarding claim 3, the processor's access operations to the cache or disk drive are distinct from the dedicated I/O channel, as shown by Figure 1 (annotated) below:



(Ex. 1001, Fig. 1.)

295. Thus, Heil in view of Nakayama renders obvious element 24.e.

D. Ground 4: Obviousness of Claim 4 Over Heil, Nakayama, and Gulick

296. In my opinion, when viewed and considered considering the knowledge possessed by one of ordinary skill in the art, Heil in view of Nakayama and Gulick discloses all the limitations of, and renders obvious, claim 4 of the '092 Patent.

297. As described in detail below, one of ordinary skill in the art would have been motivated to combine and collectively consider the teachings of these references to arrive at the subject matter claimed by the '092 Patent and would have had a reasonable expectation of success when doing so.

1. Motivation to Combine

298. It would have been obvious for a POSITA to combine Heil, Nakayama, and Gulick with a reasonable chance of success in doing so. A POSITA would have combined Heil and Nakayama as explained in Section VIII.C.1 above.

299. A POSITA would further look to Gulick to interface SCSI adapters with PCI buses, as Gulick explicitly discloses that it is possible to apply SCSI adapters to PCI buses to interface with SCSI devices, including drive storage. (Ex. 1008, 5:37-51.) This is done by a “SCSI adapter” that is coupled to the “PCI bus”

and to “various SCSI devices,” including drive storage. (*Id.*) As I explained in Section VI, SCSI is the most frequently encountered channel for small form disk drives in the mid to late 90s and offered advantages such as treating peripheral devices in a largely independent fashion. A POSITA wishing to interface a device having a PCI bus to peripheral small form disk drives would therefore be interested in interfacing a PCI bus with a SCSI adapter to a SCSI compatible drive. A POSITA reading Gulick would know that in 1995 such adapters were available. In my opinion, using such adapters would produce predictable results.

2. Claim 4

“The apparatus of claim 1, wherein the dedicated I/O channel comprises a small computer system interface (SCSI) channel.”

300. In my opinion, Heil in view of Nakayama and Gulick, when considered in view of the knowledge possessed by one of ordinary skill in the art, discloses and renders obvious all of the limitations required by claim 4.

301. As described above in Section VIII.B.1, in my opinion claim 1 is anticipated by Heil, and as described above in Section VIII.C.2, claim 1 is obvious over Heil in view of Nakayama.

302. Claim 4 depends from claim 1 and further requires that “the dedicated I/O channel comprises a small computer system interface (SCSI) channel.”

303. In my opinion, claim 4 is directed to an embodiment where the dedicated I/O channel from the host device to the apparatus is a SCSI channel.

304. As described above in Section VIII.C.2, Heil and Nakayama disclose the limitations of claim 1, but Heil’s host servers are connected to the HBAs via a PCI bus. Heil therefore utilizes a PCI channel as the dedicated I/O channel rather than a SCSI channel.

305. Gulick describes methods of connecting computer systems to peripheral devices. (Ex. 1008, 1:15-26, 5:36-51.) Specifically, Gulick teaches that computer systems may be connected to peripheral devices via a PCI bus (a “host/PCI/cache bridge [that] interfaces to a peripheral expansion bus, preferably a peripheral component interconnect (PCI) bus”), which a POSITA would understand to be the same type of PCI bus connection as between Heil’s host servers and HBAs. (*Id.*, 5:36-51.) Gulick further teaches that a SCSI adapter may be coupled to the PCI bus to interface the bridge with various SCSI devices, such as a CD-ROM drive and a tape drive. (*Id.*) A POSITA would understand from this

disclosure that Gulick provides a SCSI channel between the computer system and the SCSI compatible device.

306. As I discussed above, a POSITA would have been aware that SCSI channels are a known alternative to PCI channels and would have been motivated to combine Heil, Nakayama, and Gulick. Additionally, as also discussed above, a POSITA would have been aware that use of SCSI channels provides benefits over use of PCI, and would have been motivated to replace Heil's PCI channel between the host server and HBA by a SCSI channel as the dedicated I/O channel as taught by Gulick. It is therefore my opinion that a POSITA would have been motivated by Gulick's teachings to use a SCSI channel as taught by Gulick between Heil's host server and HBA. Gulick explicitly teaches the POSITA that SCSI channels may be used to interface computer systems, like Heil's host server, with other devices. (*Id.*, 5:36-51.) In my opinion, this renders obvious claim 4's SCSI channel.

E. Ground 5: Obviousness of Claim 5 Over Heil, Nakayama, and Berman

307. In my opinion, when viewed and considered considering the knowledge possessed by one of ordinary skill in the art, Heil in view of Nakayama

and Berman discloses all the limitations of, and renders obvious, claim 5 of the '092 Patent.

308. As described in detail below, one of ordinary skill in the art would have been motivated to combine and collectively consider the teachings of these references to arrive at the subject matter claimed by the '092 Patent and would have had a reasonable expectation of success when doing so.

1. Motivation to Combine

309. It would have been obvious for a POSITA to combine Heil, Nakayama, and Berman with a reasonable chance of success in doing so. A POSITA would have combined Heil and Nakayama as explained in Section VIII.C.1 above.

310. A POSITA in 1999 would understand that fibre channel networks have improved performance over Ethernet networks. However, the POSITA in 1999 would know that, for a variety of reasons including cost and existing infrastructure, Ethernet networks are more widely deployed than fibre channel networks. A POSITA would therefore further look to Berman to interface ethernet networks with Heil's fibre channel network. Berman explicitly teaches that it is advantageous to do so, and a POSITA would understand that interfacing a fibre

channel with a legacy network such as ethernet would yield predictable results. (Ex. 1009, 12:9-11, claims 13-15, 3:23-30.) The POSITA reading Berman would have known that a router can be used to connect fiber channel with Ethernet and that building such connections would yield predictable results. Thus, the POSITA would be motivated to use fibre channel networks where they are available and interface fibre channel networks with existing (“legacy”) Ethernet networks where fibre channel networks have not yet been deployed. (*Id.*)

2. Claim 5

“The apparatus of claim 1, wherein the second interface comprises an Ethernet interface or an asynchronous transfer mode (ATM) interface.”

311. In my opinion, Heil in view of Nakayama and Berman, when considered in view of the knowledge possessed by one of ordinary skill in the art, discloses and renders obvious all of the limitations required by claim 5.

312. As described above in Section VIII.B.1, in my opinion claim 1 is anticipated by Heil, and as described above in Section VIII.C.2, claim 1 is obvious over Heil in view of Nakayama.

313. Claim 5 depends from claim 1 and further requires that “the second interface comprises an Ethernet interface or an asynchronous transfer mode (ATM) interface.” Berman discloses this limitation.

314. Claim 5 is directed to an embodiment where the second interface from the apparatus to a network is an Ethernet interface or an ATM interface. (Ex. 1001, 4:25-27 (“The preferred embodiment of this invention uses SCSI channels to connect to local hosts and uses standard network links card such as Ethernet, or ATM to connect to remote hosts.”).)

315. As described above in Section VIII.C.2, Heil and Nakayama disclose the limitations of claim 1, but Heil’s HBAs are connected via a Fibre Channel network. Heil therefore utilizes a Fibre Channel interface as the second interface rather than an Ethernet or ATM interface.

316. Berman describes networking systems utilizing the same Fibre Channel technology as used by Heil. (Ex. 1009, abstract, 1:34-47.) Berman further describes that it would be advantageous to interface Fibre Channels with “legacy networks” such as Ethernet and discloses systems such as routers that connect Fiber Channel to Ethernet. (Ex. 1009, 12:9-11, claims 13-15, 3:23-30.) A POSITA

would recognize that a network interface to an Ethernet network is an Ethernet interface.

317. In my opinion, a POSITA would have been motivated by Berman's teachings to interface Heil's Fibre Channels with Ethernet networks as taught by Berman. Berman explicitly teaches the POSITA to do so to provide compatibility with Ethernet "legacy networks." (Ex. 1009, 12:9-11.) A POSITA would be further motivated to combine Heil, Nakayama, and Berman for the reasons discussed above. Interfacing Heil's Fibre Channel to an Ethernet network as taught by Berman discloses a network interface between the apparatus and the network that is an Ethernet interface, and therefore renders obvious claim 5's Ethernet interface.

F. Secondary Considerations of Non-Obviousness

318. Based on my review of the '092 Patent's prosecution history, I note that no secondary considerations of non-obviousness were identified in connection with the patent's claims.

319. I have also been informed that Patent Owner has not identified any evidence related to secondary considerations in the district court litigation.

320. I am not separately aware of any evidence of commercial success, skepticism, failure of others, industry praise, unexpected results, or the like. Instead, as discussed above, it is my opinion that the subject matter of the '092 Patent's claims were known in the art and would have been expected to function as part of a wireless communication device with location determining capabilities.

IX. CONCLUSION

321. I hereby declare that all statements made herein of my own knowledge are true and all statements made herein on information and belief were and are believed by me to be true, and that all statements herein were and are made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that any such willful false statements may jeopardize the validity of the application or any patents issued thereon.



Paul D. Franzon

November 2, 2022