

United States District Court,  
N.D. California.

**SUN MICROSYSTEMS, INC,**  
Plaintiff.

v.

**NETWORK APPLIANCE, INC,**  
Defendant.

No. C-07-05488 EDL

**Dec. 22, 2008.**

**Background:** Owner of patents related to disk drive technology brought infringement action against competitor. Competitor counterclaimed for patent infringement. Parties sought claim construction.

**Holdings:** The District Court, Elizabeth D. Laporte, United States Magistrate Judge, held that:

(1) term "first available memory space" meant one or more logical tracks, each of which must be empty, i.e., recognized by the system as available free space;

(2) corresponding structure of term "means, responsive to the receipt of a stream of data records from said associated data processor, for writing said received stream of data records in available memory space in one of said disk drives" was "control unit 101 programmed to execute the steps 701 through 712 illustrated in Figure 7 and described under the heading 'Data Write Operation' at column 16:26 through column 17:25" of patent specification;

(3) term "means responsive to said reading of said memory controller for periodically verifying the integrity of data currently stored in each of said identified dedicated partitions" had corresponding structure of "the control unit 101 and the disk drive subassembly programmed to execute steps 1202 to 1207 illustrated in Figure 12 and described at col. 15:47-62, and step 1008 illustrated in Figure 10 and described at col. 15:63-16:25";

(4) term "means for reading data in each of said identified dedicated partitions" had corresponding structure of "the disk drive manager and the control unit 101 programmed to execute step 1202 illustrated in Figure 12 and described at col. 15:49"; and

(5) term "active links between said active nodes" meant operational paths capable of transferring data between operational nodes capable of sending and receiving data, that are separate from the storage devices.

Claims construed.

Court-Filed Expert Resumes

5,721,937, 6,049,528, 6,574,591, 6,868,417, 7,107,385, 7,130,873, 7,313,720. Cited.

Carrie Lynne Williamson, David L. Alberti, Eliza Behold, Mark Fowler, Yakov M. Zolotorev, DLA Piper, US, LLP, East Palo Alto, CA, for Plaintiff.

# ORDER CONSTRUING CLAIM TERMS OF THE '987, '855, '012 AND '787 PATENTS

ELIZABETH D. LAPORTE, **United States Magistrate Judge.**

On November 10, 2008, the Court held a hearing to construe the disputed terms of United States Patent Numbers 5,124,987 (the "'987 patent"), 5,430,855 ("'855 patent"), 5,632,012 ("'012 patent"), and 6,421,787 ("'787 patent") pursuant to *Markman v. Westview Instruments, Inc.*, 517 U.S. 370, 116 S.Ct. 1384, 134 L.Ed.2d 577 (1996). Having read the papers and considered the arguments of counsel and the relevant legal authority, the Court hereby rules as follows.

## I. BACKGROUND

On October 29, 2007, Sun Microsystems Inc. ("Sun") filed its complaint against Network Appliance Inc. ("NetApp") for Patent Infringement, Unfair Competition Under the Lanham Act and Unfair Competition under California Business and Professions Code s. 17200. On November 2, 2007, Sun filed its First Amended Complaint, alleging that NetApp infringes the four patents noted above, as well as United States Patent Nos. 6,049,528 and 5,721,937 (collectively, the "Sun patents"), through its Fabric Attached Storage, V-series and NearStore products.

On December 21, 2007, NetApp filed its Answer and Counterclaim to First Amended Complaint. NetApp denies infringing any of the Sun patents and alleges that Sun infringes United States Patent Numbers 6,574,591 (the "'591 patent"), 6,868,417 (the "'417 patent"), 7,107,385 (the "'385 patent"), and 7,130,873 (the "'873 patent"). NetApp also alleges that Sun violated the Lanham Act and California Business and Professions Code s. 17200 et seq. On January 14, 2008, Sun filed its Reply to NetApp's Answer and Counterclaims, denying the allegations and asserting a number of affirmative defenses and counterclaims.

On February 19, 2008, NetApp filed its Supplemental Answer and Counterclaims, expanding the allegations of infringement to include United States Patent Number 7,313,720 (the "'720 patent") in addition to the patents noted above (collectively, the "NetApp patents"). On March 7, 2008, Sun filed its reply to NetApp's Supplemental Answer and Counterclaims.

The parties now seek construction of five disputed terms contained in the Sun patents.

## II. LEGAL STANDARD

[1] [2] [3] In construing claims, the court must begin with an examination of the claim language itself. The terms used in the claims are generally given their "ordinary and customary meaning." *See Phillips v. AWH Corp.*, 415 F.3d 1303, 1312-13 (Fed.Cir.2005); *see also Renishaw PLC v. Marposs Societa' per Azioni*, 158 F.3d 1243, 1248 (Fed.Cir.1998) ("The claims define the scope of the right to exclude; the claim construction inquiry, therefore, begins and ends in all cases with the actual words of the claim."). This ordinary and customary meaning "is the meaning that the terms would have to a person of ordinary skill in the art in question at the time of the invention ...." *Phillips*, 415 F.3d at 1313. A patentee is presumed to have intended the ordinary meaning of a claim term in the absence of an express intent to the contrary. *York Products, Inc. v. Central Tractor Farm & Family Ctr.*, 99 F.3d 1568, 1572 (Fed.Cir.1996).

[4] Generally speaking, the words in a claim are to be interpreted "in light of the intrinsic evidence of record, including the written description, the drawings, and the prosecution history, if in evidence." *Teleflex, Inc. v. Ficosa North Am. Corp.*, 299 F.3d 1313, 1324-25 (Fed.Cir.2002) (citations omitted); *see also Medrad, Inc. v. MRI Devices Corp.*, 401 F.3d 1313, 1319 (Fed.Cir.2005) (court looks at "the ordinary meaning in the context of the written description and the prosecution history"). "Such intrinsic evidence is the most significant source of the legally operative meaning of disputed claim language." *Vitronics Corp. v.*

[5] [6] With regard to the intrinsic evidence, the court's examination begins with the claim language. *See id.* Specifically, "the context in which a term is used in the asserted claim can be highly instructive." Phillips, 415 F.3d at 1314. As part of that context, the court may also consider the other patent claims, both asserted and unasserted. *Id.* For example, as claim terms are normally used consistently throughout a patent, the usage of a term in one claim may illuminate the meaning of the same term in other claims. *Id.* The court may also consider differences between claims as a guide to understanding the meaning of particular claim terms. *Id.*

[7] [8] Second, the claims "must [also] be read in view of the specification, of which they are a part." *Id.* at 1315. When the specification reveals a special definition given to a claim term by the patentee that differs from the meaning it would otherwise possess, the inventor's lexicography governs. *Id.* at 1316. Indeed, the specification is to be viewed as the "best source" for understanding a technical term, informed as needed by the prosecution history. *Id.* at 1315. As the Federal Circuit stated in Phillips, the specification is "the single best guide to the meaning of a disputed term," and "acts as a dictionary when it expressly defines terms used in the claims or when it defines terms by implication." 415 F.3d at 1321.

[9] Limitations from the specification, however, such as from the preferred embodiment, cannot be read into the claims absent a clear intention by the patentee to do. *Altiris v. Symantec Corp.*, 318 F.3d 1363, 1372 (Fed.Cir.2003) ("resort to the rest of the specification to define a claim term is only appropriate in limited circumstances"); *Teleflex*, 299 F.3d at 1326 ("The claims must be read in view of the specification, but limitations from the specification are not to be read into the claims.") (citations omitted); *CCS Fitness, Inc. v. Brunswick Corp.*, 288 F.3d 1359, 1366 (Fed.Cir.2002) ("a patentee need not describe in the specification every conceivable and possible future embodiment of his invention").

[10] "[T]here is sometimes a fine line between reading a claim in light of the specification, and reading a limitation into the claim from the specification.... [A]ttempting to resolve that problem in the context of the particular patent is likely to capture the scope of the actual invention more accurately than either strictly limiting the scope of the claims to the embodiments disclosed in the specification or divorcing the claim language from the specification." *Decisioning.com, Inc. v. Federated Dept. Stores, Inc.*, 527 F.3d 1300, 1307-08 (Fed.Cir.2008) (quoting *Comark Comm'ns, Inc. v. Harris Corp.*, 156 F.3d 1182, 1186 (Fed.Cir.1998)). There is therefore "no magic formula or catechism for conducting claim construction," and the court must "read the specification in light of its purposes in order to determine whether the patentee is setting out specific examples of the invention to accomplish those goals, or whether the patentee instead intends for the claims and the embodiments in the specification to be strictly coextensive." *Id.* (internal citations omitted).

[11] Finally, as part of the intrinsic evidence analysis, the court "should also consider the patent's prosecution history, if it is in evidence." Phillips, 415 F.3d at 1317. The court should take into account, however, that the prosecution history "often lacks the clarity of the specification" and thus is of limited use for claim construction purposes. *Id.*

[12] [13] [14] In most cases, claims can be resolved based on intrinsic evidence. *See Vitronics*, 90 F.3d at 1583. Only if an analysis of the intrinsic evidence fails to resolve any ambiguity in the claim language may the court then rely on extrinsic evidence, such as expert and inventor testimony, dictionaries, and learned treatises. *See Vitronics*, 90 F.3d at 1583 ("In those cases where the public record unambiguously describes the scope of the patented invention, reliance on any extrinsic evidence is improper"). "Within the class of extrinsic evidence, the court has observed that dictionaries and treatises can be useful in claim construction." Phillips, 415 F.3d at 1318. While expert testimony can be useful to a court for a variety of purposes, conclusory assertions by experts are not useful to a court. *Id.* The court generally views extrinsic evidence

as less reliable than the patent and its prosecution history in determining how to read claim terms, even if though consideration is within the court's sound discretion. *See id.* at 1318-19.

### III. DISCUSSION

The parties dispute five terms contained in four of the Sun patents: (1) "first available memory space" in the '987 patent; (2) "means, responsive to the receipt of a stream of data records from said associated data processor, for writing said received stream of data records in available memory space in one of said disk drives" in the '855 patent; (3) "means responsive to said reading of said memory controller for periodically verifying the integrity of data currently stored in each of said identified dedicated partitions" and (4) "means for reading data in each of said identified dedicated partitions" in the '012 patent; and (5) "active links between said active nodes" in the '787 patent. Following the November 10, 2008, hearing, the parties submitted a Joint Table of Final Proposed Constructions for the Sun patents.

#### A. '987 Patent

The '987 patent, "Logical Track Write Scheduling System for a Parallel Disk Drive Array Data Storage Subsystem," is directed to a data storage system. The data storage subsystem stores data on a number of small hard disk drives, but emulates the format and operation of a large form factor disk drive: "This invention relates to an inexpensive, high performance, high reliability parallel disk drive array data storage subsystem that includes an efficient data storage management system to dynamically map virtual data storage devices to logical data storage devices and schedule the writing of data on these devices." '987 patent, col. 1:7-12.

#### 1. "First available memory space"

**Disputed Claim Term:** "first available memory space" ('987 patent, claims 9 and 57)

NetApp's construction	Sun's construction
"one or more logical tracks, each of which must be empty, <i>i.e.</i> , recognized by the system as available free space"	"one or more empty logical tracks each of which is empty, <i>i.e.</i> , one or more empty stripes each of which is empty"
	"empty" means "available to be written to"

[15] The parties agree on the wording of the construction of "first available memory space" as "one or more logical tracks, each of which is empty," but dispute whether "empty" means "recognized by the system as available free space," as NetApp contends, or "available to be written to," as Sun proposes. The dispute centers on the question of whether the obsolete data must go through a process of being tagged as obsolete and the logical track thus recognized as available free space in order for the track to be available to be written to.

To begin its analysis, the Court first turns to the claims themselves. The term "first available memory space" appears in claims 9 and 57 of the '987 patent:

9. In a disk memory system, having a plurality of disk drives, a number of said plurality of disk drives being configured into at least two redundancy groups, each redundancy group consisting of at least two disk drives, a method of storing data records for at least one associated data processor comprising the steps of:

selecting, in response to the receipt of a stream of data records from said associated data processor, **first available memory space** in one of said redundancy groups to store said received stream of data records thereon;

writing said received stream of data records and redundancy data associated with said received stream of data records in said selected **first available memory space** in said selected redundancy group;

writing, in response to the subsequent receipt of modifications to one of said data records stored in said **first available memory space** in said selected redundancy group from said associated data processor, said modified data record, exclusive of the rest of said received stream of data records and said redundancy data associated with said received stream of data records written in said **first available memory space**, in second available memory space in one of said redundancy groups by including said modified data record with a stream of data records subsequently received for said step of writing; and

converting said **first available memory space** used to store said originally received data record to available memory space.

...

57. In a disk memory system having a plurality of disk drives, a number of said plurality of said disk drives configured into at least two redundancy groups, each redundancy group consisting of  $n+m$  disk drives, where  $n$  and  $m$  are both positive integers, with  $n$  being greater than 1 and  $m$  being equal to or greater than 1, for storing data records for at least one associated data processor comprising the steps of:

storing, in response to the receipt of  $n$  streams of data records from said associated data processor, said  $n$  received streams of data records;

generating  $m$  redundancy segments using said  $n$  received streams of data records;

selecting one of said redundancy groups having **first available memory space**, addressable at the same relative address for each of said  $n+m$  disk drives, for storing said  $n$  received streams of data records and said  $m$  generated redundancy segments;

writing said  $n$  received streams of data records and said  $m$  redundancy segments on said  $n+m$  disk drives in said **first available memory space** in said selected redundancy group, each stream of data records and each redundancy segment at said relative address on a respective one of said  $n+m$  disk drives;

writing, in response to the subsequent receipt of modifications to one of said data records, stored at said relative address on one of said  $n+m$  disk drives in said **first available memory space** in said selected redundancy group, from said associated data processor, said modified data record, exclusive of the rest of said  $n$  received streams of data records and said  $m$  redundancy segments associated with said  $n$  received streams of data records written in said **first available memory space**, in second available memory space in one of said redundancy groups, exclusive of said relative address on said one of said  $n+m$  disk drives in said selected redundancy group by including said modified data record with  $n$  streams of data records subsequently received for said step of writing; and

converting said **first available memory space**, addressable at said relative address on said one of said  $n+m$  disk drives in said selected redundancy group, used to store said originally received data record to available memory space.

'987 patent, col. 23:40-24:3, 31:53-32:57 (emphasis added).

The claim language teaches that the "first available memory space" is a space in a redundancy group selected as the location for writing the received streams of data records and their associated redundancy

information. Looking to the specification for further guidance, the specification states that "[e]ach redundancy group, also called a logical disk drive, is divided into a number of logical cylinders," and "[e]ach logical track is comprised of N+M physical tracks, one physical track from each disk drive in the redundancy group." '987 patent, col. 3:50-54; 3:55-57.

In the "Solution" section of the specification, the patent describes the logical track write scheduling system of the claimed invention:

This system avoids the parity update problem of the prior art by never updating the parity in a data redundancy group. Instead, **all new or modified data is written on empty logical tracks and the old data is tagged as obsolete**. The resultant "holes" in the logical tracks caused by old data are removed by a background **free-space collection process that creates empty logical tracks** by collecting valid data into previously emptied logical tracks....

**The original, unmodified data is simply flagged as obsolete**. Obviously, as data is modified, the redundancy groups increasingly contain numerous virtual tracks of obsolete data. The remaining valid virtual tracks in a logical cylinder are read to the cache memory in a background "**free space collection**" process. They are then written to a previously emptied logical cylinder and the "**collected**" **logical cylinder is tagged as being empty**.

'987 patent, col. 3:33-41, 5:13-21 (emphasis added). The patent specification thus describes a write process whereby new or modified data is written on empty logical tracks, and old data is tagged as obsolete and removed by a background free-space collection process that creates empty logical tracks. In the free space collection process, the valid data is read to the cache memory, "then written" to a previously emptied logical cylinder and the collected logical cylinder is tagged as being empty, so that it can be written to. *Id.* at 5:13-21. *See also id.* at 5:57-63 (Detailed Description of the Drawing: "all new or modified data is written on empty logical tracks and the old data is tagged as obsolete. The resultant 'holes' in the logical tracks caused by old data are removed by a background free-space collection process that creates empty logical tracks by collecting valid data into previously emptied logical tracks."); 20:64-70 (Free Space Collection: "In order to create completely empty logical cylinders for destaging, valid data in partially valid cylinders must be read into cache memory 113 and rewritten into new previously emptied logical cylinders."). The plain language of the specification thus favors NetApp's proposed construction requiring the logical tracks to be "tagged as being empty" and thereby recognized by the system as available free space.

Sun contends that NetApp's proposed construction improperly reads the "free space collection process" into the claim term. The specification, however, describes "the present invention," whereby "holes" in the logical tracks caused by old data are removed by a free-space collection process that creates empty logical tracks. '987 patent, col. 5:50, 59-62. Reading the claim in light of the specification, the Court concludes that one of ordinary skill in the art would understand empty logical tracks to have undergone a process by which old data has been tagged obsolete and the logical track is recognized as empty. By contrast, Sun's proposed construction of "empty" as "available to be written to" is vague with respect to whether the old, unmodified data has been tagged as obsolete and the logical tracks thus recognized as being empty.

The parties also dispute whether the term requires writing only to empty logical tracks and never writing to logical tracks that are not empty. As Sun conceded, whenever claim 9 or 57 is practiced, those claims require writing to an empty logical track. However, the competing claims constructions proposed by the parties do not require the Court to decide an issue that emerged at the hearing: whether an accused product must practice exclusively the method of claims 9 and 57 to infringe or, as Sun characterizes the dispute, whether an accused system that fully practices the claimed method most of the time nevertheless does not infringe the patent if it ever practices a different method. The Court does, however, make the following observations. The express language of these method claims by itself does not compel NetApp's proposed

limitation that would require an accused product to always practice these claims in order to infringe. For example, claim 9 is a method claim as recognized by dependent claims 10 through 12 ("The method of claim 9 further comprising ..."). The claim preamble specifies the configuration of the system in which the method operates, i.e., "[i]n a disk memory system." NetApp suggested for the first time at the hearing that the preamble of claim 9 imports the limitation of claim 1 and other apparatus and memory system claims in the patent. Even assuming (without deciding) that the preamble is a limitation, *see* NetApp Suppl. Cl. Const. Br. at 4 n. 1, the preamble only describes a system of multiple disk drives configured into two or more redundancy groups each consisting of at least two disk drives, without stating whether the method must be exclusive. Furthermore, the claim 9 preamble does not refer to "[t]he system of claim 1," unlike, e.g., claims 2, 4, or 8. Sun is correct that an accused product need not always practice a method to infringe. *See* Bell Commun. Research, Inc. v. Vitalink Commun., 55 F.3d 615, 622-23 (Fed.Cir.1995).

Even though the claim language by itself would not compel NetApp's interpretation, there is strong support for it in the specification and prosecution history. The specification makes clear that the invention itself, unlike the prior art, never updates parity data in the redundancy group and avoids the negative impact on performance in the prior art, using strong and unequivocal language:

"A performance improvement is obtained by *eliminating redundancy data updates* in the redundancy group by writing modified virtual track instances into previously emptied logical tracks and marking the data contained in the previous virtual track instance location as invalid. Logical cylinders containing a mixture of valid and invalid virtual tracks are emptied by writing all the valid virtual tracks into a previously emptied logical cylinder as a background process." '987 patent, Abstract (emphasis added);

"This system avoids the parity update problem of the prior art by *never* updating the parity in a data redundancy group." *Id.* at 3:33-37 [Solution] (emphasis added);

"Thus, *all* redundancy data creation, writing and free space collection occurs in background, rather than on-demand processes. This arrangement avoids the parity update problem of existing disk array systems and improves the response time versus access rate performance of the data storage subsystem by transferring these overhead tasks to background processes." *Id.* at 5:21-28 [Solution] (emphasis added);

"The data storage subsystem of *the present invention* uses a plurality of small form factor disk drives in place of a single large form factor disk drive to implement an inexpensive, high performance, high reliability disk drive memory that emulates the format and capability of large form factor disk drives. This system avoids the parity update problem of the prior art by *never* updating the parity. Instead, all new or modified data is written on empty logical tracks and the old data is tagged as obsolete." *Id.* at 5:55-59 [Detailed Description of the Drawing] (emphasis added).

NetApp also relies on excerpts from the prosecution history to demonstrate that the applicants made a clear and unambiguous disavowal of claim scope by distinguishing the invention from the prior art disk array systems as "*never* updating parity in any redundancy group at any time." NetApp Br. re Pros. Hist. at 2-3 (quoting '987 Prosecution History, Sept. 23, 1991 Amendment at 7 (NAB0013338)) (emphasis added). The applicants' statements were made in response to the examiner's objection to the specification and rejection of the claims noting that the written description was "unclear why writing data to a new location is an advantage. This simply changes which parity bits need to be regenerated." *Id.* at 3 (NAB0013334). While the applicants' statements in the prosecution history arguably do not by themselves amount to a clear and unambiguous statement that a system must always practice the claimed method in order to practice the patent, they lend further support to the strong language in the specification describing the feature of the invention as a whole, and thus may well limit the scope of the invention to require always writing only to an empty track. At this juncture, however, the Court reserves this question.

The briefs submitted by the parties raised a dispute as to whether "first available memory space" could include more than one logical track. During oral argument, however, NetApp conceded that there was no dispute about whether "first available memory space" could be construed as one or more logical tracks, but sought clarification that each of the tracks must be empty. Sun's proposed construction, "each of which is empty," also reflects this understanding, and the Court construes the term to reflect that each logical track is empty.

With respect to Sun's proposal that "empty logical tracks" be further construed as "i.e., one or more empty stripes, each of which is empty," NetApp argued that every logical track is a stripe, but not every stripe is a logical track. NetApp was concerned that the jury would be confused by using the terms stripe and logical track interchangeably. Other than offering expert opinion that a logical track is the same as a stripe in the context of this patent, Sun has not demonstrated that logical track should always be construed as a stripe. The Court declines to include the phrase "i.e., one or more empty stripes."

For the reasons set forth above, the Court adopts NetApp's proposed construction of "first available memory space" as "one or more logical tracks, each of which must be empty, i.e., recognized by the system as available free space."

**B. '855 Patent**

The '855 patent, "Disk Drive Array Memory System Using Nonuniform Disk Drives," is directed to a data storage subsystem that emulates the format of a large form factor disk drive. This patent focuses on enabling the system to work with a set of nonuniform small form factor disk drives, so that the small form factor disk drives do not need to share common physical characteristics such as track capacity and disk rotational speed. '855 patent, col. 1:41-51, 1:58-2:22.

**2. "Means, responsive to the receipt of a stream of data records from said associated data processor, for writing said received stream of data records in available memory space in one of said disk drives"**

**Disputed Claim Term:** "means, responsive to the receipt of a stream of data records from said associated data processor, for writing said received stream of data records in available memory space in one of said disk drives" ('855 patent, claim 1)

<b>NetApp's construction</b>	<b>Sun's construction</b>
This term is subject to 35 U.S.C. s. 112, para. 6.	This term is subject to 35 U.S.C. s. 112, para. 6.
Function: "writing said received stream of data records in available memory space in one of said disk drives"	Function: "for writing said received stream of data records in available memory space in one of said disk drives"
Corresponding structure: "disk drive manager 102-1 and control unit 101" "responsive to": "after and in reaction to"	Structure: "processor 204 programmed to execute the software subroutine that consists of steps 706 to 711 illustrated in Figure 7 and described at col. 16:65 through col. 17:21 of the ' 855 patent specification"

	Sun does not believe the term "responsive to" needs to be construed. However, if the Court decides to construe this term, Sun contends NetApp's "after and in reaction" construction cannot properly be read to preclude beginning a write operation prior to the receipt of the last data record in the stream of data records.
--	--



[16] [17] The parties agree that this term is a means-plus-function limitation governed by 35 U.S.C. s. 112, para. 6. That paragraph states: "An element in a claim for a combination may be expressed as a means or step for performing a specified function without the recital of structure, material, or acts in support thereof, and such claim shall be construed to cover the corresponding structure, material, or acts described in the specification and equivalents thereof." Claim construction of a means-plus-function limitation requires that the court identify the function of the limitation, and ascertain the corresponding structure in the written description that is necessary to perform that function. *Altiris, Inc. v. Symantec Corp.*, 318 F.3d 1363, 1375 (Fed.Cir.2003).

The parties agree that the function is "writing said received stream of data records in available memory space in one of said disk drives," but dispute whether the corresponding structure includes the entire control unit 101 and disk drive manager 102-1, as NetApp proposes, or only processor 204 programmed to execute steps 706 to 711, illustrated in Figure 7, as Sun contends.

The disputed term appears in claim 1 of the '855 patent:

1. A data storage system, having a plurality of disk drives each having a single set of at least two predefined data storage characteristics, for storing data records for at least one associated data processor, comprising:

a controller, responsive to one of said disk drives having at least one of said at least two data storage characteristics different than a corresponding one of said at least two data storage characteristics of the remaining ones of said plurality of disk drives, for selecting a common disk drive format comprising a set of said data storage characteristics that are emulatable by all of said plurality of disk drives, comprising;

...

means, responsive to the receipt of a stream of data records from said associated data processor, for writing said received stream of data records in available memory space in one of said disk drives.

'855 patent, col. 21:7-38.

[18] [19] " 'Structure disclosed in the specification is "corresponding" structure only if the specification or prosecution history clearly links or associates that structure to the function recited in the claim.' " *Altiris*, 318 F.3d at 1375 (quoting *B. Braun Med. v. Abbott Labs.*, 124 F.3d 1419, 1424 (Fed.Cir.1997)). Section 112 paragraph 6 does not " 'permit incorporation of structure from the written description beyond that necessary to perform the claimed function.' " *Asyst Technologies, Inc. v. Empak, Inc.*, 268 F.3d 1364, 1369-70 (Fed.Cir.2001) (quoting *Micro Chem., Inc. v. Great Plains Chem. Co.*, 194 F.3d 1250, 1257-58 (Fed.Cir.1999)). "Structural features that do not actually perform the recited function do not constitute corresponding structure and thus do not serve as claim limitations." *Id.* at 1370 (citing *Chiuminatta Concrete Concepts, Inc. v. Cardinal Indus., Inc.*, 145 F.3d 1303, 1308-09 (Fed.Cir.1998)). By contrast, structure that is "integral to performing the stated function" does constitute corresponding structure. *Gemstar-TV Guide Intern., Inc. v. International Trade Com'n*, 383 F.3d 1352, 1362 (Fed.Cir.2004).

In the section entitled "Data Write Operations," the specification details the write operation, identifies the operational steps corresponding to the function of writing the stream of data records to a disk drive, and identifies certain structure as carrying out the steps. The parties agree that steps 706 to 711 are involved in performing the claimed function, as described in the specification:

This scheduling is accomplished by the subroutine that consists of steps 706-711. At step 706, the control unit 101 determines whether the virtual track instance as updated fits into an available open logical cylinder. If it does not fit into an available open logical cylinder, then at step 707 then this presently open logical

cylinder must be closed out and written to the physical layer and another logical cylinder selected from the most free logical device or redundancy group in the disk drive array data storage subsystem 100. At step 708, the selection of a free logical cylinder from the most free logical device takes place. This ensures that the data files received from host processor 11 are distributed across the plurality of redundancy groups in the disk drive array data storage subsystem 100 in an even manner to avoid overloading certain redundancy groups while underloading other redundancy groups. Once a free logical cylinder is available, either being the presently open logical cylinder or a newly selected logical cylinder, then at step 709, the control unit 101 writes the updated virtual track instance into the logical cylinder and at step 710 the new location of the virtual track is placed in the virtual to logical map in order to render it available to the host processors 11-12. At step 711, control returns to the main routine, where at step 712 the control unit 101 cleans up the remaining administrative tasks to complete the write operation and return to an available state at 712 for further read or write operations from host processor 11.

'855 patent, col. 16:65-17:25.

Sun contends that the specification clearly links only the subroutine of steps 706 to 711 to the actual function of writing the data records to the disks. However, the specification includes additional operational steps within the writing function. The control unit 101 performs the set up for a write operation (step 701) and assures that the virtual track containing the data record to be rewritten is located in the cache memory 113 (step 702). The control unit 101 transfers the modified record data received from the host processor into the cache memory 113 (step 704). The specification explains that once the virtual track is updated with the modified data record, "the control unit 101 must schedule this updated virtual track instance to be written onto a redundancy group somewhere in the disk drive array data storage subsystem 100." '855 patent, col. 16:61-64. The scheduling for writing the data to the disks is accomplished by steps 706 to 711, but the specification identifies steps 701 to 712 as necessary to perform the entire writing function. The plain language of this "Data Write Operation" section of the specification clearly and repeatedly links "the control unit 101" to performing the write operation.

Sun argues, however, that the structure does not include the entire control unit 101 because it has multiple components, other than processor 204 and the associated control software, that are not associated with the function of writing data records, such as channel interface units 201, channel interface control 202, channel data compression circuit 203 and optical device interface 205. '855 patent, col. 7:19-45. Sun contends that the only structure that is clearly linked to the function of writing data records is processor 204-0 of control unit 101, plus the associated software, that together write the data records. Sun relies on the following specification language to demonstrate the role of processor 204-0 in the writing function:

As can be seen from the architecture illustrated in FIG. 2, all data transfers between a host processor 11 and a redundancy group in the disk drive subsets 103 are routed through cache memory 113. **Control of cache memory 113 is provided in control unit 101 by processor 204-0.** The functions provided by processor 204-0 include initialization of the cache directory and other cache data structures, cache directory searching and management, cache space management, cache performance improvement algorithms as well as other cache control functions. **In addition, processor 204-0 creates the redundancy groups from the disk drives in disk drive subsets 103 and maintains records of the status of those devices. Processor 204-0 also causes the redundancy data across the N data-disks in a redundancy group to be generated within cache memory 113 and writes the M segments of redundancy data onto the M redundancy disks in the redundancy group. The functional software in processor 204-0 also manages the mappings from virtual to logical and from logical to physical devices.** The tables that describe this mapping are updated, maintained, backed up and occasionally recovered by this functional software on processor 204-0. The free space collection function is also performed by processor 204-0 as well as management and scheduling of the optical fiber backend channels 104. Many of these above functions are well known in the data processing art and are not described in any detail herein.

'855 patent, col. 7:67-8:28 (emphasis added). *See* Levy Decl. para. 32. In support of its argument that the corresponding structure is limited to processor 204 executing the software associated with steps 706 to 711, Sun relies on *Intel Corp. v. VIA Tech., Inc.*, 319 F.3d 1357, 1365-66 (Fed.Cir.2003), where the Federal Circuit construed the corresponding structure for the functions "to selectively write data" and "determine whether data is able to be written directly" as the "core logic" (a controller chip executing software) modified by a particular software protocol to perform Fast Write. *Id.* Sun suggests that the writing function here is similar to the writing function in *Intel*, and that the corresponding structure here is similarly limited to processor 204 executing the control software. However, the issue in *Intel* differed from the one presented here. There, the Federal Circuit held that the patent was not indefinite even though the specification did not disclose circuitry to show how the core logic was modified, but the parties did not dispute whether a more general purpose computer, rather than the processor, constituted the corresponding structure. *Id.* at 1366.

Other authorities on which Sun relies are no more compelling. In *WMS Gaming v. Int'l Game Tech.*, 184 F.3d 1339, 1348-49 (Fed.Cir.1999), the Federal Circuit held, "In a means-plus-function claim in which the disclosed structure is a computer, or microprocessor, programmed to carry out an algorithm, the disclosed structure is not the general purpose computer, but rather the special purpose computer programmed to perform the disclosed algorithm." 184 F.3d at 1349 and n. 4. In *WMS Gaming* the parties stipulated that the patent disclosed "a microprocessor, or computer, to control the operation of the slot machine, including the operation of the machine in the assignment of numbers to reel stop positions." *Id.* at 1347. The Federal Circuit noted that it did not find anything in the patent limiting the "means for assigning" to a microprocessor or computer, but declined to decide an issue that the parties did not present. *Id.* at 1347 n. 2. Similarly, in *Tehrani v. Hamilton Med., Inc.*, 331 F.3d 1355, 1362 (Fed.Cir.2003), the parties did not dispute whether the microprocessor formed part of the structure, but rather precisely which algorithm it was programmed to perform. The Federal Circuit noted, "We agree with the parties that the structure corresponding to the processing function is the disclosed microprocessor that is programmed to perform the disclosed algorithm. The specification teaches that the 'first means, which preferably comprises a programmable microcomputer, is controlled by a software algorithm to operate upon the input data and provide digital output data representing the amount and optimum frequency of ventilation required for the next breath.' ... The district court, however, has not determined the precise algorithm that is part of the recited structure." 331 F.3d at 1362.

While these cases demonstrate that Sun is correct that a processor executing a software routine may in appropriate circumstances be sufficient to identify corresponding structure, Sun's proposed structure is too narrow for performing the claimed function here because the specification identifies as integral to and clearly links broader structure to the function of writing data records. For example, the control unit performs the subroutine that consists of steps 706 to 711, as well as other operational steps to perform a data write operation. '855 patent, col. 16:65-17:25. Furthermore, the specification describes the integral role of the cache memory 113, which is contained in the control unit, in the data write operation: "all of the data updating is performed in the cache memory 113." '855 patent, col. 16:32-33. Thus, this structure is also integral to the write operation. Sun is correct that the specification describes the control unit 101 as performing other functions, such as reconstruction of data for a failed disk drive and switching a spare disk drive. *See* '855 patent, col. 5:52-55. However, the fact that the control unit 101 performs these additional control functions does not preclude the control unit 101 from serving as the structure necessary to perform the stated function of writing the stream of data records to a disk drive as the specification states. The specification identifies the control unit 101, not processor 204, with the control software associated with the data write operation, as performing the function of writing data records. *See, e.g., id.* at 18:64-67 ("Control software in control unit 101 queries all disk drives in the redundancy group to identify the cylinder, track and sector format of all the disk drives in the redundancy group.").

NetApp contends that disk drive manager 102 is also directly involved in performing the claimed function

and therefore part of the corresponding structure, citing the specification:

Thus, the control and drive circuits 121 in disk drive manager 102-1 perform the data and control signal interface and transmission function between the commodity disk drives of disk drive subset 103-1 and control unit 101.

'855 patent, col. 8:63-67. The specification demonstrates that the disk drive manager interconnects disk drives with optical fiber backend channels and includes an input/output circuit that provides a hardware interface to interconnect the optical fiber backend channels with the data paths that serve control and drive circuits. '855 patent, col. 5:61-65.

Whether the disk drive manager 102 (or at least its control and drive circuits 121) is part of the corresponding structure is determined by whether the writing function includes the data and control signal interface and transmission function. *See Gemstar-TV Guide Intern., Inc. v. International Trade Com'n*, 383 F.3d 1352, 1363 (Fed.Cir.2004). In *Gemstar*, the parties disputed whether the video switcher was part of the corresponding structure to perform the recited function of "displaying the television schedule on the television screen as a grid," or whether the corresponding structure comprised only the CPU and video display generator, and the video switcher was merely a conduit for coupling the means for displaying and the television. *Id.* at 1361. The Federal Circuit determined that the video switcher was "integral to performing the stated function," reasoning as follows:

the combination of a CPU, video display generator, and video switcher is required to perform the function of displaying the television schedule in a grid format on the television screen. Without the transmission of electrical signals by the video display generator to enable the video switcher, the television schedule would not be selectively displayed on the television screen and would not be displayed in grid format, as are required by the functional statement of the claim limitation.

383 F.3d at 1362.

In *Gemstar*, the Federal Circuit distinguished *Asyst Tech., Inc. v. Empak, Inc.*, 268 F.3d 1364, 1369-70 (Fed.Cir.2001), on which Sun relies here. In *Asyst*, the Federal Circuit determined that the functions of receiving and processing digital information recited in claim 1 of the '421 patent were performed by the local control processor, but that the corresponding structure did not include external cables or devices that were connected to the processor, reversing the district court's holding that the corresponding structure included the communication line 51 between the processor, or microcomputer means, and the communication means. 268 F.3d at 1370-71. The Federal Circuit noted in *Asyst* that the specification did not refer to line 51 in the description of the "receiving" and "processing" functions of the processor, and determined that the patent did not clearly link line 51 to the recited functions of claim 1. *Id.* While communication line 51 "enabled" the microcomputer means to perform the recited functions, it did not actually perform the functions. *Id.* As *Gemstar* explained, that construction in *Asyst* was compelled by the identification of the functions as receiving and processing the data because only the microcomputers performed these functions, not communication wires. 383 F.3d at 1362-63.

By contrast, the Federal Circuit also construed another means-plus-function element in a different claim in *Asyst*, namely, the "fourth means ... for controlling [the receipt of the transportable containers and the processing of the articles within the containers] and for transmitting information related to the processing performed [to the transportable container]." 268 F.3d at 1372. As to this "fourth means," the court held that communication line 51 was part of the corresponding structure because line 51 connected the local process controller that performed the first recited function, controlling activities on the workstation, with the communication means that performed the second function, transmitting information to the container. *Id.* The Federal Circuit concluded that because the means that performed those two functions consisted of the "entire

complex comprising the local process controller 20 and the communication means 50[,] ... it also necessarily encompasses structure that connects the two, i.e., communication line 51." Id. As the court noted, the language of this "fourth means" limitation was "significantly different" from the language of the microcomputer means limitation which the court construed to exclude the communication line 51. Id.

Gemstar and Asyst teach the importance of focusing on the exact function that the structure is performing. NetApp has not demonstrated that the disk drive manager 102 is necessary to accomplish the function of writing data in available memory space in a disk drive. Significantly, the "Data Write Operation" section of the specification never even mentions the disk drive manager 102. By contrast, in Gemstar, the Federal Circuit deemed the video switcher integral to performing the function of displaying the television schedule on a television screen where the written description indicated that the video switcher, in combination with a CPU and video display generator, was required to perform the stated function. 383 F.3d at 1361-62. Here, the disk drive manager 102 is necessary for performing interface and transmission function, '855 patent, col. 8:63-67, but not the writing function itself, which as described in the specification is performed by control unit 101. The disk drive manager 102 connects the disk drives with control unit 101, id., but only control unit 101 actually writes the data records. The data records are written in the disk drives, but the disk drives themselves do not perform the write operation. Unlike the "fourth means" construed in Asyst, which required a structure connecting two discrete components to perform two functions, 268 F.3d at 1372, the writing means claimed in the '855 patent does not assign functions to separate structures that must be linked. The disk drive manager 102 enables the control unit 101 to write data records to the disk drives, but neither the disk drive manager nor the disk drives perform that function. Id. at 1371.

During oral argument, NetApp conceded that there was no dispute that the term "one of said disk drives" did not require that all streams of data records received over time be written to the same disk drive. NetApp also conceded that the dispute over the phrase "responsive to the receipt of a stream of data records from said associative data processor" concerned a separate limitation and was not submitted here for claim construction. The Court further deferred ruling on the parties' dispute over the term "stream of data records," which the parties acknowledged was a separately identified term for claim construction.

For the reasons set forth above, the Court construes "means, responsive to the receipt of a stream of data records from said associated data processor, for writing said received stream of data records in available memory space in one of said disk drives" as a means-plus-function limitation, having the function of "writing said received stream of data records in available memory space in one of said disk drives" and the corresponding structure of "control unit 101 programmed to execute the steps 701 through 712 illustrated in Figure 7 and described under the heading 'Data Write Operation' at column 16:26 through column 17:25 of the '855 patent specification."

### C. '012 Patent

The '012 patent, "Disk Scrubbing System," is directed to a data storage subsystem that emulates the format and operation of a large form factor disk drive. The patent addresses the process of "disk scrubbing," or verifying the integrity of data written to disk drives. What is newly claimed in the '012 patent is identifying and selecting memory locations containing customer or redundancy data and verifying those locations on a priority basis. '012 patent, col. 2:61-67. In the claimed invention, the disk drive array is divided into a plurality of logical partitions, including "dedicated partitions" that are storing data and "free partitions" that are available to store data in the data storage subsystem. The claimed invention determines whether the partitions are dedicated or free. Id. at col. 17:38-63, col. 20:40-62. The parties dispute two claim terms in the '012 patent.

### **3. "Means responsive to said reading of said memory controller for periodically verifying the integrity of data currently stored in each of said identified dedicated partitions"**

**Disputed Claim Term:** "means responsive to said reading of said memory controller for periodically verifying the integrity of data currently stored in each of said identified dedicated partitions" (Claim 1, '012 patent)

NetApp's construction	Sun's construction
Periodically verifying the integrity of data currently stored in each of said identified dedicated partitions: "periodically verifying only the integrity of data currently stored in each of said identified partitions" Responsive to: "after and in reaction to"	Sun contends this term is not a means-plus-function limitation. If the Court concludes this is a means-plus-function limitation, then the function is: "periodically verifying the integrity of data currently stored in each of said identified dedicated partitions."
Corresponding structure of nested "means" terms:	Sun contends NetApp's construction of the functional language incorrectly adds an only limitation and that the term "only" should not be added to the existing claim structure.
- "means ... for reading data": "control unit 101 and > disk drive manager 102-1" - "means ... for generating error check information": > "disk drive subassembly"  - "means ... for detecting errors": "disk drive > subassembly"	If the Court concludes this is a means-plus-function limitation, then the corresponding structure is: "processor 204 of control unit 101 programmed to execute the software processes that read data from selected track (item 1202 of Figure 12), generate CRC for read data (item 1203 of Figure 12), determine whether generated CRC and read CRC
-> "means ... for correcting said data containing errors": "control unit 101"	match (item 1209 of Figure 12), and reconstruct data of selected track having mismatched CRC (item 1008 of Figure 10)."

[20] While Sun initially disputed whether this term is a means-plus-function limitation, both parties now recognize that it is a means-plus-function nesting term for the four nested means-plus-function limitations. Although the parties did not address this issue in the briefs, the Court may ultimately need to conduct an analysis under 35 U.S.C. s. 112, para. 6 to determine whether the structures corresponding to each of the nested terms are together sufficient to perform the claimed function of this term. *See* Rodime PLC v. Seagate Tech., Inc., 174 F.3d 1294, 1303-04 (Fed.Cir.1999).

The parties dispute whether the function is limited to periodically verifying "only" the integrity of data currently stored in each of said identified partitions; construction of the "responsive to" limitation; and the corresponding structure.

The disputed claim terms of the '012 patent appear in claim 1:

1. A data storage subsystem that receives data from at least one connected data processor and stores said data on a plurality of disk drives divided into a plurality of logical partitions, said logical partitions comprising dedicated partitions currently storing data and free partitions available to store data, said data storage subsystem comprising:

a memory controller independent of said plurality of disk drives comprising means for identifying said dedicated and said free partitions;

means for reading said memory controller to determine the identity of said dedicated partitions; and

means responsive to said reading of said memory controller for periodically verifying the integrity of data currently stored in each of said identified dedicated partitions, said means for periodically verifying

comprising:

means for reading data in each of said identified dedicated partitions;

means responsive to said reading of data from each said identified dedicated partition for generating error check information from said read data;

means responsive to said generation of said error check information for detecting errors in data in each said identified dedicated partition; and

means responsive to said detection of errors for correcting said data containing errors.

'012 patent, col. 17:38-63.

NetApp contends that the function requires verifying **only** the integrity of partitions containing active, or current, data. The language of claim 1 itself distinguishes between "dedicated partitions currently storing data" and "free partitions available to store data." Claim 1 recites the means-plus-function limitation "for periodically verifying the integrity of *data currently stored in each of said identified dedicated partitions.*" '012 patent, col. 17:40-43, 50-52 (emphasis added). Turning to the specification, NetApp is correct that the specification refers only to verifying, or scrubbing, the active memory:

This system avoids the data integrity problems of the prior art by periodically verifying the integrity of the data stored on the disk drives of the data storage system. This is accomplished by one or more background processes that cycle through predetermined segments of active memory to verify the integrity of the data stored on the disk drives of the data storage subsystem.

'012 patent, col. 2:18-21.

As Sun points out, claim 1 must be construed broadly enough to permit the practice of dependent claims, including dependent claim 9, which recites a means for verifying administrative data in free partitions:

9. The data storage subsystem of claim 1 further comprising:

...

means responsive to said reading of said memory controller for periodically verifying the integrity of administrative data stored in each of said identified free partitions, ...

'012 patent, col. 18:63-66. NetApp contends that the doctrine of claim differentiation does not apply here to construe independent claim 1 so broadly because dependent claim 9 does not refine the verifying means recited in independent claim 1, which is focused only on the means for processing dedicated partitions, but rather adds to claim 1 by requiring another means for processing free partitions. NetApp argues that it would not be inconsistent to limit the "means for periodically verifying" in claim 1 to dedicated partitions and allow claim 9 to recite a separate means for processing free partitions. Although NetApp is correct that the "further comprising" language of claim 9 indicates that a means is being added to the data storage subsystem of claim 1, claim 9 is dependent on claim 1, and thereby incorporates by reference all the limitations of that claim. *See Monsanto Co. v. Syngenta Seeds, Inc.*, 503 F.3d 1352, 1357-58 (Fed.Cir.2007) (holding that "further comprising" claim was in dependent form and incorporated the limits of the overarching independent claim), *petition for cert. dismissed*, --- U.S. ----, 129 S.Ct. 394, 171 L.Ed.2d 809 (2008). Thus, NetApp's proposed "only" limitation would preclude the verifying means from scrubbing the administrative data stored in a free partition, thereby effectively reading claim 9 out of the patent. Claim 1 uses the open

term "comprising," thus allowing for additional structures to be added to the data storage subsystem that the data subsystem of claim 1 does not itself include. Although claim 1 does not claim the function of verifying the administrative data in the free partitions, neither does claim 1 preclude adding that functionality in a dependent claim as in claim 9.

[21] [22] [23] NetApp also relies on the patent prosecution history, in which the applicants distinguished the claimed invention from the prior art by stating that the claimed invention scrubbed only active data. The Court may rely on the prosecution history to inform the meaning of the claim language: "Like the specification, the prosecution history provides evidence of how the PTO and the inventor understood the patent. Furthermore, like the specification, the prosecution history was created by the patentee in attempting to explain and obtain the patent. Yet because the prosecution history represents an ongoing negotiation between the PTO and the applicant, rather than the final product of that negotiation, it often lacks the clarity of the specification and thus is less useful for claim construction purposes." *Phillips v. AWH Corp.*, 415 F.3d 1303, 1317 (Fed.Cir.2005) (citations omitted). The Federal Circuit has "declined to apply the doctrine of prosecution disclaimer where the alleged disavowal of claim scope is ambiguous." *Omega Eng'g, Inc. v. Raytek Corp.*, 334 F.3d 1314, 1324 (Fed.Cir.2003). "But where the patentee has unequivocally disavowed a certain meaning to obtain his patent, the doctrine of prosecution disclaimer attaches and narrows the ordinary meaning of the claim congruent with the scope of the surrender." *Id.*

To demonstrate that the applicants distinguished the invention over the prior art as verifying only partitions containing active data, NetApp cites the Preliminary Amendment dated March 6, 1996, which was submitted before claims 1 and 9 and the terms "dedicated partitions" and "free partitions" were added to the application: "the claimed invention periodically verifies the integrity of (e.g. scrubs) only logical partitions identified as containing '**active**' data." '012 Patent File History, 3/6/96 Amendment at 4 (emphasis in original). The Preliminary Amendment further states, "The prior art therefore teaches away from a *fundamental concept of the claimed invention* which improves performance of the scrubbing process in a data storage subsystem by scrubbing (periodically verifying the integrity of) **only** those logical partitions which are identified as containing **active data**." *Id.* at 6 (italicized emphasis added; bolded emphasis in original). NetApp argues that "one of ordinary skill in the art would have understood from the patentees' statements to the Patent Office that this invention, and specifically this function, requires scrubbing *only* the dedicated partitions." Ganger Decl. para. 34. While the applicants stated that "the specification at page 29, line 13-24, recites (in discussing figure 10) that free cylinders (those not containing **active data**) are skipped in the scrub processing," in the same submission to the PTO they also spelled out a single, limited exception: "Figure 10 so indicates that freed cylinders are not scrubbed ( *other than the last track thereof which contain administrative data* ) by the scrub processes of the claimed invention." 3/6/96 Amendment at 5 (emphasis added). Taken in context, these statements in the Preliminary Amendment demonstrate a clear and unmistakable disavowal of scrubbing free partitions with the sole exception of scrubbing the track containing administrative data.

Further, in the October 1, 1996 Amendment, the applicants reacted to the examiner's objections and rejection of the claims by "not reciting 'active data' in the new and amended claims," among other things. '012 Patent File History, 10/1/96 Amendment at 15. The applicants added claim 35, which was issued as claim 1, and dependent claim 40, which was issued as claim 9, and stated:

The present invention periodically verifies the data stored in the system. In Belsan, the data is only error checked when a read or write event occurs. The present invention teaches "periodically verifying" selected data, not the *sequential* "periodically verifying" of *every memory location* in a memory system of Bowden.

Dependent claims 40-44 have been added to recite the "periodically verifying" of "administrative data" disclosed in the specification.... When a "free partition:" is identified, the "administrative data" in the "free partition" is then identified. The "administrative data" is then "periodically verified." These claims are added



to recite the process disclosed in the specification but not included in the new base claims.

Id. at 15-16. Specifically, the applicants added dependent claim 9, explaining that it is dependent on, " *but not included in,*" claim 1. Id. at 16 (emphasis added). These statements show that the applicants did not include scrubbing administrative data in free partitions in claim 1, but rather preserved it as the sole exception to otherwise scrubbing only dedicated partitions. Sun is correct that this portion of the prosecution history shows that the applicants were distinguishing the prior art in part by selecting memory locations for verifying, as distinguished from verifying sequentially as in Bowden or verifying when a read/write event occurs as in Belsan. However, the applicants drew two separate distinctions, both stating that the invention " 'periodically verifies' *selected* memory locations containing data," id. at 15, and that another aspect of the invention-which they described as "fundamental"-was only scrubbing active data (with one exception), as set forth above. Thus, the prosecution history as a whole shows a clear disavowal of claim scope beyond verifying only dedicated partitions, with the sole exception of verifying free partitions containing administrative data.

However, because NetApp's proposed construction could be understood to preclude that exception, the Court declines to adopt as written NetApp's proposed "only" limitation of the function "periodically verifying the integrity of data currently stored in each of said identified dedicated partitions." There may be alternative language that would better convey this distinction.

Unlike the "responsive to" language in the '855 patent, NetApp has not expressly conceded that the phrase "responsive to said reading of said memory controller" in the '012 patent is a separate limitation. NetApp further contends, as it did with respect to the '855 patent claim term, that "responsive to" means "after and in reaction to," in that the claim describes a system that determines the identity of the dedicated partitions by reading the memory controller, and then, "after and in reaction to" identification of the dedicated partitions, it verifies the integrity of the dedicated partitions identified by reading the memory controller. *See Ganger Reply Decl.* para. 13.

Sun acknowledges that the '012 claim recites that the means must be "responsive to said reading of said memory controller," but contends that this limitation is not part of the identified function. Sun relies on *Medtronic, Inc. v. Guidant Corp.*, 2004 WL 1179338 at \*16-17 (D.Minn. May 25, 2004), where the district court construed the claim language "stabilizing means *responsive to said* atrial fibrillation *detecting means* for stabilizing the cardiac rate of the heart when the atria of the heart are in need of cardioversion," to identify the function as "stabilizing the cardiac rate of the heart when the atria of the heart are in need of cardioversion," without the "responsive to" claim language.

The only authority cited by NetApp for its proposed construction to include the "responsive to" limitation in the function is the Court's Order Construing Claims in the related case, *Network Appliance Inc. v. Sun Microsystems Inc.*, C07-6053 EDL. There, the Court construed the terms "in response to writing a data record to said one redundancy group"/ "responsive to writing a data record to one of said redundancy groups" and "in response to the receipt of a stream of data records from said data processor"/ "responsive to the receipt of a stream of the data records from said data processor" in the ' 857 patent. The "responsive to" phrase was central to the meaning of the disputed terms that were construed by the Court in that case. Here, by contrast, the parties have submitted for construction a means-plus-function limitation where the function is introduced by the claim term "for," i.e., "means ... for periodically verifying the integrity of data currently stored in each of said identified dedicated partitions." *See Greenberg v. Ethicon Endo-Surgery, Inc.*, 91 F.3d 1580, 1584 (Fed.Cir.1996) ("the use of the term 'means' (particularly as used in the phrase 'means for') generally invokes section 112(6)"). The "responsive to" phrase here does not identify any function performed by the structure, but rather identifies a separate limitation.

The Court is persuaded that the "responsive to" phrase in the '012 patent term is a claim limitation, but is not

part of the function. The Court therefore declines to construe the function as "responsive to said reading of said memory controller, periodically verifying the integrity of data currently stored in each of said identified dedicated partitions."

The parties do not dispute that the corresponding structure includes the control unit 101, FN1 as the specification confirms: "[t]he periodic scrub process is based on the logical architecture imposed by control unit 101, which establishes redundancy groups and conducts memory management;" the control unit 101 "keeps track of free and dedicated space in the storage subsystem and creates empty logical cylinders through the free space collection process." '012 patent, col. 5:14-19, 11:64-65. In briefing, NetApp proposed that the corresponding structure includes not only the control unit 101, but also the disk drive subassembly and the scrubbing procedures. In response to the Court's concern that NetApp's proposed "scrubbing procedures" was too vague for construing the corresponding structure, NetApp now proposes that the corresponding structure comprises the structure necessary to perform the nested means-plus-function limitations, namely, the control unit 101, the disk drive manager 102-1 and the disk drive subassembly. Because the disk drive manager is part of the disk drive subassembly, *see* '012 patent, Fig. 1, the Court first considers whether the entire disk drive subassembly, including the disk drive subset, is part of the corresponding structure.

The specification describes the role of the disk drive subassembly in reading and verifying data: "The hardware and software in the disk drive subassembly typically performs the track read and verify operation." '012 patent col. 15:48-50. The specification also recites, "The reading of the data from the disk drives and the generation of the CRC [cyclic redundancy check] data is accomplished by the disk drive subsystem." '012 patent, col. 13:6-8. Sun contends that the disk drive subassembly only typically performs the verifying function, and that this disclosure provides that the relevant operations may, in alternative embodiments, be performed either by the control unit, or within the disk drive subassembly and reported back to the control unit to be reviewed. Although Sun argues that the disk drive subassembly is not necessary to perform the recited function, Sun does not cite patent language that discloses these alternative embodiments. Moreover, the specification clearly links the disk drive subassembly to the verification function. At best, Sun argues based on extrinsic evidence that a person of ordinary skill in the art would know that the scrubbing process could be performed by control unit 101 instead of the disk drive subsystem, because the patent describes the role of the control unit 101 in running the software for the disk scrubbing architecture, '012 patent, col. 12:31-13:31, and executing the steps of the track scrubbing process, *id.* at col. 15:45-16:46, as illustrated in Fig. 10. Sun acknowledges, as it must, that the disk drive subassembly is an alternative (indeed, typical) method for verifying data, but contends that NetApp's proposed construction would require the disk drive subassembly to be used even though control unit 101 could perform the function. To account for alternate methods for performing the recited function, the Court will construe the corresponding structure as control unit 101 and the disk drive subassembly.

In *Callicrate v. Wadsworth Mfg., Inc.*, 427 F.3d 1361, 1368-70 (Fed.Cir.2005), the Federal Circuit held that the district court's construction of corresponding structure did not account for additional embodiments for performing the cutting function disclosed in the patent. The district court identified "a pivotally mounted cutting mechanism as the disclosed structure that performs the cutting function," but the Federal Circuit found that the patent disclosed additional embodiments: "[a]ny device for cutting the band 82 may be used [,]" including a cutting assembly 80 with a razor 88 slidably mounted within a housing 90 or hand-held cutting tools such as scissors and hand-held razors. 427 F.3d at 1369. Sun argues that unlike *Callicrate*, where the court found alternative structures, the '012 patent does not require the disk drive subassembly because the control unit 101 is sufficient. The Court does not find *Callicrate* distinguishable on this point. The specification identifies the disk drive subsystem as a method for reading and verifying data from the disk drive, '012 patent, col. 13:6-8. At most, Sun has demonstrated that the control unit 101 works with, but is not a complete alternative to, the disk drive subassembly to perform the track read and verify operation. The Court concludes that the corresponding structure includes the disk drive subassembly.

Sun proposes that the software associated with the verifying function be limited to steps 1202 (read data from selected track), 1203 (generate CRC for read data), 1209 (determine whether generated CRC and read CRC match), and 1008 (reconstruct data of selected track having mismatched CRC), as illustrated in Figure 10, which "illustrates in flow diagram form the operation of the logical cylinder scrubbing operation," and Figure 12, which "illustrates in flow diagram form the operation of the track scrubbing procedure." The Court notes that the specification describes three of these steps, but does not describe step 1209. *See* '012 patent, col. 15:45-46, 51-66 ("at step 1202 the data is read from the selected track;" "data that is read from the selected track is processed by generating a cyclic redundancy code CRC over this data at step 1203;" and "[i]f there is an error noted in this track, processing advances to step 1008 where the reconstruct track process is activated.") NetApp contends that these steps merely recite the function, and do not disclose a structure or an algorithm. However, as the Court has recognized with respect to the '855 patent, the corresponding structure of a means-plus-function limitation may consist of a processor executing a software routine. *Intel*, 319 F.3d at 1365-66; *WMS Gaming*, 184 F.3d at 1348-49 and n. 4.

[24] NetApp identifies periodic disk scrubbing procedures 501, 502, which include operational steps for periodic scrubbing other than the four steps identified by Sun, but also perform functions other than the function of periodically verifying the integrity of data currently stored in each of said identified dedicated partitions:

Each of the independent disk scrubbing procedures 501, 502 ... manages disk scrubbing for one of said device range partitions on all configured volumes. Each of the at least one periodic disk scrubbing procedures 501, 502 manages disk scrubbing for one of these partitions while another operationally independent priority disk scrubbing procedure 510 is available to sequence through the entire memory space, on all devices of all configured volumes in which data has newly been written since the last scrubbing operation was performed on that memory space. If the workload so dictates, a plurality of concurrently operational priority scrubbing procedures 510 can be used, but for simplicity of description, the preferred embodiment discloses n+1 processes, consisting of n periodic disk scrubbing procedures 501, 502 and 1 priority disk scrubbing procedure 510.

'012 patent, col. 13:49-63. The specification further describes the disk scrubbing procedures 501, 502, as illustrated in Figures 9 through 12:

The periodic disk scrubbing procedures 501, 502 function to sequence through all of their device range partitions on all of the configured volumes for a selected cylinder before proceeding to initiate the scrubbing operation for the next cylinder. This can be seen by reference to the memory map of Fig. 5, as described above. The operation of the periodic disk scrubbing procedures 501, 502 is illustrated in flow diagram form in Fig. 9, with additional details being disclosed in Figs. 10-12. The periodic disk scrubbing procedure (ex-501) begins a[sic] step 901 where the cylinder number is set to 0. At step 902, the volume number is set to 0.

'012 patent, col. 14:20-31. In light of the claims and the specification, NetApp's proposal to include the software associated with disk scrubbing procedures 501, 502 as part of the corresponding structure is too broad because the "scrubbing procedures" include functions outside the scope of the periodic verifying function. " 'When construing the functional statement in a means-plus-function limitation, we must take great care not to impermissibly limit the function by adopting a function different from that explicitly recited in the claim.' " *Omega Eng.*, 334 F.3d at 1322 (quoting *Generation II Orthotics, Inc. v. Med. Tech., Inc.*, 263 F.3d 1356, 1364-65 (Fed.Cir.2001)).

The Court agrees with Sun that not every aspect of Figures 9 through 12 is required to perform the recited function. First, the specification refers to disk scrubbing procedures 501, 502 as managing the disk scrubbing, not verifying the integrity of data. *Id.* at 13:49-63. Second, at step 905, where the selected logical

cylinder is scrubbed, as detailed at Figure 10, the operation includes a determination whether the selected logical cylinder is a free cylinder, a function that is not covered by the disputed "periodically verifying" term. Third, Sun offered expert testimony that, for example, step 904 ("Are scrubbing disable flags set for selected logical cylinder?") was not necessary to perform the verifying function, and that scrubbing disable flags could be set for reasons other than periodically verifying, i.e., either for copying in the free space collection, which is not related to scrubbing, or for priority scrubbing, which is a separate instance. *See id.* at col. 15:1-3 ("At step 904, the initialization procedures are completed by determining whether the selected logical cylinder should be scrubbed."). Inclusion of steps 501, 502 and Figures 9 through 12 in the corresponding structure would impermissibly import limitations that are not recited in the claim term.

While NetApp's proposed structure is overbroad, Sun's proposed structure is too restrictive in light of the specification. In the section entitled "Track Scrubbing Process," the specification refers to Figure 12 detailing the operation of the track scrubbing procedure 501, initiated at step 1201 and exiting at step 1208. '012 patent, col. 15:45-62. This section of the specification recites steps in the verifying process in addition to the four steps identified by Sun:

At step 1204, the previously generated CRC is read from the track and at step 1205 compared to the presently generated CRC. If the two sets of CRC data match, then there are no errors in the data read from the selected track and processing advances to step 1206 where a successful scrub is acknowledged and processing exits at step 1208. If there is a mismatch between the two sets of CRC data, the presence of an error is noted at step 1207 and processing exits at step 1208.

*Id.* at 15:54-62. At the very least, the corresponding structure includes the control unit 101 programmed to carry out the steps corresponding to the recited function as described in this section of the specification.

The specification further refers to steps 1006 to 1014 in the track scrubbing process:

At step 1007 of the logical cylinder scrubbing process, the results of the track scrubbing operation are reviewed. If there is an error noted in this track, processing advances to step 1008 where the reconstruct track process is activated. This operation uses the volume number, device position number, cylinder number and track number to specify the location of the data read error and thereby attempts to reconstruct correct track data from redundancy data on the corresponding tracks of the other drives in the redundancy group. At step 1010 the results of the reconstruction operation are reviewed and, if successful, the selected track is re-scrubbed by returning to step 1006. This process repeats until the track scrubbing operation is successful or at step 1010 an unrecoverable error is detected. The presence of an unrecoverable error is logged at step 1012 and processing advances to step 1009, as does the result of a successful track scrubbing operation as noted at step 1007. At step 1009, a determination is made of whether this is the last track in the selected logical cylinder. If not, at step 1011 the track number is incremented and processing returns to step 1006. This process, steps 1006-1011, repeat until the last track in the selected logical cylinder is reached, at which time processing advances to step 1013 where it is determined whether this device position is the last device position in the device range partition. If it is not, the device position number is incremented and processing returns to step 1003. Steps 1003-1013 are repeated until all of the device positions in the device range partition are scrubbed for this selected logical cylinder number, at which time processing exits at step 1014 and processing returns to step 906.

'012 patent, col. 15:63-16:25 and Fig. 10. This section refers not only to step 1008, which Sun contends, and the Court agrees, relates to the corresponding structure, but also to other steps of the logical cylinder scrubbing process. Neither party directly addressed the question of whether any of these other steps is necessary for performing the function of periodically verifying. Thus, the Court construes the corresponding structure to include steps 1202 to 1207 and 1008, while leaving open the question of whether additional steps shown in Figure 10 also come into play here.

For the reasons set forth above, the Court construes the term "means responsive to said reading of said memory controller for periodically verifying the integrity of data currently stored in each of said identified dedicated partitions" as a means-plus-function limitation, having the function of "periodically verifying the integrity of data currently stored in each of said identified dedicated partitions," and the corresponding structure of "the control unit 101 and the disk drive subassembly programmed to execute steps 1202 to 1207 illustrated in Figure 12 and described at col. 15:47-62, and step 1008 illustrated in Figure 10 and described at col. 15:63-16:25."

#### 4. "Means for reading data in each of said identified dedicated partitions"

**Disputed Claim Term:** "means for reading data in each of said identified dedicated partitions" (Claim 1, '012 patent)

NetApp's construction	Sun's construction
This term is subject to 35 U.S.C. s. 112, para. 6.	This term is subject to 35 U.S.C. s. 112, para. 6.
Function: "reading data in each of said identified dedicated partitions," meaning "reading only the data stored on each of said identified dedicated partitions"	Function: "for reading data in each of said identified dedicated partitions" Sun contends NetApp's construction of the functional language incorrectly adds an "only" limitation.
Corresponding structure: See above regarding the first nested means-plus-function term. ["control unit 101 and disk drive manager 102-1"]	Structure: "processor 204 of control unit 101 programmed to execute the software process that reads data from selected track (item 1202 of Figure 12)."

[25] The parties agree that this term is the first of the four nested terms in claim 1 of the '012 patent, that this is a means-plus-function limitation, and that the function is "reading data in each of said identified dedicated partitions." As with the "periodically verifying" term construed above, NetApp contends that this term should be construed as "reading **only** the data stored on each of said identified dedicated partitions." NetApp argues that this limitation is appropriate based on the patentees' explicit disavowal of claim scope, citing *Astrazeneca AB v. Mutual Pharmaceutical Co.*, 384 F.3d 1333, 1339-40 (Fed.Cir.2004) ("Where the general summary or description of the invention describes a feature of the invention (here, micelles formed by the solubilizer) and criticizes other products (here, other solubilizers, including co-solvents) that lack that same feature, this operates as a clear disavowal of these other products (and processes using these products)."). For the reasons stated above regarding reading administrative data in the free partitions, the Court declines to adopt NetApp's proposed language regarding "only" the data stored on the identified dedicated partitions.

As to the corresponding structure, Sun proposes that the structure is the processor 204 of control unit 101 and the software associated with executing step 1202 to read data from selected track; NetApp contends that control unit 101 and disk drive manager 102-1 are necessary to perform the function. NetApp reiterates its argument that step 1202 is not software but only recites the function, "READ DATA FROM SELECTED TRACK." NetApp contends that step 1202 is not an algorithm and does not recite what steps are used to read the data, and is therefore insufficient to comprise the structure. As the Court has ruled with respect to the previous two terms, corresponding structure of a means-plus-function limitation may consist of a processor executing a software routine. *Intel*, 319 F.3d at 1365-66; *WMS Gaming*, 184 F.3d at 1348-49 and n. 4. Here, the software routine for reading data is identified in the specification as step 1202. '012 patent, col. 15:49 ("at step 1202 the data is read from the selected track").

As discussed above with respect to the nesting term "periodically verifying," the specification states plainly that "[t]he reading of the data from the disk drives ... is accomplished by the disk drive subsystem," and that "[t]he hardware and software in the disk drive subassembly typically performs the track read and verify operation." '012 patent, col. 13:6-8, 15:49-51. Unlike the "means ... for writing" term of the '855 patent

construed above, the structure that corresponds to the "means ... for reading" of the '012 patent includes the disk drive manager because the specification identifies it as the structure that typically performs the read operation. FN2

The section of the specification describing the "Track Scrubbing Process" does not refer to the control unit 101 or processor 204 in describing the read operation. However, the "Data Read Operation" section of the specification refers to "the operational steps taken by processor 204 in control unit 101 of the data storage subsystem 100 to read data from a data redundancy group 122-1 to 122-n+m in the disk drive subsets 103." '012 patent, col. 9:19-22. That section of the specification further describes the role of the control unit 101 in the operational steps taken to perform a data read operation, illustrated in Figure 6:

At step 601, the control unit 101 prepares to read a record from a virtual track. At step 602, the control unit 101 branches to the cache directory search subroutine to assure that the virtual track is located in the cache memory 113 since the virtual track may already have been staged into the cache memory 113 and stored therein in addition to having a copy stored on the plurality of disk drives (122-1 to 122-n+m) that constitute the redundancy group in which the virtual track is stored. At step 603, the control unit 101 scans the hash table directory of the cache memory 113 to determine whether the requested virtual track is located in the cache memory 113. If it is, at step 604 control returns back to the main read operation routine and the cache staging subroutine that constitutes steps 605-616 is terminated.

'012 patent, col. 9:31-44. Raising arguments similar to those raised with respect to the "means ... for writing" term of the '855 patent, Sun contends that control unit 101 performs all major data storage subsystem functions, not only reading data, and that not all components of control unit 101 perform the reading data function. Sun's expert, however, offered his opinion that "[t]he specification makes it clear to one ordinary skill in the art that it is control unit 101 that performs the function of reading data from the dedicated partitions." Levy Decl. ISO Sun Resp. Br. para. 59. He further opined that "the portion of the control unit 101 ( *e.g.*, processor 204) running the software represented by element 1202 is the proper structure." *Id.* para. 61 (emphasis added). Thus, even Sun's expert does not limit the corresponding structure to the particular embodiment represented by Figure 6, showing the operational steps taken by processor 204. '012 patent, col. 9:19-22 ("the operational steps taken by processor 204 in control unit 101 of the data storage subsystem 100 to read data from a data redundancy group 122-1 to 122-n+m in the disk drive subsets 103") and Fig. 6. Based on the specification, read in light of the claims, and with the benefit of expert testimony, the Court will construe the corresponding structure as the disk drive manager and the control unit 101 programmed to execute step 1202.

For the reasons set forth above, the Court construes the term "means for reading data in each of said identified dedicated partitions" as a means-plus-function limitation, having the function of "reading data in each of said identified dedicated partitions" and the corresponding structure of "the disk drive manager and the control unit 101 programmed to execute step 1202 illustrated in Figure 12 and described at col. 15:49."

#### **D. '787 Patent**

The '787 patent, "Highly Available Cluster Message Passing Facility," is directed to a virtual disk system composed of a cluster of nodes and storage devices, such as disks. The system described in the '787 patent includes a collection of interconnected nodes and storage devices. The patent defines a node as "a group of circuitry designed to perform one or more computing tasks" and a storage device as "a persistent device capable of storing large amounts of data." '787 patent, col. 1:13-15, 29-31. Two or more interconnected nodes that have the capability of exchanging data are called a "cluster." *Id.* at 1:16-18. In the patented storage system, physical disk drives need not be directly attached to system nodes, but can be indirectly attached to system nodes via data communication links. *Id.* at 1:37-45. The resulting clusters of physical disks appear to a user as a collection of equally accessible logical or virtual disks. "Whether a storage

device is directly accessed or accessed via another node is transparent to the client accessing the storage device." '787 patent, col. 2:57-59.

### 5. "active links between said active nodes"

**Disputed Claim Term:** "active links between said active nodes" (Claims 16 and 27, '787 patent)

<b>NetApp's construction</b>	<b>Sun's construction</b>
"Active point-to-point links or active physical connections to a multi-drop link between devices in a cluster that are not storage devices"	"operational paths capable of transferring data between operational nodes capable of sending and receiving data"
Active: "designated as a member of the cluster"	

[26] The parties dispute: whether the links referred to in this claim are limited to point-to-point and multi-drop connections or may also encompass switches between the nodes; how the terms "active" nodes and links should be construed; and whether "nodes" should be defined to exclude storage devices. Although Sun initially contended that the term did not require construction because the term is clear on its face, it appears from oral argument and the final proposed construction that Sun has waived this argument. The Court agrees that this term needs construction.

The disputed term appears in Claims 16 and 27 of the '787 patent:

16. A method of transporting data in a distributed computing system comprising a plurality of nodes and a data communication bus, the method comprising:

determining physical resources in said distributed computing system, wherein said physical resources include active nodes of said distributed computing system and **active links between said active nodes;**

establishing a connection over each of said active links;

receiving a data access request to convey data to a first of said active nodes;

conveying said data over one or more of said active links to said first active node;

determining that said physical resources have changed; and

reestablishing connections to said changed physical resources;

wherein said determination of changed resources and said reestablishing of links are transparent to a client.

...

27. A computer-readable storage medium comprising program instructions for transporting data in a distributed computing system comprising a plurality of nodes and a data communication link, wherein said program instructions execute on a said plurality of nodes of said distributed computing system and said program instructions are operable to implement the steps of:

determining physical resources in said distributed computing system, wherein said physical resources include active nodes of said distributed computing system and **active links between said active nodes;**

establishing a connection over each of said active links;

receiving a data access request to convey data to a first of said active nodes;

conveying said data over one or more of said active links to said first active node;

determining that said physical resources have changed; and

reestablishing connections to said changed physical resources;

wherein said determination of changed resources and said reestablishing of connections are transparent to a client.

'787 patent, col. 19:12-29, 19:59-20:14 (emphasis added).

### **a. Switches**

NetApp contends that the "active links between said active nodes" do not involve switches, but only point-to-point and multi-drop links. NetApp argues that the specification uses the term "data communication link" to refer to all hardware that connects the nodes in the cluster. '012 patent, col. 4:1-3 ("The present invention further contemplates a method of transporting data in a distributed computing system including a plurality of nodes and a data communication link."). The term "switch" does not appear anywhere in the disclosure concerning the distributed computing system.

NetApp notes that the specification explicitly identifies the contemplated types of interconnects, i.e., multi-drop (or bus-based) links or point-to-point links, but never mentions switches. *Id.* at col. 5:26-27 ("Data communication link 102 contemplates a multi-drop link or point-to-point links."). This section of the specification and Figure 1, which shows a data communication link 102 between nodes, is merely one embodiment of the invention. '012 patent, col. 5:21-23 ("according to one embodiment of the present invention"). NetApp offers expert testimony that a multi-drop link does not require a switch: nodes are connected to a common data transmission medium called a "bus." A node broadcasts its message onto the bus, and all the other nodes on the bus can see the message, Ganger Decl. para. 53, similar to an intercom system. The bus does not switch the data traffic. NetApp offers further extrinsic evidence to show that a "point-to-point" configuration is defined differently from switches: "A communications link in which dedicated links exist between individual origins and destinations, *as opposed to* a point-to-multipoint configuration, in which the same signal goes to many destinations (such as a cable TV system), or a *switched configuration*, in which the signal moves from the origin to a switch that routes the signal to one of several possible destinations." Nathan Reply Decl. Ex. 8 at 371 (Microsoft Press technical dictionary) (emphasis added). There is no clear intrinsic evidence, however, that "active links between said active nodes" precludes switches; the patent is simply silent as to switches. NetApp's proposed construction of "active links between said active nodes" as multi-drop or point-to-point links improperly limits the claims to one disclosed embodiment.

Sun points to alternate embodiments. In one, the "data communication link 102 implements a scalable coherent interface (SCI)." '787 patent, col. 5:33-35. In another embodiment, "the cluster implements a TCP/IP protocol for transferring data over the SCI." *Id.* at col. 5:35-37. Sun offers extrinsic evidence that the SCI may use intermediate switches. Kaliski Decl. para. 18, Ex. C at 4 (ScIzzL Overview) ("The Scalable Coherent Interface (Local Area MultiProcessor) is effectively a combination computer backplane bus, processor memory bus, I/O bus, high performance switch, packet switch, ..."). Sun further argues that internet protocols, such as TCP/IP, inherently include intermediate switches, such that limiting "active links" to paths without switches would contradict the express teachings of the specification. Kaliski Resp. Decl. para. 19. Sun further contends that neither point-to-point nor multi-drop links necessarily preclude



intermediate switches: point-to-point and multi-drop links are defined by the endpoints of a communication transmitted over the link, not whether there are intermediate switches, bridges, routers or other devices between the end points. *Id.* para. 20. NetApp replies that the SCI network and high-level protocols such as TCP/IP can be used on many types of topologies and do not necessarily require a switch. Ganger Open. Decl. para. 53; Ganger Reply Decl. para. para. 23, 24. Not necessarily requiring a switch is not, however, equivalent to never using a switch in SCI and TCP/IP protocols.

NetApp argues that it is clear to one of ordinary skill in the art that the patent does not contemplate the use of a switch. Rather, NetApp contends, the specification describes a cluster networking scheme in which nodes know about the direct physical links between nodes and determine for themselves which ones to use to communicate with a particular other node: "CTI 322A (cluster transport interface) determines which links are connected to the destination node and conveys the message over one of those links." '012 patent, col. 14:47-49. NetApp offers expert opinion that the nodes are performing their own switching, which would be unnecessary if they were interconnected via a switch. Ganger Decl. para. 54. NetApp cites specification language that describes nodes in the cluster as aware of which physical link goes to which node, and that based on this knowledge, each node is able to select the outgoing link for the destination node. *Id.* para. para. 54, 55 (citing '787 patent, col. 12:39-47, 14:18-23, 14:47-49, Fig. 6B). NetApp contrasts this cluster networking scheme from a switched network, in which a node would not need to know which links are connected to which nodes because it would simply specify the receiving node's network address. *Id.* para. 55. While NetApp's argument has some appeal, it would restrict the claim to particular embodiments of the invention. *See* '787 patent, col. 12:38 ("In one embodiment"); 14:13-14 ("according to one embodiment of the present invention"); 14:46 ("In one embodiment"); 4:48-50 ("FIG. 6 is a flowchart diagram illustrating the operation of a virtual disk system according to one embodiment of the present invention.").

At oral argument, NetApp raised a new argument based on the file history, contending that in the October 3, 2000 Amendment, the applicants distinguished Wolff by distinguishing the claimed invention from a switching network. In supplemental briefing, NA cites applicants' statements referring to the Wolff prior art as disclosing "a switching network" and lacking "active links between said active nodes." The applicants distinguished Wolff on several grounds, but nowhere did the applicants state or imply that Wolff was not anticipatory because it used switches. *See* '787 Patent File History, 10/3/00 Amendment at 14-15. NetApp concedes that the prosecution history is "not strictly speaking a clear disavowal of scope of coverage, representing a clear disavowal of claim scope." NetApp Br. re Pros. Hist. at 3. Rather, NetApp suggests that the file history demonstrates that applicants understood their own invention not to include switching networks. *Id.* The Court declines to infer prosecution disclaimer based on a singular reference to a prior art switching network, taken out of context. *Omega Eng'g*, 334 F.3d at 1325 (requiring disavowing statements to be "so clear as to show reasonable clarity and deliberateness" and "so unmistakable as to be unambiguous evidence of disclaimer").

The Court has no reason to disagree with NetApp's expert that in 1998, switches were well-known, commonly used mechanisms for increasing scalability relative to multi-drop links and simplifying communication relative to collections of point-to-point links. *See* Ganger Decl. ISO NetApp Open. Br. para. 54. However, the Court will not infer from the absence of the term "switch" that the patent excludes the possibility of using switches as links. Nor will the Court import a limitation from a preferred embodiment to limit data communication links to multi-drop or point-to-point links. The Court therefore declines to adopt NetApp's proposed construction of "active links between said active nodes" as "active point-to-point links or active physical connections to a multi-drop link."

#### **b. Active**

[27] [28] The parties dispute whether "active" should be construed as "designated as a member of the cluster," as NetApp proposes, or as "capable of sending and receiving data," as Sun proposes. The

specification explains that data access requests are routed to active nodes:

In the illustrated embodiment, when client 312A accesses data from a storage device, it sends a data access request to ND 318A. In one embodiment, client 312A specifies the destination storage device, the type of operation and the location to retrieve or store the data to ND 312A. The rest of the operation is transparent to client 312A. ND 318A, based on the mapping and current membership information, determines to which node to convey the data access request. In one embodiment, the mapping information obtained from CCD 311A includes a primary and secondary node physically connected to the storage device. ND 318A may route the data access request to the primary node if the primary node is active. Alternatively, if the primary node is not active, then ND 318A may route the data access request to the secondary node. Which node is used to access the storage device is transparent to client 312A.

'787 patent, col. 9:18-35. In this embodiment, the system may route the data access request to the primary node attached to a storage device if it is active; if a primary node is not active, the system will route the request to a secondary node. *Id.* at 9:30-35. Sun thus proposes that an "active link" be construed as a "an operational path capable of transferring data," that "active node" be construed as "an operational node capable of sending and receiving data," and that the term "active links between active nodes" be construed as "operational paths capable of transferring data between operational nodes capable of sending and receiving data."

NetApp challenges Sun's proposed construction of "active links" to mean "operational paths capable of transferring data" on the ground that it captures within its scope structures that are beyond what is contemplated by the patent. NetApp contends, for example, that one way to transfer data between two nodes is to connect the nodes to a shared disk, which would satisfy Sun's proposed definition of "active links," but a shared-disk arrangement for data transfer is not covered by the patent. Sun points out correctly, however, that its proposed construction expressly states that the paths exist between nodes, not between a node and a disk, as in the shared disk arrangement. Kaliski Decl. para. 17.

Rather than equating "active" with "capable of transferring data" and "capable of sending and receiving data," as Sun proposes, NetApp contends that the patent refers to an "active" node or link as one that is considered a member of the cluster. NetApp cites Claim 1 of the '787 patent which recites: "wherein the membership information includes a list of active nodes of the plurality of nodes." '787 patent, col. 18:5-7. NetApp interprets this language as equating "active" nodes with those identified in the cluster's configuration information as members of the cluster, not just the nodes or links in the cluster that are functional. The claim and specification only state, however, that the membership information of a cluster *includes* a list of active nodes. *Id.* *See also* '787 patent, col. 8:66-9:1 ("NDD 314A may also query CMM 310A to obtain membership information, *such as* a list of active nodes.") (emphasis added). Thus, the specification does not necessarily equate membership in a cluster with whether the node is "active." *See* '787 patent, col. 8:41-49 ("Examples of membership changes include a node joining or leaving the cluster."). Rather, the language upon which NetApp relies could just as easily mean that the cluster membership information specifies which members are operational, i.e., active, and which members are not, such as temporarily malfunctioning members. For example, the specification refers to a particular embodiment where the cluster membership data may indicate that primary node 104B is not operational, and the system may route the data access requests to the secondary node. *Id.* at 10:59-63. That does not necessarily mean the cluster has been reconfigured to exclude node 104B. On the other hand, the specification also provides that "[r]econfiguration may occur when a node joins or leaves a cluster, or when a node fails," *id.* at 8:51-53, and "reconfiguration data [may be received] indicating that a destination node is not active," *id.* at 10:37-41. These passages confirm that a node failure can lead to its exclusion from the cluster. Yet they do not establish that node failure necessarily leads to exiting the cluster, as opposed to being an inactive member, at least for some period of time.

Finally, NetApp argues that Sun's proposed construction would read the term active out of the claims, offering as an example claim 11 which recites "a message received from another active node of the plurality of nodes." Because claim 11 teaches that the active node is capable of sending a message to another node, NetApp contends that construing "active" as "capable of sending and receiving data" renders the term "active" redundant. In the absence of any inconsistency between Sun's proposed construction and claim 11, the Court does not find NetApp's argument to be persuasive.

The Court accepts Sun's proposed construction of "active" as "capable of sending and receiving data." The specification describes an embodiment of the virtual disk system designed such that data access requests are performed in the presence of a node failure:

If a data access request is not successfully completed, the virtual disk system retries the data access request possibly to another node. The requesting node may detect an incomplete data access request by receiving a negative acknowledge signal or it may receive reconfiguration data indicating that a destination node is not active. When the data access request is successfully complete, it is removed from the list of pending data access requests.

For example, node 104B may be a primary node for a storage device and node 104C may be a secondary node for that storage device. When ND 318A conveys a data access request to the storage device, it may convey the data access request to the primary node, which is node 104B. If node 104B is unable to successfully complete the data access request, for example if the storage device path between disk driver 326B and the storage device is non-functional, node 104A may receive a negative acknowledgement signal indicating that the data access request was not successfully completed. Node 104A may then resend the data access request to the secondary node, which is node 104C. Node 104A may store information indicating that node 104B is not able to communicate with the storage device and subsequently send new data access requests to other nodes.

In an alternative example, node 104B may be non-operational. In one embodiment, the cluster membership data acquired by node 104A from CMM 310A may indicate that the node is not operational. Accordingly, ND 318A may route data access requests to the secondary node. In the above manner, data access requests are successfully completed even in the presence of a failure.

'787 patent, col. 10:35-65. The specification thus primarily describes active nodes in a functional capacity, i.e., as being able to communicate or being operational in order to complete a data access request. While it appears that active nodes are members of the same cluster, it does not follow that the cluster only includes active nodes and cannot include at least temporarily nonfunctioning nodes as members.

The Court adopts Sun's construction of "active link" as "operational path capable of transferring data" and "active nodes" as "operational nodes capable of sending and receiving data."

### **c. Storage devices**

NetApp proposes that "nodes" should be further construed as "devices in a cluster that are not storage devices." NetApp cites the language of Claim 1 to show that "nodes" are different than "storage devices:" "A distributed computing system comprising: a plurality of nodes ...; a storage device ... physically connected to at least one of the subset of the plurality of nodes, wherein the storage device is not physically connected to the first node; ...." '787 patent, col. 17:59-18:2. NetApp also relies on the specification, which defines a node and storage device in the disclosure: "Generally speaking, a node is a group of circuitry designed to perform one or more computing tasks.... Generally speaking, a storage device is a persistent device capable or storing large amounts of data." '787 patent, col. 1:13-15, 1:29-31. NetApp also points out that the Abstract refers to the nodes and storage devices as distinct parts of the cluster: "A cluster

implements a virtual disk system that provides each node of the cluster access to each storage device of the cluster." *Id.* at Abstract.

Sun contends that NetApp's proposed construction of the term "node" to exclude storage devices contradicts the intrinsic evidence. Sun's expert opines that in the abstract, a group of circuitry designed to perform one or more computing tasks can include a storage device, e.g., a computer typically includes a disk and other persistent devices for storing data. Kaliski Decl. para. 21. Sun also argues that the specification discloses that in one embodiment, each node includes a memory or cache for storing data from attached disks, citing '787 patent, col. 5:46-49, 15:1-4; nodes also store "mapping information" and "consistent data even in the presence of failure," *id.* at 8:33-36; and nodes store "messages and data received from other nodes," *id.* at 15:1-4. Those passages do not, however, equate "storage device" with "node," as used in the patent, although they disclose that data may be temporarily stored, or cached, in the memory of a node. Sun argues that although the specification defines "nodes" and "storage devices" separately, both definitions are broad enough to include a subset of devices that satisfy both definitions; in other words, they are not mutually exclusive terms. Sun's argument, however, is inconsistent with the usage of those terms in the patent. NetApp's proposed construction of node does not exclude the possibility that nodes may have storage capabilities but recognizes correctly that nodes with memory or cache for storing data are not the same as the long-term "storage devices" referred to in the patent. *See* '787 patent, col. 5:46-49 ("Data may be stored in storage devices 108-112 which is not currently being used by a node, and data from the storage device may be stored, or cached, in the memory of the node when the data is needed."). Nodes may temporarily store data, but a separate device stores data long-term.

The Court agrees with NetApp that the patent defines a node as "a group of circuitry designed to perform one or more computing tasks," '787 patent, col. 1:13-15, and is distinct from storage devices. The Court thus agrees with NetApp that nodes are not the same as storage devices as those terms are used in the patent.

For the reasons set forth above, the Court construes the term "active links between said active nodes" as "operational paths capable of transferring data between operational nodes capable of sending and receiving data, that are separate from the storage devices."

#### **IV. CONCLUSION**

In accordance with the foregoing, and for the reasons discussed above, the Court construes the disputed terms of the Sun patents as follows:

1. "First available memory space" is "one or more logical tracks, each of which must be empty, i.e., recognized by the system as available free space."

2. "Means, responsive to the receipt of a stream of data records from said associated data processor, for writing said received stream of data records in available memory space in one of said disk drives" is a means-plus-function limitation, in which the function is "writing said received stream of data records in available memory space in one of said disk drives" and the corresponding structure is "control unit 101 programmed to execute the steps 701 through 712 illustrated in Figure 7 and described under the heading 'Data Write Operation' at column 16:26 through column 17:25 of the '855 patent specification."

3. "Means responsive to said reading of said memory controller for periodically verifying the integrity of data currently stored in each of said identified dedicated partitions" is a means-plus-function limitation, in which the function is "periodically verifying the integrity of data currently stored in each of said identified dedicated partitions," and the corresponding structure is "the control unit 101 and the disk drive subassembly programmed to execute the steps 1202 to 1207 illustrated in Figure 12 and described at col. 15:47-62, and step 1008 illustrated in Figure 10 and described at col. 15:63-16:25."

4. "Means for reading data in each of said identified dedicated partitions" is a means-plus-function limitation, in which the function is "reading data in each of said identified dedicated partitions" and the corresponding structure is "the disk drive manager and the control unit 101 programmed to execute step 1202 illustrated in Figure 12 and described at col. 15:49."

5. "Active links between said active nodes" are "operational paths capable of transferring data between operational nodes capable of sending and receiving data, that are separate from the storage devices."

**IT IS SO ORDERED.**

FN1. Although Sun's final proposed construction limits the corresponding structure to processor 204, Sun conceded in briefing and at oral argument that control unit 101 is the corresponding structure.

FN2. While the disk drive subsystem also includes the disks themselves, these store the data that is read, rather than perform the reading function. Further, NetApp did not argue to the contrary.

Produced by Sans Paper, LLC.