

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28

IN THE UNITED STATES DISTRICT COURT
FOR THE NORTHERN DISTRICT OF CALIFORNIA

NETWORK APPLIANCE INC,

No. C-07-06053 EDL

Plaintiff,

ORDER CONSTRUING CLAIMS

v.

SUN MICROSYSTEMS INC,

Defendant.

On August 27, 2008, the Court held a claim construction hearing to construe the disputed terms of United States Patent Numbers 5,925,106 (the “106 patent”), 5,459,857 (the “857 patent”), 5,749,095 (the “095 patent”), 6,873,630 (the “630 patent”) (collectively, “Sun patents”), and U.S. Patent Nos. 5,819, 292 (the “292 patent”), 6,892, 211 (the “211 patent”), and 7,200,715 (the “715 patent”) (collectively, “NetApp patents”) pursuant to Markman v. Westview Instruments, Inc., 517 U.S. 370 (1996). Having read the parties’ papers and carefully considered their arguments and the relevant legal authority, the court hereby rules as follows.

I. BACKGROUND

On September 5, 2007, Network Appliance, Inc. (“NetApp”) filed its Complaint in the Eastern District of Texas, alleging that Sun Microsystems, Inc. (“Sun”) infringed and is infringing, directly and indirectly under 35 U.S.C. § 271, certain patents, including the ’292 patent, the ’211 patent, and the ’715 patent, by making, using, selling, or offering for sale data processing systems and related software, primarily Zettabyte File System (“ZFS”). NetApp further seeks a declaratory judgment that certain patents owned by Sun are each not infringed, are invalid and/or are unenforceable, and it seeks a permanent injunction to prevent future infringements as well as damages adequate to compensate it for Sun’s past infringement and trebling of damages pursuant to

1 35 U.S.C. § 284.

2 On October 25, 2007, Sun filed its Answer and Counterclaim, denying the material
3 allegations of the Complaint and asserting a number of affirmative defenses and counterclaims.
4 Sun denies infringing any of the NetApp Patents and alleges that NetApp infringes a number of its
5 patents, including its four patents at issue in this claim construction. Sun further seeks a declaratory
6 judgment that the NetApp Patents are each not infringed, are invalid and/or are unenforceable, as
7 well as a permanent injunction to prevent future infringements as well as damages adequate to
8 compensate it for NetApp's past infringement and trebling of damages pursuant to 35 U.S.C. § 284.

9 The case was transferred to this District on November 30, 2007. On December 7, 2007,
10 NetApp filed its Reply to Sun's Answer and Counterclaims, denying the material allegations of
11 Sun's Answer and Counterclaim and asserting a number of affirmative defenses and counterclaims.

12 The parties now seek construction of fourteen disputed terms and/or phrases, which are
13 contained in various claims in the seven patents.

14 **II. DISCUSSION**

15 **A. Legal Standard**

16 In construing claims, the court must begin with an examination of the claim language itself.
17 The terms used in the claims are generally given their "ordinary and customary meaning." See
18 Phillips v. AWH Corp., 415 F.3d 1303, 1312-13 (Fed. Cir. 2005); see also Renishaw PLC v.
19 Marposs Societa' per Azioni, 158 F.3d 1243, 1248 (Fed. Cir. 1998) ("The claims define the scope of
20 the right to exclude; the claim construction inquiry, therefore, begins and ends in all cases with the
21 actual words of the claim."). This ordinary and customary meaning "is the meaning that the terms
22 would have to a person of ordinary skill in the art in question at the time of the invention...".
23 Phillips, 415 F.3d at 131. A patentee is presumed to have intended the ordinary meaning of a claim
24 term in the absence of an express intent to the contrary. York Products, Inc. v. Central Tractor Farm
25 & Family Ctr., 99 F.3d 1568, 1572 (Fed. Cir. 1996).

26 Generally speaking, the words in a claim are to be interpreted "in light of the intrinsic
27 evidence of record, including the written description, the drawings, and the prosecution history, if in
28 evidence." Teleflex, Inc. v. Ficosa North Am. Corp., 299 F.3d 1313, 1324-25 (Fed. Cir. 2002)

1 (citations omitted); see also Medrad, Inc. v. MRI Devices Corp., 401 F.3d 1313, 1319 (Fed. Cir.
2 2005) (court looks at “the ordinary meaning in the context of the written description and the
3 prosecution history”). “Such intrinsic evidence is the most significant source of the legally
4 operative meaning of disputed claim language.” Vitronics Corp. v. Conceptronic, Inc., 90 F.3d
5 1576, 1582 (Fed. Cir. 1996).

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

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

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

28 “[T]here is sometimes a fine line between reading a claim in light of the specification, and

1 reading a limitation into the claim from the specification. . . . [A]ttempting to resolve that problem in
2 the context of the particular patent is likely to capture the scope of the actual invention more
3 accurately than either strictly limiting the scope of the claims to the embodiments disclosed in the
4 specification or divorcing the claim language from the specification.” Decisioning.com, Inc. v.
5 Federated Dep’t Stores, Inc., 527 F.3d 1300, 1308 (Fed. Cir. 2008) (citations omitted). There is
6 therefore “no magic formula or catechism for conducting claim construction,” and the court must
7 “read the specification in light of its purposes in order to determine whether the patentee is setting
8 out specific examples of the invention to accomplish those goals, or whether the patentee instead
9 intends for the claims and the embodiments in the specification to be strictly coextensive.” Id.
10 (internal citations omitted).

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

15 In most cases, claims can be resolved based on intrinsic evidence. See Vitronics, 90 F.3d at
16 1583. Only if an analysis of the intrinsic evidence fails to resolve any ambiguity in the claim
17 language may the court then rely on extrinsic evidence, such as expert and inventor testimony,
18 dictionaries, and learned treatises. See Vitronics, 90 F.3d at 1583 (“In those cases where the public
19 record unambiguously describes the scope of the patented invention, reliance on any extrinsic
20 evidence is improper”). “Within the class of extrinsic evidence, the court has observed that
21 dictionaries and treatises can be useful in claim construction.” Phillips, 415 F.3d at 1318. While
22 expert testimony can be useful to a court for a variety of purposes, conclusory assertions by experts
23 are not useful to a court. Id. The court generally views extrinsic evidence as less reliable than the
24 patent and its prosecution history in determining how to read claim terms, even if though
25 consideration is within the court’s sound discretion. See id. at 1318-19.

26 **B. Construction of Disputed Terms and Phrases**

27 The parties dispute fourteen terms and/or phrases in the seven patents. The first seven
28 disputed terms and/or phrases are contained in the Sun patents, and the latter seven terms are

1 contained in the NetApp patents.

2 **Patent '106**

3 This patent concerns an invention for obtaining and displaying information about a server on
4 a network. It describes a mechanism for helping a user of the internet better understand the source
5 of a webpage or server by adding a descriptive reference or parenthetical phrase after the web
6 address on the user's computer screen to help the user of the internet identify the owner of the
7 website or server. For example, this invention would add the parenthetical "Northern District of
8 California" after the Court's website www.cand.uscourts.gov. The patent describes the invention
9 within the context of a preferred embodiment using a world wide web browser and server, but notes
10 that the internet browser and servers are representative of the technology. See '106 Patent at 1:31-
11 36.

12 **1. "Domain Name"**

NetApp's Proposed Construction	Sun's Proposed Construction
A third-party-approved name of a website on the Internet, i.e. a registered domain name	A name that has a numerical IP address associated with it

16 The key issue here is whether domain name is limited to a registered name of an internet
17 website, or is more broadly defined. To begin its analysis, the court first turns to the claims
18 themselves. The term "domain name" appears in each independent claim of the '106 patent, claims
19 1, 9, 14, 22, 27, and 35. As Sun notes, the independent claims containing the term "domain name"
20 describe systems and methods for providing server information to a client over a network. See, e.g.,
21 claims 1 ("provide access to a network"), 9 ("A server apparatus having . . . access to a network). In
22 contrast, some dependent claims add the use of the www, urls, and http. See, e.g., claim 5 ("access
23 mechanism further comprises a world wide web (www) browser apparatus configured to use . . . http
24 to utilizing a uniform resource locator (url)", claim 11 (requires usage of www, http, and urls).
25 Accordingly, it appears that NetApp's proposed construction, which limits domain names to
26 registered internet domain names and would therefore imposes www, url, and http requirements in
27 the independent claims, may render much of the language in the dependent claims superfluous.
28 Such a reading would violate the principle of claim differentiation. See Uniroyal Inc. v. Rudkin-

1 Wiley Corp., 837 F.2d 1044, 1054-55 (Fed. Cir. 1988). However, it is not entirely clear from a plain
2 reading of the claims whether the dependent claims are rendered completely superfluous, as the
3 internet is broader than the world wide web, and encompasses other public network applications.
4 See Almeroth Decl. at 10. Therefore, because the claim language is not dispositive, the court turns
5 to the specification for further guidance.

6 The specification explains that “domain names are requested by maintainers of the website
7 and are approved by a third party.” ’106 Patent at 1:56-57. In addition, the figures in the patent are
8 cutouts of Netscape, a program known to look up websites on the internet based on domain names.
9 See ’106 Patent at 5:15-21; Almeroth Decl. at 6. While this language at first blush may seem to
10 suggest that domain names are limited to their use in websites and require third party approval, as
11 Sun notes, this portion of the specification is discussing the www preferred embodiment, which the
12 patent explicitly states is nonexclusive. See ’106 Patent at 1:32-37 (“Although the invention covers
13 information access and information provider apparatus, www browser and www server applications
14 are representative of the technology. As such, the majority of this application describes the
15 invention within the context of a preferred embodiment utilizing www browser and www server
16 applications.”); 4:24-29 (“Although the invention covers information access and information
17 provider apparatus, methods, and computer program products, www browser and www server
18 apparatus and applications are representative of the technology.”). Therefore, importing such
19 definitions into the claim terms would improperly import limitations from the preferred
20 embodiment.

21 The summary of the invention provides further evidence that the term should not be limited
22 to the www and internet. It describes the invention as applying broadly to networks, servers, and
23 client computers, which are not limited to the internet and www. ’106 Patent at 2:23-59. In
24 addition, the operation of the invention does not always require third party approval or registered
25 domain names. For instance, the patent explains that server identification data can be retrieved
26 using a conventional database that associates domain names with server identification data. ’106
27 Patent at 7:56-8:4. This reference in the specification is consistent with Sun’s extrinsic expert
28 testimony, which, while less important to the analysis, provides further support for the

1 understanding that such a database need not be maintained by a third party or use only registered
2 domain names. Kaliski Decl. ¶ 15.

3 NetApp argues that the references incorporated by the patent support its construction of the
4 “domain name” term, as they describe the domain name by defining the host component of the URL
5 as a legal internet host domain name or as a fully qualified domain name of a network host. NetApp
6 Op. Brief at 6 n.2; Almeroth Decl. at 7. However, these definitions are limited to the context of the
7 preferred embodiment and consequently are not persuasive. They also do not specifically refer to a
8 registration requirement and are limited to “domain name systems” functions. Walter Decl., Ex. K §
9 2.1. Consequently, such definitions are not particularly relevant here.

10 In analyzing the specification as a whole, therefore, while NetApp’s construction is
11 supported by some direct language in the specification, that language refers to domain names only in
12 the context of websites, a limitation which cannot be read into the term “domain name” in all
13 contexts of the patent. In addition, at the hearing, NetApp acknowledged that its proposed
14 construction effectively excludes private networks. This exclusion is improper, as the invention
15 refers broadly to networks, and unique identifying information might be desirable on large private
16 networks, such as may be maintained by businesses. According to Sun’s expert’s unrebutted
17 testimony, for example, it is common for private networks to use domain names that are not third
18 party approved or registered. See Kaliski Decl. ¶ 13. Imposing a uniqueness requirement for
19 domain names and for IP addresses, however, would not exclude private networks or other
20 embodiments of the invention, and Sun has not argued otherwise. The specification notes that the IP
21 addresses and domain names must be unique. Id. at 1:57-58. While this is in the context of the
22 preferred embodiment, logically speaking, it would be impossible to provide descriptive information
23 about the server identified by the domain name if the domain name were not unique in any way.
24 The specification, therefore, indicates that a uniqueness requirement should be incorporated into the
25 definition of domain name.

26 The parties also rely on the patent prosecution history in support of their constructions,
27 which is less important and less reliable in claim construction than the specification. See Phillips,
28 415 F.3d at 1317. The applicant explicitly defined “domain name” as a term “used in accordance

1 with standard usage in the field. It is provided to a nameserver to return an IP address.” Walter
2 Decl., Ex. E at 6 (Amendment A, Dec. 22, 1997). A nameserver is a mechanism for associating
3 names with IP addresses. *Id.* at 4. These statements do not by themselves support NetApp’s
4 construction. NetApp notes, however, that the patentee also stated with regard to other information,
5 that “unlike domain names, this additional information need not be unique nor approved by a third
6 party.” *Id.* at 3, 7. However, this latter statement was not directed to the terms or claims at issue.
7 The explicit definition set forth for “domain name” in the prosecution history is more relevant than
8 the inventor’s comments about domain names regarding different patent claims. Thus, the
9 prosecution history is too ambiguous to overcome the teachings of the claim language and the
10 specification.

11 As to the parties’ other extrinsic evidence, NetApp contends that the parties’ dictionary
12 definitions confirm that domain names are registered with third parties, as identifying the owner of
13 the domain name only makes sense if the domain name is registered with a third party. *See, e.g.*,
14 Webster’s New World Dictionary of Computer Terms, 8th Ed. (2000) (“domain name: in the system
15 of domain names used to identify individual internet computers, a single word or abbreviation that
16 makes up part of a computer’s unique name”) and Microsoft Press Computer Dictionary, 3d Ed.,
17 (1997) (defining domain name to mean “an address of a network connection that identifies the
18 owner of that address in a hierarchical format: server.organization.type”). While these definitions
19 might imply that internet addresses have owners, as Sun points out, there is nothing in these
20 definitions that requires third party approval or registration. At the same time, Sun’s own proffered
21 Webster’s definition requires that the domain name makes up part of a computer’s *unique* name,
22 which supports the specification’s focus on this requirement. *See* ’106 Patent at 1:57-58.

23 The parties discuss additional extrinsic evidence, none of which is persuasive given the
24 intrinsic evidence discussed above. For example, while Professor Almeroth opines for NetApp that
25 the term domain name is understood in the art to refer to an address of a computer on the internet,
26 where the address is approved by a third party and stored in the internet’s registry of domain names,
27 this is not persuasive given the language used in the specification to encompass networks more
28 broadly, as discussed above. Almeroth Decl. at 5:8-10. Sun also argues that NetApp’s regular use

1 of domain names to describe a “name that has a numerical IP address associated with it” and is not
2 registered or third party approved, for example, in documentation describing the accused products,
3 supports Sun’s proposed construction. See Williamson Decl., Ex. A at 101, 103, 104, 105, and 107
4 (referring to NIS domain names, which according to Sun’s expert, need not be third party approved);
5 Kaliski Decl. ¶ 16. However, this extrinsic evidence is fairly attenuated and, as NetApp notes, the
6 use of the phrase “NIS domain name” may just as easily be said to confirm that “domain name” by
7 itself has a different meaning from NIS domain name.

8 Finally, NetApp cites Federal Circuit cases from 2003-2005 that discuss domain names as
9 being premised on uniqueness of ownership and a centralized system of recording title, which
10 indicate that they are registered. See Resonate Inc. v. Alteon Websystems, Inc., 338 F.3d 1360,
11 1361-62 (Fed. Cir. 2003) (“Every web page is identified by a unique Uniform Resource Locator
12 (URL). . . . Every web site has a home page, which is identified by a URL and is the first document
13 users see when they first connect to the web site. Also associated with each web site is a domain
14 name, usually part of the URL.”); In re Oppedahl & Larson LLP, 373 F.3d 1171, 1176-77 (Fed. Cir.
15 2004) (“The simple fact that domain names can only be owned by one entity does not of itself make
16 them distinctive or source identifying.”); and Teva Pharmaceuticals USA, Inc. v. Pfizer, Inc., 405
17 F.3d 990, 997 (Fed. Cir. 2005). However, the application for the ’106 patent was filed in April
18 1996, over seven years before these cases were decided, so these cases are not extrinsic evidence of
19 what one of ordinary skill in the art understood at the relevant time. In addition, none of these cases
20 provide strong evidence that the term is limited to registered internet addresses. And while 15
21 U.S.C. § 1127, a statute regarding trademarks, not patents, defines domain name as “any
22 alphanumeric designation which is registered with or assigned by any domain name registrar,
23 domain name registry, or other domain name registration authority as part of an electronic address
24 on the Internet,” this definition was not enacted until November 1999. Nor is there any showing that
25 one of ordinary skill in the art would know of the legislation.

26 In sum, and for all the reasons set forth above, the court adopts Sun’s proposed construction
27 of “domain name,” modifying that term to include a uniqueness requirement and construes the term
28 as: “a unique name that has a unique numerical IP address associated with it.”

2. “Server Identification Data”

NetApp’s Proposed Construction	Sun’s Proposed Construction
Human-friendly information identifying a specific web server designed not to be intimidating to inexperienced users of the world wide web ¹	Information that uniquely identifies one server from other servers and can be seen by a user.

NetApp’s construction limits server identification data to that describing web servers to users of the world wide web, and further limits the term to human friendly information. The term “server identification data” appears in claims 1, 4, 9-10, 14, 17-18, 23-24, 27, 30, and 35-36 of the ’106 patent. In claim 1, the “said server identification data including descriptive information about a server” is contained in the invention, along with “a display mechanism configured to display said server identification data on said display device.” This claim language thus renders Sun’s proposed construction somewhat redundant. For example, Sun’s definition notes that the data can be seen by a user, but the separate claim terms about displaying said data on a display device already describes that function on its own.

The parties largely rely on the specification and prosecution history to construe this term. The specification explains that displaying “human-friendly server identification information to a user to better indicate the origin of the information” is the goal of the invention. ’106 Patent at 1:15-18. The specification explains that the information is needed to help inexperienced users avoid becoming disoriented and intimidated when surfing on the world wide web. *Id.* at 1:45-47, 1:60-2:3. The specification gives the example in which the server’s identification data includes the text string “Sun Microsystems” as descriptive information, which is enclosed in parentheses and separated from the URL by a space. *Id.* at 5:40-44. As noted above, the patent also specifically disclaims the www example as being the sole and exclusive focus of the invention, instead couching it as the context of a preferred embodiment, so importing an explicit web limitation into the claim term does

¹ At the hearing, without prior notice to Sun or the Court, NetApp proposed the following revised construction: “Information readily understood by an inexperienced user of the Internet that describes the identity of a specific server.”

1 not seem warranted.² Id. at 1:32-39.

2 The broad purpose of the invention is, however, to provide “human-friendly server
3 identification information to a user to better indicate the origin of the information,” and such
4 language describing “features of the present invention” limits a patent’s scope. See Verizon
5 Services Corp. v. Vonage Holdings Corp., 503 F.3d 1295, 1308 (Fed. Cir. 2007) (“When a patent
6 thus describes the features of the ‘present invention’ as a whole, this description limits the scope of
7 the invention.”). It is true that if a patent refers to the present invention and preferred embodiment
8 interchangeably, the use of the “present invention” language does not limit a patent’s scope. Karlin
9 Tech., Inc. v. Surgical Dynamics, Inc., 177 F.3d 968, 973 (Fed. Cir. 1999). However, the patent
10 does not do so here; instead, the inventor referred to “this invention” when broadly discussing the
11 purpose of “the invention” in the field of the invention section. See ’106 Patent at 1:15-18. In
12 addition, claim terms should be construed in a manner consistent with the “clear purpose of the
13 invention.” Cordis Corp. v. Medtronic Ave., Inc., 511 F.3d 1157, 1179 (Fed. Cir. 2008) (citation
14 omitted); see also CVI/Beta Ventures v. Tura LP, 112 F.3d 1146, 1160 (Fed. Cir. 1997) (“In
15 construing claims, the problem the inventor was attempting to solve, as discerned from the
16 specification and the prosecution history, is a relevant consideration.”).

17 Sun counters that the summary of the invention section of the specification broadly describes
18 the invention as a way of providing server specific information to a computer user. ’106 Patent at
19 2:24-28. The server identification data is obtained using a domain name and can be displayed using
20 a display device, and includes “descriptive information about the server identified by the domain
21 name.” Id. at 2:28-31. Sun argues that this language directly supports its proposed claim
22 construction. However, because the claims themselves refer to “server identification data including
23 descriptive information about a server and accessible by use of said domain name,” for example in
24 claim 14, Sun’s definition renders the other claim terms superfluous and repetitive, indicating that
25 “server identification data” has a more specific meaning. In sum, the specification supports
26 NetApp’s requirement of human-friendly information, but not its reference to inexperienced users of
27

28 ² NetApp’s revised construction omits reference to the www, but still refers to users of the internet.

1 the world wide web, as such limitations describe the preferred embodiment only.

2 Turning to the patent prosecution history, the inventor stated that the invention “addresses
3 the problem of identifying and/or describing a server site to a human user by providing additional
4 information about the server to the user,” distinguishing other forms of information like domain
5 names, hypertext links, URLs, and bookmarks. Walter Decl., Ex. E at 3-4. The prosecution history
6 explicitly defines “server identification information” as a term “used throughout the application to
7 mean descriptive information about a server that can be seen by the user – in particular text
8 [citations to text omitted]. This descriptive information is not an IP address that is returned by a
9 nameserver.” *Id.* at 5-6. From the prosecution history, it is clear that the term “server identification
10 data” excludes IP addresses, domain names, hypertext links, URLs, and bookmarks. NetApp argues
11 that Sun’s proposed construction must be wrong, because it would broadly cover IP addresses and
12 URLs, which were specifically disclaimed in the file history. At the hearing, Sun conceded that IP
13 addresses, domain names, and the like, in and of themselves did not constitute server identification
14 data. While Sun argued that the term “server identification data” could still include such items, this
15 distinction is meaningless: if such items cannot constitute the data, then it must be whatever else is
16 provided, along with the IP address, and the like, that constitutes the identification data. The clear
17 purpose of the invention – to provide more understandable, less technical information identifying the
18 server – is supported by the unambiguous exclusion from the definition of IP addresses returned by a
19 nameserver, domain names, hypertext links, URLs, and bookmarks.

20 The intrinsic evidence, therefore, teaches that the “server identification data” excludes IP
21 addresses, domain names, hypertext links, and the like, as it is descriptive data that is more user-
22 friendly. In sum, and for all the reasons set forth above, the court adopts a combination of the
23 parties’ proposed constructions that omits limitations to the world wide web, but includes the user-
24 friendly information limitation, which distinguishes the invention from the prior art, and construes
25 “server identification data” as: “information that uniquely identifies one server from other servers,
26 that is more human user-friendly than the server’s IP address, domain name, URL, or hypertext
27 link.” This definition stays true to the intrinsic evidence and addresses Sun’s concern that by itself
28 using the phrase “human-friendly information” may not be sufficiently clear to assist the jury.

Patent '857

This patent is directed to a fault tolerant data storage system designed to minimize the risk of failure. The patent does so through redundancy – that is, storing data on two separate storage subsystems, such that the failure of a single subsystem would not lead to the loss of data. The storage control units of both data storage subsystems are synchronized to maintain identical images of both subsystems, so that the data records are stored in available memory on both storage subsystems. Data is exchanged over the data link to maintain consistency in the two sets of information. See '857 Patent Abstract.

3. **“In response to writing a data record to said one redundancy group”/“Responsive to writing a data record to one of said redundancy groups”**
4. **“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”**

The parties address terms three and four together, and the Court will do the same. The following are the parties’ constructions for the two phrases, respectively:

NetApp’s Proposed Constructions	Sun’s Proposed Constructions
<p>After and in reaction to the writing of a data record to a single redundancy group;</p> <p>After and in reaction to the receipt of data records from a processor</p>	<p>In response to writing the data record to memory associated with the one redundancy group/Responsive to writing the data record to memory associated with one of the redundancy groups;</p> <p>Sun contends that the latter phrase does not require construction.</p>

To begin its analysis, the court turns to the claims themselves. These phrases appear in claims 6 and 11 of the '857 patent. Claim 6 states “transmitting, in response to writing a data record to said one redundancy group, said written data record to the other of said data storage subsystems via said data link to maintain duplicate data records in both said data storage subsystems.” NetApp argues that the claim language on its face requires that the transmission of a written data record from one data storage subsystem to another occur only after the record is written to a single redundancy

1 group,³ because “in response” to something means “after” according to its plain language. NetApp
2 also argues that the claim language further requires that the written data record be the thing that is
3 transmitted to the other redundancy group, and that use of the past-tense of “written” requires that
4 the writing be completed before the transmission.

5 It is true that “in response to” may mean “after.” More importantly, the inclusion of the
6 phrase “written data record” weighs heavily in favor of NetApp’s interpretation that the transmission
7 occurs after the complete writing of the data record, as the past tense term “written” has a plain
8 meaning, indicating that the data record has already been written to a redundancy group before it is
9 transmitted.

10 Turning to the specification, Sun argues that NetApp’s construction fails to include the
11 preferred embodiment, in which the writing occurs after the step of transmitting data to the other
12 storage subsystem. The ’857 patent describes a single preferred embodiment for writing duplicate
13 data records through data storage subsystems.⁴ The undisputed sequence is as follows. First, the
14 host processor transfers the data record to the first data storage subsystem. ’857 Patent at 11:51-54.
15 Then, the first data storage subsystem writes the data record into its cache memory. *Id.* at 11:54-62.
16 Then, the first data storage subsystem transfers the data record to the second data storage subsystem.
17 *Id.* at 11:62-12:1. Finally, the data record is written from cache memory to disk. *Id.* at 12:3-9.
18 According to Sun, because the data record is transmitted from the first data storage subsystem to the
19 second data storage subsystem after it has been written to cache memory, but before it is written to a
20 redundancy group of disks, the data storage subsystem can notify the processor that a write

21
22 ³ As explained in more detail below, “written to a redundancy group” means written to disk and
cannot mean written to cache memory.

23 ⁴ NetApp also argues somewhat cursorily that the patent may disclose additional possible
24 embodiments, as the patent describes both the writing step and transmitting step at one point without
25 specifying any particular sequence for performance of these steps. ’875 Patent at 4:39-50. That portion
26 of the patent states: “Host processors 101, 121 transmit data record write and read requests to storage
27 units 105 and 107 in conventional manner. The one of storage control units 105 and 107 that receives
28 these requests . . . communicates with its associated disk controller 112 and disk drives 109 to execute
the write and read requests. In addition, in response to a received data record write command, storage
control unit 105 transmits the received data record over data link 106 to storage control unit 107 to
maintain identical virtual device images in both storage control units 105 and 107 and identical data
records in data storage subsystems 1,2.” But this part of the specification does not disclose a different
sequence from the preferred embodiment discussed above. Rather, it is simply silent as to timing, and
lacks the specificity to constitute an alternative embodiment.

1 operation is complete when a data record is stored in cache memory, even if the data has not been
2 transferred to disk. '857 Patent at 11:62-12:7. This frees a host processor from waiting until data
3 has been written to disk and allows it to perform other tasks to increase performance and improve
4 efficiency. Id. at 6:63-7:1.

5 NetApp, however, argues that not all claims must cover the preferred embodiment, especially
6 where, as here, the patent has 54 claims including 16 independent claims, and only one detailed
7 description of the preferred embodiment. Generally, a claim term should not be interpreted “in a
8 way that excludes the preferred embodiment from the scope of the invention. . . Specifically, [the
9 Federal Circuit] has cautioned against interpreting a claim term in a way that excludes disclosed
10 embodiments, when that term has multiple ordinary meanings consistent with the intrinsic record.”
11 Helmsderfer v. Bobrick Washroom Equip., Inc., 527 F.3d 1379, 1383 (Fed. Cir. 2008). In
12 Helmsderfer, the appellant argued that the district court’s construction of the term was erroneous
13 because it excluded both the preferred embodiment *and every illustrated embodiment* from the two
14 particular claims at issue. The Court, however, noted that every single claim need not cover the
15 preferred embodiment in the patent where at least some claims covered that embodiment. See 527
16 F.3d at 1383. Moreover, as Helmsderfer recognized, the claim term should only be construed to
17 cover the preferred embodiment if one of the ordinary meanings consistent with the intrinsic record
18 so permits. Cf. Sinorgchem Co. v. ITC, 511 F.3d 1132, 1138 (Fed. Cir. 2007) (“Where, as here,
19 multiple embodiments are disclosed, we have previously interpreted claims to exclude embodiments
20 where those embodiments are inconsistent with unambiguous language in the patent’s specification
21 or prosecution history.”). Here, “written” does not have an alternative ordinary meaning.

22 NetApp does have a high hurdle to overcome to support its construction, as claims 6 and 11,
23 if construed as NetApp proposes, would not encompass any disclosed embodiment. However, in
24 Helmsderfer, the Court noted that some claims need not cover *any* of the embodiments if other
25 claims cover those embodiments. This reasoning seems particularly applicable here, where there are
26 so many claims and only one disclosed embodiment. In addition, the specific sequence of the
27 response is not required in every claim in this patent. For example, claim 31 covers a data storage
28 system comprising in part a “means in said first data storage control operable in response to the

1 receipt of said data record over one of said channel interface paths for transmitting said data record
2 over a data link transmission path to said second data storage control” and does not refer to
3 transmission of a written data record. ’857 Patent at 20:37-41. Other claims are similar. See, e.g.,
4 id. at Claims 24, 29, and 33.

5 Here, the claim language is clear and unambiguous. Construing the term according to Sun’s
6 construction to include the preferred embodiment would result in a construction that is contrary to
7 any ordinary meaning of the claim language. In order to have its construction cover cache memory,
8 Sun inserts the language “memory associated with” said one redundancy group into the claim term.
9 However, this improperly rewrites the claim to salvage it and does so in a manner that is inconsistent
10 with the specification. The claim language does not refer to transmitting a written data record in
11 response to writing said record to cache memory. Further, the patent makes clear that a redundancy
12 group consists of disk drives. See ’857 Patent at 12:9 (referring to “redundancy group of the disk
13 drives”); 2:5-11 (“data storage subsystem consists of a plurality of small form factor disk drives
14 which are configured into a plurality of redundancy groups”); 3:2-9 (“plurality of disk drives . . . are
15 configured into a plurality of variable size redundancy groups”). Inserting the term “memory
16 associated with” before the reference to redundancy group creates a construction at odds with the
17 plain and consistent definition in the specification of a redundancy group as “a redundancy group
18 consists of N+M disk drives.” Id. at 7:18. Sun’s proposed construction would take the plain
19 language of the claim, transmitting in response to writing a data record to a redundancy group (or
20 disk drives), and rewrite it into transmitting in response to writing a data record to cache memory
21 associated with those disk drives. But cache memory cannot be equated with or vaguely attached to
22 disk drives. For example, the specification states that “[a] virtual track is staged from a redundancy
23 group into cache,” which shows that the two are separate things. Id. at 3:65-66. It also states that a
24 storage control unit writes the data record from *its* cache memory into a selected redundancy group
25 of the disk drives, showing that cache memory pertains to the storage control unit and is separate
26 from the redundancy group. See id. at 12:7-9. For these reasons, Sun’s explanation that the term
27 “written data record” refers to data that has been written to cache memory, but not to disk, is directly
28 contrary to the plain language of the claim and to the specification.

1 The same reasoning applies to construction of the term “in response to the receipt of a stream
2 of data records,” which contains parallel language and indicates that the selection happens after the
3 receipt of the data. The claim language refers to “selecting, in response to the receipt of a stream of
4 data records from said data processor, available memory space in one of said redundancy groups to
5 store said received stream of data records thereon.” ’857 Patent at Claim 6. It also refers to a
6 “means, responsive to the receipt of a stream of data records . . . for selecting available memory
7 space in one of said redundancy groups to store said received stream of data records thereon.” Id. at
8 Claim 11. The plain language of these claims – namely the use of the term “received . . . data
9 records” – makes clear that the data records are already received when the selection of available
10 memory space occurs for storage purposes.

11 In sum, and for all the reasons set forth above, the court adopts NetApp’s construction of
12 these claim terms, construing “in response to writing a data record to said one redundancy
13 group”/“responsive to writing a data record to one of said redundancy groups” as “after and in
14 reaction to the writing of a data record to a single redundancy group” and construing “in response to
15 the receipt of a stream of data records from said data processor”/“responsive to the receipt of a
16 stream of the data records from said data processor” as “after and in reaction to the receipt of data
17 records from a processor.”

Patent ’095

19 This patent allows computer systems with multiple processors to accelerate performance by
20 providing a mechanism for rapidly writing data from a processor to system memory. This system is
21 expandable and comes with distributed shared memory, which overcomes the bottleneck of a single
22 shared bus memory structure. The patent is directed to the use of “fast write” operations in
23 multiprocessing systems where issues of “coherency” and stale memory data can arise. In these
24 systems, main memory is distributed amongst multiple processors. A coherency problem is created,
25 for example, if one processor modifies its copy of the data it is holding in its cache, but other
26 processors have an unmodified view of the data. Microprocessors achieve coherency by: 1)
27 informing the other processors that the copies of the data they are holding are invalid; or 2)
28 providing an updated copy of the data to the other processors. The ’095 patent claims a “fast write”

1 technique designed to expedite those operations that are sufficiently independent of other on-going
 2 operations that they do not implicate immediate coherency issues. Such operations can be processed
 3 “fast,” meaning that they do not need to await resolution of the coherency process to be performed.
 4 By performing the coherency operation after writing the modified data and making it accessible,
 5 performance is improved by freeing processor resources more rapidly. However, the patent notes
 6 that the fast write protocol is not suitable for all of the write operations within the computer system,
 7 and the invention codes for both fast write and slow write operations. ’095 Patent at 4:60-5:38.

8 **5. “Completing [a] Write Operation Within [a] Local Processing Node” /**
 9 **“Completing [a] Write Operation With Respect to [a] Processor”**

NetApp’s Proposed Constructions	Sun’s Proposed Constructions
<p>11 Claim 1: Transferring the write data from an 12 initiating processor to a system interface;</p> <p>13 Claims 11 & 17: Transferring the write data 14 from an initiating processor to a system 15 interface</p>	<p>16 Claim 1: Sun contends that the phrase need not 17 be construed because it is clear on its face. 18 However, if the Court construes the phrase, 19 Sun’s construction is: Data for the write 20 operation (1) is provided to subsequent read 21 operations within the local processing node to the same address as the write operation and (2) is or will be coherent within distributed shared memory;</p> <p>Claims 11 & 17: Sun contends that the phrase need not be construed because it is clear on its face. However, if the Court construes the phrase, Sun’s construction is: data for the write operation (1) is provided to subsequent read operations by a processor to the same address as the write operation and (2) is or will be coherent within distributed shared memory</p>

22 Here, Sun contends that the phrases need not be construed. Should the Court construe the
 23 claims, however, the parties assert that both phrases should be construed the same way, except that
 24 Sun’s proposed construction of the second phrase contains one variation from its proposed
 25 construction of the first phrase, where bolded above. Thus, the disputes between the parties are
 26 whether construction is necessary to assist the jury, and if so, what is the endpoint for a write
 27 operation to be considered complete with respect to a processor or within a processing node.
 28 NetApp argues that a write operation is completed for purposes of these claims when the write data

1 is transferred to the system interface.

2 As to Sun’s argument that the disputed language needs no construction because the meaning
3 of completing a write operation is clear to a person of ordinary skill in the art, while for some
4 commonly understood words, claim construction may be resolved by judges using the application of
5 the widely accepted meaning of those words, Phillips, 415 F.3d at 1314, an expert purporting to
6 comprehend a non-commonly understood phrase may not impose that meaning on a term. Rather,
7 “[b]ecause the meaning of a claim term as understood by persons of skill in the art is often not
8 immediately apparent, and because patentees frequently use terms idiosyncratically, the court looks
9 to those sources available to the public that show what a person of skill in the art would have
10 understood disputed claim language to mean.” Id. (internal quotations omitted). Such sources, as
11 discussed above, include claim language, the claim specification, the prosecution history, and
12 extrinsic evidence. This claim term is not clear on its face, and the Court turns to these sources to
13 construe the phrase.

14 To begin its analysis, the court first turns to the claims themselves. These phrases appear in
15 independent claims 1, 11, and 17 of the ’095 patent. Claim 1 is for a “method for performing write
16 operations in a multiprocessing computer system, comprising: initiating a write operation by a
17 processor within a local processing node of said multiprocessing computer system; performing a
18 coherency operation to at least one remote processing node in response to said write operation;
19 completing said write operation within said local processing node prior to completion of said
20 coherency operation if said write operation includes a specific predefined encoding; and completing
21 said write operation within said local processing node subsequent to completion of said coherency
22 operation if said write operation includes an encoding different than said specific predefined
23 encoding.” ’095 Patent at 31:63-32:10.

24 As Sun notes, this claim makes it clear that write operations within the processing node can
25 be completed prior to or subsequent to the completion of the coherency operations. The
26 specification also recognizes this. See ’095 Patent at 19:4-12; 27:17-23. The former constitutes a
27 “fast write” operation, and the latter constitutes a “slow write” operation. The specification contains
28 examples of slow-write operations completed subsequent to the completion of coherency activity,

1 such as where data is transferred to other caches in the system and where no data is transferred at all.
2 See '095 Patent at 3:44-50. NetApp's proposed construction, however, requires completion of a
3 write operation only upon transfer of the write data to the system interface. This ignores the write
4 invalidation procedure in which some write operations are completed without updating data in a
5 remote cache or in the memory, and hence without transfer to a system interface.⁵ See, e.g., id. at
6 17:57-67; 19:4-20:23 & Fig. 5. The processor need not transfer the data to the system interface in
7 operations where updated data is transferred to the other caches or is not transferred at all, as would
8 be required by NetApp's construction. Alpert Decl. ¶¶ 49-50. Therefore, this claim language, when
9 read along with the specification, indicates that NetApp's construction is too narrow.

10 In addition, independent claim 11 reads in part: "wherein said system interface is configured
11 to complete said write operation with respect to said processor . . ." '095 Patent at 32:41-43. This
12 claim therefore includes an express limitation providing that the system interface completes the
13 write operation. This limitation is not included in the other independent claims 1 or 17. This
14 indicates that the narrow claim limitation from claim 11 should not be read into the broader claims.
15 Uniroyal Inc. v. Rudkin-Wiley Corp., 837 F.2d 1044, 1054-55 (Fed. Cir. 1988). Because the two
16 claims use different language, a presumption exists that the claims have different meanings. CAE
17 Screenplates, Inc. v. Heinrich Fielder GmbH, 224 F.3d 1308, 1317 (Fed. Cir. 2000). However,
18 while the Court cannot read into a claim a limitation that is lacking, Curtiss-Wright Flow Control
19 Corp. v. Velan, Inc., 438 F.3d 1374, 1380-1381 (Fed. Cir. 2006), because these are two independent
20 claims (as opposed to an independent and dependent claim), claim differentiation is characterized
21 more generally as creating the "presumption that each claim in a patent has a different scope." Id.
22 Nonetheless, two claims with different terminology can define the exact same subject matter. Id.
23 Under the doctrine of claim differentiation, therefore, this Court still must look to the intrinsic
24 evidence to define the proper claim scope, but the claim language creates a presumption that
25 NetApp's proposed system interface limitation is not proper. See also Renishaw PLC v. Marposs
26 Societa Per Azioni, 158 F.3d 1243, 1249 (Fed. Cir. 1998) (court is not authorized to impose a

27
28 ⁵ At the hearing, NetApp argued that even in such circumstances, a command to invalidate is sent to the system interface. However, it is unclear how an invalidation command constitutes "write data," and NetApp has not offered any persuasive evidence showing that it is.

1 limitation that is not required by the specification).

2 The specification’s summary of the invention contains a general broad description of the
3 present invention without referring to system interface. It then notes, in a subsequent paragraph, that
4 the invention “further contemplates an apparatus for performing write operations in a
5 multiprocessing computer system comprising a processor and a system interface.” This second
6 paragraph explains that the system interface is “configured to complete the write operation.” ’095
7 Patent at 5:58. In a third and final paragraph, the specification explains that the present invention
8 further contemplates a computer system comprising a first processing node and a second processing
9 node. The third paragraph also notes that the first processing node is configured to complete the
10 write operation. Id. at 6:4-5. As Sun’s expert observed, the broader first paragraph in the
11 specification (id. at 5:39-51) generally tracks the language of claim 1. The second paragraph in the
12 description, limited to the system interface hardware (id. at 5:52-65), generally tracks the language
13 of claim 11. The third paragraph (id. at 5:66-6:12) tracks the language of claim 17. See also Alpert
14 Decl. ¶ 52. The patent also confirms that the invention may be implemented in alternate forms. See,
15 e.g., ’095 Patent at 31:54-59. The specification and the plain language of the claims demonstrate
16 that the inventor contemplated requiring the system interface architecture in only one of the three
17 independent claims.

18 NetApp points to examples showing that completion with respect to a processor within a
19 local processing node occurs when the write data is transferred to the system interface. However,
20 these examples merely show that fast-write operations may require data transfer to a system
21 interface to be complete. See, e.g., ’095 Patent at 27:15-27, 27:43-47 (stating that fast write
22 operations allow transfer of data to system interface before completion of coherency operations,
23 which causes the processor to view the write operation as being complete; further describing this
24 completion to be in local node); 31:36-41 (“[p]rocessor resources are freed upon transmission of the
25 write operation and corresponding data to the system interface, before an appropriate coherency
26 state is acquired by the node containing the processor”); 28:46-52 (describing completing a write
27 operation as follows: “Since ignore signal 70 is not asserted upon the fast write transaction, the
28 corresponding data is subsequently provide by processor 16 upon data bus 60 (shown in FIG 2).

1 During step 322, the data is received and stored by the system interface 24. The write operation is
2 thereby complete with respect to the initiating processor 16.”); 29:45-52 (noting that “since the fast
3 write stream transactions are completed from processors 16 by storing the transaction into SMP in
4 queue 94 and the corresponding data within output data queue 90, processors 16 may continue with
5 other operations while system interface 24 completes the write stream operations”). However, the
6 patent is not limited to fast write operations. Rather, the specification teaches that fast-write
7 operations are not always safe or proper, because in certain circumstances, coherency may be
8 necessary sooner. See id. at Abstract. As discussed above, the claims clearly contemplate that write
9 operations within the processing node can be completed prior to or subsequent to the completion of
10 the coherency operations, and may be coded differently as appropriate to do one or the other.

11 In light of the above, NetApp’s proposed construction, which requires a complete write
12 operation to transfer write data to a system interface, is far too narrow. As to Sun’s proposed
13 alternative construction, the specification requires that coherency operations be performed either
14 prior to or after performing the write, which provides support for the second prong of Sun’s
15 proposed constructions. Sun’s construction also provides that a completed write operation requires
16 that the written data be capable of being read, which is how one skilled in the art would define the
17 term, given that the specification does not support NetApp’s improper limitation. See Alpert Decl.
18 ¶¶ 39-41 (complete in context of write operation requires that written data be capable of being read
19 and this is well understood in the art) and Ex. D (citing NetApp’s expert’s 1987 publication). The
20 second prong of Sun’s constructions, however, are redundant of the express language of the claims,
21 since the claims themselves already specify that one class of writes completes prior to the coherency
22 and the other class of writes completes subsequent to coherency. Finally, Sun’s “is or will be
23 coherent” language would not be particularly helpful to a jury.

24 In sum, and for all the reasons set forth above, the court is deferring adopting a construction
25 of these terms at this time. The parties shall meet and confer and propose a new joint construction in
26 light of this claim construction order that does not import the system interface limitation and is along
27 the lines of Sun’s proposed construction, but addresses the redundancy and jury comprehension
28 issues discussed above. The parties shall submit an agreed upon construction no later than

1 **September 19, 2008.** In doing so, the parties do not waive any rights to preserve their objections
2 that their original proposals are correct.

3 **Patent '630**

4 This patent is directed at an ethernet network architecture that enables ethernet networks to
5 operate at high speed data transmission rates (multiple gigabits per second). The patent enables
6 increased transmission rates by dividing communications into smaller parts and distributing the
7 smaller parts among a plurality of channels. By sending the communication across multiple
8 channels rather than a single channel, the overall data transmission rate is increased to the sum of the
9 transmission rates for each of the multiple channels, rather than being limited to the rate for a single
10 channel. '630 Patent at 2:28-47. In one embodiment of the invention, the communication is
11 divided for transmission across multiple channels at a point below the medium access control
12 (“MAC”) layer of operation, so that the individual bytes of each frame or packet of the
13 communication are separated and sent across each one of the channels in a round robin fashion. *Id.*
14 at 2:36-42. When the portions of the data reach the destination computer, they are combined into a
15 single data stream and passed to the destination computer’s MAC layer. *Id.* at 2:33-35.

16 **6. “Portion [of a] Communication”**

17 **7. “Element [of a] Communication”**

Term	Agreed Construction
Portion of a Communication	Part of a data stream, where the part includes “elements” of the data stream.
Element of a Communication	A constituent part of a “portion.”

23 The term “portion [of a] communication” appears in claims 3, 5, 12-15, 21-22, 45, 48, 50,
24 52-53, 73, and 76 of the '630 patent. NetApp originally contended that this phrase should be
25 construed as “the fraction or portion of a frame carried by one channel.” Sun argued that the term
26 should not be limited to frames, because the frames did not include certain control data that was
27 clearly part of the data stream. In its reply brief, NetApp changed the focus of the dispute, arguing
28 instead that these terms had to be limited to at or below the medium access control (“MAC”) layer.

1 At the hearing, it became clear that the parties were unsure about the exact nature of their dispute.
2 For example, the parties agreed that the technology only includes ethernet technology, which
3 operates with frames, and agreed that the control data was part of the communication. The Court
4 agrees that such conclusions are supported by the specification. Sun also conceded that frames are
5 only found at or below the MAC layer. The parties therefore agreed to meet and confer to see if
6 they could come to an agreement about the construction of the terms. The parties subsequently
7 submitted the above agreed-to constructions, which the Court adopts. The Court construes “portion
8 of a communication” as “part of the data stream, where the part includes elements of the data
9 stream,” and construes “element of a communication” as “a constituent part of a portion.”

10 **'292 Patent**

11 The '292 patent describes a method for maintaining consistent states of a file system, and for
12 creating snapshots that are read-only copies of the file system. '292 Patent Abstract. The invention
13 uses a write anywhere file-system layout (“WAFL”). The file system progresses from one self-
14 consistent state to another self-consistent state. The set of self-consistent blocks on disk that is
15 rooted by a root inode is referred to as a consistency point. The root inode is stored in a file system
16 information structure. *Id.* In the detailed description of the invention, the invention is described as a
17 system that utilizes storage blocks on the disk, and disk drives provide the storage space for the file
18 system and maintain the structure and content of the file system. The system also uses an
19 in-memory “buffer” for holding changes to the file system prior to storing them on the disk. The
20 data blocks are organized, described, and pointed to by inode blocks. *Id.* at 5-6. The inode files
21 themselves are pointed to by a “root inode.” *Id.* at 9:34-35. When changes to the file system occur
22 during use, WAFL writes new or modified data to unallocated blocks on disk so that it never
23 overwrites existing data. *Id.* at 12:2-4. The file system in this patent can also maintain a copy of
24 past or old data in a consistent state by retaining it in a read-only form called a snapshot. *Id.* at
25 4:20-21.

26 ///

27 ///

28 ///

1 **8. “Non-volatile Storage Means”**

NetApp’s Proposed Construction	Sun’s Proposed Construction
A storage device that can retain information in the absence of power	Sun contends that the term non-volatile storage means should be construed as a means-plus-function limitation governed by 35 U.S.C. § 112 ¶ 6, and that the <u>function</u> is “storing blocks of data of a file system so that the data is not lost in the absence of power” and that the <u>structure</u> is “one or more disks with a block-based format (i.e., 4KB blocks that have no fragments), where the disk storage blocks are the same size as the data blocks of the file system.”

10 To begin its analysis, the court first turns to the claims themselves. The term “non-volatile storage means” appears in the two independent claims (claims 4 and 8) of the ’292 patent. Claim 4 provides a method for maintaining a file system comprising blocks of data stored in blocks of a non-volatile storage means at successive consistency points comprising the steps of:

11 storing a first file system information structure for a first consistency point in said non-volatile storage means, said first file system information structure comprising data describing a layout of said file system at said first consistency point of said file system;

12 writing blocks of data of said file system that have been modified from said first consistency point as the commencement of a second consistency point to free blocks of said non-volatile storage means;

13 storing in said non-volatile storage means a second file system information structure for said consistency point, said second file information structure comprising data describing a layout said file system at said second consistency point of said file system. ’292 Patent at 25:11-29.

14 Claim 8 provides a method for creating a plurality of read-only copies of a file system stored in blocks on a non-volatile storage means, said file system comprising meta-data identifying blocks of said non-volatile storage means used by said file system, comprising the steps of:

15 storing meta-data for successive states of said file system in said non-volatile storage means;

16 making a copy of said meta-data at each of a plurality of said states of said file system;

1 for each of said copies of said meta-data at a respective state of said file system, marking said
2 blocks of said non-volatile storage means identified in said meta-data as comprising a respective
3 read-only copy of said file system. Id. at 26:1-15.

4 The first critical issue is whether the term is a means-plus-function limitation under 35
5 U.S.C. § 112, ¶ 6, which states: “An element in a claim for a combination may be expressed as a
6 means or step for performing a specified function without the recital of structure, material, or acts in
7 support thereof, and such claim shall be construed to cover the corresponding structure, material, or
8 acts described in the specification and equivalents thereof.” In Altiris, Inc. v. Symantec Corp., 318
9 F.3d 1363, 1375 (Fed. Cir. 2003) (internal citations omitted), the Federal Circuit described the
10 analysis of means-plus-function claims:

11 Where a claim uses the word means to describe a limitation, we presume that
12 the inventor used the term advisedly to invoke the statutory mandates for
13 means-plus- function clauses. This presumption can be rebutted where the
14 claim, in addition to the functional language, recites structure sufficient to
15 perform the claimed function in its entirety. Once the court has concluded
16 the claim limitation is a means-plus-function limitation, the court must first
17 identify the function of the limitation. The court next ascertains the
18 corresponding structure in the written description that is necessary to perform
19 that function. Structure disclosed in the specification is corresponding
20 structure only if the specification or prosecution history clearly links or
21 associates that structure to the function recited in the claim.

22 Id. (internal citations and quotations omitted).

23 NetApp first contends that there is no presumption here of a means-plus-function limitation,
24 despite the use of the word “means” in the claim term. NetApp argues that because the claims here
25 are process claims, not apparatus claims, the presumption does not apply. See NetApp Opp. at 2
26 (citing O.I. Corp. v. Tekmar Co., 115 F.3d 1576, 1582-1583 (Fed. Cir. 1997)). This argument is not
27 persuasive. O.I. Corporation does not hold that the use of “means” in a method claim falls outside
28 of § 112 ¶ 6. See 115 F.3d at 1582-83 (“The word ‘means’ clearly refers to the generic description
of an apparatus element, and the implementation of such a concept is obviously by structure or
material. We interpret the term ‘steps’ to refer to the generic description of elements of a process,
and the term ‘acts’ to refer to the implementation of such steps. This interpretation is consistent
with the established correlation between means and structure. In this paragraph, structure and
material go with means, acts go with steps.”). NetApp cites no case so holding, and in fact, the

1 Federal Circuit has construed “means” terms in method claims as means-plus-function limitations.
2 See, e.g., On Demand Mach. Corp. v. Ingram Indus., 442 F.3d 1331, 1334, 1340-41 (Fed. Cir. 2006)
3 (construing “means for” term in a patent for “system and method of manufacturing a single book
4 copy” in a method claim as a means-plus-function limitation); J & M Corp. v. Harley-Davidson,
5 Inc., 269 F.3d 1360, 1364 n.1 (Fed. Cir. 2001) (“Claim 17, which is a method claim, contains a
6 nearly identical means-plus-function limitation.”).

7 In addition, NetApp argues that reciting “means” without reciting the phrase “means for”
8 does not give rise to the presumption of § 112 ¶ 6. However, “means for” language is not necessary
9 to invoke the means-plus-function presumption. See, e.g., Eaton Corp. v. Rockwell Int’l Corp., 323
10 F.3d 1332, 1343 n.1 (Fed. Cir. 2003) (finding that “means (26 and 34) [associated with said
11 driveline system] are written in means-plus-function format”); Unidynamics Corp. v. Automatic
12 Prods. Int’l, 157 F.3d 1311, 1318 (Fed. Cir. 1998) (“spring means tending to keep the door closed”
13 written in means-plus-function format); Signtech USA, Ltd. v. Vutek, Inc., 174 F.3d 1352, 1355
14 (Fed. Cir. 1999) (“ink delivery means” as well as some “means for” terms found to be means-plus-
15 function elements). The means-plus-function presumption therefore applies, albeit perhaps
16 somewhat less strongly than if the term “means” had been followed by “for.” See, e.g., Greenberg
17 v. Ethicon Endo-Surgery, 91 F.3d 1580, 1583 (Fed. Cir. 1996) (“use of the terms ‘means’
18 (particularly as used in the phrase ‘means for’) generally invokes section 112(6) and that the use of a
19 different formulation generally does not”). However, the presumption is rebuttable. The Court must
20 also examine the language surrounding the term “means” to determine whether or not a function
21 corresponds to the term at issue and whether the claim recites sufficient structure for performing that
22 function.

23 The presence of a function corresponding to the non-volatile storage means is readily
24 apparent. Turning to the claim language, the parties do not dispute the well-understood meaning of
25 “non-volatile” as maintained in the absence of power. The functions of the “non-volatile storage
26 means” include storing data blocks of a file system (claims 4 and 8), storing file information
27 structures (claim 4), storing read-only copies of a file system (claim 8) and storing meta-data for
28 successive states of said file system (claim 8), so that the data is not lost in the absence of power.

1 '292 Patent at 25:11-29, 26:1-15. In light of the plain claim language, the recited function at a
2 minimum is storing blocks of data for the file system such that information is retained in the absence
3 of power, as both claims 4 and 8 describe this function at the outset of the claim.

4 NetApp argues nonetheless argues that the term “non-volatile storage means” is not
5 explicitly linked to the function in an active way. For example, in the first limitation of claim 4, the
6 “first file system information structure” is stored in the “non-volatile storage means,” which,
7 according to NetApp, shows that the means is just a passive component into which something else is
8 stored. However, storing is a function, just as “to store” is a verb, albeit a stative rather than an
9 active verb. In other words, it is quite common to refer to something being stored in an object,
10 rather than referring to that object’s function in a more active way. But this active/passive
11 distinction does not undermine the conclusion that storing is a function. Nor does NetApp cite any
12 authority that functions can only correspond to active verbs in means-plus-function claims. In
13 addition, NetApp contradicts itself because it admits that every dictionary definition defines non-
14 volatile storage in “functional terms,” as a storage device that retains information in the absence of
15 power. NetApp Opp. at 5:13-15. Storing blocks of data of a file system so that the data is not lost in
16 the absence of power, therefore, is the corresponding function. The Court now turns to the more
17 difficult question of whether or not the claim recites sufficient structure for the means-plus-function
18 limitation to apply.

19 NetApp argues that the presumption does not apply because the claim recites sufficient
20 structure to perform the claimed function in its entirety. NetApp relies on a number of cases in
21 support of its argument, but each is distinguishable from this case. In Greenberg v. Ethicon
22 Endo-Surgery, 91 F.3d 1580, 1583 (Fed. Cir. 1996), the Court looked to dictionary definitions to
23 determine whether the noun “detent” disclosed sufficient structure. While finding that the term did
24 not call to mind a single well-defined structure, the Court held that the term, as the name for
25 structure, had a reasonably well understood meaning in the art. However, in Greenberg,
26 the Federal Circuit construed the term “detent mechanism,” not “detent means.” The Court
27 specifically distinguished a case in which that patentee specifically chose to invoke § 112 ¶ 6 by
28 invoking the “detent means” language. Id. at 1583-84. Specifically, the Court noted that “the use of

1 the term ‘means’ has come to be so closely associated with ‘means-plus-function’ claiming that it is
2 fair to say that the use of the terms ‘means’ (particularly as used in the phrase ‘means for’) generally
3 invokes section 112(6) and that the use of a different formulation generally does not.” Id. at 1584.

4 The fact that the “means” language is used here makes this case quite unlike Greenberg.

5 NetApp also relies on Cole v. Kimberly-Clerk Corp., 102 F.3d 524, 531 (Fed. Cir. 1996)
6 (holding that term “perforation means . . . for tearing” was not means-plus-function limitation
7 because the term described the structure supporting the tearing function (perforations) and went even
8 further to describe the location of the structure and extent of that structure); Envirco Corp. v. Clestra
9 Cleanroom, Inc., 209 F.3d 1360, 1365 (Fed. Cir. 2000) (concluding that claims recited sufficient
10 structure, where dictionary definition of the word baffle was “a device (as a plate, wall or screen) to
11 deflect, check, or regulate flow” so that term itself imparted structure (a surface which deflects air),
12 and where claims described details about location and structure of baffle); and Keithley v.
13 Homestore.com, Inc., 2007 WL 2701337, *20 (N.D. Cal. Sept. 12, 2007) (looking to dictionary
14 definition of terms at issue, and finding that server definitions including “A computer that manages
15 centralized data storage or network communications resources. A server provides and organizes
16 access to these resources for other computers linked to it.” disclosed inherent structure in the term,
17 and that claim language further detailed particularized structure). Those cases, however, differ from
18 this one in two important respects. First, the claims there detailed a more particularized structure,
19 and usually a specific location for that structure. Such claim language is not present here. Second,
20 the definitions in those cases outlined a more detailed structure for the term, such as a perforation,
21 which is a structure with a series of holes designed to allow tearing, and a baffle, which is a plate,
22 wall or screen to deflect, check, or regulate flow. The definitions for “non-volatile storage”
23 proffered here are more general and encompass a broader range of structures.

24 NetApp argues that Sun’s dictionary entries for “non-volatile storage” consistently define it
25 as a storage device that retains data in the absence of power, which recites sufficient structure to
26 perform the claim function and is the construction NetApp proposes. Furthermore, NetApp argues
27 that certain definitions point to specific electronic components, demonstrating that structure is
28 inherent in the term. The Microsoft Press Computer Dictionary (1991), for example, explains that

1 “non-volatile memory” is intended to refer to “core ROM, EPROM, bubble memory, or battery
2 backed CMOS RAM” and that the “term is occasionally used in reference to disk subsystems as
3 well.” Ganger Decl., Ex. F; see also Brandt Decl., Ex. 3 (Webster’s New World Dictionary of
4 Computer Terms, 376 (8th ed. 2000) (explaining that read-only memory is non-volatile, as are all
5 secondary storage units such as disk drives). NetApp relies on numerous other dictionary
6 definitions, many of which define the term “non-volatile storage” in term of its function, i.e., as a
7 “storage device which can retain information in the absence of power.” See Ganger Decl. ¶ 12
8 (citing IEEE Standard Dictionary of Electrical and Electronics Terms 699 (6th ed. 1996)).

9 Based on these definitions, it is clear that at a minimum, the term is understood as a storage
10 system that does not lose data when its power is cut off. This ability to maintain data in the loss of
11 power, however, is more functional than structural. Further, the myriad definitions encompass a
12 broad array of different types of structures that may include different types of ROM or disk drives,
13 among other things. However, the definitions do not make clear what type of structures are
14 necessarily contemplated by the term, and the definitions themselves are highly variable. See
15 Ganger Decl. ¶ 12 (citing various dictionary definitions which define non-volatile memory or
16 storage to include memory chips such as read-only memory in one definition; core, ROM EPROM,
17 bubble memory, or battery backed CMOS random access memory (RAM) and occasionally with
18 subsystems in another definition; read-only memory and secondary storage units like disk drives in
19 yet a separate definition; and magnetic tape, drum, or core in a different definition). NetApp’s own
20 expert includes “disks, disk arrays, flash memory drives, *and the like*” in the set of devices providing
21 non-volatile storage. Ganger Decl. ¶ 8 (emphasis added). However, flash memory drives are very
22 different structurally from disks, in that they are specifically designed to be small, readily detachable
23 by users without technical expertise, and highly portable. In addition, Dr. Ganger’s use of “and the
24 like” is extremely open-ended, indicating that the set of structures corresponding to non-volatile
25 storage are not ascertainable. While Sun’s expert seems to overreach in arguing that the structure
26 contemplated by these claims in the ’292 patent filed in 1995 would be understood by one of
27 ordinary skill in the art to include paper tape and punch cards (Brandt Decl. ¶ 65), even setting these
28 outliers aside, the dictionary definitions and NetApp’s own expert provide too broad a range of

1 devices that are not sufficiently similar in character to provide the corresponding structure. In
2 addition, hard disk drives are the relevant structure disclosed in the '292 patent, yet the dictionary
3 definitions do not focus on this structure.⁶ Therefore, while NetApp argues that the term has a broad
4 meaning in the form of many corresponding structures, citing Linear Tech. Corp. v. Impala Linear
5 Corp., 379 F.3d 1311, 1322 (Fed. Cir. 2004) (“That the disputed term is not limited to a single
6 structure does not disqualify it as a corresponding structure, as long as the class of structures is
7 identifiable by a person of ordinary skill in the art.”),⁷ here, that class of structures is not sufficiently
8 identifiable to defeat the means-plus-function presumption.

9 This case is similar to Altiris, Inc. v. Symantec Corp., 318 F.3d 1363 (Fed. Cir. 2003), in
10 which the Court analyzed whether “means of booting” was a means-plus-function limitation, where
11 the claim described “a means of booting said digital computer, said means of booting including a
12 first set of commands, said first set of commands resident on said storage device of said digital
13 computer for booting said digital computer, and a second set of commands resident on a storage
14 device external to said digital computer for booting said digital computer.” The district court
15 concluded that this was a means-plus-function limitation because “the language referring to two sets
16 of commands states only the location of the commands and is insufficient to define the structure that
17 performs the function of booting.” The Federal Circuit agreed. See id. at 1375-76. It noted that
18 while “commands” represented structure (in the form of software), “commands” did not provide
19 sufficient structure to perform the entirety of the function. Because that language was “so broad as
20 to give little indication of the particular structure used here and [wa]s described only functionally,
21 one must still look to the specification for an adequate understanding of the structure of that

22
23 ⁶ NetApp also argues that Sun’s own expert’s testimony provides additional evidence
24 confirming that non-volatile storage refers to a structural component, such as a disk, and that Sun’s
25 expert Marshall McKusick has equated non-volatile storage with disk in 2000 and 2007 publications.
26 See Ganger Decl., Exs. M at 12 (“non-volatile storage (i.e., disk)”) and N at 128 (same)). NetApp also
27 argues that Sun’s own patent portfolio reveals many patents in which non-volatile storage is equated
28 with a structure, generally a computer disk. Granger Decl. ¶ 16. But again, this disk limitation is in
contrast to the broader dictionary definitions discussed above, further indicating that the class of
structures is not sufficiently disclosed in the claims themselves to defeat the means-plus-function
presumption.

⁷ In addition, this cited passage in Linear pertains to identifying the corresponding structure
of a claim term that was already held to be governed by § 112 ¶ 6. Id.

1 software.” The Court distinguished cases holding that the claims disclosed sufficient structure as
2 construing patents in which “the exact structure used to accomplish the function appears in the claim
3 language.” *Id.* at 1376. Similarly, here, the term “non-volatile storage means” does not disclose
4 sufficient structure to accomplish the function in the claim language, and the dictionary definitions
5 do not disclose a clear indication of the particular class of structures used. See also Laitram Corp. v.
6 Rexnord, Inc., 939 F.2d 1533, 1536 (Fed. Cir. 1991) (“The recitation of some structure in a means
7 plus function element does not preclude the applicability of section 112(6). For example, in this
8 case, the structural description in the joining means clause merely serves to further specify the
9 function of that means. The recited structure tells only what the means-for-joining does, not what it
10 is structurally.”); Unidynamics Corp. v. Automatic Prods. Int’l, 157 F.3d 1311, 1319 (Fed. Cir.
11 1998) (addressing the term “spring means tending to keep the door closed” and concluding that
12 where spring was the only recitation of structure, it did not provide sufficient structure to overcome
13 presumption that § 112 ¶ 6 applied).

14 The parties both rely on cases that analyzed whether the term “storage means” was a means-
15 plus-function limitation. None of these cases alters the analysis. NetApp notes that in Lottotron,
16 Inc. v. Scientific Games Corp., 2003 WL 22075683 (S.D.N.Y. Sept. 8, 2003), the district court did
17 not find a means-plus-function limitation in the term “storage means,” where the dictionary
18 definition explained that it was a “device consisting of electronic, electrostatic, electrical, hardware,
19 or other elements into which data may be entered. . . .” *Id.* at *7. However, that Court did not
20 address the myriad of different dictionary definitions that exist in this case. That case also involved
21 a lottery wagering system – a completely different technology. In addition, neither party argued that
22 the claim was a means-plus-function limitation. *Id.* Finally, other cases have held that the term
23 “storage means” is a means-plus-function limitation. See, e.g., Apple Computer v. Burst.com, Inc.
24 2007 U.S. Dist. LEXIS 33863, *65-66 (N.D. Cal. May 8, 2007) (Patel, J.) (concluding that § 112 ¶ 6
25 applied to “storage means” where most specific dictionary definition provided that term was “any
26 device in which information can be stored, sometimes called a memory device,” noting that
27 “description of storage as a ‘memory device’ underscores the conclusion that ‘storage’ is a
28 functional term,” as a “memory device does not connote a particular structure-such as an input

1 port”).⁸

2 For these reasons, the claims do not recite sufficient structure sufficient to perform the
3 claimed function in its entirety. The term, therefore, is a means-plus-function limitation, and the
4 Court must determine the corresponding structure. The construction of a means-plus-function
5 limitation involves two steps. “A disclosed structure is corresponding only if the specification or the
6 prosecution history clearly links or associates that structure to the function recited in the claim.”
7 Omega Eng’g, Inc. v. Raytek Corp., 334 F.3d 1314, 1322 (Fed. Cir. 2003). The structure must be
8 necessary to perform the claimed function. Id.; see also Wenger Mfg., Inc. v. Coating Mach. Sys.,
9 Inc., 239 F.3d 1225, 1233 (Fed. Cir. 2001).

10 As discussed above, the function is “storing blocks of data for a file system so that the data is
11 not lost in the absence of power.” As to the corresponding structure disclosed in the specification,
12 the specification discloses only one corresponding structure storing data blocks of a file system. Sun
13 notes that the specification repeatedly and exclusively describes the invention in the context of the
14 structures and operation of a write anywhere file-system layout (WAFL) system. See, e.g., ’292
15 Patent at 5:49-52 (“The present invention uses a Write Anywhere File-system (WAFL).”); 5:62
16 (“WAFL inodes are distinct from prior art inodes.”); 6:53-56 (describing WAFL inodes); 8:57-9:17
17 (describing WAFL directories). Sun notes that fundamental to WAFL is its use of a disk format
18 system that is “block based (i.e., 4 KB blocks that have no fragments).” Id. at 5:48-53. The use of
19 disks with a block based format is the only structure disclosed for storing blocks of data in a file
20 system where the data is not lost in the absence of power. Brandt Decl. ¶ 67.

21 NetApp argues that Sun’s definition improperly imports the characteristics of the overlying
22 file system into the non-volatile storage, which “confuses jars and apple sauce.” Ganger Decl. ¶ 18.
23 Specifically, the requirements of 4KB blocks, blocks with no fragments, and disk storage blocks
24 with the same size as the data blocks of the file system are not characteristics of the non-volatile

25
26 ⁸ The Court in Apple, however, also construed the term “random access storage means” and
27 found that it was not a means-plus-function limitation. In so finding, the Court noted that both experts
28 essentially conceded that the functional language “storing the time compressed representation” was
specific enough to connote a particular class of structures, id. at *67. By contrast, this issue is hotly
contested here. Moreover, RAM storage appears to be at least slightly narrower than the “non-volatile”
storage at issue here. In addition, the Federal Circuit cases discussed above require more specificity in
the type of structure disclosed.

1 storage device. Rather, these are characteristics of the file system, as is evidenced by the fact that
2 Sun draws its proposed structure from a section of the specification describing the file system
3 layout. Ganger Decl. ¶¶ 18-19; '292 Patent at 5. This part of the specification does not focus on the
4 non-volatile storage device, but rather on how the overlying file system organizes the raw capacity.
5 See, e.g., '292 Patent at 5:49-52 (describing inodes, files, and directories, but not non-volatile
6 storage). And, indeed, block fragments are a file system concept. Ganger Decl. ¶¶ 19-20. At the
7 hearing, Dr. Ganger further explained that disks do not understand fragments, and that storage
8 devices read and write data, which they understand in units of KB. Sun's expert Dr. Brandt agreed,
9 but noted that when WAFL uses storage blocks, it accesses that data in 4 KB blocks. Dr. Ganger
10 conceded that while not all disks in this invention had to store data in multiples of 4 KB, it would be
11 far less convenient if the disks were not divisible by 4 KB for the purposes of this invention.

12 In sum, in looking at the specification in combination with the experts' explanation of the
13 invention, the structure of the file system is not the same as the structure of the non-volatile storage,
14 so Sun's proposed file system limitations are not necessary structure to perform the claimed storage
15 function. Specifically, fragments are a file system concept and should not be imported into the
16 structure of the non-volatile storage. And while the disks store 4 KB blocks, it is not clear from the
17 specification that the disks themselves must have a 4KB block structure. As Dr. Ganger notes, the
18 patent specification repeatedly explains that the thing in which the various file system structures are
19 stored is a disk, and the corresponding structure, therefore, must be disks and disk arrays.⁹ See
20 Ganger Decl. ¶ 17; see, e.g., '292 Patent at 9:48-49 ("written to disk"); 12:2-3 (WAFL always writes
21 new data to unallocated blocks on disk). It is also clear that for purposes of this invention, such
22 disks and disk arrays must be suitable for storing data in 4 KB blocks, since in order for WAFL to be
23 a disk format system that is block based (i.e. 4 KB blocks that have no fragments), the disk must be
24 formatted such that it can readily store 4 KB blocks. Supp. Brandt Decl. ¶ 7; '292 Patent at 5:48-53.

25 In sum, and for the reasons stated above, the Court adopts the following construction of this
26 term, which it construes as a means-plus-function limitation governed by 35 U.S.C. § 112, ¶ 6.

27
28 ⁹ At the hearing, Dr. Brandt did not contest inclusion of disk arrays, and noted that the multiple
disks could be used.

1 Function: storing blocks of data of a file system so that the data is not lost in the absence of power;

2 Structure: disk or disk arrays that are suitable for storing 4 KB blocks.

3 **9. “Meta-Data for Successive States of Said File System”**

NetApp’s Proposed Construction	Sun’s Proposed Construction
Information that describes “successive states of a file system” (as construed herein)	A block map file for recording snapshots of the file system

4
5
6
7
8 The parties dispute whether “successive states of a file system” refers only to snapshots and
9 whether “meta-data” in this context must be a block map file. The term “meta-data for successive
10 states of said file system” appears in claim 8 of the ’292 patent, which is set forth above, and
11 pertains to creating a plurality of read-only copies of a file system stored in blocks of a non-volatile
12 storage means. Both parties agree that such read-only copies are snapshots, and that this claim
13 pertains to creating snapshots. Claim 8 recites multiple steps for creating snapshots and recording
14 these snapshots as a copy of meta-data. Specifically, the claim requires that the meta-data first be
15 stored for successive states of said file system, then be copied “at each of a plurality of said states of
16 said file system,” then be marked. ’292 Patent at 26:1-15. At the end of this process, a snapshot has
17 been created. However, Sun’s construction renders this process circular and collapses the separate
18 steps. Under Sun’s construction, the meta-data that is stored in the first step of this process would
19 be the snapshot – more particularly, a file for recording the snapshots. However, the meta-data itself
20 cannot be the snapshot, because the meta-data is what the invention uses to make the snapshot.

21 In addition, the claim language shows that the snapshots are only a subset of the meta-data
22 for successive states that are stored in claim 8. The plain language of the claim notes that after the
23 meta-data is stored for successive states of the said file system, copies of that meta-data are made “at
24 each of a plurality of said states of said file system.” While “plurality” means that copies are made
25 at more than one of the successive states, it does not require that copies be made at every successive
26 state. In other words, the plain meaning of plurality does not require that copies be made at *all*
27 successive states. This, in turn, means that not all meta-data is marked. The inventor knew how to
28 refer to the entirety of a set, but chose not to do so. For example, in the same claim, the inventor

1 required that “each of said copies of said meta-data” be marked. Thus, all copies at this step are
2 marked. However, at the earlier copying stage, the inventor chose not to use this all-inclusive
3 language, and instead opted to use the phrase “at each of a plurality.” This claim, therefore, does not
4 require that all meta-data in these claims be stored as snapshots.

5 As to whether or not meta-data in this claim must be in the form of a block map file, NetApp
6 argues that the doctrine of claim differentiation mandates rejection of Sun’s proposed construction.
7 Specifically, dependent claims 9 and 10 add the limitations “means for recording multiple usage bits
8 per block” and “blockmap comprising multiple bit entries for each block,” respectively. NetApp
9 argues that because the “blockmap” limitation is the only difference between claims 9 and 10, that
10 limitation may not be read into independent claim 8. NetApp Opp. at 15; Ganger Decl. ¶¶ 37-38.
11 “In the most specific sense, ‘claim differentiation’ refers to the presumption that an independent
12 claim should not be construed as requiring a limitation added by a dependent claim.” Curtiss-Wright
13 Flow Control Corp. v. Velan, Inc., 438 F.3d 1374, 1380 (Fed. Cir. 2006). While claims 9 and 10
14 contain additional language that is more specific than just requiring a block map alone, and thus
15 Sun’s construction would not render these dependent claims entirely superfluous, these dependent
16 claims indicate that the block map requirement should not be imported into the independent claim.

17 In addition, the narrowing limitations of dependent claims 11-13 and 18-19 might be
18 excluded under Sun’s definition, because the proper construction for “meta-data” must be at least
19 broad enough to include different types of meta-data such as pointers to a hierarchical tree of blocks,
20 structures representing files of the file system, and a root structure referencing structures
21 representing files of said file system. See, e.g., ’292 Patent at Claim 11 (“wherein said meta-data
22 comprises pointers to a hierarchical tree of blocks”); Claim 12 (“wherein said meta-data comprises
23 structures representing files of said file system”); Claim 18 (“wherein said meta-data comprises a
24 root structure referencing structures”). While Sun argues that the requirement of a block map file
25 for recording snapshots of the file system in the independent claim does not exclude the presence of
26 the additional structures listed in the dependent claim terms, see Brandt Supp. Decl. ¶ 24, the use of
27 the “comprising” language in the dependent terms indicates that the meta-data in these dependent
28 terms must include the pointers, structures, etc. referred to in those dependent claims. Therefore, the

1 claim language indicates that the blockmap limitation should not be included in this construction.

2 Turning to the specification, Sun seems to blur the distinction between what meta-data is and
3 where it is stored. Data is stored in a file. Meta-data may be stored in a block map file. The
4 specification makes clear that meta-data does not comprise the files themselves. See '292 patent at
5 5:52-56 (WAFL uses files to store meta-data that describes layout of file system); 9:21-24 (meta-
6 data is kept in files). The blockmap file indicates which disk blocks are allocated. Id. at 5:57-58.
7 The block map file indicates whether a block is part of the active file system or a snapshot by storing
8 either a one or a zero in the bit position in the block map file. Id. at 10:7-15; 9:66-10:1; 4:39-43.
9 These files do not define what meta-data is. Therefore, Sun's proposed construction, insofar as it
10 equates meta-data with block map file, is incorrect.

11 Sun relies on the "Summary of the Invention" section, which states:

12 The present invention prevents new data written to the active file system
13 from overwriting "old" data that is part of a snapshot(s). It is necessary
14 that old data not be overwritten as long as it is part of a snapshot. This is
15 accomplished by using a multi-bit free-block map. Most prior art file
16 systems use a free block map having a single bit per block to indicate
17 whether or not a block is allocated. The present invention uses a block
18 map having 32-bit entries.

19 '292 Patent at 4:33-40. As Sun notes, these unambiguous statements, which refer to the invention
20 and not just a single embodiment, make clear that a block map is a necessary feature of the
21 invention. However, this summary does not discuss meta-data at all and does not state that it is
22 important that the meta-data be a block map file. Rather, it merely states that the present invention's
23 use of a block map is key. This description does not weigh in favor of incorporating these
24 limitations into the term at hand.

25 In addition, as NetApp argues, even in the preferred embodiment, the term at issue is not
26 limited to a block map file, because the meta-data that is copied to make a new snapshot is a root
27 inode, not a block map file. '292 Patent at 18:13-16; 18:65-19:6. In this preferred embodiment, the
28 process of creating a snapshot includes a snapshot inode that is similar to the root inode and is stored
in a fixed location inside the inode file. In addition, a copy is stored in the block map file.

In conclusion, Sun's proposed construction is not proper. At the same time, NetApp's
proposed construction simply repeats much of the language in the disputed claim term, which is not

1 particularly helpful to the jury. The specification explains that WAFL keeps information that
2 describes a file system in files known as meta-data. '292 Patent at 9:19-20. A snapshot is a read-
3 only copy of an entire file system at the particular instant when the snapshot is created. *Id.* at 17:66-
4 18:1. Claim 8 describes a method for creating snapshots. Therefore, the Court proposes construing
5 “meta-data for successive states of said file system” as “information describing a copy of the
6 structure of the active file system (such as a copy of a root inode) at a series of successive points in
7 time.” The parties shall meet and confer about the propriety of this proposed construction in light of
8 this claim construction order. The parties shall notify the Court as to whether they agree to this
9 construction, or submit a different agreed upon construction that comports with the guidance in this
10 order no later than **September 19, 2008**. In doing so, the parties do not waive any rights to preserve
11 their objections that their original proposed constructions were correct.

12 **10. “File System Information Structure”**

NetApp’s Proposed Construction	Sun’s Proposed Construction
A data structure containing information about the layout of a file system	Data structure that contains the root inode of a file system in a fixed location on disk

17 Here, the parties dispute whether or not the file system information structure contains a root
18 inode, and whether the data structure is in a fixed location on disk.

19 The term file system information structure appears in asserted claim 4, as well as non-
20 asserted claims 2, 5, and 7 of the '292 patent. NetApp’s proposed construction is “a data structure
21 containing information about the layout of a file system.” If the Court were to replace the term with
22 this proposed language, however, claim 4 would refer to “storing a first ‘data structure containing
23 information about the layout of a file system’ for a first consistency point in said non-volatile
24 storage means, said first data structure containing information about the layout of a file system’
25 comprising data describing a layout of said file system at said first consistency point of said file
26 system.” NetApp’s construction thus improperly repeats other language within the claim and
27 ascribes meaning to the term that is already implicit (indeed, explicit) in the remainder of the claim.
28 See, e.g., Mangosoft, Inc. v. Oracle Corp., 525 F.3d 1327, 1330-1331 (Fed. Cir. 2008) (refusing to

1 adopt proposed construction that ascribed no meaning to claim term that was “not already implicit in
2 the rest of the claim”). NetApp’s construction, therefore, would create a strange and awkward
3 redundancy in the claim language.

4 NetApp argues that Sun’s construction requiring a fixed location reads a limitation from
5 dependent claim 6 into independent claim 4, because claim 6 requires that the non-volatile storage
6 means “comprise fixed predetermined locations” of said means. However, this dependent claim
7 does not require that the single file information structure recited in claim 4 be fixed, so the doctrine
8 of claim differentiation is inapplicable. In addition, dependent claims 5 and 6 contain additional
9 limitations, such as overwriting copies of file system information structure, so Sun’s interpretation
10 would not render the dependent claims superfluous.¹⁰ As a whole, therefore, the claim language
11 indicates that NetApp’s proposed construction is redundant, and does not weigh against Sun’s fixed
12 location limitation. However, the Court must turn to the specification in order to construe the
13 meaning of the term.

14 As a preliminary matter, the specification seems to equate the file system information
15 structure with the “fsinfo” by placing “fsinfo” next to or within the phrase “file information
16 structure.” For example, the specification twice refers to the “file system information (fsinfo)
17 structure” and twice refers to the “file system information (fsinfo) block.” ’292 Patent at 9:33-36,
18 10:57-65, 13:63-64. The specification also equates the terms file information structure and block by
19 using the terms interchangeably. For example, the specification notes that Figure 15 illustrates a
20 fsinfo structure 1510, but then Figure 15 labels structure 1510 “fsinfo block.” *Id.* at Figure 15 and
21 10:57-58.

22 The specification states that the root inode is kept in a fixed location on disk referred to as
23 the fsinfo block. ’292 Patent at 9:33-35. See also id. at 10:57-64. Based on this language of the
24 specification, Sun argues that a person of ordinary skill in the art would recognize that file system
25 information structure must include the root inode and must be located at a fixed location. See
26 Brandt Decl. ¶ 83. According to Sun, the fixed, known location of the fsinfo block is necessary to

27 ¹⁰ Insofar as NetApp argues that the fact that claim 19 expressly recites a root inode
28 supports its argument, corresponding independent claim 8 does not contain the term “file system
information structure,” so these claims are not particularly helpful to its construction.

1 claim 4. Brandt Decl. ¶ 84. For example, the specification teaches the file system changes
2 consistency points after the root inode is written to disk. '292 Patent at 4:11-18, 12:16-20; see also
3 Brandt Decl. ¶ 83. Therefore, according to Sun's expert, the file system information structure
4 cannot operate to change consistency points absent the presence of the root inode in that structure.
5 Brandt Dec. ¶ 83. Because recovery of the file system is only possible if the root inode can be
6 located by the file system, and the specification teaches that the file system must be located at a
7 fixed location, it appears that the structure must include a root inode and be at a fixed location. See
8 '292 Patent at 9:27-35 (noting that the inode file may be written anywhere on disk, unlike the prior
9 art that writes inode tables to a fixed location on disk; the inode file is pointed to by a root inode that
10 is kept in a fixed location on disk referred to as the file system information block.); 10:57-65 (root
11 inode of file system is kept in a fixed location on disk so that it can be located during booting of the
12 file system); 11:3-5 (noting that two identical copies of the fsinfo structure are kept in fixed
13 locations on disk). Thus, according to the specification, the root inode is in a fixed location on disk,
14 which is referred to as the file system information block. The specification discloses this in
15 definitional language, and this requirement does not appear to be confined to the preferred
16 embodiment only.

17 NetApp notes that the specification also describes the file system information structure as
18 one that "also contains information including the number of blocks in the file system, the creation
19 time of the file system, etc. . . . [and] a check-sum." '292 Patent at 10:65-11:1. While this is true,
20 this portion of the specification seems to include such data as optional, while the root inode must be
21 included in the fsinfo structure. See above; see also '292 Patent at 10:57-11:3 (noting that root
22 inode is kept in fixed location on disk, and that except for the root inode, the miscellaneous
23 information such as number of blocks in the file system, creation time of the file system, etc., can be
24 kept in a meta-data file in an alternative embodiment).

25 However, while the specification indicates that the root inode is kept in a fixed location on
26 disk – the file system information block – the specification expressly teaches that the file system
27 need not necessarily use a root inode as the data structure when copying *snapshots*: "Snapshots are
28 created by duplicating the root data structure of the file system. In the preferred embodiment, the

1 root data structure is the root inode. However, any data structure representative of an entire file
2 system could be used.” *Id.* at 18:13-16. NetApp therefore argues that any data structure
3 representative of an entire file system could be stored in the fsinfo block in lieu of a root inode.
4 Ganger Decl. ¶ 48. However, this portion of the specification does not contain an explicit reference
5 to the file system information structure. Rather, this portion of the specification addresses snapshots
6 and describes “snapshot inodes” as only “similar” to root inodes, not identical. ’292 Patent at 18:12.
7 By contrast, the specification’s discussion of the file system information structure addressed above
8 is contained in the portion of the specification addressing the WAFL layout, which is fundamental to
9 the invention: “The present invention uses a Write Anywhere File-system Layout (WAFL),” *id.* at
10 5:48-49, although some “specific details” pertaining to disks might be varied. *Id.* at 5:40-43.
11 Therefore, even though the snapshot part of the specification discusses copying a root inode to form
12 a snapshot and then alludes to a possible alternative data structure that could be used to make a
13 snapshot, it does not state that the file system information block or structure can do without a root
14 inode. The abstract also discusses a file system that progresses from one self-consistent state to
15 another self-consistent state, in which the blocks on disk are rooted by a root inode. It further notes
16 that the method also creates snapshots, but does not limit the snapshot creation process to a root
17 inode data structure in the general discussion. In other words, the discussion of the file system and
18 snapshot creation are related, but separate.

19 In sum, the claim language and specification indicate that the “file system information
20 structure” contains a root inode and is on a fixed location on disk. The Court, therefore, adopts
21 Sun’s construction of this term, and construes it as “data structure that contains the root inode of a
22 file system in a fixed location on disk.”

23 **’211 Patent**

24 This patent is a continuation of the ’292 patent and is also directed to a method for keeping a
25 file system in a consistent state. It shares the same specification as the ’292 patent. The file system
26 here utilizes a root inode which contains pointers that directly point to a special meta-data file, the
27 inode file, that contains the inodes of all of the other files in the file system. ’211 Patent at 9:25-33.

28 ///

1 **11. “Root Inode”**

NetApp’s Proposed Construction	Sun’s Proposed Construction
An inode that points directly and/or indirectly to all the blocks in a consistent state of a ‘file system’ (as construed herein)	The index node data structure stored in a fixed location that roots a set of self-consistent blocks on the storage system that comprise the file system

7 The phrase “root inode” appears in independent claims 1, 9, and 17 of the ’211 patent. These
8 claims have similar limitations. Claim 9 is for: “a device comprising:

9 a processor;

10 a memory; and

11 a storage system including one or more hard disks;

12 wherein said memory and said storage system store a file system; and

13 wherein said memory also stores information including instructions executable by said

14 processor to maintain said file system, the instructions including steps of (a) maintaining an on-disk
15 **root inode** on said storage system, said on-disk **root inode** pointing directly and indirectly to a first
16 set of blocks on said storage system that store a first consistent state of said file system, and (b)
17 maintaining an incore **root inode** in said memory, said incore root inode pointing directly and
18 indirectly to buffers in said memory and a second set of blocks on said storage system, said buffers
19 and said second set of blocks storing data and meta-data for a second consistent state of said file
20 system, said second set of blocks including at least some blocks in said first set of blocks, with
21 changes between said first consistent state and said second consistent state being stored in said
22 buffers and in ones of said second set of blocks not pointed to by said on-disk inode.” ’211 Patent at
23 24:39-62 (emphasis added).

24 The fundamental difference between the parties’ constructions is whether the root inode is
25 required to reside in a fixed location. NetApp’s proposed construction, which includes pointing
26 directly and/or indirectly to all the blocks” is somewhat redundant, as the claim language itself
27 describes which sets of blocks or buffers the on-disk or incore root inodes point to. However, since
28 the claim language describes more specific “pointing,” it does not render NetApp’s claim language

1 completely superfluous, and the Court turns to the specification for guidance.

2 Multiple portions of the patent specification teach that the root inode roots a set of blocks
3 that comprise the file system in a consistent state. See, e.g., '211 Patent at 11:20-27, 11:65-67. The
4 specification teaches that the “root inode 1510B of a file system is kept in a fixed location on disk so
5 that it can be located during booting of the file system.” Id. at 10:59-61; see also id. at 9:33-35 (root
6 inode is kept on a fixed location on disk referred to as the file system information block described
7 below). The root inode is the inode that locates the inode file, which stores inodes for the other files
8 in the file system. Id. at 9:32-33. The invention teaches that the on-disk root inode is stored in a
9 fixed location on the fsinfo block to enable the location of the inode files, which are written to any
10 available locations on disk and may move around, in accordance with the “write anywhere” nature
11 of the file system. Id. at 9:19-24. See also Brandt Decl. ¶¶ 127-29. The abstract itself notes that the
12 “set of self-consistent blocks on disk is rooted by a root inode.” The use of the terms “root” and
13 “rooted” suggest that the root inode is not rootless but rather is fixed or stored at some set location,
14 although their primary meaning is serving as the base of the tree. '211 Patent at 11:24-25. These
15 portions of the specification address the on-disk root inode.

16 The claims reveal that there is also another type of root inode – the incore root inode. See,
17 e.g., Claim 1 of '211 Patent. The parties agree that “incore” means “in memory” as opposed to on-
18 disk. Incore memory contains contents not written to disk and not stored in a fixed location. See
19 '211 Patent at Fig. 8 (showing an incore inode that includes incore information, a copy of the on-
20 disk inode, pointers, etc.). An incore root inode, like other information temporarily stored in
21 memory, is often kept in whatever location in memory is allocated by the memory management
22 system, not in a fixed location. See Ganger Decl. ¶¶ 75-76. Sun concedes that the incore root inode
23 is stored in a buffer structure that is not located in a fixed location in memory and includes structures
24 absent from the on-disk root inode. See Sun Reply at 22-23. Sun, however, argues that while a
25 temporary copy of a root inode is stored in memory, the corresponding root inode is always stored in
26 a fixed position on the disk, so the incore copy is really still in a fixed location. However, the fact
27 that the incore’s original predecessor exists on a fixed location does not mean that the incore copy is
28 on a fixed location. In any event, Sun’s construction does not make this point clear. Consequently,

1 Sun’s construction improperly limits the root inode to a fixed location, because it does not
2 adequately distinguish the incore copy which need not be in a fixed location on a disk. However,
3 Sun’s construction would be remedied by adding that the in-memory copy of the root inode (the
4 incore copy) is not stored in a fixed position on disk.

5 NetApp also argues that even an on-disk root inode is not necessarily stored at a fixed
6 location. NetApp largely relies on Dr. Ganger’s declaration, but Dr. Ganger does not point to the
7 specification in making his observations, so this extrinsic evidence is not particularly persuasive.
8 See Ganger Decl. ¶¶ 79-80. Dr. Ganger notes that there are other mechanisms for ensuring that the
9 root inode can be located. For example, the file system could store a pointer to the root inode in a
10 fixed location. But a pointer that points to other inodes in non-fixed locations is really just another
11 way of describing a root inode. Dr. Ganger also notes that a file system could have a set of
12 predetermined locations that might hold the root inode. However, Dr. Ganger has not demonstrated
13 how the specification might contemplate such an embodiment, and the specification itself repeatedly
14 refers to a “fixed location,” as discussed above. However, if the set were small enough that the root
15 inode could be readily located, it might constitute a fixed location (or at least its equivalent), as
16 Sun’s expert acknowledged at the hearing.

17 NetApp also argues that the specification shows that at the precise moment a new snapshot is
18 created, at least one copy of the root inode for the active file system is not stored in a fixed location.
19 However, these portions of the specification are describing snapshot inodes, not root inodes. See,
20 e.g., ’211 patent at 17:6-63, 18:1-10, 18:17-19, 18:57-67, 19:5-8. This copy of the root inode is
21 referred to as a snapshot inode, which is created by duplicating the root data structure of the file
22 system. Id. at 18:1-15. This copy, therefore, is not the root inode itself. To the extent that the
23 specification, in discussing snapshot creation, notes that another data structure representative of the
24 entire file system could be used in lieu of a root inode, the term being defined here is “root inode,”
25 not “any data structure representative of an entire file system.” See id. at 18:7-10.

26 As to the language in Sun’s construction that the root inode “roots a set of self-consistent
27 blocks on the storage system that comprise the file system,” the summary of the invention teaches
28 that the “file system progresses from one self-consistent state to another self-consistent state. The

1 set of self-consistent blocks on disk that is rooted by the root inode is referred to as a consistency
2 point (CP).” Id. at 4:15-18; 11:64-65 (file system progresses from one self-consistent state to
3 another self-consistent state, and set of self-consistent block on disks is rooted by root inode 1510B).
4 The specification, therefore, supports this aspect of Sun’s construction.

5 To the extent that NetApp relies on nCube Corp. v. SeaChange Int’l Inc., 436 F.3d 1317,
6 1322 (Fed. Cir. 2006), for the proposition that it is inappropriate to read a limitation into a claim that
7 is “wholly apart from any need to interpret what the patentee meant by particular words or phrases in
8 that claim,” that case is inapposite. There, the Federal Circuit rejected a construction that imported a
9 limitation from a dependent claim into an independent claim. Here, Sun’s construction does not
10 read in limitations from a dependent claim, but describes the scope of the invention set forth in the
11 written description.

12 To the extent that NetApp incorporates its “and/or” construction of the term “directly and
13 indirectly” (discussed below), and the “file system” (discussed above), these aspects of NetApp’s
14 construction fail for the same reasons.

15 In sum, and for the reasons set forth above, the Court adopts Sun’s construction as modified
16 below, and finds that the term “on-disk root inode” means “the index node data structure stored in a
17 fixed location on disk that roots a set of self-consistent blocks on the storage system that comprise
18 the file system” and that the term “incore root inode” means “the copy of the on-disk root inode that
19 is contained in memory; this copy need not be in a fixed location in memory.”

20 **12. “Pointing Directly and Indirectly to Buffers in Said Memory and a Second Set of**
21 **Blocks on Said Storage System”**

NetApp’s Proposed Construction	Sun’s Proposed Construction
Pointing directly to blocks and/or buffers, and/or indirectly to blocks and/or buffers	Pointing directly and indirectly to buffers in said memory and pointing directly and indirectly to a second set of blocks on said storage system

26
27 The main dispute between the parties is whether or not this term requires both direct and
28 indirect pointing to both blocks and buffers, or can be satisfied by either indirect or direct pointing to

1 either blocks or buffers. The phrase appears in independent claims 1, 9, and 17 of the '211 patent.

2 These claims have similar limitations. Again, Claim 9 is for: “a device comprising:

3 a processor;

4 a memory; and

5 a storage system including one or more hard disks;

6 wherein said memory and said storage system store a file system; and

7 wherein said memory also stores information including instructions executable by said

8 processor to maintain said file system, the instructions including steps of (a) maintaining an on-disk

9 root inode on said storage system, said on-disk root inode pointing directly and indirectly to a first

10 set of blocks on said storage system that store a first consistent state of said file system, and (b)

11 maintaining an incore root inode in said memory, said incore root inode **pointing directly and**

12 **indirectly to buffers in said memory and a second set of blocks on said storage system**, said

13 buffers and said second set of blocks storing data and meta-data for a second consistent state of said

14 file system, said second set of blocks including at least some blocks in said first set of blocks, with

15 changes between said first consistent state and said second consistent state being stored in said

16 buffers and in ones of said second set of blocks not pointed to by said on-disk inode.” ’211 Patent at

17 24:39-62 (emphasis added).

18 A plain reading of the claim language, “root inode pointing directly and indirectly to buffers
19 in said memory and a second set of blocks on said storage system,” supports Sun’s construction.

20 The use of “and” rather than “or” indicates that the root inode points both directly to some buffers

21 and indirectly to some other buffers, and it also requires that the inode point directly to some blocks

22 and indirectly to other blocks. As the Federal Circuit has explained, “[a] claim must be read in

23 accordance with the precepts of English grammar.” See In re Hyatt, 708 F.2d 712, 714 (Fed. Cir.

24 1983). The inventor chose to use the language “and,” which has a conjunctive meaning, rather than

25 “or.” NetApp’s proposed construction, however, replaces “and” with “and/or,” which disregards the

26 plain claim language. Under NetApp’s construction, any one of at least eight different alternatives

27 would satisfy the claim language, which ignores the inventor’s use of the conjunctive rather than

28 disjunctive.

1 NetApp argues that the language surrounding the disputed term supports its position that
2 there is no requirement for both direct and indirect pointers to both buffers and blocks. NetApp
3 notes that there is a parallel structure in the claim in which the first set of blocks (limitation a) is the
4 group that represents the first consistent state of the file system, and the second set of blocks
5 together with buffers (limitation b) is the group that represents the second consistent state of the file
6 system. NetApp argues that each of these groups is to be treated as a whole. But since the disputed
7 term only appears in the second group, this structure instead indicates that the “and” used to link the
8 “buffers and second set of blocks” describes those things as a “single combined group.” In contrast,
9 NetApp’s use of “and/or” would break up this single combined group into two parts. Nor does
10 Sun’s construction apply the adverb “directly and indirectly” separately to the blocks and buffers, as
11 NetApp argues. While Sun separates out the requirements in this way for clarity’s sake (directly and
12 indirectly to blocks and directly and indirectly to buffers), Sun’s construction merely means that the
13 term applies “directly and indirectly” to both “blocks and buffers,” whereas NetApp’s construction
14 does not.

15 At the hearing, NetApp provided some examples of colloquial speech in which “and” is used
16 disjunctively, for example: “the stars and stripes are red, white and blue,” “we were eating and
17 drinking milk and cookies,” and “children and adults were drinking milk and juice.” However, these
18 examples are not persuasive for two reasons. First, these examples are of casual speech, not the
19 formal written language of patents, where the inventors choose their words in a much more
20 deliberate manner. In addition, these examples are not parallel to the term at issue, because they
21 combine subject nouns and their transitive verbs with their direct noun objects or adjectives,
22 whereas the disputed term combines subject nouns and verbs with adverbs and prepositional objects.

23 NetApp’s reliance on Ortho-McNeil Pharm., Inc. v. Mylan Labs., Inc., 520 F.3d 1358, 1362
24 (Fed. Cir. 2008) is misplaced, as the claim language here is quite dissimilar. There, the Court noted
25 that “[t]he claim also does not use and in isolation but in a larger context that clarifies its meaning.
26 Specifically, *and* appears in conjunction with the adverbs *independently* and *together*. As the district
27 court explained, these terms signal that *and* links alternatives that occur under the different
28 conditions of independence or togetherness.” *Id.* (construing *and* in claim that stated “R2, R3, R4,

1 and R5 are independently hydrogen or lower alkyl *and* R2 and R3 and/or R4 and R5 together may be
2 a group of formula”) (emphasis in original). The Court found that it would not make sense to
3 require two mutually exclusive subsets of compounds, one in which the compounds were
4 independent and one in which the same compounds were not independent and constituted a group.
5 Here, there is no similar use of the phrase “independently and together.” In addition, unlike here,
6 the Court in Ortho-McNeil found that construing “and” in the conjunctive would have rendered
7 dependent claims meaningless. Id.

8 The claim language, therefore, supports Sun’s construction. Turning to the specification, the
9 specification teaches that inodes may point to “indirect and/or direct blocks.” ’211 Patent at 7:55-
10 57. This part of the specification, however, is not referring to the root inode, which plays the special
11 role of rooting the entire consistent file system. In fact, the inventor’s use of “and/or” in the
12 specification with regard to non-root inodes undermines NetApp’s argument that the word “and” in
13 the claims has the same meaning as “and/or,” as the inventor clearly knew how to use that broader
14 language.

15 Generally, the specification makes clear that the buffers and blocks to which the incore root
16 inode points serve a specific purpose: to store data and meta-data for a second consistent state of the
17 file system. ’211 Patent at 24:4-6, 55-57, 25:37-38. The file system maintains at least three
18 bookkeeping files in the form of meta-data files, to keep data in a consistent state: the inode file, the
19 block map file, and the inode map file. ’211 Patent at 9:19-24; Brandt Decl. ¶ 105. The inode file
20 contains inodes describing all other files in the file system, including the block map and inode map
21 files. ’211 Patent at 9:26-27. The inode file is pointed to by an inode referred to as the root inode.
22 Id. at 9:32-33. The block map file keeps track of which blocks are allocated in the current file
23 system, which blocks are used by snapshots, and which blocks are free for the file system to write
24 data onto. Id. at 9:50-65. The inode map file is used by the file system to keep track of which
25 inodes in the inode file are unallocated. Id. at 10:20-49.

26 The specification teaches that the root inode points to the inode file directly and to the other
27 meta-data files and data files indirectly. Id. at 11:20-27 (describing root inode that describes inode
28 file that in turn describes the rest of the files in the file system including meta-data files, so that the

1 root inode is viewed as the root of the tree of blocks). Thus, the root inode points directly to the
2 inode, and indirectly to the other data files in order to serve the claimed function of rooting “buffers
3 and said second set of blocks storing data and meta-data for a second consistent state of said file
4 system.” Brandt Decl. ¶ 111. The incore root inode, on the other hand, points to the inode file
5 buffer holding inodes for modified files and the block map file, and also points indirectly to buffers
6 holding modified data for these files. ’211 Patent at 16:4-17:57 & Fig. 17L. Once these buffers are
7 sent to disk, the incore root inode directly and indirectly points to a new set of blocks comprising the
8 changed data and inodes (*id.* at 17:35-37) as well as directly to the unchanged blocks of the inode
9 file and indirectly to the inode’s corresponding unchanged data blocks. Thus, according to Sun’s
10 expert, the specification teaches that the incore root inode points directly and indirectly to buffers
11 and directly and indirectly to a second set of blocks on the said storage system. Brandt Decl. ¶ 112-
12 120.

13 While NetApp argues that there is no requirement that the incore root inode have both direct
14 and indirect pointers, and that, for example, a file system can be in a state where the incore root
15 inode points to all buffers directly, NetApp relies heavily on extrinsic expert opinion and does not
16 point to any clear language in the specification supporting its arguments. Rather, Dr. Ganger notes
17 that Figures 8 and 3 show that if all 16 blocks to which the incore root inode points directly are
18 modified, then the 16 buffer pointers will point directly to the contents of the second consistent state,
19 but that no block pointers will do so. Thus, he concludes, if all 16 blocks of the root inode of the
20 first consistent state change, the incore root inode for the second consistent state will not point
21 directly to any blocks in the second consistent state. Ganger Decl. ¶ 64. However, Dr. Brandt
22 contests this argument, maintaining that the root inode includes a copy of the on-disk inode, which
23 points to the blocks comprising the first consistent state even when all such blocks have undergone
24 changes and are represented by buffers. Supp. Brandt Decl. ¶ 41 (citing ’211 patent at 7:58-60 and
25 Fig. 10 Element 1010D).

26 Dr. Ganger also argues that there would not be any indirect pointers to buffers by the root
27 inode group in a situation in which the only changes between the first consistent state and second
28 consistent state were to inodes. Ganger Decl. ¶ 65. However, while Dr. Ganger refers to the

1 specification’s description of the general file system, he does not point to any support in the
 2 specification for his inferences on which he bases his conclusion. In addition, Dr. Brandt contests
 3 this conclusion on the ground that Ganger incorrectly assumes that the incore inodes are the same as
 4 on-disk inodes, among other things. Supp. Brandt Decl. ¶ 42. In sum, the experts contest whether
 5 the specification discloses systems in which the incore root inode may not have both direct and
 6 indirect pointers. However, NetApp’s expert Dr. Ganger has not pointed to any persuasive intrinsic
 7 evidence in support, so his extrinsic evidence does not counsel interpreting this claim term in a
 8 manner inconsistent with the clear claim language.

9 In sum, and for the reasons set forth above, the Court adopts Sun’s construction and finds
 10 that the term means “pointing directly and indirectly to buffers in said memory and pointing directly
 11 and indirectly to a second set of blocks on said storage system.”

12 **13. “Consistent State”/”State of a File System”**

NetApp’s Proposed Construction	Sun’s Proposed Construction
A set of blocks on disk, rooted by a ‘root inode’ (as construed herein), that includes all the blocks required for the data and file structure of a ‘file system’ (as construed herein)	A set of storage blocks for that file system that includes all blocks required for the data and [file] structure of the file system

17
 18 The parties originally appeared to dispute whether it was proper to add the word “file” before
 19 the phrase “structure of the file system.” However, Sun has now agreed to add this language to its
 20 proposed construction, so this is no longer at issue. The dispute now boils down to whether or not
 21 the set of blocks needs to be rooted by a root inode. To begin its analysis, the court first turns to the
 22 claims themselves. The term “consistent state” appears in the ’211 patent, while the term “state of a
 23 file system” appears in the ’292 patent. The parties agree that the construction of these two claim
 24 terms should be the same. The two patents share the same specification. The terms appear in claims
 25 1-3, 9-11, and 17-19 of the ’211 patent and claim 8 of the ’292 patent.

26 Turning to the claim language, NetApp’s proposed construction requiring that the consistent
 27 state be rooted by a root inode is quite repetitive of the language already in the claims. See
 28 Mangosoft, Inc. v. Oracle Corp., 525 F.3d 1327, 1330-1331 (Fed. Cir. 2008) (refusing to adopt

1 proposed construction that ascribed no meaning to claim term that was “not already implicit in the
2 rest of the claim”). For example, Claim 1 of the ’211 patent refers to an “on-disk root inode pointing
3 directly and indirectly to a first set of blocks on said storage system that store a first consistent state
4 of said file system.” Claim 9 refers to an “on-disk root inode pointing directly and indirectly to a
5 first set of blocks on said storage system that store a first consistent state of said file system.” In
6 other words, the claims themselves already describe the root inode’s relationship to the consistent
7 states of the file system. NetApp’s inclusion of the phrase “rooted by a root inode (as construed
8 herein)” in its proposed construction of “consistent state” creates a confusing repetition, because
9 “root inode” is a separate claim limitation. Specifically, if one inserts NetApp’s construction of
10 “consistent state” into claim 1, for example, the following repetitive statement results: “maintaining
11 an on-disk root inode on said storage system, said on-disk root inode pointing directly and indirectly
12 to a first set of blocks on said storage system that store a first [set of blocks on disk, rooted by a root
13 inode, that includes all the blocks required for the data and file structure of a file system] of said file
14 system.” ’211 Patent at 23:64-67 (bracketing NetApp’s proposed construction). If one were to also
15 insert NetApp’s proposed construction of “root inode” into the claim, an even more confusing
16 tautology results, as NetApp’s construction of “root inode” itself refers to consistent states:
17 “maintaining an on-disk [inode that points directly and/or indirectly to all the blocks in a consistent
18 state of a file system] on said storage system, said on-disk [inode that points directly and/or
19 indirectly to all the blocks in a consistent state of a file system] pointing directly and indirectly to a
20 first set of blocks on said storage system that store a first [set of blocks on disk, rooted by an inode
21 that points directly and/or indirectly to all the blocks in a consistent state of a file system, that
22 includes all blocks required for the data and file structure of a file system] of said file system.” Id.
23 (bracketing NetApp’s proposed construction).

24 In support of its proposed construction, Sun points to NetApp’s 7,174,352 patent (’352
25 Patent), which shares most of the inventors of the ’292 patent and was filed roughly eight years after
26 the ’292 and ’211 specifications were drafted. The ’292 and ’352 patents are continuations of patent
27 application number 71, 643, which was abandoned. The ’352 patent’s specification concerns
28 NetApp’s WAFL file system which is disclosed in the ’292 and ’211 patents, and the ’352 patent

1 incorporates the disclosures of the '292 patent. NetApp acknowledged that the patents were
2 "cousins." The '352 patent includes an express definition of consistent: "As used herein, the term
3 'consistent,' referring to a file system (or to storage blocks in a file system), means a set of storage
4 blocks for that file system that includes all blocks required for the data and file structure of that file
5 system." '352 Patent at 4:24-27 (Williamson Decl., Ex. D). "Thus, a consistent file system stands
6 on its own and can be used to identify a state of the file system at some point in time that is both
7 complete and self-consistent. *Id.* at 4:28-30. This definition is identical to Sun's proposed
8 construction, which now includes the "file" language. Sun's definition, therefore, finds support in
9 this related patent's specification.

10 There are passages in the '292 patent specifications that support NetApp's construction and
11 describe consistent states or consistency points that are rooted by a root inode. For example, the
12 specification notes that "[c]hanges to the file system are tightly controlled to maintain the file system
13 in a consistent state. The file system progresses from one self-consistent state to another self-
14 consistent state. The set of self-consistent blocks on disk that is rooted by the root inode is referred
15 to as a consistency point." '292 Patent at 4:9-13. Figure 16 and the accompanying description of a
16 file system in a "consistent state" shows the consistent state of the system as being rooted by a root
17 inode. *Id.* at 11:6-27 & Fig. 16. However, while consistency point and consistent state are closely
18 related concepts, it is not entirely clear that the terms are identical. For example, "consistency
19 point" is claimed in claim 4 of the '292 patent, and "consistent state" appears in claim 8 of the '292
20 patent. Sun notes that NetApp's construction equating consistent states with consistency points
21 cannot be correct, because the second consistent state is achieved while the file system is still at a
22 first consistency point. However, as NetApp points out, the '352 patent lends some support to
23 NetApp's position, as it refers to a "state of the file system at some point in time that is both
24 complete and self-consistent," which appears to refer to both a consistent state and consistency point
25 in a similar way. '352 Patent at 4:24-37. Regardless of whether consistency state and consistency
26 point are synonymous, however, NetApp's proposed construction is too redundant of the
27 surrounding claim language, as discussed above.

28 Finally, to the extent that NetApp's construction incorporates its definitions of "file system"

1 and “root inode” that the Court repeated above, this construction is similarly flawed. In sum, and for
2 the reasons set forth above, the Court adopts Sun’s construction and finds that the term means “a set
3 of storage blocks for that file system that includes all blocks required for the data and file structure
4 of the file system.”

5 **’715 Patent**

6 The ’715 patent is directed to a technology for controlling data storage over arrays of disk
7 drives. ’715 patent at 1:7-9. Specifically, the technology improves the read and write efficiency of
8 a “Redundant Array of Inexpensive Disks” (“RAID array”) in combination with a file system.
9 RAID systems can provide greater storage capacity than a traditional single disk drive system
10 because additional disk drives can be added to the array, and RAID systems have high data transfer
11 rates. The techniques used in RAID systems generally allow data to be protected even when
12 individual disks storing data experience failure, and when a single disk fails, it can be replaced
13 without shutting down the entire system. The RAID array accomplishes this using a “stripe”
14 technique. Data blocks are written across the array of disk drives in “stripes.” The patent defines a
15 “stripe” as including “one storage block on each disk drive in an array of drives in the system.” *Id.*
16 at 1:37-39. Essentially, data is written over an array of disks as a series of stripes, and at least one of
17 these disks is used to store parity (also called redundancy data), and the others are used to store data.
18 The parity information can be used to reconstruct a data file from a failed disk.

19 There is a conflict between optimizing the functionality of file systems and individual disks
20 on the one hand, and optimizing the performance of a RAID on the other. File systems and disks are
21 optimized by writing large amounts of related data contiguously on an individual disk. RAIDS, on
22 the other hand, tend to be optimized by writing a full stripe all at once in order to minimize the
23 amount of time spent reading data from disks in order to recalculate parity. However, when
24 organized in this way, related data tends to be broken up into small pieces and performance suffers.
25 The invention of the ’715 patent is directed to optimizing both disk storage and parity recalculation.
26 Specifically, the patent discloses a data structure, called an “association” for associating data blocks
27 with storage blocks over the stripes. The multiple requests are ultimately written on multiple disks
28 in the RAID in order to write whole stripes at a time.

1 **14. “Associating the Data Blocks with One or More Storage Blocks Across the Plurality**
2 **of Stripes as an Association”/”the Association to Associate the Data Blocks with One or**
3 **More Storage Blocks Across the Plurality of Stripes”**

NetApp’s Proposed Construction	Sun’s Proposed Construction
NetApp contends that neither phrase requires construction given their plain and ordinary meaning. In the event the Court construes the terms, NetApp proposes slightly different constructions for the phrases: “creating a data structure that relates data blocks to locations on more than one stripe,” and “the data structure relating data blocks to locations on more than one stripe,” respectively.	Sun contends both of these phrases are indefinite under 35 U.S.C. § 112 ¶ 2, and cannot be reasonably construed because they fail to particularly point out and distinctly claim the subject matter which the applicants regard as their invention. However, to the extent the Court construes the phrases, they mean “associating each data block with a respective one of the storage blocks across the plurality of stripes as an association.”

10 To begin its analysis, the court first turns to the claims themselves. The first term appears in
11 claims 21 and 52 of the ’715 patent, and the second claim appears in claim 39 of that patent.

12 Claim 21 is for “a method for controlling storage of data, comprising:

13 receiving one or more write requests associated with data blocks; receiving topological
14 information associated with storage blocks configured in a plurality of parallel stripes of a storage
15 system;

16 **associating the data blocks with one or more storage blocks across the plurality of**
17 **stripes as an association;**

18 and writing the data blocks, in response to the association, to the one or more storage devices
19 in a single write request.” ’715 Patent at 21:40-52 (emphasis added).

20 Claim 39 is for a “storage system, comprising:

21 a file system, the file system to receive one of more write requests associated with data
22 blocks;

23 a storage device manager, the storage device manager to generate topological information of
24 storage blocks configured in a plurality of parallel stripes of one or more storage devices, and to
25 send the topological information to the file system; and

26 an association generated in the file system, **the association to associate the data blocks**
27 **with one or more storage blocks across the plurality of stripes** of the one or more storage
28 devices, the association to be sent to the storage device manager, the storage device manager to

1 write the data blocks, in response to the association, to the one or more storage blocks as a single
2 write request.” ’715 Patent at 22:25-40 (emphasis added).

3 Sun first argues that the claim is indefinite under 35 U.S.C. § 112 ¶ 2. A claim is indefinite
4 under this paragraph, if “it would be apparent to one of skill in the art, based on the specification,
5 that the invention set forth in the claim is not what the patentee regarded as his invention.” Allen
6 Eng’g Corp. v. Bartell Indus., Inc., 299 F.3d 1336, 1349 (Fed. Cir. 2002). A claim is also indefinite
7 under that paragraph when it is “inherently inconsistent” or when its “legal scope is not clear enough
8 that a person of ordinary skill in the art could determine whether a particular [product] infringes or
9 not.” Geneva Pharms., Inc. v. GlaxoSmithKline PLC, 349 F.3d 1373, 1384 (Fed. Cir. 2003).

10 Of course, claims are not indefinite merely because they present a difficult task of claim
11 construction. Instead, “[i]f the meaning of the claim is discernible, even though the task may be
12 formidable and the conclusion may be one over which reasonable persons will disagree,” the claim
13 is “sufficiently clear to avoid invalidity on indefiniteness grounds.” Exxon Research & Eng’g Co. v.
14 United States, 265 F.3d 1371, 1375 (Fed. Cir. 2001) (citations omitted). Proof of indefiniteness
15 requires an exacting standard because claim construction often poses a difficult task over which
16 “expert witnesses, trial courts, and even the judges of this court may disagree.” Id. “Nevertheless,
17 this standard is met where an accused infringer shows by clear and convincing evidence that a
18 skilled artisan could not discern the boundaries of the claim based on the claim language, the
19 specification, and the prosecution history, as well as her knowledge of the relevant art area.”
20 Halliburton Energy Servs. v. M-I LLC, 514 F.3d 1244, 1250 (Fed. Cir. 2008).

21 Sun argues that the phrase is indefinite because it is flatly inconsistent with and impossible
22 under the teaching in the specification. The claim language requires “one or more storage blocks
23 across a plurality of stripes,” but under the invention, it is impossible for any “one” storage block to
24 be “across a plurality of stripes.” The specification provides that “a stripe includes one storage
25 block on each disk drive in an array of disk drives.” ’715 Patent at 1:37-39. A storage block is a
26 component of a single stripe, is located on a single disk, and does not traverse a plurality of stripes.
27 Id. at 6:35-38 & Figs. 1A&B. Each square in Figure 1B is a separate storage block, and no one
28 storage block is located in more than one stripe. See also id. at Fig. 8.

1 NetApp concedes that one storage block cannot exist on multiple stripes at the same time.
2 Yet NetApp’s proposed construction recognizes that a single storage block only exists on one stripe,
3 as NetApp expunges the infeasible “one or more” language from the claims so that “data blocks” are
4 related to storage block locations on more than one stripe. The Court, however, cannot rewrite
5 claims to preserve their validity. Allen Eng’g Corp., 299 F.3d at 1349 (“It is not our function to
6 rewrite claims to preserve their validity.”). NetApp argues that in the context of the specification,
7 the claims mean only that the majority of the time, storage blocks will be written across more than
8 one stripe. According to NetApp, the rare write request that is too small to occupy blocks on
9 multiple stripes should be included in the claim. Including this single stripe scenario in the claim,
10 however, contradicts the plain claim language, which indicates that one or more than one storage
11 block will be across a *plurality* of stripes. While “[i]f one skilled in the art would understand the
12 bounds of the claim when read in light of the specification, then the claim satisfies section 112
13 paragraph 2,” Exxon Research & Eng’g Co. v. United States, 265 F.3d 1371, 1375 (Fed. Cir. 2001),
14 the problem here is that NetApp’s proposed construction both rewrites the plain language of the
15 claims and contradicts the specification.

16 Nor is NetApp’s extrinsic evidence particularly persuasive, even if it could trump the
17 intrinsic evidence. NetApp’s expert claims that “one of ordinary skill in the art would understand
18 that the term ‘one or more’ covers a degenerate (abnormal) case only [in which the association is
19 only a single storage block on a single stripe], because, in practical terms using more than one stripe,
20 and therefore more than one storage block, is an important feature of the data structure . . .” Ganger
21 Decl. ¶ 94. Thus, he concedes that using more than one stripe is an important feature. Dr. Ganger
22 also argues that in the “degenerate case,” even though the data block associated with a single storage
23 block itself does not itself extend across a plurality of stripes, the storage blocks as a whole exist
24 across a plurality of stripes, so one data block can be associated with one storage block across a
25 plurality of stripes. But Dr. Ganger cites no intrinsic evidence in support of this convoluted and
26 strained interpretation, which, as discussed above, contradicts the plain language of the claim. And,
27 of course, Dr. Brandt disputes this conclusion. Hence, one of skill in the art would not know if a
28 single storage block in one stripe would be within the claim scope or not. “That is the epitome of

1 indefiniteness.” Geneva Pharms., Inc. v. GlaxoSmithKline PLC, 349 F.3d 1373, 1384 (Fed. Cir.
2 2003). Moreover, NetApp undermines its own argument by noting that the specification
3 distinguishes prior art on the basis that the prior art does not generally carry out writes to multiple
4 stripes. See, e.g., Ganger Decl. ¶ 95 (“Indeed, the specification frequently describes multi-stripe
5 writes and distinguishes itself from the prior art on the basis of the prior art not performing writes to
6 multiple stripes.”); Ho Decl., Ex. L at 20 (Preliminary Feb. 22, 2005 Amendment) (“applicant goes
7 one step further by mapping each data block of the single write request with a storage block across a
8 plurality of stripes of the storage system before transmitting the buffered write request. . .”).

9 Consequently, it does not appear from the prosecution history that this “degenerate case,” in which
10 only one single storage block is associated with a single stripe, was contemplated by the invention.

11 In sum, because it is unclear (at best) to one of ordinary skill in the art as to whether the
12 claim could encompass the so-called “degenerate case,” and because NetApp’s interpretation
13 contradicts the plain meaning of the actual claim language, the claim is indefinite. Sun has shown
14 by clear and convincing evidence that a skilled artisan could not discern the boundaries of the claim
15 based on the claim language, the specification, the prosecution history, and knowledge of the
16 relevant art area. Halliburton Energy Servs., 514 F.3d at 1250.

17 Because the claim is indefinite, the Court need not determine which of the parties’ proposed
18 constructions is proper. Sun’s construction would have required a one-to-one correspondence
19 between each data block and an associated storage block, and NetApp’s would not. As this issue
20 relates to the indefiniteness issue discussed above, however, the Court notes that it agrees that the
21 intrinsic evidence requires a one-to-one correspondence. The abstract clearly states that “[e]ach data
22 block of the data to be written is associated with a respective one of the storage blocks. . .” The
23 specification states that the block layout information “associates each data block of the buffered
24 write requests 41 with a storage block in a group of storage blocks.” ’715 Patent at 9:16-19; see also
25 id. at 13:37-39 (referring to one-to-one association with data blocks); 13:20-22 (“each data block is
26 stored at its associated storage block in the group”); 13:2-5 (“unallocated storage blocks are selected
27 in the array 20A for storage of a corresponding number of data blocks associated with buffered write
28 requests.”). The specification, therefore, indicates that each data block is associated with a

1 respective one of the storage blocks. NetApp does not contradict this evidence, nor does it point to
2 any examples in the specification describing something other than a one-to-one correspondence.
3 Therefore, the specification strongly supports Sun’s construction.

4 The prosecution history supports this requirement as well, as the patent applicant
5 distinguished prior art on the basis that this invention “associat[ed] each data block with a respective
6 one of the storage blocks, for transmitting the association to a storage device manager for processing
7 of the single write transaction.” Williamson Decl., Ex. F at 17; see also id. at 17-18 (“Applicant
8 goes one step further by associating each data block of the single write request with a storage block
9 of the storage system.”); Ex. H at 20 (distinguishing prior art after subsequent rejection, noting that
10 applicant claimed novel “mapping each data block with a respective one of the storage blocks across
11 a plurality of stripes”); Ex. I at 15 (“associating each data block with a respective one of the storage
12 blocks”). In allowing the claims, the examiner noted that the prior art did not teach buffering write
13 requests, “associating each data blocks to be stored with a respective one of the storage blocks
14 across a plurality of stripes for a single write operation.” Williamson Decl., Ex. J at 2.¹¹

15 However, as discussed above, the phrase “associating the data blocks with one or more
16 storage blocks” in the claim language indicates that there is not a one-to-one correspondence, as the
17 claim language requires multiple data blocks, but allows only one storage block in certain situations.
18 Because this claim language cannot be reconciled with the specification, which clearly requires a
19 one-to-one correspondence, this finding that there is a one-to-one correspondence in this claim term
20 further supports the finding that the claim is indefinite. In sum, and for the reasons discussed above,
21 the Court finds that the claim is indefinite under 35 U.S.C. § 112 ¶ 2.

22 **III. CONCLUSION**

23 In accordance with the foregoing, and for the reasons discussed above, the court construes
24 the parties’ disputed terms as follows:

25 ¹¹ Because the claim is indefinite, the Court need not resolve the parties’ other disputes
26 regarding construction of the term. However, the Court briefly notes that NetApp does not provide any
27 support for its use of the term “locations” in lieu of “storage blocks,” and replacing this claim language
28 with broader language seems improper. NetApp also argues that its construction makes clear that the
data structure called an “association” associates data blocks with storage blocks, while Sun’s
construction does not. At the hearing, however, Sun conceded that should the term be construed, it did
not object to the definition “data structure” for the term “association.”

1 1. “Domain name” means “a unique name that has a unique numerical IP address
2 associated with it.”

3 2. “Server identification data” means “information that uniquely identifies one server
4 from other servers, that is more human user-friendly than the server’s IP address, domain name,
5 URL, or hypertext link.”

6 3. “Responsive to writing a data record to said one redundancy group”/“Responsive to
7 writing a data record to one of said redundancy groups” means “after and in reaction to the writing
8 of a data record to a single redundancy group.”

9 4. “Responsive to the receipt of a stream of data records from said data processor”/
10 “Responsive to the receipt of a stream of the data records from said data processor” means “after and
11 in reaction to the receipt of data records from a processor.”

12 5. “Completing [a] Write Operation Within [a] Local Processing Node” / “Completing
13 [a] Write Operation With Respect to [a] Processor” is not yet construed. The parties shall submit an
14 agreed upon construction no later than **September 19, 2008**, as discussed above.

15 6. “Portion [of a] Communication” means “part of the data stream, where the part
16 includes elements of the data stream.”

17 7. “Element [of a] Communication” means “a constituent part of a portion.”

18 8. “Non-volatile storage means” is a means-plus-function limitation governed by 35
19 U.S.C. § 112, ¶ 6, wherein the function is storing blocks of data of a file system so that the data is
20 not lost in the absence of power, and the structure is a disk or disk arrays that are suitable for storing
21 4 KB blocks.

22 9. “Meta-data for successive states of said file system” tentatively means “information
23 describing a copy of the structure of the active file system (such as a copy of a root inode) at a series
24 of successive points in time.” The parties are to meet and confer about the propriety of this
25 proposed construction and may submit a different agreed upon construction that comports with the
26 guidance in this order no later than **September 19, 2008**.

27 10. “File system information structure” means “data structure that contains the root inode
28 of a file system in a fixed location on disk.”

1 11. “On-disk root inode” means “the index node data structure stored in a fixed location
2 on disk that roots a set of self-consistent blocks on the storage system that comprise the file system,”
3 and “incore root inode” means “the copy of the on-disk root inode that is contained in memory; this
4 copy need not be in a fixed location in memory.”

5 12. “Pointing directly and indirectly to buffers in said memory and a second set of blocks
6 on said storage system” means “pointing directly and indirectly to buffers in said memory and
7 pointing directly and indirectly to a second set of blocks on said storage system.”

8 13. “Consistent state”/”State of a file system” means “a set of storage blocks for that file
9 system that includes all blocks required for the data and file structure of the file system.”

10 14. “Associating the Data Blocks with One or More Storage Blocks Across the Plurality
11 of Stripes as an Association”/”the Association to Associate the Data Blocks with One or More
12 Storage Blocks Across the Plurality of Stripes” are indefinite phrases under 35 U.S.C. § 112 ¶ 2

13 All of these terms shall be so construed where they appear in the claim preambles, as well as
14 in the body of the claims.

15 IT IS FURTHER ORDERED that a case management conference is set for **October 3, 2008**
16 at 3:00 p.m. to discuss setting further dates in the case. The parties shall file a joint case
17 management statement no later than September 26, 2008. Proposals for scheduling summary
18 judgment motions should take into account the timing of the claim construction in related case
19 number 08-1641, and whether there should be a limit on the number of such motions and whether
20 such motions should be staggered and prioritized.

21 **IT IS SO ORDERED.**

22 Dated: September 10, 2008

Elizabeth D. Laporte

ELIZABETH D. LAPORTE
United States Magistrate Judge

23
24
25
26
27
28