IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re patent of: Baek et al.	§	Petition for *Inter Partes* Review
	§	
U.S. Patent No. 6,978,346	§	Attorney Docket No.: 47415.430
	§	
Issued: December 20, 2005	§	Customer No.: 112792
	§	
Title: APPARATUS FOR REDUNDANT INTER-CONNECTION BETWEEN MULTIPLE HOSTS AND RAID	§	Real Parties in Interest: Dell Inc., Hewlett-Packard Company, and NetApp, Inc.
	§	
	§	
	§	

PETITION FOR INTER PARTES REVIEW

Mail Stop Patent Board
Patent Trial and Appeal Board
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Sir:

Dell Inc., Hewlett-Packard Company, and NetApp, Inc. ("Petitioners") petition the Patent Trial and Appeal Board to institute an *inter partes* review under 35 U.S.C. §§ 311-319 of claims 1-9 of United States Patent No. 6,978,346 ("the '346 Patent," Exhibit DHPN-1001) that issued on December 20, 2005, to Sung-Hoon Baek et al., resulting from U.S. Patent Application No. 09/753,245, filed on December 29, 2000, and claiming priority to Korean Patent Application No. 2000-54807 filed on September 19, 2000. According to USPTO records, the '346 Patent is assigned to the Electronics and Telecommunication Research Institute.

TABLE OF CONTENTS

I. Mandatory Notices

A. Real Parties-in-Interest

The real parties-in-interest are Dell Inc., Hewlett-Packard Company, and NetApp, Inc.

B. Related Matters

As of the filing date of this petition and to the best knowledge of the petitioner, the '346 Patent is involved in the following litigation:

Safe Storage LLC v. StoneFly, Inc., 1-13-cv-01152; *Safe Storage LLC v. Int'l Business Machines Corp.,* 1-13-cv-01151; *Safe Storage LLC v. Emulex Corporation et al,* 1-13-cv-01150; *Safe Storage LLC v 3PAR Inc.,* 1-13-cv-01088; *Safe Storage LLC v Oracle America Inc.et al,* 1-13-cv-01089; *Safe Storage LLC v ATTO Technology Inc. et al.,* 1-13-cv-01090; *Safe Storage LLC v. VMware Inc.,* 1-13-cv-00928; *Safe Storage LLC v. Promise Technology Inc.,* 1-13-cv-00927; *Safe Storage LLC v. Nexsan Corporation,* 1-13-cv-00931; *Safe Storage LLC v. Overland Storage Inc.,* 1-13-cv-00932; *Safe Storage LLC v. IQSS LLC,* 1-13-cv-00930; *Safe Storage LLC v. Infortrend Corporation,* 1-13-cv-00929; *Safe Storage LLC v. Cisco Systems Inc.,* 1-13-cv-00926; *Safe Storage LLC v. Silicon Graphics Int'l Corp.,* 1-12-cv-01629; *Safe Storage LLC v. Dot Hill Systems Corp.,* 1-12-cv-01625; *Safe Storage LLC v. Hitachi Data Systems Corp. ,* 1-12-cv-01627; *Safe Storage LLC v. Dell Inc.,* 1-12-cv-01624; *Safe Storage LLC v. NetApp Inc.,* 1-12-cv-01628; *Safe Storage LLC v. Hewlett-Packard Company,* 1-12-cv-01626

Each of these proceedings is located in the District Court for Delaware.

C. Lead and Back-up Counsel and Service Information

Lead Counsel David L. McCombs HAYNES AND BOONE, LLP 2323 Victory Ave. Suite 700 Dallas, TX 75219	Phone: (214) 651-5533 Fax: (214) 200-0853 david.mccombs.ipr@haynesboone.com USPTO Reg. No. 32,271
Back-up Counsel Andrew S. Ehmke HAYNES AND BOONE, LLP 2323 Victory Ave. Suite 700 Dallas, TX 75219	Phone: (214) 651-5116 Fax: (214) 200-0853 Andy.ehmke.ipr@haynesboone.com USPTO Reg. No. 50,271
Back-up Counsel Thomas W. Kelton HAYNES AND BOONE, LLP 2323 Victory Ave. Suite 700 Dallas, TX 75219	Phone: (972) 739-6923 Fax: (214) 200-0853 Thomas.kelton.ipr@haynesboone.com USPTO Reg. No. 54,214
Back-up Counsel John Russell Emerson HAYNES AND BOONE, LLP 2323 Victory Ave. Suite 700 Dallas, TX 75219	Phone: (214) 651-5328 Fax: (214) 200-0884 Russ.emerson.ipr@haynesboone.com USPTO Ref. No. 44,098

II. Grounds for Standing

Petitioners certify that the '346 Patent is available for *inter partes* review and that Petitioners are not barred or estopped from requesting *inter partes* review challenging the patent claims on the grounds identified in the petition.

III. Relief Requested

Petitioners ask that the Board review the accompanying prior art and

analysis, institute a trial for *inter partes* review of claims 1-9 of the '346 Patent, and cancel those claims as unpatentable.

IV. The Reasons for the Requested Relief

The full statement of the reasons for the relief requested is as follows:

A. Summary of the '346 Patent

The '346 Patent relates to a system with "redundant interconnections between multiple hosts and a RAID [redundant array of independent discs]." The system includes two RAID controllers 460 and 461. Each RAID controller has two network interface controllers (NICs), so RAID controller 460 includes NICs 470 and 471, and RAID controller 461 includes NICs 480, 481. The system also has two hub/switch devices 440 and 441. Fig. 4 of the '346 Patent is illustrative:

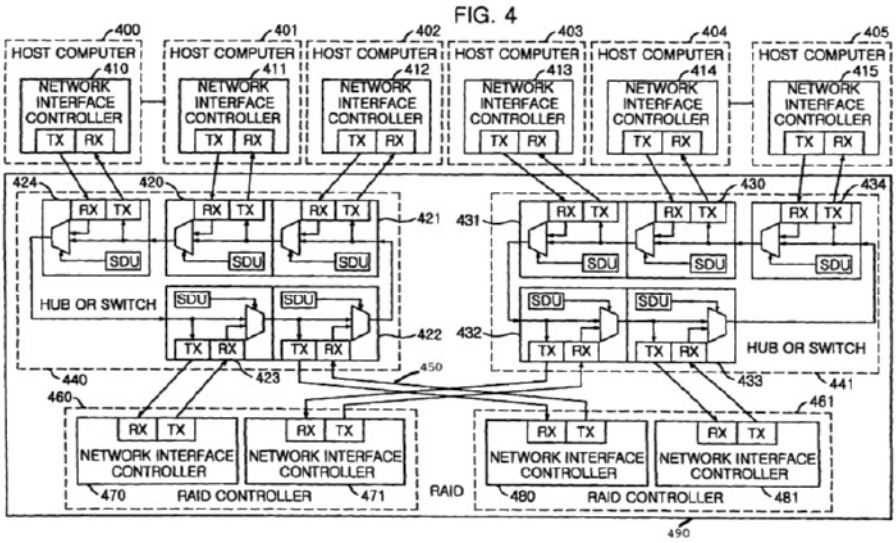

DHPN-1001, Fig. 4

Each RAID controller is connected to each hub/switch device by one of its NICs. RAID controller 460, on the left, connects to hub/switch 440 by NIC 470 and to hub/switch 441 by NIC 471. Similarly, RAID controller 461, on the right,

connects hub/switch 441 by NIC 481 and to hub/switch 440 by NIC 480.

This structure provides a "communication passage between two RAID controllers." DHPN-1001, 3:64-65. For instance, the RAID controller 460, on the left, can send information to the RAID controller 461, on the right, via NIC 471, switch/hub 441, and NIC 481 (at RAID controller 461). DHPN-1001, 3:66 – 4:12. In the same way, communication can also flow from NIC 481 to NIC 471. DHPN-1001, 4:6-9. Additionally, RAID controller 460, on the left, can send information to RAID controller 461, on the right, via NIC 470, hub/switch 440, and NIC 480. DHPN-1001, 4:9-12. In the same way, communication can also flow from NIC 480 to NIC 470. DHPN-1001, 4:2-6.

This system of redundant RAID controllers and NICs purports to provide a "fault tolerant function." DHPN-1001, 3:63-66. The specification describes that a RAID controller "having [an] error occurrence is removed from the network." Then a NIC from other RAID controller "takes over a function" of a NIC on the RAID controller with the error. DHPN-1001, 4:19-25.

B. The Prosecution History

The '346 Patent's prosecution included two rejections and two amendments. The application of the '346 Patent was filed with claims that eventually issued as claims 1-8 after amendment. DHPN-1002, pp.123-125. The Examiner initially rejected the claims over US 5,812,754 (hereinafter, "Lui"). DHPN-1002, pp.65-

69. In response, the Applicant amended claim 1 and added a claim that would eventually issue as claim 9. DHPN-1002, p.49. The amendment to claim 1 is not discussed further herein because it was substantially undone in the Applicant's next response. The Examiner issued a Final Office Action rejecting all claims over Lui. In response, the Applicant amended claims 1 and 9 and argued that Lui does not teach "two network interface controlling units included in each RAID controller." DHPN-1002, p.26. Applicant also argued that Lui does not teach "the first network controlling unit exchanges information with the fourth network controlling unit and the second network controlling unit exchanges information with the third network controlling unit." DHPN-1002, pp.26-27. The Examiner allowed the claims without providing reasons for allowance. DHPN-1002, pp.8-11.

V. Identification of Challenges and Claim Construction

A. Challenged Claims

Claims 1-9 of the '346 Patent are challenged in this petition.

B. Claim Construction

This petition analyzes the claims consistent with the broadest reasonable interpretation in light of the specification. *See* 37 C.F.R. § 42.100(b):

1. "*network controlling unit,*" "*network interface controlling unit*" and "*network interface controller*" (claims 1 and 9 and various dependent claims)

The terms "*network controlling unit,*" "*network interface controlling unit,*"

and "*network interface controller*" should be construed to be the same, namely, as any component allowing a device to communicate over a network (e.g., Fibre Channel, ATM, or other network). DHPN-1006 at pp.10-12.

2. "*the second network interface controlling unit and the fourth network controlling unit are used for executing a function of the first network interface controlling unit and the third network controlling unit when one of the first RAID controlling unit and the second RAID controlling unit is faulty*" (claim 4)

The above-recited limitation should be construed as "if either one of the first or second RAID controlling unit has an occurrence of an error, the apparatus uses a network controlling unit of the RAID controlling unit not having the error occurrence." DHPN-1006 at pp.12-13.

3. "*hub*" (claim 5)

The term "*hub*" should be construed as "hub or switch" as defined in the specification. DHPN-1006 at pp.16-17.

4. "*the rest of the connection ports being provided as a [hub equipment/network switch equipment/switch] connected with the numerous host computers*" (claims 5, 6, 7)

First, as a matter of grammar, the phrase "*connected with...*" modifies the "*hub equipment, network switch equipment,* or *switch*" and does not modify the "*connection ports*," because the phrase should modify the immediately preceding noun. This construction is consistent with Figs. 4, 5, and 6 showing a hub or switch connected with the host computers. DHPN-1006 at p.16. Second, the

phrase, "*the rest of…*" is not used in the specification to show a different meaning. DHPN-1006 at p.16. Third, the above-recited words do not logically exclude that the other ports, coupled to the *network controlling units*, are also provided as the *hub equipment*, *network switch equipment*, or *switch*.

5. "*coupled to*" (claims 3 and 5-7)

The term "*coupled to*" should be construed to be broader than the phrase "connected to." For example if entity A is *coupled to* entity B, then entity A is connected, directly or indirectly, in order to allow the transfer of signals between entities A and B. DHPN-1006 at pp.17-18.

6. *"[X] of the at least [Y] connection ports is [are] coupled to one of the first network interface controlling unit and the third network controlling unit*" (claims 5-7)

The term "[X] of the at least [Y] connection ports is [are] coupled to one of the first network interface controlling unit and the third network controlling unit" should be construed as "a hub (or switch) that has at least [Y] ports where at least [X] of the ports are connected directly or indirectly with the first network interface controlling unit or the third network controlling unit." DHPN-1006 at pp.13-16.

7. "*host computers*" (claims 1, 4-7 and 9)

The term "*host computers*" should be construed as "network connected computers." DHPN-1006 at pp.18-19.

8. "*RAID controlling unit*" and "*RAID controller*" (claims 1 and 9 and various dependent claims)

The terms "*RAID controlling unit*" and "*RAID controller*" should be construed as "a functional component including hardware that may be controlled by computer code, the functional component providing control to implement RAID storage in an array of storage drives." DHPN-1006 at pp.19-21.

9. "*RAID*" (claims 1 and 9)

The term "*RAID*" should be construed as "at least a redundant array of independent disks." DHPN-1006 at pp.21-23. A RAID may include RAID controllers. DHPN-1006 at p.21.

C. Statutory Grounds for Challenges

Challenge #1: Claims 1-3, 5, and 8 are anticipated under 35 U.S.C. § 102(b) by Peter Weygant, *Clusters for High Availability: A Primer of HP-UX Solutions*, 1996 ("Weygant") (DHPN-1003). Weygant published in 1996 and qualifies as prior art under 35 U.S.C. § 102(b).

Challenge #2: Claims 1-3 and 8 are obvious under 35 U.S.C. § 103(a) over Weygant in view of PCT/US99/01282, filed by Mylex Corporation ("Mylex") (DHPN-1007). Mylex published on July 29, 1999, and thus qualifies as prior art under 35 U.S.C. § 102(b).

Challenge #3: Claims 4 and 9 are obvious under 35 U.S.C. § 103(a) over Weygant and Mylex, further in view of *Managing MC/ServiceGuard*, Hewlett-Packard Company, ("ServiceGuard") (DHPN-1004). ServiceGuard published in 1998 and

is prior art under 35 U.S.C. § 102(b).

Challenge #4: Claims 5-7 are obvious under 35 U.S.C. § 103(a) over Weygant and Mylex, further in view of *American National Standard for Information Technology – Fibre Channel Arbitrated Loop (FC-AL-2)* ("ANSI") (DHPN-1008). ANSI was published June 28, 1999, and thus qualifies as prior art under 35 U.S.C. § 102(b).

Challenge #5: Claims 1-3 and 5-8 are anticipated under 35 U.S.C. § 102(b) by Hathorn et al., U.S. Pat. No. 5,574,950 ("the '950 Patent") (DHPN-1005). The '950 Patent issued in 1996 and qualifies as prior art under 35 U.S.C. § 102(b).

D. Identification of How the Construed Claims Are Unpatentable

1. Challenge #1: Claims 1-3, 5, and 8 are anticipated under 35 U.S.C. § 102(b) over Weygant

Claims 1-3, 5, and 8 of the '346 Patent are anticipated by Weygant under 35 U.S.C. § 102(b). Weygant discloses a RAID system having redundant RAID controllers and network interface controllers to provide fault tolerance functionality. For example, in Fig. 4.1 (below), Weygant shows two RAID controllers, each with two LAN cards. DHPN-1003 at p.131; DHPN-1006 at p.23.

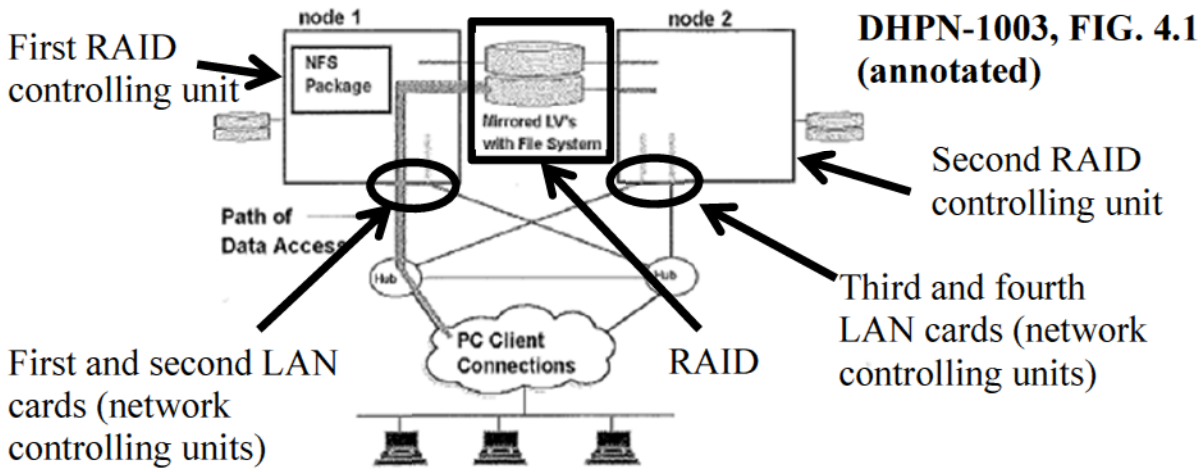

Weygant also discloses hubs that connect the RAID units with numerous computers. For example, in Fig 4.1, Weygant discloses two hubs that couple the PC clients to the RAID controlling units. DHPN-1006 at p.24; DHPN-1003, Fig. 4.1. Further, Fig. 4.1 of Weygant shows the same topology as in Fig. 4 of the '346 Patent. Specifically, at node 1, one LAN card is connected to the hub on the left, and the other LAN card is connected to the hub on the right. Also, in node 2, one LAN card is connected to the hub on the left, and the other LAN card is connected to the hub on the right. The topology of Fig. 4.1 provides for the same "communication passages" described in the '346 Patent's specification ('346 Patent, 3:62 – 4:12. DHPN-1006 at p.24. Weygant also discloses that a LAN card in one RAID controlling unit exchanges information (such as heartbeat messages) with a LAN card in a second RAID controlling unit. *See* DHPN-1003 at p.139-140.

Thus, Weygant discloses every element of claim 1, as well as dependent claims 2, 3, 5 and 8, as shown in detail below:

Claim 1

[1.0] ***An apparatus for a redundant interconnection between multiple hosts and a RAID, comprising:***

Weygant discloses this feature because it discloses a system having multiple "PC clients" connected redundantly to a set of mirrored disks in a RAID configuration. DHPN-1003, Fig 4.1; DHPN-1006 at pp.26-28. Nodes 1 and 2 share a set of mirrored disks arranged in a "RAID level 1" configuration using a "Logical Volume Manager and the separate MirrorDisk/UX subsystem." DHPN-1003 at p.67; DHPN-1006 at p.27. The PC clients have multiple connections to the RAID through redundant hubs. DHPN-1006 at p.26-27.

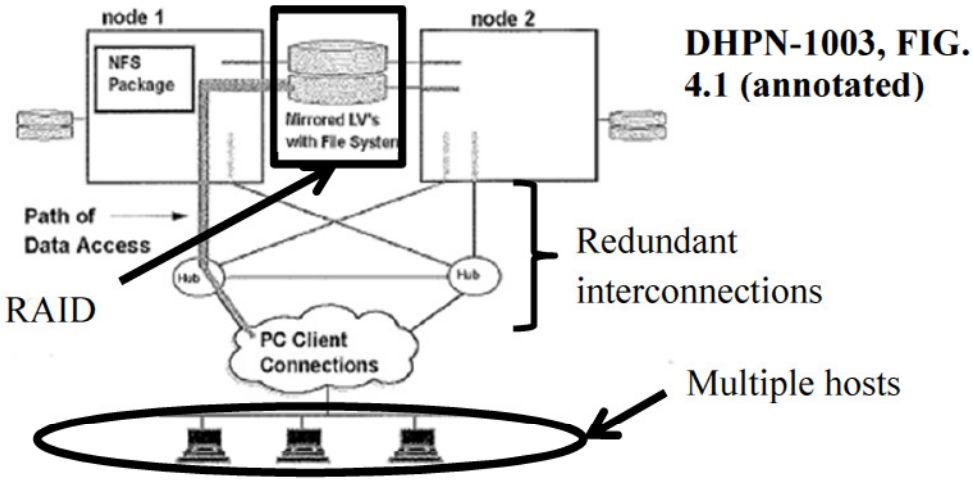

DHPN-1003, FIG. 4.1 (annotated)

Thus, the system that has redundant interconnections between multiple PC clients and a set of mirrored disks in a RAID configuration discloses "an apparatus for a redundant interconnection between multiple hosts and a RAID."

[1.1] ***a first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers***

Weygant discloses this limitation because it teaches a first node having

software that implements a "RAID level 1" configuration and a second node having software that implements a "RAID level 1" configuration. DHPN-1003 at p.67, 111, 172; DHPN-1006 at pp.28-30. Further, each node processes a requirement of numerous "PC clients" by providing a "facility for accessing file systems remotely" so that "writers and editors do not lose access" should a failure occur. DHPN-1003 at p.128; DHPN-1006 at pp.29-30. Thus, the two nodes that implement a RAID to provide a facility so that numerous clients can access remote file systems discloses "a first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers."

[1.2] ***the first RAID controlling unit including a first network controlling unit and a second network controlling unit***

Weygant discloses this limitation because it teaches that each node contains "two LAN cards," which transmit "file server requests from clients and also the cluster's own heartbeat messages" over a local area network. DHPN-1003 at p.131; DHPN-1006 at pp.30-31. Thus, the first node including two LAN cards that communicate over a network discloses "the first RAID controlling unit including a first network controlling unit and a second network controlling unit."

[1.3] ***and the second RAID controlling unit including a third network controlling unit and a fourth network controlling unit***

Weygant discloses this limitation because it teaches each node contains "two LAN cards," which transmit "file server requests from clients and also the cluster's

own heartbeat messages" over a local area network. DHPN-1003 at p.131; DHPN-1006 at pp.31-32. Thus, the second node including two LAN cards that communicate over a network discloses "the second RAID controlling unit including a third network controlling unit and a fourth network controlling unit."

[1.4] ***a plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers***

Weygant discloses this limitation because it teaches multiple hubs that provide communication paths between the nodes (RAID controlling units) and the PC clients (host computers). DHPN-1006 at pp.32-33. Thus, the multiple hubs that connect the first and second nodes (the RAID controlling units) to the PC clients (host computers) disclose "a plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers."

[1.5] ***wherein the first RAID controlling unit and the second RAID controlling unit directly exchange information with the numerous host computers through the plurality of connecting units***

Weygant discloses this limitation because it teaches that the first node and the second node communicate with the PC clients through the hubs to process file system requests for writing and editing. DHPN-1003 at p.128; DHPN-1006 at pp.33-35. Thus, the first node (a RAID controlling unit) and the second node (a RAID controlling unit) that exchange information with the PC clients through the hubs discloses "the first RAID controlling unit and the second RAID controlling

unit directly exchange information with the numerous host computers through the plurality of connecting units."

[1.6] ***and the first network controlling unit exchanges information with the fourth network controlling unit***

Weygant discloses this limitation because it teaches that the nodes (RAID controlling units) exchange "heartbeat messages" via their LAN cards (network controlling units). DHPN-1003 at p.66, 131; DHPN-1006 at pp.35-38. Weygant also teaches an active LAN card (network controlling unit) communicates with an active LAN card of another node (RAID controlling unit). While this concept is shown in Figs. 2.10 and 2.12, it also applies to Figs. 4.1-4.4. DHPN-1006 at p.37.

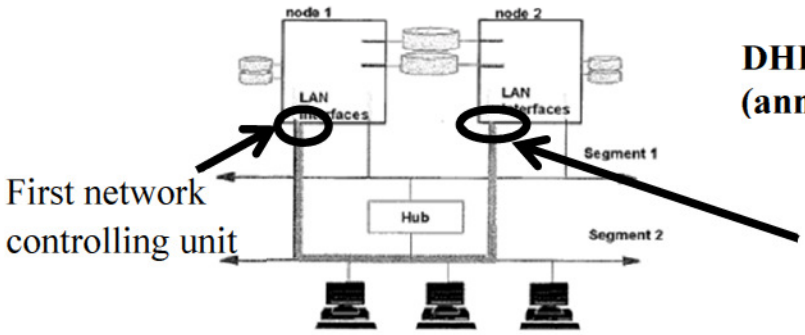

DHPN-1003, FIG. 2.10 (annotated)

Figure 2.10 *Ethernet LANs in a Grouped Subnet*

The "LAN interfaces" in the figures of Weygant include the LAN cards. DHPN-1006 at pp.25-26. Thus, the first LAN card that sends heartbeat messages to the fourth LAN card (and the fourth LAN card that sends heartbeat messages to the first LAN card) discloses "the first network controlling unit exchanges information with the fourth network controlling unit."

[1.7] ***and the second network controlling unit exchanges information with the***

third network controlling unit.

Weygant discloses this limitation because it teaches that the nodes (RAID controlling units) exchange "heartbeat messages" via their LAN cards (network controlling units). DHPN-1006 at pp.38-40; DHPN-1003 at p.131.

Also, Weygant teaches an active LAN card (network controlling unit) communicates with an active LAN card of another node (RAID controlling unit). This concept is shown in Figs. 2.10 and 2.12 and applies to the examples of Figs. 4.1-4.4. DHPN-1006 at p.39.

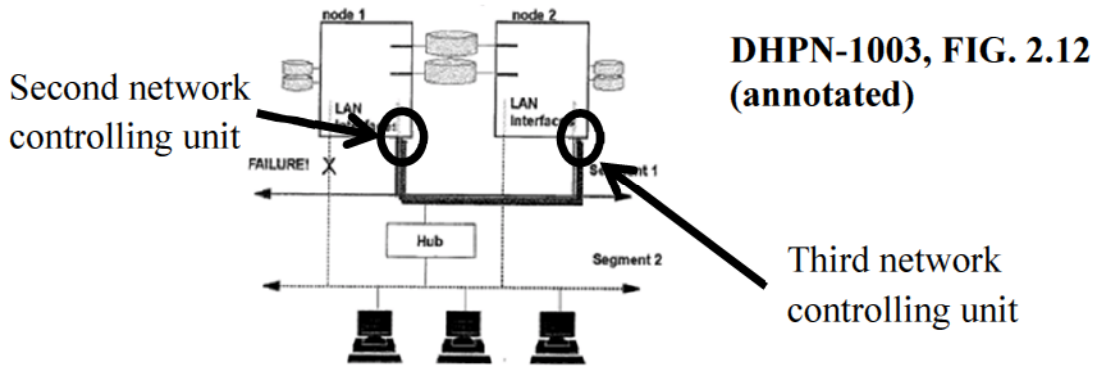

DHPN-1003, FIG. 2.12 (annotated)

***Figure 2.12** Grouped Net Following LAN Cable Failure*

The first, second, third, and fourth designations in the annotations above are exemplary, as either LAN card in node 1 can be considered first or second and any LAN card in node 2 can be considered third or fourth. DHPN-1006 at p.40.

Thus, the second LAN card that sends heartbeat messages to the third LAN card (and the third LAN card that sends heartbeat messages to the second LAN card) discloses "the second network controlling unit exchanges information with the third network controlling unit."

Claim 2

[2.1] ***The apparatus as recited in claim 1, wherein said respective RAID controlling units are connected to the plurality of individual connecting units.***

As discussed above, Weygant discloses the apparatus of claim 1. Weygant further discloses "said respective RAID controlling units are connected to the plurality of individual connecting units" because it teaches that node 1 and node 2 (the first and second RAID controlling units) are connected to two hubs. DHPN-1006 at pp.40-42. Thus, the first node and second node (the first and second RAID controlling units) that are connected to two hubs disclose "said respective RAID controlling units are connected to the plurality of individual connecting units."

Claim 3

[3.1] ***The apparatus as recited in claim 2, wherein the first network interface controlling unit is coupled to the connecting unit of one side and the second network interface controlling unit is coupled to the connecting unit of another side.***

As discussed above, Weygant discloses the apparatus of claim 2. Weygant further discloses the limitation of [3.1] because it teaches, in Fig. 4.1, that the LAN card on the left of node 1(first network controlling unit) is connected to the hub on the left side, and the LAN card on the right of node 1 (second network controlling unit) is connected to the hub on the right side. DHPN-1006 at pp.42-43.

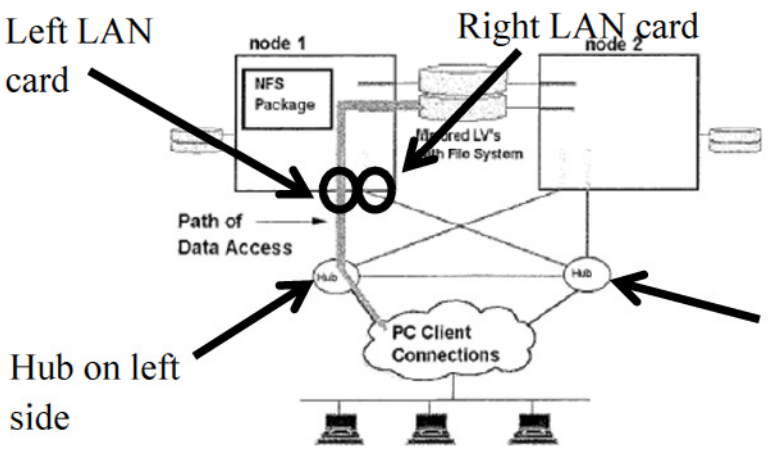

DHPN-1003, FIG. 4.1 (annotated)

Thus, the left LAN card connected to the hub on the left and the right LAN card connect to the hub on the right, discloses "the first network interface controlling unit is coupled to the connecting unit of one side and the second network interface controlling unit is coupled to the connecting unit of another side."

Claim 5

[5.1] ***The apparatus as recited in claim 1, wherein said plurality of connecting units have at least three connection ports,***

As discussed above, Weygant discloses the apparatus of claim 1. Weygant discloses this limitation because it teaches each hub (a connecting unit) having at least four ports—one port in communication with node 1, one port in communication with node 2, one port to the other hub, and one port to the PC client connections. DHPN-1006 at pp.44-45. Since each hub has at least four ports, the total number of ports in the plurality of connecting units is at least eight. DHPN-1006 at p.44. Thus, the ports of Weygant's hubs disclose "wherein said

plurality of connecting units have at least three connection ports," as recited in the claims.

[5.2] ***two of the at least three connection ports is coupled to one of the first network interface controlling unit and the third network controlling unit***

Weygant discloses this feature. Specifically, Weygant's Fig. 4.1 shows a connection port in the hub on the left coupled to a first LAN card at node 1 and a connection port at the hub on the right coupled to a third LAN card at node 2.

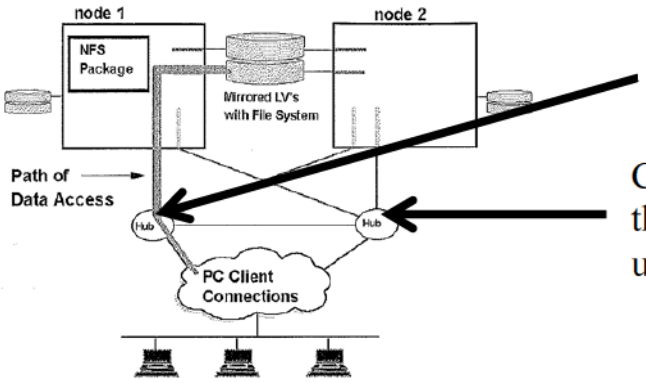

DHPN-1003, FIG. 4.1 (annotated)

Thus, the hub port coupled to the LAN card in node 1 and the hub port coupled to the LAN card in node 2 discloses "two of the at least three connection ports is coupled to one of the first network interface controlling unit and the third network controlling unit." DHPN-1006 at pp.44-47.

[5.3] ***and the rest of the connection ports being provided as a hub equipment connected with the numerous host computers.***

Weygant discloses this limitation. Weygant at Fig. 4.1 shows ports as hub equipment, where the hub equipment is connected with the PC clients (hosts). Weygant shows all ports in the hubs to be provided as hub equipment, therefore

"the rest of the connection ports" are provided as hub equipment, and the hub equipment is connected to the host computers. DHPN-1006 at p.47. Thus, the ports provided by the hubs in Weygant, where the hubs are connected with the PC clients, disclose "the rest of the connection ports being provided as a hub equipment connected with the numerous host computers."

Claim 8

[8.1] ***The apparatus as recited in claim 1, wherein the first network interface controlling unit of the first RAID controlling unit being connected to a first connecting unit, the second network interface controlling unit of said first RAID controlling unit being connected to a second connecting unit, the third network interface controlling unit of the second RAID controlling unit being connected to the second connecting unit, and the fourth network interface controlling unit of the second RAID controlling unit being connected to the first connecting unit.***

Weygant discloses this limitation because it teaches, in Fig. 4.1, that each of the two LAN cards (network interface controlling units) are connected to the hubs (connecting units). DHPN-1006 at pp.47-50. As shown below, (i) the first LAN card is connected to the first hub, (ii) the second LAN card is connected to the second hub, (iii) the third LAN card is connected to the second hub, and (iv) the fourth LAN card is connected to the first hub. DHPN-1006 at pp.47-50.

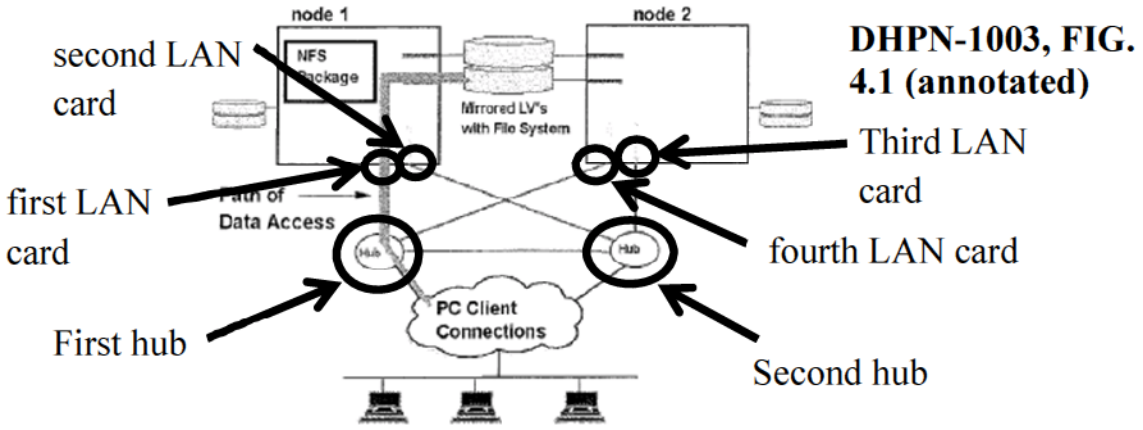

DHPN-1003, FIG. 4.1 (annotated)

The first, second, third, and fourth designations are exemplary, as either LAN card in node 1 can be considered a first or second network interface controlling unit, and either LAN card in node 2 can be considered a third or fourth network interface controlling unit. DHPN-1006 at pp.48-49. Therefore, Weygant's connections illustrated in Fig. 4.1 discloses this limitation.

2. Challenge #2: Claims 1-3 and 8 are obvious under 35 U.S.C. § 103(a) over Weygant in view of Mylex

Claims 1-3 and 8 are obvious under 35 U.S.C. § 103(a) over Weygant in view of Mylex. To the extent that a node having software that implements a RAID level 1 configuration on a set of mirrored disks, as disclosed in Weygant, does not disclose "the [first/second] RAID controlling unit," it would have been obvious to a person of ordinary skill in the art that RAID controlling functionality as recited in the claims could be implemented in a variety of software and hardware configurations. DHPN-1006 at pp.51-53. For example, Mylex discloses that RAID controllers could be implemented internal or external to a host system (see

Fig 1A, 1B). DHPN-1006 at p.51. Mylex also discloses exemplary physical components of a RAID controller. DHPN-1007, 6:24-7:5; DHPN-1006 at p.51.

The combination of Weygant with Mylex provides the details of a component to perform RAID controlling functionality, such as the nodes disclosed in Weygant. DHPN-1006 at pp.52-53. A person of ordinary skill in the art would implement such a combination in order to satisfy various design preferences and because employing any particular internal or external RAID controller configuration is merely a simple substitution of one known element for another to obtain predictable results. DHPN-1006 at p.53.

Thus, Weygant in view of Mylex discloses every element of claims 1-3 and 8. DHPN-1006 at pp.51-74.

Claim 1

[1.0] ***An apparatus for a redundant interconnection between multiple hosts and a RAID, comprising:***

As discussed in [1.0] of Challenge #1, Weygant discloses this limitation.

[1.1] ***a first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers***

Weygant in view of Mylex discloses this limitation. Weygant teaches a first node having software that implements a "RAID level 1" and a second node having software that implements a "RAID level 1". DHPN-1003 at p.67, 111, 172; DHPN-1006 at pp.53-55. Further, each node processes a requirement of numerous "PC clients" by providing a "facility for accessing file systems remotely" so that

"writers and editors do not lose access" should a failure occur. DHPN-1003 at p.128; DHPN-1006 at pp.55-57. Mylex discloses that a RAID controller could be implemented in a variety of hardware and software configurations. DHPN-1007 at Fig. 1A, 1B; DHPN-1006 at pp.56-57.

Thus, the two devices that implement a RAID to provide a facility so that numerous clients can access remote file systems, as disclosed in Weygant, combined with the RAID controller implementations of Mylex renders obvious "a first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers."

[1.2] ***the first RAID controlling unit including a first network controlling unit and a second network controlling unit***

As discussed in [1.2] of Challenge #1, Weygant discloses this limitation.

[1.3] ***and the second RAID controlling unit including a third network controlling unit and a fourth network controlling unit***

As discussed in [1.3] of Challenge #1, Weygant discloses this limitation.

[1.4] ***a plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers***

As discussed in [1.4] of Challenge #1, Weygant discloses this limitation.

[1.5] ***wherein the first RAID controlling unit and the second RAID controlling unit directly exchange information with the numerous host computers through the plurality of connecting units***

As discussed in [1.5] of Challenge #1, Weygant discloses this limitation.

[1.6] ***and the first network controlling unit exchanges information with the fourth network controlling unit***

As discussed in [1.6] of Challenge #1, Weygant discloses this limitation.

[1.7] ***and the second network controlling unit exchanges information with the third network controlling unit.***

As discussed in [1.7] of Challenge #1, Weygant discloses this limitation.

Claim 2

[2.1] ***The apparatus as recited in claim 1, wherein said respective RAID controlling units are connected to the plurality of individual connecting units.***

As discussed in [2.1] of Challenge #1, Weygant discloses this limitation.

Claim 3

[3.1] ***The apparatus as recited in claim 2, wherein the first network interface controlling unit is coupled to the connecting unit of one side and the second network interface controlling unit is coupled to the connecting unit of another side.***

As discussed in [3.1] of Challenge #1, Weygant discloses this limitation.

Claim 8

[8.1] ***The apparatus as recited in claim 1, wherein the first network interface controlling unit of the first RAID controlling unit being connected to a first connecting unit, the second network interface controlling unit of said first RAID controlling unit being connected to a second connecting unit, the third network interface controlling unit of the second RAID controlling unit being connected to the second connecting unit, and the fourth network interface controlling unit of the second RAID controlling unit being connected to the first connecting unit.***

As previously discussed in Challenge #1, Weygant discloses this limitation.

3. Challenge #3: Claims 4 and 9 are obvious over Weygant and Mylex, further in view of ServiceGuard under 35 U.S.C. § 103

Claims 4 and 9 are obvious over Weygant in view of Mylex, further in view

of ServiceGuard. The ServiceGuard reference provides further teaching and disclosures for implementing redundant LAN interfaces.

One skilled in the art would have multiple reasons to combine the teachings of Weygant and Mylex with ServiceGuard. For instance, ServiceGuard specifically states that its software package is specifically "designed to work in conjunction with other HP high availability products, such as MirrorDisk/UX," which is the software disclosed in Weygant. *See, e.g.,* ServiceGuard at 18; DHPN-1006 at p.76. Thus, ServiceGuard itself provides an explicit suggestion to combine. DHPN-1006 at p.76.

Further, it would have been obvious to combine Weygant and Mylex with ServiceGuard because it is a combination of known elements to achieve the predictable result of network controlling units exchanging information with each other over a pathway. *See* KSR International Co. v. Teleflex Inc. (KSR), 550 U.S. 82 USPQ2d 1385 (2007); MPEP §2141; *see also* DHPN-1006 at p.77. Weygant teaches the known element of using LAN interfaces (a type of network controlling unit) in high-availability computing nodes to communicate among the nodes. DHPN-1006 at p.77. Weygant also teaches that at each node, a LAN interface is active and another LAN interface is a standby interface. DHPN-1003 at p.131.

ServiceGuard teaches the known element that any given LAN interface in one node can exchange information (e.g., heartbeat signals) with an active LAN

interface in another node or communicate with a host computer; it just depends on which LAN interfaces are configured to be active. DHPN-1006 at p.76.

Thus, for a given subset of LAN interfaces (such as those in Weygant), ServiceGuard teaches different configurations for the LAN interface. There is no functional distinction between which given LAN interface is configured to provide a particular functionality. DHPN-1006 at p.77. Implementing this combination would have been a simple configuration decision within the abilities of a person of ordinary skill in the art in the 1999-2000 period. DHPN-1006 at p.77.

Therefore, because ServiceGuard specifically suggests a combination with the RAID software disclosed in Weygant, and that implementing such a combination would have been a simple configuration decision that would have enabled improved redundancy and fault tolerance, it is proper to combine ServiceGuard and Weygant. DHPN-1006 at pp.75-77.

Claim 4

[4.1] ***The apparatus as recited in claim 3, wherein the first network interface controlling unit and the third network interface controlling unit process the requirement of the numerous host computers;***

Weygant in view of ServiceGuard renders obvious this limitation. First, Weygant teaches that the LAN cards (network interface controlling units) process the requirements of the host computers by carrying file server requests. DHPN-1003 at p.131; DHPN-1006 at 78.

Second, ServiceGuard teaches that, whether the first or second LAN card

(network controlling unit) is active at node 1 (first RAID controlling unit), or whether the third or fourth LAN card (network controlling unit) is active at node 2 (second RAID controlling unit), is a matter of configuration. DHPN-1006 at pp.78-79. Therefore, in a scenario in which the first and third LAN cards (network controlling unit) are configured to be active, the first and third LAN cards process requirements of the clients (host computers). DHPN-1006 at p.78.

Thus, the LAN cards of Weygant that carry file server requests, in view of the LAN card configuration teachings of ServiceGuard render obvious "wherein the first network interface controlling unit and the third network interface controlling unit process the requirement of the numerous host computers."

[4.2] ***and the second network interface controlling unit and the fourth network controlling unit are used for communication between the first RAID controlling unit and the second RAID controlling unit when the first and second RAID controlling units are not faulty and***

Weygant in view of ServiceGuard renders obvious this limitation. First, Weygant teaches that the LAN cards (network interface controlling units) communicate between node 1 and node 2 (first and second RAID control units) at least by passing heartbeat signals. DHPN-1003 at p.131; DHPN-1006 at pp79-80.

Second, ServiceGuard teaches that whether the first or second LAN card (network controlling unit) is active at node 1 (first RAID controlling unit) is a matter of configuration; similarly, whether the third or fourth LAN card (network controlling unit) is active at node 2 (second RAID controlling unit) is a matter of

configuration. DHPN-1006 at p.80. Therefore, in a scenario in which the second and fourth LAN cards (network controlling units) are configured to be active, the second and fourth LAN cards would exchange the heartbeat signals. DHPN-1006 at p.80.

Thus, the LAN cards of Weygant that exchange heartbeat messages in view of the LAN card configuration teachings of ServiceGuard render obvious "the second network interface controlling unit and the fourth network controlling unit are used for communication between the first RAID controlling unit and the second RAID controlling unit when the first and second RAID controlling units are not faulty."

[4.3] ***the second network interface controlling unit and the fourth network controlling unit are used for executing a function of the first network interface controlling unit and the third network controlling unit when one of the first RAID controlling unit and the second RAID controlling unit is faulty.***

Weygant in view of ServiceGuard renders obvious this limitation. First, Weygant teaches that when a node (the first RAID controlling unit) fails, "the node will shut down, and MC/ ServiceGuard on the other node will start the package in its alternate location." DHPN-1003 at p.135. When that failure occurs, "the applications will continue running on the alternate node until the appropriate repair can be made on the failed node." DHPN-1003 at p.141; DHPN-1006 at 81. Fig. 3.1 shows the scenario in which the first network controlling unit and the fourth

network controlling unit are active. Fig. 3.2 then shows a failure of the first RAID controlling unit resulting in the second RAID controlling unit (and its active network controlling unit—the fourth) taking over. DHPN-1006 at pp.81-84.

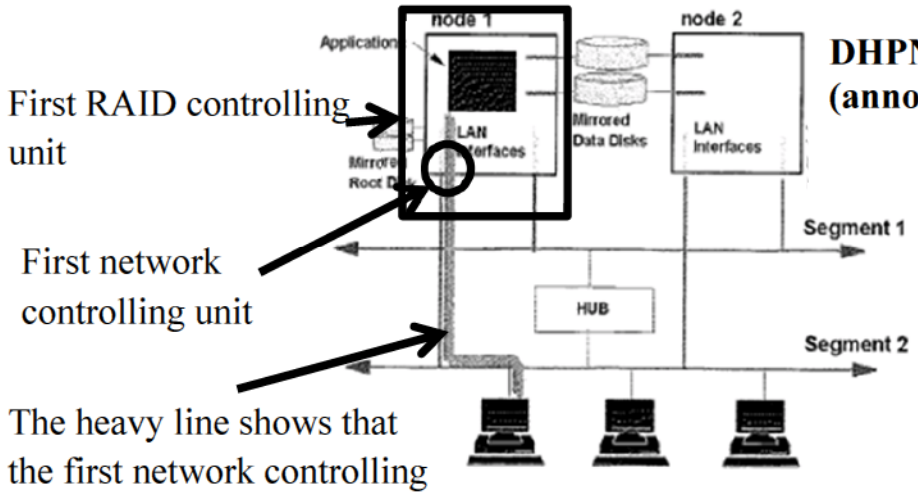

DHPN-1003, FIG. 3.1 (annotated)

Figure 3.1 *Active/Standby Cluster Before Failover*

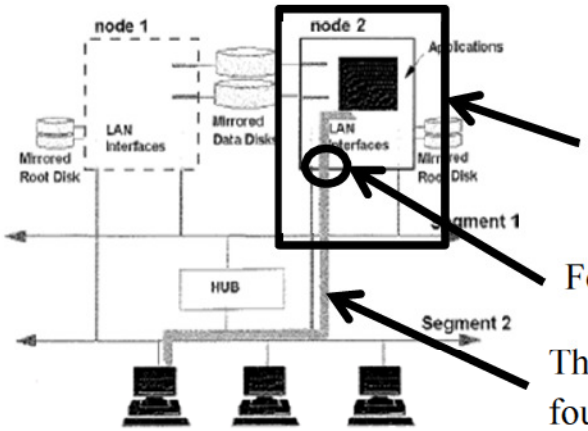

DHPN-1003, FIG. 3.2 (annotated)

Figure 3.2 *Active/Standby Cluster After Failover*

Thus, the operation of Weygant's LAN cards in the context of node failover renders obvious "the second network interface controlling unit and the fourth network controlling unit are used for executing a function of the first network interface controlling unit and the third network controlling unit when one of the

first RAID controlling unit and the second RAID controlling unit is faulty."

Claim 9

[9.0] ***An apparatus for a redundant interconnection between multiple host computers and a RAID, the apparatus comprising:***

Weygant discloses this limitation because it discloses a system having "PC clients" connected redundantly to a set of mirrored disks in a RAID configuration. DHPN-1003, Fig. 4.1; DHPN-1006 at pp.84-85. Nodes 1 and 2 share a set of mirrored disks arranged in a "RAID level 1" configuration using a "Logical Volume Manager and the separate MirrorDisk/UX subsystem." DHPN-1003 at p.67; DHPN-1006 at pp.85-86. The PC clients have multiple connections to the RAID through redundant hubs. DHPN-1006 at pp.85-86.

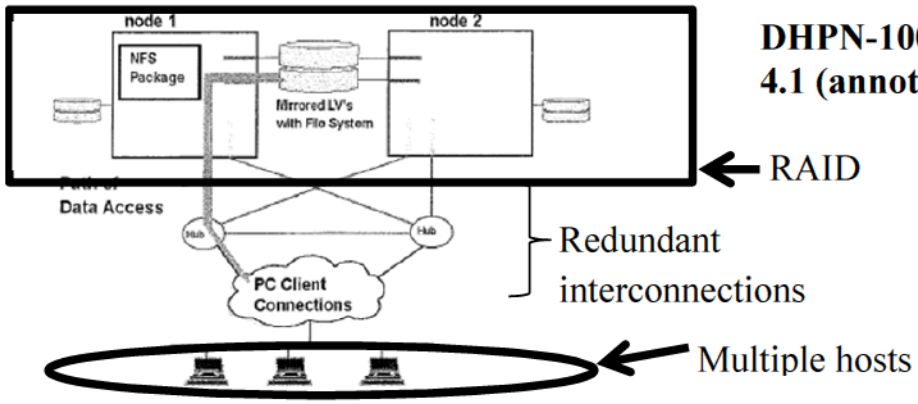

DHPN-1003, FIG. 4.1 (annotated)

Thus, the system that has redundant interconnections between multiple PC clients and a set of mirrored disks in a RAID configuration discloses "an apparatus for a redundant interconnection between multiple host computers and a RAID."

[9.1] ***a plurality of connection units for connecting the host computers and the RAID;***

Weygant discloses this limitation because it teaches multiple hubs that provide communication paths between the RAID (the array of disks and the two RAID controller nodes) and the PC clients (host computers). DHPN-1006 at pp.86-87. Thus, the multiple hubs that connect PC clients (host computers) to the RAID discloses "a plurality of connection units for connecting the host computers and the RAID."

[9.2] ***a first and a second RAID controllers, included in the RAID,***

Weygant and Mylex render obvious this limitation because Weygant teaches a RAID composed of a set of mirrored disks and two nodes that implement the RAID on the set of mirrored disks. DHPN-1003, p.67, 111 and 172; DHPN-1006 at pp.87-90.

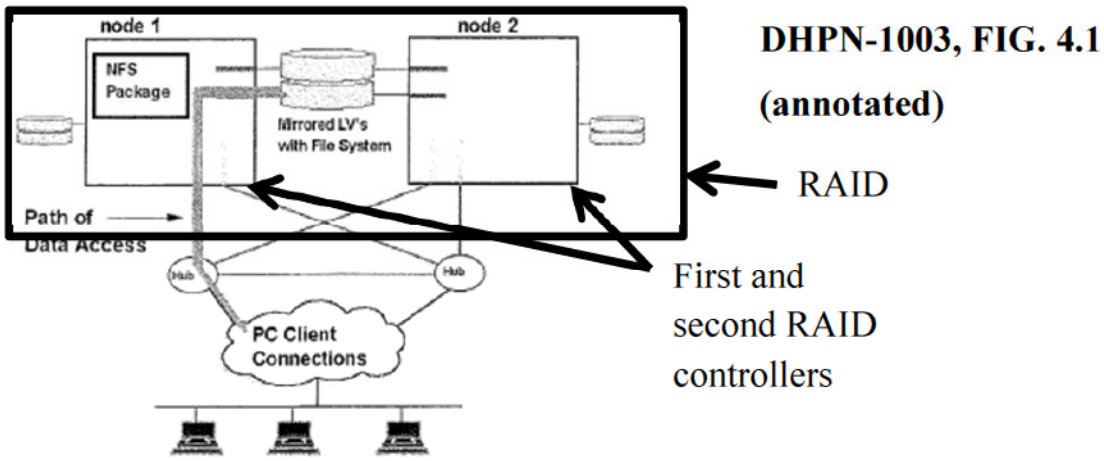

DHPN-1003, FIG. 4.1 (annotated)

Furthermore, Mylex teaches using a separate RAID controller with a node, and the concepts of Mylex can be applied to the nodes of Weygant. DHPN-1006 at pp.89-90. Thus, the first node and second node that implement a level 1 RAID

configuration with the set of disks, in view of Mylex, discloses "a first and a second RAID controllers, included in the RAID."

[9.3] ***each of which having a first network interface controller and a second network interface controller for processing requests from the plurality of the host computers connected through the plurality of the connection units,***

Weygant discloses this limitation because it teaches that each of the nodes contain "two LAN cards," which transmit "file server requests from clients and also the cluster's own heartbeat messages" over a local area network. DHPN-1003 at p.131; DHPN-1006 at pp.90-91.

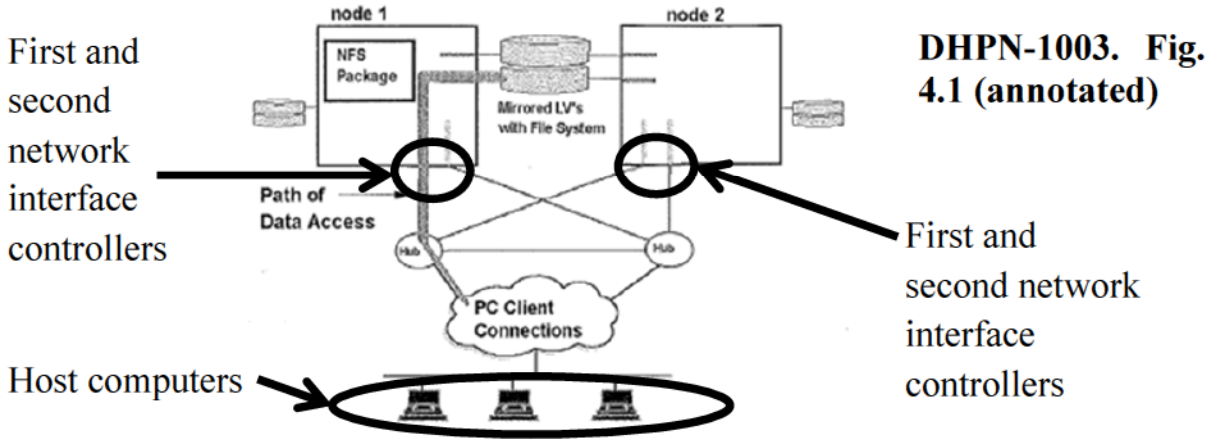

DHPN-1003. Fig. 4.1 (annotated)

Thus, the nodes that each include two LAN cards that communicate over a network disclose "each of which having a first network interface controller and a second network interface controller for processing requests from the plurality of the host computers connected through the plurality of the connection units."

[9.4] ***wherein the first network interface controller in the first RAID controller supplies data to the host computers connected through the plurality of connection units and***

Weygant discloses this limitation because it teaches that the first network

interface controller in the first RAID controller supplies data to the clients (host computers) by carrying "file server requests from clients" and provide access to a network file system. DHPN-1003 at p.131, 128; DHPN-1006 at p.91-93. Thus, the first LAN card in the first RAID controller that carries file server requests from PC clients and provides access to a network file system discloses "wherein the first network interface controller in the first RAID controller supplies data to the host computers connected through the plurality of connection units," as recited in the claims. DHPN-1006 at pp.91-93.

[9.5] ***processes information transmitted from the second network interface controller in the second RAID controller,***

Weygant in view of ServiceGuard renders obvious that the first network interface controller "processes information transmitted from the second network interface controller in the second RAID controller." Weygant teaches that the LAN cards (network interface controllers) send "heartbeat messages" among the nodes (RAID controllers). DHPN-1003 at p.76, 131, and 139-140; DHPN-1006 at pp.94-95. In the example of Fig. 4.1, when the first network interface controller in the first RAID controller is active, it processes information (heartbeat signals) from the active network interface controller in the second RAID controller. DHPN-1006 at pp.95-96. ServiceGuard discloses that, whether the active network interface controller is the first or the second network interface controller in the second RAID controller is simply a matter of configuration. DHPN-1006 at pp.96-

97.

Thus, the LAN cards of Weygant that exchange heartbeat signals, in view of the LAN card configuration teachings of ServiceGuard renders obvious that the first network interface controller "processes information transmitted from the second network interface controller in the second RAID controller."

[9.6] ***wherein the first network interface controller in the second RAID controller supplies data to the host computers connected through the plurality of connection units and***

Weygant discloses this limitation because it teaches that the first network interface controller in the second RAID controller supplies data to the PC clients (host computers) by carrying "file server requests from clients" and provides access to a network file system. DHPN-1003 at p.131, 128; DHPN-1006 at p.98.

Whichever LAN card at node 2 is the "first" network interface controller in the system of Weygant is exemplary, as active and standby status of a LAN card is a matter of configuration (as evidenced by ServiceGuard), and a person of ordinary skill in the art would have recognized that either LAN card in node 2 could have been designated as active and used to supply data to the PC clients (host computers). DHPN-1006 at 99-100.

Thus, the first LAN card in the second RAID controller that carries file server requests from clients provides access to an NFS document system discloses "wherein the first network interface controller in the second RAID controller

supplies data to the host computers connected through the plurality of connection units," as recited in the claims. DHPN-1006 at pp.98-100.

[9.7] ***processes information transmitted from the second network interface controller in the first RAID controller,***

Weygant in view of ServiceGuard renders obvious that the first network interface controller in the second RAID controller "processes information transmitted from the second network interface controller in the first RAID controller." First, Weygant teaches that the LAN cards (network interface controllers) send "heartbeat messages" among the nodes (RAID controllers): *See* DHPN-1003 at p.76, 131, 139-140. DHPN-1006 at p.101. In the example of Fig. 4.1, when the first network interface controller in the second RAID controller is active, it processes information (heartbeat signals) from the active network interface controller in the first RAID controller. DHPN-1006 at 101. Second, ServiceGuard discloses that, whether the active network interface controller is the first or the second network interface controller in the second RAID controller is simply a matter of configuration. DHPN-1006 at p.101.

Thus, the LAN cards of Weygant that exchange heartbeat signals, in view of the LAN card configuration teachings of ServiceGuard renders obvious that the first network interface controller in the second RAID controller "processes information transmitted from the second network interface controller in the first RAID controller."

[9.8] ***wherein the second network interface controller in the first RAID controller is used for fault tolerance by performing functions of the first network interface controller in the second RAID controller when the second RAID controller is faulty, and***

Weygant in view of ServiceGuard renders obvious this limitation. First, Weygant teaches that when a node (*e.g.*, the second RAID controller) fails, "the node will shut down, and MC/ServiceGuard on the other node will start the package in its alternate location." DHPN-1003 at p.135; DHPN-1006 at p.101. Further, Weygant teaches that, in the event of such a failure, "the applications will continue running on the alternate node until the appropriate repair can be made on the failed node." DHPN-1003 at p.141; DHPN-1006 at p.102. The following examples using Figs. 3.1 and 3.2 of Weygant to show a failover operation.

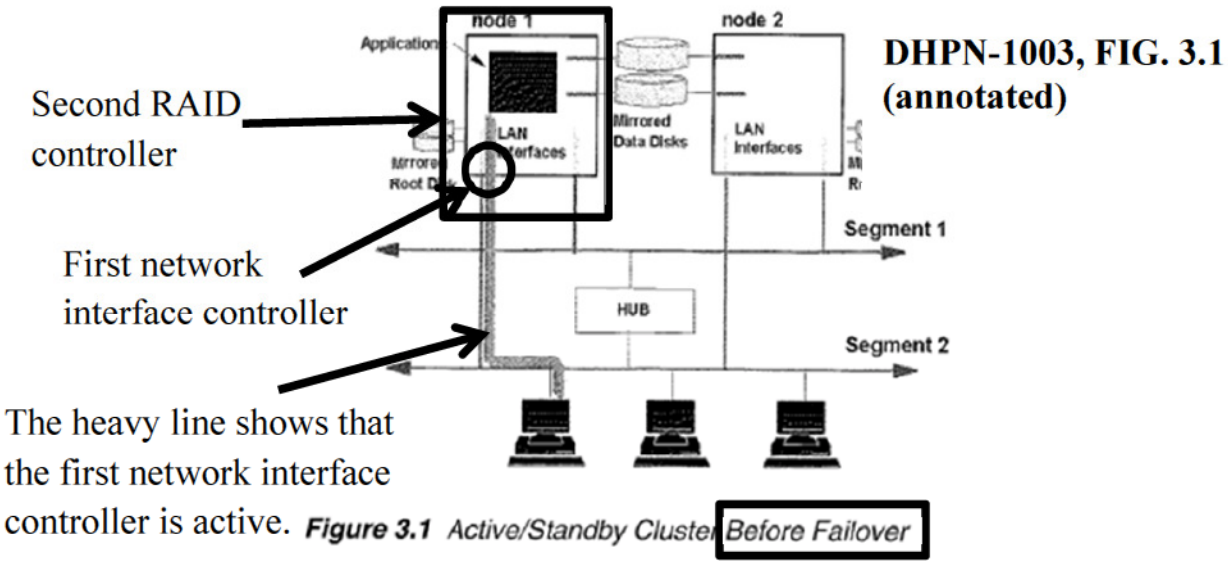

DHPN-1003, FIG. 3.1 (annotated)

Fig. 3.2 shows a failure of the second RAID controller resulting in the first RAID controller (and its active network interface controller) taking over. DHPN-1006 at p.104-105.

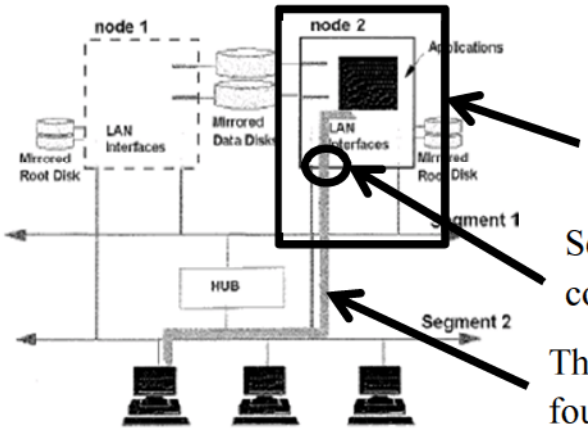

DHPN-1003, FIG. 3.2 (annotated)

First RAID controller

Second network interface controller

The heavy line shows that the fourth network interface controller is active.

Figure 3.2 *Active/Standby Cluster After Failover*

As shown above at Figs. 3.1 and 3.2, in a scenario in which the second network interface controller in the first RAID controller is active and in which the first network interface controller in the second RAID controller is active, a failure of the second RAID controller will result in the first RAID controller (and its active network interface controller) taking over. DHPN-1006 at p.104-105. Further, which network interface controllers are active is a matter of system configuration, as evidenced by ServiceGuard. DHPN-1006 at p.105. Thus, the operation of Weygant's LAN cards in the context of node failover renders obvious "the second network interface controller in the first RAID controller is used for fault tolerance by performing functions of the first network interface controller in the second RAID controller when the second RAID controller is faulty."

[9.9] ***wherein the second network interface controller in the second RAID controller is used for fault tolerance by performing functions of the first network interface controller in the first RAID controller when the first RAID controller is faulty, and***

Weygant in view of ServiceGuard renders obvious this feature. As

previously discussed with reference to Figs. 3.1 and 3.2 of Weygant, a scenario in which the second network interface controller in the second RAID controller is active and in which the first network interface controller in the first RAID controller is active, a failure of the first RAID controller will result in the second RAID controller (and its active network interface controller) taking over. DHPN-1006 at p.108. Also, as explained above, which network interface controller is active is a matter of system configuration, as evidenced by ServiceGuard. DHPN-1006 at p.108.

Thus, the operation of Weygant's LAN cards in the context of node failover renders obvious "wherein the second network interface controller in the second RAID controller is used for fault tolerance by performing functions of the first network interface controller in the first RAID controller when the first RAID controller is faulty."

[9.10] ***wherein the first network controlling unit in the first RAID controlling unit exchanges information with the second network controlling unit in the second RAID controlling unit, and***

Weygant discloses this feature because it teaches that the nodes (RAID controlling units) exchange "heartbeat messages" via their LAN cards (network controlling units). DHPN-1003 at p.76, 131; DHPN-1006 at pp.109-110. Weygant also teaches an active LAN card (network controlling unit) communicates with a LAN card of another node (RAID controlling unit). DHPN-

1006 at p.111. While this concept is shown in Figs. 2.10 and 2.12, it also applies to Figs. 4.1-4.4. DHPN-1006 at pp.111-113. In this example using Fig. 2.10, node 1 discloses the first RAID controlling unit, and node 2 discloses a second RAID controlling unit.

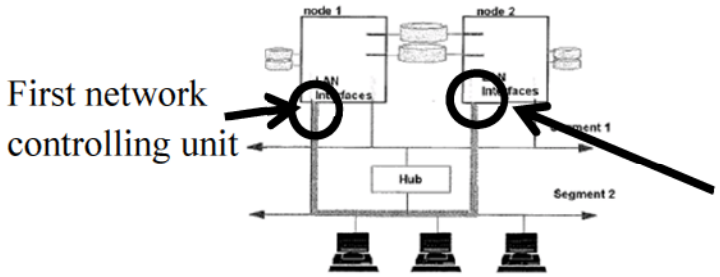

DHPN-1003, FIG. 2.10 (annotated)

Figure 2.10 *Ethernet LANs in a Grouped Subnet*

Thus, the first LAN card at node 1 that sends heartbeat messages to, and receives heartbeat messages from, the second LAN card at node 2 discloses "wherein the first network controlling unit in the first RAID controlling unit exchanges information with the second network controlling unit in the second RAID controlling unit."

[9.11] ***the second network controlling unit in the first RAID controlling unit exchanges information with the first network controlling unit in the second RAID controlling unit.***

Weygant discloses this feature because it teaches that the nodes (RAID controlling units) exchange "heartbeat messages" via their LAN cards (network controlling units). DHPN-1003 at p.76, 131; DHPN-1006 at p.113. Also, Weygant teaches an active LAN card (network controlling unit) communicates with an active LAN card of another node (RAID controlling unit). DHPN-1006 at

pp.113-115. While this concept is shown in Figs. 2.10 and 2.12, it also applies to Figs. 4.1-4.4. DHPN-1006 at pp.113-115. In this example using Fig. 2.10, node 1 discloses the first RAID controlling unit, and node 2 discloses a second RAID controlling unit.

Second network controlling unit

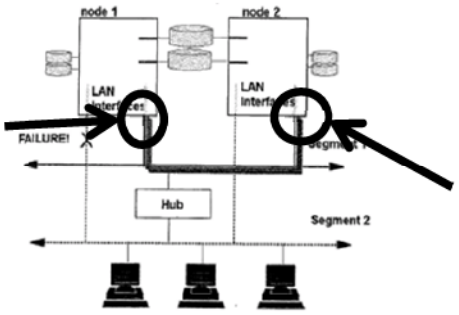

First network controlling unit

DHPN-1003, FIG. 2.12 (annotated)

Figure 2.12 Grouped Net Following LAN Cable Failure

The "first" and "second" designations in the annotations are exemplary, as either LAN card in node 1 can be considered first or second and any LAN card in node 2 can be considered first or second. DHPN-1006 at, e.g., pp.73-74.

Thus, the second LAN card that sends heartbeat messages to the third LAN card (and the third LAN card that sends heartbeat messages to the second LAN card) discloses "the second network controlling unit in the first RAID controlling unit exchanges information with the first network controlling unit in the second RAID controlling unit."

4. Challenge #4: Claims 5-7 are obvious under 35 U.S.C. § 103(a) over Weygant and Mylex further in view of ANSI

The term, "[X] of the at least [Y] connection ports is [are] coupled to one of the first network interface controlling unit and the third network controlling unit"

does not require that all [X] connection ports must be coupled to the first network controlling unit OR all [X] connection ports must be coupled to the third network controlling unit. However, to the extent that such a construction might be argued, the workings of hubs and switches would have provided direct and indirect connections among all of the ports of a given hub or a switch so that a connection to a given hub or switch port serves as a connection to all of the devices coupled to the other ports on such hub or switch (thereby satisfying even such a construction). Challenge #4 is presented to supplement the teachings of Weygant and Mylex with ANSI and to show that, in a given hub or switch, a connection to a given port serves as a connection to all of the devices coupled to the other ports on such hub or switch. A person of ordinary skill in the art would implement such a combination in order to ensure proper interconnection functionality, among other reasons. DHPN-1006 at pp.117-118.

Additionally, with respect to claims 6 and 7, it would have been obvious to a person of ordinary skill in the art to have used network switch equipment or a switch to provide network interconnections in the system of Weygant as an added function to "take advantage of these natural pauses in communication" and because a switch fabric is "robust." DHPN-1006 at pp.117-118. Also, using a network switch equipment or switch in place of Weygant's hubs is merely a simple substitution of one known element for another to obtain predictable results (the

result of providing network interconnections). DHPN-1006 at pp.117-118.

Claim 5

[5.1] ***The apparatus as recited in claim 1, wherein said plurality of connecting units have at least three connection ports,***

As discussed above, Weygant discloses the apparatus of claim 1. Also, Weygant discloses this limitation for the reasons given at [5.1] at Challenge #1.

[5.2] ***two of the at least three connection ports is coupled to one of the first network interface controlling unit and the third network controlling unit***

Weygant in view of ANSI renders obvious this limitation. First, Weygant teaches that its two hubs are coupled to each other, and that there are a total of eight connection ports. DHPN-1003, Fig. 4.1; DHPN-1006 at p.120. Second, ANSI discloses that a hub connects each of its ports to all of the other ports on the hub. DHPN-1008 at p.20. Notably, ANSI discloses the same single loop connecting each port as depicted in the Fig. 4 of the '346 Patent. DHPN-1006 at pp.120-121.

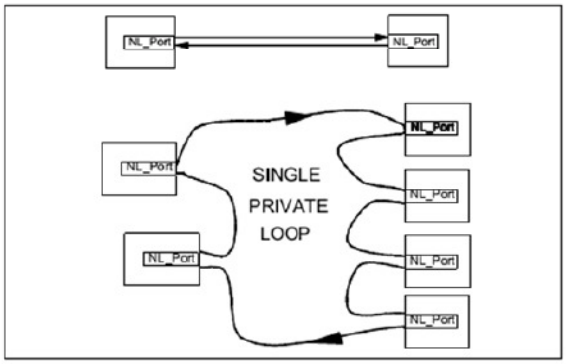

DHPN-1008, FIG. 1(a)

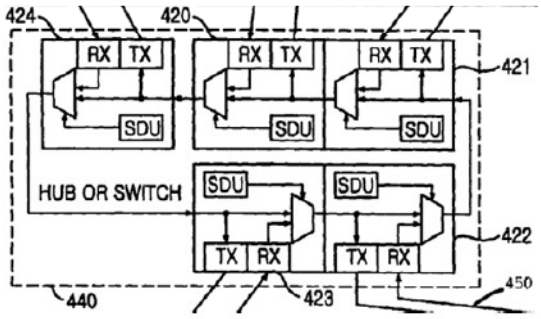

'346 PATENT, FIG. 4

Therefore, in a given hub, a connection to a given port serves as a

connection to all of the devices coupled to the other ports on such hub. DHPN-1006 at p.121. Thus, Weygant in view of ANSI discloses eight connection ports, in which each port is coupled, via a loop in the hub, to all of the devices on the network (including the respective LAN cards on the nodes), thereby rendering obvious "two of the at least three connection ports is coupled to one of the first network interface controlling unit and the third network controlling unit."

[5.3] ***and the rest of the connection ports being provided as a hub equipment connected with the numerous host computers.***

Weygant in view of ANSI renders obvious this limitation. Weygant at Fig. 4.1 shows ports as hub equipment, where the hub equipment is connected with the PC clients (hosts). DHPN-1006 at p.122. ANSI discloses that a hub provides "a single private loop," so the rest of the ports are connected to the host computers, directly or indirectly. DHPN-1006 at pp.122-123. Thus, the ports provided by the hubs in Weygant in view of ANSI, where the hubs are connected with the PC clients, render obvious "the rest of the connection ports being provided as a hub equipment connected with the numerous host computers."

Claim 6

[6.1] ***The apparatus as recited in claim 1 wherein said plurality of connecting units have at least three connection ports,***

This limitation is identical to the limitation of element [5.1] discussed above, and for those same reasons is disclosed by Weygant.

[6.2] ***two of the at least three connection port are coupled to one of the first***

network interface controlling unit and the third network controlling unit

Weygant in view of ANSI renders obvious this limitation. Specifically, ANSI discloses that a switch is an alternative to a hub and that a switch connects each of its ports to all of the other ports on the switch. DHPN-1008 at pp.18-19; DHPN-1006 at p.125. Thus, in a switch, a connection to a given port serves as a connection to all of the devices coupled to the other ports on such switch. DHPN-1006 at p.125. Therefore, Weygant in view of ANSI discloses eight connection ports, in which each of the eight connection ports is coupled, via the functionality of the switch, to all of the devices on the network (including the respective LAN cards on the nodes), thereby rendering obvious "two of the at least three connection port are coupled to one of the first network interface controlling unit and the third network controlling unit."

[6.3] ***and the rest of the connection ports being provided as a network switch equipment connected with the numerous host computers.***

Weygant in view of ANSI renders obvious this limitation. Weygant at Fig. 4.1 shows ports as hub equipment, where the hub equipment is connected with the PC clients (hosts). DHPN-1006 at p.126. Also, ANSI discloses that network switch equipment could be used instead of a hub and that the switch connects each of its ports to all of the other ports on the switch, so the rest of the ports are connected to the host computers, directly or indirectly. DHPN-1006 at p.126.

Thus, the connections of Weygant in view of the switches of ANSI render

obvious "the rest of the connection ports being provided as a network switch equipment connected with the numerous host computers."

Claim 7

[7.1] ***The apparatus as recited in claim 1, wherein said plurality of connecting units have at least five connection ports,***

As discussed above, Weygant discloses the apparatus of claim 1. Weygant further discloses this limitation because Weygant teaches that the plurality of hubs has eight ports (i.e., "at least five connection ports"). DHPN-1006 at pp.127-128.

Thus, the eight ports of the plurality of hubs in Weygant discloses "wherein said plurality of connecting units have at least five connection ports."

[7.2] ***four of the at least five connection ports is coupled to one of the first network interface controlling unit and the third network controlling unit***

Weygant in view of ANSI renders obvious this limitation. Specifically, ANSI discloses that a switch is an alternative to a hub and that a switch connects each of its ports to all of the other ports on the switch. DHPN-1008 at pp.18-19; DHPN-1006 at pp.128-129. Thus, in a switch, a connection to a given port serves as a connection to all of the devices coupled to the other ports on such hub. DHPN-1006 at p.129.

Thus, Weygant in view of ANSI discloses eight connection ports, in which each of the ports is coupled, via the switch, to all devices on the network (including the respective LAN cards on the nodes), thereby rendering obvious this limitation.

[7.3] ***and the rest of the connection ports being provided as a switch connected***

with the numerous host computers.

Weygant in view of ANSI renders obvious this limitation. Weygant at Fig. 4.1 shows ports as hub equipment, where the hub equipment is connected with the PC clients (hosts). DHPN-1006 at p.130. Also, ANSI discloses that a switch could be used instead of a hub and that the switch connects each of its ports to all of the other ports on the switch, so the rest of the ports are connected to the host computers, directly or indirectly. DHPN-1006 at p.130.

Thus, the connections of Weygant in view of the switches of ANSI render obvious "the rest of the connection ports being provided as a switch connected with the numerous host computers," as recited in the claim.

5. Challenge #5: Claims 1-3 and 5-8 are anticipated by the '950 Patent under 35 U.S.C. § 102(b)

Claims 1-3 and 5-8 of the '346 Patent are anticipated by the '950 Patent under 35 U.S.C. § 102(b). The '950 Patent describes storage controllers (RAID controlling units) that communicate with each other via network controlling units. *See* DHPN-1004, '950 Patent, 8:11-15. Fig. 3 of the '950 Patent is illustrative.

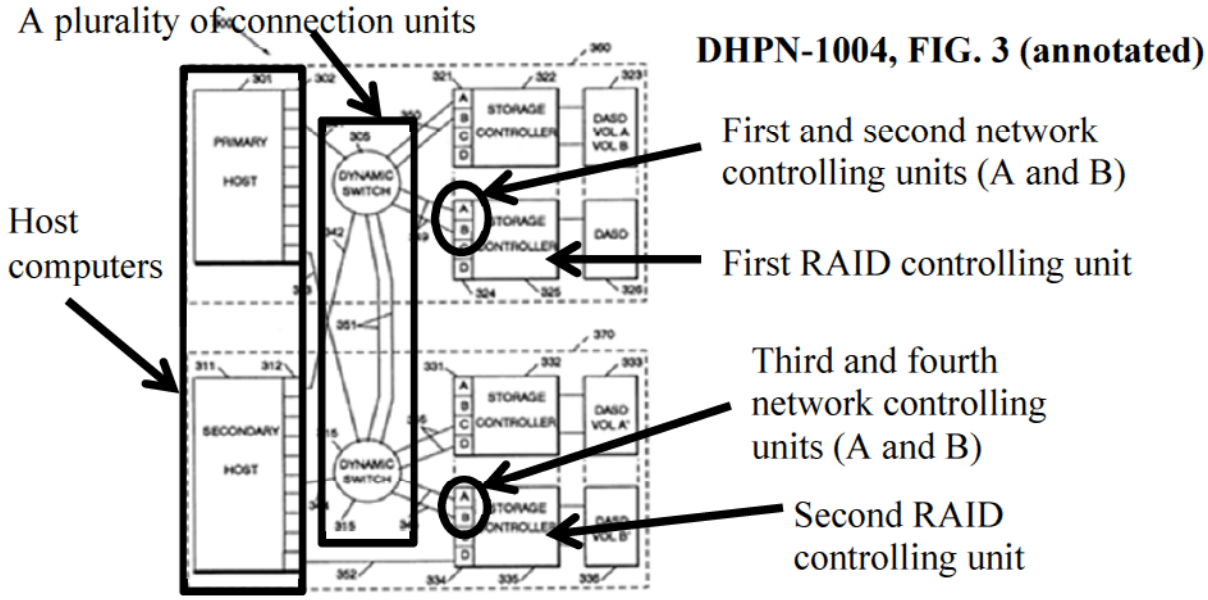

In the '950 Patent, "primary storage controller 322, via the same port A 321, can communicate with secondary storage controller 332 by communication links 350, dynamic switch 305, communication links 351, dynamic switch 315, and communication links 346, wherein port A 321 acts as a channel link-level facility." DHPN-1005, 8:11-15. Therefore, as shown element-by-element below, the '950 Patent anticipates claims 1-3 and 5-8.

Claim 1

[1.0] ***An apparatus for a redundant interconnection between multiple hosts and a RAID, comprising:***

The '950 Patent discloses this limitation because it discloses a system having two dynamic switches that provide redundant interconnections between multiple hosts and RAIDs. DHPN-1006 at pp.132-134. Fig. 3 of the '950 Patent illustrates hosts and RAIDs:

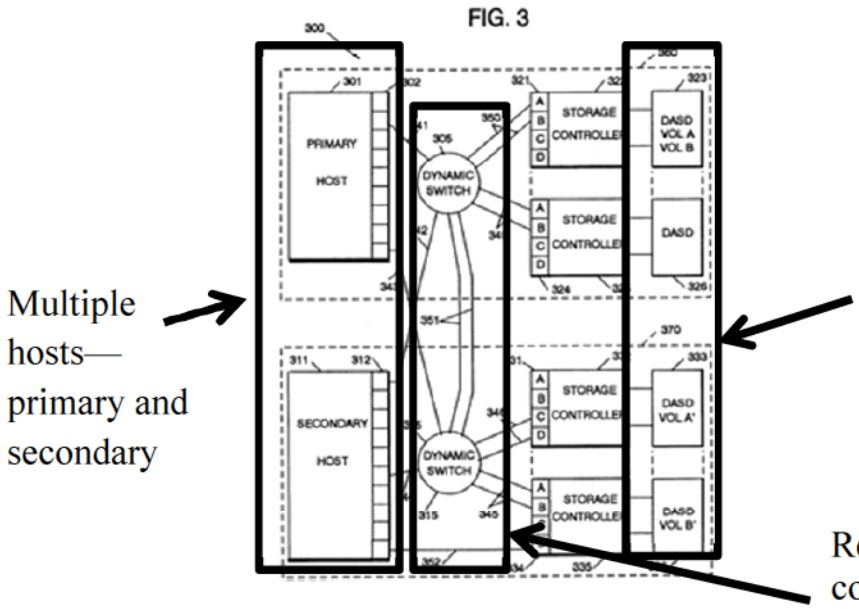

DHPN-1005, FIG. 3 (annotated)

Thus, the system with two dynamic switches that provide redundant interconnections between primary and secondary hosts with a RAID discloses "an apparatus for a redundant interconnection between multiple hosts and a RAID."

[1.1] ***a first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers***

The '950 Patent discloses this limitation because it teaches two storage controllers that control direct access storage devices (DASDs). DHPN-1006 at pp.134-136. In particular, the '950 Patent teaches that the DASDs could be implemented in a RAID configuration. DHPN-1005, 2:5-11; DHPN-1006 at p.134.

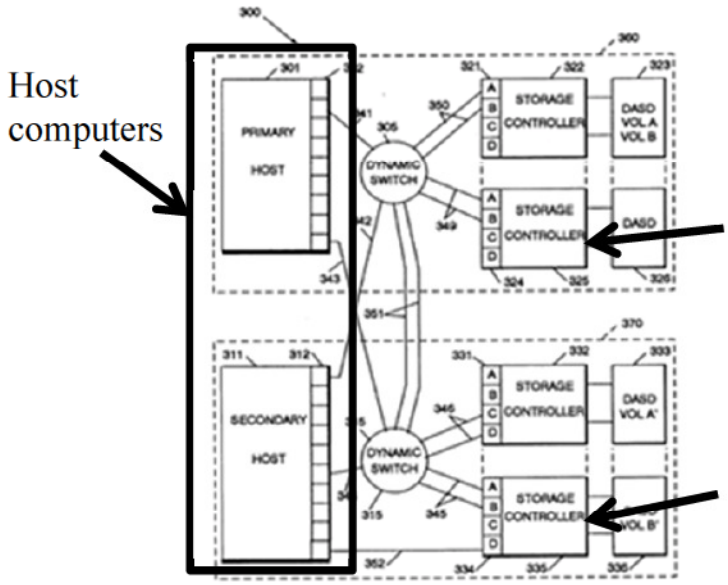

DHPN-1005, '950 Patent, FIG. 3 (annotated)

Further, the '950 Patent teaches that the storage controllers process requests from the primary host and secondary host for transferring data or records from the DASDs. See, e.g., DHPN-1005, 7:28; see DHPN-1006 at pp.135-136 (the functionality of the hosts communicating with the storage controllers in Fig. 2 applies as well to Fig. 3). Thus, the two storage controllers that control DASDs in a RAID configuration and that process requests from the multiple hosts disclose "a first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers."

[1.2] ***the first RAID controlling unit including a first network controlling unit and a second network controlling unit***

The '950 Patent discloses this limitation because it teaches that each storage controller has four ports that can be "dynamically set … either as a channel or control unit link-level facility" for communicating on the network. DHPN-1005, 8:3-15; DHPN-1006 at pp.137-139.

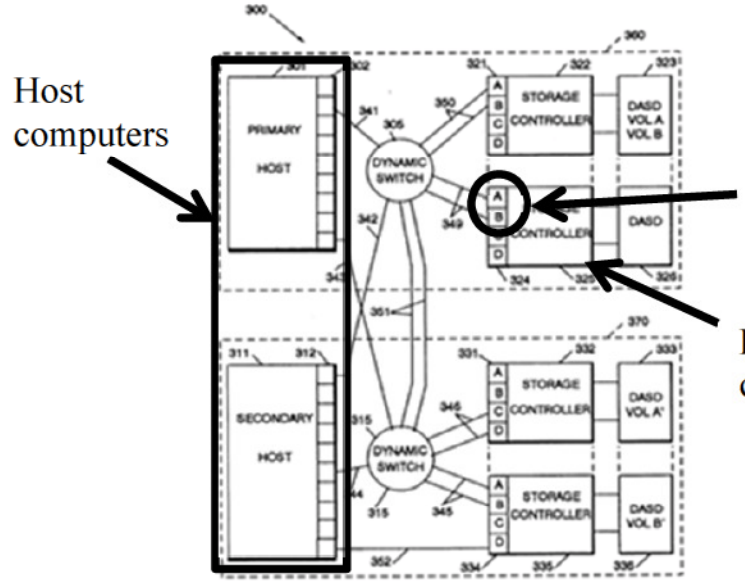

DHPN-1005, FIG. 3 (annotated)

First and second network controlling units (A and B)

First RAID controlling unit

Thus, the first storage controller having at least two channel or control unit link-level facilities discloses "the first RAID controlling unit including a first network controlling unit and a second network controlling unit."

[1.3] ***and the second RAID controlling unit including a third network controlling unit and a fourth network controlling unit***

The '950 Patent discloses this limitation because it teaches that each storage controller has four ports that can be "dynamically set … either as a channel or control unit link-level facility" for communicating on the network. DHPN-1005, 8:3-15; DHPN-1006 at p.139.

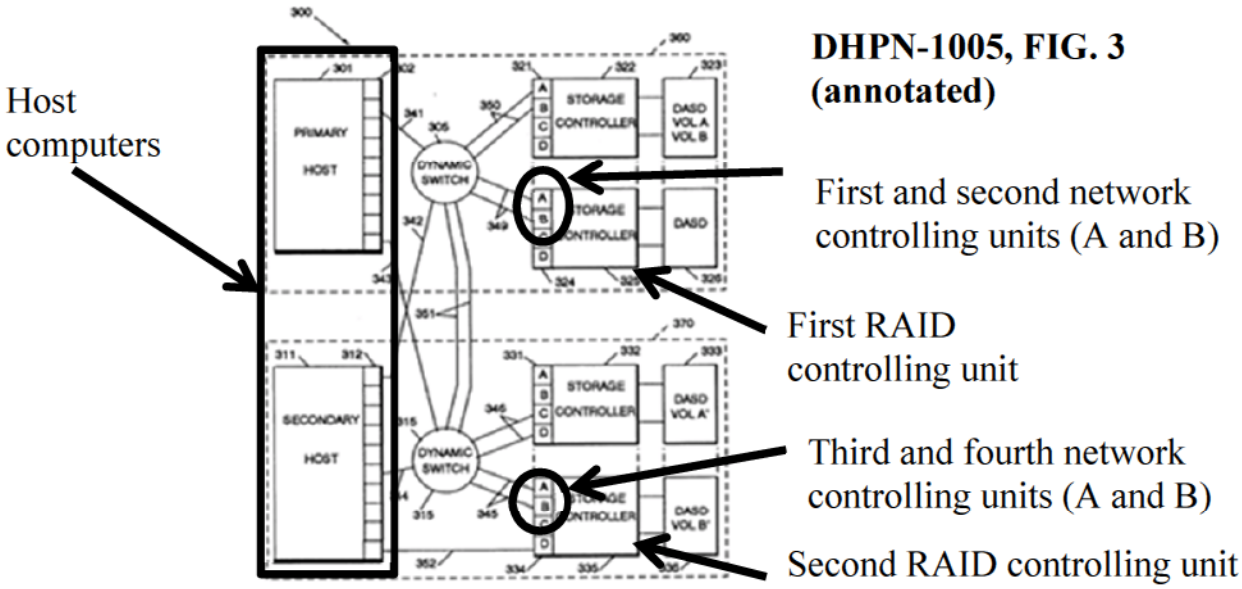

Thus, the second storage controller having at least two channel or control unit link-level facilities discloses "the second RAID controlling unit including a third network controlling unit and a fourth network controlling unit."

[1.4] ***a plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers***

The '950 Patent discloses this limitation because it teaches two dynamic switches that connect the two storage controllers with the primary and secondary hosts. DHPN-1006 at pp.139-141.

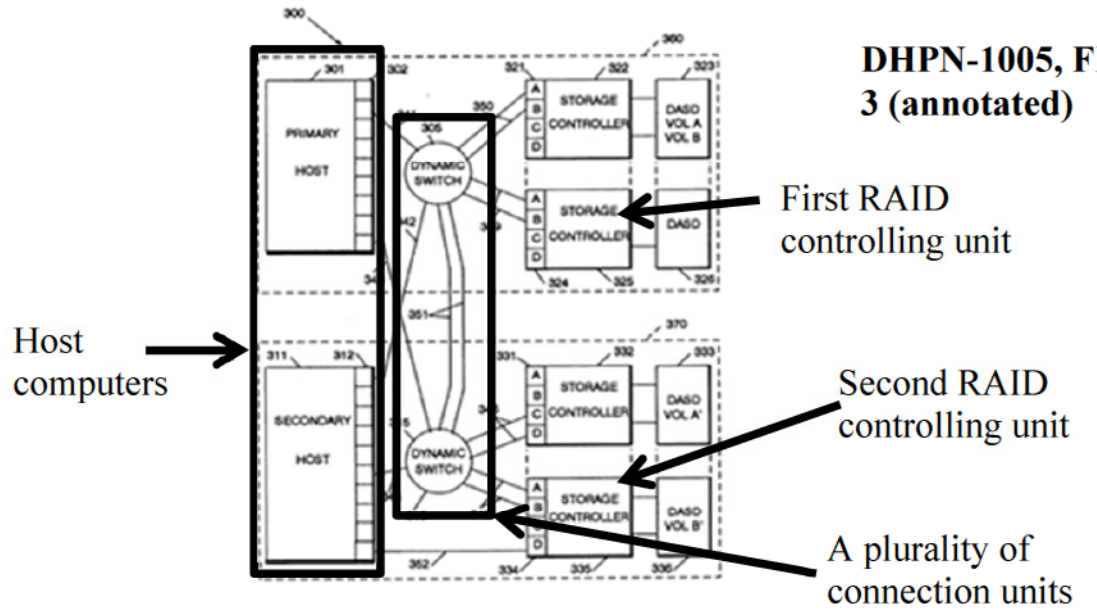

DHPN-1005, FIG. 3 (annotated)

Thus, the two dynamic switches that connect the first storage controller and the second storage controller to the primary and secondary hosts discloses "a plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers."

[1.5] ***wherein the first RAID controlling unit and the second RAID controlling unit directly exchange information with the numerous host computers through the plurality of connecting units***

The '950 Patent discloses this limitation because it teaches that the "primary storage controller" and the "secondary storage controller" (RAID controlling units) directly exchange data and records with the host computers through the dynamic switches. DHPN-1005, 7:28-35; s*ee also* 8:3-15 and Fig. 6, step 601; *see also* DHPN-1006 at pp.141-142.

Thus, the first storage controller and the second storage controller exchange records and data with each of the primary and secondary hosts through the

dynamic switches, thereby disclosing wherein the first RAID controlling unit and the second RAID controlling unit directly exchange information with the numerous host computers through the plurality of connecting units

[1.6] ***and the first network controlling unit exchanges information with the fourth network controlling unit***

The '950 Patent discloses this limitation because it teaches that the ports (network controlling units) of a primary storage controller communicate with the ports (network controlling units) of a secondary storage controller via switches 305 and 315:

> [P]rimary storage controller 322, via the same port A 321, can communicate with secondary storage controller 332 by communication links 350, dynamic switch 305, communication links 351, dynamic switch 315, and communication links 346, wherein port A 321 acts as a channel link-level facility.

DHPN-1005, 8:3-15; DHPN-1006 at pp.142-144.

The '950 Patent also explains that ports A and B 334 (the third and fourth network controlling units) initiate the operation of Fig. 4. DHPN-1006 at pp.143-144. Similarly, ports A and B 324 (first and second network controlling units) perform the data mirroring of Fig. 5. DHPN-1005, 8:61-63, 9:49-51; DHPN-1006 at pp.143-144.

Thus, the '950 Patent teaches that during normal operation, ports A and B 324 (first and second network controlling units) communicate with ports A and B

334 (third and fourth network controlling units). DHPN-1006 at p.144. Further, "the communication links between primary and secondary processors and between primary and secondary storage controllers may vary." DHPN-1005, 13:13-16; DHPN-1006 at p.144. Thus, the '950 Patent teaches that the first network controlling unit (e.g., port A 324) exchanges information with either or both of the third and fourth network controlling units (ports A and/or B 334). DHPN-1006 at p.144. Also, the second network controlling unit (e.g., port B 324) exchanges information with either or both of the third and fourth network controlling units (ports A and/or B 334). DHPN-1006 at p.144.

Therefore, the information exchange between and among any of the network controlling units of the storage controllers of the '950 Patent teaches "the first network controlling unit exchanges information with the fourth network controlling unit."

[1.7] ***and the second network controlling unit exchanges information with the third network controlling unit.***

As discussed above at [1.6], the second network controlling unit (e.g., port B 324) exchanges information with either or both of the third and fourth network controlling units (ports A and/or B 334). DHPN-1006 at p.145. Thus, the '950 Patent teaches "the second network controlling unit exchanges information with the third network controlling unit."

Claim 2

[2.1] ***The apparatus as recited in claim 1, wherein said respective RAID controlling units are connected to the plurality of individual connecting units.***

As discussed above, the '950 Patent discloses the apparatus of claim 1. The '950 Patent further discloses the limitation of [2.1] because, in Figure 3, it teaches the storage controllers 322, 325, 332, 335 (RAID controlling units) connected to the dynamic switches 305 and 315 (plurality of connecting units). DHPN-1006 at p.145-146.

Thus, the '950 Patent's storage controllers connected to the dynamic switches discloses "wherein said respective RAID controlling units are connected to the plurality of individual connecting units," as recited in the claim.

Claim 3

[3.1] ***The apparatus as recited in claim 2, wherein the first network interface controlling unit is coupled to the connecting unit of one side and the second network interface controlling unit is coupled to the connecting unit of another side.***

As discussed above, the '950 Patent discloses the apparatus of claim 2. The '950 Patent further discloses the limitation of [3.1] because it teaches that ports A and B 324 (first and second network controlling units) are both connected via links 349 to dynamic switch 305 (connecting unit) on an upper side and are also both connected to dynamic switch 315 (connecting unit) on a lower side via links 349, switch 305, and links 351. DHPN-1005, Fig. 3; DHPN-1006 at pp.146-148. Thus, the '950 Patent's storage controller ports coupled to respective switches discloses

"wherein the first network interface controlling unit is coupled to the connecting unit of one side and the second network interface controlling unit is coupled to the connecting unit of another side."

Claim 5

[5.1] ***The apparatus as recited in claim 1, wherein said plurality of connecting units have at least three connection ports,***

As discussed above, the '950 Patent discloses the apparatus of claim 1. The '950 Patent further discloses the limitation of [5.1] because it teaches, in Figure 3, dynamic switch 305 having at least eight ports—one port for link 341, two ports for links 350, two ports for links 349, two ports for links 351, and one port for link 342. DHPN-1005 at Fig. 3; DHPN-1006 at p.148. Dynamic switch 315 has a similar number of connection ports. DHPN-1006 at p.148.

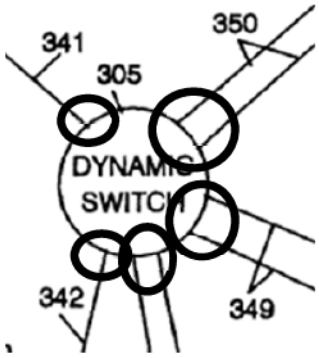

DHPN-1005, Figure 3 (truncated and annotated)

Switch 305 with eight ports circled.

The total number of links for the plurality of connection units is at least 16. DHPN-1006 at pp.148-150. Thus, the 16 links for the plurality of connection units discloses "said plurality of connecting units have at least three connection ports."

[5.2] ***two of the at least three connection ports is coupled to one of the first network interface controlling unit and the third network controlling unit***

The '950 Patent discloses this limitation because it teaches one port of switch 305 connected to the first network interface controlling unit by link 349. DHPN-1006 at p.151. The other seven ports of dynamic switch 305 are coupled to the first network interface controlling unit by virtue of being included in switch 305 (all ports of a switch are connected to all other ports of the switch virtue of the connections provided by the switch). DHPN-1006 at p.151. Therefore, all of the at least eight ports of switch 305 are coupled to the first network controlling unit either directly or indirectly. DHPN-1006 at p.151. In a similar way, all of the at least eight ports of dynamic switch 315 are coupled to the third network controlling unit, either directly or indirectly. DHPN-1006 at p.151. Thus, the ports of switch 305 being coupled to ports at storage controller 325 disclose "two of the at least three connection ports is coupled to … the first network interface controlling unit" and the ports of switch 315 being coupled to ports at storage controller 335 disclose "two of the at least three connection ports is coupled to … the third network controlling unit."

[5.3] ***and the rest of the connection ports being provided as a hub equipment connected with the numerous host computers.***

The '950 Patent discloses this limitation because it teaches dynamic switches 305 and 315 connected to the primary and secondary hosts. Since the '346 Patent defines "hub" to include a hub or switch (see '346 Patent at 3:15-18), the dynamic switches 305 and 315 (plurality of connecting units) disclose the

claimed hub equipment. DHPN-1006 at p.152. The ports discussed above at [5.1] and [5.2] are thus provided as a hub equipment connected to the host computers 301 and 311. DHPN-1006 at p.152. Thus, the '950 Patent's switches having ports discloses "the rest of the connection ports being provided as a hub equipment connected with the numerous host computers."

Claim 6

[6.0] ***The apparatus as recited in claim 1, wherein said plurality of connecting units have at least three connection ports,***

This limitation is identical to the limitation of element [5.1] discussed above, and for those same reasons is disclosed by the '950 Patent.

[6.2] ***two of the at least three connection port are coupled to one of the first network interface controlling unit and the third network controlling unit***

This limitation is identical to the limitation of element [5.2] discuss above, and for those same reasons is disclosed by the '950 Patent.

[6.3] ***and the rest of the connection ports being provided as a network switch equipment connected with the numerous host computers.***

The '950 Patent discloses this limitation because it teaches dynamic switches 305 and 315 (plurality of connecting units) that are provided as network switch equipment and which are connected to the primary and second hosts. DHPN-1006 at p.156. Thus, the '950 Patent's switches having ports discloses "the rest of the connection ports being provided as a network switch equipment connected with the numerous host computers."

Claim 7

[7.1] ***The apparatus as recited in claim 1, wherein said plurality of connecting units have at least five connection ports,***

As discussed above, the '950 Patent discloses the apparatus of claim 1. The '950 Patent further discloses the limitation of [7.1] because it teaches that the switches 305, 315 have a total of at least sixteen ports. DHPN-1006 at p.157. Thus, the sixteen ports of the plurality of switches in the '950 Patent discloses "wherein said plurality of connecting units have at least five connection ports."

[7.2] ***four of the at least five connection ports is coupled to one of the first network interface controlling unit and the third network controlling unit***

The analysis of [7.2] is the same as that given above for [5.2], noting that all of the at least eight ports of switch 305 are coupled to the first network controlling unit, and all of the at least eight ports of dynamic switch 315 are coupled to the third network controlling unit. DHPN-1006 at pp.158-160.

[7.3] ***and the rest of the connection ports being provided as a switch connected with the numerous host computers***.

The '950 Patent discloses this limitation because it teaches dynamic switches 305 and 315 (plurality of connecting units) that are connected to the primary and second hosts. DHPN-1006 at p.160. Thus, the '950 Patent's switches having ports connected to the primary and second hosts discloses "the rest of the connection ports being provided as a switch connected with the numerous host computers."

Claim 8

[8.1] ***The apparatus as recited in claim 1, wherein the first network interface controlling unit of the first RAID controlling unit being connected to a first connecting unit, the second network interface controlling unit of said first RAID controlling unit being connected to a second connecting unit, the third network interface controlling unit of the second RAID controlling unit being connected to the second connecting unit, and the fourth network interface controlling unit of the second RAID controlling unit being connected to the first connecting unit.***

The '950 Patent discloses this limitation because it teaches, in Fig. 3, that:

- *the first network interface controlling unit of the first RAID controlling unit being connected to a first connecting unit*: Port A 324 is connected to dynamic switch 305 by link 349.
- *the second network interface controlling unit of said first RAID controlling unit being connected to a second connecting unit:* Port B 324 is connected to the dynamic switch 315 by links 351, switch 305, and links 349.
- *the third network interface controlling unit of the second RAID controlling unit being connected to the second connecting unit*: Port A 334 is connected to the dynamic switch 315 by link 345.
- *the fourth network interface controlling unit of the second RAID controlling unit being connected to the first connecting unit*: Port B 334 is connected to switch 305 by link 345, switch 315, and links 351.

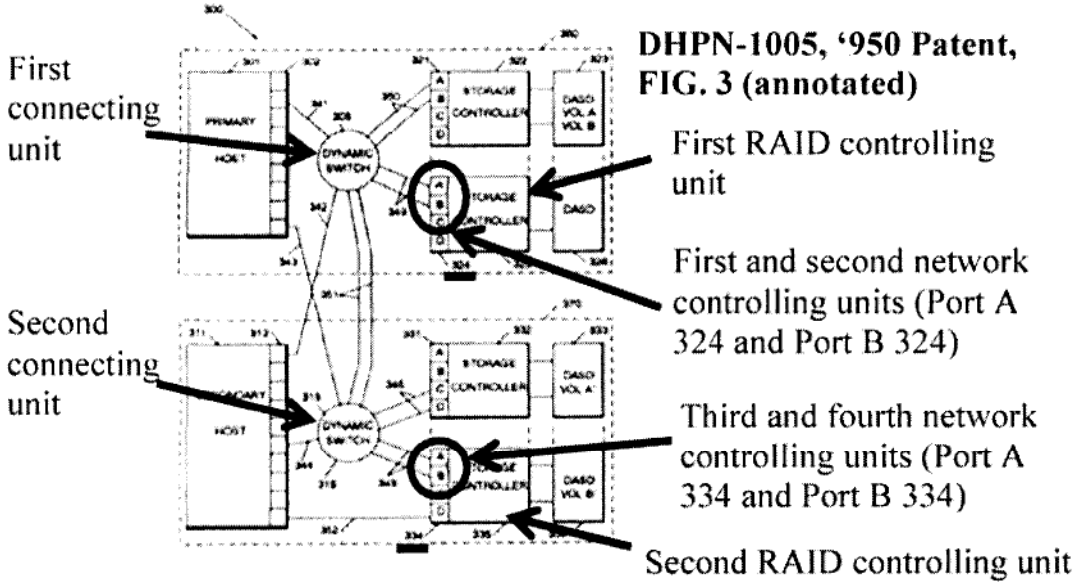

DHPN-1005, '950 Patent, FIG. 3 (annotated)

DHPN-1006 at pp.160-163. Therefore, the '950 Patent's connections illustrated in Fig. 3 discloses this limitation

VI. Conclusion

For the reasons set forth above, Petitioners have established a reasonable likelihood of prevailing with respect to at least one claim of the '346 Patent and ask that the Patent Office order an *inter partes* review trial and then proceed to cancel claims 1-9.

Respectfully submitted,

Dated: September 27, 2013

David McCombs
Registration No. 32,271

HAYNES AND BOONE, LLP
Customer No. 112792
Telephone: (214) 651-5533
Facsimile: (214) 200-0853
Attorney Docket No.: 47415.430

PETITIONER'S EXHIBIT LIST

September 27, 2013

DHPN-1001	Baek et al., U.S. Patent No. 6,978,346
DHPN-1002	Prosecution History of U.S. Patent No. 6,978,346
DHPN-1003	Peter Weygant, *Clusters for High Availability: A Primer of HP-UX Solutions*, 1996 ("Weygant")
DHPN-1004	*Managing MC/ServiceGuard*, Hewlett-Packard Company, 1998 ("ServiceGuard")
DHPN-1005	Hathorn et al., U.S. Pat. No. 5,574,950 ("the '950 Patent")
DHPN-1006	Declaration of Dr. Ray Mercer
DHPN-1007	Surugguchi et al., International Publication No. WO 99/38067 ("Mylex")
DHPN-1008	*American National Standard for Information Technology – Fibre Channel Arbitrated Loop (FC-AL-2)*, June 28, 1999 ("ANSI")

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re patent of: Baek et al.	§	Petition for *Inter Partes* Review
	§	
U.S. Patent No. 6,978,346	§	Attorney Docket No.: 47415.430
	§	
Issued: December 20, 2005	§	Customer No.: 112792
	§	
Title: APPARATUS FOR REDUNDANT INTER-CONNECTION BETWEEN MULTIPLE HOSTS AND RAID	§	Real Parties in Interest: Dell Inc., Hewlett-Packard Company, and NetApp, Inc.
	§	
	§	

CERTIFICATE OF SERVICE

The undersigned certifies, in accordance with 37 C.F.R. § 42.205, that service was made on the Patent Owner as detailed below.

Date of service September 27, 2013

Manner of service FEDERAL EXPRESS

Documents served Petition for *Inter Partes* Review Transmittal Form;

Petition for *Inter Partes* Review;

Petitioner's Exhibit List;

Exhibits DHPN-1001 through DHPN-1008

Persons served Blakely Sokoloff Taylor & Zafman
1279 Oakmead Parkway
Sunnyvale, CA 94085-4040

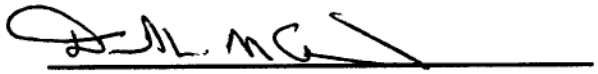

David McCombs
Registration No. 32,271