

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

VMware, Inc.
Petitioner

v.

Safe Storage LLC
Patent Owner

Case IPR2014-_____
Patent 6,978,346

DECLARATION OF ROBERT HORST, PH.D IN SUPPORT OF *INTER PARTES* REVIEW OF U.S. PATENT NO. 6,978,346

TABLE OF CONTENTS

I.	INTRODUCTION	4
II.	BACKGROUND	4
III.	LEVEL OF ORDINARY SKILL IN THE ART	7
IV.	APPLICABLE LEGAL STANDARD	8
	A. Claim Construction	8
	B. Anticipation And Obviousness	9
V.	TECHNOLOGY OF THE '346 PATENT	11
VI.	CHALLENGE #1 – CLAIMS 1-9 ARE RENDERED OBVIOUS BY THE MYLEX PAPER (Exhibit VMWARE-1006) IN VIEW OF THE TEACHINGS OF THE HATHORN PATENT (Exhibit VMWARE-1005)	17
	A. Brief Summary of the Mylex Paper	20
	B. Brief Summary of the Hathorn Patent	25
	C. The Mylex Paper In View Of The Teachings Of The Hathorn Patent Compared To The '346 Patent, Claims 1-9.....	32
VII.	CHALLENGE # 2 - CLAIMS 1-9 ARE RENDERED OBVIOUS BY THE HATHORN PATENT (Exhibit VMWARE-1005) IN VIEW OF THE TEACHINGS OF THE MYLEX PAPER (Exhibit VMWARE-1006)	78
	A. The Hathorn Patent In View Of The Teachings Of The Mylex Paper Compared To The '346 Patent, Claims 1-9.....	81
VIII.	CHALLENGE # 3 – CLAIMS 1-9 ARE RENDERED OBVIOUS BY THE DEITZ PATENT (Exhibit VMWARE-1008) OR THE MYLEX PAPER (Exhibit VMWARE-1006) IN VIEW OF THE TEACHINGS OF THE GRIFFITH PATENT (Exhibit VMWARE-1007) OR THE DEKONING PATENT (Exhibit VMWARE-1010)	112
	A. Brief Summary of the Griffith and DeKoning Patents.....	113
	B. The Mylex Paper In View Of The Teachings Of The Griffith and/or DeKoning Patent Compared To The '346 Patent, Claims 1-9.....	119
	C. Brief Summary of the Deitz Patent	124

D. The Deitz Patent In View Of The Teachings Of The Griffith and/or DeKoning Patents Compared To The '346 Patent, Claims 1-9..... 126

IX. STATEMENT UNDER U.S.C. SECTION 1001 OF TITLE 18.....165

I. INTRODUCTION

1. I have been retained on behalf of the Petitioner VMware, Inc. to provide this Declaration concerning technical subject matter relevant to the *inter partes* review of U.S. Patent No. 6,978,346 (the “346 patent”). I reserve the right to supplement this Declaration in response to additional evidence that may come to light.

2. I am over 18 years of age. I have personal knowledge of the facts stated in this Declaration and could testify competently to them if asked to do so.

II. BACKGROUND

3. My name is Robert W. Horst. I am an independent consultant with more than 30 years expertise in the design and architecture of computer and storage systems. I hold a Bachelor of Science degree in Electrical Engineering from Bradley University, a Master of Science in Electrical Engineering from the University of Illinois at Urbana-Champaign, and a Ph.D. in Computer Science, also from the University of Illinois.

4. Currently, I am an independent consultant for HT Consulting, where my work includes technology consulting and serving as an expert witness in patent and technology litigation. I am also Chief Technology Officer of Robotics for AlterG, Inc., where I am working on the design of orthotic devices to assist those with impaired mobility.

5. I have testified as an expert witness and served as a consultant in patent and intellectual property litigation and in *inter partes* review and re-examination proceedings.

6. I have worked as a Technical Director at Network Appliance, Inc., where I worked on processor and interconnect options for future generations of network-attached storage systems. I also served as Vice President of Research & Technology at 3ware, Inc., where I developed low-cost RAID controllers and initiated and led a project on one of the industry's first Ethernet Storage Area Network RAID storage systems. Prior to 3Ware, I worked as Technical Director at Tandem Computers/Compaq Computers where I contributed to the design and architecture of several generations of fault-tolerant systems. This work included development of CPUs, system-area networks, I/O systems, and storage systems. I have also been published in the areas of Networks, Storage, CPU Architecture, Fault Tolerance, and Bionics. In 2001, I was elected as an IEEE Fellow "for contributions to the architecture and design of fault tolerant systems and networks." I have worked with patent attorneys on numerous patent applications, and I am a named inventor on 78 issued U.S. patents.

7. My qualifications and experience are set forth in more detail in my Curriculum Vitae, which is being filed as Exhibit VMWARE-1004.

8. I am being compensated for my work preparing this report. My compensation is not contingent upon the outcome of this proceeding or the opinions I develop in this matter.

9. My preliminary opinions expressed herein are based on review and analysis of certain information obtained in connection with my work in this matter, together with my training, education, and experience. The opinions expressed herein are my own.

10. The following identifies the information relied upon to date in connection with my work:

- (1) VMWARE-1001: U.S. Patent No. 6,978,346 to Baek et al., foreign application priority date 9/19/2000 (“the ’346 patent”);
- (2) VMWARE-1002: Excerpts from the Prosecution History of the ’346 Patent;
- (3) VMWARE-1005: U.S. Patent No. 5,574,950 to Hathorn et al., issued 11/12/1996 (“Hathorn”);
- (4) VMWARE-1006: Smith, Kevin J., “Storage Area Networks: Unclogging LANs and Improving Data Accessibility,” Mylex Corporation White Paper (published 5/29/1998) (“Mylex paper”);
- (5) VMWARE-1007: U.S. Patent No. 6,401,170 to Griffith et al., filed on 8/18/1999 (“Griffith”);
- (6) VMWARE-1008: U.S. Patent No. 6,578,158 to Deitz et al., filed on 10/28/1999 (“Deitz”);
- (7) VMWARE-1009: Affidavit of Mr. Chris Butler, on behalf of Internet Archive;

- (8) VMWARE-1010: U.S. Patent No. 6,073,218 to DeKoning, et al., filed on 12/23/1996 (“DeKoning”);
- (9) VMWARE-1011: Clark, “Designing Storage Area Networks,”1st Edition, Addison-Wesley Professional (1999);
- (10) VMWARE-1012: Spainhower, “Design for Fault-Tolerance in System ES /9000 Model 900,” IEEE (1992);
- (11) VMWARE-1013: IEEE 100: Authoritative Dictionary of IEEE Standards Terms, 7th Edition (2000); and
- (12) VMWARE-1014: Siewiorek, D and Swarz R., “Reliable Computer Systems, Design and Evaluation,” Digital Press (1992).

III. LEVEL OF ORDINARY SKILL IN THE ART

11. Based on my education and extensive experience relating to RAID storage systems and fault-tolerant systems, I believe I am qualified to provide opinions about the understanding and qualifications of a person of ordinary skill in the art of the technology at issue in this proceeding.

12. In my opinion, a person of ordinary skill in the art of the '346 patent, as of 2000, would have had a B.S. in Electrical Engineering or Computer Science and at least two years of experience in designing storage systems.

13. My opinions below explain how a person of ordinary skill in the art would have understood the technology described in the references I have identified below around the 2000 time period.

IV. APPLICABLE LEGAL STANDARD

A. Claim Construction

14. I understand that in an *inter partes* review proceeding, the claims of a patent are to be given their broadest reasonable meaning as they would be understood by one of ordinary skill in the art, consistent with the specification of the patent.

15. It is my understanding that the Patent Trial and Appeal Board previously construed certain '346 patent claim terms in an *inter partes* review filed by Dell Inc., Hewlett-Packard Company, and NetApp. For the purposes of my opinions set forth herein, I have used the Patent Trial and Appeal Board's constructions identified below. I reserve the right to offer an opinion as to the proper construction of other claim terms in this proceeding. At this time, I have no opinion as to whether these constructions would be the proper constructions for any district court litigation involving the '346 patent.

Claim Term	Construction¹
"RAID controlling unit" and "RAID controller"	"A component that controls operation of the RAID"
"RAID"	"Redundant array of inexpensive disks"
"exchange/exchanges information"	"To transmit and receive information reciprocally"
"connection unit"	"a hub or switch"

¹ See IPR2013-00635, Paper 19 at pp. 8-11.

16. Further, I understand that the Patent Owner stated in a prior IPR proceeding (IPR2013-00635, Paper 14 at p. 19) that a “network interface controller is the part of a RAID controller that allows the RAID controller to communicate with the ‘connection units.’” Therefore, for purposes of this proceeding, I incorporate the construction of the claim terms “network interface controller,” “network controlling unit,” and “network interface controlling unit,” as “the part of a RAID controller that allows the RAID controller to communicate with the ‘connection units.’”

B. Anticipation And Obviousness

17. I have been informed that a patent claim is invalid as anticipated under 35 U.S.C. § 102 if each and every element of a claim, as properly construed, is found either explicitly or inherently in a single prior art reference. Under the principles of inherency, if the prior art necessarily functions in accordance with or includes the claimed elements, it anticipates.

18. I have been informed that a claim is invalid under 35 U.S.C. § 102(a) if the claimed invention was known or used by others in the U.S., or was patented or published anywhere, before the applicant’s invention. I further have been informed that a claim is invalid under 35 U.S.C. § 102(b) if the invention was patented or published anywhere, or was in public use, on sale, or offered for sale in this country, more than one year prior to the filing date of the patent application

(critical date). I further have been informed that a claim is invalid under 35 U.S.C. § 102(e) if an invention described by that claim was disclosed in a U.S. patent granted on an application for a patent by another that was filed in the U.S. before the date of invention for such a claim.

19. I have been informed that a patent claim is invalid as “obvious” under 35 U.S.C. § 103 in light of one or more prior art references if it would have been obvious to one of ordinary skill in the art, taking into account (1) the scope and content of the prior art, (2) the differences between the prior art and the claims, (3) the level of ordinary skill in the art, and (4) any so called “secondary considerations” of non-obviousness, which include: (i) “long felt need” for the claimed invention, (ii) commercial success attributable to the claimed invention, (iii) unexpected results of the claimed invention, and (iv) “copying” of the claimed invention by others. For purposes of my analysis above, and because I know of no indication from the patent owner or others to the contrary, I have applied a date of September 19, 2000, as the date of invention in my obviousness analyses, although in many cases the same analysis would hold true even at an earlier time than September 19, 2000.

20. I have been informed that a claim can be obvious in light of a single prior art reference or multiple prior art references. To be obvious in light of a single prior art reference or multiple prior art references, there must be a reason to

modify the single prior art reference, or combine two or more references, in order to achieve the claimed invention. This reason may come from a teaching, suggestion, or motivation to combine, or may come from the reference or references themselves, the knowledge or “common sense” of one skilled in the art, or from the nature of the problem to be solved, and may be explicit or implicit from the prior art as a whole. I have been informed that the combination of familiar elements according to known methods is likely to be obvious when it does no more than yield predictable results. I also understand it is improper to rely on hindsight in making the obviousness determination. *KSR Int’l Co. v. Teleflex Inc.*, 550 U.S. 398, 421 (2007).

V. TECHNOLOGY OF THE ’346 PATENT

21. The ’346 patent relates to interconnections between host computers and storage systems. The storage systems referenced in the patent and claims are those known by the acronym RAID, which stands for Redundant Array of Inexpensive (or sometimes, Independent) Disks. The term RAID was first used in a 1987 paper by David Patterson and Randy Katz to describe storage systems built from multiple low-cost disk drives and configured to improve the reliability and/or performance of the storage system. Over time, multiple different levels of RAID were developed, including disk striping (RAID 0), disk mirroring (RAID 1), and various forms of parity protection across groups of drives (RAID 2 through RAID

6). The '346 patent does not specifically address a particular type of RAID system and instead focuses on connections between the RAID and the host computers.


22. By the 2000 time frame, several different types of networks were being used to connect RAID systems to hosts. Networks providing block-level access to storage were called SANs (for storage-area networks) and these networks included FCAL (fibre channel arbitrated loop), switched fibre channel and IBM's ESCON (Enterprise Systems Connection). Other networks used for block-level or file level access to storage included ATM (Asynchronous Transfer Mode), ServerNet, InfiniBand and Ethernet. The '346 patent gives examples of some networks for connecting storage, but does not make claims specific to any particular network. The '346 patent groups them together, calling them "industrial standard communications networks," and also describes them using the general term "network" ('346 patent at 3:25-29).

23. The terms "fault tolerance" and "fault tolerant system" describe systems that continue to function when part of the system encounters a fault. The concept of fault tolerant systems dates back to fault tolerant telephone switching systems introduced by AT&T in the 1960s and commercial fault-tolerant systems

first introduced by Tandem Computers in the 1970s.² The focus of the '346 patent is on the configuration of redundant network connections between hosts and RAID storage to assure that fault tolerance and performance are maintained when a RAID controller experiences a fault.

24. I reproduced Figure 4 of the '346 patent below:

² For a summary of commercial fault tolerant systems, *see* Siewiorek, D and Swarz R., *Reliable Computer Systems*, 1992. (Exhibit VMWARE-1014). Page 568 shows redundant host to storage connections in the AT&T 3B20D system (1981). Page 589 shows redundant host to storage connections in the Tandem NonStop (1976). Page 619 shows redundant host to storage connections in the Tandem Integrity S2 (1991).



25. The '346 patent describes Figure 4 as “a block diagram showing one embodiment of a host interface system as an internal installment system between a RAID and host computers in accordance with the present invention.” ('346 patent at 2:46-49.) Figure 4 identifies a system where three “host computer[s]” (labeled 400-402) are connected to a first “hub or switch” (labeled 440), and further connected to a RAID (labeled 490) through a network interface controlling unit port (labeled 470) on a first RAID controller (labeled 460) and a network interface controlling unit port (labeled 480) on a second RAID controller (labeled 461). Figure 4 also illustrates that three other “host computers” (labeled 403-405) are connected to a second “hub or switch” (labeled 441), and further connected to the RAID through a network interface controlling unit port (labeled 471) on the first

RAID controller and a network interface controlling unit port (labeled 481) on the second RAID controller.

26. As shown in Figure 4 above, network interface controlling unit port 470 is networked through “hub or switch” 440 to network interface controlling unit port 480, and network interface controlling unit port 471 is networked through “hub or switch” 441 to network interface controlling unit port 481. The ’346 patent discloses that “information” can be transmitted on these networks, but does not describe what type of information is transmitted.

27. The ’346 patent does not disclose any specific modifications to the network interface controlling unit ports, the RAID controllers, or the hubs or switches that need to be made in order to transmit information on the networks between network interface controlling units. At most, the ’346 patent discloses that the mere addition of communication lines allowing both “hubs or switches” to connect both RAID controllers (*e.g.*, in Figure 4, see lines labeled 450) is sufficient to allow communication between the network interface controlling units on two RAID controllers. The ’346 patent describes these communication lines broadly as follows:

a communication line, representatively shown as 450 in the drawing, for connecting the network interface controller to the hub is a copper line or an optical fibre, which is matched to a corresponding standard. (’346 patent at 3:39-42.)

28. I understand that the claims at issue in this proceeding are claims 1-9, with claims 1 and 9 being independent.

29. Claim 1 of the '346 patent recites the following:

[1a] An apparatus for a redundant interconnection between multiple hosts and a RAID, comprising:

[1b] a first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers, the first RAID controlling unit including a first network controlling unit and a second network controlling unit, and the second RAID controlling unit including a third network controlling unit and a fourth network controlling unit;

[1c] a plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers, wherein the first RAID controlling unit and the second RAID controlling unit directly exchange information with the connecting units, and 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.

30. Claim 9 of the '346 patent recites the following:

[9a] An apparatus for a redundant interconnection between multiple host computers and a RAID, the apparatus comprising:

[9b] a plurality of connecting units for connecting the host computers and the RAID;

[9c] a first and a second RAID controllers, included in the RAID, 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,

[9d] 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 processes information transmitted from the second network interface controller in the second RAID controller,

[9e] 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 processes information transmitted from the second network interface controller in the first RAID controller,

[9f] 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

[9g] 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

[9h] 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 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.

VI. CHALLENGE #1 – CLAIMS 1-9 ARE RENDERED OBVIOUS BY THE MYLEX PAPER (Exhibit VMWARE-1006) IN VIEW OF THE TEACHINGS OF THE HATHORN PATENT (Exhibit VMWARE-1005)

31. It is my opinion that the Mylex paper in view of the teachings of the Hathorn patent renders obvious claims 1-9 of the '346 patent.³ I provide a brief description of the Mylex paper and the Hathorn patent below and then a more

³ In setting forth my opinions regarding obviousness within this declaration, I am not setting forth the opinion that any reference does not anticipate any of the '346 patent's claims.

detailed discussion identifying the disclosures in the Mylex paper and teachings of the Hathorn patent that support my opinion.

32. As described below, a person of ordinary skill would understand that the Mylex paper discloses every element of the '346 patent's claims 1-9, with the exception of a direct exchange of information between network interface controlling units. Instead, the Mylex paper discloses a direct "heartbeat" communication path between controllers for exchanging information. The Hathorn patent, on the other hand, teaches that communication paths are expensive, and that this expense can be reduced by modifying network interface controlling unit ports to use the existing switch network for communications between RAID controllers (instead of using a direct "heartbeat" path).

33. A person of ordinary skill in the art would have been motivated to apply these Hathorn patent teachings to the system disclosed in the Mylex paper in order to render every claim in the '346 patent obvious. For example, both the Mylex paper and the Hathorn patent are in the same field of endeavor. Both references disclose redundant RAID systems that connect multiple hosts to switches or hubs, which in turn connect to RAID controllers with two or more ports. Both references disclose redundancy in terms of sending communications between two or more RAID controllers and/or network interface controlling unit ports. Additionally, both references are concerned with RAID 1 (disk

mirroring/shadowing). (See Mylex paper at 12 (“SAN-attached RAID arrays should support disk mirroring”); see Hathorn patent at 1:9-12 (“The present invention relates generally to remote data shadowing...”).) One of ordinary skill in the art would have been motivated to study multiple instances of systems for disk mirroring when designing a new RAID system. Further, both references disclose redundant RAID systems and disclose using off-the-shelf components for constructing the RAID system, and, as such, their combination is merely the use of known techniques to achieve predictable results. (See, e.g., Mylex paper at 15 (marketing “Mylex controllers”); Hathorn patent at 6:25-34 (describing an IBM Enterprise Systems/9000 (ES/9000) processor running DFSMS/MVS operating software, IBM 3990 Model 6 storage controllers, and an IBM ESCON Director dynamic switch).)

34. Finally, one of ordinary skill in the art would be motivated to combine the teachings of the Hathorn patent with the Mylex controllers because there was a close relationship between IBM, the assignee of the Hathorn patent, and Mylex Corporation. In September of 1999, IBM completed the acquisition of Mylex. Storage system designers at both companies in that timeframe would have been strongly motivated to combine and leverage storage technology from the other company. In fact, later IBM products were partly based on the acquired Mylex

technology, demonstrating that the motivation to combine the features actually resulted in new products.

A. Brief Summary of the Mylex Paper

35. The whitepaper titled “Storage Area Networks: Unclogging LANs and Improving Data Accessibility” by Kevin J. Smith of Mylex Corporation (“the Mylex paper”) generally discloses Mylex’s Fibre Channel RAID controllers and the use of storage area networks to configure reliable and high-performance pools of storage. The Mylex paper was published on May 29, 1998, and made available on the Mylex public web site (www.Mylex.com). I understand that Petitioner is submitting a declaration by a business records custodian for *archive.org* identifying that the Mylex paper was publically available on the Mylex web site (www.Mylex.com) at least as early as February 4, 1999. (Exhibit VMWARE-1009).

36. The Mylex paper includes illustrations of several configurations of the “seamless product line of external Raid Controllers” and explains that “Mylex array controllers are available in simplex configurations for network servers and duplex (dual) configurations for SAN’s and clusters. In duplex mode, advanced features are implemented to accelerate performance, protect data and guarantee data accessibility.” (Mylex Paper at 14). While some of the figures focus on particular features of one configuration, one of ordinary skill in the art would

understand that the features could be combined in a single system. The Mylex paper discloses Fibre Channel arbitrated loop hub topologies and switched topologies, and combination hub and switch topologies that include redundant connections between hosts and RAID arrays to allow host-independent failover.

For example, the Mylex paper states:

Mylex controllers have dual SAN ports which *doubles the bandwidth to controllers and allows redundant paths from other SAN devices to the controllers* to increase the resiliency of the SAN topology. As described later in this paper, dual host ports are particular critical for controller failover in Fibre Channel topologies. *The SAN ports can be connected* directly to UNIX and NT servers or indirectly *through hubs and switches*. (Mylex Paper at 16, emphasis added.)

37. For example, the Mylex paper's Figures 6 and 20-21 are reproduced below, showing (i) a combination switch and hub cascade topology (Figure 6), (ii) hub topologies (Figure 20), and (iii) switch topologies (Figure 21):




Figure 6. SAN With Switched and Shared (Loop) Interconnects




Figure 20. Mylex External Array Controllers in a Loop SAN Topology




Figure 21. Mylex External Array Controllers in a Switched SAN Topology

38. As shown in in the figures above, the Mylex paper discloses that multiple hosts (each associated with a fibre channel host-bus adapter “FC HBA” in Figures 20 and 21) may be connected to at least two RAID controllers (labeled, e.g., FL Controllers 0 and 1 in Figures 20 and 21) through a plurality of connecting

units (labeled as hubs and switches). While only one hub or switch is shown in Figures 20 and 21, one of ordinary skill in the art would understand the SANs implemented with hubs and switches are often implemented and extended with multiple hubs or switches. The RAID controllers include at least two network interface controlling units (labeled Port 0 and Port 1 in Figures 20 and 21).

39. The Mylex paper also discloses a feature of the Mylex controller that provides high availability by heartbeat monitoring and transparent controller failover/failback. For example, the Mylex paper's Figures 17 and 18 are reproduced below:

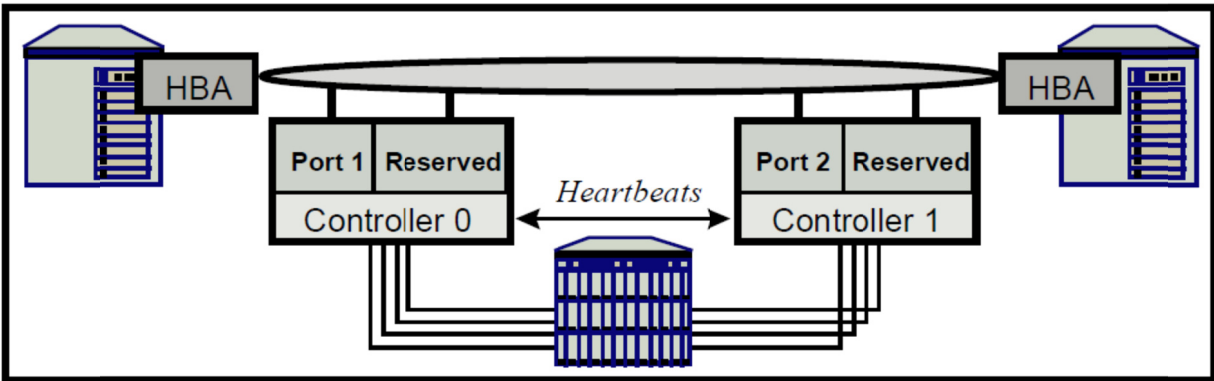


Figure 17. Normal Operating Mode of Dual DAC SF and FL Fibre Controllers




Figure 18. DAC SF and FL Fibre Array Controller Failover

40. As shown in Figures 17 and 18, the Mylex paper teaches that RAID controllers can provide fault tolerance. Two Mylex RAID controllers exchange information by sending “heartbeat” signals through a path linking the two controllers. When one RAID controller has a fault, that fault is detected by the reserved network controller port on the non-faulty RAID controller, which assumes the network ID of the faulty network controller port in order to process the host I/O requests being sent to the faulty network controller port. As such, in the event of a fault occurrence, there is no loss in bandwidth to the RAID storage.

41. The Mylex paper’s Figure 19 reproduced below illustrates that “[w]hen the failed controller is replaced, it is detected by the surviving controller which allows it to restart and returns the failed controller’s port ID’s, and then it starts processing I/O” (Mylex paper at 18):




Figure 19. DAC SF and FL Fibre Array Controller Failback

B. Brief Summary of the Hathorn Patent

42. U.S. Patent No. 5,574,950 (the “Hathorn patent”) is titled “Remote Data Shadowing Using A Multimode Interface To Dynamically Reconfigure Control Link-Level And Communication Link-Level.” The Hathorn patent issued on November 12, 1996 and is assigned to IBM.

43. The Hathorn patent generally discloses the use of dynamic switches to couple primary and secondary hosts to local and remote storage controllers. The Hathorn patent discloses “data shadowing” (*see, e.g.*, 1:9-10), which one of ordinary skill in the art would understand to be synonymous with “data mirroring,” a type of RAID architecture known as RAID level 1.⁴

⁴ *See, e.g.*, IEEE 100: Authoritative Dictionary of IEEE Standards Terms, 7th Edition, 2000 (Exhibit VMWARE-1013) at pp. 915-16, definition of “RAID storage”: “... Six basic architectures of RAID storage, referred to as levels 0 through 5 have been defined; see corresponding figure [reproduced below].”


44. The Hathorn patent generally discusses DASDs (direct access storage devices) and discloses that multiple DASDs can be configured as a RAID. For example, the Hathorn patent states:

Another data back-up alternative that overcomes the need to double the storage devices involves writing data to a redundant array of inexpensive devices (RAID) configuration. In this instance, the data is written such that the data is apportioned amongst many DASDs. If a single DASD fails, then the lost data can be recovered by using the remaining data and error correction procedures. Currently there are several different RAID configurations available. (Hathorn patent at 2:4-11.)

45. The Hathorn patent's Figure 3 is reproduced below:

Level	Description
0	Data striping without parity
1	Mirrored disk array
2	
3	Parallel disk array
4	Independent disk array
5	Independent disk array

FIG. 3



46. As shown in Figure 3, the Hathorn patent discloses that multiple hosts (labeled 301 and 311) are connected to two switches (labeled 305 and 315), then connected to multiple storage controllers (labeled 322, 325, 332, and 335), then connected to multiple DASDs (which, as described above, can be configured as a RAID, including DASDs 323, 326, 333, and 336). More specifically, each switch is connected to at least two network interface controlling unit ports (labeled 321A&B ports, 324 A&B ports, 331 C&D ports, and 334 A&B ports).

47. The Hathorn patent provides examples of standard components for this system. For example:

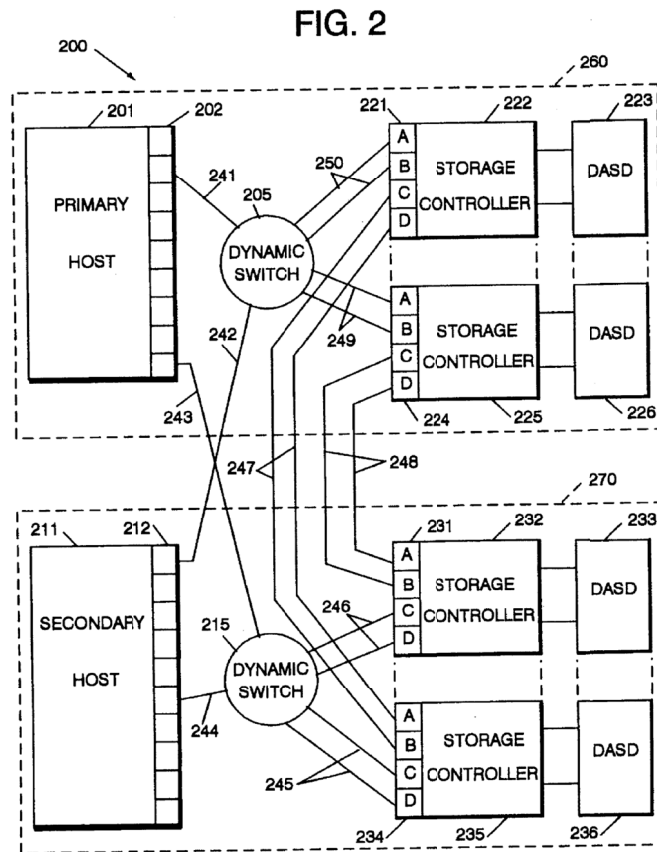
The primary site 260 includes a host or primary processor 201 (herein after referred to as primary host 201), for example, an IBM Enterprise Systems/9000 (ES/9000) processor running DFSMS/MVS operating software ... A plurality of primary storage controllers 222, 225, for example, IBM 3990 Model 6 storage controllers, are coupled to the primary host 201 via a dynamic switch 205, for example, an IBM ESCON Director. (*Id.* at 6:25-34.)

48. The Hathorn patent teaches that the network interface controlling unit ports can be dynamically modified in order to allow the ports to communicate with each other on the existing switch network. For example, the Hathorn patent discloses “***dynamically modifiable ports*** on the storage controllers, ***such that those ports can operate*** either as a control unit link-level facility or ***as a channel link-level facility***.” (Hathorn patent at Abstract (emphasis added).) The Hathorn patent explains that when “configured as a channel link-level facility, a primary storage controller can appear as a host processor to a secondary storage controller.” (*Id.*)

49. More specifically, one of ordinary skill would understand that the “modifiable ports on the storage controllers” correspond to the ’346 patent’s claimed network interface controlling units. When the network interface controlling units are operating as a “control unit link-level facility,” they exchange information with the hosts. When the network interface controlling units are

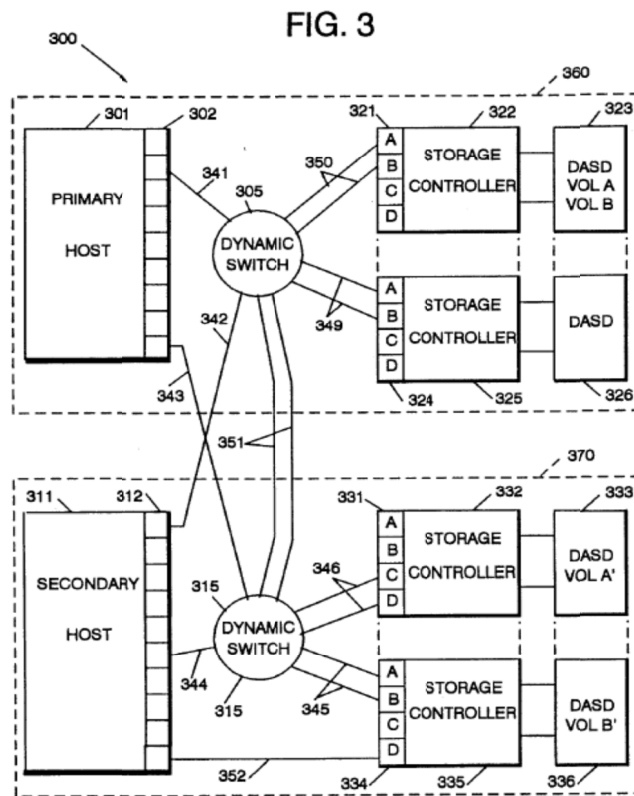
operating as a “channel link-level facility,” they exchange information with network interface controlling units on other storage controllers, e.g., RAID controllers.

50. It is useful to compare Hathorn’s Figures 2 and 3. Figure 2 of the Hathorn patent (reproduced below) discloses *direct paths* (labeled as lines 247 and 248)—*not utilizing the switch network*—between network interface controlling unit ports in order to facilitate the exchange of information.



51. The Hathorn patent teaches that these extra communication paths (247 and 248) are expensive. (*Id.* at 4:1-11.) Therefore, the Hathorn patent teaches that an improvement over Figure 2’s “remote dual copy system 200” involves reducing

the number of required communication links and their associated expense. (*Id.* at 7:45-47.) In order to do this, the Hathorn patent refers to Figure 3 and teaches that “communication links 247 and 248 are reduced to and replaced by communication links 351 ... [which] are connected between dynamic switches 305 and 315.” (*Id.* at 7:59-63.)



52. The Hathorn patent teaches that the reduction of communication links between storage controllers “*is enabled by modification of the storage controller ports* or link-level facilities into dual function link-level facilities ... [which allows] the primary and secondary storage controller ports 321, 324, 331, and 334 [to] be dynamically set to communicate either as a channel or control unit link-

level facility.” (*Id.* at 8:1-6 (emphasis added); *see also id.* at 10:41-45.) Use of a “channel link-level facility” allows the network interface controlling unit ports on two different RAID controllers to exchange information. (*Id.* at 5:8-15.)

53. The Hathorn patent teaches that storage controller ports (*i.e.*, network interface controlling units) can be modified to exchange information by “[e]stablishing logical paths between storage controllers ... with a combination of an Establish Logical Path (ELP) link-level frame and a device level control frame for indicating that the logical path supports peer-to-peer protocols.” (*Id.* at 10:54-58.) This modification allows a primary storage controller to “perform any functions on a peer-to-peer logical path that a channel is allowed to perform....” (*Id.* at 10:64-67.) The Hathorn patent also teaches that the modified storage controller ports, in the process of establishing the peer-to-peer protocol, will exchange information with each other. (*See, e.g., id.* at 11:25-43 (“The primary storage controller 325, acting as host with the ports 324 enabled as channel link-level facility, sends an EPC frame to the secondary storage controller 335 ... the secondary storage controller 335 processes the EPC frame and returns an acknowledgement (ACK) frame.”).)

54. The Hathorn patent discloses this method of avoiding the need for expensive direct paths between storage controllers as “using ***shared communication links*** that can dynamically interface either a host processor to a

storage controller, or can interface one storage controller to another storage controller....” (*Id.* at 4:31-37 (emphasis added).)

55. As a comparison, the ’346 patent includes a Figure 2 labeled “prior art” which identifies two “communication controllers” 221 and 222 that communicate directly, bypassing the “hub or switch” 210—*similar to the system disclosed in the Hathorn patent’s Figure 2*. The ’346 patent’s Figure 4 is described as an improvement over the system illustrated in its Fig. 2, where inter-controller communications use the switch network and avoid direct paths between controllers—*similar to the system disclosed in the Hathorn patent’s Figure 3*.

C. The Mylex Paper In View Of The Teachings Of The Hathorn Patent Compared To The ’346 Patent, Claims 1-9

1. Claim 1

1a) An apparatus for a redundant interconnection between multiple hosts and a RAID, comprising:

56. The Mylex paper discloses this claim element. For example, the Mylex paper discloses connections between servers and storage devices in stating that a “Storage Area Network can be designed with a specialized or standard networking technology, e.g., Fibre Channel. Its purpose is to provide high bandwidth connections between servers and storage devices, and between storage devices, e.g., storage arrays and tape libraries.” (Mylex paper at 4.)

57. The servers disclosed by the Mylex paper function as hosts, and the Mylex paper uses both the terms “servers” and “hosts” to refer to structures having

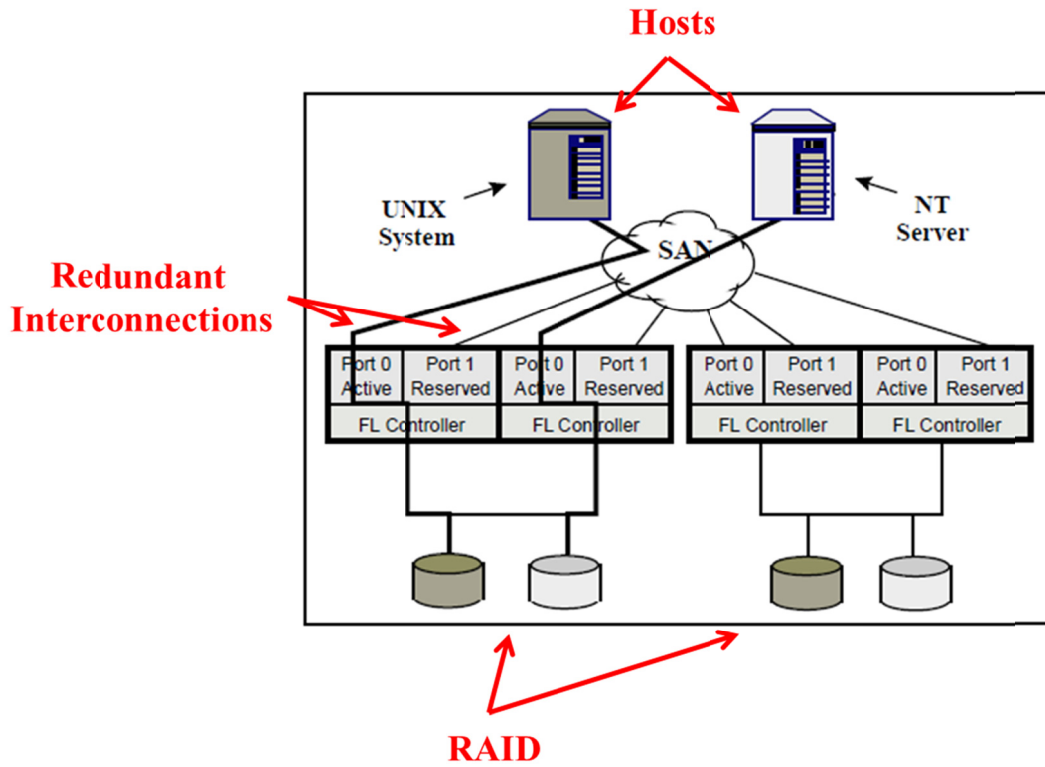
the same function. For example, the Mylex paper discloses that “dual host ports are particular critical for controller failover in Fibre Channel topologies. The SAN ports can be connected directly to UNIX and NT servers or indirectly through hubs and switches.” (*Id.* at 16; *see also id.* at 11.)

58. The Mylex paper further discloses that a storage device can be a RAID. For example, the Mylex paper states that “[t]he Fibre Channel standard is widely supported and a broad range of Fibre Channel interconnect devices (hubs and switches) and storage devices (RAID arrays, tape and optical libraries, and disk, tape, and optical drives) will be available....” (*Id.* at 11; *see also id.* at 12.)


59. Furthermore, the Mylex paper discloses that the connections between servers and storage devices are redundant interconnections. For instance, the Mylex paper discloses that “[e]ach controller has redundant paths to host systems and pairs of controllers provide redundant paths to disks.” (*Id.* at 15; *see also id.* at 11, 16.)

60. The Mylex paper’s Figures 12 and 17 (reproduced and annotated below) illustrate the redundant interconnections between multiple servers and a RAID.

Annotated Figure 12




Annotated Figure 17



1b) a first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers

61. The Mylex paper discloses this element. For example, see the components labeled “Controller 0” and “Controller 1” in annotated Mylex Figure 17 below.

Annotated Figure 17



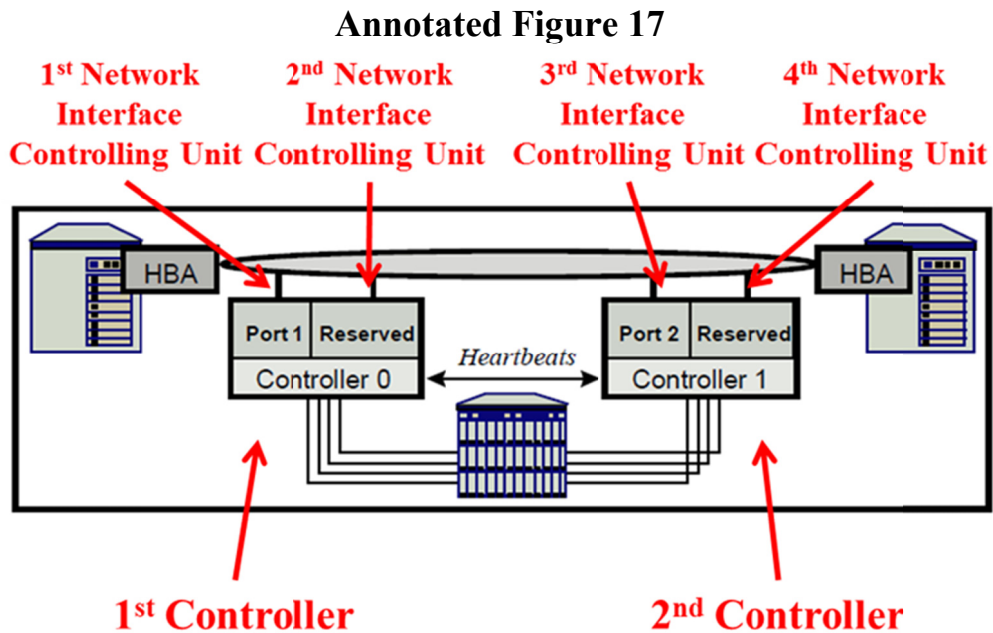
62. The Mylex paper discloses the use of duplex RAID controllers for fault tolerance when stating, for example, that “SAN-attached arrays should be configured with duplex controllers and with the disks connected to both controllers. Multiple SAN interfaces (on each controller) and duplex controllers with shared disks provide the level of fault tolerance required in SAN configurations.” (*Id.* at 11; *see also id.* at 15.)

63. The Mylex paper further discloses that the duplex RAID controllers are used for processing a requirement of host computers. For example, the Mylex paper discloses that “Mylex controllers support active-active operation; both controllers simultaneously satisfy I/O requests from SAN nodes. Some vendors offer active-passive controllers which is similar in concept to a hot spare disk. The passive controller waits for the active controller to fail and then assumes the I/O load of the failed controller. With active-active, both controllers service I/O requests and hence, deliver up to twice the performance of active-passive controllers.” (*Id.* at 16; *see also id.* at 11.)

*1c) the first RAID controlling unit including a **first network controlling unit** and a **second network controlling unit** and the second RAID controlling unit including a **third network controlling unit** and a **fourth network controlling unit***

64. The Mylex paper discloses a system having RAID controlling units that include first, second, third and fourth network interface controlling units.

65. For example, as shown in the annotated version of Figure 17 below, the Mylex system's "Controller 0" and "Controller 1" function as a first and second controller respectively, where the controllers can specifically be RAID controllers. Port 1, Port 2 and the "Reserved" ports on each controller interface with the hosts as indicated by the horizontal lines between the dual-ported HBAs (host bus adapters). (Note that in other Mylex figures, ports are numbered as Port 0 and Port 1.) One of ordinary skill in the art would understand that the active ports and the "Reserved" ports on each RAID controller are network interface controlling units.



1d) a plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers

66. The Mylex paper discloses this element by describing numerous connection units for connecting the first and second controllers to the numerous servers.

67. Specifically, the Mylex paper discloses that “[SAN] Ports can be directly attached to SAN servers or indirectly through hubs and switches.” (*Id.* at 19; *see also, e.g., id.* at 4 (“Storage Area Networks can be configured in fabric topologies with switches to interconnect servers and devices or implemented in loop topologies with hubs to simplify cable management and increase loop resiliency.”).)

68. The Mylex paper shows this illustratively, for example, at FIG. 12 (reproduced below), where multiple servers are connected to controller ports through a SAN network. As disclosed by the Mylex paper, “[s]witches, hubs and routers are interconnect devices that can be employed to construct SAN networks.” (*Id.* at 5; *see also id.* at 8 (describing hubs and switches), 11.)

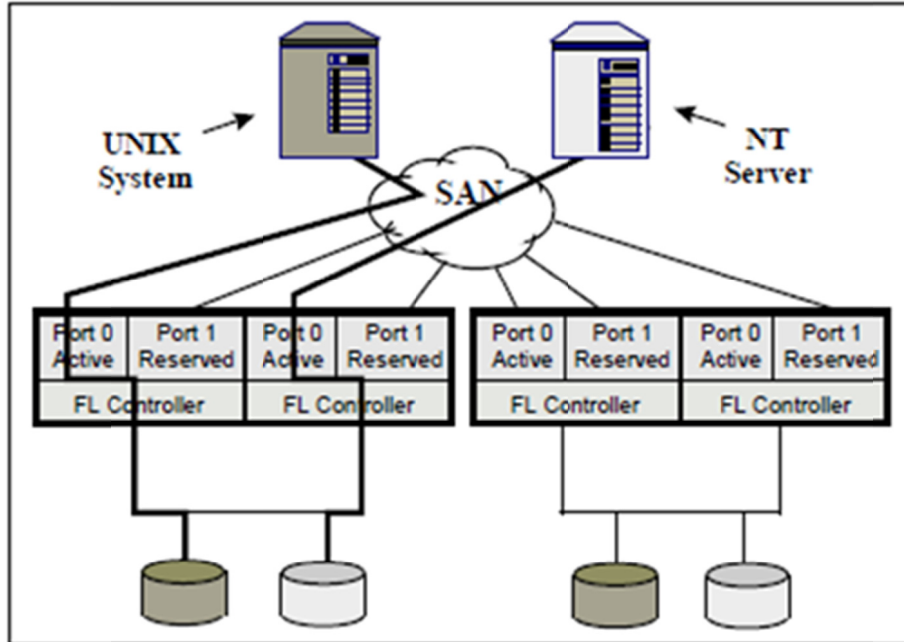


Figure 12. DAC FL Attached to a Fibre Channel SAN

69. The Mylex paper additionally discloses cascading hubs and switches to create a network. For example, the Mylex paper discloses that “[h]ubs and switches are interconnect devices used in Storage Area Networks...Hubs and switches can be cascaded to increase node connectivity.” (*Id.* at 8.) The Mylex paper shows cascading of hubs and switches at FIG. 6 (reproduced below):




Figure 6. SAN With Switched and Shared (Loop) Interconnects

70. Additionally, while the Mylex paper's Figure 17, reproduced above, does not explicitly label the hubs and switches used to connect the hosts and RAID controllers, one of ordinary skill would understand that hubs and/or switches could be used in Mylex's Figure 17. (*Compare, e.g.,* Mylex Figure 17 (disclosing RAID controllers with dual-SAN ports) with Mylex at 19 ("Since Mylex controller have dual SAN ports for increased performance and resiliency in the face of interconnect failures, they are well suited for SAN applications. ***Ports can be directly attached to SAN servers or indirectly through hubs and switches.***") (emphasis added).)

71. Therefore, the Mylex paper discloses a plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers," as claimed by the '346 patent.

72. Additionally, to the extent that the Mylex paper is not found to disclose this claim element, this element would have been obvious to one of ordinary skill in the art because using multiple switches and hubs in RAID systems was well-known in the 2000 time frame, in particular to achieve the redundancy directed by the Mylex paper. For example, Figures 2 and 3 of the Hathorn patent teach using more than one switch in a RAID system to permit hosts to interact with multiple RAID controllers in a redundant fashion. As such, applying the multiple-switch teachings of the Hathorn patent to the system disclosed in the Mylex paper,


which one of skill in the art would have been motivated to do to achieve the desired redundancy would render this claim element obvious.

*1e) 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*

73. As discussed above, the Mylex paper discloses this feature when describing first and second controllers exchanging information with servers through a plurality of switches and hubs. (*See, e.g.*, Mylex paper at 16 (“Mylex controllers have dual SAN ports which doubles the bandwidth to controllers and allows redundant paths from other SAN devices to the controllers to increase the resiliency of the SAN topology... The SAN ports can be connected directly to UNIX and NT servers or indirectly through hubs and switches.”).)

74. This exchange of information between RAID controllers and hosts is illustrated in the annotated version below of the Mylex paper Figure 17. The oval in this figure represents the SAN as referenced in the heading “Clusters use SAN Technologies” with a bullet items “• I/O bandwidth is scaleable (with switches); storage capacity is scaleable, • Redundant links can be used for fault tolerance and higher data availability” (*Id.* at 9) disclosing the connecting units of Claim 1:

Annotated Figure 17



75. The Mylex paper discloses that an exchange of information can include input/output (I/O) operations, including read and write operations, performed by one or more servers and/or storage devices. For example, the Mylex paper discloses that “Mylex external RAID array controllers use sophisticated caching algorithms for both read and write operations.” (*Id.* at 20; *see also id.* at 16.)

1f) and 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

76. The Mylex paper in view of the teachings of the Hathorn patent renders this claim element obvious.

77. The Mylex paper, as shown in the annotated version of Mylex Figure 17 above, discloses a “heartbeat path” between two RAID controllers (Controller 0 and 1). The Mylex paper discloses that the controllers use this path to provide high availability and transmit/receive “I’m alive heartbeat messages.” (*Id.* at 16.) The Mylex paper discloses the process as follows:

The absence of heartbeat messages signals that one of the controllers is off-line and the remaining controller immediately initiates a failover operation and then begins servicing I/O requests directed to itself and its off-line partner to provide non-stop access to data. Controller failover / failback operations are host independent and transparent to SAN nodes. During the failover / failback process, SAN nodes simply continue sending I/O requests to the same ID’s across the SAN interconnect. As far as the nodes are concerned, these commands are processed identically whether both controllers are functional or one has failed. (*Id.* at 16.)

78. The Mylex paper further discloses that the RAID controllers exchange information using the heartbeat path for fault tolerance:

If a controller fails, the surviving controller senses the absence of heartbeats, fails over the ID of the active port on the failed controller to its reserved port, and updates its data structures with configuration information stored on disk. The failover process is transparent since the nodes still see the same fibre port ID’s on the SAN interconnect. ... When the failed controller is replaced, it is detected by the surviving controller which allows it to restart and returns the failed controller’s port ID’s, and then it starts processing I/O. (*Id.* at 18.)

79. However, as discussed above (see above, ¶¶48-55), the Hathorn patent teaches reducing the expense associated with direct communication paths between RAID controllers by replacing them with modified storage ports (i.e., network interface controlling units) that act as “channel link facilities,” sending and receiving information between ports using an existing switch network. For example, the Hathorn patent teaches:

dynamically modif[ying] ports on the storage controllers such that those ports can operate either as a control unit link-level facility or as a channel link-level facility. When configured as a channel link-level facility, a primary storage controller can appear as a host processor to a secondary storage controller. Using dynamic switches coupled between primary and secondary sites, fewer ESCON communication links are required since the ESCON communication links can function either as a channel or as Storage controller Communication link. (Hathorn patent at Abstract.)

80. Applying these teachings to the system disclosed by the Mylex paper, one of ordinary skill in the art would have found it obvious to modify the system disclosed in the Mylex paper so that the expense of the heartbeat path is avoided by modifying the network interface controlling units to exchange information with each other, as claimed by the '346 patent. An annotated version below of Figure 17 of the Mylex paper illustrates this configuration:

**Annotated Figure 17
(Mylex System in view of Hathorn Teaching
to Modify Network Controlling Units To Communicate)**




Figure 17. Normal Operating Mode of Dual DAC SF and FL Fibre Controllers

network
controlling
unit 1

network
controlling
unit 2

network
controlling
unit 3

network
controlling
unit 4

2. Claim 2

*2a) The apparatus as recited in claim 1, wherein said respective **RAID controlling units** are connected to the plurality of individual connecting units*

81. The Mylex paper discloses controllers connected to hubs and switches for communication with one or more servers, and thus discloses this claim element.

82. For example, the Mylex paper discloses that “SAN ports can be connected directly to UNIX and NT servers or indirectly through hubs and switches.” (Mylex paper at 16; *see also id.* at 4, 5, 19.)

83. Additionally, the Mylex paper at Figures 6, 20, and 21 (reproduced below) shows controllers connected to hubs or switches, alone or in a cascade, via the ports of the controllers. As discussed above with respect to claim element 1(b), the controllers disclosed by the Mylex paper are RAID controllers.




Figure 6. SAN With Switched and Shared (Loop) Interconnects




Figure 20. Mylex External Array Controllers in a Loop SAN Topology




Figure 21. Mylex External Array Controllers in a Switched SAN Topology

84. Therefore, the Mylex paper discloses that said respective RAID controlling units are connected to the plurality of individual connecting units. Further, to the extent that the Mylex paper is not found to disclose this claim element, this element would have been obvious to one of ordinary skill in the art because using multiple switches and hubs in RAID systems was well-known in the 2000 time frame, in particular to achieve the redundancy directed by the Mylex paper. For example, Figures 2 and 3 of the Hathorn patent teach using more than one switch in a RAID system to permit hosts to interact with multiple RAID controllers in a redundant fashion. As such, applying the multiple-switch teachings of the Hathorn patent to the system disclosed in the Mylex paper, which one of skill in the art would have been motivated to do to achieve the desired redundancy would render this claim element obvious.

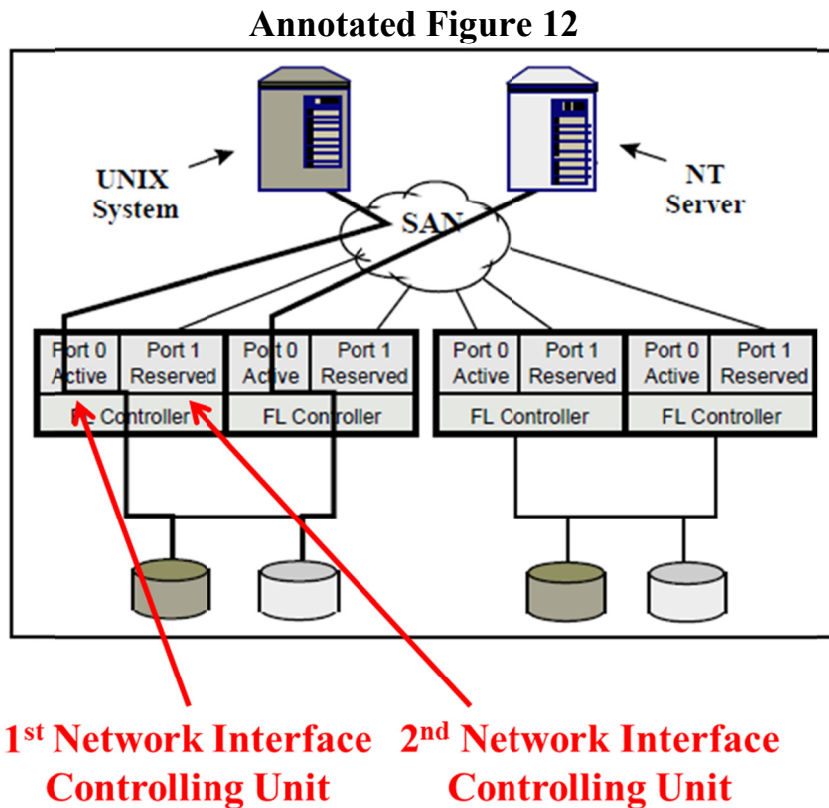
3. Claim 3

*3a) 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***

85. It is my opinion that the Mylex paper in view of the teachings of the Hathorn patent renders this claim element obvious. However, I note that in describing a real world system, the language requiring a “connecting unit *of one side*” is not specific. I also note that the ’346 patent does not use the term “side” except within Claim 3. However, for purposes of this discussion, and under the broadest reasonable interpretation standard, I interpret the “sides” being referred to as the sides of the Figures included in the ’346 patent, and as such, I refer below to the sides of the Figures in the Mylex paper and the Hathorn patent in my analysis of this claim element.

86. As described with respect to claim element 1(d), the Mylex paper discloses a SAN network that consists of hubs, switches, and/or a cascade of hubs and switches. As such, the Mylex paper’s Figure 12 (an annotated version of which appears below) illustrates a SAN network in which the first network interface controlling unit is coupled to one section of the SAN network ports, which, e.g., may be coupled to the Unix system on the left side of the figure, and the second network interface controlling unit is coupled to a different section of the SAN network ports, which, e.g., may be coupled to the NT server on the right side

of the figure. As such, in Figure 12, the first network interface controller is connected to at least the left side host via a left side hub/switch indicated by the line through the SAN and the second network interface controller is connected to at least the right side host via a right side hub/switch indicated by the right line through the SAN.

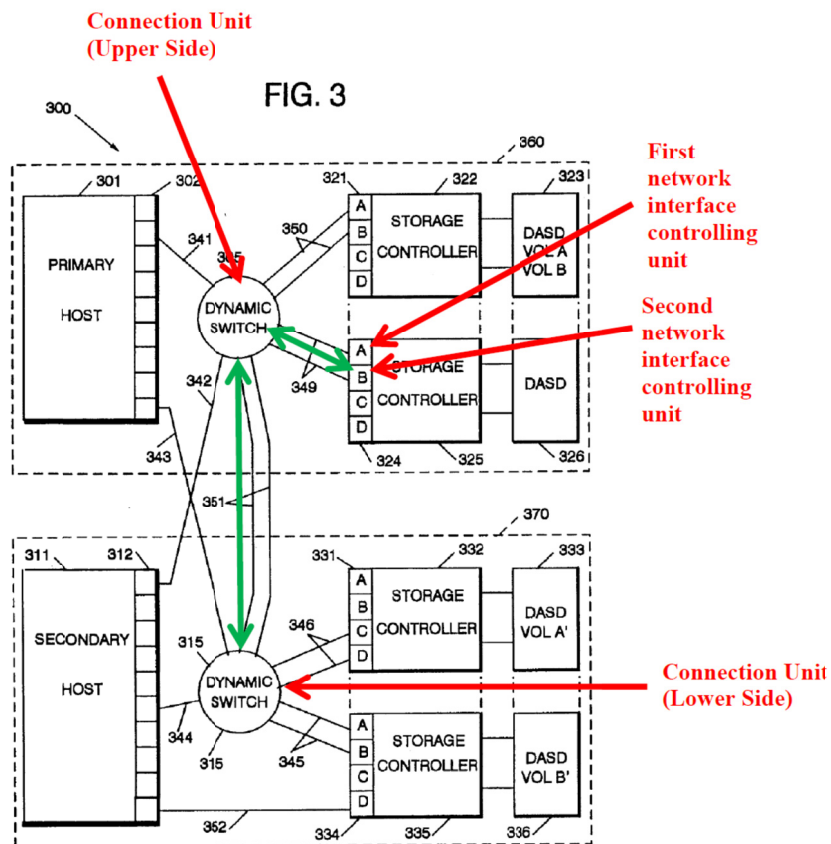


87. Additionally, the Mylex paper’s Figure 17, reproduced above in claim [1e], illustrates that all network interface controllers are connected to the hubs on the left and right sides of Figure 17.

88. Additionally, the Hathorn patent includes a Figure 3 (an annotated version of which appears below) that teaches that first and second network

interface controlling unit ports are coupled to connecting units of two sides (upper and lower side of Figure 3) and a second network interface controlling unit can be coupled to a connecting unit of a second side. As shown below, the first network interface controlling unit (324A) is connected to connecting unit 305 on the upper side of Figure 3, via path 349, and network interface controlling unit (324B) is connected to connecting unit 315 on the lower side of Figure 3, via path 349, connecting unit 305, and path 351:

Annotated Hathorn Patent Figure 3



4. Claim 4

*4a) The apparatus as recited in claim 1, wherein the **first** network interface controlling unit and the **third** network interface controlling unit **process the requirement of the numerous host computers***

89. The Mylex paper discloses this claim element. As described with respect to claim element 1(e), requirements of host computers can include I/O requests, e.g., read and write requests, processed by controllers and/or servers.

90. The Mylex paper discloses a redundant system where the first and third network interface controlling units are active and process I/O requests of the servers. (*Id.* at 18 (“fibre arrays are configured with one port active and the other reserved”); *id.* at 16 (“Mylex controllers support active-active operation; both controllers simultaneously satisfy I/O requests from SAN nodes.”).) This is illustrated by the annotated version of Mylex Figure 17 below.

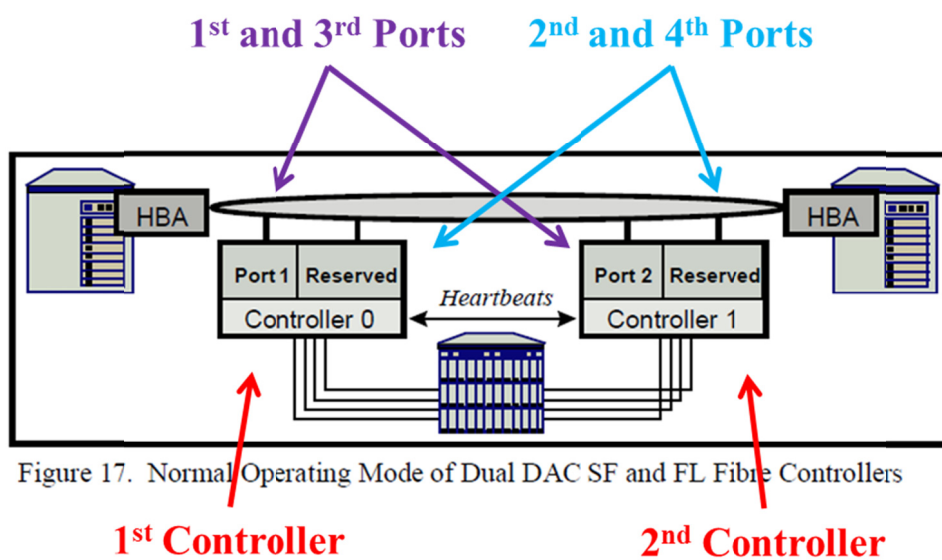


Figure 17. Normal Operating Mode of Dual DAC SF and FL Fibre Controllers

*4b) the **second** network interface controlling unit and the **fourth** network interface 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 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***

91. The Mylex paper in view of the teachings of the Hathorn patent renders this claim element obvious.

92. The Mylex paper discloses the use of duplex RAID controllers for fault tolerance when stating, for example, that “SAN-attached arrays should be configured with duplex controllers and with the disks connected to both controllers. Multiple SAN interfaces (on each controller) and duplex controllers with shared disks provide the level of fault tolerance required in SAN configurations.” (*Id.* at 11; *see also id.* at 12, 15.)

93. The Mylex paper discloses that the RAID controllers can exchange heartbeat signals during normal operation and use the reserved network interface controlling unit ports for processing host requests when the alternate RAID controller experiences a fault. (*Id.* at 18 (“If a controller fails, the surviving controller senses the absence of heartbeats, fails over the ID of the active port on the failed controller to its reserved port”).) Figures 17 and 18 of the Mylex paper (annotated and reproduced below) illustrate this failover functionality.

Requests with Faulty Controller 1 (Green and Red)

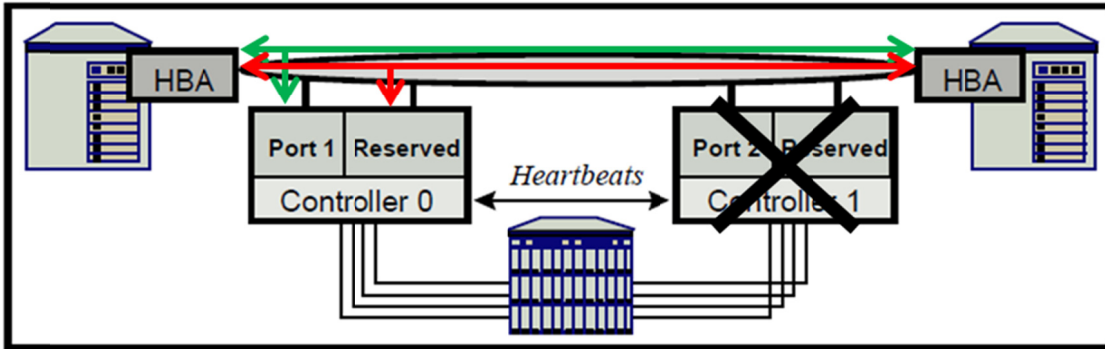


Figure 17. Normal Operating Mode of Dual DAC SF and FL Fibre Controllers

Requests with Faulty Controller 0 (Purple and Red)




Figure 17. Normal Operating Mode of Dual DAC SF and FL Fibre Controllers




Figure 18. DAC SF and FL Fibre Array Controller Failover

94. The Mylex paper at Figure 19 illustrates the return process: “When the failed controller is replaced, it is detected by the surviving controller which

allows it to restart and returns the failed controller's port ID's, and then it starts processing I/O.” (*Id.* at 18.)




Figure 19. DAC SF and FL Fibre Array Controller Failback

95. Additionally, the Hathorn patent, as discussed above (see ¶¶48-55), teaches that storage controller ports can be modified to exchange information. As such, one of ordinary skill in the art would have found it obvious to configure the claimed element by avoiding use of the heartbeat paths disclosed in the Mylex paper, and instead modifying Mylex's storage controller ports to communicate for the recited fault tolerance function. Further, one of ordinary skill would have found it obvious to configure the Mylex paper's disclosed "reserved ports" (*i.e.*, the second and fourth network interface controlling units) for communication between the first RAID controlling unit and the second RAID controlling unit in a non-fault state, e.g., to effectuate the high-bandwidth goal disclosed in the Mylex paper (*see, e.g., id.* at 4 ("A [SAN] is a specialized network ... that provides both high bandwidth and low latency")), as the other (first and third active) network

interface controlling units are active and thus in use during a non-fault state. This is illustrated below in an annotated version of the Mylex paper's Figure 17:

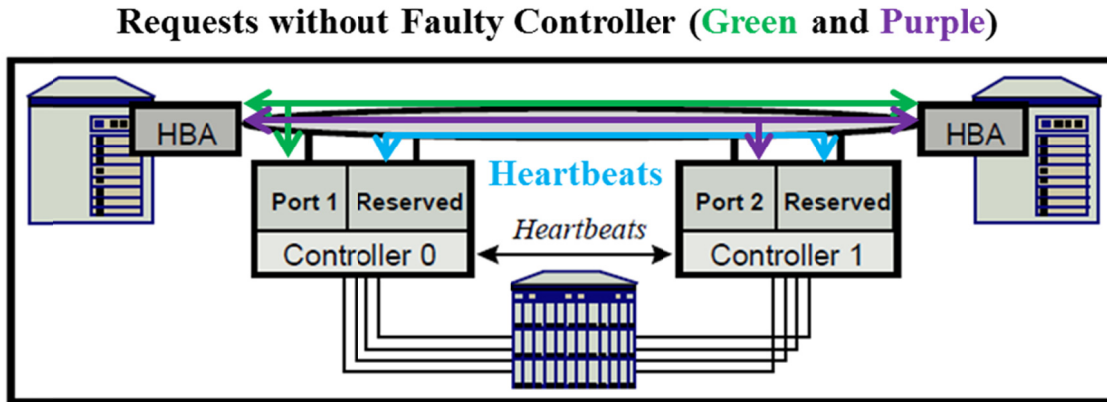


Figure 17. Normal Operating Mode of Dual DAC SF and FL Fibre Controllers

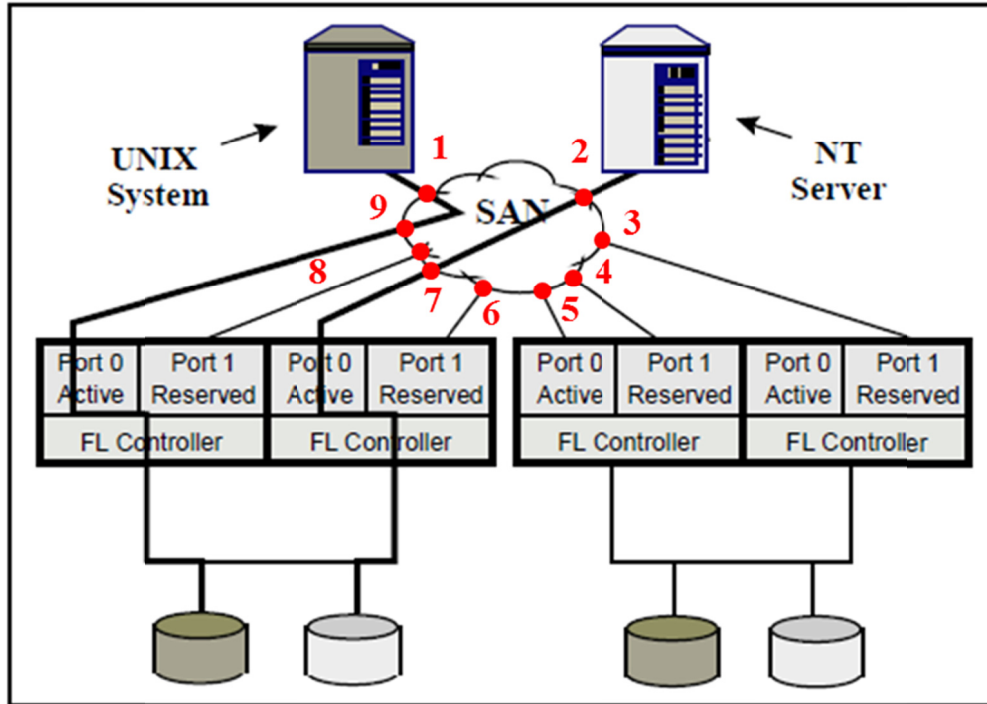
96. For this section, I am interpreting the claim language under the broadest reasonable interpretation standard as reciting "... the second network interface controlling unit ~~and~~ or the fourth network interface controlling unit are used for executing a function of the first network interface controlling unit ~~and~~ or the third network controlling unit when one of the first RAID controlling unit and the second RAID controlling unit is faulty." I use this interpretation because claim 1 requires that the second and fourth network interface controlling units are on different RAID controllers, and as such, if one RAID controller is faulty, either the second or the fourth network interface controlling unit will not be used to execute any function. I reserve the right to offer a different interpretation in the district court litigation, which I have been informed uses a different legal standard.

5. Claim 5

5a) The apparatus as recited in claim 1, wherein said plurality of connecting units have at least three connection ports

97. The Mylex paper discloses this claim element.

98. Specifically, as described at claim element 1(d), the switches and hubs, and/or cascaded switches and hubs, disclosed in the Mylex paper correspond to the claimed connecting units disclosed in the '346 patent. A SAN includes these hubs and switches. For example, the Mylex paper discloses that "SAN ports can be connected directly to UNIX and NT servers or indirectly through hubs and switches." (*Id.* at 16.) The Mylex paper at Figure 12 (reproduced below) shows such a SAN with at least nine connection ports connecting multiple servers and multiple controllers. "In Figure 12...controllers are SAN-attached. Each controller has redundant paths to host systems and pairs of controllers provide redundant paths to disks." (*Id.* at 15.)



99. Additionally, the Mylex paper's Figure 6 (reproduced below) discloses a cascaded hub-switch system with at least four connection ports on the switch.





Figure 6. SAN With Switched and Shared (Loop) Interconnects

100. Furthermore, the Mylex paper at Figures 20 and 21 discloses that hubs and switches can have up to twelve connection ports. One of ordinary skill would

understand that the system disclosed in Mylex's Figure 17 could be configured with any of the above SAN-connection configurations.

101. Additionally, to the extent this element is not explicitly disclosed in the Mylex paper, the Mylex paper in view of the teachings of the Hathorn patent renders this claim element obvious. For example, Figure 3 of the Hathorn patent, reproduced below, teaches using dynamic switches with at least eight ports. As such, applying the multi-ported, multi-connection unit teachings of the Hathorn patent to the system disclosed in the Mylex paper, which one of ordinary skill would have been motivated to do to achieve the desired redundancy, would render this claim element obvious.

FIG. 3



5b) *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*

102. The Mylex paper discloses that two of the connection ports of a hub or switch are connected to one of a first network interface controlling unit port of a first controller and a third network interface controlling unit port of a second controller, and therefore discloses this claim element. This is illustrated by the annotated version below of the Mylex paper's Figure 12 (where the SAN is made up of hubs, switches, or hub-switch cascade):

First RAID controller

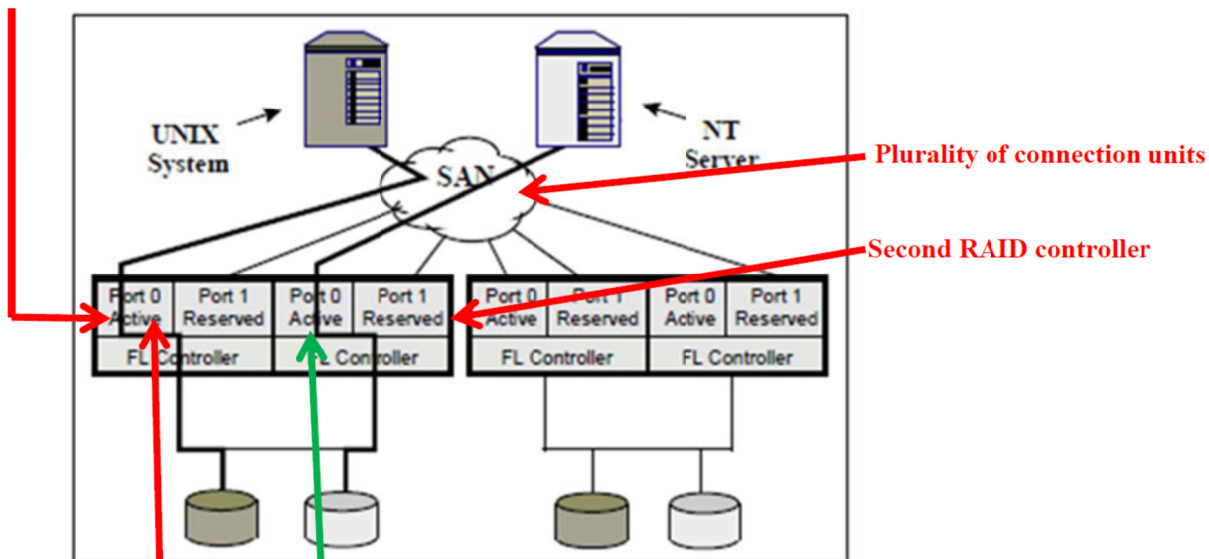


Figure 2. DAC FL Attached to a Fibre Channel SAN

First network interface controlling unit
Third network interface controlling unit

103. One of ordinary skill in the art would understand that in switches and hubs, every port is connected to every other port. For example, I note that the Mylex paper discloses that in hubs and switches, every port is connected to every other port. (See, e.g., *id.* at 8 (disclosing, in the caption of Figure 6, “SAN With Switched and Shared (Loop) Interconnects”); *id.* at Figure 17 (disclosing all nodes connected); see also, e.g., DeKoning patent, Exhibit VMWARE-1010, at 11:25-31 (“Specifically, a Fibre Channel Arbitrated Loop (FC-AL) connection provides a high speed communication medium in which any of the nodes on the multipoint loop may exchange information with any other node on the loop.”); “Designing Storage Area Networks” by Tom Clark (Exhibit VMWARE-1011), at 50

(disclosing that every port of a hub or switch are connected to every other port, for example, when disclosing, “Arbitrated Loop is a true physical loop...In this way, a continuous data path exist through all the NL_Ports, allowing any device to access any other device on the loop...”).) As such, the annotated version of Figure 12 above shows that *all* of the SAN connection ports are coupled to one of the first network interface controlling unit and the third network controlling unit.

5c) and the rest of the connection ports being provided as a hub equipment connected with the numerous host computers

104. The Mylex paper discloses this claim element. First, I note that this claim element appears to be inconsistent with claim 1, in that claim 5 as a whole only recites that the two ports on the connecting units are connected to the first and third network interface controlling units and the “rest” are connected to the hosts—as such, apparently no ports are connected to the second and fourth network controlling ports. However, as discussed in claim 5(b), one of ordinary skill would understand that in a switch or hub, every port is connected to every other port. Therefore, under the broadest reasonable interpretation standard, the Mylex paper discloses a system that meets this claim element because all ports on the connecting units are connected to every network interface controlling unit and the hosts.

105. Further, the Mylex paper at Figures 6 and 20 illustrates that the disclosed RAID system can use hubs.

6. Claim 6

6a) The apparatus as recited in claim 1, wherein said plurality of connecting units have at least three connection ports

106. This claim element is identical to the corresponding element in claim

5. Thus, as discussed above, the Mylex paper discloses this claim element.

6b) two of the at least three connection port are coupled to one of the first network controlling unit and the third network controlling unit

107. This claim element is identical to the corresponding element in claim

5. Thus, as discussed above, the Mylex paper discloses this claim element.

6c) and the rest of the connection ports being provided as a network switch equipment connected with the numerous host computers

108. This claim element is similar to the corresponding element in claim 5.

109. Further, the Mylex paper at Figures 6 and 21 illustrates that the disclosed RAID system can use switches.

110. Therefore, the Mylex paper discloses that the rest of the connection ports are provided as network switch equipment connected with the numerous host computers, as claimed by the '346 patent.

7. Claim 7

7a) The apparatus as recited in claim 1, wherein said plurality of connecting units have at least five connection ports

111. The Mylex paper discloses this claim element. For example, in discussing claim 5(a) above, I have shown how the Mylex paper discloses using

SAN connection units (e.g., hubs and switches) with at least twelve connection ports.

7b) 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

112. The Mylex paper discloses this claim element. For example, as discussed above with respect to claim 5(b), one of ordinary skill in the art would understand that all connection ports on a switch are coupled to each other as well as all devices on the SAN. For example, the Mylex paper discloses that “[i]n a FC-AL, nodes arbitrate to gain access to the loop and then pairs of nodes establish a logical point-to-point connection to exchange data; the other nodes on the loop act as repeaters.” (*Id.* at 7.)

113. As such, the twelve-ported switches in the Mylex paper would have at least five connecting ports coupled to the first and third network controller units, as all connecting ports are coupled to all network controller units. This is illustrated by Figure 21 of the Mylex paper, reproduced below: [Note the dark and light lines to the Switch box from ports 0 and 1 of each RAID controller. To one of ordinary skill in the art, the different line types suggest that the Switch may be implemented with separate port 0 switches and port 1 switches.]




Figure 21. Mylex External Array Controllers in a Switched SAN Topology

7c) and the **rest of the connection ports** being provided as a **switch** connected with the numerous host computers

114. This claim element is similar to claim 6(c) discussed above. For the same reasons, the Mylex paper discloses this claim element.

8. Claim 8

8a) *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***

115. I first note that the '346 patent broadly uses the term "connected" to encompass indirect connections, e.g., through a hub or switch. (See, e.g., '346 patent at 4:33-34 ("the host computers 500, 501, 502, 503, 504 and 505 are connected to the RAID 530 by using external hubs 510 and 520").) As such, under

the broadest reasonable interpretation standard, the Mylex paper discloses that the first and second ports of a first controller are connected to each of two different switches or hubs, and that the first and second ports of a second controller are also connected to each of the two different switches or hubs.

116. Specifically, Figure 17 of the Mylex paper discloses a first controller (“Controller 0”) with two ports (“Port 1” and “Reserved”) and a second controller (“Controller 1”) with two ports (“Port 2” and “Reserved”). Each of these ports is connected to multiple servers over a SAN. Further, Figure 6 of the Mylex paper discloses that the SAN can be connected by a plurality of hubs and switches.

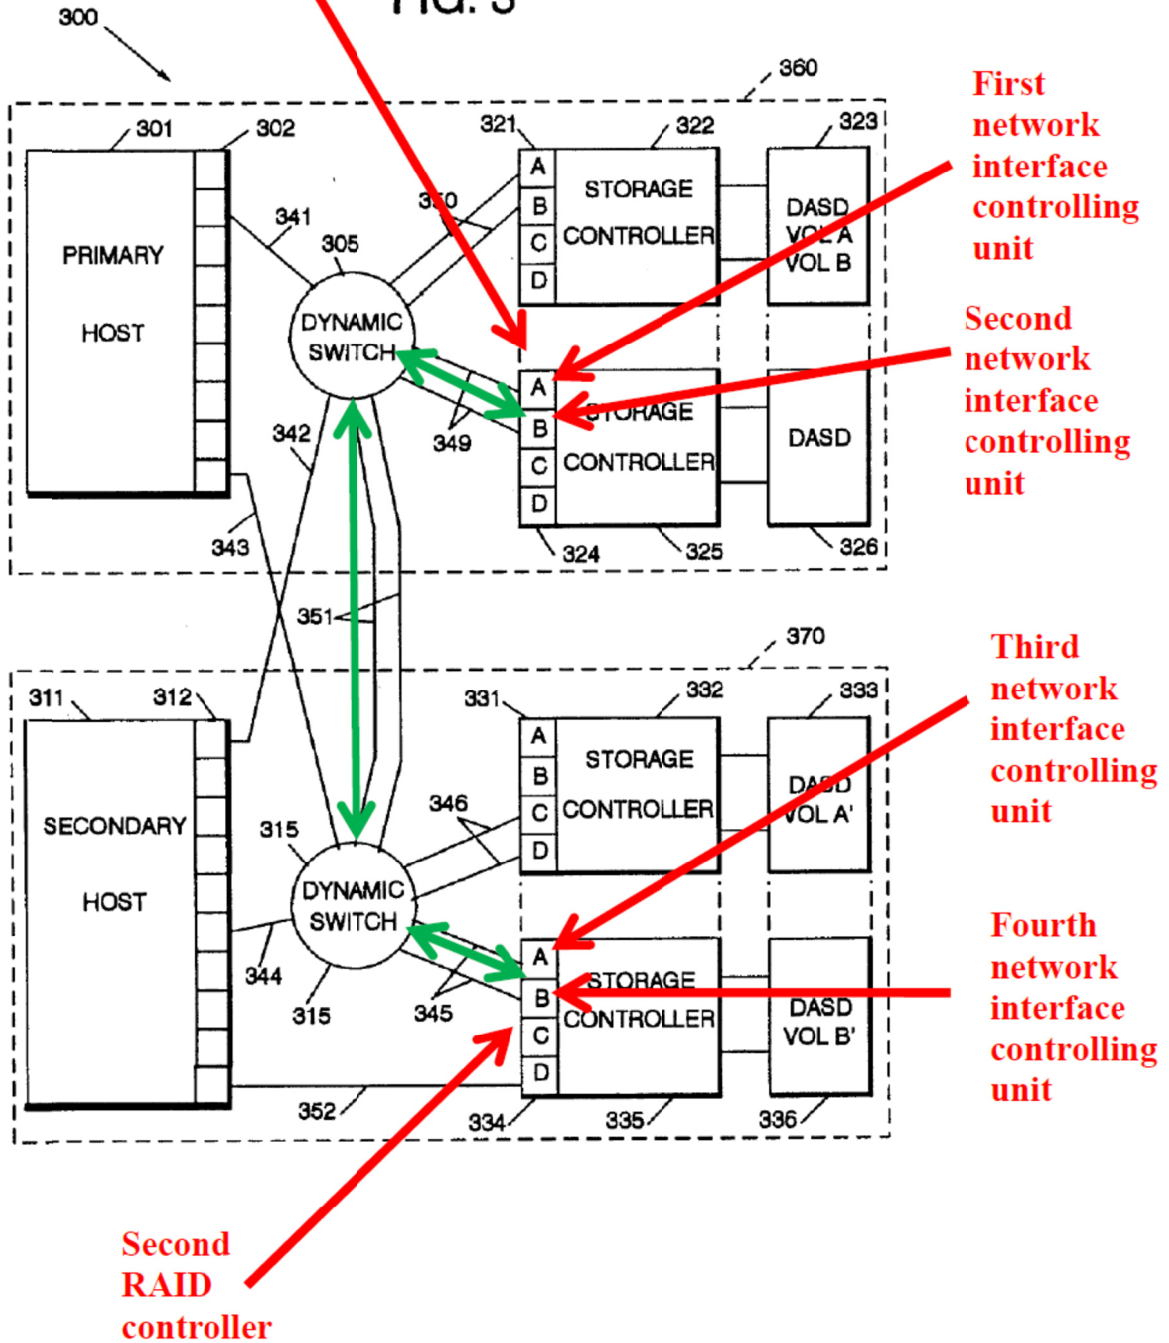
117. As such, the Mylex paper teaches that two ports for each of two controllers can be connected to a plurality of hubs and switches. Since the ports of a hub or switch are coupled to one another, including to the ports of other hubs or switches in a cascade, the Mylex paper teaches that (i) a first port of a first controller has a connection to a first hub or switch, (ii) a second port of the first controller has a connection to a second hub or switch, (iii) a first port of a second controller has a connection to the first hub or switch, and (iv) a second port of the second controller has a connection to the second hub or switch.

118. Additionally, the Mylex paper in combination with the teachings of the Hathorn patent renders this claim element obvious. For example, as illustrated by the annotated version below of the Hathorn patent’s Figure 3, (i) the first

network interface controlling unit is connected to a first connecting unit, switch 305, through connection path 349; (ii) the second network interface controlling unit is connected to a second connecting unit, switch 315, through connection path 349, switch 305, and connection path 351; (iii) the third network interface controlling unit is connected to the second connecting unit, switch 315, through connection path 345; and (iv) the fourth network interface controlling unit is connected to the first connecting unit, switch 305, through connection path 345, switch 315, and connection path 351. One of ordinary skill in the art would have been motivated to use this configuration to achieve the redundancy goals set forth in the Mylex paper, and to avoid the expensive direct connection for “heartbeats” described in that paper.

First RAID controller

FIG. 3



9. Claim 9

*9a) An apparatus for a **redundant interconnection** between **multiple host** computers and a **RAID**, the apparatus comprising*

119. This claim element is identical to the corresponding element in claim 1. Thus, as discussed above, the Mylex paper discloses “an apparatus for a redundant interconnection between multiple hosts and a RAID.”

*9b) a **plurality of connecting units** for connecting the host computers and the RAID*

120. This claim element is similar to an element in claim 1 that recites: “a plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers.” Thus, as discussed above, the Mylex paper discloses “a plurality of connecting units for connecting the host computers and the RAID.” Further, to the extent that the Mylex paper is not found to disclose this claim element, this element would have been obvious to one of ordinary skill in the art because using multiple switches and hubs in RAID systems was well-known in the 2000 time frame, in particular to achieve the redundancy directed by the Mylex paper. For example, Figures 2 and 3 of the Hathorn patent teach using more than one switch in a RAID system to permit hosts to interact with multiple RAID controllers in a redundant fashion. As such, applying the multiple-switch teachings of the Hathorn patent to the system disclosed in the Mylex paper, which one of skill in the art would have been

motivated to do to achieve the desired redundancy would render this claim element obvious

9c) a first and a second RAID controllers, included in the RAID, 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

121. This claim element is similar to language in claim 1 that recites: “a first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers, the first RAID controlling unit including a first network interface controlling unit and a second network interface controlling unit, and the second RAID controlling unit including a third network interface controlling unit and a fourth network interface controlling unit; and a plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers.”

122. Thus, as discussed above, the Mylex paper discloses first and second RAID controllers, included in the RAID, each of which has a first network interface controlling unit and a second network interface controlling unit for processing requests from the plurality of the host computers connected through the plurality of connection units. Figure 6 of the Mylex paper, reproduced below, shows that the Mylex controllers can be used with a SAN consisting of a plurality of connection units, and the annotated versions of the Mylex paper’s Figures 12

and 17 below show the claimed RAID controllers, network interface controllers, and processing of host requests by the first network interface controllers:




Figure 6. SAN With Switched and Shared (Loop) Interconnects

First RAID controller




Figure 12. DAC FL Attached to a Fibre Channel SAN

1st Network Interface Controlling Units

2nd Network Interface Controlling Units

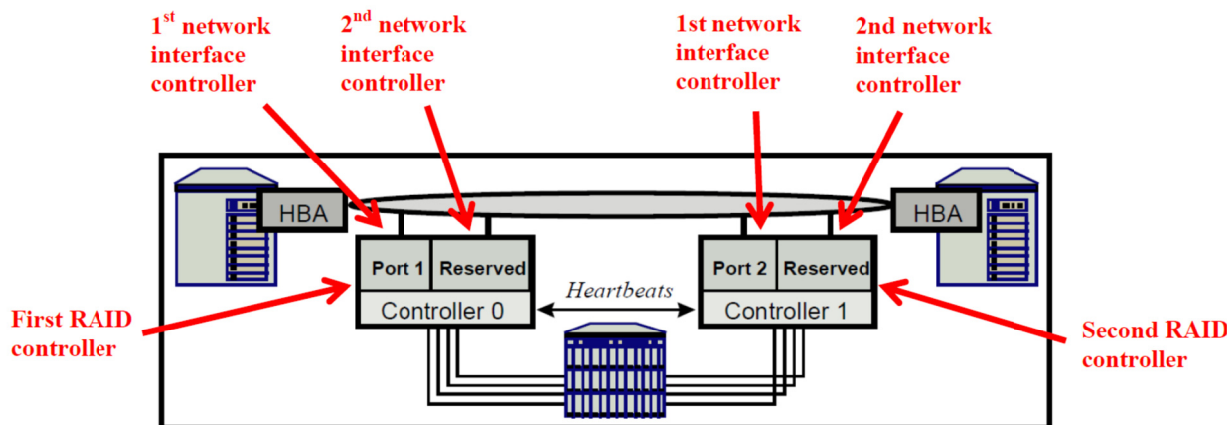


Figure 17. Normal Operating Mode of Dual DAC SF and FL Fibre Controllers

9d) 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 processes information transmitted from the second network interface controller** in the second RAID controller

123. The Mylex paper in combination with the teachings of the Hathorn patent renders this claim element obvious.

124. As explained above with respect to claim element 9(c), the Mylex paper discloses that the first network interface controllers (labeled as “active” in the annotated version of the Mylex paper’s Figure 12 above) supply data to the host computers in a non-faulty mode.

125. The Mylex paper, for example at Figure 17, discloses a direct “heartbeat path” between two RAID controllers (Controller 0 and 1). The Mylex paper discloses that the RAID controllers use this path to provide high availability and transmit/receive “I’m alive heartbeat messages.” (Mylex paper at 16.) The Mylex paper discloses this process as follows:

The absence of heartbeat messages signals that one of the controllers is off-line and the remaining controller immediately initiates a failover operation and then begins servicing I/O requests directed to itself and its off-line partner to provide non-stop access to data. Controller failover / failback operations are host independent and transparent to SAN nodes. During the failover / failback process, SAN nodes simply continue sending I/O requests to the same ID's across the SAN interconnect. As far as the nodes are concerned, these commands are processed identically whether both controllers are functional or one has failed. (*Id.* at 16.)

126. The Mylex paper further discloses that the RAID controllers exchange information using the heartbeat path to achieve fault tolerance:

If a controller fails, the surviving controller senses the absence of heartbeats, fails over the ID of the active port on the failed controller to its reserved port, and updates its data structures with configuration information stored on disk. The failover process is transparent since the nodes still see the same fibre port ID's on the SAN interconnect. ... When the failed controller is replaced, it is detected by the surviving controller which allows it to restart and returns the failed controller's port ID's, and then it starts processing I/O. (*Id.* at 18.)

127. Additionally, as discussed above (see ¶¶48-55), the Hathorn patent teaches reducing the expense of direct communication paths between RAID controllers by replacing them with modified storage ports (*i.e.*, network interface controllers) that act as “channel link facilities” for sending and receiving information between the ports using an existing switch network. For example, the Hathorn patent teaches:

dynamically modif[ying] ports on the storage controllers such that those ports can operate either as a control unit link-level facility or as a channel link-level facility. When configured as a channel link-level facility, a primary storage controller can appear as a host processor to a secondary storage controller. Using dynamic switches coupled between primary and secondary sites, fewer ESCON communication links are required since the ESCON communication links can function either as a channel or as storage controller communication link. (Hathorn patent at Abstract.)

128. The Hathorn patent also teaches that when the ports are modified to act in channel link-level facility mode, two RAID controllers, through their respective storage ports, can transmit “an EPC frame” and receive an “acknowledgement (ACK) frame” (at least). (*Id.* at 11:25-44.)

129. Applying these teachings of Hathorn to the system disclosed in the Mylex paper, one of ordinary skill in the art would have found it obvious to modify the system disclosed in the Mylex paper so that the first network interface controller in the first RAID controller would supply data to the host computers, connected through the plurality of connection units, and process heartbeat communications, e.g., the ACK frame, transmitted from the second network interface controller in the second RAID controller.

*9e) 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 processes information transmitted from the second network interface controller** in the first RAID controller*

130. For similar reasons as those described with respect to claim element 9(d), the combination of the Mylex paper and the Hathorn patent disclose that “the first network interface controller in the second RAID controller supplies data to the host computers connected through the plurality of connection units and processes information transmitted from the second network interface controller in the first RAID controller.”

*9f) 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*

131. The Mylex paper discloses this claim element.

132. For example, the Mylex paper discloses that a second port in a first RAID controller is used for fault tolerance by performing functions of a first port of a second RAID controller when the second controller experiences a fault.

133. For example, the Mylex paper discloses that SANs are capable of performing failover and failback operations for fault tolerance. The Mylex paper discloses, for example, that “[t]he absence of heartbeat messages signals that one of the controllers is off-line and the remaining controller immediately initiates a failover operation and then begins servicing I/O requests directed to itself and its

off-line partner to provide non-stop access to data. Controller failover / failback operations are host independent and transparent to SAN nodes. During the failover / failback process, SAN nodes simply continue sending I/O requests to the same ID's across the SAN interconnect.” (Mylex paper at 16; *see also id.* at 11.)

134. Implementing failover in the system disclosed by the Mylex paper includes using a second port of a non-faulty RAID controller to perform operations of a port of a faulty controller. For example, Figures 17, 18 and 19 of the Mylex paper show one implementation of failover being performed by configuring a port of a second controller that is not faulty to perform functions of a port of a first, faulty controller, and the return process once the fault is resolved.

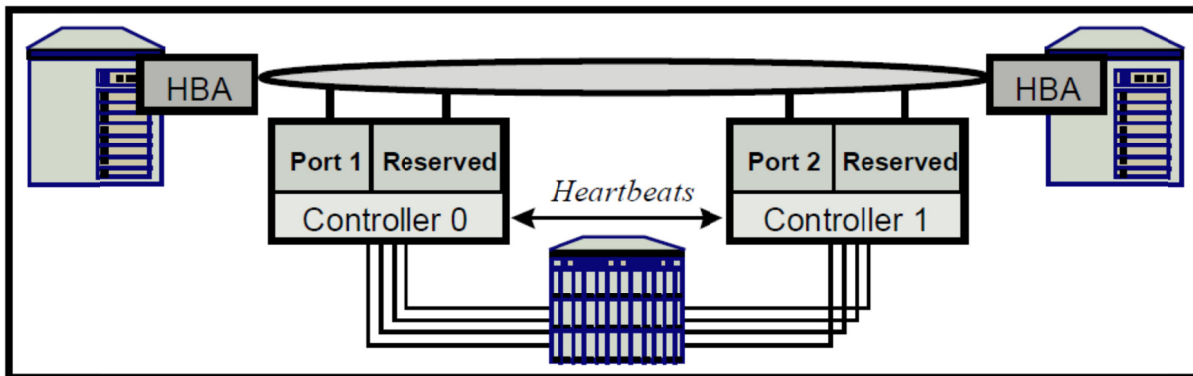


Figure 17. Normal Operating Mode of Dual DAC SF and FL Fibre Controllers




Figure 18. DAC SF and FL Fibre Array Controller Failover

**Performs
Functions of
Port 1 of
Controller 0**




Figure 19. DAC SF and FL Fibre Array Controller Failback

135. As the Mylex paper discloses, “[i]f a controller fails, the surviving controller senses the absence of heartbeats, fails over the ID of the active port on the failed controller to its reserved port, and updates its data structures with configuration information stored on disk. The failover process is transparent since the nodes still see the same fibre port ID’s on the SAN interconnect.” (*Id.* at 18.) Mylex Figures 17 through 19 show that the second network interface controller in the first RAID controller (Controller 1 Port 2) is used for fault tolerance by

performing functions of the first network interface controller in the second RAID controller (Controller 0 Port 1) when the second RAID controller (Controller 0) is faulty.

136. Therefore, the Mylex paper teaches “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,” as claimed by the ’346 patent.

9g) 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

137. One of ordinary skill in the art would recognize that the Mylex failover discussion applies equally to failure of either controller. For example, with reference to Mylex’s Figure 17, when one of the first or second RAID controllers are faulty, one of the second network interface controllers (labeled “Reserved”) will perform the functions of the first network interface controllers in the faulty RAID controller.

*9h) 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 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*

138. This claim element is similar to language in claim 1 that recites: “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.” Thus, as discussed above, the Mylex paper in view of the teachings of the Hathorn patent renders this element obvious.

VII. CHALLENGE # 2 - CLAIMS 1-9 ARE RENDERED OBVIOUS BY THE HATHORN PATENT (Exhibit VMWARE-1005) IN VIEW OF THE TEACHINGS OF THE MYLEX PAPER (Exhibit VMWARE-1006)

139. As an alternative view to Challenge 1 (discussed above), it is my opinion that the Hathorn patent in view of the teachings of the Mylex paper renders obvious claims 1-9 of the '346 patent. I provided brief summaries of the Hathorn patent and the Mylex paper above. (See Sections VI(A) and VI(B).) Below, I provide a more detailed discussion identifying the disclosures in the Hathorn patent and teachings of the Mylex paper that support my opinions.

140. As described below, a person of ordinary skill in the art would understand that the Hathorn patent discloses every element of the '346 patent's claims 1-9. However, to the extent that the fault tolerance elements in claims 4 and

9 are found to not be disclosed in the Hathorn patent,⁵ the Mylex paper teaches the claimed fault tolerance function of a storage controller port (*i.e.*, network interface controller) on one RAID controller taking over the identity and function of a storage controller port on a different RAID controller, and vice versa. This is illustrated by the Mylex paper's Figures 17 and 18, reproduced below:

⁵ Clam 4 recites: “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 recites: “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 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.***”




Figure 17. Normal Operating Mode of Dual DAC SF and FL Fibre Controllers




Figure 18. DAC SF and FL Fibre Array Controller Failover

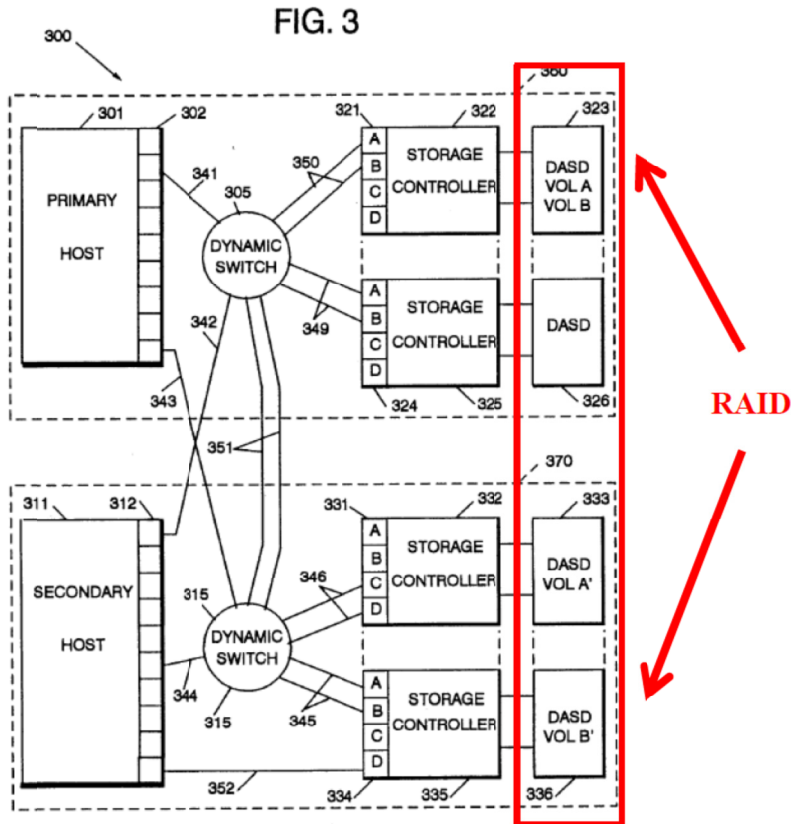
141. A person of ordinary skill in the art would have been motivated to apply these teachings of the Mylex paper to the system disclosed by the Hathorn patent in order to render every claim in the '346 patent obvious. (See ¶¶ 33-34 above.)

A. The Hathorn Patent In View Of The Teachings Of The Mylex Paper Compared To The '346 Patent, Claims 1-9

1. Claim 1

1a) An apparatus for a redundant interconnection between multiple hosts and a RAID, comprising:

142. The Hathorn patent discloses this claim element. For example, as shown in the annotated version of Hathorn Figure 3 below, the Hathorn patent discloses multiple hosts, including primary and secondary hosts 301 and 302. The Hathorn patent further discloses redundant connections such as 341, 342, 343 and 344, respectively, to dynamic switches 305 and 315.



143. One form of RAID is RAID 1, in which the data from one disk is mirrored (or “shadowed”) on another disk. Figure 5 of the Hathorn patent discloses this RAID configuration of shadowing data across multiple disks to create a remote dual copy. (*Id.* at 8:64-9:51; *see also id.* at 12:54-60.) The Hathorn patent discloses that in a RAID configuration, data is written to “many DASDs.” (*Id.* at 2:4-8.) As such, one of ordinary skill in the art would understand that the DASDs in Hathorn Figure 3 (323, 326, 333, and 336) could be configured as a RAID. Additionally, Hathorn Figure 3 above illustrates a DASD 323 that includes “Vol. A” and “Vol. B” being mirrored (as RAID 1) to DASD 333 and DASD 336. Since the Hathorn patent shows data mirroring of one DASD 323 to two other DASDs 333 and 336, one of ordinary skill would understand that DASD 326 can perform similar data mirroring with DASDs 333 and 336 as part of the RAID configuration.

1b) a first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers

144. The Hathorn patent discloses “a first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers.”

145. Figure 3 of the Hathorn patent, for instance, discloses storage controllers 322, 325, 332 and 335.


146. The storage controllers disclosed in the Hathorn patent are RAID controllers because they perform at least the RAID-1 function of disk mirroring. Thus, storage controller 325 is an instance of a first RAID controller and storage controller 335 is an instance of a second RAID controller.

147. The Hathorn patent discloses “processing a requirement of numerous host computers” because it discloses connections to hosts and the steps in performing a requirement (processing a write command) for the hosts. Specifically, the Hathorn patent discloses that the storage controllers can be configured as a “control unit link-level facility” for “communicating with the host processor,” or as a “channel link-level facility for communicating with the another storage controller.” (*Id.* at 5:9-13.) These communications involve a storage controller processing write commands (i.e., requirements) from a host. (*See, e.g., id.* at Fig. 6; *id.* at 9:52-59 (“the primary storage controller 325 will interface acting as a host to the secondary storage controller 335...while operating in duplex pair mode, the primary storage controller 325 intercepts primary host 301 write commands to the primary storage controller 325 duplex pair devices”).)

*1c) the first RAID controlling unit including a **first network controlling unit** and a **second network controlling unit**, and the second RAID controlling unit including a **third network controlling unit** and a **fourth network controlling unit***

148. The Hathorn patent discloses this claim element.

149. For example, the Hathorn patent discloses that the “A-D” blocks illustrated on each storage controller in Figure 13 are “controller ports.” (*Id.* at 8:3-5.) One of ordinary skill in the art would understand that when the four DASDs blocks illustrated in Figure 13 are configured as a RAID, then the ports A-D on RAID controllers 325 and 335 meet the claimed requirement for the first through fourth network interface controlling units. The annotated version of Hathorn Figure 3 below illustrates this:



1d) a plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers

150. The Hathorn patent discloses this claim element.

151. As shown in the annotated version of Hathorn Figure 3 above, two dynamic switches 305 and 315 connect the first and second RAID controlling units to the host computers 301 and 311.

*1e) 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*

152. The Hathorn patent discloses this claim element. For example, Figure 3 of the Hathorn patent discloses storage controllers 325 and 335 connected to dynamic switches 305 and 315 by communication links 349 and 345. The dynamic switches 305 and 315 are connected to hosts 301 and 311 through communication links 341, 342, 343, and 344.

153. Figure 6 of the Hathorn patent is a flow chart that further illustrates the exchange of information between the hosts 301 and 311 and the storage controllers 325 and 335. The flowchart shows the steps performed by the storage controllers (primary and secondary control units (CU)) in order to process a write command from a host. The features taught by Hathorn Figure 6 can be applied equally to either the storage controller 325 or the storage controller 335.

154. The Hathorn patent further discloses that the “primary storage controller 322, via port A 321, can communicate with primary host 301 by communication links 350, dynamic switch 305 and communication link 341, wherein port A 321 is a control unit link-level facility. Alternately, 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.” (*Id.* at 8:6-15; *see also id.* at 4:46-51 (“According to a first embodiment of the present invention, a method of communicating between a host processor, a first storage subsystem, and a second storage subsystem, the host processor, and first and second storage subsystems coupled together by at least one communication link and at least one dynamic switch.”).)


155. Thus, the Hathorn patent discloses the first RAID controlling unit and the second RAID controlling unit directly exchanging information with the numerous host computers through the plurality of connecting units, as claimed by the ’346 patent.

1f) and 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

156. As described in more detail above (see ¶48-55), the Hathorn patent discloses modifying storage controller ports so that they operate as channel link-level facilities where the ports (*i.e.*, network controlling units) exchange information with other ports. For example, Figure 3 of the Hathorn patent discloses that the “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.”

(*Id.* at 8:3-15; *see also, e.g. id.* at 11:25-43 (“The primary storage controller 325, acting as host with the ports 324 enabled as channel link-level facility, sends an EPC frame to the secondary storage controller 335 ... the secondary storage controller 335 processes the EPC frame and returns an acknowledgement (ACK) frame.”); *id.* at Fig. 7 (flowchart disclosing this process).)

157. The Hathorn patent teaches that modifying the storage controller ports in this fashion enables the ports to communicate using the existing switch network instead of using direct communication paths between RAID controllers. (*Id.* at 10:41-67; *compare* Hathorn’s Figure 2 (including direct communication paths 247 and 248) and Figure 3 (not including direct communication paths 247 and 248).) The Hathorn patent calls this establishing a “peer-to-peer logical path” between storage controller ports. (*Id.* at 10:64-67.) As such, one of ordinary skill in the art would have understood that the Hathorn patent discloses network interface controllers exchanging information as claimed by the ’346 patent, and as illustrated in the annotated version of Hathorn Figure 3 below.



158. Furthermore, Figure 7 of the Hathorn patent discloses a process for defining peer-to-peer communications between storage controllers 325 and 335, in which an information exchange includes storage controller 335 responding to EPC (establish pathing control) frames transmitted by storage controller 325. For instance, the Hathorn patent states that “the primary storage controller 325, acting

as host with the ports 324 enabled as channel link-level facility, sends an EPC frame to the secondary storage controller 335 at step 711.... At step 712 the secondary storage controller 335 processes the EPC frame and returns an acknowledgement (ACK) frame.” (*Id.* at 11:25-43.)

2. Claim 2

*2a) The apparatus as recited in claim 1, wherein said respective **RAID controlling units are connected to the plurality of individual connecting units.***

159. The Hathorn patent discloses that storage controllers are connected to dynamic switches through multiple communication links.

160. As I explained in my discussion of claim element 1(d), Figure 3 of the Hathorn patent discloses RAID storage controllers 325 and 335 connected to dynamic switches 305 and 315 by communication links 345 and 349.

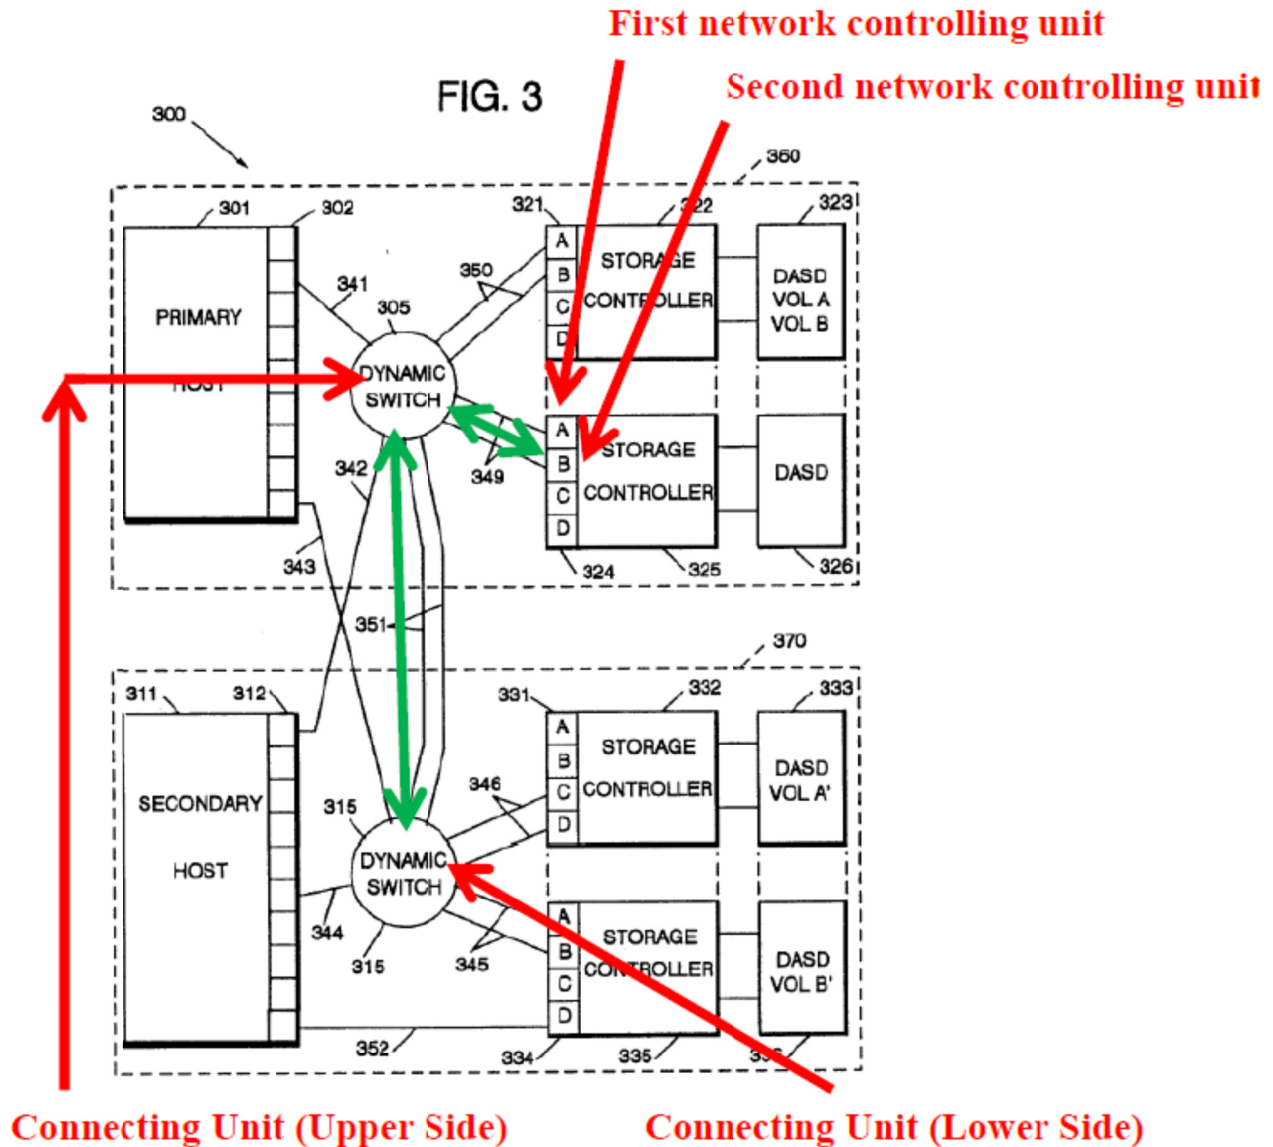
3. Claim 3

*3a) 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***

161. The Hathorn patent discloses this claim element. First, I note that in describing a real world system, the language requiring a “connecting unit of one side” is not specific. I also note that the ’346 patent does not use the term “side” except within Claim 3. However, for purposes of this discussion, and under the broadest reasonable interpretation standard, I assume that the “sides” referenced by Claim 3 are the sides of the figures included in the ’346 patent, and as such, I refer

below to the sides of the figures in the Hathorn patent in my analysis of this claim element.

162. Under this interpretation of Claim 3, the Hathorn patent discloses this claim element. As shown by the annotated version below of Figure 3 of the Hathorn patent, one of ordinary skill in the art would understand that Figure 3 discloses that the first network interface controlling unit (port 324A) is coupled to dynamic switch 305 on the upper “side” of Hathorn Figure 3 by connection link 349. One of ordinary skill in the art would also understand that the second network interface controlling unit (port 324B) is coupled to dynamic switch 315 on the lower “side” of Hathorn Figure 3 by connection link 349, dynamic switch 305, and connection link 351.



4. Claim 4

*4a) 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***

163. The Hathorn patent discloses this claim element. For example, while referring to Figure 2, the Hathorn patent discloses that the “primary host 201 can

thus communicate with any secondary storage controller 232, 235, or the secondary host 211 via the dynamic switch 205 or 215.” (*Id.* at 7:28-30.) One of ordinary skill in the art would have understood that this disclosure also applies to Hathorn Figure 3. As such, one of ordinary skill in the art would have understood, as shown by the annotated version of Hathorn Figure 3 (shown above in my discussion of claim 1(f), that the Hathorn patent discloses that first and third network interface controlling units are used to process the requirements of at least two host computers.

164. Further, to the extent this element is not disclosed by the Hathorn patent itself, the Hathorn patent discloses that in the case of a fault, storage controller ports on two different RAID controllers can process host requirements. (*See, e.g., id.* at 2:47-50 (“The secondary or remote location, in addition to providing a back-up data copy, must also have enough system information to take over processing for the primary system should the primary system become disabled.”).)

165. Therefore, the Hathorn patent discloses the first network interface controlling unit and the third network interface controlling unit processing the requirement of numerous host computers.

*4b) the **second** network interface controlling unit and the **fourth** network interface 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 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***

166. It is my opinion that a person of ordinary skill in the art would understand that the Hathorn patent discloses this claim element. Figure 3 of the Hathorn patent, for example, discloses multiple storage controllers with multiple ports, and multiple paths between the switches and ports, thus strongly suggesting a fault tolerant system. Further, the Hathorn patent discloses use of “IBM 3990 storage controller[s].” (*Id.* at 5:59.) One of ordinary skill in the art would understand that IBM 3990 storage controllers were fault tolerant controllers. (*See, e.g.,* Exhibit VMWARE-1012, Spainhower, *Design for Fault-Tolerance in System ES /9000 Model 900*, IEEE (1992) at p. 44, Figure 5 (illustrating that the IBM 3990 storage controller would dynamically select a different path (to a different storage controller port), in the event of a fault); *id.* (“Dynamic reconnection, the ability for an I/O device to continue an in-process I/O operation with an attached channel other than the one which initiated the operation, also provides fault tolerance from permanent failures in a single channel.”).)

167. Further, as explained above, the Hathorn patent discloses that storage controller ports can be modified to exchange communications with other controller ports. Therefore, one of ordinary skill in the art would understand that the second and fourth network interface controlling units could be configured as channel link facilities to exchange information when the RAID controllers are not faulty. The annotated version of Hathorn Figure 3 shown above in my discussion of Claim 1(f) illustrates this configuration. One of ordinary skill would have found it obvious to use the second and fourth network interface controlling units for communicating in a non-fault state, if the first and third network interface controlling units, e.g., on the fault tolerant IBM 3990 controllers, are being used for processing host requests, as this configuration will have the least performance impact on the active ports. Additionally, one of ordinary skill in the art would understand that in this configuration, the second and fourth network interface controlling units could be configured to execute a function of the first and third network interface controlling units when one of the RAID controlling units is faulty.

168. Further, the Hathorn patent in combination with the Mylex paper renders this element claim obvious. As discussed above (see ¶¶39-41, 140-141), the Mylex paper teaches the claimed fault tolerance functionality when disclosing that a reserved port on one RAID controller takes over the function of the active

port on a second RAID controller in the event of a fault. This is shown in the annotated versions below of the Mylex paper's Figures 17 through 19 :




Figure 17. Normal Operating Mode of Dual DAC SF and FL Fibre Controllers

1st Controller **2nd Controller**




Figure 18. DAC SF and FL Fibre Array Controller Failover




Figure 19. DAC SF and FL Fibre Array Controller Failback

169. The Mylex paper teaches the use of duplex RAID controllers for fault tolerance, for example, when stating that “SAN-attached arrays should be configured with duplex controllers and with the disks connected to both

controllers. Multiple SAN interfaces (on each controller) and duplex controllers with shared disks provide the level of fault tolerance required in SAN configurations.” (*Id.* at 11; *see also id.* at 12, 15, 21.)

170. The Mylex paper further discloses that “Mylex controllers have dual SAN ports which doubles the bandwidth to controllers and allows redundant paths from other SAN devices to the controllers to increase the resiliency of the SAN topology. As described later in this paper, dual host ports are particular critical for controller failover in Fibre Channel topologies. The SAN ports can be connected directly to UNIX and NT servers or indirectly through hubs and switches.” (*Id.* at 16.)

171. As such, one of ordinary skill would find it obvious to modify the system disclosed in the Hathorn patent’s Figure 3 with the fault tolerance teachings of the Mylex paper to configure a system where the second and fourth network interface controlling units exchange fault tolerance information, and where the second and fourth network interface controlling units are used to execute a function of the active first and third network controlling units in the event of a fault.

172. I note that for purposes of this claim element, I am interpreting the claim language under the broadest reasonable interpretation standard to recite “... the second network interface controlling unit ~~and~~ or the fourth network interface

controlling unit are used for executing a function of the first network interface controlling unit ~~and~~ or the third network controlling unit when one of the first RAID controlling unit and the second RAID controlling unit is faulty.” I use this interpretation because claim 1 requires that the second and fourth network interface controlling units are on different RAID controllers, and as such, if one RAID controller is faulty, either the second or the fourth network interface controlling unit will not be used to execute any function. I reserve the right to offer a different interpretation in the district court litigation, which I have been informed uses a different legal standard.

5. Claim 5

5a) The apparatus as recited in claim 1, wherein said plurality of connecting units have at least three connection ports

173. The Hathorn patent discloses this claim element.

174. For example, Figure 3 of the Hathorn patent shows dynamic switch 305 connected to eight ports: communication links 341, 342, 349 (two ports), 350 (two ports), and 351 (two ports). Similarly, dynamic switch 315 is shown connected to eight ports: communication links 343, 344, 345 (two ports), 346 (two ports), and 351 (two ports).

175. Thus, the Hathorn patent discloses the plurality of connecting units having at least three connection ports.

5b) 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

176. One of ordinary skill in the art would understand that in a switch or hub, each port is coupled to every other port.⁶ For example, in a four-port switch, with one port connected to the first network interface controlling unit and one port connected to the third network interface controlling unit, the remaining two connection ports are coupled to both the first and the third network interface controlling units. As such, the Hathorn patent discloses that at least two of the eight connection ports for each switch are coupled to the first and the third network interface controlling units. For example, Figure 3 of the Hathorn patent discloses that each of the eight ports on each switch is coupled to each port on each storage controller.

5c) and the rest of the connection ports being provided as a hub equipment connected with the numerous host computers

177. The Hathorn patent in combination with the Mylex paper renders this claim element obvious.

178. First, I note that this claim element appears to be inconsistent with claim 1, in that claim 5 as a whole only recites that the two ports on the connecting units are connected to the first and third network interface controlling units and the

⁶ See above, ¶103.

“rest” are connected to the hosts—as such, no ports are connected to the second and fourth network controlling ports. However, one of ordinary skill in the art would understand that in switches, every port is connected to every other port, and therefore that the Hathorn patent’s Figure 3 discloses a system where each of the eight ports on each switch is coupled to both hosts.

179. Additionally, the Mylex paper teaches that RAID systems can use switches and/or hubs. (*See, e.g.*, Mylex paper at Figures 20 and 21; *see also id.* at 5 (“Switches, hubs and routers are interconnect devices that can be employed to construct SAN networks.”).)

180. Additionally, the ’346 patent discloses that a “hub” can be a hub or a switch when it states that “the hubs 440, 442 are provided to connect a system connected to these hubs by one network and maintain the network even though one system has an occurrence of a trouble or a short of a line, and it can be as a hub or a switch. Hereinafter, they are named a “hub” altogether.” (’346 Patent at 3:13-18.)

181. As such, the Hathorn patent either discloses this claim element, or renders it obvious in combination with the teachings of the Mylex paper.

6. Claim 6

6a) The apparatus as recited in claim 1, wherein said plurality of connecting units have at least three connection ports

182. This claim element is identical to claim element 5(a), and therefore my analysis as to how the Hathorn patent discloses this claim element is identical.

6b) two of the at least three connection port are coupled to one of the first network controlling unit and the third network controlling unit

183. This claim element is identical to the corresponding element in claim 5(b). Thus, as discussed above, the Hathorn patent discloses this claim element.

6c) and the rest of the connection ports being provided as a network switch equipment connected with the numerous host computers

184. The Hathorn patent discloses this element, as explained above in my discussion of similar claim limitation 5(c).

7. Claim 7

7a) The apparatus as recited in claim 1, wherein said plurality of connecting units have at least five connection ports

185. The Hathorn patent discloses this claim element. For example, in my discussion above of claim element 5(a), I have shown how the Hathorn patent discloses using switches with at least eight connection ports.

7b) 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

186. The Hathorn patent discloses this claim element. For example, in my discussion above of claim element 5(b), I have shown how the Hathorn patent discloses a system in which all eight connection ports are coupled to the first and third network controlling units.

7c) and the rest of the connection ports being provided as a switch connected with the numerous host computers


187. This claim element is similar to claim element 6(c) discussed above. For the same reasons, the Hathorn patent discloses this claim element.

8. Claim 8

8a) 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

188. I first note that the '346 patent broadly uses the term "connected" to encompass indirect connections, e.g., through a hub or switch. (See, e.g., '346 patent at 4:23-34 ("the host computers 500, 501, 502, 503, 504 and 505 are connected to the RAID 530 by using external hubs 510 and 520").)

189. As such, under the broadest reasonable interpretation standard, the Hathorn patent discloses that the first and second ports of a first controller are connected to each of two different switches, and that first and second ports of a second controller are also connected to each of the two different switches. For example, one of ordinary skill in the art would understand that Figure 3 of the Hathorn patent shows that (i) the first network interface controlling unit is connected to a first connecting unit, switch 305, through connection path 349; (ii) the second network interface controlling unit is connected to a second connecting unit, switch 315, through connection path 349, switch 305, and connection path 351; (iii) the third network interface controlling unit is connected to the connecting unit, switch 315, through connection path 345; and (iv) the fourth network interface controlling unit is connected to the first connecting unit, switch 305, through connection path 345, switch 315, and connection path 351.



190. Additionally, the Hathorn patent discloses that the communication links between storage controllers can vary. (Hathorn patent at 13:13-16.) As such,

one of ordinary skill in the art would understand that a direct connection from any storage controller port on any RAID controller could have been established with any dynamic switch.

9. Claim 9

*9a) An apparatus for a **redundant interconnection** between **multiple host** computers and a **RAID**, the apparatus comprising*

191. This claim element is identical to the corresponding element in claim 1. Thus, as discussed above, the Hathorn patent discloses “an apparatus for a redundant interconnection between multiple hosts and a RAID.”

*9b) a **plurality of connecting units** for connecting the host computers and the RAID*

192. This claim element is similar to an element in claim 1 which recites: “a plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers.” Thus, as discussed above, the Hathorn patent discloses “a plurality of connecting units for connecting the host computers and the RAID.”

*9c) a **first and a second RAID controllers**, included in the RAID, 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***

193. This claim element is similar to language in claim 1 which recites: “a first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers, the first RAID controlling unit including

a first network controlling unit and a second network controlling unit, and the second RAID controlling unit including a third network controlling unit and a fourth network controlling unit; and a plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers.”

194. Thus, as discussed above, the Hathorn patent discloses first and second RAID controllers, included in the RAID, each of which has a first network interface controlling unit and a second network interface controlling unit for processing requests from the plurality of the host computers connected through the plurality of connection units.

*9d) 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 processes information transmitted from the second network interface controller** in the second RAID controller*

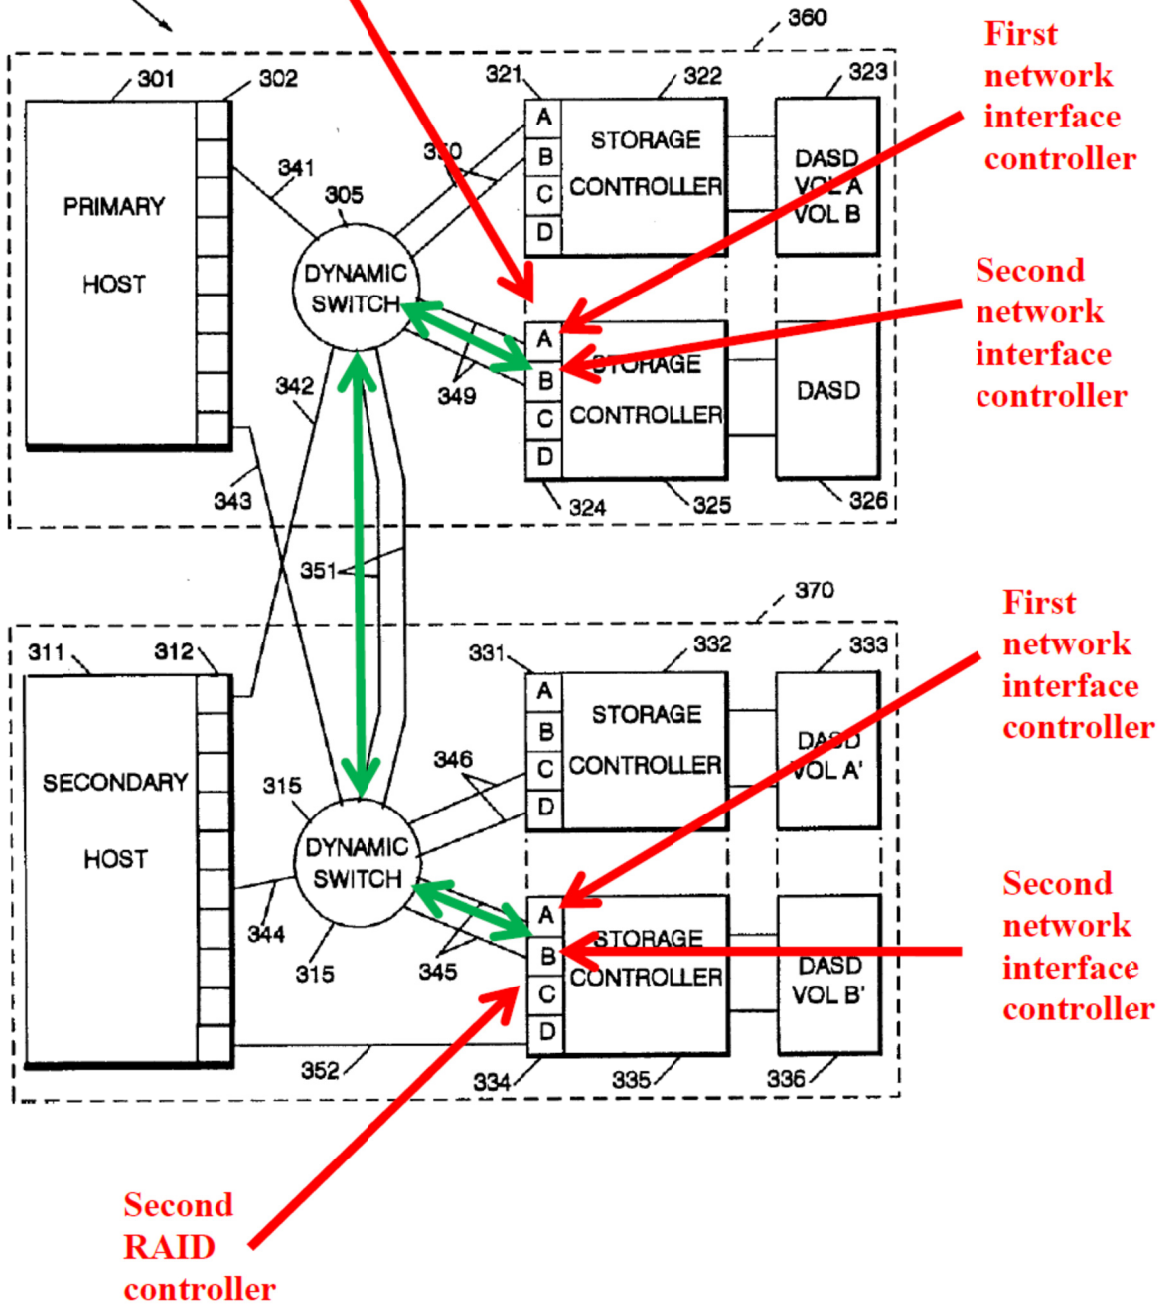
195. The Hathorn patent discloses this claim element. For example, as discussed with respect to claim elements 1(e) and 1(f), the Hathorn patent discloses that both RAID controllers, through their respective storage controller ports (i.e. network interface controllers), can exchange information with the host computers and each other. For example, the Hathorn patent discloses that when the ports are configured as a channel link-level facility, “primary storage controllers 222, 225 [can] send data or records for back-up directly to secondary storage controllers

232, 235, respectively.” (*Id.* at 7:38-41.) As such, and as shown in the annotated version of Hathorn Figure 3 below, one of ordinary skill would understand that, for example, the first network interface controller (324A) in the first RAID controller (325) could be used to supply data to the hosts and process information received from the second network interface controller (334B) in the second RAID controller (335). (*See also id.* at 5:8-15 (disclosing that when ports are modified as channel link-level facility “the storage controller acts as host to the another storage controller”); *id.* at 8:3-15 (“primary storage controller 322, via port A 321, can communicate with primary host 301 by communication links 350, dynamic switch 305 and communication link 341, wherein port A 321 is a control unit link-level facility. Alternately, 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.”); 7:28-38 (“The primary host 201 can thus communicate with any secondary storage controller 232, 235, or the secondary host 211 via the dynamic switch 205 or 215. Likewise, the secondary host can communicate with any primary storage controller 222, 225, or the primary host 201 via the dynamic switch 205 or 215. Additionally, primary storage controllers 222, 225 can communicate with secondary storage controllers 232, 235, respectively. Thus, the

primary host 201 could send data or records for back-up directly to the secondary storage subsystem (however, this may be undesirable due to the required primary host resources).”).)

First RAID controller

FIG. 3



*9e) 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 processes information transmitted from the second network interface controller** in the first RAID controller*

196. The Hathorn patent discloses this claim limitation. As explained in my discussion above of claim element 9(d), the Hathorn patent discloses that both RAID controllers can supply data to either host and that any controller port can be modified to act as a channel link-level facility to process the data of any other controller port. (*See also, e.g., id.* at 7:28-33 (“The primary host 201 can thus communicate with any secondary storage controller 232, 235, or the secondary host 211 via the dynamic switch 205 or 215. Likewise, the secondary host can communicate with any primary storage controller 222, 225, or the primary host 201 via the dynamic switch 205 or 215.”).)

*9f) 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*

197. The Hathorn patent alone or in view of the Mylex paper anticipates or renders this claim element obvious. As explained above in my discussion of claim element 4(b), the Hathorn patent discloses fault tolerant IBM 3990 RAID controllers. Additionally, the network topology of the Hathorn patent’s Figure 3 strongly suggests a fault tolerant architecture.

198. Moreover, as explained above in my discussion of claim element 4(b), the Mylex paper discloses fault tolerance in Figures 17 through 19, which show that ports on alternate RAID controllers (i.e., network interface controllers) perform the function of a RAID controller port on a faulty RAID controller.

9g) 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

199. For similar reasons as those discussed above with respect to claim element 9(f), the Hathorn patent discloses “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,” as claimed by the ’346 patent.

9h) 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 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

200. This claim element is similar to language in claim 1 which recites: “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.” Thus, as discussed above, the Hathorn patent discloses this claim element.

VIII. CHALLENGE # 3 – CLAIMS 1-9 ARE RENDERED OBVIOUS BY THE DEITZ PATENT (Exhibit VMWARE-1008) OR THE MYLEX PAPER (Exhibit VMWARE-1006) IN VIEW OF THE TEACHINGS OF THE GRIFFITH PATENT (Exhibit VMWARE-1007) OR THE DEKONING PATENT (Exhibit VMWARE-1010)

201. It is my opinion that either the Deitz patent, or the Mylex paper, in view of the teachings of the Griffith patent or the DeKoning patent, renders obvious claims 1-9 of the '346 patent. I provided a brief summary of the Mylex paper above. (*See* Section VI(A).) I provide a brief summary of the Griffith, DeKoning and Deitz patents below, and then a more detailed description identifying the disclosures in the Deitz patent, the Mylex paper, and the teachings of the Griffith and DeKoning patents that support my opinion.

202. As discussed below, the Griffith patent teaches that (i) RAID controllers can communicate either by a direct path between controllers *or* by using the existing switch network, and (ii) dual-ported RAID controllers can act both as primary controller for their associated disks and as secondary controller for the disks of another RAID controller in the case of a fault, and vice versa. The DeKoning patent teaches that several communication mediums, e.g., the host-side communication bus, can be used to allow RAID controllers to exchange information

203. The Mylex paper and the Deitz patent both disclose every element of the '346 patent's claims 1-9, with the exception that they disclose RAID

controllers that communicate via a direct “heartbeat” path. One of ordinary skill in the art would have been able to apply the teachings of the Griffith and/or the DeKoning patent to either the Mylex paper or the Deitz patent to render the ’346 patent’s claims obvious.

A. Brief Summary of the Griffith and DeKoning Patents

204. The Griffith patent is titled “RAID Systems During Non-Fault And Faulty Conditions On A Fiber Channel Arbitrated Loop SCSI Bus Or Switch Fabric Configuration.” The application for the Griffith patent was filed on August 18, 1999, and the patent issued on June 4, 2002. Digi-Data Corporation is the assignee identified on the face of the Griffith patent. The DeKoning patent is titled “Methods And Apparatus For Coordinating Shared Multiple Raid Controller Access To Common Storage Devices.” The application for the DeKoning patent was filed on December 23, 1996, and the patent issued on June 6, 2000. LSI Logic Corp. is the assignee identified on the face of the DeKoning patent.

205. As explained below, a primary teaching of the Griffith patent is that network interface controlling unit ports can exchange information either by (i) a direct path between controllers, or (ii) using the existing switch network, as disclosed in the ’346 patent. (*See* Griffith patent at 9:15-22.)


206. Indeed, the ability to use an existing communications network instead of a direct heartbeat path was well known to one of ordinary skill in the art in the

2000 time frame. For example, the DeKoning patent, Exhibit VMWARE-1010—cited by the Griffith patent in the “References Cited” section—discloses an “invention [that] provides inter-controller communications to obviate the need for host system intervention to control failover operations among the controllers... [so that a plurality of RAID controllers] communicate among themselves to permit continued operations in case of failures.” (DeKoning patent at 3:15-21.) The DeKoning patent further discloses that:

information and the cache data and meta-data **are exchanged between the plurality of shared controllers through any of several communication mediums.** A dedicated communication bus interconnecting all RAID controllers may be preferred for performance criteria, but may present cost and complexity problems. Another preferred approach is where the information is exchanged via the communication bus which connects the plurality of controllers to the common subset of disk drives in the common LUN. This communication bus may be any of several industry standard connections, including, for example, ... Fibre Channel ... *Similarly the host connection bus which connects the plurality of RAID controllers to one or more host computer systems may be utilized as the shared communication medium.* In addition, the communication medium may be a shared memory architecture in which the plurality of controllers share access to a common, multiported memory subsystem (such as the cache memory subsystem of each controller). (*Id.* at 4:58-5:10 (emphasis added).)

207. The Griffith patent discloses a “RAID system” that “uses arbitrated fiber channels or switch fabric to connect multiple host computers and storage array controllers (SAC).” (Griffith patent at Abstract). Figure 5 of the Griffith

patent discloses an embodiment of an “ACTIVE-ACTIVE redundant RAID system ... which incorporates a switch fabric configuration.” (*Id.* at 4:52-55.) An annotated version of Griffith Figure 5 appears below:



208. The annotations above are supported by the disclosure in the Griffith patent. For example, the Griffith patent discloses that 210 and 310 (and 510 and 610) are “host computers.” (*Id.* at 7:47-48.) The Griffith patent discloses that 230, 330, 430, and 530 are SACs (storage area controllers). (*Id.* at 7:48-49.) The Griffith patent discloses that “the components in Fig. 5 are connected by

connecting means to a switch fabric device 24.” (*Id.* at 8:22-24.) While the number “24” does not appear in Griffith Figure 5, one of ordinary skill in the art would recognize that the block annotated above as “switch” is the switch connecting the hosts and SACs. Additionally, the Griffith patent discloses that lines 212, 236, 234, 336, 334, and 312 (and others) are “connectors” connecting the hosts, switch, and SACs. (*Id.* at 7:50-52.) These “[c]onnectors may be fiber optics or copper wires.” (*Id.* at 8:49.) As such, one of ordinary skill in the art would understand that the entry points for the connectors on the hosts, switch, and SACs are ports (i.e. network interface controller ports). While the Griffith patent does not number the SAC ports, one of ordinary skill in the art would understand that Griffith Figure 5 could be labeled as shown above.

209. The Griffith patent discloses a redundant RAID system. For example, the Griffith patent discloses that the switch fabric connecting the host computers and the controllers “provides redundancy in the case of any single computer or controller failure.” (*Id.* at 2:35-38; *see also id.* at 8:63-64.) The Griffith patent also discloses that “each SAC is designated a primary SAC for an array of storage units, which it normally serves as controller, and as a secondary SAC for another array of storage units.” (*Id.* at Abstract.)

210. As such, as shown by the annotated version of Griffith Figure 5 above, one of ordinary skill in the art would understand that the Griffith patent

teaches that when a RAID controller is faulty, one network interface controller port on a non-faulty RAID controller will detect that fault and then one of the network interface controlling units on the non-faulty RAID controller will take over processing the host requests for the faulty RAID controller, while the other network interface controller on the non-faulty RAID controller maintains its role in processing the host requests for its primary storage.

211. For example, the Griffith patent discloses that “[i]n the event of a failure of a primary SAC or its associated host computer, the secondary SAC, as a member of the storage array set, assumes the identity of the primary SAC, identifies the array of DASD in the storage array set, and controls both the array of the storage array set of which it was secondary SAC as well as its own array of DASD.” (*Id.* at 9:5-11.) One of ordinary skill in the art would understand that this failover process occurs by one network interface controlling unit (on the non-faulty SAC) assuming the identity/address for the failed network interface controlling unit (on the faulty SAC). (*See id.* at 9:25-31 (“the secondary SAC of the storage array set of which SAC 230 is primary SAC, in this case SAC 330, detects the absence of heartbeat from SAC 230. SAC 330 then uses its interface chip 332 to assume the identity of SAC 230, so that instructions and requests which are directed to SAC 230 are intercepted by SAC 330.”).)

212. The Griffith patent teaches that two primary SACs will monitor each other in order to determine if there is a fault warranting the non-faulty SAC using one of its network interface controlling units to assume the identity of the faulty SAC. For example, as shown by the annotated version of Griffith Figure 5 above, the Griffith patent teaches that “heartbeat signals” can be sent reciprocally over the network connections labeled as lines 823 and 824. (*Id.* at 9:11-15.)

213. The Griffith patent also teaches that the SACs (i.e., RAID controllers) may exchange information reciprocally through their network interface controlling units using the existing switch network instead of the private network between RAID controllers 823 and 824:

In an alternative configuration, the heartbeat is emitted to the loop connecting means 22 which might be a fiber channel loop. If such a loop is used, the heartbeat signal would have to follow the established priority and arbitration procedures for use of the loop, a requirement which is avoided by use of a direct connections between the SACs. (*Id.* at 9:15-21.)

214. One of ordinary skill in the art would understand that the Griffith patent’s teaching of using the existing redundant switch network to exchange information across RAID controllers, instead of a direct connection between RAID controllers, applies regardless of whether a switch or a fiber channel loop is used. (*See id.* at 8:25-26 (“Loop connecting means may be a SCSI bus, fibre channel arbitrated loop, or a switch fabric device.”); *id.* at 9:37-40 (“The second

embodiment of FIG. 5 functions in exactly the same manner except the loop connecting means is replaced by a switch fabric chip.”.)

B. The Mylex Paper In View Of The Teachings Of The Griffith and/or DeKoning Patent Compared To The '346 Patent, Claims 1-9

215. It is my opinion that the Mylex paper in view of the teachings of the Griffith and/or DeKoning patent renders obvious claims 1-9 of the '346 patent. A brief summary of the teachings of the Griffith and DeKoning patents is provided above. (*See* Section VIII(A).) A brief summary and detailed comparison of the Mylex paper to the '346 patent's claims 1-9 also are provided above. (*See* Sections VI(A) and VI(C).)

216. As discussed above, a person of ordinary skill in the art would understand that the Mylex paper discloses every element of the '346 patent's claims 1-9, with the exception of a direct exchange of information between network interface controlling units. Instead, the Mylex paper discloses a direct “heartbeat” communication path between RAID controllers for exchanging information. However, the Griffith patent teaches that fault tolerance can be implemented by either exchanging information on a direct path between controllers (like in the Mylex paper) or, alternatively, by using the existing switch network. (*See, e.g.*, Griffith patent at 9:15-21.) Additionally, the DeKoning patent teaches that several communication mediums can be used to exchange information

between RAID controllers, including using the existing host-side communication bus. (See, e.g., DeKoning patent at 4:58-5:10.)

217. As such, a person of ordinary skill in the art would have found it obvious, in view of the teachings of the Griffith and/or DeKoning patents, to configure the Mylex paper's redundant RAID system to send failover communications on the existing switch network, and thus exchange information between the failover network interface controlling units on separate RAID controllers (e.g., in the Mylex paper's Figure 17, the respective RAID controllers 0 and 1 and their respective network interface controlling unit ports).

218. Below is an annotated version of Figure 17 of the Mylex paper:




Figure 17. Normal Operating Mode of Dual DAC SF and FL Fibre Controllers

network
controlling
unit 1

network
controlling
unit 2

network
controlling
unit 3

network
controlling
unit 4

219. As shown in annotated Mylex Figure 17 above, using the teachings of the Griffith and/or DeKoning patents, it would be obvious to configure RAID

Controller 0 to act as (i) a primary controller for its associated disks, and (ii) a secondary backup controller for the disks associated with RAID controller 1, and vice versa. Additionally, using the teachings of Griffith discussed above, all network interface controlling unit ports on each RAID Controller 0 can exchange heartbeat signals with all network interface controlling unit ports on RAID Controller 1. The particular ports used could be ports that are normally reserved, or ports also used for active data traffic. This configuration is strongly suggested by Figure 17 in the Mylex paper and its associated text. Similarly, it would be obvious to configure RAID Controller 1 to act as (i) a primary controller for its associated disks, and (ii) a secondary backup controller for the disks associated with RAID controller 0.

220. The result of this failover is shown, e.g., in the Mylex paper Figure 18, reproduced below, where one network interface controller port is used for executing a function of the network interface controlling unit port on the faulty RAID controller:




Figure 18. DAC SF and FL Fibre Array Controller Failover

221. A person of ordinary skill in the art would have been motivated to apply the teachings of the Griffith and/or DeKoning patents to the system disclosed in the Mylex paper in order to render every claim in the '346 patent obvious. The Mylex paper, the Griffith patent, and the DeKoning patents are in the same field of endeavor, and discloses redundant RAID systems. While the Griffith patent only discloses using one switch or hub, and the DeKoning patent discloses using a host-side communication bus, the concept of using multiple switches or hubs in RAID systems was well known at the time of the alleged invention. (*See, e.g.*, Hathorn patent at Figure 3.) Each reference discloses fault tolerance in terms of sending communications between two or more RAID controllers and/or network interface controlling unit ports. One of ordinary skill in the art would have been motivated to study multiple instances of fault tolerant RAID systems when designing a new RAID system. Additionally, one of ordinary skill in the art would have been motivated to combine the teachings of the Griffith patent with the Mylex paper's

system because both Mylex Corporation and Digi-Data Corporation (assignee of the Griffith patent) were RAID providers. A person of ordinary skill in the art would have known to look at the teachings of these RAID providers when configuring redundant RAID systems. Furthermore, both the Griffith patent and the Mylex paper disclose redundant RAID systems constructed from off-the shelf components, and as such their combination is merely the use of known techniques to achieve predictable results. (*See, e.g.*, Mylex paper at 15 (marketing “Mylex controllers”); Griffith patent at 5:33-35 (“A preferred SAC is the Z-9100 Ultra-Wide SCSI RAID controller manufactured by Digi-Data Corporation, Jessup, Md.”).)

222. Finally, one of ordinary skill in the art would be motivated to combine the teachings of the Griffith patent with the Mylex controllers because the Griffith patent discloses that its “preferred dual-port disk is the 3.5-Inch Ultrastar2 XP, available from IBM” (Griffith patent at 8:38-39), and there was a close relationship between IBM and Mylex Corporation at the time of the alleged invention. In September of 1999, IBM completed the acquisition of Mylex Corporation. Storage system designers using the IBM 3.5-Inch Ultrastar2 XP disclosed in the Griffith patent in that timeframe would have been strongly motivated to combine and leverage the teachings from other IBM and Mylex Corporation storage technology.

C. Brief Summary of the Deitz Patent

223. U.S. Patent No. 6,578,158 (the “Deitz patent”) is titled “Method And Apparatus For Providing A Raid Controller Having Transparent Failover And Fallback.” The Deitz patent was filed on October 28, 1999 and issued on June 10, 2003. IBM is the assignee identified on the face of the Deitz patent.

224. The Deitz patent discloses redundant connections to RAID systems. The Deitz patent discloses multiple host computers that are connected to a plurality of hubs, where 1) one hub is connected to (i) an active RAID controller port on a first RAID controller and (ii) an inactive RAID controller port on a second RAID controller, and 2) a second hub is connected to (i) an inactive RAID controller port on a first RAID controller and (ii) an active RAID controller port on a second RAID controller. The Deitz patent also discusses heartbeat signals being transmitted between the RAID controllers. Figures 1 and 2 of the Deitz patent disclose separate embodiments for the heartbeat path: a connection line 205 existing as an inter-RAID-controller path (Figure 1), and a connection line 205 existing as a storage-side path (Figure 2). A comparison of the Deitz patent’s Figure 2 and the ’346 patent’s Figure 4 appears below:

Deitz Patent Figure 2





FIG. 2

'346 Patent, Figure 4



225. I have been informed that challenges involving both the Deitz patent and the Mylex paper are being included in the petition because (i) the Patent Owner has not indicated whether it intends to seek a priority date earlier than the Deitz patent, and (ii) the Patent Owner cannot swear behind the priority date of the Mylex paper.

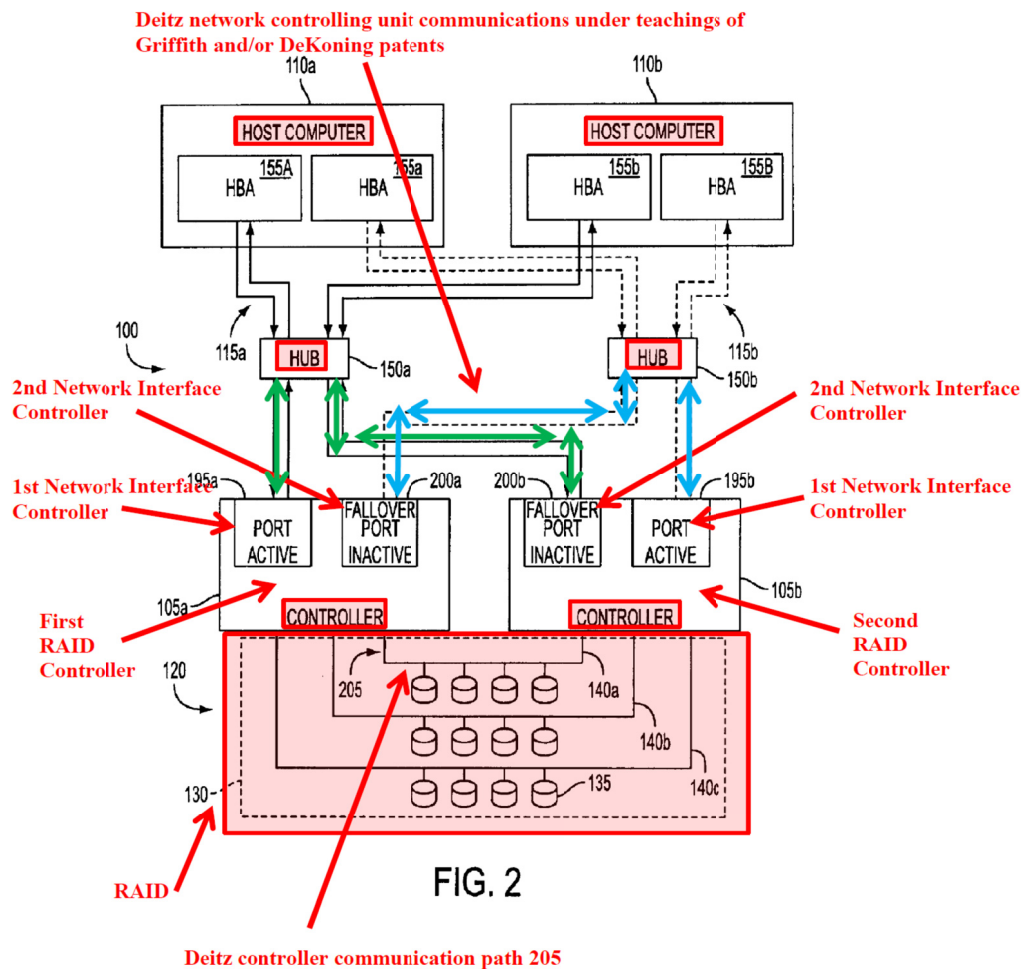
D. The Deitz Patent In View Of The Teachings Of The Griffith and/or DeKoning Patents Compared To The '346 Patent, Claims 1-9

226. As discussed above, the Griffith patent teaches that RAID controllers may exchange fault tolerance information reciprocally through their network interface controlling units using the existing switch network instead of the private network between RAID controllers. (*See, e.g.*, Griffith patent at 9:15-21.) Additionally, the DeKoning patent teaches that several communication mediums can be used to exchange information between RAID controllers, including using the existing host-side communication bus. (*See, e.g.*, DeKoning patent at 4:58-5:10.)

227. As discussed above (*see* Section VIII(C)), the Deitz patent discloses multiple host computers attached to at least two dual-ported RAID controllers through at least two hubs. The Deitz patent discloses that the network interface controlling units on the respective RAID controllers communicate and monitor one another's health by exchanging passive or dynamic signals (*see* Deitz patent at

6:63-7:5) across a communication path (e.g., path 205 in Deitz patent Figure 2) connecting the RAID controllers.

228. To illustrate this, an annotated version of Figure 2 from the Deitz patent appears below:



229. As shown in the annotated Figure 2 above, using the teachings of the Griffith patent, it would be obvious to configure both RAID controllers to act as (i) primary controllers for their own associated disks, and (ii) secondary backup controllers for the disks associated with the other RAID controllers. This

configuration is strongly suggested by Figure 2 in the Deitz patent and its associated text. Further, it would be obvious using the teachings of the DeKoning patent to configure both RAID controllers to exchange information via their NICs using the host-side communication network. As such, the network interface controlling unit ports on one RAID controller will exchange failover communications with the network interface controlling unit ports on the other RAID controllers.

230. One of ordinary skill in the art would have been motivated to combine the teachings of the Griffith and/or DeKoning patents with the Deitz patent's disclosure to configure a RAID system that renders obvious every claim in the '346 patent. For example, the Griffith patent, DeKoning patent, and the Deitz patent are concerned with creating redundancy in a RAID system to allow for failover and continuous operations in the event of a component fault. Additionally, each reference is in the same field of endeavor, disclosing redundant RAID systems and sending signal communications across RAID controllers for redundancy purposes. While the Griffith patent discloses a configuration using one switch or hub loop, and the DeKoning patent discloses using a host-side communication bus, the concept of using multiple connection units in a network was well-known, multiple hubs are explicitly disclosed in the Deitz patent. Additionally, one of ordinary skill in the art would have been motivated to

combine the teachings of the Griffith patent with the Deitz system because both IBM (assignee of the Deitz patent) and Digi-Data Corporation (assignee of the Griffith patent) were RAID providers. A person of ordinary skill in the art would have known to look at the teachings of these RAID providers when configuring redundant RAID systems. Furthermore, the Griffith and Deitz patents disclose redundant RAID systems constructed from off-the shelf components, and as such their combination is merely the use of known techniques to achieve predictable results. (*See, e.g.*, Deitz patent at 5:33-36 (“controllers 105 can be any suitable fibre channel compatible controller that can be modified to operate according to the present invention, such as for example the DAC960SF, commercially available from Mylex, Inc., Boulder, Colo.”); Griffith patent at 5:33-35 (“A preferred SAC is the Z-9100 Ultra-Wide SCSI RAID controller manufactured by DigiData Corporation, Jessup, Md.”).)

1. Claim 1

*1a) An apparatus for a **redundant interconnection** between **multiple hosts** and a **RAID**, comprising*

231. The Deitz patent discloses this claim element.

232. For example, the Deitz patent discloses multiple host computers and a RAID when stating that “[t]he present invention is directed to a memory system having a number of controllers adapted to transfer data between at least one host computer and a data storage system, such as one or more Redundant Array of

Independent Disks (RAID) storage systems.” (Deitz patent at 4:44-48; *see also id.* at 1:30-35.) The Deitz patent also discloses redundant interconnections between multiple host computers and a RAID, as shown in the annotated version below of Deitz Figure 2:




FIG. 2

233. The Deitz patent further discusses redundant connections between the host computers and the RAID. For example, the Deitz patent discloses that “a pair of the controllers 105 can be configured to operate as dual-active controllers as described above, or as dual-redundant controllers wherein one controller serves as an installed spare for the other, which in normal operation handles all I/O requests

from the 45 host computer 110. Preferably, the controllers 105 operate as dual-active controllers to increase the bandwidth of the memory system 100.” (*Id.* at 5:41-48; *see also id.* at 1:30-35; 10:22-27.)


234. Therefore, the Deitz patent discloses apparatus for a redundant interconnection between multiple hosts and a RAID, as claimed by the '346 patent.

1b) a first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers

235. The Deitz patent discloses first and second RAID controllers, and therefore teaches this claim element.

236. First, the title of the Deitz patent directly references RAID controllers: “Method And Apparatus For Providing A **RAID Controller** Having Transparent Failover And Failback.” (*Id.* at Title (emphasis added).) The Deitz patent discloses multiple RAID controllers in its Abstract, stating that the invention consists of “[a] method and apparatus for controlling a memory system 100 comprising a plurality of controllers 105 connected by a fibre channel arbitrated loop 145 to provide transparent failover and failback mechanisms for failed controllers. The controllers 105 are adapted to transfer data between a data storage system 120 and at least one host computer 100 in response to instructions therefrom.” (*Id.* at Abstract.)

237. The Deitz patent further illustrates first and second RAID controllers in Figure 2, where each controller 105a and 105b is a RAID controller. Specifically, the Deitz patent discloses that “in a preferred embodiment, shown in FIG. 2, reliability is further enhanced by providing a clustered environment in which two host computers (110a and 110b) each have direct access to both controllers 105 through a number of HBAs 155a-d.” (*Id.* at 6:35-39; *see also id.* at 4:64-5:5; 2:3-11; 3:3-12; 4:60-64.)



238. The Deitz patent further discloses that the RAID controllers process a requirement of the host computers. For example, the Deitz patent discloses “a

plurality of controllers 105 connected by a fibre channel arbitrated loop 145 to provide transparent failover and failback mechanisms for failed controllers. The controllers 105 are adapted to transfer data between a data storage system 120 and at least one host computer 110 in response to instructions therefrom.” (*Id.* at Abstract; *see also id.* at 3:43-49; 4:44-48; Claim 1.)

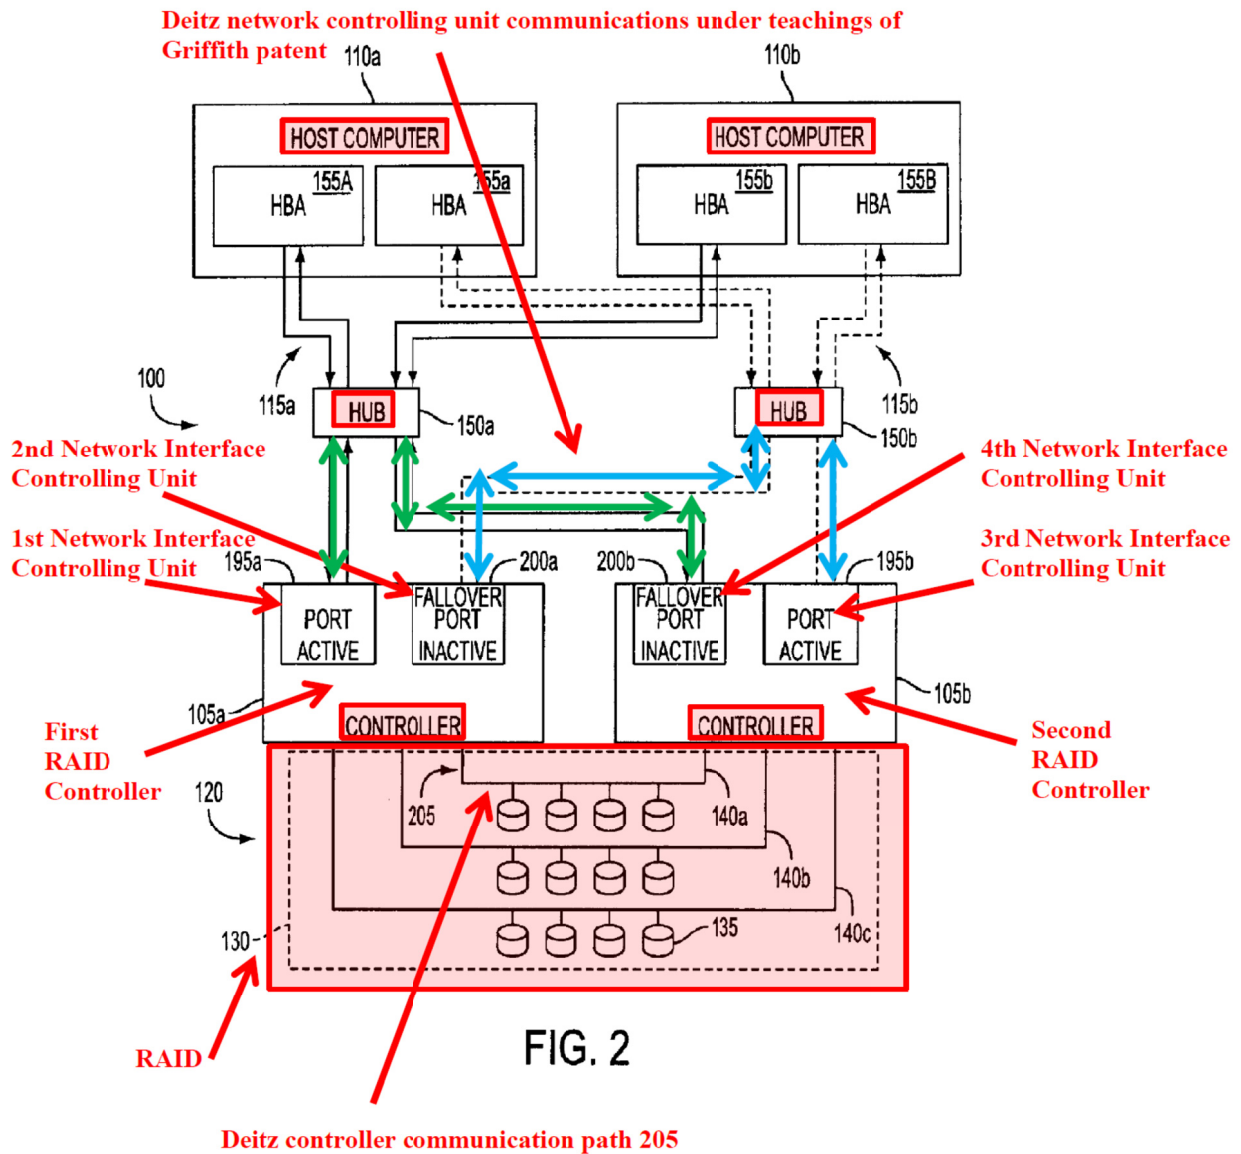
239. Therefore, the Deitz patent discloses “first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers,” as claimed by the ’346 patent.

1c) the first RAID controlling unit including a first network controlling unit and a second network controlling unit and the second RAID controlling unit including a third network controlling unit and a fourth network controlling unit

240. The Deitz patent discloses that each of the RAID controllers has an active and an inactive port, and therefore discloses this claim element.

241. Specifically, the Deitz patent teaches that, “as shown in FIG. 2, each of the controllers 105 have at least one active port 195a, 195b and one inactive port 200a, 200b. The active ports 195a, 195b receive and process I/O requests sent by the host computers 110 on the host-side loops 115. The inactive ports 200a, 200b, also known as a failover ports, can process I/O requests only when the active port 195a, 195b on the same host-side loop 115a, 115b, has failed.” (*Id.* at 6:42-49; *see also id.* at 3:63-66; 5:37-45; 7:52-58; Claim 5; Claim 14; Claim 20; Claim 21.)

242. The active and failover ports disclosed by the Deitz patent are the network interface controllers disclosed in the '346 patent. This is shown by the annotated version below of Figure 2 of the Deitz patent, where the active and failover ports 195a, 195b, 200a and 200b correspond to the network interface controllers 470, 471, 480 and 481 of Figure 4 of the '346 patent:




243. Therefore, the Deitz patent discloses that the first RAID controlling unit includes a first and second network interface controlling unit and the second RAID controlling unit includes a third and fourth network interface controlling unit, as claimed by the '346 patent.

Id) a **plurality of connection units** for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers

244. The Deitz patent discloses a plurality of hubs for connecting the host computers to the RAID controllers, and therefore discloses this claim element.

245. For example, the Deitz patent discloses that the “host-side loops 115 are made up of several fibre channels 145 and a hub 150a, 150b.... Each of the host-side loops 115 connect to three nodes or ports, including a single server port known as a host bus adapter HBA 155a, 155b, on the host computer 110 and to two controller ports 160a, 160b, on each of the controllers 105. The host-side loops 115 are adapted to enable data and input/output (I/O) requests from the host computer 110 to be transferred between any port on the loop 115.” (*Id.* at 5:17-32; *see also id.* at Figure 2.)

246. Figure 2 of the Deitz patent, an annotated version of which appears below, also shows the hubs 150a and 150b connecting the first and second RAID controllers to the host computers 110a and 110b:



247. Therefore, the Deitz patent discloses “a plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers,” as claimed by the ’346 patent.

*1e) 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***

248. The Deitz patent discloses this claim element. For example, the Deitz patent discloses a data exchange between the host computers and the RAID controllers, stating that “host-side loops 115 are adapted to enable data and

input/output (I/O) requests from the host computer 110 to be transferred between any port on the loop 115.” (*Id.* at 5:29-32.)

249. Figure 2 of the Deitz patent, reproduced above, illustrates this and shows that the information exchange occurs between the controllers 105a and 105b and the hosts 110a and 110b through the hubs 150a and 150b.

250. Therefore, the Deitz patent 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,” as claimed by the ’346 patent.

*1f) and 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*****

251. The Deitz patent discloses the use of a direct communication path between RAID controllers for exchanging failover information.

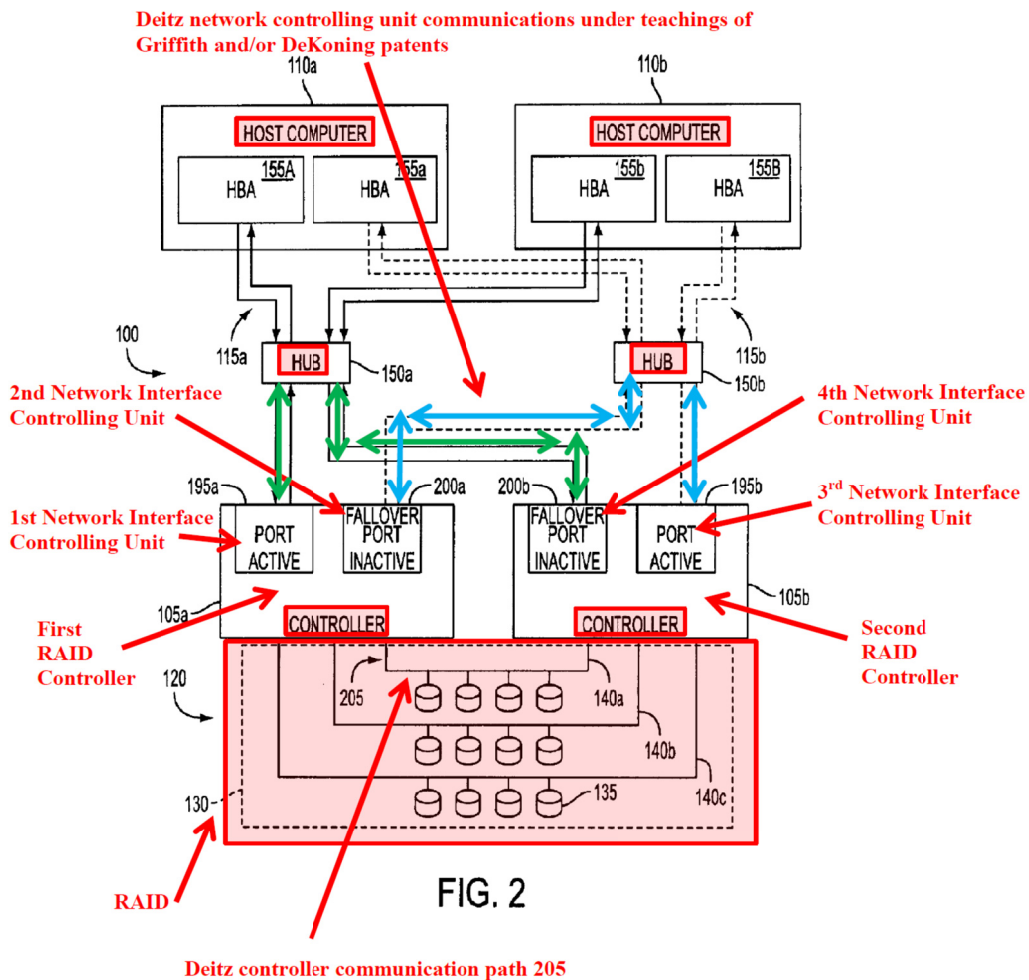
252. For example, the Deitz patent discloses that “[t]he signal passed between the controllers 105 to indicate controller failure can be a passive signal, such as for example the lack of a proper response to a polling or pinging scheme in which each controller interrogates the other at regular, frequent intervals to ensure the other controller is operating correctly. Alternatively, the signal can be a dynamic signal transmitted directly from a failed or failing controller 105a, 105b to the surviving controller 105b, 105a instructing it to initiate a failover process or

mechanism. Optionally, the communication path 205 is also adapted to enable the controllers 105 to achieve cache coherency in case of controller failure.” (*Id.* at 6:63-7:7; *see also id.* at 6:54-63; 7:35-39; 9:44-47; Claim 11.)

253. Furthermore, Figure 2 of the Deitz patent shows a potential communication path for such signals, where the failover port 200b is in connection with the active port 195a through the hub 150a, and the active port 195b is in connection with the failover port 200a through the hub 150b. Such pathways may enable these signals to be passed between the controllers 105a and 105b without the need for additional or dedicated pathways for the transmission of such signals.

254. One of ordinary skill in the art would understand that the system disclosed in the Deitz patent in combination with the teachings of the Griffith and/or DeKoning patents renders this claim element obvious. For example, the Griffith patent teaches that failover communications can (i) be exchanged between RAID controllers over a direct communication path, as disclosed by the Deitz patent, **or** (ii) by using the existing switch network connecting network interface controlling unit ports, without a direct communication path between controllers. (*See, e.g.*, Griffith patent at 9:15-21.) Additionally, the DeKoning patent teaches that several communication mediums can be used to exchange information between RAID controllers, including using the existing host-side communication bus. (*See, e.g.*, DeKoning patent at 4:58-5:10.) As such, one of ordinary skill

would find it obvious to modify the system disclosed in Deitz's Figure 2 with the teachings of the Griffith patent to configure the claimed RAID system as shown in the annotated version of the Deitz patent's Figure 2 below.



255. Therefore, the Deitz patent in combination with the teachings of the Griffith and/or DeKoning patents renders this claim element obvious.

2. Claim 2

2a) *The apparatus as recited in claim 1, wherein said respective **RAID controlling units are connected to the plurality of individual connecting units***

256. The Deitz patent discloses the RAID controllers communicating with the host computers through hubs, and therefore discloses this claim.

257. For example, the Deitz patent discloses that “host-side loops 115 are adapted to enable data and input/output (I/O) requests from the host computer 110 to be transferred between any port on the loop 115.” (*Id.* at 5:29-32.)

258. Figure 2 of the Deitz patent further discloses this claim. Specifically, Deitz Figure 2 shows the controllers 105a and 105b connected to the hubs 150a and 150b:





FIG. 2

3. Claim 3

*3a) 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***

259. The Deitz patent discloses this claim element. However, I note that in describing a real world system, the language requiring a “connecting unit of one side” is not specific. I also note that the ’346 patent does not use the term “side” except within Claim 3. However, for purposes of this discussion, and under the broadest reasonable interpretation standard, I assume that the “sides” being referred to in Claim 3 are the sides of the figures included in the ’346 patent, and as such, I refer below to the sides of the figures in the Deitz patent in my analysis of this claim element.

260. Specifically, the Deitz patent discloses active port 195a connected to the hub on the left side of Figure 2, and failover port 200a connected to the hub on the right side of Figure 2.



4. Claim 4

*4a) 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***

261. The Deitz patent discloses this claim element.

262. As explained above with respect to claim 1, ports 195a, 195b, 200a and 200b in Deitz Figure 2 correspond to the first, second, third, and fourth network interface controlling units of the '346 patent. This is shown below in an annotated version of Deitz Figure 2:




FIG. 2

263. The Deitz patent further discloses that active ports 195a and 195b correspond to the '346 patent's first and third network interface controlling units and process I/O requests of host computers 110a and 110b. For example, the Deitz patent discloses that "as shown in FIG. 2, each of the controllers 105 have at least one active port 195a, 195b and one inactive port 200a, 200b. The active ports 195a, 195b receive and process I/O requests sent by the host computers 110 on the host-side loops 115... The inactive ports 200a, 200b, also known as a failover ports,

can process I/O requests only when the active port 195a, 195b on the same host-side loop 115a, 115b, has failed.” (*Id.* at 6:41-49; *see also id.* at 7:32-35.)

264. Therefore, the Deitz patent teaches that “the first network interface controlling unit and the third network interface controlling unit process the requirement of the numerous host computers,” as claimed by the ’346 patent.

*4b) the **second** network interface controlling unit and the **fourth** network interface 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 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***

265. The Deitz patent in view of the teachings of the Griffith and/or DeKoning patents renders this claim element obvious.

266. Specifically, the Deitz patent in view of the teachings of the Griffith and/or DeKoning patents renders obvious the network interface controlling units on primary and secondary RAID controllers exchanging failover information, as explained above in my discussion of claim 1(f).

267. Additionally, the Deitz patent teaches that active ports on a first RAID controller can fail over to inactive ports on a second RAID controller when the first RAID controller is faulty. (*See, e.g., id.* at Figure 2.) The Deitz patent discloses such communication between RAID controllers when it states, for example, that

“[t]he signal passed between the controllers 105 to indicate controller failure can be a passive signal, such as for example the lack of a proper response to a polling or pinging scheme in which each controller interrogates the other at regular, frequent intervals to ensure the other controller is operating correctly.

Alternatively, the signal can be a dynamic signal transmitted directly from a failed or failing controller 105a, 105b to the surviving controller 105b, 105a instructing it to initiate a failover process or mechanism.” (*Id.* at 6:63-7:7; *see also id.* at 6:54-63; 7:35-39; 9:44-47; Claim 11.)

268. The Deitz patent further discloses that when one of the two RAID controllers experiences a fault, a first and second port of the remaining (e.g., non-faulty) RAID controller executes a function of both RAID controllers and their ports. For example, the Deitz patent discloses that “[a] failover unit is adapted to enable a surviving controller to respond to instructions addressed to it and to instructions addressed to the failed controller.... The failover unit also includes a loop initialization unit, which is adapted to instruct a surviving controller to assume the identity of the failed controller and to instruct the surviving controller to respond to instructions addressed to it and to the failed controller as well as instructions addressed to the surviving controller.” (*Id.* at 3:52-61.) The Deitz patent also discloses that the “firmware of each controller is modified to support the failover and a failback mechanism of the present invention.” (*Id.* at 5:55-57.)

The Deitz patent further discloses how such operations are performed when the system is in a failover mode, disclosing, for example, that “[i]n the assuming identity step 240, the failover port 200a, 200b of the surviving controller 105a, 105b, begins accepting and processing I/O requests addressed by the host computers 110a, 110b, to the failed controller 105b, 105a.” (*Id.* at 7:55-58.)

269. Additionally, in view of the teachings of the Griffith patent that each storage controller port can exchange fault tolerance information across the existing switch network, or the teachings of the DeKoning patent that that several communication mediums can be used to exchange information between RAID controllers, including existing host-side communication buses, one of ordinary skill in the art would understand that the second and fourth network interface controlling units in the Deitz patent could be configured to send failover monitoring communications when the RAID controllers are not faulty. For example, the Griffith patent discloses that each of the RAID controllers in a redundant RAID system is able to communicate over a switch or hub. (*See* Griffith patent at 9:16-22 (“In an alternative configuration , the heartbeat is emitted to the loop connecting means 22 which might be a fiber channel loop. If such a loop is used, the heartbeat signal would have to follow the established priority and arbitration procedures for use of the loop, a requirement which is avoided by use of a direct connections...”); *see also id.* at 8:22-26.) The Griffith patent discloses that

the “loop connecting means 22” can be implemented as a “fibre channel arbitrated loop” (which one of ordinary skill would understand is implemented with a hub) or a “switch fabric device.” (*Id.* at 8:25-26). Additionally, the Griffith patent teaches that a switch can connect to all RAID controller ports in the system. (*Id.* at Figs. 4-5.) Therefore, combining the disclosure of the Griffith patent with the disclosure of the Deitz patent would result in a configuration similar to that of Figure 6 of the ’346 patent, in which the second and fourth network interface controlling units in the Deitz patent communicate failover messages through the hubs 150a and 150b, where each network interface controlling unit is connected to all other RAID controller ports, when the RAID controllers 105a and 105b are not faulty. This configuration would be obvious to one of ordinary skill in the art, *e.g.*, to avoid sending monitoring communications over the active first and third network interface controlling units, thus reducing the performance impact on the active ports. Additionally, as shown in Figure 2 of the Deitz patent, the second and fourth network interface controlling units are configured to execute a function of the first and third network interface controlling units when one of the RAID controlling units experiences a fault.

270. I note that for purposes of my discussion of this claim element, I am interpreting the claim language to recite “... the second network interface controlling unit ~~and~~ or the fourth network interface controlling unit are used for

executing a function of the first network interface controlling unit ~~and~~ or the third network controlling unit when one of the first RAID controlling unit and the second RAID controlling unit is faulty.” I use this interpretation because claim 1 requires that the second and fourth network interface controlling units are on different RAID controllers, and as such, if one RAID controller is faulty, either the second or the fourth network interface controlling unit will not be used to execute any function. I reserve the right to offer a different interpretation in the district court litigation, which I have been informed uses a different legal standard.

5. Claim 5

5a) The apparatus as recited in claim 1, wherein said plurality of connecting units have at least three connection ports

271. The Deitz patent discloses hubs with at least three connection ports, and therefore discloses this claim element.

272. For example, Figure 2 of the Deitz patent discloses that each of the hubs 150a and 150b has connections to two host computers 110a and 110b, as well as connections to the two controllers 105a and 105b. Thus, each of the hubs 150a and 150b has at least four connection ports, as shown below in an annotated version of Deitz Figure 2.

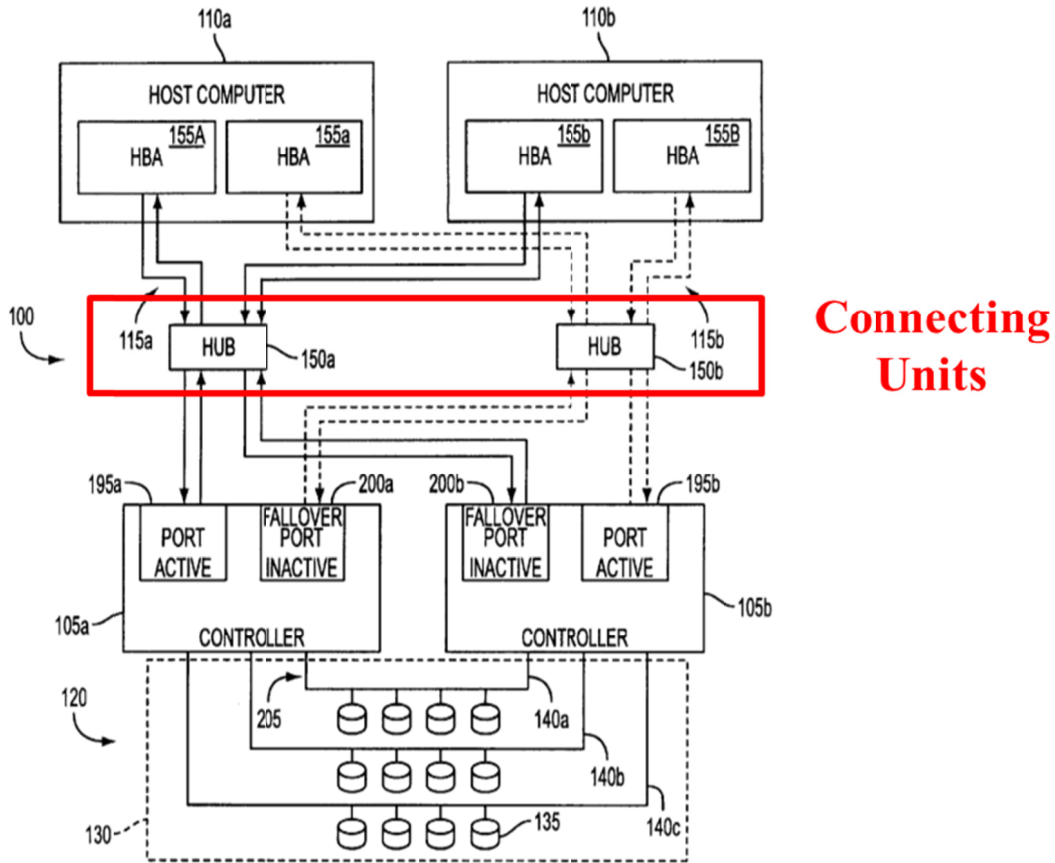


FIG. 2

5b) *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*

273. The Deitz patent discloses that two of the connection ports of the hubs are connected to one of the first network interface controlling unit and the third network controlling unit in the controllers 105a, 105b, and therefore discloses this claim element.

274. For example, Figure 2 of the Deitz patent shows connections from the hubs 150a and 150b to each of the ports 195a, 195b, 200a and 200b (*i.e.*, network controlling units) of the RAID controllers 105a and 105b. Additionally, one of

ordinary skill in the art would understand that in hubs, every port is connected to every other port.⁷ As such, under the broadest reasonable interpretation standard, the Deitz patent discloses a system that meets this claim element because all ports on the hubs' network controlling units are connected to every other network controlling unit and the hosts. Since the Deitz patent discloses connections to each of these ports 195a, 195b, 200a and 200b, the Deitz patent discloses at least two couplings from the ports of the connecting units to the “first network controlling unit and the third network controlling unit.” (See also *id.* at 5:17-29.)




FIG. 2

⁷ See above, ¶103.

5c) and the rest of the connection ports being provided as a hub equipment connected with the numerous host computers

275. The Deitz patent discloses this claim element.

276. As explained above with respect to claim element 5(b), each of the ports of a hub is coupled to every other port of the hub. As such, the Deitz patent discloses that all of the connection ports on hubs 150a and 150b (shown in Deitz Figure 2) are connected to the host computers 110a and 110b.




FIG. 2

6. Claim 6

6a) The apparatus as recited in claim 1, wherein said plurality of connecting units have at least three connection ports

277. This claim element is identical to claim element 5(a), and therefore my analysis as to how the Deitz patent discloses this claim element is identical.

6b) two of the at least three connection port are coupled to one of the first network controlling unit and the third network controlling unit

278. This claim element is identical to claim element 5(b), and therefore my analysis as to how the Deitz patent discloses this claim element is identical.

6c) and the rest of the connection ports being provided as a network switch equipment connected with the numerous host computers

279. The Deitz patent in view of the teachings of the Griffith patent discloses this claim element.

280. As explained in my discussion of claim element 5(c), Figure 2 of the Deitz patent discloses two hubs 150a and 150b, and one of ordinary skill in the art would understand that all ports on these hubs are connected to every other port. The Griffith patent teaches that RAID controlling units, e.g., 105a and 105b in Figure 2 of the Deitz patent, can be connected to host computers via hubs or switches. (See Griffith patent at 8:25-26 (“Loop connecting means may be a SCSI bus, fibre channel arbitrated loop, or a switch fabric device.”); see also *id.* at Figs. 4 and 5 (showing loops and/or switches being connected to hosts and RAID

controllers).) It is also my opinion that one of ordinary skill in the art would understand that every port on a switch can communicate with and is thus connected to every other port on a switch. As such, the system disclosed in the Deitz patent in view of the teachings of the Griffith patent renders this claim element obvious.

7. Claim 7

7a) The apparatus as recited in claim 1, wherein said plurality of connecting units have at least five connection ports

281. To the extent this claim element is read to require the “plurality of connecting units” to collectively have at least five connection ports, Deitz discloses this element. (See Deitz Fig. 2 (disclosing two hubs, each with four connection ports, for a collective total of eight connection ports).) Otherwise, the Deitz patent does not include a figure showing five or more connections to each hub, but the specification many times refers to “at least one host” – six times in the specification, and eight times in the claims. (Deitz patent at Abstract; *see also id.* at 3:7-10; 3:26-29; 3:43-49; 4:5-8; 4:44-48; Claim 1; Claim 4; Claim 5; Claim 6; Claim 7; Claim 9; Claim 10; Claim 13; Claim 14.) The Deitz patent also refers to the “host-side loops” connecting the hosts, disclosing that “host-side loops 115 are made up of several fibre channels 145 and a hub 150a, 150b.” (*Id.* at 5:18-19.) As such, one of ordinary skill in the art would have recognized that the Deitz patent does not restrict the number of hosts on the host-side loop, and thus would have known to configure systems with three or more hosts connected to the two RAID

controllers using connecting units consisting of hubs and/or switches with at least five ports.

282. Additionally, the Deitz patent in view of the teachings of the Griffith patent renders this claim element obvious. For example, as explained in my discussion of claim element 5(a), Figure 2 of the Deitz patent discloses hubs with four connection ports. The Griffith patent, e.g., at Figures 4 and 5, teaches that hubs or switches with at least twelve connection ports can be used in redundant RAID systems. Using multi-ported switches was well known to one of ordinary skill in the art at the time of the alleged invention. As such, combining the teachings of the Griffith patent with the system of the Deitz patent renders this claim element obvious.

7b) 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

283. The Deitz patent alone discloses this element, or in view of the multi-ported hub or switch teachings in the Griffith patent, renders this claim element obvious. For example, in my discussion of claim 5(b) above, I explained how one of ordinary skill in the art would understand that all ports on a hub or switch can communicate with all other ports on the same hub or switch. Therefore, one of ordinary skill in the art would understand that the system disclosed in Figure 2 of the Deitz patent in combination with the multi-ported teachings of the Griffith

patent would render obvious a system as claimed by the '346 patent, in which four of the at least five connection ports are coupled to one of the first network interface controlling unit and the third network controlling unit.

7c) and the rest of the connection ports being provided as a switch connected with the numerous host computers

284. The Deitz patent in combination with the teachings of the Griffith patent discloses this claim element.

285. As explained with respect to claim element 7(b), one of ordinary skill in the art would understand that all ports of a network switch or hub are coupled to each of the other ports of the switch or hub. The use of a network switch in lieu of a hub is a design decision that would be understood by one of ordinary skill in the art at the time of the '346 patent, as shown by the teachings of Figures 2 and 3 of the Griffith patent. See also my discussion of claim element 6(c).

8. Claim 8

8a) 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

286. The Deitz patent discloses a port of a first controller and a port of a second controller both connected to one hub, and another port of the first controller

and another port of the second controller both connected to another hub. The Deitz patent therefore discloses this claim element.

287. This is shown in Figure 2 of the Deitz patent, for example, where hub 150a has connections to both the active port 195a of controller 105a and to failover port 200b of controller 105b. Figure 2 of the Deitz patent also shows hub 150b connected to both failover port 200a and active port 195b.




FIG. 2

9. Claim 9

*9a) An apparatus for a **redundant interconnection** between **multiple host** computers and a **RAID**, the apparatus comprising*

288. This claim element is identical to the corresponding element in claim 1. Thus, as discussed above, the Deitz patent discloses “an apparatus for a redundant interconnection between multiple hosts and a RAID,” as claimed by the ’346 patent.

*9b) a **plurality of connecting units** for connecting the host computers and the RAID;*

289. This claim element is similar to an element in claim 1 which recites: “a plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers.” Thus, as discussed above, the Deitz patent discloses “a plurality of connecting units for connecting the host computers and the RAID,” as claimed by the ’346 patent.

*9c) a **first and a second RAID controllers**, included in the RAID, 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***

290. This claim element is similar to language in claim 1 which recites: “a first RAID controlling units and a second RAID controlling unit for processing a requirement of numerous host computers, the first RAID controlling unit including a first network controlling unit and a second network controlling unit, and the second RAID controlling unit including a third network controlling unit and a

fourth network controlling unit; and a plurality of connection units for connecting the first RAID controlling units and the second RAID controlling unit to the numerous host computers.”

291. Thus, as discussed above, the Deitz patent discloses first and second RAID controllers, included in the RAID, each of which has a first network controlling unit and a second network controlling unit for processing requests from the plurality of the host computers connected through the plurality of the connection units, as claimed by the '346 patent.

*9d) 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 processes information transmitted from the second network interface controller** in the second RAID controller*

292. The Deitz patent in view of the teachings of the Griffith patent renders this claim element obvious.

293. For example, the Deitz patent discloses that a first port of a first RAID controller supplies data to the host computers through hubs. Specifically, the Deitz patent discloses that, “as shown in FIG. 2, each of the controllers 105 have at least one active port 195a, 195b and one inactive port 200a, 200b. The active ports 195a, 195b receive and process I/O requests sent by the host computers 110 on the host-side loops 115.” (*Id.* at 6:42-45; *see also id.* at 4:64-5:5; 5:37-46; Claim 5; Claim 14; Claim 20; Claim 21.) As explained above with respect to claim element

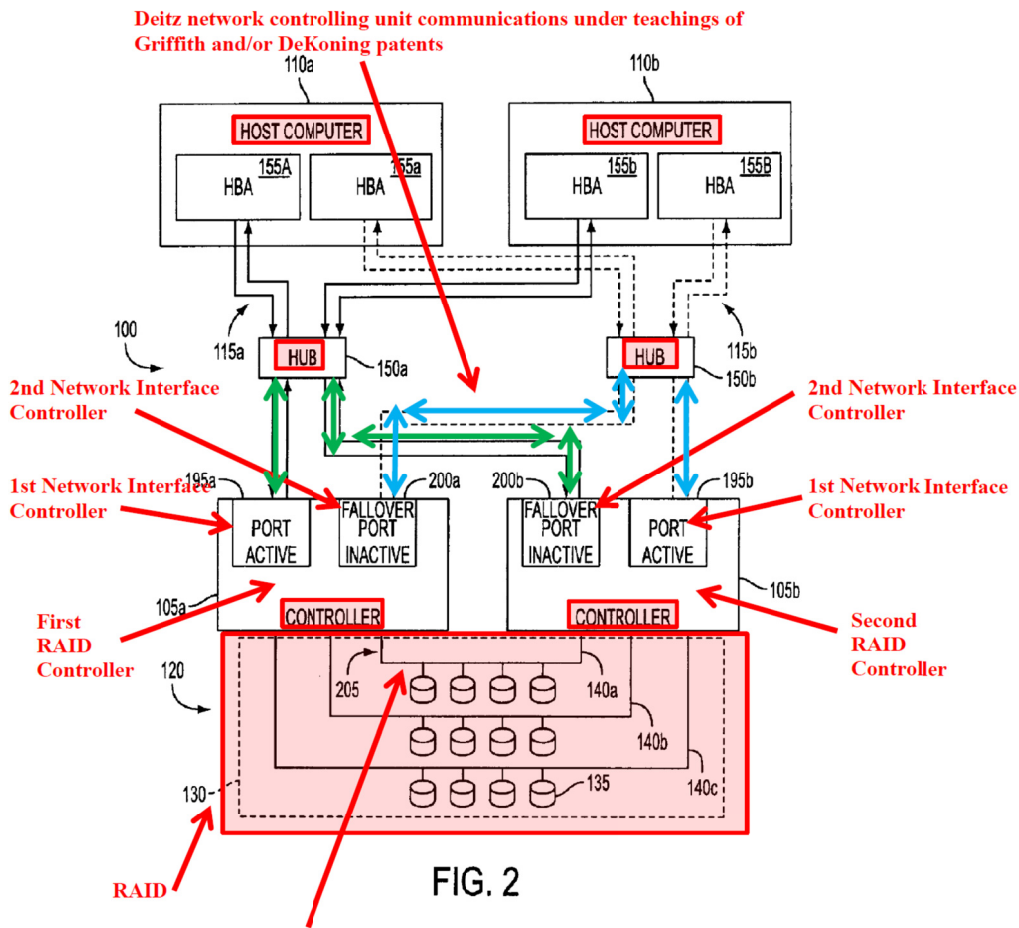
9(c), the ports 195a, 195b, 200a and 200b of the controllers 105a and 105b are connected to the host computers through the hubs 150a and 150b.

294. The Deitz patent further discloses communication between a port of a RAID controller and a host computer when discussing the use of dual-active mode. For example, the Deitz patent discloses that “the memory system 100 is then ready to begin regular operations in a dual-active operation step 225 in which the controllers 105 both simultaneously receive and process I/O requests from the host computers 110.” (*Id.* at 7:31-35; *see also id.* at 3:29-30; 5:41-48; 5:58-61.)

295. Additionally, the Deitz patent discloses signals transmitted between the RAID controllers for performing failover and failback operations. The Deitz patent discloses such communication between the RAID controllers when stating, for instance, that “[t]he signal passed between the controllers 105 to indicate controller failure can be a passive signal, such as for example the lack of a proper response to a polling or pinging scheme in which each controller interrogates the other at regular, frequent intervals to ensure the other controller is operating correctly. Alternatively, the signal can be a dynamic signal transmitted directly from a failed or failing controller 105a, 105b to the surviving controller 105b, 105a instructing it to initiate a failover process or mechanism. Optionally, the communication path 205 is also adapted to enable the controllers 105 to achieve

cache coherency in case of controller failure.” (*Id.* at 6:63-7:7; *see also id.* at 6:54-63; 7:35-39; 9:44-47; Claim 11.)

296. The Griffith patent teaches that storage controller ports can exchange information without the need of a direct path between RAID controllers as disclosed in the Deitz patent. (Griffith patent at 9:16-22; *see also id.* 7:47-52.). As such, the Deitz patent in view of the teachings of the Griffith patent renders obvious the claimed system 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 processes information transmitted from the second network interface controller in the second RAID controller. This is shown by the annotated version below of Figure 2 of the Deitz patent, in which port 195a on the first RAID controller 105a is active and used to supply data to the hosts, and the same port 195a is used to process failover communications with the inactive port 200 on the second RAID controller 105b:



9e) 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 processes information transmitted from the second network interface controller** in the **first RAID controller**

297. The Deitz patent in view of the teachings of the Griffith and/or DeKoning patents renders this claim element obvious. This is explained above in my discussion of claim element 9(d), which notes that a system as claimed where both RAID controllers can supply data to either host and any controller port can communicate with any other controller port would be obvious. This is also shown

by the annotated version of Deitz Figure 2 above, in which the port 195b on the second RAID controller 105b is active and used to supply data to the hosts, and the same port 195b is used to process failover communications with the inactive port 200a on the first RAID controller 105a.

*9f) 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*

298. The Deitz patent discloses this claim element.

299. As explained in my discussion of claim element 4(b), the Deitz patent discloses failover processes that enable a failover port in a remaining (e.g., non-faulty) RAID controller to perform functions of an active port in a failed RAID controller. For example, the Deitz patent discloses that “[o]n detection of a controller failure, a failover procedure is performed on the surviving controller 105a, 105b, the failover procedure involves the steps of disabling the failed controller (step 235) and assuming the identity of the failed controller (step 240). In the disabling step 235, the surviving controller 105a, 105b asserts a reset signal, which disables the failed controller 105b, 105a by resetting its local processor 185a, 185b, and the active port 195a, 195b, fibre protocol chip (not shown). Resetting the fibre protocol chip causes the hub 150a, 150b to automatically bypass the primary port 195a, 195b, on the failed controller 105a, 105b. In the assuming identity step 240, the failover port 200a, 200b of the surviving controller 105a,

105b, begins accepting and processing I/O requests addressed by the host computers 110a, 110b, to the failed controller 105b, 105a.” (*Id.* at 7:44-58; *see also id.* at 7:24-31; Claim 10; Claim 14.)

300. The Deitz patent further discloses failover operations with respect to Figure 2 when it states that “[t]he inactive ports 200a, 200b, also known as a failover ports, can process I/O requests only when the active port 195a, 195b on the same host-side loop 115a, 115b, has failed. For example, in case of failure of controller 105a, inactive port 200b on surviving controller 105b assumes the identity of the active port 195a on failed controller 105a and begins accepting and processing I/O requests directed to the failed controller 105a.” (*Id.* at 6:44-53.)

301. Furthermore, Figure 2 of the Deitz patent, reproduced below, illustrates that the storage ports labeled “failover port inactive” (which I understand to mean “failover port inactive”) on both RAID controllers are being used for fault tolerance by performing functions of the “active” ports on the other RAID controllers when they experience a fault.




FIG. 2

9g) 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

302. For similar reasons as those explained with respect to claim element 9(f), the Deitz patent discloses “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,” as claimed by the ’346 patent.

*9h) 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 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*

303. This claim element is similar to language in claim 1 which recites: “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.” Thus, as discussed above, the Deitz patent in view of the teachings of the Griffith and/or DeKoning patents renders this element obvious.

IX. STATEMENT UNDER U.S.C. SECTION 1001 OF TITLE 18

304. I declare under penalty of perjury under the laws of the United States of America that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like are punishable by fine, imprisonment, or both under Section 1001 of Title 18 of the United States Code.

Executed this 3rd day of June, 2014 in San Jose, CA.

Robert Horst

Robert Horst, Ph.D.

