IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re patent of Hoese	§	Petition for *Inter Partes* Review
	§	
U.S. Patent No. 7,934,041	§	Attorney Docket No.: 31046.8
	§	Customer No.: 27683
Issued: April 26, 2011	§	
	§	Real Parties in Interest:
Title: STORAGE ROUTER AND METHOD FOR PROVIDING VIRTUAL LOCAL STORAGE	§	Cisco Systems, Inc.
	§	Quantum Corporation
	§	

<u>Declaration of Andrew Hospodor, Ph.D.</u>
<u>Under 37 C.F.R. § 1.68</u>

Table of Contents

I. Introduction

I, Andrew Hospodor, Ph.D., declare:

1. I am making this declaration at the request of Cisco Systems, Inc. and Quantum Corporation in the matter of the *Inter Partes* Review of U.S. Patent No. 7,934,041 ("the '041 Patent") to Hoese.

2. I am being compensated for my work in this matter. My compensation in no way depends upon the outcome of this proceeding.

3. In the preparation of this declaration, I have studied:

(1) The '041 Patent, CQ-1001;

(2) The prosecution history of the '041 Patent, CQ-1002;

(3) CMD Technology, Inc., *CRD-5500 SCSI Raid Controller User's Manual*, Revision 1.3, November 21, 1996 ("CRD-5500 Manual"), CQ-1004;

(4) *CMD Technology CRD-5500 RAID*, http://web.archive.org/web/19961226091552/http://www.cmd.com/brochure/crd5500.htm, archived December 26, 1996 by archive.org ("CRD-5500 Data Sheet"), CQ-1005;

(5) Hewlett-Packard Journal, Volume 47, Number 5, October 1996 ("HP Journal"), CQ-1006;

(6) American National Standard for Information Systems, *Fibre Channel Physical and Signaling Interface (FC-PH) X3.230*, Rev. 4.3, June 1, 1994 ("ANSI Fibre Channel FC-PH Standard"), CQ-1007; and

(7) Sun Microsystems Computer Company, *SPARCstorage Array Configuration Guide*, Revision A, March 1995 ("SPARCstorage Guide"), CQ-1008;

(8) *SPARCstorage Array – Product Brief*, http://web.archive.org/web/19961220045017/http://www.sun.com/products-n-solutions/hw/peripherals/array.html, archived December 20, 1996 by archive.org ("SPARCstorage Product Brief"), CQ-1009; and

(9) ORDER of November 8, 2011, *Crossroads Systems, Inc. v. 3PAR, Inc*., *et. al*., case no. 1-10-cv-00652 (W.D. Tex. 2010), CQ-1010;

4. In forming the opinions expressed below, I have considered:

(1) The documents listed above,

(2) The relevant legal standards, including the standard for obviousness provided in *KSR International Co. v. Teleflex, Inc.*, 550 U.S. 398 (2007), and

(3) My knowledge and experience based upon my work in this area, as described below.

II. Qualifications and Professional Experience

5. My complete qualifications and professional experience are described in my curriculum vitae, a copy of which is attached as an exhibit to this declaration. Following is a brief summary of my relevant qualifications and professional experience:

6. I received a Bachelor of Science degree in Computer Engineering from Lehigh University in 1981, a Master of Science degree in Computer Science from Santa Clara University in 1986, and a Ph.D. in Computer Engineering from Santa Clara University in 1994. My Ph.D. emphasis was in storage architecture and systems. My dissertation was entitled: "A Study of Prefetch in Caching SCSI Disk Drive Buffers."

7. I have been part of the data storage industry for over 25 years and involved in firmware engineering for disk drive and tape drive controllers, including implementation of command processing, error correction, and buffer management. I have also focused on simulation and implementation of disk and tape drives at Quantum. I have been involved in the architecting of network storage devices that included disk drives, tape drives, network switches, routers and software. I have also been involved in the simulation and implementation of disk interfaces, including ATA, SCSI, and Fibre Channel.

8. I have taught graduate and undergraduate courses at Santa Clara University. After receiving my Master's degree in 1986, I joined the Institute for Information Storage Technology as an Adjunct Lecturer, then later as a Research Fellow. I have taught courses in Computer Architecture, Storage Architecture, Hard Disk and Floppy Disk Controller Design, and Grid Computing. I am currently the Executive Director of the Storage Systems Research Center at University of California, Santa Cruz. Here, I oversee the research of faculty, graduate students, post-doctoral scholars and continue to work with industrial sponsors in the data storage industry as well as the National Science Foundation.

9. I am a named inventor on twelve U.S. patents related to data storage that have been cited as prior art in 183 other patents. I have authored numerous publications in reference journals, industry periodicals, and am often cited by my peers in textbooks and journal publications. I have presented to the American National Standards Institute (ANSI) committee on the Small Computer Systems Interface (SCSI), the National Association of Broadcasters (NAB), the SCSI Forum, the Institute of Electrical and Electronic Engineers (IEEE) Systems Design and Network Conference, and many other storage related conferences

10. In summary, I have a deep familiarity with data storage devices, systems, interfaces, networks, and architectures, and had first-hand experience with these technologies at the relevant time of the '041 Patent invention and before.

III. Level of Ordinary Skill in the Art

11. I am familiar with the knowledge and capabilities possessed by one of ordinary skill in the data storage field in 1997, the year in which the parent patent application of the ’041 was filed. Specifically, my extensive experience (i) in the industry and (ii) with engineers practicing in the industry allowed me to become personally familiar with the level of skill of individuals and the general state of the art. Additionally, I personally possessed the knowledge and capabilities of one of ordinary skill in the data storage field at the relevant time.

12. In my opinion, the level of ordinary skill in the art needed to have the capability of understanding the scientific and engineering principles applicable to the ’041 Patent is (i) a Master of Science (M.S.) degree in Computer Science or Computer Engineering or a Bachelor of Science (B.S.) degree in Computer Engineering or equivalent training, and (ii) at least five years of direct experience in developing data storage technologies. Relevant industry experience would include experience with network-based data storage, including block-level storage protocols, logical addressing, storage virtualization, and access controls. Such skills and experience would have been necessary in order to appreciate what was obvious and/or anticipated in the industry and what a person having ordinary skill in the art would have thought at the time. Unless otherwise stated, my testimony below refers to the knowledge of one of ordinary skill in the data storage field in

1997.

IV. Relevant Legal Standards

13. I have been asked to provide my opinions regarding whether the claims 1-53 of the ’041 Patent are anticipated or would have been obvious to a person having ordinary skill in the art at the time of the alleged invention, in light of the prior art. It is my understanding that, to anticipate a claim under 35 U.S.C. § 102, a reference must teach every element of the claim. Further, it is my understanding that a claimed invention is unpatentable under 35 U.S.C. § 103 if the differences between the invention and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which the subject matter pertains. I also understand that the obviousness analysis takes into account factual inquiries including the level of ordinary skill in the art, the scope and content of the prior art, and the differences between the prior art and the claimed subject matter.

14. It is my understanding, based on my review of *KSR*, that the Supreme Court has recognized several rationales for combining references or modifying a reference to show obviousness of claimed subject matter. Some of these rationales include the following: combining prior art elements according to known methods to yield predictable results; simple substitution of one known element for another to obtain predictable results; use of a known technique to improve a similar device (method, or product) in the same way; applying a known technique to a known device (method, or product) ready for improvement to yield predictable results; choosing from a finite number of identified, predictable solutions, with a reasonable expectation of success; and some teaching, suggestion, or motivation in the prior art that would have led one of ordinary skill to modify the prior art reference or to combine prior art reference teachings to arrive at the claimed invention.

V. The '041 Patent

A. Overview

15. The '041 Patent generally relates to network-based storage and describes a "storage router" that routes storage requests between workstations and storage devices. (CQ-1001, Abstract). The '041 Patent has three independent claims (claims 1, 20, and 37) and a total of 53 claims. Claim 1 provides a basic overview of the teachings of the '041 Patent:

> 1. A storage router for providing virtual local storage on remote storage devices, comprising:
>
> a first controller operable to interface with a first transport medium, wherein the first medium is a serial transport media; and
>
> a processing device coupled to the first controller, wherein the

processing device is configured to:

> maintain a map to allocate storage space on the remote storage devices to devices connected to the first transport medium by associating representations of the devices connected to the first transport medium with representations of storage space on the remote storage devices, wherein each representation of a device connected to the first transport medium is associated with one or more representations of storage space on the remote storage devices;
>
> control access from the devices connected to the first transport medium to the storage space on the remote storage devices in accordance with the map; and
>
> allow access from devices connected to the first transport medium to the remote storage devices using native low level block protocol.

16. As noted by the background section of the ’041 specification, various types of communication links supporting different communication distances were known in the art at the time of the ’041 invention. For example, the ’041 specification notes that communication links based on the Small Computer System Interface (SCSI) standard provide for “relatively short distances” between devices (*e.g.*, less than 25 meters) while communication links based the Fibre Channel standard provide for “large distances” between devices (*e.g.*, more than 10 kilometers). (CQ-1001, 1:51-64). According to the specification, the storage router of the ’041 Patent utilizes the large distances provided for by Fibre Channel to provide workstations with access to “significantly remote” storage devices. (CQ-1001, 2:55-61).

17. Figure 3 of the ’041 Patent, annotated below, illustrates the architecture of the storage network in which the storage router operates. Specifically, in Fig. 3, workstations on the Fibre Channel serial communication link (*i.e.*, transport medium) are connected to one side of the storage router (the “host side”) while storage devices on the SCSI parallel communication bus are connected to the other side of the storage router (the “disk side”). (CQ-1001, 4:27-34). The ’041 specification describes the storage router as “a bridge device that connects a Fiber Channel link directly to a SCSI bus.” (CQ-1001, 5:59-61). The storage router enables the exchange of SCSI commands and data between the workstations and the storage devices. (CQ-1001, 5:59-63). According to the specification, a “SCSI command” is one example of a native low level block protocol command. (CQ-1001, 5:59-63). Additionally, the specification states that Fibre Channel-based workstations on one side of the storage router may communicate with SCSI-based storage devices on the other side of the storage router by encapsulating SCSI commands into Fibre Channel Protocol (FCP) requests. (CQ-1001, 6:56-67). I note that encapsulating low level commands such as SCSI commands inside of a Fibre Channel request was a feature of the Fibre

Channel standard, and was thus well known in the art at the time of the '041 invention. (CQ-1006, pp. 94-95).

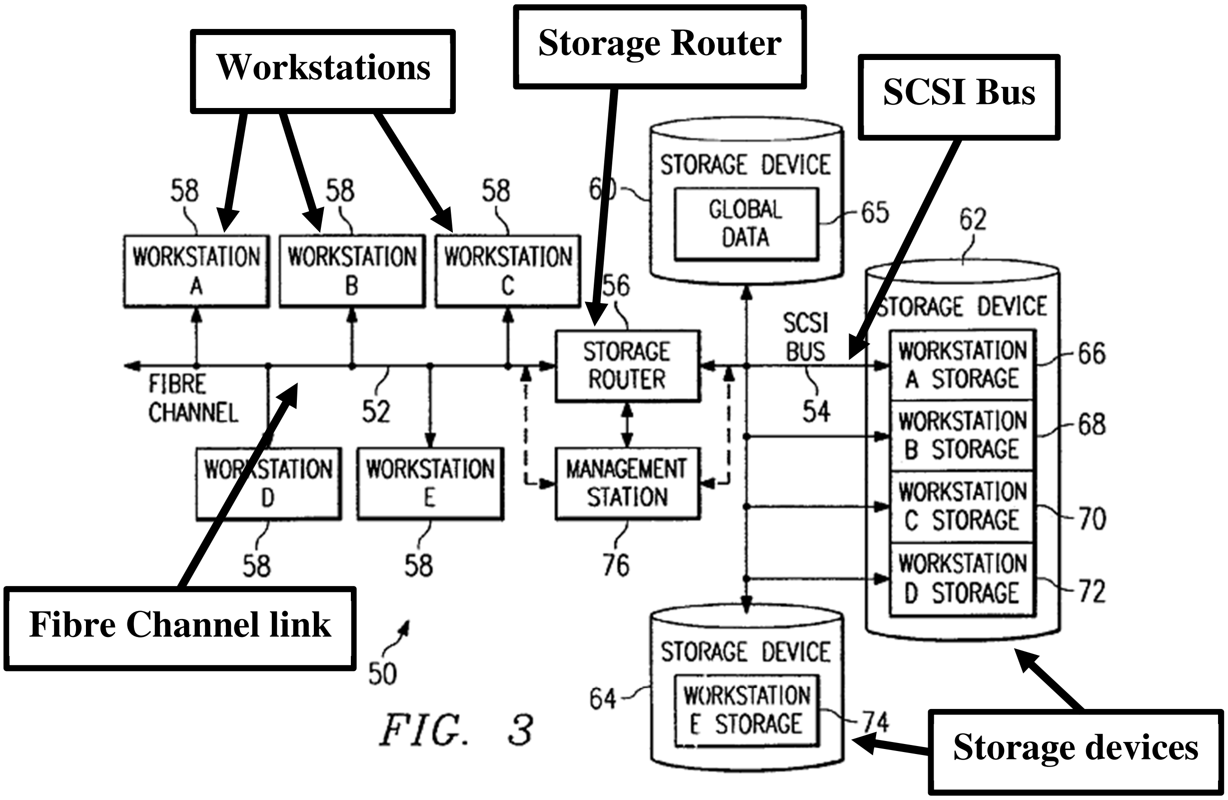

CQ-1001, Fig. 3 (annotated)

18. In addition to the mode of operation shown in association with Fig. 3, specification also notes that the "storage router has various modes of operation." (CQ-1001, 6:43). For example, the storage router may route data between (i) a Fibre Channel host and a SCSI storage device, (ii) a SCSI host and a Fibre Channel storage device, (iii) a SCSI host and SCSI storage device, and (iv) a Fibre Channel host and a Fibre Channel storage device. (CQ-1001, 6:43-46). With regard to the Fibre Channel-to-Fibre Channel mode of operation, the specification states that the storage router may "act as a bridge between two FC loops." (CQ-1001, 6:54-55).

19. The specification further states that the storage router uses "mapping tables" to allocate subsets of storage space (*i.e.*, partitions) on the storage devices to particular workstations. (CQ-1001, 4:41-44). For example, with reference to Fig. 3, "[s]torage device **62** can be configured to provide partitioned subsets **66**, **68**, **70** and **72**, where each partition is allocated to one of the workstations **58** (workstations A, B, C and D)." (CQ-1001, 4:47-50). Further, the '041 specification states that the storage router provides "virtual local storage" such that a partition mapped to a workstations 58 is "considered by the workstation 58 to be its local storage"—*i.e.*, the mapped partition "has the appearance and characteristics of local storage." (CQ-1001, 4:35-41, 4:4-14). As discussed below in more detail, it was well known in the art at the time of the '041 invention to map workstations on one side of a storage router to partitions on the other side of the storage router, and to make the partitions appear local. (CQ-1004, p. 1-2, 4-5).

20. According to the '041 specification, the storage router uses the mapping functionality to facilitate both routing and access control. (CQ-1001, 5:52-54). With respect to routing, the specification states that the map between the initiators and the specific subsets of storage allows the storage router to determine "what partition is being addressed by a particular request," thus enabling it to "distribute[] requests and data" to storage devices. (CQ-1001, 9:21-24, 4:16-20).

With respect to access control, the specification states that the storage router prevents a workstation from accessing a subset of storage not allocated to that workstation in the map. (CQ-1001, 9:15-27). For example, in Fig. 3, "subsets 66, 68, 70 and 72 can only be accessed by the associated workstation 58." (CQ-1001, 4:50-51). I have annotated Fig. 3 below to illustrate the general flow of I/O commands in the storage network of the '041 Patent. As discussed below in more detail, the concept of using a map to facilitate routing to and access control of storage devices was well known in the art at the time of the '041 invention. (CQ-1004, p. 1-2, 4-5).

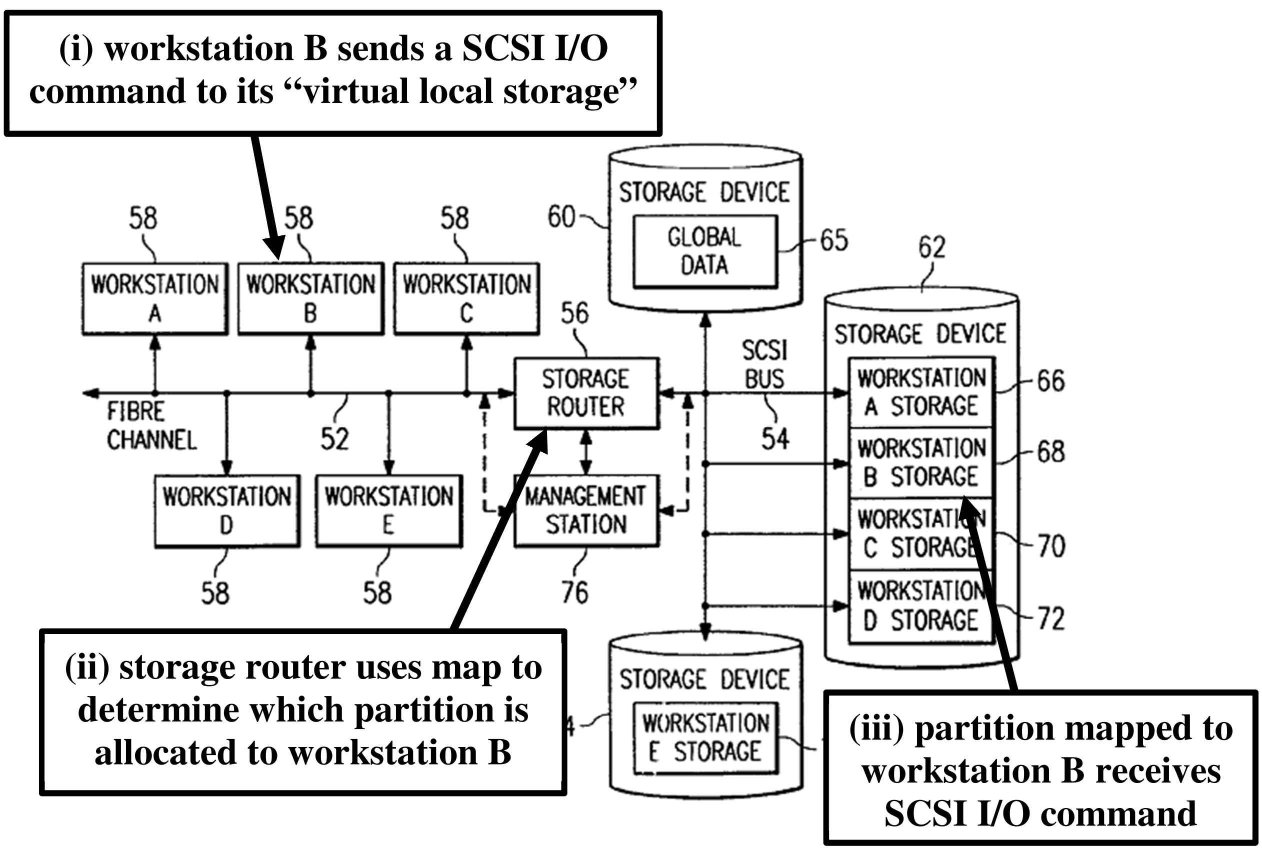

CQ-1001, Fig. 3 (annotated)

21. Because routing SCSI commands over Fibre Channel, mapping workstations to storage partitions, and using such a mapping for routing and access control were all well known in the art at the time of the ’041 invention, it is my opinion that the storage network described by the ’041 Patent is simply a collection of components that were well known in the art at the time of the ’041 invention. And, as shown below, these well-known components are arranged in a manner that would have been obvious to one of ordinary skill in the art.

B. History of the ’041 Patent

22. The ’041 Patent issued on April 26, 2011, from U.S. Patent Application No. 12/690,592 (“the ’592 application”) filed on January 20, 2010, by Geoffrey B. Hoese and Jeffry T. Russell. The ’041 Patent claims priority back to U.S. Patent No. 5,941,972 (“the ’972 Patent”), filed on December 31, 1997, via a long string of intervening continuation applications.

23. Based on my review of the prosecution history of the ’592 application, it appears the Patent Office did not substantively consider the CRD-5500 Manual, the HP Journal, or the ANSI Fibre Channel FC-PH Standard in view of the claims of the ’041 Patent. (*See* CQ-1002).

VI. Claim Construction

24. It is my understanding that in order to properly evaluate the ’041

Patent, the terms of the claims must first be interpreted. It is my understanding that for the purposes of this *inter partes* review the claims are to be given their broadest reasonable interpretation in light of the specification. It is my further understanding that claim terms are given their ordinary and accustomed meaning as would be understood by one of ordinary skill in the art, unless the inventor has set forth a special meaning for a term. As such, any claim term not construed below should be given its ordinary and customary meaning.

25. In order to construe the following claim terms, I have reviewed the entirety of the '041 Patent, as well as its prosecution history. I have also reviewed the constructions given to the claim terms in previous litigations involving patents related to the '041 Patent. In particular, I reviewed the claim construction order from the *Crossroads v. 3Par* district court case[1], that construed claim terms recited in U.S. Patent No. 6,425,035. I note that the '035 Patent is a family member of the '041 Patent and recites similar claim language. It is my understanding that the constructions given to the claim terms for the purposes of this *inter partes* review may be broader than or equal in scope to the constructions given by the District Court. In any case, I believe the following constructions are consistent with the

---

[1] *Crossroads Systems, Inc. v. 3PAR, Inc., et. al.*, case no. 1-10-cv-00652 (W.D. Tex. 2010).

District Court's constructions.

26. The following table summarizes my claim constructions. My analysis of each claim term follows.

Claim Term	Claim Construction
"native low level block protocol"	*A protocol in which storage space is accessed at the block level, such as the SCSI protocol*
"remote"	*Indirectly connected through a storage router to enable connections to storage devices at a distance greater than allowed by a conventional parallel network interconnect*

A. "native low level block protocol"

27. The claim term "native low level block protocol" is found in claims 1, 6, 20, 25, 37, and 42.

28. I note that in the *3Par* case, the District Court construed "native low level block protocol" to mean "a set of rules or standards that enable computers to exchange information and do not involve the overhead of high level protocols and file systems typically required by network servers." (CQ-1010, p. 13).

29. Based on my review of the '041 Patent, it appears that the specification does not provide an explicit definition for "native low level block protocol." I note, however, that the '041 specification discloses—with reference to Fig. 1—that:

> In network 10, each workstation 12 has access to its local storage

device as well as network access to data on storage devices 20. The access to a local storage device is typically through native low level, block protocols. On the other hand, access by a workstation 12 to storage devices 20 requires the participation of network server 14 which implements a file system and transfers data to workstations 12 only through high level file system protocols. Only network server 14 communicates with storage devices 20 via native low level, block protocols. (CQ-1001, 3:43-52, emphasis added).

Similarly, with reference to Fig. 3, the specification states that:

> Storage device **62** can be configured to provide partitioned subsets **66**, **68**, **70** and **72**, where each partition is allocated to one of the workstations **58** (workstations A, B, C and D). These subsets **66**, **68**, **70** and **72** can only be accessed by the associated workstation **58** and appear to the associated workstation **58** as local storage accessed using native low level, block protocols. (CQ-1001, 4:47-53, emphasis added).

> This is accomplished without limiting the performance of workstations 58 because storage access involves native low level, block protocols and does not involve the overhead of high level protocols and file systems required by network servers. (CQ-1001, 5:29-33, emphasis added).

Further, the specification states that:

> The storage router of the present invention is a bridge device that connects a Fiber Channel link directly to a SCSI bus and enables the exchange of SCSI command set information between

application clients on SCSI bus devices and the Fiber Channel links. (CQ-1001, 5:59-63, emphasis added).

> The FC Initiator to SCSI Target mode provides for the basic configuration of a server using Fiber Channel to communicate with SCSI targets. This mode requires that a host system have an FC attached device and associated device drivers and software to generate SCSI-3 FCP requests. This system acts as an initiator using the storage router to communicate with SCSI target devices. The SCSI devices supported can include SCSI-2 compliant direct or sequential access (disk or tape) devices. The storage router serves to translate command and status information and transfer data between SCSI-3 FCP and SCSI-2, allowing the use of standard SCSI-2 devices in a Fibre Channel environment. (CQ-1001, 6:56-67, emphasis added).

30. Based on the above, it appears that in the context of the '041 Patent native low level block protocols are storage protocols that access storage space on storage devices at the block level. The specification contrasts these low level storage protocols with "high level file system protocols" that address storage space at the file level. (CQ-1001, 3:43-52). Further, according to the specification, a "SCSI command" is one example of a native low level block protocol command. (CQ-1001, 5:59-63). Additionally, the specification describes a scenario in which Fibre Channel initiators may communicate with SCSI storage devices by encapsulating low level SCSI commands into Fibre Channel Protocol (FCP) requests. (CQ-1001, 6:56-67).

31. It is therefore my opinion that a person of ordinary skill in the art would understand the broadest reasonable interpretation of "native low level block protocol" in view of the specification to include at least: "*a protocol in which storage space is accessed at the block level, such as the SCSI protocol*."

B. "remote"

32. The claim term "remote" is found in claims 1-4, 7, 12, 13, 19-23, 26, 31, 32, 37-40, 43, 48, 49, and 53.

33. I note that in the *3Par* case, the District Court construed "remote" to mean "indirectly connected through at least one serial network transport medium." (CQ-1010, p. 12).

34. Based on my review of the '041 Patent, it appears that the specification does not provide an explicit definition for "remote." I note, however, that the '041 specification discloses that:

> Typical storage transport mediums provide for a relatively small number of devices to be attached over relatively short distances. One such transport medium is a Small Computer System Interface (SCSI) protocol, the structure and operation of which is generally well known as is described, for example, in the SCSI-1, SCSI-2 and SCSI-3 specifications. High speed serial interconnects provide

> enhanced capability to attach a large number of high speed devices to a common storage transport medium over large distances. (CQ-1001, 1:51-59).
>
> A technical advantage of the present invention is the ability to centralize local storage for networked workstations without any cost of speed or overhead. Each workstation access its virtual local storage as if it work [sic] locally connected. Further, the centralized storage devices can be located in a significantly remote position even in excess of ten kilometers as defined by Fibre Channel standards. (CQ-1001, 2:55-61).

With reference to Fig. 3, the specification states that:

> Storage router 56 provides centralized control of what each workstation 58 sees as its local drive, as well as what data it sees as global data accessible by other workstations 58. Consequently, the storage space considered by the workstation 58 to be its local storage is actually a partition (i.e., logical storage definition) of a physically remote storage device 60, 62 or 64 connected through storage router 56. (CQ-1001, 5:8-14, emphasis added).

35. I note that, when describing the storage network shown in Fig. 3, the specification states that storage devices on the SCSI bus connected to one side of the storage router are "physically remote" from workstations on the Fibre Channel transport medium connected to the other side of the storage router. (CQ-1001, 5:8-14). Further, the specification states that SCSI provides for "relatively short distances" between devices (*e.g.*, less than 25 meters) while Fibre Channel provides for "large distances" between devices (*e.g.*, more than 10 kilometers). (CQ-1001, 1:51-64). One of ordinary skill in the art reading the specification at the time of the '041 Patent invention would understand that the SCSI-1, SCSI-2, and SCSI-3 transport media referred to therein are conventional parallel network interconnects, and that the Fibre Channel transport medium referred to is a serial network interconnect.

36. It is therefore my opinion that a person of ordinary skill in the art would understand the broadest reasonable interpretation of "remote" in view of the specification to include at least: "*indirectly connected through a storage router to enable connections to storage devices at a distance greater than allowed by a conventional parallel network interconnect*."

VII. Challenge #1: Claims 1-14, 16-33, 35-50, and 53 are obvious over the CRD-5500 Manual in view of the HP Journal

37. It is my opinion that the CRD-5500 Manual in view of the HP Journal renders obvious each and every element of at least claims 1-14, 16-33, 35-50, and 53 of the '041 Patent.

A. The CRD-5500 Manual

38. The CRD-5500 Manual describes the features and operation of the CRD-5500 SCSI Raid Controller ("CRD-5500 RAID controller"). In general, the

CRD-5500 RAID controller routes commands and data between hosts (*i.e.*, initiators) and storage devices (*i.e.*, targets) coupled to the controller. (CQ-1004, pp. 1-1, 1-4). Figure 1-2 in the CRD-5500 Manual, annotated below, illustrates the architecture of the storage network in which the CRD-5500 RAID controller operates. Specifically, in Figure 1-2, hosts attached to SCSI buses (*i.e.*, transport media) are connected to one side of the CRD-5500 RAID controller and storage devices attached to SCSI buses are connected to the other side of the CRD-5500 RAID controller. (CQ-1004, p. 2-4). The CRD-5500 RAID controller enables the exchange of SCSI commands and data between the hosts and the storage devices. (CQ-1004, pp. 1-1, 1-4, 2-1, 2-4). To improve performance, the CRD-5500 RAID controller includes an onboard cache (also commonly known as a buffer) with "512 megabytes of memory in the form of standard 72-pin, 60-nanosecond SIMMs" that temporarily buffers data flowing between the hosts and the storage devices. (CQ-1004, p. 1-4.)

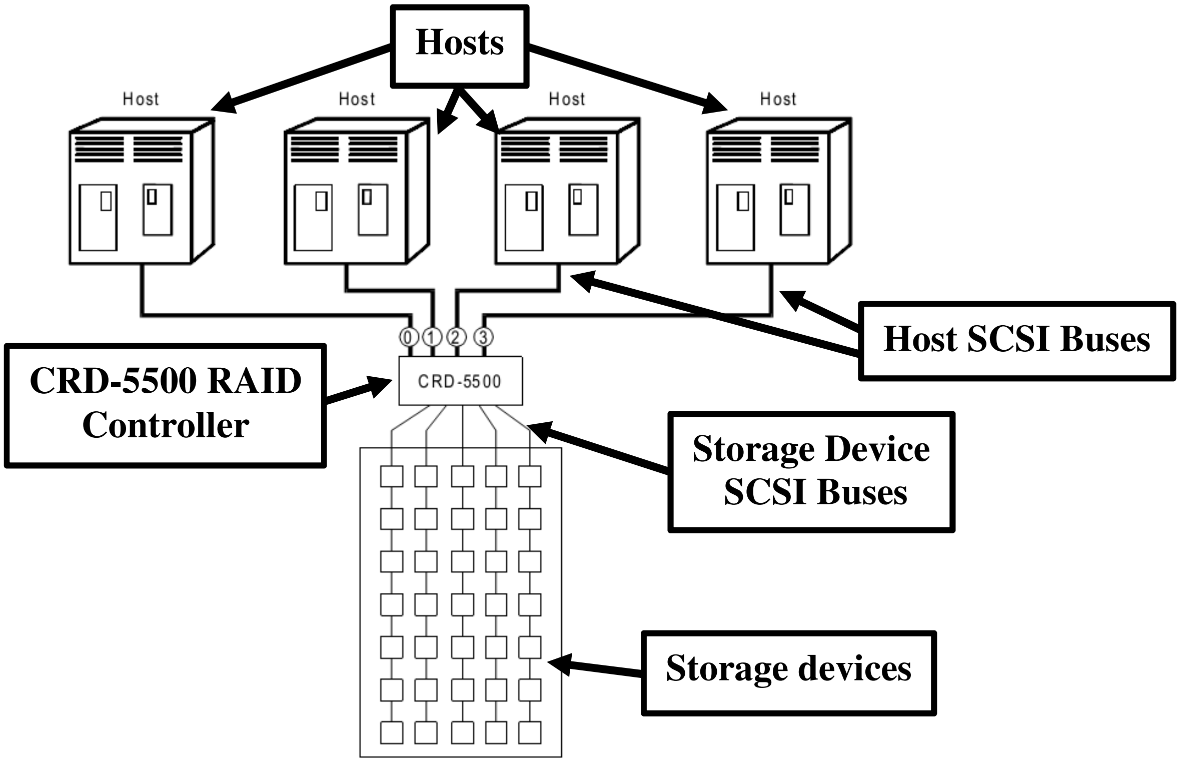

CQ-1004, Figure 1-2 (annotated)

39. The CRD-5500 RAID controller includes a Monitor Utility in its firmware that gives a user "complete control over the configuration and operation of the controller." (CQ-1004, p. 4-1). The Monitor Utility includes a "Host LUN Mapping" feature that allows a user to map subsets of storage space on the storage devices (referred to as "redundancy groups") to specific hosts, where each host has its own mapping table. (CQ-1004, pp. 1-2, 1-11, 4-2, 4-5). (A "LUN" is a logical unit number used to represent storage space). As shown in the screenshot of a mapping table below, each host has a set of addressable virtual LUNs to which these redundancy groups are mapped. (CQ-1004, pp. 1-2, 1-11, 4-2, 4-5). A host accesses the physical redundancy groups via its virtual host LUNs (*i.e.*, a host points a SCSI I/O command to a host LUN rather than a physical redundancy group). (CQ-1004, pp. 1-2, 1-11, 4-2, 4-5). Each mapped redundancy group "will appear to the host as a different disk drive." (CQ-1004, p. 3-6). The following is a screenshot of the Host LUN Mapping feature of the Monitor Utility:

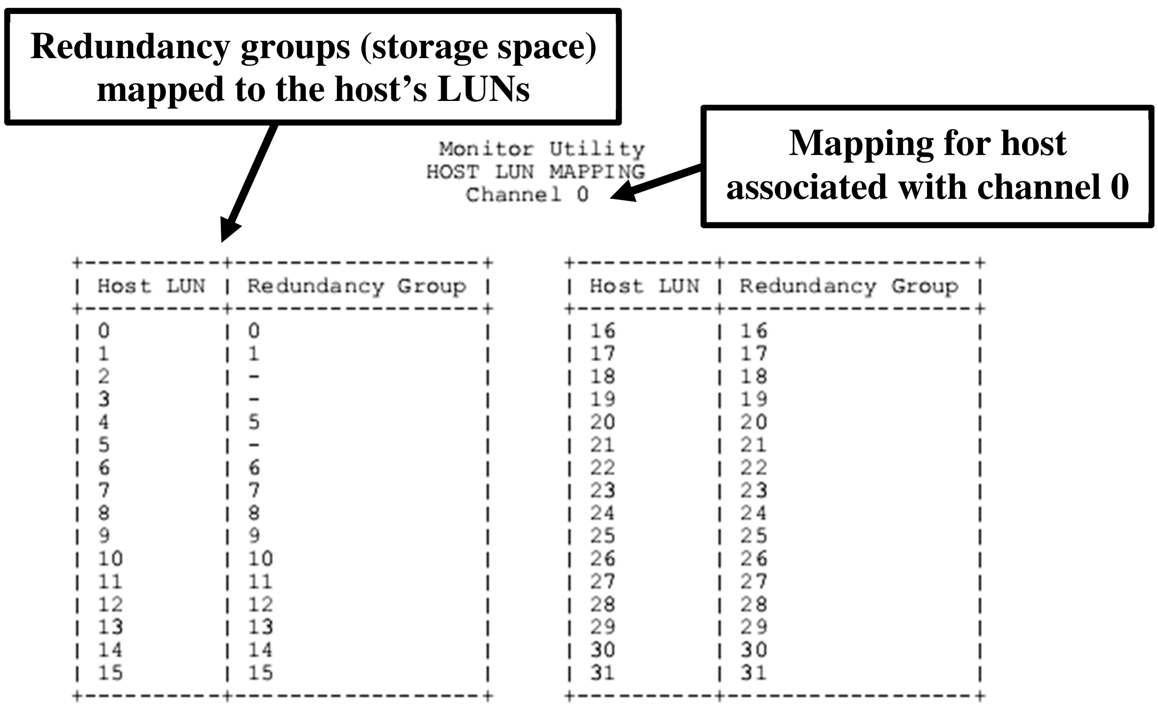

Host LUN	Redundancy Group
0	0
1	1
2	-
3	-
4	5
5	-
6	6
7	7
8	8
9	9
10	10
11	11
12	12
13	13
14	14
15	15

Host LUN	Redundancy Group
16	16
17	17
18	18
19	19
20	20
21	21
22	22
23	23
24	24
25	25
26	26
27	27
28	28
29	29
30	30
31	31

CQ-1004, p. 4-5 (annotated)

40. The CRD-5500 RAID controller uses the Host LUN mapping tables to facilitate routing and access control. In particular, the CRD-5500 RAID controller may "make a redundancy group visible to one host but not to another" by having a host transmits SCSI commands to a set of virtual LUNs rather than the physical storage devices. (CQ-1004, p. 1-1). For example, the CRD-5500 RAID controller “may make redundancy group 8 available on LUN 4 on host channel 0 and block access to it on host channel 1.” (CQ-1004, p. 4-5). I have annotated Figure 1-2 below to illustrate the general flow of I/O commands in the CRD-5500 RAID controller system.

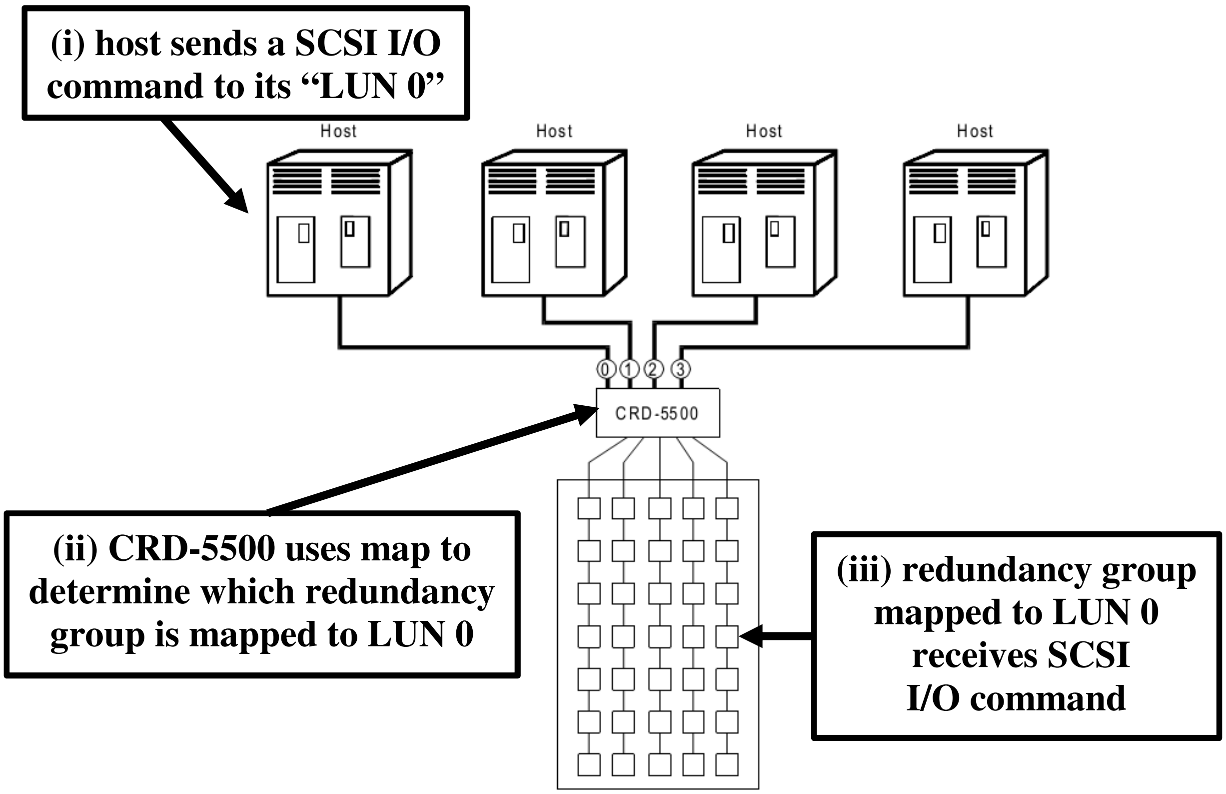

CQ-1004, Figure 1-2 (annotated)

41. As mentioned above, the Monitor Utility which implements the Host LUN mapping is stored in firmware on the CRD-5500 RAID controller. The CRD-5500 Manual states that the firmware is user-upgradable: “To update the flash EPROM containing the controller's firmware, follow the directions in this section. (The version number of your current firmware is displayed in the title box of the monitor utility's opening screen.)” CQ-1004, pp. 4-14 – 4-16.

42. Further, the CRD-5500 RAID controller includes numerous slots for I/O adapter cards that connect to SCSI buses on which the hosts and storage devices communicate. (CQ-1004, pp. 1-1, 2-1, 2-4). A notable feature of the CRD-5500 RAID controller is that it “employs a modular design for maximum flexibility.” (CQ-1004, p. 1-1). Specifically, the controller’s slots accept different types of I/O adapter cards that connect to different types of SCSI buses. (CQ-1004, pp. 1-1, 2-1, 2-4). Figure 2-1 of the CRD-5500 Manual illustrates the modular nature of the controller:

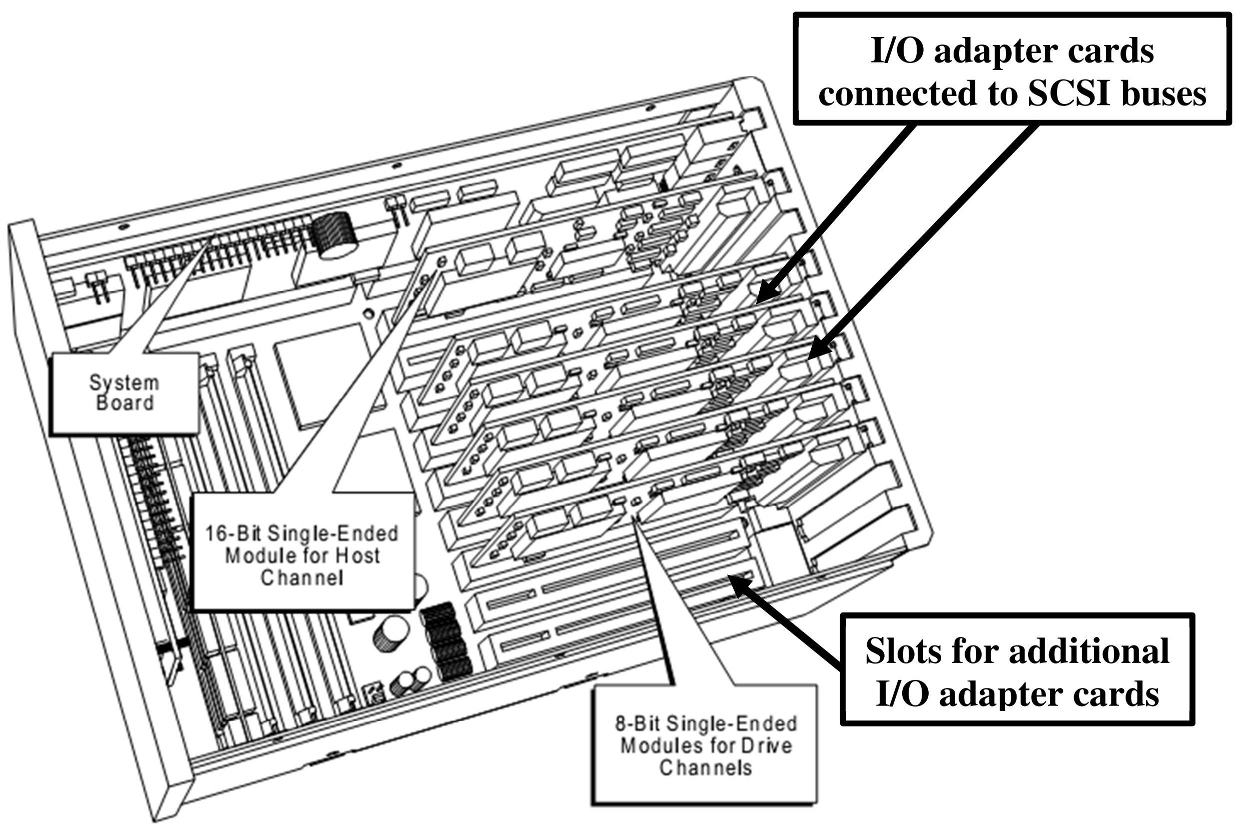

CQ-1004, Figure 2-1 (annotated)

43. It is important to note that the designers of the CRD-5500 RAID controller intended for the CRD-5500 to work not only with SCSI buses but also with other types of communication links. A data sheet advertising the features of the CRD-5500 RAID controller states that the controller's "RAID architecture and ASICs were designed to support tomorrow's high speed serial interfaces, such as Fiberchannel (FCAL) and Serial Storage Architecture (SSA)." (CQ-1005, p. 1). Further, the data sheet states that "[t]o meet the performance requirements of tomorrow's high-speed serial interfaces, CMD designed the CRD-5500 to use the MIPS R3000 family of 32-bit RISC processors and proprietary CMD custom ASIC logic components." (CQ-1005, p. 2). Further still, the CRD-5500 Manual details the high-performance components that allow the CRD-5500 RAID controller to support serial interfaces such as Fibre Channel:

1.3.1 Custom Components

To increase performance and reliability, the CRD-5500's core functions have been encapsulated in four custom ASIC (Application Specific Integrated Circuits) components.

XOR ASIC: Used in the Exclusive-Or parity calculations employed by RAID levels 4 and 5.

DMA ASIC: Controls the data path hardware for the various I/O ports.

CPU Interface ASIC: Supports the controller's MIPS R3000 RISC central processing unit.

Memory Controller ASIC: Controls the memory system and supports data movement on the internal bus at a maximum burst rate of 80 MB/second and a maximum sustainable rate of 60 MB/second.

(CQ-1004, p. 1-3).

44. The front cover of the CRD-5500 Manual is dated November 21, 1996 and the manual itself was available for public download from the CMD

Technologies website, http://www.cmd.com, at least by December 26, 1996. I was able to download the manual from an archived version of the website captured on December 26, 1996 by the Internet Archive Wayback Machine at http://web.archive.org/web/19961226085953/http://www.cmd.com/ftproot/pub/raid/5500/manual/crd5500user.pdf. Exhibit CQ-1004 reflects the document I downloaded.

45. Additionally, the CRD-5500 Data Sheet was publically available on the CMD Technologies website at least by December 26, 1996. I was able to download the CRD-5500 Data Sheet from an archived version of the CMD website captured on December 26, 1996 by the Internet Archive Wayback Machine at http://web.archive.org/web/19961226091552/http://www.cmd.com/brochure/crd5500.htm. Exhibit CQ-1005 reflects the document I downloaded.

B. <u>The HP Journal</u>

46. Volume 47, issue 5 of the Hewlett-Packard Journal includes a number of articles that address the growing problem in 1997 of "I/O channels becom[ing] bottlenecks to system performance." (CQ-1006, p. 5). Specifically, one article in the issue provides an introduction to the Fibre Channel I/O interface that is described as "a flexible, scalable, high-speed data transfer interface that can operate over a variety of both copper wire and optical fiber at data rates up to 250 times faster than existing communications interfaces." (CQ-1006, p. 94). The article additionally provides many reasons a Fibre Channel communication link is superior to SCSI bus (*e.g.*, longer distances and higher bandwidth, smaller connectors, reduction in the number of I/O slots needed). (CQ-1006, pp. 94, 95, 100). The same article further notes that SCSI commands may be "encapsulated and transported within Fibre Channel frames" to support existing storage hardware. (CQ-1006, pp. 94, 95, 100).

47. A second article in the same issue of the HP Journal describes a Fibre Channel protocol chip made by HP called "Tachyon." The article states that the Tachyon chip implements the Fibre Channel standard and "enables low-cost gigabit host adapters on industry-standard buses." (CQ-1006, p. 101). Additionally, the article provides details about how to implement a Fibre Channel I/O adapter card that uses the Tachyon chip. (CQ-1006, p. 111).

C. <u>Reasons to Combine the CRD-5500 Manual and the HP Journal</u>

48. It is my opinion that one of ordinary skill in the art would have been motivated to combine the teachings of the CRD-5500 Manual and the HP Journal for the reasons set forth below.

49. First, the CRD-5500 Manual describes the operation of the CRD-5500 RAID controller, which routes data between hosts and SCSI disk arrays connected to SCSI buses. The CRD-5500 Manual further teaches that the CRD-5500 RAID controller has a modular design that accepts many different types of I/O modules to interface with different transport media. (CQ-1004, pp. 1-1, 2-1). More importantly, the CRD-5500 RAID Controller was specifically "designed to support tomorrow's high speed serial interfaces, such as Fiberchannel." (CQ-1005, p. 1).

50. Second, the HP Journal—in describing the contents of the issue—teaches that "as the number of interconnects between systems and I/O devices continues to increase, I/O channels become bottlenecks to system performance. For all these reasons, today's parallel bus architectures are reaching their limits." (CQ-1006, p. 5). More specifically, the HP Journal teaches that there are a number of inherent limitations in the SCSI bus architecture that prevent it from "keeping pace with ever-increasing processor speeds and data rate requirements." (CQ-1006, p. 99).

51. The HP Journal further teaches that the Fibre Channel serial transport medium provides solutions to these limitations because its "increased bandwidth provides distance flexibility, increased addressability, and simplified cabling." (CQ-1006, p. 99). In particular, the HP Journal teaches that Fibre Channel "can operate from 2.5 to 250 times faster than existing communications interfaces" and that "[a] single 100-Mbyte/s Fibre Channel port can replace five 20-Mbyte/s SCSI ports, in terms of raw through put." (CQ-1006, p. 94). Additionally, "Fibre

Channel resolves the 'slots and watts' problem that current symmetric multiprocessing systems have" because Fibre Channel supports the same I/O services with fewer number of slots. (CQ-1006, pp. 100, 101). Further, the HP Journal notes that Fibre Channel provides backwards compatibility with SCSI-based hardware because SCSI commands may be "encapsulated and transported within Fibre Channel frames." (CQ-1006, pp. 94-95).

52. The HP Journal additionally describes the structure and functionality of HP's Tachyon Fibre Channel protocol chip and teaches one of ordinary skill in the art how to implement a generic Fibre Channel I/O adapter board using the Tachyon chip. (CQ-1006, pp. 101-111; Fig. 14). Notably, the Tachyon chip was designed to integrate with a variety of system types: "Tachyon's host attach enables low-cost gigabit host adapters on industry-standard buses including PCI, PMC, S-Bus, VME, EISA, Turbo Channel, and MCA. It is easily adaptable both to industry-standard and proprietary buses through the Tachyon system interface (a generic interface) and provides a seamless interface to GLM-compliant modules and components." (CQ-1006, p. 101).

53. Accordingly, one of ordinary skill in the art would have been motivated to replace the SCSI I/O host modules in the CRD-5500 RAID Controller with a Fibre Channel I/O module based on the Tachyon chip. In some instances, one of ordinary skill in the art would have also been motivated to replace the SCSI

I/O disk modules with one or more Fibre Channel I/O modules. Both substitutions would have been simple because (i) the CRD-5500 RAID controller has a modular design that accepts different types of I/O modules and (ii) the Tachyon Fibre Channel chip is "easily adaptable" to different systems. Using a Fibre Channel-based I/O module instead of SCSI-based I/O modules on the "host side" of the CRD-5500 RAID Controller would have allowed the CRD-5500 to communicate with the hosts via a Fibre Channel transport medium rather than via SCSI buses, and using one or more Fibre Channel-based I/O modules instead of SCSI-based I/O modules on the "disk side" of the CRD-5500 RAID Controller would have allowed the CRD-5500 to communicate with the disk drives via Fibre Channel transport media. Using Fibre Channel to communicate on the "host side" and/or "disk side" of the controller would eliminate many of the known limitations of SCSI buses, as described in the HP Journal.

54. Additionally, the beneficial result of modifying the CRD-5500 RAID Controller to interface with a Fibre Channel transport medium would have been predictable because the CRD-5500 RAID Controller was specifically "designed to support tomorrow's high speed serial interfaces, such as Fiberchannel." (CQ-1005, p. 1). Moreover, the HP Journal specifically contemplates that a single Fibre Channel port may replace multiple SCSI ports, thereby saving system slots. (CQ-1006, pp. 94, 100, 101). Additionally, the creators of the Fibre Channel standard specifically intended for Fibre Channel links to carry SCSI commands so that they would be backwards compatible with SCSI-based devices, such as those attached to the CRD-5500 RAID controller. (CQ-1006, pp. 94-95).

55. Further, based on my knowledge of the storage industry at the time of the '041 invention, it would have been predictable to communicate with the SCSI disks attached to the CRD-5500 RAID controller via a Fibre Channel link instead of a SCSI bus. In that regard, one of ordinary skill in the art would know that legacy SCSI disks could be aggregated and placed into a storage enclosure coupled to a Fibre Channel link enclosure to allow migration from a SCSI-based architecture to a Fibre Channel-based architecture. Such a Fibre Channel-enclosure could house a greater number of SCSI disks than could be attached to a SCSI bus because of Fibre Channel's improved addressing capabilities. Furthermore, a Fibre Channel-enclosure would offer data transfer rates far in excess of those permitted on a SCSI bus. One example of such a storage enclosure was the Sun Microsystems SPARCstorage Array, which housed SCSI disks and communicated with other network elements via a Fibre Channel optical link. (CQ-1008, pp. 1-1 - 1-4, 3-1, 3-2; CQ-1009, pp. 1-4).

56. Moreover, to the extent any modifications would have been needed to the teachings of the CRD-5500 Manual in order to accommodate the teachings of the HP Journal, such modifications would have been within the level of ordinary skill in the art of network storage. For example, because the CRD-5500 RAID controller was *specifically intended to be used with Fibre Channel*, it already includes high-performance hardware components capable of supporting high rates of data transfer. And, because the firmware in the CRD-5500 RAID controller is user upgradeable, one of ordinary skill in the art could have easily made any software modifications necessary to accommodate Fibre Channel-based devices. Accordingly, any hardware or software modifications to the components of the CRD-5500 necessary to keep them operating in their intended manner would have been well within the skills of one of ordinary skill in the art.

57. Thus, the CRD-5500 RAID Controller's ability to accept different I/O modules and compatibility with high speed serial interfaces such as Fibre Channel, as taught by the CRD-5500 Manual and Data Sheet, in view of the advantages of using a Fibre Channel I/O module (*e.g.*, increased bandwidth), as taught by the HP Journal, together produce the obvious, beneficial, and predictable result of utilizing a Fibre Channel I/O module in the CRD-5500 to communicate with hosts connected to a Fibre Channel transport medium.

D. <u>Detailed Analysis</u>

58. The following claim chart describes how the CRD-5500 Manual in view of the HP Journal renders obvious each and every element of at least claims

1-14, 16-33, 35-50, and 53.

U.S. Patent No. 7,934,041	
Claim 1	
[1.0] A storage router for providing virtual local storage on remote storage devices, comprising:	[1.0] ***"A storage router for providing virtual local storage on remote storage devices"*** **First**, the CRD-5500 Manual describes the operation of a CRD-5500 RAID controller (storage router), which routes data between hosts and SCSI disk arrays (remote storage devices): 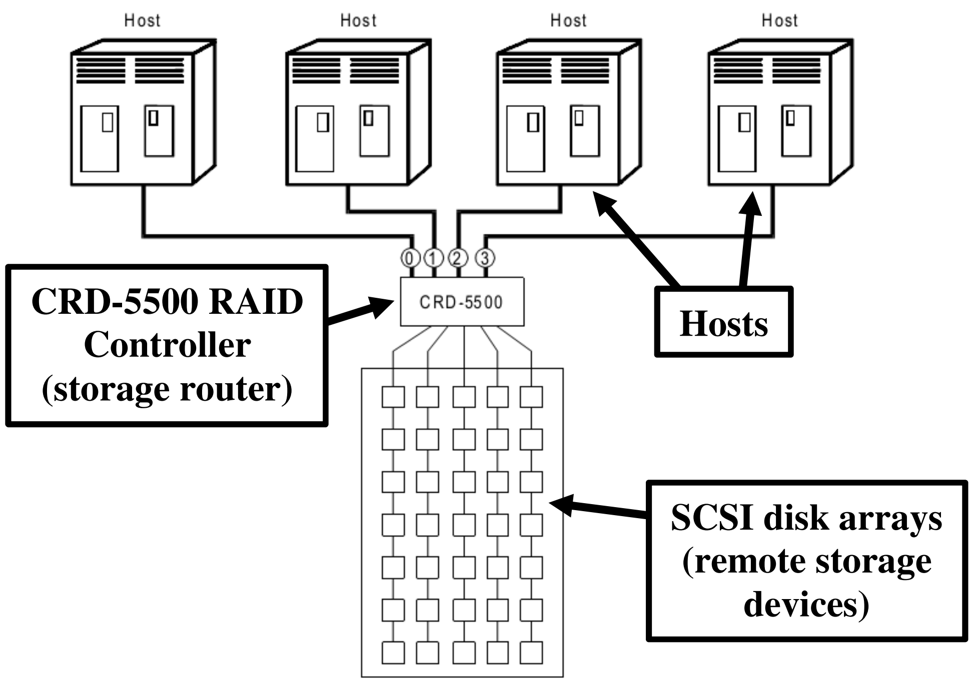 CQ-1004, Figure 1-2 (annotated) "Figure 1-2 shows how you can <u>connect as many as four hosts to the CRD-5500</u>. By using the controller's Host LUN Mapping feature, you can assign redundancy groups to a particular host." CQ-1004, p. 1-2 (emphasis added). "The CRD-5500 RAID controller <u>provides high-performance,</u>

	high-availability access to SCSI disk array subsystems along a Fast/Wide SCSI bus. With a modular hardware design and an intuitive configuration utility, the controller may be tailored to suit a wide range of storage needs, now and in the future." CQ-1004, p. 1-1 (emphasis added). **Second**, the CRD-5500 Manual teaches that the CRD-5500 RAID controller provides virtual local storage on disk drives coupled to the controller by grouping the drives into RAID sets, partitioning the RAID sets into redundancy groups and assigning the redundancy groups to Logical Unit Numbers (LUNs), which are logical partitions that appear to a host as disk drives: "The CRD-5500 affords great flexibility in creating RAID sets. Drives of different sizes and manufacturers may be combined in a RAID set." CQ-1004, p. 3-3 (emphasis added). "The CRD-5500 supports the partitioning of RAID sets. A partitioned RAID set will have multiple redundancy groups associated with it. Each redundancy group will have its own redundancy group number, which corresponds to the Logical Unit Number (LUN) that the host will use to address the partition, unless you map the redundancy group to another LUN with the Host LUN Mapping screen." CQ-1004, p. 4-2 (emphasis added). "Since **each partition will appear to the host as a different disk drive,** it must have its own LUN." CQ-1004, p. 3-6 (emphasis added). "In this example, Redundancy Group numbers 1 through 4 are associated with RAID Set number 1. From the controller's point of view in this example, Redundancy Group number 0 belongs to RAID Set 0, which is configured for RAID Level 4, and Redundancy Group numbers 1 through 4 belong to RAID Set 1, which is configured for Level 5. From the host's point of view, there is no such thing as RAID Sets 0 and 1, only (in this example) LUNs 0 through 4, which the CRD-5500 sees as Redundancy Groups 0 through 4." CQ-1004, p. 3-6 (emphasis

added).

Accordingly, the CRD-5500 Manual teaches that a logical partition of storage space on remote disks will "appear to the host" as a local disk drive (*i.e.*, "virtual local storage").

Third, with respect to the recitation in the preamble that the storage devices are "remote," the CRD-5500 Manual teaches that the storage devices are indirectly connected to the hosts via the CRD-5500 RAID controller:

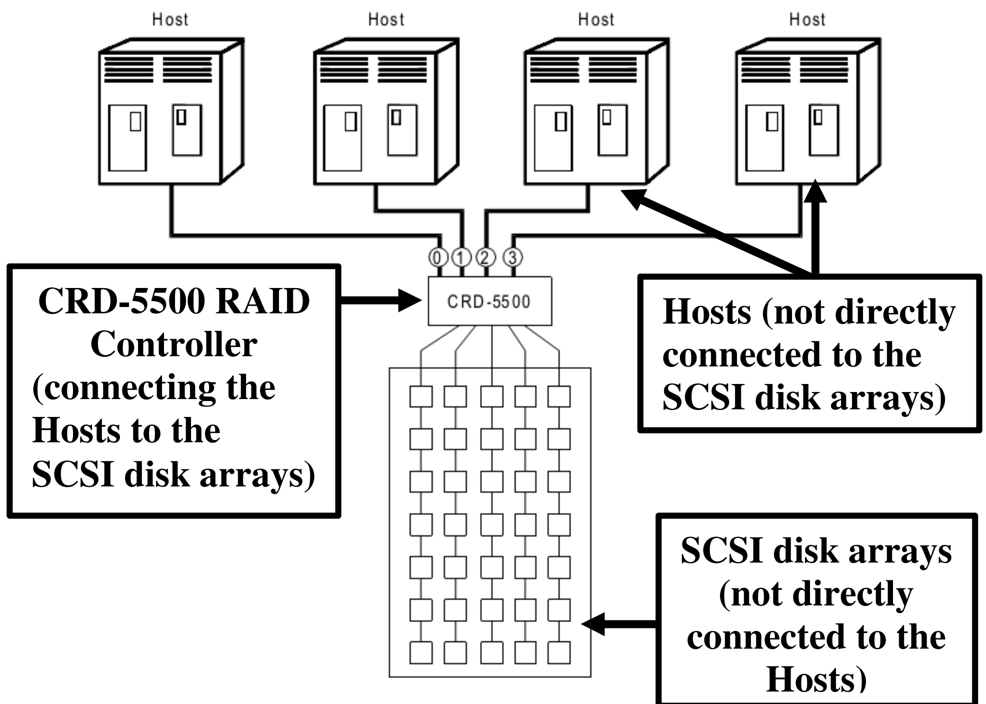

CQ-1004, Figure 1-2 (annotated)

Further, if "remote" is interpreted to mean that either (i) the storage devices are separated from the hosts by a distance greater than allowed by a conventional parallel network interconnect or (ii) that the storage devices are indirectly connected through at least one serial network transport medium, then section [1.1] of this claim chart describes how one of ordinary skilled in the art would have been motivated by the teachings of the HP Journal to utilize a Fibre Channel serial transport medium in conjunction with the CRD-5500 RAID controller. To that end, the HP Journal

	teaches that Fibre Channel "allows sustained gigabit data throughput at distance options from ten meters on copper to ten kilometers over single-mode optical fiber." CQ-1006, p. 99 (emphasis added). **Thus**, the CRD-5500 RAID controller that routes data between hosts and remote disk drives, and (ii) provides virtual partitions associated with the remote disk drives so that the drives appear to be local to the host computers, as taught by the CRD-5500 Manual, discloses "a storage router for providing virtual local storage on remote storage devices" as recited in the claim.
[1.1] a first controller operable to interface with a first transport medium, wherein the first medium is a serial transport media; and	[1.1] ***"a first controller operable to interface with a first transport medium, wherein the first medium is a serial transport media"*** **First**, the CRD-5500 Manual teaches that the CRD-5500 RAID controller includes nine slots for I/O modules (controllers) that interface with SCSI buses (transport media). One or more of the I/O modules may be designated as host modules that communicate with host computers on SCSI buses: "The controller has nine I/O module slots. One of these slots is reserved for use as a host channel, and five are reserved for use as drive channels." CQ-1004, p. 1-1. "All controllers come with a system board. The rest of the slots are I/O channels, either for I/O to and from the host computer or to and from the disk drives connected to the controller. CQ-1004, p. 2-1 (emphasis added). "In this example, we will configure the 16-bit I/O channels in slots 0 and 1 as host modules. Connect the host SCSI cables to these modules." CQ-1004, p. 3-1 (emphasis added). "The CRD-5500 involves several SCSI buses, which can take the form of 8-bit single-ended, 16-bit single-ended, or 16-bit differential buses. On top of this, one or more of the buses may be designated as a host bus for communicating with the host

	computer, and the others may be populated with disk drives to form the RAID set pool." CQ-1004, p. 2-4 (emphasis added). 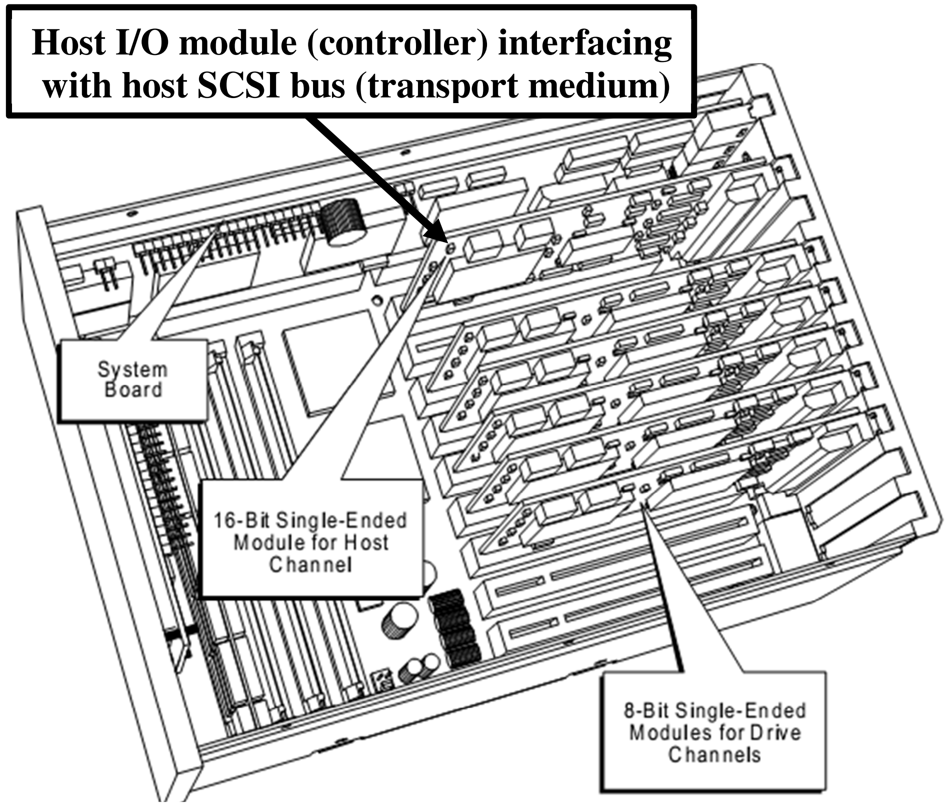 CQ-1004, Figure 2-1 (annotated) **Second**, the CRD-5500 Manual teaches that the CRD-5500 RAID controller has a modular design. Specifically, the controller's slots accept many different types of I/O modules to interface with different buses (transport media): "The CRD-5500 employs a modular design for maximum flexibility. You may customize your controller using a combination of 8-bit single-ended SCSI I/O modules, 16-bit single-ended modules, and 16-bit differential modules." CQ-1004, p. 1-1 (emphasis added). "The CRD-5500's modular design offers great flexibility in configuring the controller. All controllers come with a system board. The rest of the slots are I/O channels, either for I/O to and from the host computer or to and from the disk drives connected

	to the controller. CQ-1004, p. 2-1 (emphasis added). **Third**, the HP Journal teaches that there are a number of advantages of using the Fibre Channel serial transport medium instead of SCSI buses, and contemplates replacing SCSI ports with a Fibre Channel port: "A channel such as SCSI (Small Computer Systems Interface), which operates at a maximum throughput of 20 megabytes per second in fast and wide mode, simply cannot keep pace with ever-increasing processor speeds and data rate requirements." CQ-1006, p. 99. "Additionally, as the number of interconnects between systems and I/O devices continues to increase, I/O channels become bottlenecks to system performance. For all these reasons, today's parallel bus architectures are reaching their limits. In the search for a higher-performance <u>serial interface</u>, HP chose Fibre Channel because it overcomes the limitations mentioned above by supporting sustained gigabit data transfer rates." CQ-1006, p. 5 (emphasis added). "Fibre Channel is a flexible, scalable, high-speed data transfer interface that can operate over a variety of both copper wire and optical fiber at data rates up to 250 times faster than existing communications interfaces." CQ-1006, p. 94. "A single 100-Mbyte/s Fibre Channel port can replace five 20-Mbyte/s SCSI ports, in terms of raw through put. Fibre Channel provides a total network bandwidth of about one gigabit per second." CQ-1006, p. 94. "Another advantage of Fibre Channel is that it uses small connectors. The <u>serial connectors used for Fibre Channel</u> are a fraction of the size of SCSI parallel connectors and have fewer pins, thereby reducing the likelihood of physical damage. Also, depending on the topology, many more devices can be interconnected on Fibre Channel than on existing channels." CQ-1006, p. 94 (emphasis added).

	"Fibre Channel's increased bandwidth provides distance flexibility, increased addressability, and simplified cabling. Fibre Channel has versatility, room for growth, and qualified vendor support. Mass storage suppliers are using Fibre Channel to interconnect subsystems and systems and to control embedded disk drives." CQ-1006, p. 99. **Additionally**, the HP Journal teaches one of ordinary skill in the art how to implement a Fibre Channel I/O adapter board (first controller) using the Tachyon controller chip that interfaces with a Fibre Channel serial transport medium (first transport medium): "The system interconnect laboratory of the HP Networked Computing Division became interested in <u>Fibre Channel in 1993 as a method of entering the high-speed serial interconnect market</u> because Fibre Channel was the first technology that could be used for <u>both networking and mass storage</u>. . . . The <u>Tachyon chip (Fig. 3) implements the FC-1 and FC-2 layers of the five-layer Fibre Channel standard</u> (see Article 11)." CQ-1006, p. 101 (emphasis added). "<u>Tachyon's host attach enables low-cost gigabit host adapters on industry-standard buses including PCI, PMC, S-Bus, VME, EISA, Turbo Channel, and MCA. It is easily adaptable both to industry-standard and proprietary buses through the Tachyon system interface (a generic interface)</u> and provides a seamless interface to GLM-compliant modules and components." CQ-1006, p. 101 (emphasis added). "<u>A generic Fibre Channel host bus adapter board using the Tachyon chip contains the following</u>: <u>Backplane Connector</u>. Connects the backplane interface chip to the system bus. <u>Backplane Interface Chip</u>. Enables the connection of the Tachyon system interface bus to PCI, EISA, HP-HSC or other bus. <u>Tachyon Chip. HP's Fibre Channel interface controller</u>.

Physical Link Module. Tachyon interfaces to many GLM-compliant physical link modules currently in the marketplace. Types of modules include:
1063-Mbit/s DB9 copper connectors for distances up to 30 meters.
1063-Mbit/s shortwave laser physical link modules for distances up to 500 meters.
1063-Mbit/s longwave laser physical link modules for distances up to 10 kilometers.
266-Mbit/s shortwave laser physical link modules for distances up to 2 kilometers.
A block diagram of a typical host bus adapter is shown in Fig. 14:" CQ-1006, p. 111 (emphasis added).

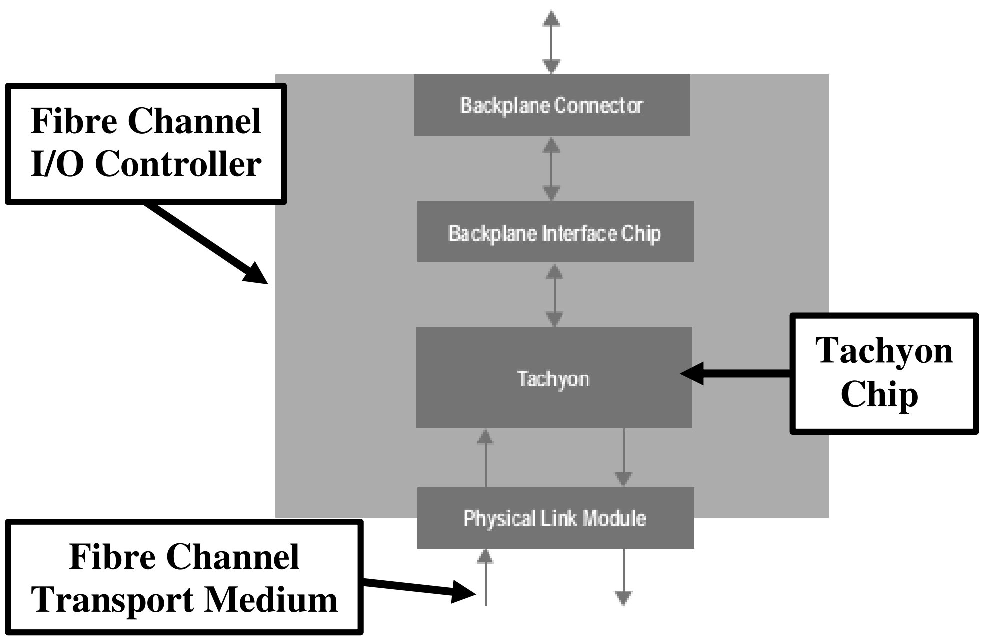

Fig. 14. Block diagram of a typical host adapter board.

CQ-1006, Fig. 14 (annotated)

Accordingly, one of ordinary skill in the art would have been motivated to replace the SCSI I/O host modules in the CRD-5500 RAID Controller with a Fibre Channel I/O host module so that the controller could interface with a Fibre Channel communication link. Such an exchange would have been a simple substitution because the CRD-5500 RAID controller has a modular design. The Fibre Channel-based I/O module would have allowed the CRD-5500 RAID Controller to communicate

	with hosts on a Fibre Channel transport medium, thereby overcoming many of the known limitations of SCSI buses described by the HP Journal. *Additionally*, the beneficial results of modifying the CRD-5500 RAID Controller to interface with a Fibre Channel transport medium would have been predictable because the CRD-5500 RAID Controller *was specifically designed to support high speed serial interfaces, such as Fibre Channel*: "Unlike other RAID controllers, CMD's advanced 'Viper' RAID architecture and ASICs were **<u>designed to support tomorrow's high speed serial interfaces, such as Fiberchannel (FCAL)</u>** and Serial Storage Architecture (SSA)." CQ-1005, p. 1 (emphasis added). **Additionally,** if the "first transport medium" is defined as a single physical link, the HP Journal teaches that multiple workstations may be connected to the same communication link via the Fibre Channel arbitrated loop method: "Fibre Channel arbitrated loop, or FC-AL, is a method for interconnecting from two to 126 devices through attachment points called L_Ports in a loop configuration. . . . Fig. 3 shows an office configured in a public arbitrated loop topology." CQ-1006, p. 95.

	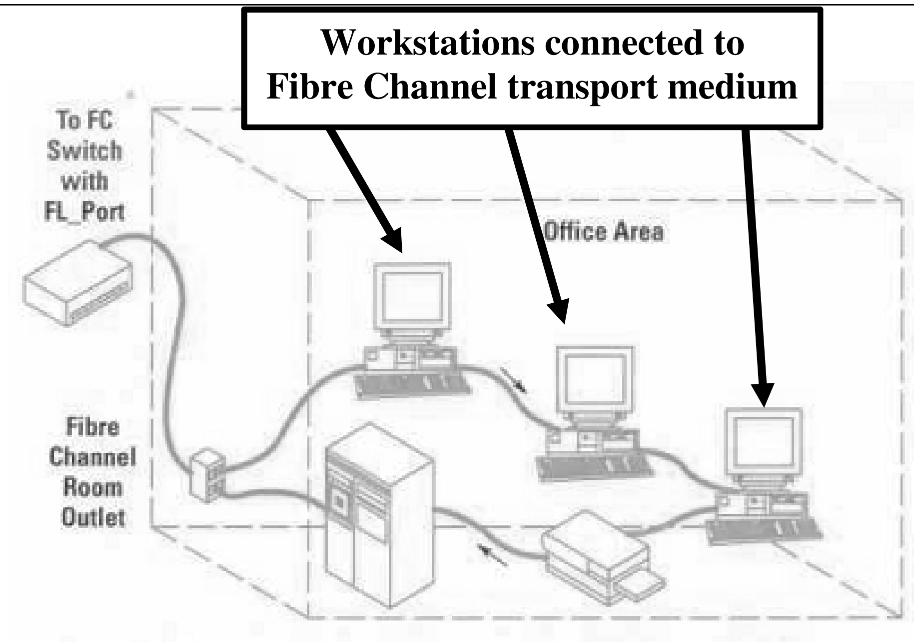 **Fig. 3.** An office configured in a public arbitrated loop topology. CQ-1006, Fig. 3, p. 96 (annotated) **Thus**, the CRD-5500 RAID controller’s ability to accept different I/O modules that interface with different transport media, as taught by the CRD-5500 Manual, in view of the Fibre Channel I/O module and serial Fibre Channel transport medium, as taught by the HP Journal, render obvious “a first controller operable to interface with a first transport medium, wherein the first medium is a serial transport media” as recited in the claim.
[1.2a] a processing device coupled to the first controller, wherein the processing device is configured to:	[1.2a] ***“a processing device coupled to the first controller”*** **First**, the CRD-5500 Manual teaches a processor coupled to the I/O modules (controllers): “To increase performance and reliability, the CRD-5500's core functions have been encapsulated in four custom ASIC (Application Specific Integrated Circuits) components. . . . CPU Interface ASIC: Supports the controller's MIPS R3000 RISC central processing unit.” CQ-1004, p. 1-3 (emphasis added).

"All controllers come with a system board. The rest of the slots are I/O channels, either for I/O to and from the host computer or to and from the disk drives connected to the controller. Each slot may take an 8-bit single-ended module, a 16-bit single-ended module, or a 16-bit differential module." CQ-1004, p. 2-1.

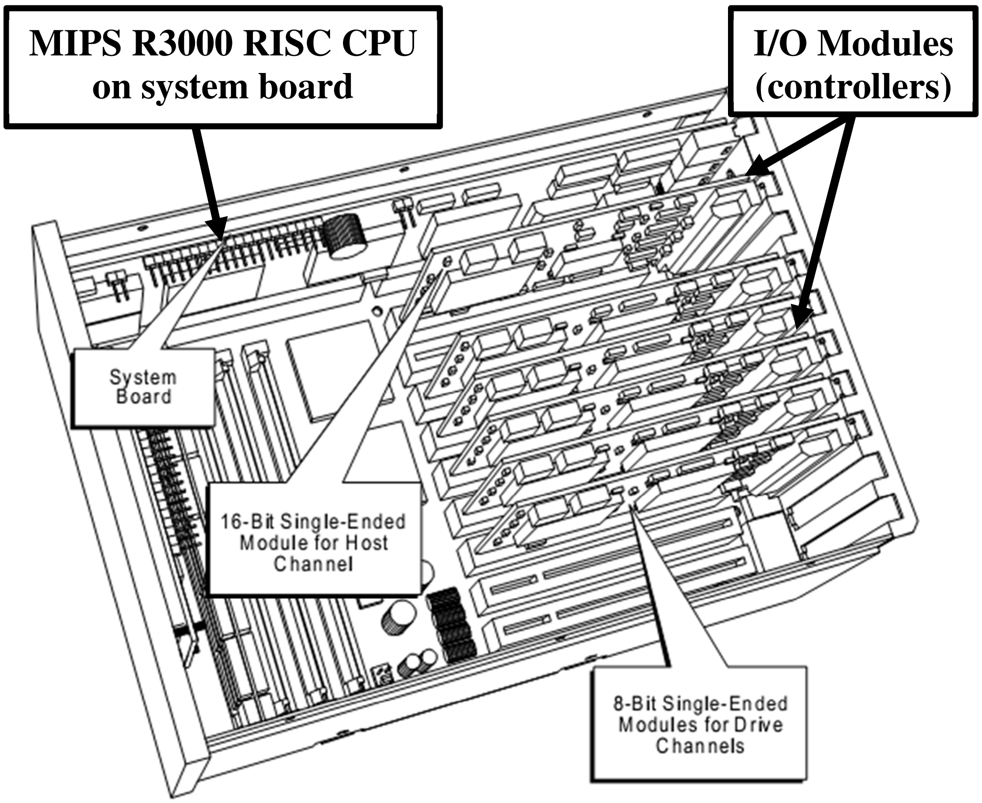

CQ-1004, Figure 2-1 (annotated)

Second, the CRD-5500 Manual teaches that the CRD-5500 RAID controller includes executable firmware with a Monitor Utility that configures the operation of the controller:

"In addition to its flexible hardware design, the CRD-5500's firmware offers the user the flexibility to configure RAID sets in many different ways . . ."CQ-1004, p. 1-1 (emphasis added).

"The CRD-5500's monitor utility provides complete control over the configuration and operation of the controller." CQ-1004, p. 4-1.

"To update the flash EPROM containing the controller's

	<u>firmware</u>, follow the directions in this section. (The version number of your current firmware is displayed in the title box of the monitor utility's opening screen.)" CQ-1004, p. 4-14 (emphasis added). **Thus**, the MIPS R3000 RISC central processing unit coupled to the I/O modules together with the executable firmware that controls the operation of the CRD-5500 RAID controller, as taught by the CRD-5500 Manual, discloses "a processing device coupled to the first controller" as recited in the claim.
[1.2b(i)] maintain a map to allocate storage space on the remote devices to devices connected to the first transport medium by associating representations of the devices connected to the first medium with representations of storage space on the remote storage devices,	[1.2b(i)] ***"maintain a map to allocate storage space on the remote devices to devices connected to the first transport medium by associating representations of the devices connected to the first medium with representations of storage space on the remote storage devices"*** The CRD-5500 Manual teaches that the CRD-5500 RAID controller maintains a map using the Monitor Utility to allocate redundancy groups (storage space) to hosts: "Figure 1-2 shows how you can connect as many as four hosts to the CRD-5500. <u>By using the controller's Host LUN Mapping feature, you can assign redundancy groups to a particular host</u>." CQ-1004, p. 1-2 (emphasis added).

Monitor Utility 02-09-96
HOST LUN MAPPING 13:14:00
Channel 0

Host LUN	Redundancy Group	Host LUN	Redundancy Group
0	0	16	16
1	1	17	17
2	-	18	18
3	-	19	19
4	5	20	20
5	-	21	21
6	6	22	22
7	7	23	23
8	8	24	24
9	9	25	25
10	10	26	26
11	11	27	27
12	12	28	28
13	13	29	29
14	14	30	30
15	15	31	31

CQ-1004, p. 4-5

In more detail, the mapping table above associates representations of hosts with representations of storage space on disk drives:

First, up to four host devices connected to the first medium are represented as “channels” in the map:

“The CRD-5500 permits the I/O modules in slots 1, 2, and 3 to be configured as host or disk channel modules. The Channel Settings screen is the place to configure these modules. Use the up and down arrow keys to maneuver to the channel you wish to configure and press Enter. Then use the arrow keys to toggle between host and disk. Press Enter again to save your selection. Channel 0 is always a host channel and channels 4 through 6 are always disk channels, so the monitor utility will restrict access to these fields.” CQ-1004, p. 4-5 (emphasis added).

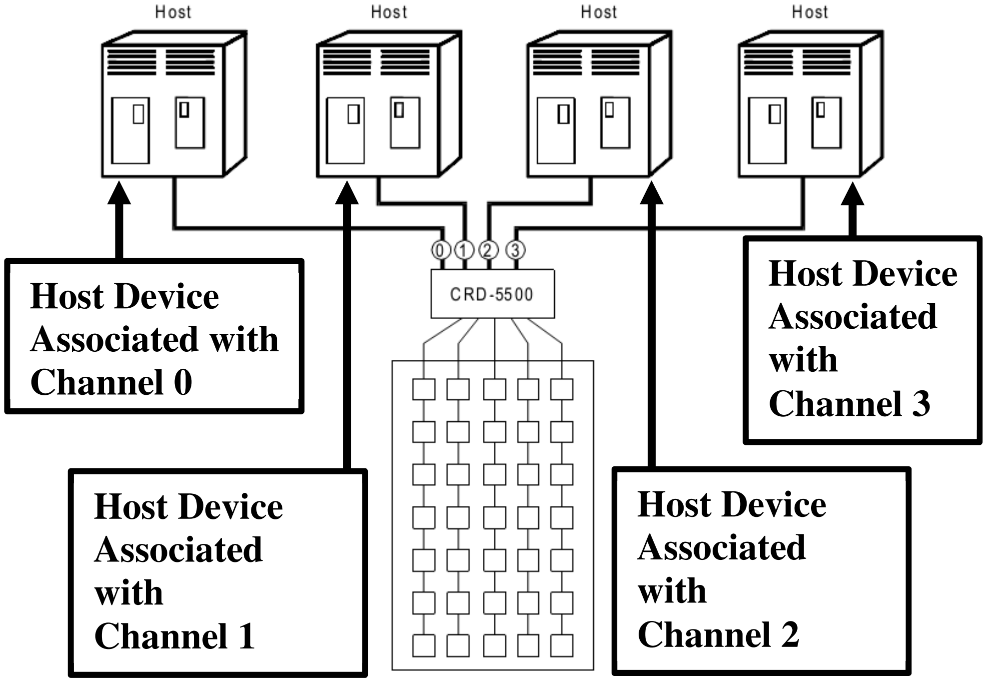

CQ-1004, Figure 1-2 (annotated)

For example, if two hosts are connected to the CRD-5500 RAID Controller, they may be assigned to channels 0 and 1:

Two hosts respectively associated with channels 0 and 1

Monitor Utility
CHANNEL SETTINGS

Channel	Module Type	Module Description
0	HOST	16-Bit Single-Ended
1	HOST	16-Bit Single-Ended
2	DISK	8-Bit Single-Ended
3	DISK	8-Bit Single-Ended
4	DISK	8-Bit Single-Ended
5	DISK	8-Bit Single-Ended
6	DISK	8-Bit Single-Ended
7	DISK	8-Bit Single-Ended

CQ-1004, p. 3-2 (annotated)

The Monitor Utility maintains a mapping table *for each* host channel number, as shown below. (The mapping tables may collectively be considered the "map").

Representation of host device (host channel 0)

```
Monitor Utility                    02-09-96
HOST LUN MAPPING                   13:14:00
Channel 0
```

Host LUN	Redundancy Group	Host LUN	Redundancy Group
0	0	16	16
1	1	17	17
2	-	18	18
3	-	19	19
4	5	20	20
5	-	21	21
6	6	22	22
7	7	23	23
8	8	24	24
9	9	25	25
10	10	26	26
11	11	27	27
12	12	28	28
13	13	29	29
14	14	30	30
15	15	31	31

CQ-1004, p. 4-5 (annotated)

Second, each mapping table includes representations of storage space associated with a particular host channel number. Specifically, in each map, storage space is represented by a set of Host LUNs (logical unit numbers) and associated redundancy groups. The redundancy groups represent physical subsets of storage on disk drives. Specifically, disk drives are grouped into RAID sets and assigned redundancy group numbers:

"RAID sets may comprise any combination of disk drives, regardless of their I/O channel or SCSI ID. In addition, the CRD-5500 defines each RAID set or partition of a RAID set as a 'redundancy group.'" CQ-1004, p. 1-11 (emphasis added).

"The CRD-5500 supports the partitioning of RAID sets. A partitioned RAID set will have multiple redundancy groups associated with it. Each redundancy group will have its own redundancy group number, which corresponds to the Logical Unit Number (LUN) that the host will use to address the partition, unless you map the redundancy group to another LUN with the Host LUN Mapping screen." CQ-1004, p. 4-2

	(emphasis added). *Additionally*, the CRD-5500 Manual teaches that after the redundancy groups have been created, they are mapped to a set of virtual Host LUNs: "These redundancy groups may be mapped to host LUNs, either in a direct one-to-one relationship or in a manner defined by the user." CQ-1004, p. 1-11 (emphasis added). Further, a host addresses the redundancy groups via the virtual Host LUNs: "Each redundancy group will have its own redundancy group number, which corresponds to the Logical Unit Number (LUN) that the host will use to address the partition, unless you map the redundancy group to another LUN with the Host LUN Mapping screen." CQ-1004, p. 4-2 (emphasis added). **Accordingly**, the Host LUNs and redundancy groups in the mapping tables are individually and collectively representations of storage space. **Third**, when configuring a multi-host environment, the CRD-5500 Manual teaches that an administrator uses the Monitor Utility to associate representations of storage space with each host by mapping physical redundancy groups to virtual Host LUNs for each host channel number, as shown below: "This screen may be used to map LUNs on each host channel to a particular redundancy group."

Representations of storage space (Host LUNs and redundancy groups) associated with a representation of a host (channel 0)

Monitor Utility 02-09-96
HOST LUN MAPPING 13:14:00
Channel 0

Host LUN	Redundancy Group	Host LUN	Redundancy Group
0	0	16	16
1	1	17	17
2	-	18	18
3	-	19	19
4	5	20	20
5	-	21	21
6	6	22	22
7	7	23	23
8	8	24	24
9	9	25	25
10	10	26	26
11	11	27	27
12	12	28	28
13	13	29	29
14	14	30	30
15	15	31	31

CQ-1004, p. 4-5 (annotated)

Thus, the CRD-5500 RAID controller's Host LUN Mapping feature that maintains a map to allocate storage space on remote disk drives to hosts by associating host channels with Host LUN/ redundancy group pairings, as taught by the CRD-5500 Manual, discloses a processing device configured to "maintain a map to allocate storage space on the remote devices to devices connected to the first transport medium by associating representations of the devices connected to the first medium with representations of storage space on the remote storage devices" as recited in the claim.

[1.2b(ii)] wherein each representation of a device connected to the first transport medium is

[1.2b(ii)] ***wherein each representation of a device connected to the first transport medium is associated with one or more representations of storage space on the remote storage devices***

The CRD-5500 Manual teaches that the CRD-5500 RAID Controller associates each representation of a host (channel number) with representations of storage space on remote storage devices (Hosts LUNs and redundancy groups):

<table>
<tr>
<td>associated with one or more representations of storage space on the remote storage devices;</td>
<td>
“Figure 1-2 shows how you can connect as many as four hosts to the CRD-5500. <u>By using the controller's Host LUN Mapping feature, you can assign redundancy groups to a particular host</u>.” CQ-1004, p. 1-2 (emphasis added).

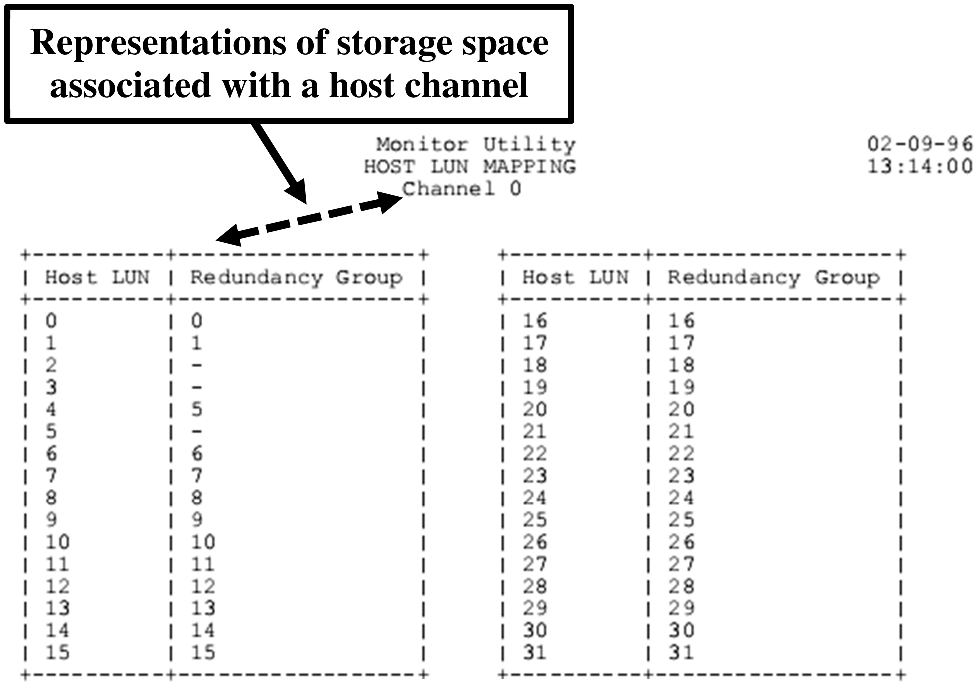

Monitor Utility
HOST LUN MAPPING
Channel 0
02-09-96
13:14:00

<table>
<tr><th>Host LUN</th><th>Redundancy Group</th></tr>
<tr><td>0</td><td>0</td></tr>
<tr><td>1</td><td>1</td></tr>
<tr><td>2</td><td>-</td></tr>
<tr><td>3</td><td>-</td></tr>
<tr><td>4</td><td>5</td></tr>
<tr><td>5</td><td>-</td></tr>
<tr><td>6</td><td>6</td></tr>
<tr><td>7</td><td>7</td></tr>
<tr><td>8</td><td>8</td></tr>
<tr><td>9</td><td>9</td></tr>
<tr><td>10</td><td>10</td></tr>
<tr><td>11</td><td>11</td></tr>
<tr><td>12</td><td>12</td></tr>
<tr><td>13</td><td>13</td></tr>
<tr><td>14</td><td>14</td></tr>
<tr><td>15</td><td>15</td></tr>
</table>

<table>
<tr><th>Host LUN</th><th>Redundancy Group</th></tr>
<tr><td>16</td><td>16</td></tr>
<tr><td>17</td><td>17</td></tr>
<tr><td>18</td><td>18</td></tr>
<tr><td>19</td><td>19</td></tr>
<tr><td>20</td><td>20</td></tr>
<tr><td>21</td><td>21</td></tr>
<tr><td>22</td><td>22</td></tr>
<tr><td>23</td><td>23</td></tr>
<tr><td>24</td><td>24</td></tr>
<tr><td>25</td><td>25</td></tr>
<tr><td>26</td><td>26</td></tr>
<tr><td>27</td><td>27</td></tr>
<tr><td>28</td><td>28</td></tr>
<tr><td>29</td><td>29</td></tr>
<tr><td>30</td><td>30</td></tr>
<tr><td>31</td><td>31</td></tr>
</table>

CQ-1004, p. 4-5 (annotated)

As shown in the above mapping, the host device represented by channel 0 is associated with a plurality of redundancy groups (redundancy groups 0, 1, 5 and 6-31) via the Host LUNs.

Thus, the CRD-5500 RAID controller’s Host LUN Mapping feature that associates a host channel with one or more Host LUN/redundancy group pairings, as taught by the CRD-5500 Manual, discloses “wherein each representation of a device connected to the first transport medium is associated with one or more representations of storage space on the remote storage devices” as recited in the claim.
</td>
</tr>
<tr>
<td>[1.2c] control access from the devices</td>
<td>[1.2c] “control access from the devices connected to the first transport medium to the storage space on the remote storage devices in accordance with the map”</td>
</tr>
</table>

connected to the first transport medium to the storage space on the remote storage devices in accordance with the map; and	The CRD-5500 Manual teaches that the CRD-5500 RAID controller controls access from the hosts to the redundancy groups (storage space on the remote storage devices) in accordance with the mapping tables: "The controller's Host LUN Mapping feature makes it possible to map RAID sets differently to each host. You make the same redundancy group show up on different LUNs to different hosts, or make a redundancy group visible to one host but not to another." CQ-1004, p. 1-1 (emphasis added). "This screen may be used to map LUNs on each host channel to a particular redundancy group. Or you may prevent a redundancy group from appearing on a host channel. Thus, for example, you may map redundancy group 1 to LUN 5 on host channel 0 and the same redundancy group to LUN 12 on host channel 1. Or you may make redundancy group 8 available on LUN 4 on host channel 0 and block access to it on host channel 1." CQ-1004, p. 4-5 (emphasis added). **Thus**, the CRD-5500 RAID controller's Host LUN Mapping feature that controls access from a host to the redundancy groups in accordance with the mapping table associated with the host, as taught by the CRD-5500 Manual, discloses "control access from the devices connected to the first transport medium to the storage space on the remote storage devices in accordance with the map" as recited in the claim.
[1.2d] allow access from devices connected to the first transport medium to the remote storage devices using native low	[1.2d] ***"allow access from devices connected to the first transport medium to the remote storage devices using native low level, block protocol"*** The CRD-5500 Manual teaches that the CRD-5500 RAID controller allows access from hosts connected to a host bus to the remote disk drives using the SCSI protocol (*i.e.*, a native low level block protocol): "The CRD-5500 RAID controller provides high-performance,

level, block protocol.	high-availability access to SCSI disk array subsystems along a Fast/Wide SCSI bus." CQ-1004, p. 1-1. "The CRD-5500 involves several SCSI buses, which can take the form of 8-bit single-ended, 16-bit single-ended, or 16-bit differential buses. On top of this, one or more of the buses may be designated as a host bus for communicating with the host computer, and the others may be populated with disk drives to form the RAID set pool." CQ-1004, p. 2-4 (emphasis added). "The CRD-5500 maintains a running log of controller-related events. These events include the initiation and completion of RAID set creations and rebuilds, drive failures, and SCSI messages emanating from the host and drive channels, among other events." CQ-1004, p. 4-18 (emphasis added) **Additionally**, one of ordinary skill in the art would recognize that if a Fibre Channel host I/O module was added to the CRD-5500 RAID controller and utilized to communicate with hosts via a Fibre Channel transport medium—as suggested by the HP Journal—the CRD-5500 RAID controller would still allow access from the hosts to disk drives using a native low level block protocol because a Fibre Channel I/O module based on the Tachyon chip supports the encapsulation of SCSI I/O commands within Fibre Channel communications: "Networking and I/O protocols, such as SCSI commands, are mapped to Fibre Channel constructs and encapsulated and transported within Fibre Channel frames. The main purpose of Fibre Channel is to have any number of existing protocols operate over a variety of physical media and existing cable plants." CQ-1006, pp. 94-95 (emphasis added). "To provide support for customer mass storage applications, Tachyon: Supports up to 16,384 concurrent SCSI I/O transactions. Can be programmed to function as either an initiator or a target. Assists the protocol for peripheral I/O transactions via SCSI

	encapsulation over Fibre Channel (FCP)." CQ-1006, p. 102 (emphasis added). **Additionally**, if "allow access from devices connected to the first transport medium to the remote storage devices using native low level, block protocol" is given the *3Par* definition of "*permit or deny reading or writing of data using the native low level, block protocol of the Virtual Local Storage without involving a translation from a high level file system command to a native low level, block protocol request*," the above teachings of the CRD-5500 Manual still disclose this limitation. Specifically, as shown in [1.2b(i)], the CRD-5500 permits or denies a host from sending a SCSI command to a specific a redundancy group based on the host's mapping. Further, as shown above, because the hosts and storage devices both communicate with the CRD-5500 RAID controller using SCSI commands, there is no need to translate between file system commands and SCSI commands. **Additionally**, if "native low level block protocol" is given the *3Par* definition of "*a set of rules or standards that enable computers to exchange information and do not involve the overhead of high level protocols and file systems typically required by network servers*," the above teachings of the CRD-5500 Manual still disclose this limitation. Specifically, as shown above, the hosts and storage devices both communicate with the CRD-5500 RAID controller using SCSI commands and not using file system commands or high level storage protocols. **Thus**, the CRD-5500 RAID controller allowing access from hosts to disk drives on a SCSI bus using SCSI I/O commands, as taught by the CRD-5500 Manual, in view of the encapsulation of SCSI I/O commands within Fibre Channel communications, as taught by the HP Journal, renders obvious "allow access from devices connected to the first transport medium to the remote storage devices using native low level, block protocol" as recited in the claim.

Claim 2	
[2.0] The storage router of claim 1,	See [1.0] – [1.2]
[2.1] wherein the map associates a representation of storage space on the remote storage devices with multiple devices connected to the first transport medium.	[2.1] ***"wherein the map associates a representation of storage space on the remote storage devices with multiple devices connected to the first transport medium"*** The CRD-5500 Manual teaches that a redundancy group (one representation of storage space on the remote storage devices) may be associated with multiple host devices: "This screen may be used to map LUNs on each host channel to a particular redundancy group. Or you may prevent a redundancy group from appearing on a host channel. Thus, for example, <u>you may map redundancy group 1 to LUN 5 on host channel 0 and the same redundancy group to LUN 12 on host channel 1</u>. Or you may make redundancy group 8 available on LUN 4 on host channel 0 and block access to it on host channel 1." CQ-1004, p. 4-5 (emphasis added). **Thus**, the CRD-5500 RAID controller's Host LUN Mapping feature that associates a redundancy group with multiple host devices, as taught by the CRD-5500 Manual, discloses "wherein the map associates a representation of storage space on the remote storage devices with multiple devices connected to the first transport medium" as recited in the claim.
Claim 3	
[3.0] The storage router of claim 1,	See [1.0] – [1.2]
[3.1] wherein the storage space on the remote storage devices	[3.1] ***"wherein the storage space on the remote storage devices comprises storage space on multiple remote storage devices"*** The CRD-5500 Manual teaches that the storage space is located on multiple remote disk drives:

comprises storage space on multiple remote storage devices.	"The CRD-5500 affords great flexibility in creating RAID sets. Drives of different sizes and manufacturers may be combined in a RAID set." CQ-1004, p. 3-3 (emphasis added). "RAID sets may comprise any combination of disk drives, regardless of their I/O channel or SCSI ID. In addition, the CRD-5500 defines each RAID set or partition of a RAID set as a 'redundancy group.'" CQ-1004, p. 1-11 (emphasis added). "Our system will have four disk drives on each of the controller's six disk channels, for a total of 24 disk drives. To connect the drives, run a shielded SCSI cable from each of the controller's drive channels, and daisy-chain the drives to each cable."CQ-1004, p. 3-1. 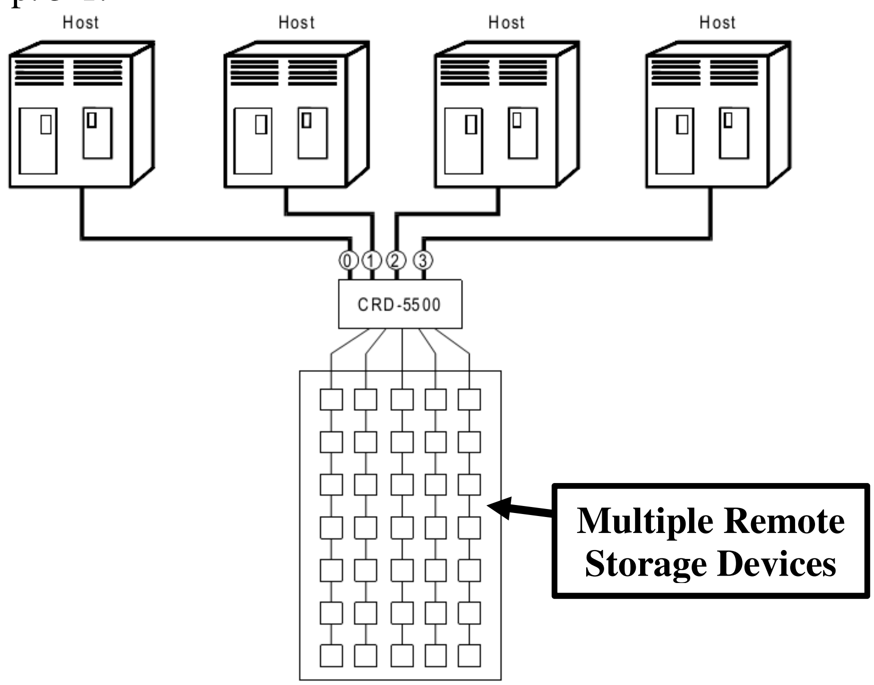 CQ-1004, Figure 1-2 (annotated) **Thus**, the storage space that is located on multiple disk drives, as taught by the CRD-5500 Manual, teaches "wherein the storage space on the remote storage devices comprises storage space on multiple remote storage devices" as recited in the claim.
Claim 4	

[4.0] The storage router of claim 1,	See [1.0] – [1.2].
[4.1] wherein the map associates a representation of a device connected to the first transport medium with a representation of an entire storage space of at least one remote storage device.	[4.1] ***"wherein the map associates a representation of a device connected to the first transport medium with a representation of an entire storage space of at least one remote storage device"*** **First,** as discussed above in [1.2b(i)], the CRD-5500 Manual teaches that a mapping table associates a host channel (representation of a device) with representations of storage space (Host LUNs and redundancy groups): **Second,** the CRD-5500 Manual further teaches that a redundancy group number (a representation of storage space) may be assigned to an entire disk drive: "JBOD, which stands for 'Just a Bunch of Disks,' makes it possible to connect one or more standalone disk drives to the controller. A JBOD disk drive is not part of a redundancy group, even though the controller <u>assigns a redundancy group number to the drive. This number becomes that logical unit number (LUN) that the host will use to address the drive</u>. (You may map any redundancy group number to another LUN using the Host LUN Mapping feature of the monitor utility.)" CQ-1004, p. 1-10 (emphasis added). **Thus**, the mapping table that associates a host with a redundancy group representing an entire disk drive, as taught by the CRD-5500 Manual, discloses "wherein the map associates a representation of a device connected to the first transport medium with a representation of an entire storage space of at least one remote storage device" as recited in the claim.
Claim 5	
[5.0] The storage router of claim 1,	See [1.0] – [1.2].

<table>
<tr>
<td>[5.1] wherein the map resides at the storage router and is maintained at the storage router.</td>
<td>

[5.1] ***"wherein the map resides at the storage router and is maintained at the storage router"***

First, the CRD-5500 Manual teaches that the CRD-5500 RAID Controller (storage router) contains a Monitor Utility for maintaining the map between the host devices and the storage space:

"The CRD-5500 contains a monitor utility that is accessible through the UART/Monitor connector on the back of the controller box." CQ-1004, p. 2-9.

"The CRD-5500's monitor utility provides complete control over the configuration and operation of the controller." CQ-1004, p. 4-1.

The below screenshot of the Monitor Utility shows the map that allocates storage space to hosts:

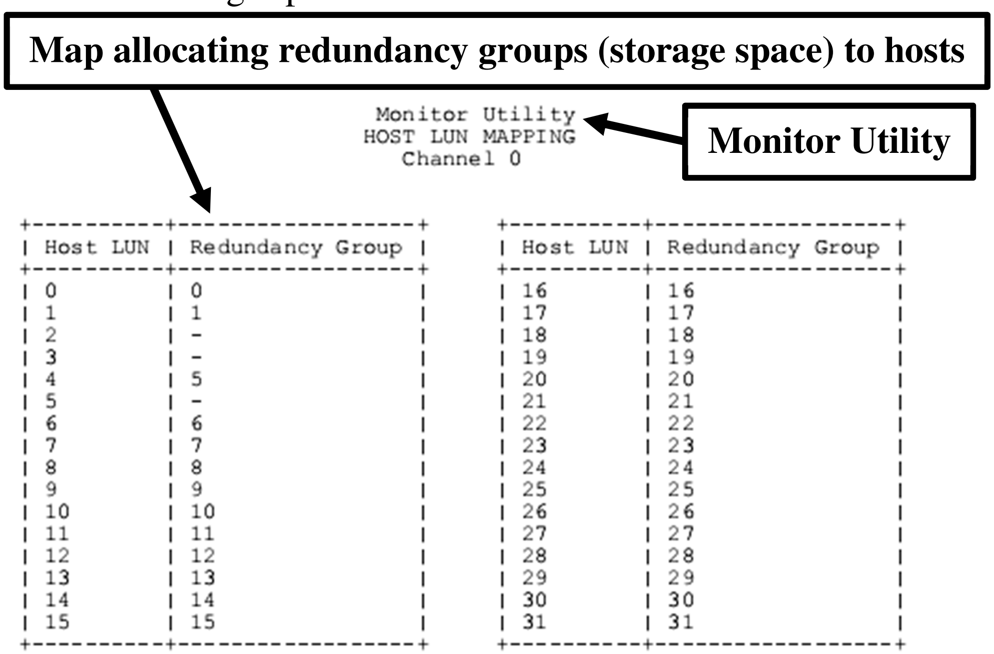

<table>
<tr><th>Host LUN</th><th>Redundancy Group</th><th>Host LUN</th><th>Redundancy Group</th></tr>
<tr><td>0</td><td>0</td><td>16</td><td>16</td></tr>
<tr><td>1</td><td>1</td><td>17</td><td>17</td></tr>
<tr><td>2</td><td>-</td><td>18</td><td>18</td></tr>
<tr><td>3</td><td>-</td><td>19</td><td>19</td></tr>
<tr><td>4</td><td>5</td><td>20</td><td>20</td></tr>
<tr><td>5</td><td>-</td><td>21</td><td>21</td></tr>
<tr><td>6</td><td>6</td><td>22</td><td>22</td></tr>
<tr><td>7</td><td>7</td><td>23</td><td>23</td></tr>
<tr><td>8</td><td>8</td><td>24</td><td>24</td></tr>
<tr><td>9</td><td>9</td><td>25</td><td>25</td></tr>
<tr><td>10</td><td>10</td><td>26</td><td>26</td></tr>
<tr><td>11</td><td>11</td><td>27</td><td>27</td></tr>
<tr><td>12</td><td>12</td><td>28</td><td>28</td></tr>
<tr><td>13</td><td>13</td><td>29</td><td>29</td></tr>
<tr><td>14</td><td>14</td><td>30</td><td>30</td></tr>
<tr><td>15</td><td>15</td><td>31</td><td>31</td></tr>
</table>

CQ-1004, p. 4-5 (annotated)

Thus, the Monitor Utility contained on the CRD-5500 RAID Controller that maintains the map, as taught by the CRD-5500 Manual, discloses "wherein the map resides at the storage router

</td>
</tr>
</table>

	and is maintained at the storage router" as recited in the claim.
Claim 6	
[6.0] The storage router of claim 1,	See [1.0] - [1.2]
[6.1] wherein the native low level block protocol is received at the storage router via the first transport medium and the processing device uses the received native low level block protocol to allow the devices connected to the first transport medium access to storage space specifically allocated to them in the map.	[6.1] ***"wherein the native low level block protocol is received at the storage router via the first transport medium and the processing device uses the received native low level block protocol to allow the devices connected to the first transport medium access to storage space specifically allocated to them in the map"*** **First,** the CRD-5500 Manual teaches that the CRD-5500 RAID controller receives SCSI (a native low level block protocol) I/O to allow host devices connected to a host bus (first transport medium) to access storage space on remote disk drives: "<u>The CRD-5500 involves several SCSI buses</u>, which can take the form of 8-bit single-ended, 16-bit single-ended, or 16-bit differential buses. On top of this, <u>one or more of the buses may be designated as a host bus for communicating with the host computer, and the others may be populated with disk drives to form the RAID set pool</u>." CQ-1004, p. 2-4 (emphasis added). "The CRD-5500 maintains a running log of controller-related events. These events include the initiation and completion of RAID set creations and rebuilds, drive failures, and <u>SCSI messages emanating from the host and drive channels</u>, among other events." CQ-1004, p. 4-18 (emphasis added). **Second,** one of ordinary skill in the art would recognize that if a Fibre Channel host I/O module was added to the CRD-5500 RAID controller and utilized to communicate with hosts via a Fibre Channel transport medium—as suggested by the HP Journal—the CRD-5500 RAID controller would allow hosts to communicate with the CRD-5500 RAID controller via SCSI encapsulation over Fibre Channel Protocol (FCP):

	"Networking and I/O protocols, such as SCSI commands, are mapped to Fibre Channel constructs and encapsulated and transported within Fibre Channel frames. The main purpose of Fibre Channel is to have any number of existing protocols operate over a variety of physical media and existing cable plants." CQ-1006, pp. 94-95 (emphasis added). "To provide support for customer mass storage applications, Tachyon: Supports up to 16,384 concurrent SCSI I/O transactions. Can be programmed to function as either an initiator or a target. Assists the protocol for peripheral I/O transactions via SCSI encapsulation over Fibre Channel (FCP)." CQ-1006, p. 102 (emphasis added). **Third**, as discussed in [1.2b(i)] – [1.2d], the CRD-5500 Manual teaches that the CRD-5500 RAID controller allocates redundancy groups to specific host devices using a map, and then allows a host device to access the specific redundancy groups allocated to it: "Figure 1-2 shows how you can connect as many as four hosts to the CRD-5500. By using the controller's Host LUN Mapping feature, you can assign redundancy groups to a particular host." CQ-1004, p. 1-2 (emphasis added). "The controller's Host LUN Mapping feature makes it possible to map RAID sets differently to each host. You make the same redundancy group show up on different LUNs to different hosts, or make a redundancy group visible to one host but not to another." CQ-1004, p. 1-1. "This screen may be used to map LUNs on each host channel to a particular redundancy group. Or you may prevent a redundancy group from appearing on a host channel. Thus, for example, you may map redundancy group 1 to LUN 5 on host channel 0 and the same redundancy group to LUN 12 on host channel 1. Or you

	may make redundancy group 8 available on LUN 4 on host channel 0 and block access to it on host channel 1." CQ-1004, p. 4-5. **Thus**, the CRD-5500 RAID controller receiving SCSI I/O commands and using the SCSI I/O commands to allow hosts to access redundancy groups (storage space) specifically allocated to the hosts in the map, as taught by the CRD-5500 Manual, in view of the encapsulation of SCSI I/O commands within Fibre Channel communications, as taught by the HP Journal, renders obvious "wherein the native low level block protocol is received at the storage router via the first transport medium and the processing device uses the received native low level block protocol to allow the devices connected to the first transport medium access to storage space specifically allocated to them in the map" as recited in the claim.
Claim 7	
[7.0] The storage router of claim 1,	See [1.0] – [1.2]
[7.1] wherein the storage router is configured to receive commands according to a first low level block protocol from the device connected to the first transport medium and forward commands according to a	[7.1] ***"wherein the storage router is configured to receive commands according to a first low level block protocol from the device connected to the first transport medium and forward commands according to a second low level block protocol to the remote storage devices"*** **First,** the CRD-5500 Manual teaches that the CRD-5500 RAID controller receives SCSI (a low level block protocol) I/O commands from host devices connected to a host bus (first transport medium) and relays those commands to a SCSI disk drive bus: "The CRD-5500 involves several SCSI buses, which can take the form of 8-bit single-ended, 16-bit single-ended, or 16-bit differential buses. On top of this, one or more of the buses may be designated as a host bus for communicating with the host computer, and the others may be populated with disk drives to

second low level block protocol to the remote storage devices.	form the RAID set pool." CQ-1004, p. 2-4 (emphasis added). "The CRD-5500 maintains a running log of controller-related events. These events include the initiation and completion of RAID set creations and rebuilds, drive failures, and SCSI messages emanating from the host and drive channels, among other events." CQ-1004, p. 4-18 (emphasis added). **Second,** one of ordinary skill in the art would recognize that if a Fibre Channel host I/O module was added to the CRD-5500 RAID controller and utilized to communicate with hosts via a Fibre Channel transport medium—as suggested by the HP Journal—the CRD-5500 RAID controller would allow hosts to communicate with the CRD-5500 RAID controller via SCSI encapsulation over Fibre Channel Protocol (FCP) (a low level block protocol): "Networking and I/O protocols, such as SCSI commands, are mapped to Fibre Channel constructs and encapsulated and transported within Fibre Channel frames. The main purpose of Fibre Channel is to have any number of existing protocols operate over a variety of physical media and existing cable plants." CQ-1006, pp. 94-95 (emphasis added). "To provide support for customer mass storage applications, Tachyon: Supports up to 16,384 concurrent SCSI I/O transactions. Can be programmed to function as either an initiator or a target. Assists the protocol for peripheral I/O transactions via SCSI encapsulation over Fibre Channel (FCP)." CQ-1006, p. 102 (emphasis added). **Accordingly**, the CRD-5500 RAID Controller—as modified by the HP Journal—would (i) receive FCP messages (first low level block protocol) containing SCSI commands from hosts over Fibre Channel, (ii) remove the SCSI commands from the FCP messages, and (iii) forward the SCSI commands (second low level block protocol) to the disk drives.

	Thus, the CRD-5500 RAID controller receiving FCP messages encapsulating SCSI commands from the hosts and forwarding the SCSI commands to the remote storage devices, as taught by the CRD-5500 Manual in view of the HP Journal, renders obvious "wherein the storage router is configured to receive commands according to a first low level block protocol from the device connected to the first transport medium and forward commands according to a second low level block protocol to the remote storage devices" as recited in the claim.
Claim 8	
[8.0] The storage router of claim 7,	See [7.0] – [7.1]
[8.1] wherein the first low level block protocol is an FCP protocol and the second low level block protocol is a protocol other than FCP.	See [7.1]
Claim 9	
[9.0] The storage router of claim 1,	See [1.0] – [1.2]
[9.1] wherein the map comprises one or more tables.	[9.1] ***"wherein the map comprises one or more tables"*** As discussed in [1.2b(i)], the CRD-5500 Manual teaches that the map maintained by the Host LUN Mapping feature of the CRD-5500 RAID Controller includes a mapping table associated with each host (represented by a channel number):

"This screen may be used to map LUNs on each host channel to a particular redundancy group. Or you may prevent a redundancy group from appearing on a host channel. Thus, for example, you may map redundancy group 1 to LUN 5 on host channel 0 and the same redundancy group to LUN 12 on host channel 1. Or you may make redundancy group 8 available on LUN 4 on host channel 0 and block access to it on host channel 1." CQ-1004, p. 4-5.

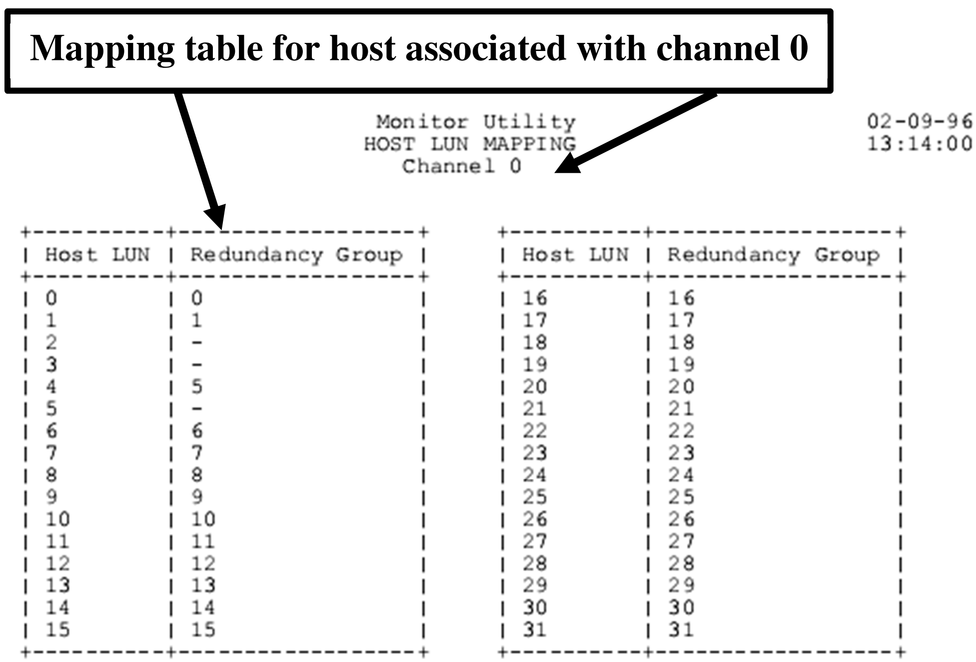

Host LUN	Redundancy Group
0	0
1	1
2	-
3	-
4	5
5	-
6	6
7	7
8	8
9	9
10	10
11	11
12	12
13	13
14	14
15	15

Host LUN	Redundancy Group
16	16
17	17
18	18
19	19
20	20
21	21
22	22
23	23
24	24
25	25
26	26
27	27
28	28
29	29
30	30
31	31

CQ-1004, p. 4-5 (annotated)

Thus, the mapping tables of the Host LUN Mapping feature in the CRD-5500 RAID Controller, as taught by the CRD-5500 Manual, discloses "wherein the map comprises one or more tables" as recited in the claim.

Claim 10	
[10.0] The storage router of claim 1,	See [1.0] – [1.2]
[10.1] wherein the virtual local	[10.1] ***"wherein the virtual local storage is provided to the devices connected to the first transport medium in a manner***

storage is provided to the devices connected to the first transport medium in a manner that is transparent to the devices and wherein the storage space allocated to the devices connected to the first transport medium appears to the devices as local storage.	***that is transparent to the devices and wherein the storage space allocated to the devices connected to the first transport medium appears to the devices as local storage"*** The CRD-5500 Manual teaches that the CRD-5500 RAID controller provides virtual local storage to the connected hosts in a transparent manner. First, each host has a set of addressable virtual LUNs to which redundancy groups (storage space) are mapped. The hosts address the storage space via the LUNs: "RAID sets may comprise any combination of disk drives, regardless of their I/O channel or SCSI ID. In addition, the CRD-5500 defines each RAID set or partition of a RAID set as a 'redundancy group.'" CQ-1004, p. 1-11 (emphasis added). "These redundancy groups may be mapped to host LUNs, either in a direct one-to-one relationship or in a manner defined by the user." CQ-1004, p. 1-11 (emphasis added). "Each redundancy group will have its own redundancy group number, which corresponds to the Logical Unit Number (LUN) that the host will use to address the partition, unless you map the redundancy group to another LUN with the Host LUN Mapping screen." CQ-1004, p. 4-2 (emphasis added). **Second**, this mapping is transparent to the hosts because, from the host's point of view, only the virtual storage associated with LUNs exist. Hence, the hosts cannot address the actual physical storage mapped to the LUNs: "In this example, Redundancy Group numbers 1 through 4 are associated with RAID Set number 1. From the controller's point of view in this example, Redundancy Group number 0 belongs to RAID Set 0, which is configured for RAID Level 4, and Redundancy Group numbers 1 through 4 belong to RAID Set 1, which is configured for Level 5. From the host's point of view, there is no such thing as RAID Sets 0 and 1, only (in this

	example) LUNs 0 through 4, which the CRD-5500 sees as Redundancy Groups 0 through 4." CQ-1004, p. 3-6 (emphasis added). **Third**, the virtual storage spaces assigned to the hosts appear to the hosts as local disk drives: "Since each partition will appear to the host as a different disk drive, it must have its own LUN." CQ-1004, p. 3-6 (emphasis added). **Thus**, CRD-5500 RAID Controller providing virtual local storage to hosts through addressable virtual LUNs that are transparently mapped to the redundancy groups (storage space) that appear as different local disk drives to the hosts, as taught by the CRD-5500 Manual, discloses "wherein the virtual local storage is provided to the devices connected to the first transport medium in a manner that is transparent to the devices and wherein the storage space allocated to the devices connected to the first transport medium appears to the devices as local storage" as recited in the claim.
Claim 11	
[11.0] The storage router of claim 1,	See [1.0] – [1.2]
[11.1] wherein the storage router provides centralized control of what the devices connected to the first transport medium see as local storage.	[11.1] ***"wherein the storage router provides centralized control of what the devices connected to the first transport medium see as local storage"*** The CRD-5500 Manual teaches that the CRD-5500 RAID controller includes a Monitor Utility that provides centralized control of the mapping of between the host devices and the storage space, wherein the mapping determines what storage space is visible to the hosts: "The CRD-5500's monitor utility provides complete control over the configuration and operation of the controller." CQ-1004, p.

	4-1 (emphasis added). "The controller's Host LUN Mapping feature makes it possible to map RAID sets differently to each host. You make the same redundancy group show up on different LUNs to different hosts, or make a redundancy group visible to one host but not to another." CQ-1004, p. 1-1 (emphasis added). "This screen may be used to map LUNs on each host channel to a particular redundancy group. Or you may prevent a redundancy group from appearing on a host channel. Thus, for example, you may map redundancy group 1 to LUN 5 on host channel 0 and the same redundancy group to LUN 12 on host channel 1. Or you may make redundancy group 8 available on LUN 4 on host channel 0 and block access to it on host channel 1." CQ-1004, p. 4-5 (emphasis added). **Thus**, the Monitor Utility of the CRD-5500 RAID Controller that provides centralized control of what storage space is visible to the hosts, as taught by the CRD-5500 Manual, discloses "wherein the storage router provides centralized control of what the devices connected to the first transport medium see as local storage" as recited in the claim.
Claim 12	
[12.0] The storage router of claim 1,	See [1.0] – [1.2]
[12.1] wherein the representations of storage space comprise logical unit numbers that represent a subset of storage on the	[12.1] ***"wherein the representations of storage space comprise logical unit numbers that represent a subset of storage on the remote storage devices"*** **First**, as discussed above in [1.2b(i)], the CRD-5500 Manual teaches that the mapping tables maintained by the Monitor Utility include representations of storage space associated with hosts. **Second**, the CRD-5500 Manual teaches that the representations

remote storage devices.	of storage space in the mapping tables include Host LUNs (logical unit numbers) that the host uses to address subsets of storage space on storage devices. "Each redundancy group will have its own redundancy group number, which corresponds to the Logical Unit Number (LUN) that the host will use to address the partition, unless you map the redundancy group to another LUN with the Host LUN Mapping screen." CQ-1004, p. 4-2 (emphasis added). "These redundancy groups may be mapped to host LUNs, either in a direct one-to-one relationship or in a manner defined by the user." CQ-1004, p. 1-11 (emphasis added).

Host LUNs (logical unit numbers)

```
                         Monitor Utility                         02-09-96
                        HOST LUN MAPPING                         13:14:00
                           Channel 0

+----------+-----------------+       +----------+-----------------+
| Host LUN | Redundancy Group |      | Host LUN | Redundancy Group |
+----------+-----------------+       +----------+-----------------+
| 0        | 0               |       | 16       | 16              |
| 1        | 1               |       | 17       | 17              |
| 2        | -               |       | 18       | 18              |
| 3        | -               |       | 19       | 19              |
| 4        | 5               |       | 20       | 20              |
| 5        | -               |       | 21       | 21              |
| 6        | 6               |       | 22       | 22              |
| 7        | 7               |       | 23       | 23              |
| 8        | 8               |       | 24       | 24              |
| 9        | 9               |       | 25       | 25              |
| 10       | 10              |       | 26       | 26              |
| 11       | 11              |       | 27       | 27              |
| 12       | 12              |       | 28       | 28              |
| 13       | 13              |       | 29       | 29              |
| 14       | 14              |       | 30       | 30              |
| 15       | 15              |       | 31       | 31              |
+----------+-----------------+       +----------+-----------------+
```

CQ-1004, p. 4-5 (annotated)

Thus, the representations of storage space in the mapping tables that include Host LUNs, which are used by the hosts to address subsets of storage on disk drives, as taught by the CRD-5500 Manual, disclose "wherein the representations of storage space comprise logical unit numbers that represent a subset of storage on the remote storage devices" as recited in the claim.

Claim 13	
[13.0] The storage router of claim 12,	See [12.0] – [12.1]
[13.1] wherein the storage router is operable to route requests to the same logical unit number from different devices connected to the first transport medium to different subsets of storage space on the remote storage devices.	[13.1] ***"wherein the storage router is operable to route requests to the same logical unit number from different devices connected to the first transport medium to different subsets of storage space on the remote storage devices"*** **First**, as discussed in [1.2b(i)], the CRD-5500 Manual teaches that the CRD-5500 RAID Controller routes I/O commands from each host to redundancy groups (subsets of storage space) based on the Host LUN Mapping table associated with the host. **Second**, the CRD-5500 Manual further teaches that each map has a set of 32 addressable virtual LUNs to which these redundancy groups are mapped: "These redundancy groups may be mapped to host LUNs, either in a direct one-to-one relationship or in a manner defined by the user." CQ-1004, p. 1-11 (emphasis added). Further, a host accesses the redundancy groups via its virtual host LUNs: "Each redundancy group will have its own redundancy group number, which corresponds to the Logical Unit Number (LUN) that the host will use to address the partition, unless you map the redundancy group to another LUN with the Host LUN Mapping screen." CQ-1004, p. 4-2 (emphasis added).

Addressable host LUNs mapped to redundancy groups

Monitor Utility
HOST LUN MAPPING
Channel 0

02-09-96
13:14:00

Host LUN	Redundancy Group	Host LUN	Redundancy Group
0	0	16	16
1	1	17	17
2	-	18	18
3	-	19	19
4	5	20	20
5	-	21	21
6	6	22	22
7	7	23	23
8	8	24	24
9	9	25	25
10	10	26	26
11	11	27	27
12	12	28	28
13	13	29	29
14	14	30	30
15	15	31	31

CQ-1004, p. 4-5 (annotated)

In the above mapping table for the host associated with channel 0, a user has mapped LUN 0 to redundancy group 0.

Third, the CRD-5500 Manual teaches that the Host LUNs for each host are independently configurable. For instance, Host LUN 0 in the mapping table for a first host attached to Channel 0 may be mapped to a different redundancy group than Host LUN 0 in the mapping table for a second host attached to Channel 1:

"You may use the Host LUN Mapping screen under Setup parameters to give each host channel a different view of the redundancy groups connected to the controller. For instance, you could have host channel 0 see a redundancy group as LUN 1, and host channel 1 see the same redundancy group as LUN 5." CQ-1004, p. 4-10 (emphasis added).

"The controller's Host LUN Mapping feature makes it possible to map RAID sets differently to each host. You make the same redundancy group show up on different LUNs to different hosts, or make a redundancy group visible to one host but not to another." CQ-1004, p. 1-1 (emphasis added).

	"This screen may be used to map LUNs on each host channel to a particular redundancy group. Or you may prevent a redundancy group from appearing on a host channel. <u>Thus, for example, you may map redundancy group 1 to LUN 5 on host channel 0 and the same redundancy group to LUN 12 on host channel 1. Or you may make redundancy group 8 available on LUN 4 on host channel 0 and block access to it on host channel 1."</u> CQ-1004, p. 4-5 (emphasis added). Accordingly, it is simply a matter of setup and configuration as to whether a particular redundancy group is assigned the same LUN for a given host. Each host has the same LUNs (0-31), and whether a particular LUN (e.g., LUN 4) for a host on channel 0 and for a host on channel 1 is assigned to the same or different redundancy is a direct option based on the the intended configuration of the CRD-5500 controller. A person of skill in the art would understand and know that the router is operable to route requests to the same LUN to different storage devices. **Thus,** the CRD-5500 RAID Controller routing I/O requests to the same Host LUN from different hosts to different redundancy groups (subsets of storage space), as taught by the CRD-5500 Manual, discloses "wherein the storage router is operable to route requests to the same logical unit number from different devices connected to the first transport medium to different subsets of storage space on the remote storage devices" as recited in the claim.
Claim 14	
[14.0] The storage router of claim 1,	See [1.0] – [1.2]
[14.1] wherein the representations of devices connected to	[14.1] ***"wherein the representations of devices connected to the first transport medium are unique identifiers"*** **First**, as discussed in [1.2b(i)], the CRD-5500 Manual teaches that hosts connected to CRD-5500 RAID Controller are

<table>
<tr>
<td>the first transport medium are unique identifiers.</td>
<td>represented by channel numbers.

Second, the CRD-5500 Manual teaches that hosts may be assigned to four unique channel identifiers:

"<u>The CRD-5500 permits the I/O modules in slots 1, 2, and 3 to be configured as host or disk channel modules</u>. The Channel Settings screen is the place to configure these modules. Use the up and down arrow keys to maneuver to the channel you wish to configure and press Enter. Then use the arrow keys to toggle between host and disk. Press Enter again to save your selection. <u>Channel 0 is always a host channel</u> and channels 4 through 6 are always disk channels, so the monitor utility will restrict access to these fields." CQ-1004, p. 4-5 (emphasis added).

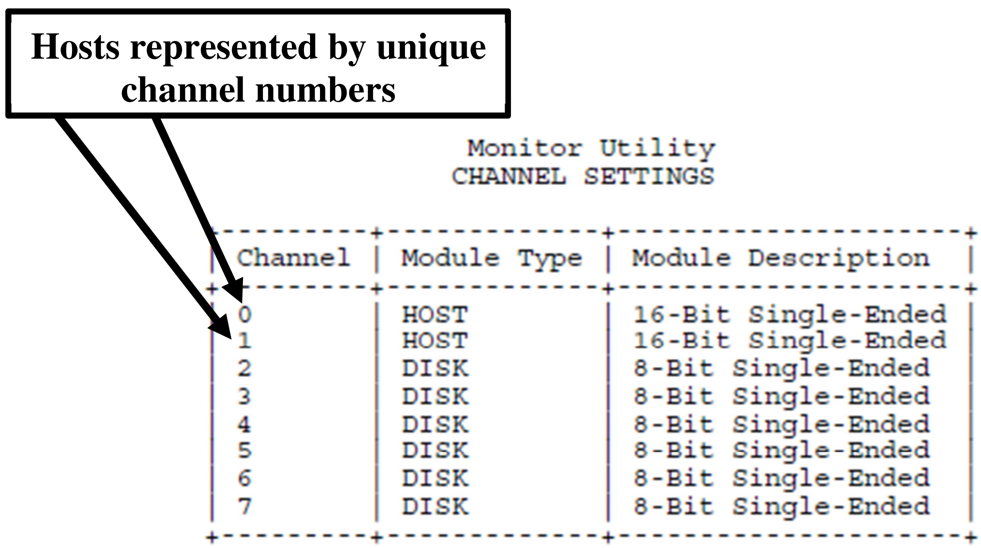

CQ-1004, p. 3-2 (annotated)
In the example above, the CRD-5500 RAID controller uniquely identifies two host devices with channels 0 and 1.

Thus, the channel numbers that uniquely identify hosts, as taught by the CRD-5500 Manual, disclose "wherein the representations of devices connected to the first transport medium are unique identifiers" as recited in the claim.</td>
</tr>
<tr>
<td colspan="2">**Claim 16**</td>
</tr>
</table>

[16.0] The storage router of claim 1,	See [1.0] – [1.2]
[16.1] wherein the storage router is configured to allow modification of the map in a manner transparent to and without involvement of the devices connected to the first transport medium.	[16.1] ***"wherein the storage router is configured to allow modification of the map in a manner transparent to and without involvement of the devices connected to the first transport medium"*** **First**, as discussed above in [10.1], the CRD-5500 Manual teaches that the mapping between hosts and storage space is transparent to hosts. As such, any modification to the mapping is also transparent to hosts. **Second**, the CRD-5500 Manual teaches that the map is modifiable through the Monitor Utility, which is accessible via a terminal or computer attached to UART/Monitor connector on the back of the CRD-5500 RAID Controller. Such a terminal or computer is independent of the hosts attached to the first transport medium: "The CRD-5500's monitor utility provides complete control over the configuration and operation of the controller. You may also use the utility to view the status of the controller's RAID sets and monitor the progress of create and rebuild operations." CQ-1004, p. 4-1. "The <u>CRD-5500 contains a monitor utility that is accessible through the UART/Monitor connector on the back of the controller box</u>. If the your [sic] CRD-5500 is part of an enclosure, there is probably a serial cable connector on the exterior of the box. <u>To use the monitor utility, connect a serial cable from the enclosure to a terminal or a computer equipped with terminal emulation software</u>." CQ-1004, p. 2-9 (emphasis added). "<u>You may connect the serial cable to a PC or laptop and access the monitor utility</u> through a communications application such as the Terminal program in Microsoft Windows." CQ-1004, p. 2-9

	(emphasis added). 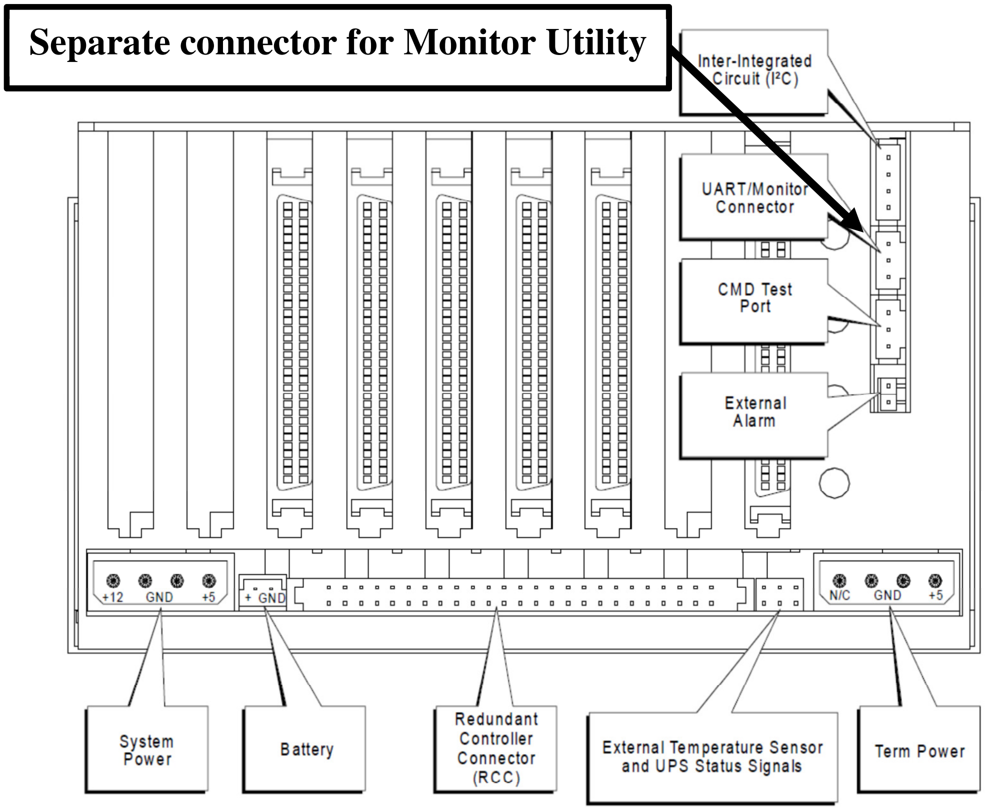 **Figure 2-6: Back panel connectors** CQ-1004, p. 3-2 (annotated) **Thus,** the CRD-5500 RAID Controller that allows modification of the map by a terminal, laptop, or PC connected to a separate UART/Monitor connector in a manner transparent to the hosts, as taught by the CRD-5500 Manual, discloses "wherein the storage router is configured to allow modification of the map in a manner transparent to and without involvement of the devices connected to the first transport medium" as recited in the claim..
Claim 17	
[17.0] The storage router of claim 1,	See [1.0] – [1.2]

[17.1] wherein the processing device is a microprocessor .	[17.1] ***"wherein the processing device is a microprocessor"*** The CRD-5500 Manual teaches that the processing device in the CRD-5500 RAID Controller is a MIPS R300 RISC central processing unit (*i.e.*, a "microprocessor"): "CPU Interface ASIC: Supports the controller's MIPS R3000 RISC central processing unit." CQ-1004, p. 1-3 (emphasis added). 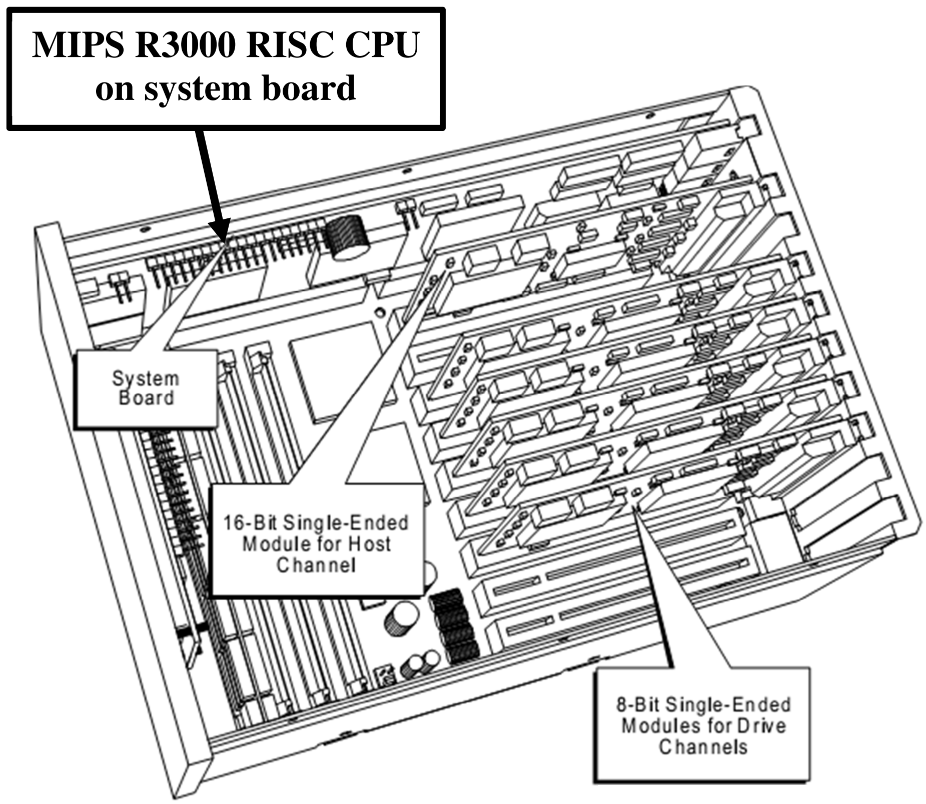 CQ-1004, Figure 2-1 (annotated) **Thus,** the MIPS R300 RISC central processing unit in the CRD-5500, as taught by the CRD-5500 Manual, teaches "wherein the processing device is a microprocessor" as recited in the claim.
Claim 18	
[18.0] The storage router	See [1.0] – [1.2]

of claim 1,	
[18.1] wherein the processing device is a microprocessor and associated logic to implement a stand-alone processing system.	[18.1] ***"wherein the processing device is a microprocessor and associated logic to implement a stand-alone processing system"*** **First**, the CRD-5500 Manual teaches that the processing device in the CRD-5500 RAID Controller is a MIPS R300 RISC central processing unit (*i.e.*, a "microprocessor"): "To increase performance and reliability, the CRD-5500's core functions have been encapsulated in four custom ASIC (Application Specific Integrated Circuits) components. . . . CPU Interface ASIC: Supports the controller's MIPS R3000 RISC central processing unit." CQ-1004, p. 1-3 (emphasis added). 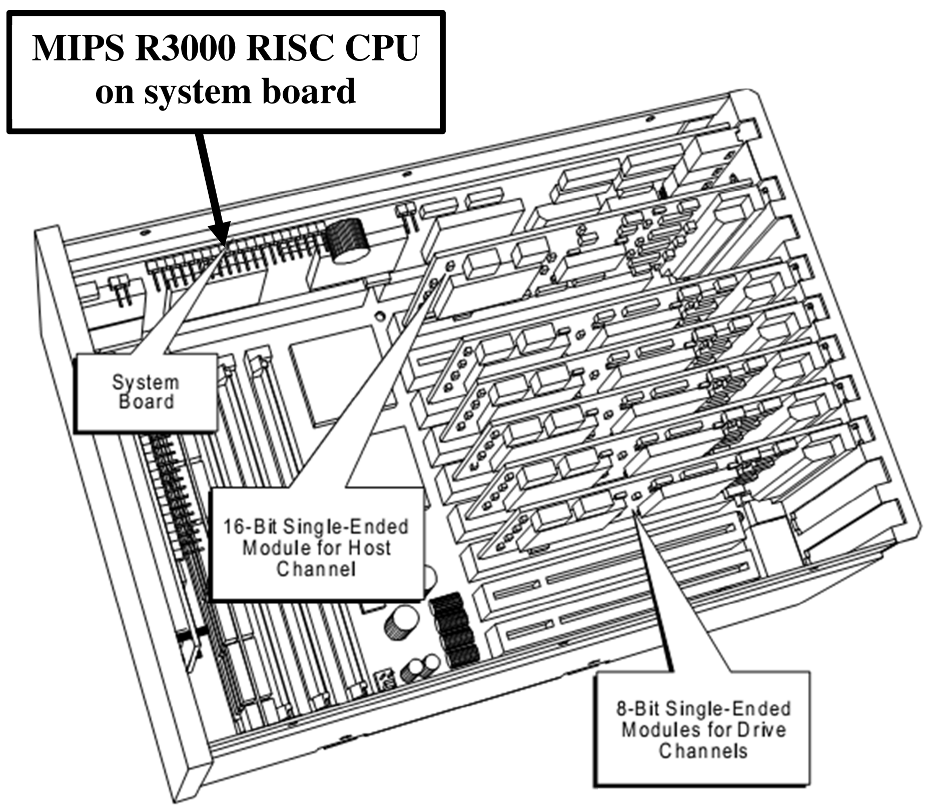 CQ-1004, Figure 2-1 (annotated)

	Second, the CRD-5500 Manual teaches that the CRD-5500 RAID controller includes executable firmware with a Monitor Utility that controls the operation of the controller: "In addition to its flexible hardware design, the CRD-5500's firmware offers the user the flexibility to configure RAID sets in many different ways . . ."CQ-1004, p. 1-1 (emphasis added). "To update the flash EPROM containing the controller's firmware, follow the directions in this section. (The version number of your current firmware is displayed in the title box of the monitor utility's opening screen.)" CQ-1004, p. 4-14 (emphasis added). "The CRD-5500's monitor utility provides complete control over the configuration and operation of the controller." CQ-1004, p. 4-1. **Accordingly**, the RISC processor and associated firmware together implement the stand-alone CRD-5500 RAID Controller storage processing system. **Thus**, the MIPS R3000 RISC central processing unit together with the associated firmware that implement the CRD-5500 RAID controller system, as taught by the CRD-5500 Manual, teaches "wherein the processing device is a microprocessor and associated logic to implement a stand-alone processing system" as recited in the claim.
Claim 19	
[19.0] The storage router of claim 1,	See [1.0] – [1.2]
[19.1] wherein the first transport medium is a fibre channel	[19.1] ***"wherein the first transport medium is a fibre channel transport medium and further comprising a second transport medium connected to the remote storage devices that is a fibre channel transport medium"***

transport medium and further comprising a second transport medium connected to the remote storage devices that is a fibre channel transport medium.	**First**, as discussed in [1.1], one of ordinary skill in the art would have been motivated to replace the SCSI I/O host modules and associated SCSI buses in the CRD-5500 RAID Controller with a Fibre Channel I/O host module and Fibre Channel link (first transport medium). **Second**, for the same reasons as discussed in [1.1] (increased bandwidth, simplified cabling, etc), one of ordinary skill in the art would have been motivated to also replace the SCSI I/O disk modules and associated SCSI buses in the CRD-5500 RAID Controller with one or more Fibre Channel I/O modules and Fibre Channel links (second transport medium). Furthermore, the Fibre Channel I/O disk modules overcome the SCSI bus limitation of addressing only 7 or 15 SCSI disk drives, and allow additional disk drives to be housed in different enclosures. **Thus**, Fibre Channel transport medium connecting the hosts to the CRD-5500 RAID Controller and the Fibre Channel transport medium connecting the disks to the controller, as taught by the CRD-5500 Manual in view of the HP Journal, renders obvious "wherein the first transport medium is a fibre channel transport medium and further comprising a second transport medium connected to the remote storage devices that is a fibre channel transport medium" as recited in the claim.

Claim 20	
[20.0] A storage network comprising:	[20.0] ***"A storage network"*** The CRD-5500 Manual describes the operation of the CRD-5500 RAID controller, which routes data between hosts and remote SCSI disk arrays in a storage network: "The CRD-5500 RAID controller provides high-performance, high-availability access to SCSI disk array subsystems along a Fast/Wide SCSI bus. With a modular hardware design and an intuitive configuration utility, the controller may be tailored to suit a wide range of storage needs, now and in the future." CQ-

	1004, p. 1-1. "Figure 1-2 shows how you can connect as many as four hosts to the CRD-5500." CQ-1004, p. 1-2. 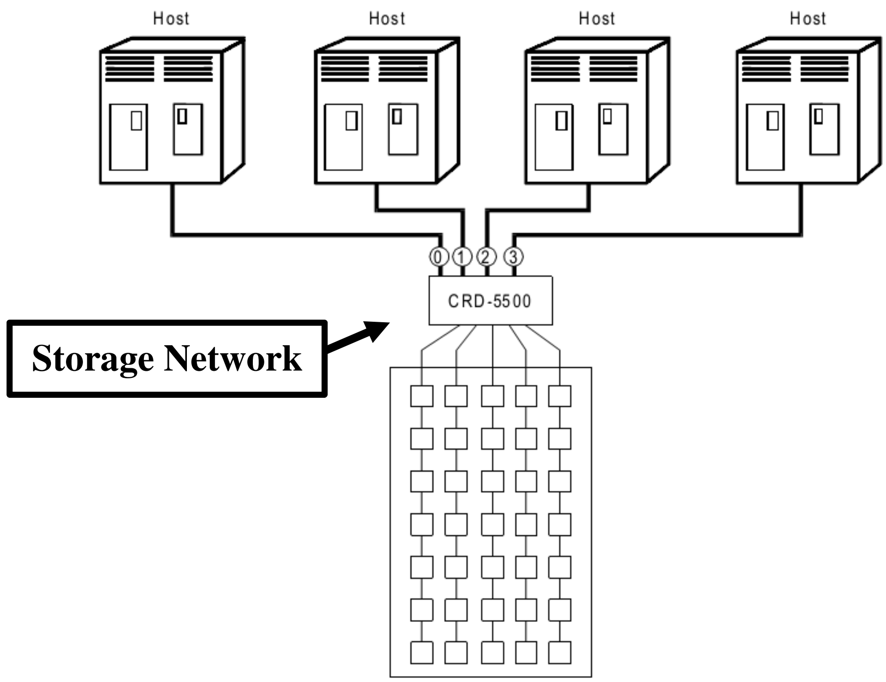 CQ-1004, Figure 1-2 (annotated) **Thus**, the storage network in which the CRD-5500 RAID controller routes data between hosts and remote disk drives, as taught by the CRD-5500 Manual, discloses "a storage network," as recited in the claim.
[20.1] a set of devices connected to a first transport medium, wherein the first transport medium is a serial transport medium;	[20.1] ***"a set of devices connected to a first transport medium, wherein the first transport medium is a serial transport medium"*** **First,** the CRD-5500 Manual teaches that a set of hosts are coupled to the CRD-5500 RAID controller via a host bus: "<u>Figure 1-2 shows how you can connect as many as four hosts to the CRD-5500</u>. By using the controller's Host LUN Mapping feature, you can assign redundancy groups to a particular host. This makes it possible for hosts running incompatible operating systems to use the same CRD-5500 controller." CQ-1004, p. 1-2

(emphasis added).

"The CRD-5500 involves several SCSI buses, which can take the form of 8-bit single-ended, 16-bit single-ended, or 16-bit differential buses. On top of this, one or more of the buses may be designated as a host bus for communicating with the host computer, and the others may be populated with disk drives to form the RAID set pool." CQ-1004, p. 2-4 (emphasis added).

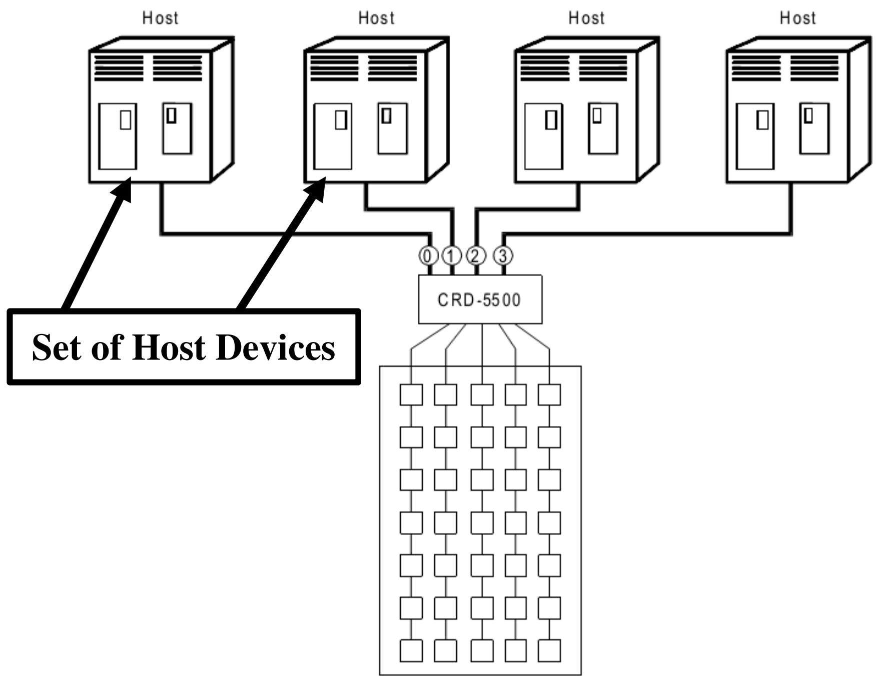

CQ-1004, Figure 1-2 (annotated)

Second, the HP Journal teaches that there are advantages to using a serial Fibre Channel transport medium instead of a SCSI bus transport medium:

"Additionally, as the number of interconnects between systems and I/O devices continues to increase, I/O channels become bottlenecks to system performance. For all these reasons, today's parallel bus architectures are reaching their limits. In the search for a higher-performance serial interface, HP chose Fibre Channel because it overcomes the limitations mentioned above by supporting sustained gigabit data transfer rates." CQ-1006, p.

	5. "A channel such as SCSI (Small Computer Systems Interface), which operates at a maximum throughput of 20 megabytes per second in fast and wide mode, simply cannot keep pace with ever-increasing processor speeds and data rate requirements." CQ-1006, p. 99. "Current peripheral interconnect protocols are limited in the number of devices they can interconnect. For example, parallel SCSI can connect eight devices and 16-bit wide SCSI can connect 16 devices. In addition, peripheral connectors are becoming too large to fit into the shrinking footprints of systems and peripherals. Other SCSI limitations include half-duplex operation only, lack of a switching capability, inability to interconnect individual buses, and the need for customized drivers and adapters for various types of attached devices." CQ-1006, p. 99. "Fibre Channel is a flexible, scalable, high-speed data transfer interface that can operate over a variety of both copper wire and optical fiber at data rates up to 250 times faster than existing communications interfaces." CQ-1006, p. 94. "A single 100-Mbyte/s Fibre Channel port can replace five 20-Mbyte/s SCSI ports, in terms of raw through put. Fibre Channel provides a total network bandwidth of about one gigabit per second." CQ-1006, p. 94. "Another advantage of Fibre Channel is that it uses small connectors. The serial connectors used for Fibre Channel are a fraction of the size of SCSI parallel connectors and have fewer pins, thereby reducing the likelihood of physical damage. Also, depending on the topology, many more devices can be interconnected on Fibre Channel than on existing channels." CQ-1006, p. 94. "Fibre Channel's increased bandwidth provides distance flexibility, increased addressability, and simplified cabling. Fibre

	Channel has versatility, room for growth, and qualified vendor support. Mass storage suppliers are using Fibre Channel to interconnect subsystems and systems and to control embedded disk drives." CQ-1006, p. 99. **Third**, the HP Journal teaches one of ordinary skill in the art how to implement a Fibre Channel I/O adapter board using the Tachyon controller chip that interfaces with a Fibre Channel transport medium: "The system interconnect laboratory of the HP Networked Computing Division became interested in Fibre Channel in 1993 as a method of entering the high-speed serial interconnect market because Fibre Channel was the first technology that could be used for both networking and mass storage. . . . The Tachyon chip (Fig. 3) implements the FC-1 and FC-2 layers of the five-layer Fibre Channel standard (see Article 11)." CQ-1006, p. 101 (emphasis added). "Tachyon's host attach enables low-cost gigabit host adapters on industry-standard buses including PCI, PMC, S-Bus, VME, EISA, Turbo Channel, and MCA. It is easily adaptable both to industry-standard and proprietary buses through the Tachyon system interface (a generic interface) and provides a seamless interface to GLM-compliant modules and components." CQ-1006, p. 101 (emphasis added). "A generic Fibre Channel host bus adapter board using the Tachyon chip contains the following: Backplane Connector. Connects the backplane interface chip to the system bus. Backplane Interface Chip. Enables the connection of the Tachyon system interface bus to PCI, EISA, HP-HSC or other bus. Tachyon Chip. HP's Fibre Channel interface controller. Physical Link Module. Tachyon interfaces to many GLM-compliant physical link modules currently in the marketplace. Types of modules include: 1063-Mbit/s DB9 copper connectors for distances up to 30

	meters. 1063-Mbit/s shortwave laser physical link modules for distances up to 500 meters. 1063-Mbit/s longwave laser physical link modules for distances up to 10 kilometers. 266-Mbit/s shortwave laser physical link modules for distances up to 2 kilometers. A block diagram of a typical host bus adapter is shown in Fig. 14:" CQ-1006, p. 111 (emphasis added).

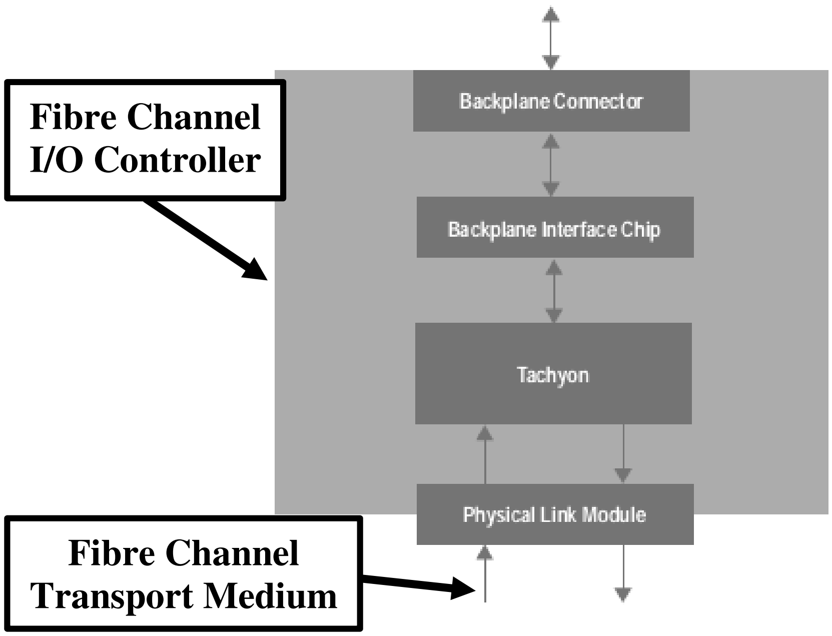

Fig. 14. Block diagram of a typical host adapter board.

CQ-1006, Fig. 14 (annotated)

Accordingly, one of ordinary skill in the art would have been motivated to replace the SCSI I/O host modules in the CRD-5500 RAID Controller with a Fibre Channel I/O host module. Such an exchange would have been a simple substitution because the CRD-5500 RAID controller has a modular design. The Fibre Channel-based I/O module would have allowed the CRD-5500 RAID Controller to communicate with hosts on a serial Fibre Channel transport medium, thereby overcoming many of the known limitations of SCSI buses described by the HP Journal.

Additionally, the beneficial results of modifying the CRD-5500 RAID Controller to interface with a Fibre Channel transport

	medium would have been predictable because the CRD-5500 RAID Controller was specifically designed to support high speed serial interfaces, such as Fibre Channel: "Unlike other RAID controllers, CMD's advanced 'Viper' RAID architecture and ASICs were **designed to support tomorrow's high speed serial interfaces, such as Fiberchannel (FCAL)** and Serial Storage Architecture (SSA)." CQ-1005, p. 1 (emphasis added). **Fourth**, the HP Journal teaches that a Fibre Channel I/O module supports multiple workstations being connected to the same communication link via the Fibre Channel arbitrated loop method: "Fibre Channel arbitrated loop, or FC-AL, is a method for interconnecting from two to 126 devices through attachment points called L_Ports in a loop configuration. L_Ports can consist of I/O devices and systems of various performance levels. . . . Arbitrated loop is the most common Fibre Channel topology." CQ-1006, p. 95 (emphasis added).

	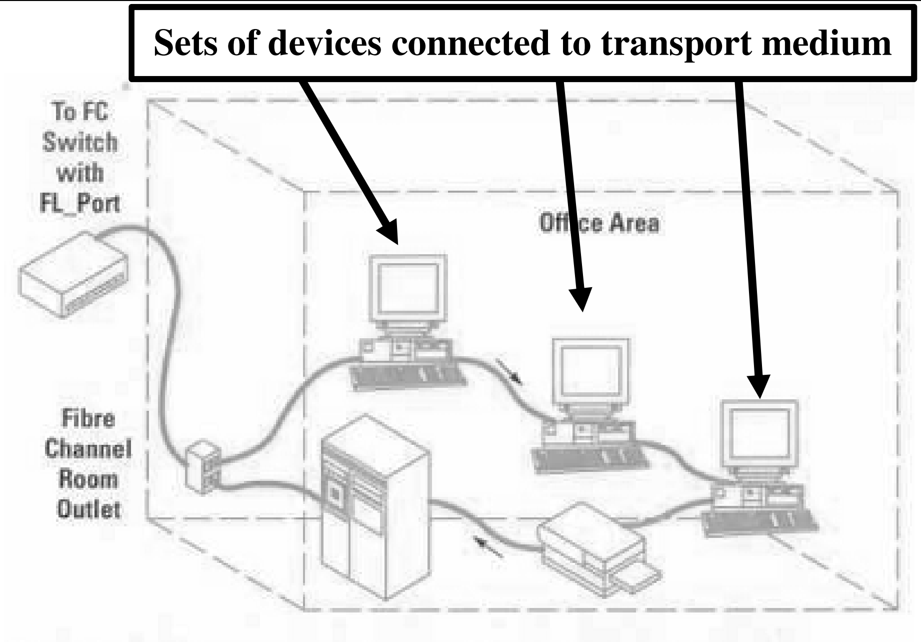 Fig. 3. An office configured in a public arbitrated loop topology. CQ-1006, Fig. 3, p. 96 (annotated) **Additionally**, the HP Journal teaches that the Tachyon controller chip supports Fibre Channel arbitrated loop: "To provide the most flexibility for customer applications, Tachyon . . . Supports Fibre Channel arbitrated loop (FC-AL), point-to-point, and fabric (crosspoint switched) topologies." CQ-1006, p. 101 (emphasis added). One of ordinary skill in the art would have been motivated to move the plurality of hosts individually connected to the CRD-5500 RAID Controller (via SCSI buses) to a single Fibre Channel arbitrated loop communication link because doing so would save I/O slots in the CRD-5500 and minimize cabling infrastructure: "In addition, Fibre Channel resolves the 'slots and watts' problem that current symmetric multiprocessing systems have. For example, in 1995, three FDDI ports and six fast and wide SCSI ports were required to use fully the I/O capabilities of a symmetric multiprocessing HP server. Fibre Channel could

	<u>support these I/O services with just three slots.</u>" CQ-1006, p. 100 (emphasis added). "<u>A single 100-Mbyte/s Fibre Channel port can replace five 20-Mbyte/s SCSI ports</u>, in terms of raw through put. Fibre Channel provides a total network bandwidth of about one gigabit per second." CQ-1006, p. 94 (emphasis added). "Tachyon host adapters save system slots, minimizing cost and cabling infrastructure." CQ-1006, p. 101. **Thus**, the set of hosts connected to the CRD-5500 RAID controller and the controller's modular design, as taught by the CRD-5500 Manual, in view of the Fibre Channel I/O module interfacing with the serial Fibre Channel transport medium to which multiple devices may be connected, as taught by the HP Journal, render obvious "a set of devices connected to a first transport medium, wherein the first transport medium is a serial transport medium" as recited in the claim.
[20.2] a set of remote storage devices connected to a second transport medium;	[20.2] ***"a set of remote storage devices connected to a second transport medium"*** The CRD-5500 Manual teaches that a set of disk drives are coupled to the CRD-5500 RAID controller via a SCSI bus (second transport medium): "Our system will have <u>four disk drives on each of the controller's six disk channels,</u> for a total of 24 disk drives. To connect the drives, <u>run a shielded SCSI cable from each of the controller's drive channels, and daisy-chain the drives to each cable</u>."CQ-1004, p. 3-1 (emphasis added). "The CRD-5500 involves several SCSI buses, which can take the form of 8-bit single-ended, 16-bit single-ended, or 16-bit differential buses. On top of this, one or more of the buses may be designated as a host bus for communicating with the host computer, <u>and the others may be populated with disk drives to form the RAID set pool</u>." CQ-1004, p. 2-4 (emphasis added).

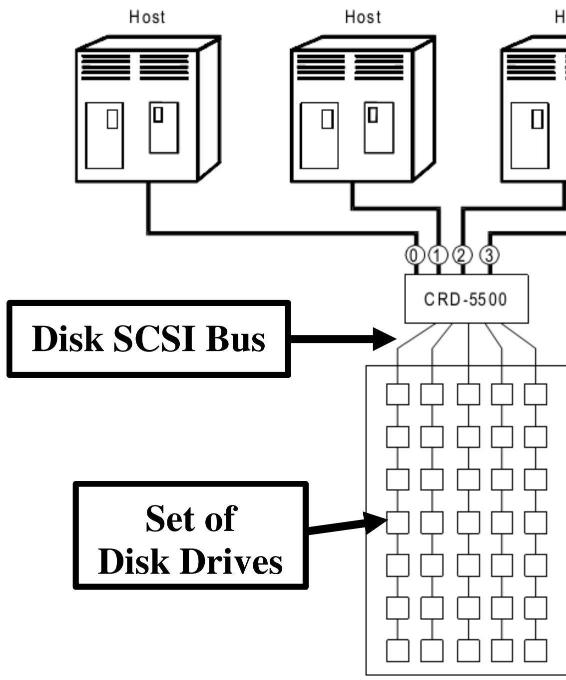

CQ-1004, Figure 1-2 (annotated)

Second, with respect to the recitation that the storage devices are "remote," the CRD-5500 Manual teaches that the storage devices are indirectly connected to the hosts via the CRD-5500 RAID controller.

Further, if "remote" is interpreted to mean that either (i) the storage devices are separated from the hosts by a distance greater than allowed by a conventional parallel network interconnect or (ii) that the storage devices are indirectly connected through at least one serial network transport medium, then section [20.1] of this claim chart describes how one of ordinary skilled in the art would have been motivated by the teachings of the HP Journal to utilize a Fibre Channel serial transport medium in conjunction with the CRD-5500 RAID controller. To that end, the HP Journal teaches that Fibre Channel "allows sustained gigabit data throughput at <u>distance options from ten meters on copper to ten kilometers</u> over single-mode optical fiber." CQ-1006, p. 99 (emphasis added).

Thus, the set of daisy-chained storage devices connected to the

	CRD-5500 RAID controller via a SCSI cable, as taught by the CRD-5500 Manual, discloses "a set of remote storage devices connected to a second transport medium" as recited in the claim.
[20.3] a storage router connected to the serial transport medium;	[20.3] ***"a storage router connected to the serial transport medium"*** As discussed in [20.1], one of ordinary skill in the art would have been motivated to connect the CRD-5500 RAID controller (storage router) to a serial Fibre Channel transport medium. **Thus**, the CRD-5500 RAID controller connected to a serial Fibre Channel transport medium, as taught by the CRD-5500 Manual in view of the HP Journal, render obvious "a storage router connected to the serial transport medium" as recited in the claim.
[20.4a] a storage router connected to the first transport medium and second transport medium to provide virtual local storage on the remote storage devices, the storage router configured to:	[20.4a] ***"a storage router connected to the first transport medium and second transport medium to provide virtual local storage on the remote storage devices"*** As discussed above in [20.2] and [20.3], the CRD-5500 RAID controller (storage router) is connected to a serial Fibre Channel transport medium (first transport medium) and at least one SCSI bus with disk drives (second transport medium). **Second**, the CRD-5500 Manual teaches that the CRD-5500 RAID controller provides virtual local storage on disk drives coupled to the controller by grouping the drives into RAID sets and redundancy groups and then mapping the redundancy groups to Logical Unit Numbers (LUNs) that the hosts use to address their assigned virtual storage spaces, which appear to the hosts as different disk drives: "The CRD-5500 affords great flexibility in creating RAID sets. Drives of different sizes and manufacturers may be combined in a RAID set." CQ-1004, p. 3-3 (emphasis added). "The CRD-5500 supports the partitioning of RAID sets. A partitioned RAID set will have multiple redundancy groups associated with it. Each redundancy group will have its own

	redundancy group number, which corresponds to the Logical Unit Number (LUN) that the host will use to address the partition, unless you map the redundancy group to another LUN with the Host LUN Mapping screen." CQ-1004, p. 4-2. "Since **each partition will appear to the host as a different disk drive**, it must have its own LUN." CQ-1004, p. 3-6 (emphasis added). "In this example, Redundancy Group numbers 1 through 4 are associated with RAID Set number 1. From the controller's point of view in this example, Redundancy Group number 0 belongs to RAID Set 0, which is configured for RAID Level 4, and Redundancy Group numbers 1 through 4 belong to RAID Set 1, which is configured for Level 5. From the host's point of view, there is no such thing as RAID Sets 0 and 1, only (in this example) LUNs 0 through 4, which the CRD-5500 sees as Redundancy Groups 0 through 4." CQ-1004, p. 3-6 (emphasis added). **Accordingly**, the CRD-5500 Manual teaches that a LUN/partition on a remote disk will "appear to the host" as a local disk drive. **Thus**, the CRD-5500 RAID controller that is connected to a serial Fibre Channel transport medium and a SCSI bus with disk drives and provides virtual partitions on the drives that appear to be disk drives to the hosts, as taught by the CRD-5500 Manual in view of the HP Journal, discloses "a storage router connected to the first transport medium and second transport medium to provide virtual local storage on the remote storage devices, the storage router configured to" as recited in the claim.
[20.4b(i)] maintain a map to allocate storage space on the remote	See [1.2b(i)]

storage devices to devices connected to the first transport medium by associating representations of the devices connected to the first transport medium with representations of storage space on the remote storage devices,	
[20.4b(ii)] wherein each representation of a device connected to the first transport medium is associated with one or more representations of storage space on the remote storage devices;	See [1.2b(ii)]
[20.4c] control access from the devices connected to	See [1.2c]

the first transport medium to the storage space on the remote storage devices in accordance with the map; and [20.4d] allow access from devices connected to the first transport medium to the remote storage devices using native low level block protocol.	 See [1.2d]
Claim 21	
[21.0] The storage network of claim 20,	See [20.0] – [20.4]
[21.1] wherein the map associates a representation of storage space on the remote storage devices with multiple devices	See [2.1]

connected to the first transport medium.	
Claim 22	
[22.0] The storage network of claim 20,	See [20.0] – [20.4]
[22.1] wherein the storage space on the remote storage devices comprises storage space on multiple remote storage devices.	See [3.1]
Claim 23	
[23.0] The storage network of claim 20,	See [20.0] – [20.4]
[23.1] wherein the map associates a representation of a device connected to the first transport medium with a representation of an entire	See [4.1]

storage space of at least one remote storage device.	
Claim 24	
[24.0] The storage network of claim 20,	See [20.0] – [20.4]
[24.1] wherein the map resides at the storage router and is maintained at the storage router.	See [5.1]
Claim 25	
[25.0] The storage network of claim 20,	See [20.0] – [20.4]
[25.1] wherein the native low level block protocol is received at the storage router via the first transport medium and the storage router uses the received native low level block protocol to	See [6.1]

allow the devices connected to the first transport medium access to storage space specifically allocated to them in the map.	
Claim 26	
[26.0] The storage network of claim 20,	See [20.0] – [20.4]
[26.1] wherein the storage router is configured to receive commands according to a first low level block protocol from the device connected to the first transport medium and forward commands according to a second low level block protocol to the	See [7.1]

remote storage devices.	
Claim 27	
[27.0] The storage network of claim 20,	See [20.0] – [20.4]
[27.1] wherein the first low level block protocol is an FCP protocol and the second low level block protocol is a protocol other than FCP.	See [8.1]
Claim 28	
[28.0] The storage network of claim 20,	See [20.0] – [20.4]
[28.1] wherein the map comprises one or more tables.	See [9.1]
Claim 29	
[29.0] The storage network of claim 20,	See [20.0] – [20.4]
[29.1] wherein the virtual local	See [10.1]

storage is provided to the devices connected to the first transport medium in a manner that is transparent to the devices and wherein the storage space allocated to the devices connected to the first transport medium appears to the devices as local storage.	
Claim 30	
[30.0] The storage network of claim 20,	See [20.0] – [20.4]
[30.1] wherein the storage router provides centralized control of what the devices connected to the first transport medium see as local storage.	See [11.1]

Claim 31	
[31.0] The storage network of claim 20,	See [20.0] – [20.4]
[31.1] wherein the representations of storage space comprise logical unit numbers that represent a subset of storage on the remote storage devices.	See [12.1]
Claim 32	
[32.0] The storage network of claim 31,	See [31.0] – [30.1]
[32.1] wherein the storage router is operable to route requests to the same logical unit number from different devices connected to the first transport medium to	See [13.1]

different subsets of storage space on the remote storage devices.	
Claim 33	
[33.0] The storage network of claim 20,	See [20.0] – [20.4]
[33.1] wherein the representations of devices connected to the first transport medium are unique identifiers.	See [14.1]
Claim 35	
[35.0] The storage network of claim 20,	See [20.0] – [20.4]
[35.1] wherein the storage router is configured to allow modification of the map in a manner transparent to	See [16.1]

and without involvement of the devices connected to the first transport medium.	
Claim 36	
[36.0] The storage network of claim 20,	See [20.0] – [20.4]
[36.1] wherein the first transport medium is a fibre channel transport medium and the second transport medium is a fibre channel transport medium.	See [19.1]

Claim 37	
[37.0] A method for providing virtual local storage on remote storage devices comprising:	[37.0] ***"A method for providing virtual local storage on remote storage devices"*** **First**, the CRD-5500 Manual describes the operation of the CRD-5500 RAID controller, which routes data between hosts and SCSI disk arrays: "The CRD-5500 RAID controller provides high-performance, high-availability access to SCSI disk array subsystems along a

Fast/Wide SCSI bus. With a modular hardware design and an intuitive configuration utility, the controller may be tailored to suit a wide range of storage needs, now and in the future." CQ-1004, p. 1-1.

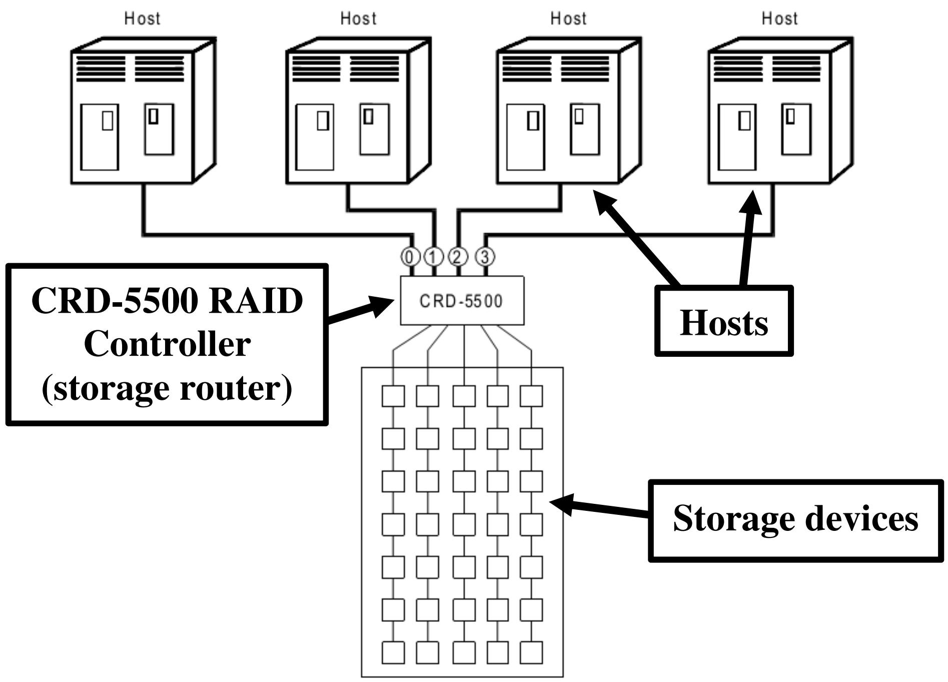

CQ-1004, Figure 1-2 (annotated)

Second, the CRD-5500 Manual teaches that the CRD-5500 RAID controller provides virtual local storage on disk drives coupled to the controller by grouping the drives into RAID sets and redundancy groups and then mapping the RAID sets to Logical Unit Numbers (LUNs) that the hosts use to address their assigned virtual storage spaces, which appear to the hosts as different disk drives:

"The CRD-5500 affords great flexibility in creating RAID sets. Drives of different sizes and manufacturers may be combined in a RAID set." CQ-1004, p. 3-3 (emphasis added).

"The CRD-5500 supports the partitioning of RAID sets. A partitioned RAID set will have multiple redundancy groups associated with it. Each redundancy group will have its own redundancy group number, which corresponds to the Logical

	Unit Number (LUN) that the host will use to address the partition, unless you map the redundancy group to another LUN with the Host LUN Mapping screen." CQ-1004, p. 4-2. "Since each partition will appear to the host as a different disk drive, it must have its own LUN." CQ-1004, p. 3-6 (emphasis added). "In this example, Redundancy Group numbers 1 through 4 are associated with RAID Set number 1. From the controller's point of view in this example, Redundancy Group number 0 belongs to RAID Set 0, which is configured for RAID Level 4, and Redundancy Group numbers 1 through 4 belong to RAID Set 1, which is configured for Level 5. From the host's point of view, there is no such thing as RAID Sets 0 and 1, only (in this example) LUNs 0 through 4, which the CRD-5500 sees as Redundancy Groups 0 through 4." CQ-1004, p. 3-6 (emphasis added). **Accordingly**, the CRD-5500 Manual teaches that a LUN/partition on a remote disk will "appear to the host" as a local disk drive. **Third**, with respect to the recitation in the preamble that the storage devices are "remote," the CRD-5500 Manual teaches that the storage devices are indirectly connected to the hosts via the CRD-5500 RAID controller. Further, if "remote" is interpreted to mean that either (i) the storage devices are separated from the hosts by a distance greater than allowed by a conventional parallel network interconnect or (ii) that the storage devices are indirectly connected through at least one serial network transport medium, then section [37.1] of this claim chart describes how one of ordinary skilled in the art would have been motivated by the teachings of the HP Journal to utilize a Fibre Channel serial transport medium in conjunction with the CRD-5500 RAID controller. To that end, the HP Journal teaches that Fibre Channel "allows sustained gigabit data throughput at distance options from ten meters on copper to ten

	kilometers over single-mode optical fiber." CQ-1006, p. 99 (emphasis added). **Thus**, the CRD-5500 RAID controller that routes data between hosts and disk drives and provides virtual partitions on the drives that appear to be disk drives to the hosts, as taught by the CRD-5500 Manual, discloses "a method for providing virtual local storage on remote storage devices" as recited in the claim.
[37.1] connecting a storage router between a set of devices connected to a first transport medium and a set of remote storage devices, wherein the first transport medium is a serial transport medium;	[37.1] ***"connecting a storage router between a set of devices connected to a first transport medium and a set of remote storage devices, wherein the first transport medium is a serial transport medium"*** **First**, the CRD-5500 Manual teaches connecting the CRD-5500 RAID controller (storage router) between a set of hosts connected to SCSI buses and a set of remote storage devices connected to SCSI buses: "Figure 1-2 shows how you can connect as many as four hosts to the CRD-5500. By using the controller's Host LUN Mapping feature, you can assign redundancy groups to a particular host. This makes it possible for hosts running incompatible operating systems to use the same CRD-5500 controller." CQ-1004, p. 1-2 (emphasis added). "Our system will have four disk drives on each of the controller's six disk channels, for a total of 24 disk drives. To connect the drives, run a shielded SCSI cable from each of the controller's drive channels, and daisy-chain the drives to each cable."CQ-1004, p. 3-1 (emphasis added). The CRD-5500 involves several SCSI buses, which can take the form of 8-bit single-ended, 16-bit single-ended, or 16-bit differential buses. On top of this, one or more of the buses may be designated as a host bus for communicating with the host computer, and the others may be populated with disk drives to form the RAID set pool." CQ-1004, p. 2-4 (emphasis added).

	"In this example, we will configure the 16-bit I/O channels in slots 0 and 1 as host modules. Connect the host SCSI cables to these modules." CQ-1004, p. 3-1 (emphasis added).

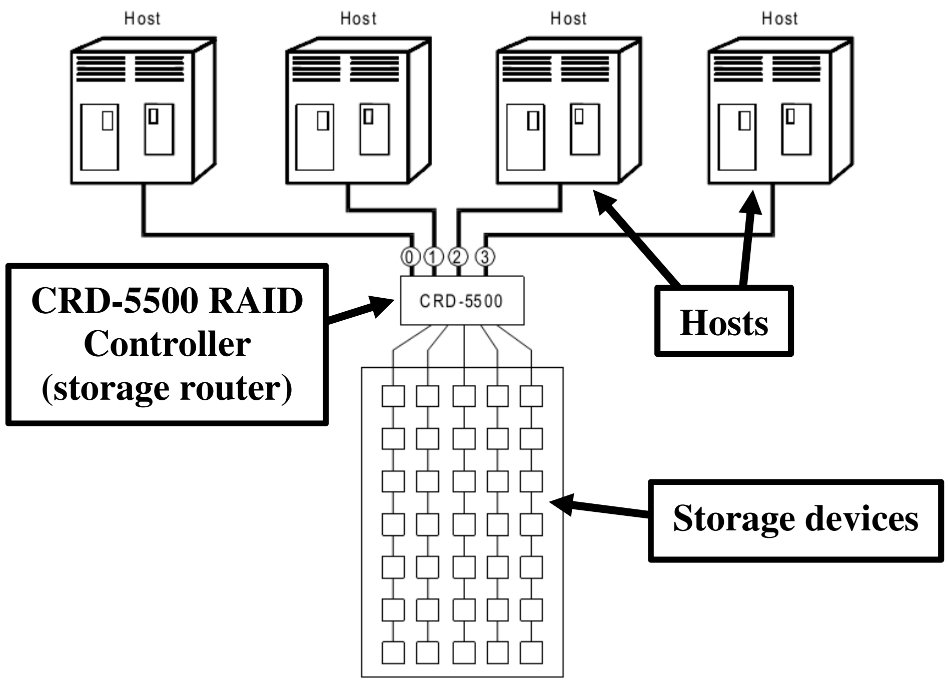

CQ-1004, Figure 1-2 (annotated)

Second, the CRD-5500 Manual teaches that the CRD-5500 RAID controller has a modular design. Specifically, the controller's slots accept many different types of I/O modules to interface with different buses (transport media):

"The CRD-5500 employs a modular design for maximum flexibility. You may customize your controller using a combination of 8-bit single-ended SCSI I/O modules, 16-bit single-ended modules, and 16-bit differential modules." CQ-1004, p. 1-1 (emphasis added).

"The CRD-5500's modular design offers great flexibility in configuring the controller. All controllers come with a system board. The rest of the slots are I/O channels, either for I/O to and from the host computer or to and from the disk drives connected

	to the controller. CQ-1004, p. 2-1 (emphasis added). **Third**, the HP Journal teaches that there are a number of advantages of using the Fibre Channel serial transport medium instead of SCSI buses: "Additionally, as the number of interconnects between systems and I/O devices continues to increase, I/O channels become bottlenecks to system performance. For all these reasons, today's parallel bus architectures are reaching their limits. In the search for a higher-performance serial interface, HP chose Fibre Channel because it overcomes the limitations mentioned above by supporting sustained gigabit data transfer rates." CQ-1006, p. 5. "A channel such as SCSI (Small Computer Systems Interface), which operates at a maximum throughput of 20 megabytes per second in fast and wide mode, simply cannot keep pace with ever-increasing processor speeds and data rate requirements." CQ-1006, p. 99. "Current peripheral interconnect protocols are limited in the number of devices they can interconnect. For example, parallel SCSI can connect eight devices and 16-bit wide SCSI can connect 16 devices. In addition, peripheral connectors are becoming too large to fit into the shrinking footprints of systems and peripherals. Other SCSI limitations include half-duplex operation only, lack of a switching capability, inability to interconnect individual buses, and the need for customized drivers and adapters for various types of attached devices." CQ-1006, p. 99. "Fibre Channel is a flexible, scalable, high-speed data transfer interface that can operate over a variety of both copper wire and optical fiber at data rates up to 250 times faster than existing communications interfaces." CQ-1006, p. 94. "A single 100-Mbyte/s Fibre Channel port can replace five 20-Mbyte/s SCSI ports, in terms of raw through put. Fibre Channel

	provides a total network bandwidth of about one gigabit per second." CQ-1006, p. 94. "Another advantage of Fibre Channel is that it uses small connectors. The serial connectors used for Fibre Channel are a fraction of the size of SCSI parallel connectors and have fewer pins, thereby reducing the likelihood of physical damage. Also, depending on the topology, many more devices can be interconnected on Fibre Channel than on existing channels." CQ-1006, p. 94. "Fibre Channel's increased bandwidth provides distance flexibility, increased addressability, and simplified cabling. Fibre Channel has versatility, room for growth, and qualified vendor support. Mass storage suppliers are using Fibre Channel to interconnect subsystems and systems and to control embedded disk drives." CQ-1006, p. 99. **Additionally**, the HP Journal teaches one of ordinary skill in the art how to implement a Fibre Channel I/O adapter board (first controller) using the Tachyon controller chip that interfaces with a Fibre Channel serial transport medium (first transport medium). As will be described below, one of ordinary skill in the art would have been motivated to replace the SCSI I/O host modules in the CRD-5500 RAID Controller with a Fibre Channel I/O host module. "The system interconnect laboratory of the HP Networked Computing Division became interested in <u>Fibre Channel in 1993 as a method of entering the high-speed serial interconnect market</u> because Fibre Channel was the first technology that could be used for <u>both networking and mass storage</u>. . . . The <u>Tachyon chip (Fig. 3) implements the FC-1 and FC-2 layers of the five-layer Fibre Channel standard</u> (see Article 11)." CQ-1006, p. 101 (emphasis added). "<u>Tachyon's host attach enables low-cost gigabit host adapters on industry-standard buses including PCI, PMC, S-Bus, VME, EISA, Turbo Channel, and MCA. It is easily adaptable both to</u>

	industry-standard and proprietary buses through the Tachyon system interface (a generic interface) and provides a seamless interface to GLM-compliant modules and components." CQ-1006, p. 101 (emphasis added). "A generic Fibre Channel host bus adapter board using the Tachyon chip contains the following: Backplane Connector. Connects the backplane interface chip to the system bus. Backplane Interface Chip. Enables the connection of the Tachyon system interface bus to PCI, EISA, HP-HSC or other bus. Tachyon Chip. HP's Fibre Channel interface controller. Physical Link Module. Tachyon interfaces to many GLM-compliant physical link modules currently in the marketplace. Types of modules include: 1063-Mbit/s DB9 copper connectors for distances up to 30 meters. 1063-Mbit/s shortwave laser physical link modules for distances up to 500 meters. 1063-Mbit/s longwave laser physical link modules for distances up to 10 kilometers. 266-Mbit/s shortwave laser physical link modules for distances up to 2 kilometers. A block diagram of a typical host bus adapter is shown in Fig. 14:" CQ-1006, p. 111 (emphasis added).

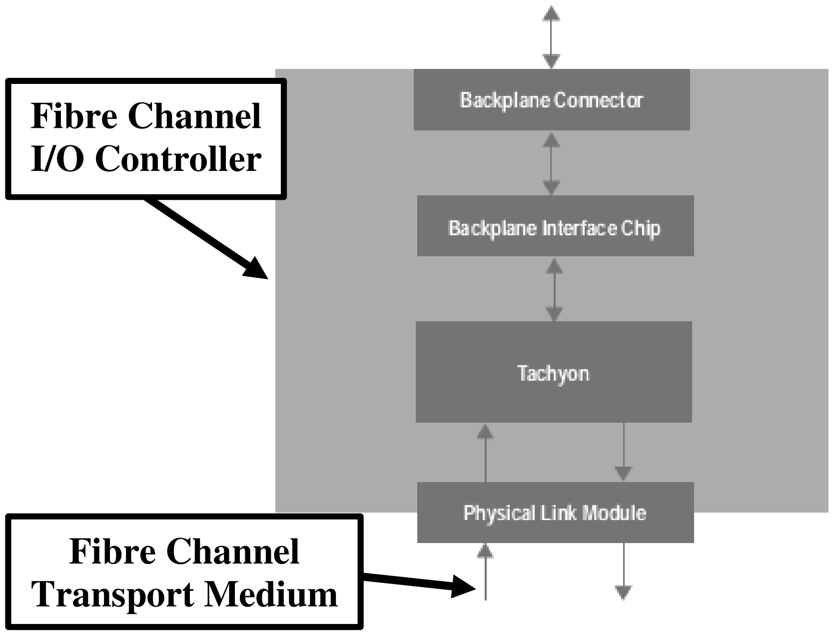

Fig. 14. Block diagram of a typical host adapter board.

CQ-1006, Fig. 14 (annotated)

Accordingly, one of ordinary skill in the art would have been motivated to replace the SCSI I/O host modules in the CRD-5500 RAID Controller with a Fibre Channel I/O host module. Such an exchange would have been a simple substitution because the CRD-5500 RAID controller has a modular design. The Fibre Channel-based I/O module would have allowed the CRD-5500 RAID Controller to communicate with hosts on a Fibre Channel transport medium, thereby overcoming many of the known limitations of SCSI buses described by the HP Journal.

Additionally, the beneficial results of modifying the CRD-5500 RAID Controller to interface with a Fibre Channel transport medium would have been predictable because the CRD-5500 RAID Controller *was specifically designed to support high speed serial interfaces, such as Fibre Channel*:

"Unlike other RAID controllers, CMD's advanced 'Viper' RAID architecture and ASICs were **designed to support tomorrow's high speed serial interfaces, such as Fiberchannel (FCAL)** and Serial Storage Architecture (SSA)." CQ-1005, p. 1 (emphasis added).

Fourth, the HP Journal teaches that a Fibre Channel I/O module supports multiple workstations being connected to the same communication link via the Fibre Channel arbitrated loop method:

"Fibre Channel arbitrated loop, or FC-AL, is a method for interconnecting from two to 126 devices through attachment points called L_Ports in a loop configuration. L_Ports can consist of I/O devices and systems of various performance levels. . . . Arbitrated loop is the most common Fibre Channel topology." CQ-1006, p. 95 (emphasis added).

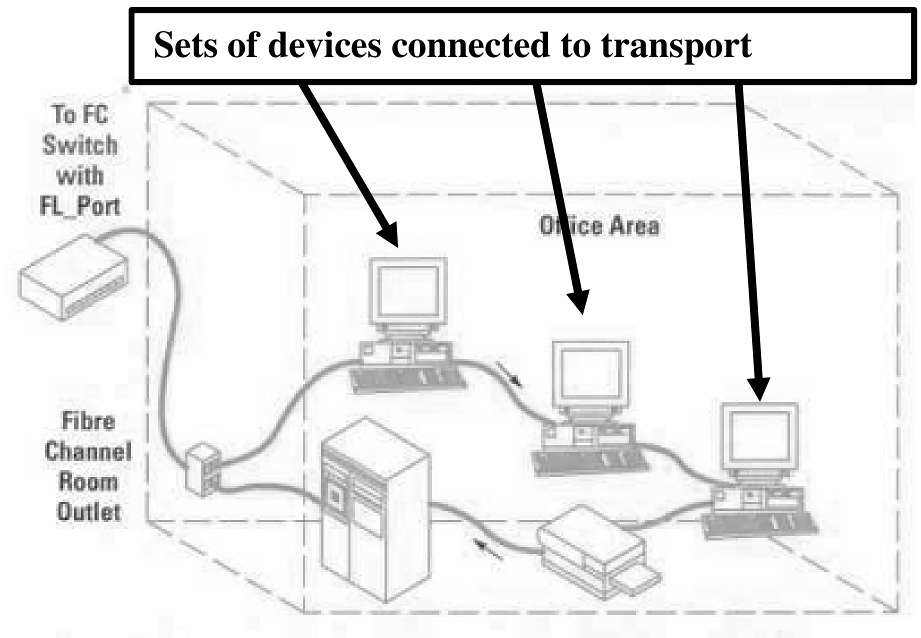

Fig. 3. An office configured in a public arbitrated loop topology.

CQ-1006, Fig. 3, p. 96 (annotated)

Additionally, the HP Journal teaches that the Tachyon controller chip supports Fibre Channel arbitrated loop:

"To provide the most flexibility for customer applications, Tachyon . . . Supports Fibre Channel arbitrated loop (FC-AL), point-to-point, and fabric (crosspoint switched) topologies." CQ-1006, p. 101 (emphasis added).

	One of ordinary skill in the art would have been motivated to move the plurality of hosts individually connected to the CRD-5500 RAID Controller (via SCSI buses) to a single Fibre Channel arbitrated loop communication link because doing so would save I/O slots in the CRD-5500 and minimize cabling infrastructure: "In addition, Fibre Channel resolves the 'slots and watts' problem that current symmetric multiprocessing systems have. For example, in 1995, three FDDI ports and six fast and wide SCSI ports were required to use fully the I/O capabilities of a symmetric multiprocessing HP server. Fibre Channel could support these I/O services with just three slots." CQ-1006, p. 100 (emphasis added). "A single 100-Mbyte/s Fibre Channel port can replace five 20-Mbyte/s SCSI ports, in terms of raw through put. Fibre Channel provides a total network bandwidth of about one gigabit per second." CQ-1006, p. 94 (emphasis added). "Tachyon host adapters save system slots, minimizing cost and cabling infrastructure." CQ-1006, p. 101. **Thus**, connecting the CRD-5500 RAID controller between a set of hosts on a serial Fibre Channel transport medium and a set of disk drives, as taught by the CRD-5500 Manual in view of the HP Journal, renders obvious "connecting a storage router between a set of devices connected to a first transport medium and a set of remote storage devices, wherein the first transport medium is a serial transport medium" as recited in the claim.
[37.2a] maintaining a map at the storage router to allocate storage space on the remote	See [1.2b(i)]

storage devices to devices connected to the first transport medium by associating representations of the devices connected to the first transport medium with representations of storage space on the remote storage devices, [37.2b] wherein each representation of a device connected to the first transport medium is associated with one or more representations of storage space on the remote storage devices;	 See [1.2b(ii)]
[37.3] controlling access from the devices	See [1.2c]

connected to the first transport medium to the storage space on the remote storage devices in accordance with the map; and	
[37.4] allowing access from devices connected to the first transport medium to the remote storage devices using native low level block protocol.	See [1.2d]
Claim 38	
[38.0] The method of claim 37,	See [37.0] – [37.4]
[38.1] wherein the map associates a representation of storage space on the remote storage devices with multiple devices	See [2.1]

connected to the first transport medium.	
Claim 39	
[39.0] The method of claim 37,	See [37.0] – [37.4]
[39.1] wherein the storage space on the remote storage devices comprises storage space on multiple remote storage devices.	See [3.1]
Claim 40	
[40.0] The method of claim 37,	See [37.0] – [37.4]
[40.1] wherein the map associates a representation of a device connected to the first transport medium with a representation of an entire storage space of at least one	See [4.1]

remote storage device.	
Claim 41	
[41.0] The method of claim 37,	See [37.0] – [37.4]
[41.1] wherein the map resides at the storage router and is maintained at the storage router.	See [5.1]
Claim 42	
[42.0] The method of claim 37, further comprising	See [37.0] – [37.4]
[42.1] receiving the native low level block protocol at the storage router via the first transport medium; using the received native low level block protocol at the storage router to allow the devices	See [6.1]

connected to the first transport medium access to storage space specifically allocated to them in the map.	
Claim 43	
[43.0] The method of claim 37, further comprising	See [37.0] – [37.4]
[43.1] receiving commands at the storage router according to a first low level block protocol from the device connected to the first transport medium and forwarding commands according to a second low level block protocol to the remote storage devices.	See [7.1]

Claim 44	
[44.0] The method of claim 43,	See [43.0] – [43.1]
[44.1] wherein the first low level block protocol is an FCP protocol and the second low level block protocol is a protocol other than FCP.	See [8.1]
Claim 45	
[45.0] The method of claim 37,	See [37.0] – [37.4]
[45.1] wherein the map comprises one or more tables.	See [9.1]
Claim 46	
[46.0] The method of claim 37,	See [37.0] – [37.4]
[46.1] wherein the virtual local storage is provided to the devices connected to the first	See [10.1]

transport medium in a manner that is transparent to the devices and wherein the storage space allocated to the devices connected to the first transport medium appears to the devices as local storage.	
Claim 47	
[47.0] The method of claim 37,	See [37.0] – [37.4]
[47.1] wherein the storage router provides centralized control of what the devices connected to the first transport medium see as local storage.	See [11.1]
Claim 48	
[48.0] The method of claim 37,	See [37.0] – [37.4]

[48.1] wherein the representations of storage space comprise logical unit numbers that represent a subset of storage on the remote storage devices.	See [12.1]
Claim 49	
[49.0] The method of claim 48,	See [48.0] – [48.1]
[49.1] wherein the storage router is operable to route requests to the same logical unit number from different devices connected to the first transport medium to different subsets of storage space on the remote storage devices.	See [13.1]

Claim 50	
[50.0] The method of claim 37,	See [37.0] – [37.4]
[50.1] wherein the representations of devices connected to the first transport medium are unique identifiers.	See [14.1]
Claim 53	
[53.0] The method of claim 1 (typo, interpreted to be claim 37),	See [37.0] – [37.4]
[53.1] wherein connecting the storage router between a set of devices connected to a first transport medium and a set of remote storage devices further comprises connecting the storage router between a first fibre channel	See [19.1]

transport medium and a second fibre channel transport medium.	

VIII. Challenge #2: Claims 15, 34, 51, 52 are obvious over the CRD-5500 Manual in view of the HP Journal and in further view of the ANSI Fibre Channel FC-PH Standard

59. It is my opinion that the CRD-5500 Manual in view of the HP Journal and in further view of the ANSI Fibre Channel FC-PH Standard renders obvious each and every element of at least claims 15, 34, 51 and 52 of the ’041 Patent.

A. ANSI Fibre Channel FC-PH Standard

60. The ANSI Fibre Channel FC-PH Standard describes the physical and signaling interface aspects of the Fibre Channel standard. Specifically, the FC-PH standard describes “physical interface, transmission protocol, and signaling protocol” of a Fibre Channel high-performance serial link. (CQ-1007, p. 1). One aspect of the standard requires that devices connected to a Fibre Channel network must be addressed by a unique identifier. (CQ-1007, p. 132). Further, when Fibre Channel devices are part of a heterogeneous network (i.e., FC devices communicating with Non-FC devices), the Fibre Channel devices may be assigned a Worldwide Name (WWN) that is a 64 bit worldwide unique address. (CQ-1007, pp. 46, 148).

61. The ANSI Fibre Channel FC-PH Standard is dated June 1, 1994 and was "published by" the American National Standards Institute (ANSI). (CQ-1007, pp. 1, 4).

B. Reasons to Combine the ANSI Fibre Channel FC-PH Standard with the CRD-5500 Manual and the HP Journal

62. It is my opinion that a person of ordinary skill in the art would have found it obvious to combine teachings of the ANSI Fibre Channel FC-PH Standard with the teachings of the CRD-5500 Manual and the HP Journal for the reasons set forth below.

63. As discussed above, one of ordinary skill in the art would have been motivated to replace the SCSI I/O host modules in the CRD-5500 RAID Controller with a Fibre Channel I/O host module. Such a modification would have created a heterogeneous network with hosts connected to the CRD-5500 RAID controller via a Fibre Channel link and disk drives connected to the controller via a SCSI bus. One of ordinary skill in the art would have looked to the Fibre Channel FC-PH standard for guidance on how to properly identify the Fibre Channel-based hosts in the heterogeneous network. Assigning unique Worldwide Names to the hosts in accordance with the Fibre Channel standard would have allowed the hosts to interoperate with any other devices complying with the standard. Such a result would have been predictable because the CRD-5500 RAID controller was designed to support Fibre Channel. Accordingly, it is my opinion that one of ordinary skill would have combined the teachings the ANSI Fibre Channel FC-PH Standard with the teachings of the CRD-5500 Manual and the HP Journal.

C. <u>Detailed Analysis</u>

64. The following claim chart describes how the CRD-5500 Manual in view of the HP Journal and in further view of the ANSI Fibre Channel FC-PH Standard renders obvious each and every element of at least claim 15, 34, 51 and 52 of the ’041 Patent.

Claim 15	
[15.0] The storage router of claim 14,	See [14.0] – [14.1]
[15.1] wherein the unique identifiers are world wide names.	[15.1] ***“wherein the unique identifiers are world wide names”*** **First**, as discussed above in [14.1], hosts are uniquely identified by channel numbers. **Second**, as discussed above in [1.1], one of ordinary skill in the art would have been motivated to replace the SCSI I/O host modules in the CRD-5500 RAID Controller with a Fibre Channel I/O host module, so as to communicate with the hosts via a Fibre Channel link. Such a modification would have created a heterogeneous network with hosts connected to the CRD-5500 RAID controller via a Fibre Channel link and disk drives connected to the controller via a SCSI bus. **Third,** the ANSI Fibre Channel FC-PH Standard teaches that in

	a heterogeneous network (FC to Non-FC) nodes (*e.g.*, hosts) are represented by unique identifiers that are worldwide names: "**Worldwide_Name**: An Name_Identifier which is worldwide unique, and represented by a 64 bit unsigned binary value." CQ-1007, p. 46. "The application of Name_Identifiers in Network_Header for heterogeneous (FC to Non-FC) networks and homogeneous (FC to FC) networks is summarized in table 42." CQ-1007, p. 148.

Table 42 - Network addresses

NAA	Name_Identifier	Network
IEEE	WWN	Heterogeneous
CCITT - individual address	WWN	Heterogeneous
CCITT - group address	WWN	Heterogeneous
IP	WWN	Heterogeneous
IEEE extended	FCN	FC Networks
Local	FCN	FC Networks

Note:
WWN - Worldwide Name (worldwide unique address)
FCN - Fibre Channel Name (Fibre Channel unique address)

CQ-1007, p. 148, Table 42

Table 43 - Fibre Channel user identifiers

NAA	Name_Identifier	Fibre Channel users			
		N_Port	Node	F_Port	Fabric
IEEE	WWN	yes	yes	yes	yes
CCITT - individual address	WWN	yes	yes	yes	yes
CCITT - group address	WWN	yes	yes	yes	yes
IEEE extended	FCN	yes	no	yes	no
Local	FCN	yes	yes	yes	yes

Note: yes - applicable to the user
no - not applicable to the user
WWN - Worldwide Name (worldwide unique identifier)
FCN - Fibre Channel Name (Fibre Channel unique identifier)

CQ-1007, p. 149, Table 43

Accordingly, the teachings of the ANSI Fibre Channel FC-PH Standard provide that in a heterogeneous network the nodes (host devices) may be represented by unique identifiers that are

	worldwide names. **Thus**, the CRD-5500 RAID controller uniquely identifying the attached hosts, as taught by the CRD-5500 Manual, in view of the Fibre Channel transport medium that couples the host devices to the controller, as taught by the HP Journal, and further in view of the assigning of unique worldwide names to nodes, as taught by the ANSI Fibre Channel FC-PH Standard, render obvious "wherein the unique identifiers are world wide names" as recited in the claim.
Claim 34	
[34.0] The storage router of claim 33,	See [33.0] – [33.1]
[34.1] wherein the unique identifiers are world wide names.	See [15.1]
Claim 51	
[51.0] The method of claim 50,	See [50.0] – [50.1]
[51.1] wherein the unique identifiers are world wide names.	See [15.1]
Claim 52	
[52.0] The method of claim 51,	See [51.0] – [51.1]
[52.1] wherein	See [16.1]

the storage router is configured to allow modification of the map in a manner transparent to and without involvement of the devices connected to the first transport medium.	

IX. <u>Declaration</u>

65. I declare that all statements made herein on 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 so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Executed:

____

By: Andrew Hospodor, Ph.D.

Date: 05 Sep 2014

Andrew David Hospodor, Ph.D.

P.O. Box 1196, Los Gatos CA. 95031-1196

408.921.5099 andy.hospodor@ieee.org

An accomplished executive and engineer experienced in both startups and Fortune 500 companies with extensive background in distributed systems, applications and storage. Strengths include industrial-academic relationship building, intellectual property development, strategic and technical leadership.

Education

- **Ph.D. Computer Engineering** , Santa Clara University, Santa Clara, CA.
  Emphasis in storage architecture and embedded systems. Minor in business administration.
  Dissertation: A Study of Prefetch in Caching SCSI Disk Drive Buffers.

- **M.S. Computer Science**, Santa Clara University, Santa Clara, CA.
  Concentrated studies in networking, communications, data storage, memory hierarchies, interfaces, computer architecture, performance measurement, and error correction coding.

- **B.S. Computer Engineering**, Lehigh University, Bethlehem PA.
  Emphasis in computer programming, architecture, physics, mathematics.

Experience

Storage System Research Center, University of California, Santa Cruz, CA. 2009 – present

Executive Director and, Project Scientist

Engaged in research and funding related activities. Worked with faculty to develop funding strategies and manage industrial sponsors. Participated in NSF and UC led events designed to attract research funding into the data storage space. Built relationships with industrial and academic contacts. Advised graduate students and reviewed their results pre-publication.

BookRenter, San Jose, CA. 2006 – 2008

CTO

Led the team that created the first nationwide book rental service. Defined the architecture of a new web 2.0 platform for e-business that combines distributed computing with Ruby on Rails (RoR), mySQL, and web services of partners like Amazon, Barnes&Noble and UPS. Formed capitalization strategy, managed fund raising activities and created partnerships to maximize equity leverage. Responsible for all technical aspects of www.bookrenter.com from hiring to operations.

GridPlan, Santa Clara, CA. 2003 - present

Architect

Introduced the first open source capacity planning tool for Grid Computing that enabled both Enterprise, e-business and Scientific environments to accurately access the value of computational grids, cloud computing, large-scale Linux clusters, blade servers and distributed compute farms. Provided the crucial ability to assess cost-performance of interconnection strategies (such as Infini-Band, 10 Gigabit Ethernet, Myrinet), processors (such as XEON, Opteron), storage, switches and middleware. Written entirely in java, GridPlan successfully simulated up to 4096 nodes and created an XML based link-map for systems integrators. Established partnerships with Grid Global Forum (GGF), hardware vendors, independent software vendors and open source providers to provide best-of-breed planning technology toIT shops.

Corosoft, Inc., Cupertino, CA. 2001 - 2003

CTO, Founder

Developed a novel approach to managing enterprise e-business applications (databases, appservers, webservers, file services, network services). Introduced Corosoft virtualization software that aggregates resources (servers, networks, storage) behind application services. Built team, raised $3.8M funding, delivered product to market. Established partnerships in enterprise management (BMC, HP), grid computing (Platform, IBM), network content management (F5) and software

(Oracle, Microsoft). Extended strategy to include power management (ACPI) middleware, streaming, clustering file systems, volume management and flexible storage architectures.

Western Digital Corp., San Jose, CA. 1999 - 2001
Vice President, Systems Architecture

Responsible for all aspects of technology in the formation of new business units. Reported to the Chief Technical Officer and VP of Business Development. Identified new technologies for Enterprise Storage Area Networking (SAN) and Audio/Video streaming storage networking systems, most notably switched fabrics. Drove strategic relationships with well-established software companies such as Microsoft and Veritas for existing technology. Identified and structured relationships with partners and performed due diligence on emerging key technology startups that led to capital investment of $2-5M. Created detailed business plans including capitalization, development and staffing requirements.

Quantum Corp., Milpitas, CA. 1993 - 1999
Storage Architect, Director, Network Storage Architecture Group
Manager, Advanced Storage Applications Group

Led team that developed the first low-cost Network Attached Storage (NAS) disk and tape products. Coordinated company wide technology direction for storage management, file systems, device drivers, BIOs, and APIs for FibreChannel (FC), Gigabit Ethernet (GbE), InfiniBand (IB), Redundant Arrays of Independent Disks, etc. Responsible for technical relationships with strategic partners, such as Microsoft, Legato, Veritas, Oracle. Supported business units with cross connects to Compatibility Lab, Design Engineering, Sales and Marketing. Participated in customer investigations of new storage applications and developed requirements for new storage markets. Managed architecture and performance labs to provide real data for product planning. Created and managed storage simulator team that laid groundwork for delivery of SCSI, IDE, ATA-33, to 133 interfaces, ultimately resulting in net savings of $50M+. Participated in architecting, planning and specification of Self-Monitoring, Automatic Reporting Technology (S.M.A.R.T.) for intelligent storage devices.

Institute for Information Storage Technology, Santa Clara University. 1990 - 1993
IIST Research Fellow and Instructor in Electrical and Computer Engineering (Adjunct Faculty)

Investigated performance of magnetic disk, optical disk, magnetic tape, and buffering in relation to computer system performance. Researched data storage interfaces and architectures and interacted with local data storage industry. Developed relationships with companies such as Quantum, Conner, Seagate, IBM, Iomega and Maxtor. Generated research proposals including budget and staffing requirements and followed proposals through federal government approval. Represented IIST as part of the $24.5M NSIC Ultra-High Density Storage project. Responsible for teaching graduate and undergradute courses.

I/O XEL Incorporated, Santa Clara, CA. 1986 - 1990
Director of Engineering

Developed and delivered performance analysis tools for the data storage industry including The SCSI Benchmark® tester. Clients included Quantum, Priam, Maxtor, Iomega.

Scientific Micro Systems, Mountain View, CA. 1982 - 1986
Design Engineer

Designed disk controller firmware for IEEE 796, LSI-11, and SCSI based disk controllers with embedded microprocessors and microcontrollers. Participated in successful IPO.

National Semiconductor, Santa Clara CA. 1981 - 1982
System Engineer

Designed and delivered the hardware and firmware of a Winchester Disk Controller for IEEE 796 bus.

Digital Equipment Corporation, Tewksbury MA. 1978 - 1979
Software Engineer
Designed conformance tests for compilers and test packages for the VAX VMS debugger

Languages

Java, XML, HTML, UML, c, c++, 8086 assembler, 2900 bit slice, FORTRAN

Operating Systems

Windows, NT, XP, ME, Linux, Ultrix, AIX, MS-DOS, VMS, CPM

Interfaces

SCSI, IDE, ATA, FibreChannel, Gigabit Ethernet and TCP/IP, 10 GbE, InfiniBand

Patents

U.S. Patent 8,343,553, "Essential element extractor " Issued January 01, 2013

U.S. Patent 7,274,659, "Providing streaming media data" Issued September 25, 2007

U.S. Patent 7,002,926, "Isochronous Switched Fabric Network" Issued February 21, 2006

U.S. Patent 6,965,563, "Resource reservation system in a computer network to support end-to-end quality-of-service constraints" Issued November 15, 2005

U.S. Patent 6,888,831 "Distributed resource reservation system for establishing a path through a multi-dimensional computer network to support isochronous data" Issued May 3, 2005

U.S. Patent 6,744,772 "Converting asynchronous packets into isochronous packets for transmission through a multi-dimensional switched fabric network" Issued June 1, 2004

U.S. Patent 6,697,914 "Switched node comprising a disk controller with integrated multi-port switching circuitry" Issued February 24, 2004

U.S. Patent 6,615,312 "Method for processing file system service to reproduce stream data" Issued September 2, 2003

U.S. Patent 6,603,625 "Spindle synchronizing in a multi-dimensional network" Issued August 5, 2003

U.S. Patent 6,470,420 "Method for designating one of a plurality of addressable storage" Issued October 22, 2002

U.S. Patent 6,012,839 "Method and apparatus to protect data within a disk drive buffer" Issued January 11, 2000

U.S. Patent 5,771,397 "SCSI Disconnect/Reconnect timing algorithm for optimal performance" Issued June 23, 1998

U.S. Patent 4,851,998 "Method to Analyze Performance of Computer Peripherals" Issued July 25, 1989

As of July 2010, the above patents have been cited as prior art 183 times in patents granted by the US Patent and Trademark Office.

Peer Reviewed Publications

Preeti Gupta, Avani Wildani, Daniel Rosenthal, Ethan L. Miller, Ian Adams, Christina Strong, Andy Hospodor, "An Economic Perspective of Disk vs. Flash Media in Archival Storage," Proceedings of the 22th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2014), September 2014. [Archival Storage]

Rekha Pitchumani, Andy Hospodor, Ahmed Amer, Yangwook Kang, Ethan L. Miller, Darrell D. E. Long, "Emulating a Shingled Write Disk," Proceedings of the 20th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2012), August 2012. [Shingled Disk]

Ziqian Wan, Alex Nelson, Tao Li, Darrell D. E. Long, Andy Hospodor, "Computer Hard Drive Geolocation by HTTP Feature Extraction," Technical Report UCSC-SSRC-12-04, May 2012. Technical Report UCSC-SSRC-12-04 [Digital Forensics]

Philippe Huibonhoa, Chris Williams, JoAnne Holliday, Andy Hospodor, Thomas Schwarz, "Redundancy Management for P2P Storage," CCGrid 2007 CCGrid 2007, the 7th IEEE International Symposium on Cluster Computing and the Grid, Rio De Jeniero, Brazil, May 2007.

Thomas J. E. Schwarz, Qin Xin, Ethan L. Miller , Darrell D. E. Long, Andy Hospodor, "Spencer Ng, Disk Scrubbing in Large Archival Storage Systems," Proceedings of the IEEE / ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS), Volendam, The Netherlands, October 2004. Awarded Best Paper.

A.D. Hospodor, E.L. Miller, "Interconnection Architectures for Petabyte-Scale High-Performance Storage Systems", Proceedings of the NASA/IEEE Conference on Mass Storage Systems and Technologies (MSST2004), College Park, Maryland, USA, April, 2004.
Conference rated among top 15% by Citeseer.

Andy Hospodor, "Design Alternatives to Improve Access Time Performance Under DOS and UNIX", *ASME Journal of Information Storage Systems*, vol. 6, 1995.

Andy Hospodor, " Mechanical Access Time, Measurement and Calculation ", *ASME Journal of Information Storage Systems*, vol. 6, 1995.

Andrew David Hospodor, "The Effect of Prefetch in SCSI Disk Drive Cache Buffers", Doctoral Dissertation, Institute for Information Storage Technology and Computer Engineering Department, Santa Clara University, 1994.

Andrew D. Hospodor, Albert S. Hoagland, "The Changing Nature of Storage Controllers," *Proceedings of the IEEE*, Special Issue on Data Storage, vol. 81. no. 4, April 1993.

A.S. Hoagland, A.D. Hospodor, "The Information Storage Technology Program at SCU," *IEEE Transactions on Education,* vol. 36, no. 1, February, 1993

Conference Proceedings, Seminars and Presentations

"The Data Storage Industry: Trends and Performance" invited seminar, Electrical and Computer Engineering Graduate Seminar, Santa Clara University, January 22, 1998

Andy Hospodor, “Video On the Desktop: High Demands on Disk Drive Performance and Capacity”, Computer Technology Review, vol. 16, no. 5, May 1996.

Andy D. Hospodor, “Storage as the Backbone of Video Server Performance,” Proceedings of the Annual Broadcast Engineering Conference, National Association of Broadcasters, Las Vegas, NV. April 1996.

"Performance alternatives under DOS and UNIX" and "Mechanical Access Time, Measurement and Calculation", *ASME Info. Storage Conf*, New Orleans, LA, November 1994.

Andy Hospodor, “Storage Interfaces & Architectures” and “RAID Technology,” seminars for Quantum Corp., Milpitas, CA, October 1994.

A. Hospodor, “Magnetic Disk Performance from Optical Disk,” *Proceedings of the IEEE Laser Electro-Optics conference*, Maui, HI, May, 1993

A. Hospodor, “Hit Ratio of Caching Disk Buffers,” *Proceedings of the IEEE CompCon conference*, San Francisco CA, Februrary, 1992

Andy Hospodor, “Measuring and Modeling Storage Access Times”, seminar for UC Berkeley Computer Science Department, January 1992.

A. Hospodor, “Storage Interfaces and Architectures”, an invited tutorial, presented IEEE Systems Design and Networks Conference, Santa Clara, CA, June, 1990.

Andy Hospodor, “SCSI Throughput, How it Affects Performance”, DiskTest Conference, San Jose CA, July, 1989.

Andy Hospodor, “Survival SCSI”, an invited tutorial, IEEE Systems Design and Networks Conference, Santa Clara, CA, May 1989.

Andy Hospodor, “Performance Analysis of Computer Peripherals”, proceedings of the IEEE Systems Design and Networks Conference, Santa Clara, CA, May 1989.

Andy Hospodor, “DCS - the Diagnostic Command Set for SCSI and SCSI-2,” *ANSI X3B7.1*, May 1989.

A. Hospodor, J. Hospodor, “U.S. Science and Technology Museums – 1”, *IEEE Spectrum*, special issue on science and technology for youth, vol. 32, no. 9, September, 1995.

Andy Hospodor, “Input, Output, and Throughput”, *Proceedings of the IEEE Systems Design and Networks Conference*, Santa Clara, CA, May 1989.

“The SCSI Disk Drive Performance Evaluation Report”, Dataquest report, November, 1988.

Andy Hospodor, “A New Look at Analyzing Peripheral Performance,” *Computer Design*, PennWell Publishing, vol. 27, no. 6, March 15, 1988.

Andy Hospodor, “The Quest for the Best,” *Mini-Micro Systems*, Cahners Publishing, vol. 22, no. 11, November, 1988.

Andy Hospodor, “Software Puts SCSI to the Test,” *Mini-Micro Systems*, Cahners Publishing, vol. 20, no. 2, February, 1987.

Professional Activities

- Member, American Association of Cancer Researchers
- Member, Society of Cannabis Clinicians
- Panelist, National Science Foundation review of Big Data Initiative 2014
- Panelist, National Science Foundation review of the National Middleware Initiative (NMI) 2004
- Senior Member, IEEE and Computer Society
- Member, IEEE P1285 standards committee
- Member, Infiniband trade association
- Member ACM, Special Interest Group for Performance Evaluation (SIGMETRICS)
- Member, Grid Global Forum (GGF), Globus
- Invitee, the 2004 International School of Grid Computing, Vico Equense, Italy
- Past program chair of Magnetics Society, Santa Clara Valley chapter
- Research affiliate of University of California, Santa Cruz, System Storage Research Center
- Research associate and Lecturer of Santa Clara University, Computer Engineering Dept.

Expert Witness

SBC v. InRange 2004 - 2005

Brooks, Kushman, Detroit, MI

Advised attorneys on technical matters in Storage Area Networks and Fibre-Channel Switches. Designed Claim Chart, wrote Expert Report, attended depositions. Identified infringing source code and settled favorably for my client, SBC.

Exabyte v. Certance 2005

Weil, Gotshall and Manges, Redwood Shores, CA

Advised attorneys on technical matters in Tape Storage, Error Correction and Recovery. Performed discovery, reviewED patents, identified various methods to achieve similar functionality without infringement. Settled favorably for my client, Certance.

EchoStar v. Tivo 2005 - 2008

Irell and Manella, Irvine, CA

Advised attorneys on technical matters in Streaming and Audio-Video Data storage. Searched literature and patents for prior art, drafted claim charts, wrote expert and rebuttal reports.

HP v. SBC 2005

Brooks, Kushman, Detroit, MI

Advised attorneys on technical matters in Storage Area Networks and Fibre-Channel Switches. Searched literature and patents for prior art. Assisted in staffing additional expert witnesses.

Axalto v. Toshiba 2006

Osha – Liang, Santa Clara, CA

Advised attorneys on technical matters in Flash Memory. Performed initial discovery and examination of products. Identified potentially infringing software, determined exact function and date of creation.

Tandberg Data v. HP 2006-2008

Bartlitt-Beck, Denver, CO

Advised attorneys on technical matters in Tape Storage, Error Correction and Recovery. Performed discovery, reviewED patents. Created claim charts, wrote expert and rebuttal reports.

Mathworks v. Comsol 2007-2008

Nixon Peabody, Boston, MA

Advised attorneys on technical matters in Object Oriented Programming and computer modeling/simulation. Searched literature and patents for prior art. Identified significant software in support of invalidity arguments.

Overland Storage v. IBM 2011

Kirkland & Ellis, Washington, DC

Advised IBM counsel on technical matters in automated tape library storage for an International Trade Commission case. Constructed a live demonstration tape libraries available to one skilled in the art more than one year prior to issuance of the patent in question. Provided deposition testimony and testified from the witness stand at trial.

Round Rock v. Dell 2012

Farella Braun + Martel LLP, San Francisco, CA

Advised Dell counsel on technical matters in data storage, networking, communication and power management. Examine source code and identified embodiments of claim terms. Assisted in preparing a tutorial for the judge. Prepared an expert report that allowed the case to settle favorably for my client, Dell.

Overland Storage v. Quantum 2013

Durie Tangri, San Francisco, CA

Advised Quantum counsel on technical matters in network storage and automated tape library storage. Disassembled both Overland and Quantum storage systems and identified apparatus relevant to patent claims. Examined source code and opined on relevance to method claims for patents in suit. Prepared a detailed expert report that allowed the case to settle favorably for my client, Quantum.