

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

SAP AMERICA, INC.
Petitioner,

v.

VALTRUS INNOVATIONS LTD. AND KEY PATENT INNOVATIONS LTD.
Patent Owner.

Case No. IPR2025-00414
U.S. Patent No. 7,152,182

**DECLARATION OF KEVIN C. ALMEROOTH, PH.D. UNDER 37 C.F.R. §
1.68 IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW**

TABLE OF CONTENTS

	<u>Page</u>
I. INTRODUCTION	9
II. QUALIFICATIONS AND PROFESSIONAL EXPERIENCE	10
III. LEVEL OF ORDINARY SKILL IN THE ART	23
IV. RELEVANT LEGAL STANDARDS	24
V. OVERVIEW OF THE '182 PATENT	26
VI. CLAIM CONSTRUCTION	29
VII. IDENTIFICATION OF HOW THE CLAIMS ARE UNPATENTABLE....	29
VIII. THE PRINCIPAL PRIOR ART	31
A. Summary of <u>Lubbers</u> (Ex-1005)	31
B. Summary of <u>Sicola</u> (Ex-1006).....	35
C. Summary of <u>Ofek</u> (Ex-1007).....	41
IX. DETAILED ANALYSIS THAT CLAIMS 1-8 ARE OBVIOUS OVER <u>LUBBERS</u> AND <u>SICOLA</u>	44
A. <i>Claim 1</i>	50
1. <i>[1preamble]: “A data redundancy system, comprising:”</i>	50
2. <i>[1a]: “a primary storage facility for storing a primary copy of data, the primary storage facility including;”</i>	52
3. <i>[1b]: “a first redundancy appliance operable in a role for receiving a sequence of write requests and for storing data for the sequence of write requests in storage associated with the primary storage facility; and”</i>	55
4. <i>[1c]: “a second redundancy appliance for shadowing the first redundancy appliance wherein the second redundancy appliance assumes the role of the first redundancy appliance;”</i>	58
5. <i>[1d]: “a secondary storage facility for storing data that is redundant of the primary copy of the data, the secondary storage facility discrete from the primary storage facility and including:”</i>	62

6.	<i>[1e]: “a third redundancy appliance operable in a role for storing redundant data for the sequence of write requests in storage associated with the secondary facility; and”</i>	64
7.	<i>[1f]: “a fourth redundancy appliance for shadowing the third redundancy appliance wherein the fourth redundancy appliance assumes the role of the third redundancy appliance in the event of a fault at the third redundancy appliance.”</i>	66
B.	<i>Claim 2:</i>	68
1.	<i>“The data redundancy system according to claim 1, wherein the secondary storage facility assumes the role of the primary storage facility in response to a failover event, ...</i>	68
2.	<i>the failover event including a communication traffic condition having a greater portion of request traffic originating closer to the secondary storage facility than to the primary storage facility.”</i>	68
C.	<i>Claim 3: “The data redundancy system according to claim 1, wherein the second redundancy appliance assumes the role of the first redundancy appliance in response to a failover event selected from the group consisting of a manually initiated event and an automatically initiated event.”</i>	71
D.	<i>Claim 4:</i>	72
1.	<i>“The data redundancy system according to claim 1, wherein the first redundancy appliance updates the second redundancy appliance with information about the data in storage by forwarding a write record to the second redundancy appliance for each write request,”</i>	72
2.	<i>“the second redundancy appliance sending acknowledgements to the first redundancy appliance in batches in response to receiving write records.”</i>	73
E.	<i>Claim 5: “The data redundancy system according to claim 4, wherein the first redundancy appliance sends an acknowledgement to an originator of the write request in response to receiving the acknowledgement from the second redundancy appliance.”</i>	76

F.	<i>Claim 6: “The data redundancy system according to claim 1, wherein the first redundancy appliance is disjoint from the second redundancy appliance.”</i>	77
G.	<i>Claim 7: “The data redundancy system according to claim 1, wherein the second redundancy appliance maintains a shadow copy of data stored by the first redundancy appliance wherein the second redundancy appliance accesses the shadow copy of the data when the second redundancy appliance assumes the role of the first redundancy appliance.”</i>	79
H.	<i>Claim 8: “The data redundancy system according to claim 1, wherein the second redundancy appliance accesses the data stored in storage by the first redundancy appliance when the second redundancy appliance acts as the first redundancy appliance in response to a failover event.”</i>	81
X.	DETAILED ANALYSIS THAT CLAIMS 9-18 ARE OBVIOUS OVER <u>LUBBERS</u> AND <u>SICOLA</u> IN VIEW OF <u>OFEK</u>	83
A.	<i>Claim 9: “The data redundancy system according to claim 1, wherein the first redundancy appliance forwards redundant data to the secondary storage facility as send batches.”</i>	92
B.	<i>Claim 10: “The data redundancy system according to claim 9, wherein a send batch size is determined according to either a time interval for collecting transactions into the send batch or a predetermined quantity of transactions to be included in the send batch.”</i>	94
C.	<i>Claim 11: “The data redundancy system according to claim 9, wherein a send batch size is determined adaptively based on available communication bandwidth between the primary storage facility and the secondary storage facility.”</i>	95
D.	<i>Claim 12: “The data redundancy system according to claim 1, wherein a first write transaction in a send batch at the primary storage facility is replaced with a second write transaction where the second write transaction overwrites data written by the first write transaction.”</i>	97
E.	<i>Claim 13: “The data redundancy system according to claim 1, wherein the third redundancy appliance stores redundant data to the storage associated with the secondary facility according to receive batches.”</i>	99

- F. *Claim 14: “The data redundancy system according to claim 13, wherein the third redundancy appliance stores a receive batch after the receive batch is determined safe to apply.” 101*
- G. *Claim 15: “The data redundancy system according to claim 14, wherein a receive batch is determined safe to apply when a data integrity test is successfully performed.” 103*
- H. *Claim 16: “The data redundancy system according to claim 13, wherein a first write transaction in a first receive batch at the secondary storage facility is replaced with a second write transaction in a second receive batch wherein the second write transaction overwrites data written by the first write transaction and the first and second receive batches are combined.” 104*
- I. *Claim 17 106*
 - 1. *[17preamble]: “A data redundancy method for storing a primary copy of data at a primary storage facility and for storing data that is redundant of the primary copy at a secondary storage facility, the method comprising” 106*
 - 2. *[17a]: “receiving a sequence of write requests at a first redundancy appliance at the primary storage facility;” 107*
 - 3. *[17b]: “storing data for the sequence of write requests in primary storage at the primary storage facility;” 108*
 - 4. *[17c]: “logging at least a sequence of write requests for transfer to the secondary storage facility;” 109*
 - 5. *[17d]: “forwarding the logged write requests as send batches to a second redundancy appliance at the secondary storage facility disjoint from the primary storage facility, the second storage facility batching the received write request as receive batches;” 109*
 - 6. *[17e]: “storing data associated with the receive batches in secondary storage at the secondary storage facility” 111*
 - 7. *[17f]: “replacing a first write transaction in a receive batch at the second storage facility with a second write transaction wherein the second write transaction overwrites data written by the first write transaction; and” ..111*

8.	<i>[17g]: “the second redundancy appliance acting as the first redundancy appliance in response to a failover event.”</i>	112
J.	<i>Claim 18</i>	112
1.	<i>[18preamble]: “A data redundancy system, comprising:”</i>	112
2.	<i>[18a]: “a primary storage facility including at least one primary redundancy appliance for receiving a sequence of write requests and for storing data for the sequence of write requests in storage associated with the primary storage facility, the primary storage facility batching the write requests as send batches for transfer to a secondary storage facility; and”</i>	113
3.	<i>[18b]: “the secondary storage facility disjoint from the primary storage facility including at least one secondary redundancy appliance for receiving the sequence of write requests and for storing data for the sequence of write requests in storage associated with the secondary storage facility, the second storage facility batching the received write request as receive batches, the second redundancy appliance acting as the primary redundancy appliance in response to a failover event;”</i>	113
4.	<i>[18c]: “wherein a first write transaction in a receive batch at the second storage facility is replaced with a second write transaction where the second write transaction over writes data written by the first write transaction.”</i>	114
XI.	<i>CONCLUSION</i>	115

EXHIBIT LIST

Ex- No.	Brief Description
1001	U.S. Patent No. 7,152,182 (“182 Patent”)
1002	File History of U.S. Patent No. 7,152,182
1003	Declaration of Kevin C. Almeroth, Ph.D.
1004	<i>Curriculum Vitae</i> of Kevin C. Almeroth, Ph.D.
1005	U.S. Pub. No. 2003/0187847 (“ <u>Lubbers</u> ”)
1006	U.S. Patent No. 6,601,187 (“ <u>Sicola</u> ”)
1007	U.S. Pat. No. 5,901,327 (“ <u>Ofek</u> ”)
1008	Ramez Elmasri and Shamkant B. Navathe, FUNDAMENTALS OF DATABASE SYSTEMS (2d ed.1994)
1009	Original Complaint for Patent Infringement in <i>Valtrus Innovations Ltd. v. SAP America, Inc. et al</i> , 2:24-cv-00021-JRG (E.D. Tex.)
1010	Original Complaint for Patent Infringement in <i>Valtrus Innovations Ltd. v. SAP America, Inc. et al</i> , 2:24-cv-00533-JRG (E.D. Tex.)
1011	Director Vidal’s Memorandum Regarding Discretionary Denial Procedure, Dated June 21, 2022
1012	First Amended Docket Control Order in <i>Valtrus Innovations Ltd. v. SAP America, Inc. et al</i> , 2:24-cv-00021-JRG (E.D. Tex.)
1013	Motion to Dismiss for Lack of Standing (Public Version) in <i>Valtrus Innovations Ltd. v. SAP America, Inc. et al</i> , 2:24-cv-00021-JRG (E.D. Tex.)

Ex- No.	Brief Description
1014	Motion to Reconsider Ruling on Motion to Dismiss for Lack of Standing (Public Version) in <i>Valtrus Innovations Ltd. v. SAP America, Inc. et al</i> , 2:24-cv-00021-JRG (E.D. Tex.)
1015	Plaintiffs' Notice Regarding <i>Valtrus II</i> in <i>Valtrus Innovations Ltd. v. SAP America, Inc. et al</i> , 2:24-cv-00021-JRG (E.D. Tex.)
1016	SAP's Invalidity Contentions (Cover Pleading) in <i>Valtrus Innovations Ltd. v. SAP America, Inc. et al</i> , 2:24-cv-00021-JRG (E.D. Tex.)

I, Kevin C. Almeroth, do hereby declare as follows:

I. INTRODUCTION

1. I am making this declaration at the request of Petitioner in the matter of the *Inter Partes* Review of U.S. Patent No. 7,152,182 (“’182 Patent”) to Ji et al.

2. I am being compensated for my work in this matter at my standard hourly rate. I am also being reimbursed for reasonable and customary expenses associated with my work and testimony in this proceeding. My compensation is not contingent on the outcome of this matter or the specifics of my testimony.

3. I have been asked to provide my opinions regarding whether the subject matter of claims 1-18 (“the Challenged Claims”) of the ’182 Patent are anticipated or otherwise would have been obvious to a person having ordinary skill in the art (“POSITA”) at the time of the alleged invention, in light of the prior art. It is my opinion that the Challenged Claims would have been obvious to a POSITA.

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

- a. The ’182 Patent, Ex-1001.
- b. The prosecution history of the ’182 Patent (“’182 File History”), Ex-1002.
- c. U.S. Pub. No. 2003/0187847 (“Lubbers”), Ex-1005.
- d. U.S. Patent No. 6,601,187 (“Sicola”), Ex-1006.
- e. U.S. Patent No. 5,901,327 (“Ofeek”), Ex-1007.

f. The additional references I refer to below and/or are identified in the exhibit list.

5. In forming the opinions expressed below, I have considered: the documents listed above; the relevant legal standards as explained to me, including the standards for anticipation and obviousness; and my own knowledge and experience based upon my work in the field of network communications and security as described below.

6. Unless otherwise noted, all emphasis in any quoted material has been added. Claim terms are *italicized*. Exhibit titles are underlined.

II. QUALIFICATIONS AND PROFESSIONAL EXPERIENCE

7. My complete qualifications and professional experience are described in my Curriculum Vitae, a copy of which can be found in Ex.1004. The following is a brief summary of my relevant qualifications and professional experience.

8. I am currently a Professor Emeritus in the Department of Computer Science at the University of California, Santa Barbara (UCSB). While at UCSB, I held faculty appointments and was a founding member of the Computer Engineering (CE) Program, Media Arts and Technology (MAT) Program, and the Technology Management Program (TMP). I also served as the Associate Director of the Center for Information Technology and Society (CITS) from 1999 to 2012. I have been a faculty member at UCSB since July 1997.

9. I hold three degrees from the Georgia Institute of Technology: (1) a Bachelor of Science degree in Information and Computer Science (with minors in Economics, Technical Communication, and American Literature) earned in June 1992; (2) a Master of Science degree in Computer Science (with specialization in Networking and Systems) earned in June 1994; and (3) a Doctor of Philosophy (Ph.D.) degree in Computer Science (Dissertation Title: Networking and System Support for the Efficient, Scalable Delivery of Services in Interactive Multimedia System, minor in Telecommunications Public Policy) earned in June 1997. During my education, I took a wide variety of courses as demonstrated by my minors. My undergraduate degree also included a number of courses more typical of a degree in electrical engineering including digital logic, signal processing, and telecommunications theory.

10. One of the major concentrations of my research over the past 30+ years has been the delivery of multimedia content and data between computing devices, including through various network architectures. In my research, I have studied large-scale content delivery systems, and the use of servers located in a variety of geographic locations to provide scalable delivery to hundreds or thousands of users simultaneously. I have also studied smaller-scale content delivery systems in which content is exchanged between individual computers and portable devices. My work has emphasized the exchange of content more efficiently across computer networks,

including the scalable delivery of content to many users, mobile computing, satellite networking, delivering content to mobile devices, and network support for data delivery in wireless networks.

11. In 1992, the initial focus of my research was on the provision of interactive functions (e.g., VCR-style functions like pause, rewind, and fast-forward) for near video-on-demand systems in cable systems; in particular, how to aggregate requests for movies at a cable head-end and then how to satisfy a multitude of requests using one audio/video stream broadcast to multiple receivers simultaneously. This research has continually evolved and resulted in the development of techniques to scalably deliver on-demand content, including audio, video, web documents, and other types of data, through the Internet and over other types of networks, including over cable systems, broadband telephone lines, and satellite links.

12. An important component of my research has been investigating the challenges of communicating multimedia content, including video, between computers and across networks including the Internet. Although the early Internet was used mostly for text-based, non-real time applications, the interest in sharing multimedia content, such as video, quickly developed. Multimedia-based applications ranged from downloading content to a device to streaming multimedia content to be instantly used. One of the challenges was that multimedia content is

typically larger than text-only content, but there are also opportunities to use different delivery techniques since multimedia content is more resilient to errors. I have worked on a variety of research problems and used a number of systems that were developed to deliver multimedia content to users. One content-delivery method I have researched is the one-to-many communication facility called “multicast,” first deployed as the Multicast Backbone, a virtual overlay network supporting one-to-many communication. Multicast is one technique that can be used on the Internet to provide streaming media support for complex applications like video-on-demand, distance learning, distributed collaboration, distributed games, and large-scale wireless communication. The delivery of media through multicast often involves using Internet infrastructure, devices and protocols, including protocols for routing and TCP/IP.

13. Starting in 1997, I worked on a project to integrate the streaming media capabilities of the Internet together with the interactivity of the web. I developed a project called the Interactive Multimedia Jukebox (IMJ). Users would visit a web page and select content to view. The content would then be scheduled on one of a number of channels, including delivery to students in Georgia Tech dorms delivered via the campus cable plant. The content of each channel was delivered using multicast communication.

14. In the IMJ, the number of channels varied depending on the capabilities of the server including the available bandwidth of its connection to the Internet. If one of the channels was idle, the requesting user would be able to watch their selection immediately. If all channels were streaming previously selected content, the user's selection would be queued on the channel with the shortest wait time. In the meantime, the user would see what content was currently playing on other channels, and because of the use of multicast, would be able to join one of the existing channels and watch the content at the point it was currently being transmitted.

15. The IMJ service combined the interactivity of the web with the streaming capabilities of the Internet to create a jukebox-like service. It supported true Video-on-Demand when capacity allowed, but scaled to any number of users based on queuing requested programs. As part of the project, we obtained permission from Turner Broadcasting to transmit cartoons and other short-subject content. We also connected the IMJ into the Georgia Tech campus cable television network so that students in their dorms could use the web to request content and then view that content on one of the campus's public access channels.

16. More recently, I have also studied issues concerning how users choose content, especially when considering the price of that content. My research has examined how dynamic content pricing can be used to control system load. By

raising prices when systems start to become overloaded (i.e., when all available resources are fully utilized) and reducing prices when system capacity is readily available, users' capacity to pay as well as their willingness can be used as factors in stabilizing the response time of a system. This capability is particularly useful in systems where content is downloaded or streamed on-demand to users.

17. As a parallel research theme, starting in 1997, I began researching issues related to wireless devices and sensors. In particular, I was interested in showing how to provide greater communication capability to "lightweight devices," i.e., small form-factor, resource-constrained (e.g., CPU, memory, networking, and power) devices. Starting in 1998, I published several papers on my work to develop a flexible, lightweight, battery-aware network protocol stack. The lightweight protocols we envisioned were similar in nature to protocols like Bluetooth, Universal Plug and Play (UPnP) and Digital Living Network Alliance (DLNA).

18. From this initial work, I have made wireless networking—including ad hoc, mesh networks and wireless devices—one of the major themes of my research. My work in wireless networks spans the protocol stack from applications through to the encoding and exchange of data at the data link and physical layers.

19. At the application layer, even before the large-scale "app stores" were available, my research looked at building, installing, and using apps for a variety of purposes, from network monitoring to support for traditional computer-based

applications (e.g., content retrieval) to new applications enabled by ubiquitous, mobile devices. For example, my research has looked at developing applications for virtually exchanging and tracking “coupons” through “opportunistic contact” among mobile wireless devices (i.e., communication among devices moving into communication range with each other). In many of the courses I have taught there is a project component. Through these projects I have supervised numerous efforts to develop new “apps” for download and use across a variety of mobile platforms.

20. Toward the middle of the protocol stack, my research has also looked to build wireless infrastructure support to enable communication among a set of mobile devices unaided by any other kind of network infrastructure. These kinds of networks are useful either in challenged network environments (e.g., when a natural disaster has destroyed existing infrastructure) or when suitable support for network communication never existed. The deployment of such networks (or even the use of traditional network support) are critical to support services like disaster relief, catastrophic event coordination, and emergency services deployment.

21. Yet another theme is monitoring wireless networks, in particular different variants of IEEE 802.11 compliant networks, to (1) understand the operation of the various protocols used in real-world deployments, (2) use these measurements to characterize use of the networks and identify protocol limitations and weaknesses, and (3) propose and evaluate solutions to these problems. I have

successfully used monitoring techniques to study wireless data link layer protocol operation and to improve performance by enhancing the operation of such protocols. For wireless protocols, this research includes functions like network acquisition and channel bonding.

22. One theme in my wireless network research has been cross-layer solutions and innovations. As mentioned above, with greater wireless device use and network support, we envisioned new application paradigms and services, for example, when mobile devices come into contact with each other. Instead of relying on existing infrastructure to relay communication, the devices are able to discover each other and communicate directly. Other examples include discovering and using location information to enhance users' experiences. Network support and novel applications span use a variety of network architectures supporting users on foot, in vehicles, and across varying terrains and environments. Finally, we studied how communication efficiency can be supported through intelligent handoffs as well as location and movement prediction.

23. Protecting networks, including their operation and content, has been an underlying theme of my research almost since the beginning of my research career. Starting in 2000, I have been involved in several projects that specifically address security, network protection, and firewalls. After significant background work, a team on which I was a member successfully submitted a \$4.3M grant proposal to the

Army Research Office (ARO) at the Department of Defense to propose and develop a high-speed intrusion detection system. Key aspects of the system included associating streams of packets and analyzing them for viruses and other malware. Once the grant was awarded, we spent several years developing and meeting the milestones of the project. A number of my students worked on related projects and published papers on topics ranging from intrusion detection to developing advanced techniques to be incorporated into firewalls. I have also used firewalls, including their associated malware detection features, in developing techniques for the classroom to ensure that students are not distracted by online content.

24. My recent work ties some of the various threads of my past research together. I have investigated content delivery in online social networks and proposed reputation management systems in large-scale social networks and marketplaces. On the content delivery side, I have looked at issues of caching and cache placement, especially when content being shared and the cache has geographical relevance. We were able to show that effective caching strategies can greatly improve performance and reduce deployment costs. Our work on reputation systems showed that reputations have economic value, and as such, creates a motivation to manipulate reputations. In response, we developed a variety of solutions to protect the integrity of reputations in online social networks. The techniques we developed for content delivery and reputation management were

particularly relevant in peer-to-peer communication and recommendations for downloadable “apps.”

25. As an important component of my research program, I have been involved in the development of academic research into available technology in the market place. One aspect of this work is my involvement in the Internet Engineering Task Force (IETF). The IETF is a large and open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet. I have been involved in various IETF groups including many content delivery-related working groups like the Audio Video Transport (AVT) group, the MBone Deployment (MBONED) group, Source Specific Multicast (SSM) group, the Inter-Domain Multicast Routing (IDMR) group, the Reliable Multicast Transport (RMT) group, the Protocol Independent Multicast (PIM) group, etc. I have also served as a member of the Multicast Directorate (MADDOGS), which oversaw the standardization of all things related to multicast in the IETF. Finally, I was the Chair of the Internet2 Multicast Working Group for seven years.

26. My involvement in the research community extends to leadership positions for several academic journals and conferences. I am the co-chair of the Steering Committee for the ACM Network and System Support for Digital Audio and Video (NOSSDAV) workshop and on the Steering Committees for the

International Conference on Network Protocols (ICNP), ACM Sigcomm Workshop on Challenged Networks (CHANTS), and IEEE Global Internet (GI) Symposium. I have served or am serving on the Editorial Boards of IEEE/ACM Transactions on Networking, IEEE Transactions on Mobile Computing, IEEE Network, ACM Computers in Entertainment, AACE Journal of Interactive Learning Research (JILR), and ACM Computer Communications Review. I have co-chaired a number of conferences and workshops including the IEEE International Conference on Network Protocols (ICNP), IEEE Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON), International Conference on Communication Systems and Networks (COMSNETS), IFIP/IEEE International Conference on Management of Multimedia Networks and Services (MMNS), the International Workshop On Wireless Network Measurement (WiNMee), ACM Sigcomm Workshop on Challenged Networks (CHANTS), the Network Group Communication (NGC) workshop, and the Global Internet Symposium, and I have served on the program committees for numerous conferences.

27. Furthermore, in the courses I taught at UCSB, a significant portion of my curriculum covered aspects of the Internet and network communication including the physical and data link layers of the Open System Interconnect (OSI) protocol stack, and standardized protocols for communicating across a variety of physical media such as cable systems, telephone lines, wireless, and high-speed

Local Area Networks (LANs). The courses I have taught also cover most major topics in Internet communication, including data communication, multimedia encoding, and mobile application design. My research and courses have covered a range of physical infrastructures for delivering content over networks, including cable, Integrated Services Digital Network (ISDN), Ethernet, Asynchronous Transfer Mode (ATM), fiber, and Digital Subscriber Line (DSL). For a complete list of courses I have taught, see my curriculum vitae (CV).

28. In addition, I co-founded a technology company called Santa Barbara Labs that was working under a sub-contract from the U.S. Air Force to develop very accurate emulation systems for the military's next generation internetwork. Santa Barbara Labs' focus was in developing an emulation platform to test the performance characteristics of the network architecture in the variety of environments in which it was expected to operate, and, in particular, for network services including IPv6, multicast, Quality of Service (QoS), satellite-based communication, and security. Applications for this emulation program included communication of a variety of multimedia-based services, including video conferencing and video-on-demand.

29. In addition to having co-founded a technology company myself, I have worked for, consulted with, and collaborated with companies for nearly 30 years. These companies range from well-established companies to start-ups and include

IBM, Hitachi Telecom, Turner Broadcasting System (TBS), Bell South, Digital Fountain, RealNetworks, Intel Research, Cisco Systems, and Lockheed Martin.

30. Through my graduate education, leadership with CITS, involvement in TMP, role in the development of the Internet2 infrastructure, and consulting with ISPs, I have gained a strong understanding in the role of the Internet in our society and the challenges of deploying large-scale production networking infrastructure. CITS, since its inception, has looked at the role of the Internet in society, including how the evolution of technology have created communication opportunities and challenges, including, for example through disruptive technologies like P2P. TMP looks to focus on non-purely technical issues, including, for example, state-of-the-art business methods, strategies for successful technology commercialization, new venture creation, and best practices for fostering innovation. Through my industry collaborations and Internet2 work, I have developed significant experience in the challenges of deploying, monitoring, managing, and scaling communication infrastructure to support evolving Internet services like streaming media, conferencing, content exchange, social networking, and e-commerce.

31. I am a Member of the Association of Computing Machinery (ACM) and a Fellow of the Institute of Electrical and Electronics Engineers (IEEE).

III. LEVEL OF ORDINARY SKILL IN THE ART

32. I understand there are multiple factors relevant to determining the level of ordinary skill in the pertinent art, including (1) the levels of education and experience of persons working in the field at the time of the invention; (2) the sophistication of the technology; (3) the types of problems encountered in the field; and (4) the prior art solutions to those problems.

33. A Person of Ordinary Skill in the Art (“POSITA”) as of June 6, 2003, would have had a working knowledge of data storage, including data redundancy systems and methods used in a fault-tolerant manner. Such a POSITA would have had a bachelor’s degree in electrical engineering, computer science, or equivalent training, and approximately two years of work experience in software development in the field of network data processing and storage. Additional education could compensate for work experience, and vice versa.

34. For purposes of this Declaration, in general, and unless otherwise noted, my statements and opinions, such as those regarding my own experience and what a POSITA would have understood or known generally (and specifically related to the references I consulted herein), reflect the knowledge that existed in the relevant field as of the priority date of the ’182 Patent.

IV. RELEVANT LEGAL STANDARDS

35. I am not an attorney. In preparing and expressing my opinions and considering the subject matter of the '182 Patent, I am relying on certain basic legal principles that counsel has explained to me.

36. I understand that prior art to the '182 Patent includes patents and printed publications in the relevant art that predate the priority date of the '182 Patent. For purposes of this Declaration, I am applying June 6, 2003, as the priority date of the '182 Patent, but I take no position on that issue. Even if the Patent Owners contend the date of invention is earlier than June 6, 2003, that would not change my opinions to the extent that would not disqualify any of the applied references as prior art.

37. I have been informed by counsel that a claimed invention is invalid if it is “anticipated” by the prior art. I understand that, for a patent claim to be “anticipated” by the prior art, each and every requirement of the claim must be found, expressly or inherently, in a single prior art reference as recited in the claim. I understand that claim limitations that are not expressly described in a prior art reference may still be there if they are “inherent” to the thing or process being described in the prior art.

38. I understand that if a reference incorporates other documents by reference in their entirety, the incorporating reference and the incorporated reference(s) should be treated as a single prior art reference.

39. I understand that it is acceptable to consider evidence other than the information in a particular prior art document to determine if a feature is necessarily present in or inherently described by that reference.

40. I have been informed by counsel that a claimed invention is unpatentable under 35 U.S.C. § 103 if the differences between the claimed 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 POSITA. I have also been informed by counsel that the obviousness analysis considers 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.

41. I have been further informed by counsel that there are several recognized rationales for combining references or modifying a reference to show obviousness. These rationales include: (a) combining prior art elements according to known methods to yield predictable results; (b) simple substitution of one known element for another to obtain predictable results; (c) use of a known technique to improve a similar device (method, or product) in the same way; (d) applying a known technique to a known device (method, or product) ready for improvement


to yield predictable results; (e) choosing from a finite number of identified, predictable solutions, with a reasonable expectation of success; and (f) some teaching, suggestion, or motivation in the prior art that would have led a POSITA to modify the prior art or to combine prior art teachings to arrive at the claimed invention.

V. OVERVIEW OF THE '182 PATENT

42. The '182 Patent issued from the standalone non-provisional Application No. 10/456,053, which was filed on June 6, 2003. I have been asked to assume that June 6, 2003 is the effective filing date, but I take no position on that issue.

43. I have reviewed the file history of the 182 Patent and note that none of the activity there is pertinent to my analysis.

44. The '182 Patent describes techniques for performing data redundancy operations in a fault-tolerant manner. Ex-1001, Abstract, 1:29-31, 2:64-66. The computer system includes a primary storage facility and a secondary storage facility at different locations. Ex-1001, 1:55-59, 2:66-3:1, 3:18-31. The secondary storage facility stores data that is redundant of the primary storage facility. Ex-1001, 1:55-59, 2:66-3:1, 4:1-5. Figure 2 shows operation of the primary and secondary storage facilities.



Ex-1001, FIG. 2, 29:29-31.

45. The computer system has two modes of operation: normal operating conditions (e.g., in the absence of a fault), and operation in the event of a fault. Ex-1001, 4:17-28, 4:50-68, 12:66-13:9, 13:32-33.

46. The primary and secondary storage facilities include data redundancy appliances, which can be implemented by hardware or software. Ex-1001, 1:57-67, 3:1-9, 4:29-41. A first data redundancy appliance of the primary storage facility receives a sequence of write requests and stores data for the sequence of write

requests in storage associated with the primary storage facility. Ex-1001, 1:57-67, 3:1-9. A data redundancy appliance of the secondary storage facility shadows the first redundancy appliance and assumes its role in the event of a fault at the first redundancy appliance. Ex-1001, 1:57-67, 3:1-9. In certain arrangements, multiple redundancy appliances can be included in each storage facility. For example, in contrast to Figure 2 above (showing a single redundancy appliance in the primary and secondary storage facilities), Figure 11 shows multiple redundancy appliances at each storage facility.




FIG. 11

47. As demonstrated below, all the features recited in the challenged claims of the '182 Patent were known in the prior art before the alleged invention date of the '182 Patent.

VI. CLAIM CONSTRUCTION

48. It is my understanding that in order to properly evaluate the '182 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 construed under the so-called *Phillips* standard, under which claim terms are given their ordinary and customary meaning as would have been understood by a POSITA in light of the specification and prosecution history, unless the inventor has set forth a special meaning for a term. I have also been informed that claim terms only need to be construed to the extent necessary to resolve the unpatentability inquiry.

49. I do not believe that any claim constructions would affect my analysis or conclusion that the claims are obvious over the asserted prior art.

VII. IDENTIFICATION OF HOW THE CLAIMS ARE UNPATENTABLE

50. The discussion in this Declaration provides a detailed analysis of how the asserted prior art teaches each limitation of the Challenged Claims.

51. As part of my analysis, I have considered, and discuss in detail, the scope and content of the prior art and any differences between the alleged invention and the prior art. As detailed below, the scope and content of the Lubbers prior art

reference alone or in combination with the Sicola and Ofek prior art references includes every element of the asserted claims without any meaningful difference between the teachings of those references and the elements of the asserted claims.

52. I am not aware of any secondary considerations. The Patent Owners have not identified any secondary considerations of nonobviousness so I have not considered any in my analysis.

53. I have considered the teachings in the prior art from the perspective of a POSITA and it is my conclusion that a POSITA would have found it obvious to combine the teachings of those known elements in the way they are recited in the Challenged Claims as of June 6, 2003.

54. Regardless of whether secondary considerations exist or not, they would not rise to the level to overcome the strong obviousness arguments presented below. I do not believe any secondary considerations would alter my overall conclusion that the claims would have been obvious to a POSITA as of June 6, 2003. Nevertheless, should the Patent Owners provide secondary considerations and argue those considerations rise to the level of maintaining the validity of the Challenged Claims, I reserve the right to respond.


VIII. THE PRINCIPAL PRIOR ART

A. Summary of Lubbers (Ex-1005)

55. The application underlying U.S. Pub. No. 2003/0187847 to Lubbers et al. ("Lubbers") was filed on March 26, 2002, and published on October 2, 2003. I understand that Lubbers is therefore prior art to the '182 Patent at least under pre-AIA 35 U.S.C. § 102(e).

56. Lubbers "relates generally to computer-based information storage systems," that are "readily extensible to provide multiple replication, load balancing, and disaster tolerance...." Ex-1005, [0001], [0010].


57. Lubbers discloses a multi-site storage area network (SAN) that includes multiple sites (Site 1, Site 2, and Site 3) in Figure 1:



Ex-1005, [0028], FIG. 1.

58. Each site includes one or more storage cells 101 accessible by host computers 102. Ex-1005, [0028]-[0029]. Storage cells 101 and hosts 102 couple to a data communication network 103. Ex-1005, [0030]. Any storage cell 101 can function as a primary storage location for any of hosts 102 while simultaneously functioning as a secondary storage location for a replica of data from another storage cell 101. Ex-1005, [0036].

59. Lubbers discloses various features in a disaster tolerance configuration with redundant components (shown in Figure 3):



Ex-1005, [0037]; FIG. 3.

60. Lubbers discloses that one goal of the disaster tolerant configuration in Figure 3 is to provide hardware redundancy for multiple levels of fault tolerance in order to keep the primary and alternate sites running. Ex-1005, [0038]. Lubbers discloses that if one component fails at a site, the data replication management (DRM) will failover to a redundant component at that site to allow continued operation. Ex-1005, [0038].

61. As far as controllers 105, Lubbers discloses that the physical storage devices 302 are coupled to controllers 105 by a high-speed redundant connection such as fibre channel arbitrated loop (FCAL) connections 303. Ex-1005, [0040].

Lubbers discloses that each of the tens or hundreds of individual hard disk drives can be accessed by either controller 105 in a storage cell. Ex-1005, [0040]. Lubbers discloses a registration and discovery process that allows each controller 105 to connect to other controllers, as well as a “failover” condition in which one controller 105 fails. Ex-1005, [0045]-[0046]. Lubbers discloses that redundant components can operate simultaneously and in parallel to provide improved throughput as well as serving during failover to replace functionality of failed systems and connections. Ex-1005, [0064]. And Lubbers explains the benefit is great flexibility in implementing a particular system to emphasize any balance of low latency, high availability, disaster tolerance, and cost. Ex-1005, [0064].

62. In operation, Lubbers discloses that the storage controllers 105 periodically exchange “heartbeats” as well as acknowledgements with peer controllers, which are pings with useful information in managing data transfer processes between peer controllers. Ex-1005, [0068]. Lubbers discloses operation can be synchronous or asynchronous and so Lubbers discloses logging and merging to track and preserve the order of write operations. Ex-1005, [0059]-[0060]. The sites may maintain logs of recent transactions (GSN, VBA, LUN, Block Count) by copying input/output (I/O) operations from the host computer into the log unit’s cache buffers. Ex-1005, [0076]. Lubbers discloses that a processor of the controller associated with the data log may monitor and maintain a ratio of

logging writes to merge writes. Ex-1005, [0080]. Those merges write I/Os are transmitted to a destination storage system where they are merged into the target LUN on the destination storage system. Ex-1005, [0080].

63. Operation can be synchronous or asynchronous, so Lubbers discloses logging and merging to track and preserve the order of write operations. Ex-1005, [0059]-[0060]. The sites may maintain logs of recent transactions (GSN, VBA, LUN, Block Count) by copying input/output (I/O) operations from the host computer into the log unit's cache buffers. Ex-1005, [0076]. Lubbers discloses that a processor of the controller associated with the data log may monitor and maintain a ratio of logging writes to merge writes. Ex-1005, [0080]. Those merges write I/Os are transmitted to a destination storage system where they are merged into the target LUN on the destination storage system. Ex-1005, [0080].

B. Summary of Sicola (Ex-1006)

64. The application underlying U.S. Patent No. 6,601,187 to Sicola et al. ("Sicola") was filed on March 31, 2000, and published upon issuance on July 29, 2003. I understand that Sicola is therefore prior art to the '182 Patent under at least 35 U.S.C. § 102(e).

65. Lubbers incorporates Sicola by reference. Ex-1005, [0008] (stating "Compaq Corporation introduced an array controller referred to as the HSG80, that implemented Data Replication Management features, as described in U.S. patent

application Ser. No. 09/539,745 assigned to the assignee of the present application and incorporated herein by reference, that implemented an architecture utilizing redundant storage controllers”).

66. Sicola recognizes “the critical role played by the storage controller with respect to computer system I/O performance” and the desirability “to minimize the potential for interrupted I/O service due to storage controller malfunctions.” Sicola teaches that “all replication functionality is totally self-contained within each array controller and automatically done without user intervention.” Ex-1006, 5:2–4.

67. Sicola discloses a system 100 that “employs a switch, dual-fabric, Fibre Channel configuration to provide a disaster tolerant storage system.” Ex-1006, 6:1–3. “FIG. 2 illustrates an exemplary configuration ... which comprises a switched dual fabric, disaster-tolerant storage system 100,” as shown below:





Fig. 2

Ex-1006, 6:64–66; FIG. 2. “[D]ata storage sites 218 and 219 each respectively comprise two hosts 101/101A and 102/102A, and two storage array controllers 201/202 and 211/212 connected to storage arrays 203, and 213, respectively.” Ex-1006, 7:1–5. Sicola discloses that “at site 218, host 101 is connected to switches 204 and 214 via respective paths 231A and 231B; host 101A is connected to the switches via paths 241A and 241B.” Ex-1006, 7:13–16.

68. Sicola discloses that “array controller pair A1/A2 comprise storage node 301, and array controller pair B1/B2 comprise storage node 302.” Ex-1006 at 8:23–25. Sicola describes the operation of two array controllers and describes


Ex-1003

operation of the two array controller heartbeat timers in FIG. 6A (inter-site controllers) and 6B (intra-site controllers):



Ex-1006, 9:35–37; FIGS 6A, 6B. In the embodiment in FIG. 6A, Sicola discloses that “host computer 101 sends requests to write data to array 203 via controller A1 (201).” Ex-1006, 9:41–43.

69. Sicola also describes synchronous system operation (FIG. 7) and asynchronous system operation (FIGS. 8A, 8B), as shown below:



70. In synchronous operation (FIG. 7), Sicola explains that host computer 101 first issues a write command to local controller A1 (201), the controller writes the data into its write-back cache through the normal cache manager path, and then calls the peer-to-peer remote copy (PPRC) manager 515. Ex-1006, 11:44–53. The PPRC manager 515 sends the write data to the remote target controller B1 (211), and across fabric 103A. Ex-1006, 11:53–57. The remote target controller B1 writes data to its write-back cache (or directly to media if a write through operation), and then controller B1 sends the completion status back to initiator controller A1. Ex-1006, 11:57–61. Once PPRC manager 515 in controller A1 receives a completion status from target controller B1, it notifies VA 530 of the completion, and VA 530 completes the write in the normal path (media write if write through), releases the cache lock, and completes the operation by sending a completion status to host 101. Ex-1006, 11:61–67.

71. In asynchronous operation, (FIGS. 8A, 8B), operation is largely the same. However, Sicola teaches that asynchronous system operation involves a micro-logging and micro-merge process in which write commands can be exchanged and mirrored between controller-specific write back caches for both controllers in a redundant controller pair. Ex-1006, 12:25–32, 12:39–43, 44–59, 13:6–13.


C. Summary of Ofek (Ex-1007)

72. The application underlying U.S. Patent No. 5,901,327 to Ofek (“Ofek”) was filed on March 17, 1997, and published upon issuance on May 4, 1999. I understand that Ofek is therefore prior art to the '182 Patent under at least 35 U.S.C. § 102(b).

73. Ofek “relates to data storage” Ex-1007, 1:15, such as “automatically providing and maintaining a copy or mirror of data stored at a location remote from the main or primary data storage device.” Ex-1007, 1:16–19.

74. Ofek discloses “a system and method for automatically providing and maintaining a copy or mirror of data stored at a location remote from the main or primary data storage device.” Ex-1007, 1:15–19.

75. Ofek discloses a data storage system 10 that includes a host computer system 12 coupled to a primary data storage system 14, which includes a primary data storage system controller 16 that receives data form the host 12 and is coupled to a storage device 20. 4:57–5:2. The data storage system 10 also includes a secondary data storage system 46 with a secondary data storage system controller 44 (shown in FIG. 1 below):



Ex-1007, 4:35–38; FIG. 1.

76. Ofek discloses that “[i]n order to more effectively use the capacity of the data link, the data storage system containing the primary (R1) volume bundles the write data for all write commands in the chain into a single write command for transmission over the data link to the secondary data storage system containing the secondary (R2) volume.” Ex-1007, 4:35–38, 34:66–35:6 (it is very desirable to ‘bundle’ the write data for all write commands in the channel command word chain into a single write command transmitted over a link to the secondary storage system 246.”).

77. Ofek also discloses that “link queue 504 contains a respective entry for each command that is transmitted over the link 240.” Ex-1007, 35:18–20.

78. Ofek discloses “the link adapter 236 services the link queue 504 by transferring the write data across the link 240,” (shown below in FIG. 18 of Ofek):




FIG. 18

Ex-1007, 35:52–54; FIG. 18.

79. Ofek discloses that “[i]n the preferred implementation, the entry in the link queue 504 includes a one-byte password for confirming the validity of the entry, a logical volume number specifying the secondary (R2) volume, a number identifying the channel adapter that created the entry, a pointer to a starting location in the link buffer 505 for additional information for the command to be sent to the secondary storage system containing the specified secondary (R2) volume, and

lock information by which a particular channel adapter or link.” Ex-1007, 35:60–36:2.

80. Ofek discloses that the data storage system “achieves nearly 100 percent data integrity by assuring that all data is copied to a geographically remote site, and in those cases when a back-up copy is not made due to an error of any sort, an indication is stored that the data has not been copied, but instead must be updated at a future time.” Ex-1007, 3:10–18. Ofek discloses the use of a one-byte password for confirming the validity of the entry, among other data integrity and verification mechanisms. Ex-1007, 35:60-36:2.

IX. DETAILED ANALYSIS THAT CLAIMS 1-8 ARE OBVIOUS OVER LUBBERS AND SICOLA

81. Lubbers teaches or suggests every element of claims 1-8. Even had Lubbers alone not expressly taught every element of claims 1-8, Sicola, which is incorporated by reference in Lubbers, provides clarifying details confirming Lubbers would have rendered obvious every element of claims 1-8. Since Sicola is incorporated by reference in Lubbers, it is my understanding that Sicola is properly considered part of the disclosure of Lubbers.

82. Regardless of whether Sicola is incorporated by reference, it would have been obvious for a POSITA to combine the teachings of Lubbers and Sicola.

83. A POSITA reading Lubbers would have been motivated to look to Sicola for additional details on the operation of the controllers described in Lubbers because Lubbers specifically references the controllers of Sicola, and Sicola provides additional details of an architecture utilizing redundant storage controllers, among other things, that would be of interest to a POSITA. Additionally, Lubbers and Sicola were originally assigned to the same company and have two common inventors: Stephen Sicola and Susan Elkington. The subject matter of Lubbers also significantly overlaps with Sicola, as demonstrated by the similar fields of the invention, similar architectures described and shown in the figures, and Lubbers expressly referring to Sicola for its array controllers. Thus, it would have been obvious for a POSITA looking at the teachings of Lubbers to study the teachings of Sicola beyond the clarifying details on the controllers. As explained below, the combined teachings of Lubbers and Sicola includes the architecture and functionality of Lubbers with additional implementation details of Sicola.

84. It would have been obvious to use clarifying details from Sicola for the operation and communication between controllers of Lubbers. First, Lubbers teaches how redundant controllers can update one another (*“first redundancy appliance updates the second redundancy appliance with information about data in storage”*). For example, Lubbers explains that if a “controller does not have the

most current data, then the controller goes offline to the host to receive a full copy from the controller that assumes the role of the source controller.” Ex-1005, [0073]. Sicola further explains how redundant controllers update one another in a multi-site storage system. For example, Sicola explains that asynchronous system operation involves micro-logging and micro-merge processes in which write commands can be exchanged and written to controller-specific write back caches. Ex-1006, 12:16-30. Sicola also describes that outstanding writes logged to initiator controller A1’s write-back cache, are mirrored in partner controller A2’s (mirrored) write-back cache. Ex-1006, 13:2-13.

85. Second, Lubbers teaches that controllers can send acknowledgements to one another in batches after receiving write requests (“*second redundancy appliance sending acknowledgments to the first redundancy appliance in batches in response to receiving write records*”). For example, Lubbers explains that each controller periodically sends “heartbeats,” or pings to its peer controllers, and the peer controllers respond to the ping with a message including information useful in managing data transfer processes between peer controllers. Ex-1005, [0068] Sicola further explains controllers communicating and sending acknowledgments in this manner. For example, Sicola describes intra-site controller heartbeat timer operation. Ex-1006, 10:22-26. Additionally, Sicola teaches that during normal operation, companion controllers periodically send pings to each other via Dual

Asynchronous Receiver/Transmitters, and the pings are received unless either controller has timed out. Ex-1006, 10:45-53.

86. Based on these teachings, a POSITA would have been motivated to enhance the teachings of Lubbers to further improve upon the stated goals of maintaining integrity and availability of data, minimizing latency and improving data transmission path performance, and minimizing controller load. A POSITA would have found it obvious to look to the teachings of Sicola for these improvements because Sicola also recognizes many of the same objectives. Ex-1006, 9:4-6 (describing that local host 101 is generally not allowed to access the storage logical unit numbers (LUN) 410' to preserve integrity of the backup), 11:35-40 (explaining that having write-back caches at the standby site can somewhat mitigate latency), 8:3-7 (explaining the use of multi-pathing software that dynamically allows failover between storage paths and static load balancing of storage volumes (LUNs) between paths).

87. A POSITA would also have been motivated to look to the teachings of Sicola to further improve the teachings of Lubbers, because Lubbers teaches that system has different operation modes to focus on certain objectives criteria and tradeoffs, and the operation mode can be changed. For example, Lubbers teaches that multi-pathing software allows for failover between storage paths as well and static load balancing of storage volumes (logical units of storage capacity) between

paths of controllers. Ex-1005, [0042]. Additionally, Lubbers explains how desired controller paths and storage are selected from criteria including path availability, path performance, path load, controller load, and the like. Ex-1005, [0042].

88. A POSITA would have been motivated to combine Lubbers and Sicola because both patents solve the same data storage problem in the same way: by providing multiple levels of redundancy. For example, Lubbers teaches a goal of the disaster tolerant configuration is to provide hardware redundancy for multiple levels of fault tolerance to keep the primary and alternate sites running at almost non-stop metrics. Ex-1005, [0038]. Sicola also teaches that it solves prior art problems and advances the field by providing a system that includes a completely redundant configuration with dual Fibre Channel fabric links interconnecting each of the components. Ex-1006, 3:63-4:6.

89. Further, the benefits of data redundancy and distributed databases were well-known as of June 6, 2003, when the application underlying the '182 Patent was filed. Ex-1008, 703-727. For example, data redundancy (1) “can improve availability remarkably as the entire system can continue to operate as long as at least one site is up;” and (2) “improves performance of retrieval for global queries, because the result of such a query can be obtained locally from any one site.” Ex-1008, 711.

90. Additionally, both Lubbers and Sicola are analogous art to the '182 Patent for at least two reasons. First, Lubbers and Sicola are both directed toward the same field of endeavor of the '182 Patent, which is redundant data storage. Ex-1001, 1:29-31; Ex-1005, [0001], [0007]; Ex-1006, 1:8-14.

91. Second, POSITAs knew that communication delays and component failures were a common data storage problem as of June 6, 2003. For example, Lubbers and Sicola each recognize that problem and describe improved techniques for redundant data storage to overcome communication delays and component failures. Ex-1005, [0038]; Ex-1006, 3:63-4:6. The '182 Patent also confirms that communication delays and component failures were a problem in data storage as of June 6, 2003. Ex-1001, 1:45-51.

92. A POSITA would have reasonably expected to succeed in incorporating the teachings of Sicola with the teachings of Lubbers because Sicola merely provides additional details for similar systems and components implementing similar techniques already taught by Lubbers. Ex-1005, [0068], [0076]; Ex-1006, 12:16–30, 13:6–13

93. Finally, Lubbers and Sicola are at the same level of generality as the '182 Patent such that a POSITA would have been able to apply that same level of skill in the art in understanding Lubbers and the ways in which the teachings could

be implemented. Thus, a POSITA would have reasonably expected to succeed in combining teachings of Lubbers and Sicola.


94. As one example, A POSITA would have reasonably expected to succeed in having redundant controllers for shadowing one another because Lubbers already contemplates the importance of hardware redundancy for multiple levels of fault tolerance to provide seamless operation and maintain performance when a component fails. Ex-1005, [0038].

A. Claim 1

1. [1preamble]: “A data redundancy system, comprising:”


95. For purposes of this analysis, I have been told to assume that the preamble is limiting for this claim.

96. Lubbers teaches a multi-site storage area network (SAN) which “relates generally to computer-based information storage systems” that is a “*data redundancy system*.” Ex-1005, [0001], [0010]. (explaining the SAN provides data redundancy functions, including “multiple replication, load balancing, and disaster tolerance....”). The SAN (“*data redundancy system*”) includes multiple sites, such as Site 1, Site 2, and Site 3 shown in the embodiment in Figure 1:



Ex-1005, [0001], [0010], [0028], FIG. 1.

97. One embodiment of the SAN (*“data redundancy system”*) is illustrated in Figure 3, which shows a disaster tolerant configuration with redundant components:




Ex-1005, [0038], FIG. 3 (annotated).

98. Accordingly, Lubbers teaches or suggests this limitation.

2. **[1a]: “a primary storage facility for storing a primary copy of data, the primary storage facility including;”**

99. Within the multi-site storage area network (SAN) (*“data redundancy system”*), there are multiple sites with storage cell(s) 101, and a storage area network (SAN) management appliance (SMA) 107 located at each site to operate storage cell(s) 101. Ex-1005, [0028] (describing a SAN with three sites physically separated from one another, and each site including one or more storage cells 101),

[0034] (describing SMAs 107 are typically located at each site to manage operation of local storage cells 101). And Lubbers teaches additional embodiments of multiple sites (including “*primary storage facilities*”) (shown in Figure 3):



Ex-1005, [0037]-[0040], FIG. 3 (annotated).

100. The primary site of Lubbers, including the controllers 105 (controllers 1 and 2), storage management appliances 107, physical storage device 302, fibre channel arbitrated loop (FCAL) connections 303, switches 301, and host 102, is a “*primary storage facility*” because it is a “local” copy of the data. Ex-1001, 1:33-38 (referring to one copy of data as a ‘primary’ or ‘local’ copy, and a redundant copy of the data as a ‘secondary’ or ‘slave’ copy); Ex-1005, [0037] (explaining

Ex-1003

storage management appliances 107 include bridge and data replication management consoles (DRMC) that manage controllers 105), [0040] (describing physical storage devices 302 coupled to controllers 105 by FCAL connections 303), [0041] (describing dual fabric implemented by switches 301, dual hosts 102, and dual storage topology).


101. The “*primary storage facility*” is “*for storing a primary copy of data*” because it includes physical storage devices 302 coupled to controllers 105 by an FCAL connection allowing “each of the tens or hundreds of individual hard disks to be accessed not only at high speed, but by either controller in a storage cell.” Ex-1005, [0040]; *see also* Ex-1005, [0030] (explaining each site includes storage cells with hundreds of gigabytes to terabytes of physical storage capacity accessible by controllers 105), [0010] (explaining that SAN systems enable storing multiple copies of data at various physical locations throughout the system), [0020] (describing storage implemented at various sites is accessed by a host computer 102 that requires data storage), [0036] (explaining storage cell 101 can simultaneously be primary storage for a host and secondary storage for a replica of data from another storage cell 101), [0038] (describing that, if a failure occurs at the primary site, data processing can be resumed at an alternate site where data is intact with little or no downtime).

102. Accordingly, Lubbers teaches or suggests this element.

3. [1b]: “a first redundancy appliance operable in a role for receiving a sequence of write requests and for storing data for the sequence of write requests in storage associated with the primary storage facility; and”

103. This element would have been obvious based on the combined teachings of Lubbers and Sicola. A POSITA would have been motivated to combine the teachings of Lubbers and Sicola for the reasons described above in ¶¶81-94.

104. Lubbers teaches a first controller 105 (controller 1) (“*first redundancy appliance*”) at the primary site (“*primary storage facility*”) and physical storage device 302 (“*storage associated with the primary storage facility*”), as shown in annotated Figure 3:



Ex-1005, FIG. 3 (annotated).

105. The '182 Patent uses the phrase “*a sequence of write requests*” to refer to sequential-in-time write requests. Ex-1001, 19:7-9 (explaining “data blocks are written to sequentially in time....”). The combined teachings of Lubbers and Sicola would have rendered obvious “*a sequence of write requests*” as that phrase is used in the '182 Patent.

106. The first controller 105 is “*a first redundancy appliance*” because the first controller 105 is “*operable in a role for receiving a sequence of write requests.*” The controllers 105 of Lubbers “are configured to handle both operational data transactions and data replication management (DRM) traffic.” Ex-1005, [0044]. After receiving sequential-in-time write requests from the host, the controllers 105 can pack “the log data from as many host writes as possible into a single buffer.” Ex-1005, [0083]; *see also* Ex-1005, [0011] (explaining that “I/O commands and the data associated with those commands are written to the data log”), [0030] (explaining “controllers 105 virtualize the physical storage capacity such that it is configurable into logical units (LUNs) of storage capacity”), [0032] (explaining “[s]torage controllers 105 manage tasks of allocating storage capacity to specified LUNs, monitoring and maintaining integrity of LUNs, moving data between physical storage devices, and other functions that maintain integrity and availability of the data stored therein”).

107. Sicola provides additional details of how a controller (*“first redundancy appliance”*) is *“in a role for receiving a sequence of write requests.”* Ex-1006, 11:45-47 (explaining, with reference to Figure 7, that “host computer 101 issues a write command to local controller A1 (201), which receives the command at host port 109 over path 221*h*”), 12:17-20 (same for Figure 8A).

108. The physical storage devices 302 (e.g., hard disks) at the primary site of Lubbers are *“storage associated with the primary storage facility.”* Ex-1005, [0040] (explaining that each physical storage device 302 can take the form of “tens or hundreds of individual hard disk drives”), [0030]-[0031] (explaining how “all data objects are identified with a universally unique identification (UUID) and UUIDs are stored persistently in metadata structures within physical storage,” which can “be mapped and presented by any storage controller 105.”), [0032] (explaining “[s]torage controllers 105 manage tasks of allocating storage capacity to specified LUNs” and “moving data between physical storage devices”), [0005] (describing that “SAN systems enable the possibility of storing multiple copies or ‘replicas’ of data at various physical locations throughout the system”), [0011] (explaining that the “data storage system maintains a data log that may be activated when the destination volume is inaccessible”).

109. The first controller (controller 1) (*“a first redundancy appliance”*) is coupled to physical storage device 302 (*“storage associated with the primary*

storage facility”) “for storing data for the sequence of write requests.” Ex-1005, [0040] (explaining “[p]hysical storage devices 302 are coupled to controllers 105 by a high-speed redundant connection such as fibre channel arbitrated loop (FCAL) connections 303”), [0040] (describing that the “FCAL connection allows each of tens or hundreds of individual hard disk drives to be accessed at high speed and by either controller 105”), [0045] (explaining how “each controller 105 is identified by a storage cell unique identification number (UUID)”), [0054] (describing that a “controller may support tens, hundreds, or thousands of copy sets”), [0076] (explaining that “[d]uring the data logging process, I/O operations from the host computer are copied into the log unit’s cache buffers”).


110. Accordingly, Lubbers and Sicola would have rendered obvious this element.

4. ***[1c]: “a second redundancy appliance for shadowing the first redundancy appliance wherein the second redundancy appliance assumes the role of the first redundancy appliance;”***

111. This element would have been obvious based on the combined teachings of Lubbers and Sicola. A POSITA would have been motivated to combine the teachings of Lubbers and Sicola for the reasons described above in ¶¶81-94.

112. Lubbers teaches second controller (controller 2) (“*a second redundancy appliance*”) at the primary site (“*primary storage facility*”), shown in annotated

Figure 3, which uses the same reference numerals as other controllers 105. A POSITA would understand that in the context of Lubbers, using the same reference numerals in a patent application implies the individual components have the same structure and function as one another.



Ex-1005, [0040], FIG. 3 (annotated).

113. The second controller (controller 2) is “*a second redundancy appliance.*” Ex-1005, [0040] (explaining that “[p]hysical storage devices 302 are coupled to controllers 105 by a high-speed redundant connection such as fibre channel arbitrated loop (FCAL) connections 303”), [0045] (describing how “each controller 105 is identified by a storage cell unique identification number (UUID)”), [0083] (explaining that “controllers 105 are configured to handle both

operational data transactions and data replication management (DRM) traffic”), [0030] (describing that the “controllers 105 virtualize the physical storage capacity such that it is configurable into logical units (LUNs) of storage capacity”).

114. The ’182 Patent uses “*shadowing*” to refer to a redundancy appliance having “the same data and state information” as another redundancy appliance for when a fault occurs (Ex-1001, 21:42-44), and the “shadowing redundancy appliance 202” takes over for the failed redundancy appliance (Ex-1001, 21:44-47). The combined teachings of Lubbers and Sicola would have rendered obvious “*shadowing*” as that term is used in the ’182 Patent.

115. The second controller (controller 2) of Lubbers (“*second redundancy appliance*”) is “*for shadowing*” and can “*assume the role of the first redundancy appliance.*” Ex-1005, [0031] (explaining that “[a]ll objects in the system are identified with a universally unique identification (UUID),” which are “stored persistently in metadata structures within the physical storage which enables all or part of a storage cell, and all the LUNs implemented therein, to be mapped and presented by any storage controller”), [0038] (describing a “goal of a disaster tolerant configuration is to provide hardware redundancy for multiple levels of fault tolerance,” allowing for “failover” where the functionality of a failed component can be taken over by a redundant component at a site).

116. Sicola further describes how a “*secondary redundancy appliance*” is “*for shadowing*” and can “*assume the role of the first redundancy appliance*.” Ex-1006, 12:16-30 (describing asynchronous operation involving a “micro-logging” process in which the controller “writes the data into its write-back cache”), 13:6-13 (explaining a “micro-merge” process during asynchronous operation in which “outstanding writes were logged to the initiator controller A1’s write-back cache, which is also mirrored in partner controller A2’s (mirrored) write-back cache, so that the cache data is available to controller A2 if controller A1 fails”).

117. Even had Lubbers not expressly taught the second redundancy appliance “for shadowing” the first redundancy appliance, a POSITA would have found it obvious to use the second controller (“*second redundancy appliance*”) for shadowing the first controller (“*first redundancy appliance*”). Ex-1008, 711. A POSITA would have reasonably expected to succeed in having the second controller (“*second redundancy appliance*”) shadow the first controller (“*first redundancy appliance*”) because it would involve merely using the controllers and the write-back caches of the controllers of Lubbers and Sicola for their intended functions to achieve predictable results. Ex-1008, 711. A POSITA would have been motivated to have the second controller shadow the first controller to allow for rapid recovery of user data from a disaster or significant error event at a data

processing facility, as taught by Lubbers, and for the reasons stated above. *See* ¶¶81-94.


118. Accordingly, Lubbers and Sicola would have rendered obvious this element.

5. *[1d]: “a secondary storage facility for storing data that is redundant of the primary copy of the data, the secondary storage facility discrete from the primary storage facility and including:”*

119. Lubbers teaches that the multi-site storage area network can include any number of sites with the same, similar, or entirely different configurations.

120. The alternate (secondary) site of Lubbers, including the controllers 105 (controllers 1 and 2), storage management appliances 107 (SMAs A and B), physical storage device 302, fibre channel arbitrated loop (FCAL) connections 303, switches 301 (switches A and B), and host 102 (host 1), at the alternate (secondary), is a “*secondary storage facility*” of Lubbers. Ex-1005, [0037], [0040], [0041].

121. The “*secondary storage facility*” of Lubbers is shown in annotated Figure 3:



Ex-1005, FIG. 3 (annotated).

122. The “secondary storage facility” of Lubbers is “for storing data that is redundant of the primary copy of data” for the same reasons discussed in element [1a]. Ex-1005, [0010], [0020], [0030], [0036], [0038], [0040].

123. The ’182 Patent explains that “discrete” storage facilities refers to the storage facilities being physically separate from one another. Ex-1001, 4:42-58 (describing the “secondary” storage facility as the “remote” storage facility, and the “primary” storage facility as the “local” storage facility), 1:32-48 (explaining the “secondary” copy of data is “usually at a remote site”). The teachings of

Lubbers would have rendered obvious “*discrete*” storage facilities as that term is used in the ’182 Patent.

124. Lubbers teaches “*the secondary storage facility discrete from the primary storage facility.*” Ex-1005, [0028] (explaining the “various sites may be physically ... separated by an arbitrary distance”), [0007] (describing a “safe physical distance between the original data and the data to [be] backed up” and that “[d]ata warehousing, ‘continuous computing’, and enterprise applications all benefit from remote copy capabilities”), [0026] (“the present invention may be implemented in a storage system that provides a peer-to-peer remote copy (backup) function”).


125. Accordingly, Lubbers teaches or suggests this element.

6. ***[1e]: “a third redundancy appliance operable in a role for storing redundant data for the sequence of write requests in storage associated with the secondary facility; and”***

126. Lubbers teaches a third controller 105 (controller 1) at an alternate site (e.g., another instance of the primary site). Ex-1005, [0037]. The third controller 105 (controller 1) is “*operable in a role for receiving a sequence of write requests*” for the same reasons discussed for element [1b].

127. Annotated Figure 3 of Lubbers shows the “*primary storage facility*” (left side) that includes a first controller (controller 1) (“*first redundancy appliance*”) and a second controller (controller 2) (“*second redundancy*”).

appliance”), and the *“secondary storage facility”* (right side) includes a third controller (controller 1) (*“third redundancy appliance”*) and fourth controller (controller 2) (*“fourth redundancy appliance”*):



Ex-1005, FIG. 3 (annotated).

128. The physical storage devices 302 (e.g., multiple individual hard disks) at the alternate site of Lubbers is *“storage associated with the secondary storage facility”* for the same reasons discussed for element [1b].

129. Additionally, the third controller (*“third redundancy appliance”*) is coupled to the physical storage device 302 (*“storage associated with the secondary storage facility”*) “for storing redundant data for the sequence of write requests”


for the same reasons discussed in element [1b]. Further, as discussed for element [1c], the controllers and other components use the same reference numerals. A POSITA would understand that in the context of Lubbers, using the same reference numerals in a patent application implies the individual components have the same structure and function as one another.

130. Accordingly, Lubbers teaches or suggests this element.

7. ***[1f]: “a fourth redundancy appliance for shadowing the third redundancy appliance wherein the fourth redundancy appliance assumes the role of the third redundancy appliance in the event of a fault at the third redundancy appliance.”***

131. Lubbers teaches a fourth controller 105 (controller 2) at the alternate site (“*secondary storage facility*”) is coupled to physical storage device 302 by an FCAL connection 303. Ex-1005, [0037]. The fourth controller is “*a fourth redundancy appliance*” for the same reasons discussed for element [1c].

132. The fourth controller 105 (controller 2) (“*fourth redundancy appliance*”) described in Lubbers is shown in annotated Figure 3 below, which uses the same reference numerals as other controllers 105. A POSITA would understand that in the context of Lubbers, using the same reference numerals in a patent application implies the individual components have the same structure and function as one another.



Ex-1005, [0040], FIG. 3 (annotated).

133. The fourth controller of Lubbers (“*fourth redundancy appliance*”) is “for shadowing” and can “assume the role of the third redundancy appliance in the event of a fault at the third redundancy appliance.” for the same reasons discussed in element [1c]. The third and fourth redundancy appliances taught by Lubbers operate in the same “shadowing” arrangement described for the first and second redundancy appliances.

134. Accordingly, Lubbers teaches or suggests this element.

B. Claim 2:

1. ***“The data redundancy system according to claim 1, wherein the secondary storage facility assumes the role of the primary storage facility in response to a failover event,***

135. Claim 2 would have been obvious based on the combined teachings of Lubbers and Sicola. A POSITA would have been motivated to combine the teachings of Lubbers and Sicola for the reasons described above in ¶¶81-94.

136. Lubbers teaches the alternate site (“*secondary storage facility*”) assumes the role of the primary site (“*primary storage facility*”) to provide system recovery after a disaster event or significant error (“*in response to a failover event*”). Ex-1005, [0007] (explaining that it “is desirable to provide the ability for rapid recovery of user data from a disaster or significant error event at a data processing facility,” termed “disaster tolerance,” which allows “providing for replicated data and redundant storage to support recovery after the event”), [0038] (explaining that in a disaster tolerant configuration, “[i]f a significant failure happens to the primary site, then data processing can be resumed at the alternate site where the data is intact”).

2. ***the failover event including a communication traffic condition having a greater portion of request traffic originating closer to the secondary storage facility than to the primary storage facility.***

137. Lubbers would have rendered obvious “*the failover event including a communication traffic condition having a greater portion of request traffic*

originating closer to the secondary storage facility than to the primary storage facility” for several reasons.

138. First, as discussed for element [1d], Lubbers teaches “*the secondary storage facility discrete from the primary storage facility.*”

139. Second, Lubbers explains the benefits of load balancing and added complexities of increasing transmission distances. Ex-1005, [0006] (describing how prior storage systems “provided minimal load balancing benefit that would enable data access requests to be directed intelligently to replicas such that resources were used more efficiently”), [0041] (explaining that longer fibre channels increases “both the latency and the variability in latency”), [0042] (describing that “[h]osts 102 preferably run multi-pathing software that dynamically allows failover between storage paths as well as static load balancing of storage volumes (LUNs) between the path of controllers 102,” which “enables a host 102 to identify multiple connection options that lead to desired storage, and select amongst the available paths based on selected criteria” such as “path availability, path performance, path load, controller load, and the like.”), [0064] (explaining that “the present invention allows great flexibility in implementing a particular system to emphasize any balance of, for example, low latency, high availability, disaster tolerance, and cost”).

140. Thus, Lubbers teaches criteria like latency, path performance, and path load in balancing loads between storage volumes at different sites. Based on these criteria, a POSITA would have found it obvious to use a failover condition in which the primary site and alternate site change roles in response to the secondary site being located closer to a host computer 102 generating write and read requests. A POSITA would have been motivated to consider this type of failover condition to minimize latency and improve data transmission path performance, as taught by Lubbers.

141. A POSITA would have reasonably expected to succeed in using this type of failover condition for the multi-site storage area network (SAN) of Lubbers because Lubbers already contemplates flexible assignment of sites as primary and secondary, Ex-1005, [0036], contemplates site failover situations, such as during failure in a communication link between controllers and where storage cells are “inaccessible” to sites, Ex-1005, [0068], and further teaches site failover can occur automatically, Ex-1005, [0022]. Thus, using this type of failover condition for the SAN of Lubbers would have been a routine matter well within a POSITA’s ability to achieve predictable results.

142. Accordingly, Lubbers and Sicola would have rendered obvious claim 2.

- C. ***Claim 3: “The data redundancy system according to claim 1, wherein the second redundancy appliance assumes the role of the first redundancy appliance in response to a failover event selected from the group consisting of a manually initiated event and an automatically initiated event.”***

143. Claim 3 would have been obvious based on the combined teachings of Lubbers and Sicola. A POSITA would have been motivated to combine the teachings of Lubbers and Sicola for the reasons described above in ¶¶81-94.

144. As discussed for element [1c], Lubbers teaches “*wherein the second redundancy appliance assumes the role of the first redundancy appliance.*”

145. Lubbers also teaches that the second controller (“*second redundancy appliance*”) assumes the role of the first controller (“*first redundancy appliance*”) automatically upon failover (“*in response to a failover event selected from the group consisting of a manually initiated event and automatically initiated event*”). Ex-1005, [0046] (explaining that when “a controller fails, the other controller presents the LUNs that were being presented on the failed controller,” and can “continue data replication operations without loss of data or requiring manual process.”).

146. Accordingly, Lubbers and Sicola would have rendered obvious claim 3.

D. Claim 4:

1. *“The data redundancy system according to claim 1, wherein the first redundancy appliance updates the second redundancy appliance with information about the data in storage by forwarding a write record to the second redundancy appliance for each write request,”*

147. Claim 4 would have been obvious based on the combined teachings of Lubbers and Sicola. A POSITA would have been motivated to combine the teachings of Lubbers and Sicola for the reasons described above in ¶¶81-94.

148. Lubbers would have rendered obvious *“wherein the first redundancy appliance updates the second redundancy appliance with information about the data in storage by forwarding a write record to the second redundancy appliance for each write request.”*

149. First, Lubbers teaches that the first controller (*“first redundancy appliance”*) updates the second controller (*“second redundancy appliance”*) *“with information about the data in storage”* by moving data between physical storage devices. Ex-1005, [0032] (explaining “[s]torage controllers 105 manage tasks of ... moving data between physical storage devices ...”), [0073] (describing that if a “controller does not have the most current data, then the controller goes offline to the host to receive a full copy from the controller that assumes the role of the source controller (step 830)”).


150. Second, Sicola teaches that the “*updates*” occur “*by forwarding a write record to the second redundancy appliance for each write request*” for the same reasons discussed above in element [1c]. Ex-1006, describing asynchronous operation involving a “micro-logging” process in which the controller “writes the data into its write-back cache”), 13:6-13 (explaining a “micro-merge” process during asynchronous operation in which “outstanding writes were logged to the initiator controller A1’s write-back cache, which is also mirrored in partner controller A2’s (mirrored) write-back cache, so that the cache data is available to controller A2 if controller A1 fails.”).

2. “*the second redundancy appliance sending acknowledgements to the first redundancy appliance in batches in response to receiving write records.*”

151. Lubbers also would have rendered obvious “*the second redundancy appliance sending acknowledgements to the first redundancy appliance in batches in response to receiving write records.*” Ex-1005, [0068] (explaining that each “controller periodically sends ‘heartbeats,’ to and receives ‘heartbeats’ from peer controllers,” which are pings “with a message including information useful in managing data transfer processes between peer controllers”), [0069] (describing “a full-duplex path between one controller and another is referred to as a ‘tunnel’”), [0072] (explaining a “controller determines whether the data under the supervision of the peer controllers is consistent ... by comparing GSNs of the various

controllers”), claim 2 (explaining a “step of sending the host computer a write completion status message prior to sending the data to the destination storage system), claim 11 (same).

152. These pings and responsive information are *acknowledgements* ... “in response to receiving write records,” as further described and shown in Figure 6B of Sicola:



Ex-1006, FIG 6B, 10:22-26 (explaining intra-site controller heartbeat timer operation), 10:45-53 (describing that during normal operation, at step 640, companion controllers “periodically send pings to each other via DUARTs (Dual

Ex-1003

Asynchronous Receiver/Transmitters),” and the pings are received unless either controller has timed out).

153. Thus, Lubbers teaches component redundancies for multiple levels of fault tolerance, allowing failover to a redundant component at the same site for continued operation, storing write commands in controller write-back caches, and using criteria to balance controller load.

154. Based on these teachings, a POSITA would have found it obvious to use the first controller (“*the first redundancy appliance*”) to update a second (redundant) controller (“*the second redundancy appliance*”) with information about the data in physical storage device 302 (“*storage*”) by forwarding a write request or log (“*write record*”) to the second controller (“*second redundancy appliance*”) for each write request, where the second controller (“*second redundancy appliance*”) sends pings (“*acknowledgements*”) to the first controller (“*first redundancy appliance*”) in batches in response to receiving write records.”

155. Further, a POSITA would have been motivated to consider forwarding write records and sending return acknowledgements in this manner to maintain integrity and availability of data, while also minimizing controller load, as taught by Lubbers.

156. Accordingly, Lubbers and Sicola would have rendered obvious claim 4.

- E. ***Claim 5: “The data redundancy system according to claim 4, wherein the first redundancy appliance sends an acknowledgement to an originator of the write request in response to receiving the acknowledgement from the second redundancy appliance.”***

157. Claim 5 would have been obvious based on the combined teachings of Lubbers and Sicola. A POSITA would have been motivated to combine the teachings of Lubbers and Sicola for the reasons described above in ¶¶81-94.

158. Lubbers would have rendered obvious “*wherein the first redundancy appliance sends an acknowledgement to an originator of the write request in response to receiving the acknowledgement from the second redundancy appliance*” by teaching controllers communicating through and responding to pings and sending write completion status messages to the host computer. Ex-1005, claim 2 (explaining a “step of sending the host computer a write completion status message prior to sending the data to the destination storage system”), claim 11 (same), [0068] (explaining that each “controller periodically sends ‘heartbeats,’ to and receives ‘heartbeats’ from peer controllers,” which are pings “with a message including information useful in managing data transfer processes between peer controllers”).

159. Sicola also teaches that when the host write request is complete, “the host is given the completion status.” Ex-1005, 12:44-46. Given the teachings of Lubbers and Sicola, a POSITA would have recognized the pings could contain the

write completion status information to synchronize the controllers (for redundancy) before relaying the write completion status to the host.


160. Accordingly, Lubbers and Sicola would have rendered obvious claim 5.

F. *Claim 6: “The data redundancy system according to claim 1, wherein the first redundancy appliance is disjoint from the second redundancy appliance.”*

161. Claim 6 would have been obvious based on the combined teachings of Lubbers and Sicola. A POSITA would have been motivated to combine the teachings of Lubbers and Sicola for the reasons described above in ¶¶81-94.

162. The '182 Patent refers to a redundancy appliance as “*disjoint*” from another redundancy appliance where “it preferably does not share common hardware or software elements, or common failure modes” (Ex-1001, 21:23-26). The combined teachings of Lubbers and Sicola would have rendered obvious “*disjoint*” redundancy appliances as that term is used in the '182 Patent.

163. The first controller (“*first redundancy appliance*”) is “*disjoint*” from the second controller (“*second redundancy appliance*”) of Lubbers because they do not share common hardware or software elements:



Ex-1005, FIG. 3 (annotated), [0037] (explaining “[t]he configuration of FIG. 3 emphasizes redundancy in physical components”), [0038] (describing a “goal of a disaster tolerant configuration in FIG. 3 is to provide hardware redundancy for multiple levels of fault tolerance ... to keep the primary and alternate sites running, almost at ‘non-stop’ metrics”), [0026] (explaining controller-based replication “between two remotely separated pairs of array controllers”), [0039]-[0040] (describing that redundant FCALs allows for failure of one FCAL “while each storage device 302 remains accessible by either controller”)

164. Accordingly, Lubbers and Sicola would have rendered obvious claim 6.

- G. Claim 7: “The data redundancy system according to claim 1, wherein the second redundancy appliance maintains a shadow copy of data stored by the first redundancy appliance wherein the second redundancy appliance accesses the shadow copy of the data when the second redundancy appliance assumes the role of the first redundancy appliance.”**

165. Claim 7 would have been obvious based on the combined teachings of Lubbers and Sicola. A POSITA would have been motivated to combine the teachings of Lubbers and Sicola for the reasons described above in ¶¶81-94.

166. The second controller (“*second redundancy appliance*”) is “*for shadowing*,” assuming the role of, and independently performing the same functions as, the first controller (“*first redundancy appliance*”) for the reasons discussed in element [1c].

167. Lubbers teaches that the second controller (“*second redundancy appliance*”) “*maintains a shadow copy of data stored by the first redundancy appliance*” and “*accesses the shadow copy of the data*” when the second controller assumes the role of the first controller. Operating independently includes the second controller (“*second redundancy appliance*”) using the “shadow copy” of data from the first controller (“*first redundancy appliance*”), particularly during a failover event where role assumption becomes necessary. Further, Sicola teaches that both controllers in a redundant controller pair can store their own copy of data in their own separate write-back cache for the same reasons discussed above in

element [1c]. Ex-1006, 12:16-30 (describing asynchronous operation involving a “micro-logging” process in which the controller “writes the data into its write-back cache”), 13:6-13 (explaining a “micro-merge” process during asynchronous operation in which “outstanding writes were logged to the initiator controller A1’s write-back cache, which is also mirrored in partner controller A2’s (mirrored) write-back cache, so that the cache data is available to controller A2 if controller A1 fails.”).

168. Thus, Lubbers teaches component redundancies for multiple levels of fault tolerance, allowing failover to a redundant component at the same site for continued operation, storing write commands in controller write-back caches, and using criteria to balance controller load. A POSITA would have found it obvious to use a second (redundant) controller to maintain a copy of the write-back cache data stored by the first controller (“*wherein the **second redundancy appliance** maintains a shadow copy of data stored by the **first redundancy appliance***”) and to use the second controller to access that write-back cache data when the second controller assumes the role of the first controller (e.g., due to failover) (“*wherein the **second redundancy appliance** accesses the shadow copy of the data when the **second redundancy appliance** assumes the role of the **first redundancy appliance***”).

Ex-1008, 711. A POSITA would have been motivated to use a second controller to maintain a copy of the write-back cache data in this manner to facilitate

asynchronous operation and maintain data integrity when the first controller fails in the window where the host write status has already returned to the host, but remote operation has not yet completed, as taught by Sicola. Ex-1008, 711.


169. Accordingly, Lubbers and Sicola would have rendered obvious claim 7.

H. ***Claim 8: “The data redundancy system according to claim 1, wherein the second redundancy appliance accesses the data stored in storage by the first redundancy appliance when the second redundancy appliance acts as the first redundancy appliance in response to a failover event.”***

170. Claim 8 would have been obvious based on the combined teachings of Lubbers and Sicola. A POSITA would have been motivated to combine the teachings of Lubbers and Sicola for the reasons described above in ¶¶81-94.

171. The second controller (“*second redundancy appliance*”) of Lubbers assumes the role of, and can independently perform the same functions as, the first controller (“*first redundancy appliance*”) in the event of a first controller failure (“*in response to a failover event*”), for the reasons discussed in element [1c].

172. Lubbers also teaches the first controller (“*first redundancy appliance*”) and second controller (“*second redundancy appliance*”) can each store data in physical storage devices 302 (“*storage associated with the primary storage facility*”) for the reasons discussed in elements [1b] and [1c], the configuration of which is shown in annotated Figure 3 of Lubbers:



Ex-1005, FIG. 3 (annotated). Lubbers teaches that the second controller (*“second redundancy appliance”*) can access data in the same physical storage device 302 (*“storage associated with the primary storage facility”*) written to by the first controller (*“first redundancy appliance”*), particularly when the controllers swap roles. Ex-1005, [0040] (describing access to the individual hard disk drives “by either controller 105”), [0038], [0046] (describing that “[w]hen a controller fails, the other controller presents the LUNs that were being presented on the failed controller”), [0020].

173. Accordingly, Lubbers and Sicola would have rendered obvious claim 8.

X. DETAILED ANALYSIS THAT CLAIMS 9-18 ARE OBVIOUS OVER LUBBERS AND SICOLA IN VIEW OF OFEK

174. As discussed above, Lubbers and Sicola would have rendered obvious every element of claims 1-8. Lubbers, Sicola, and Ofek would have rendered obvious every element of claims 9-18, as explained below.

175. It would have been obvious to combine Ofek's teachings to bundle write data for transmission to the secondary data storage system with the teachings of Lubbers and Sicola ("*redundant data to the secondary storage facility as send batches*"). Lubbers teaches "a data logging mechanism (e.g., a write history log unit) for storing commands and associated data for I/O transactions that occur when the remote backup storage device is unavailable." Ex-1005, [0026] (explaining that when "the destination site becomes available, the system performs an in-order merging of the log unit data with the data on the previously unavailable backup device to quickly return the source and destination sites to the same data state after link restoration or destination site restoration"), [0060] (describing "I/O commands and the data associated with those commands are written to the data log, and after destination becomes available the information in the data log is merged into the destination volume to conform the data in the destination volume to the data in the source volume"), [0075] (describing how the "source controller may begin logging"), [0076] (explaining that during "the data logging process, I/O

operations from the host computer are copied into the log unit's cache buffers"). In other words, Lubbers already teaches storing multiple transactions ("*batches*") of transactions in a log, which is then sent and merged with the data in the alternate site.

176. Similarly, Ofek teaches that "[i]n order to more effectively use the capacity of the data link, the data storage system containing the primary (R1) volume bundles the write data for all write commands in the chain into a single write command for transmission over the data link to the secondary data storage system containing the secondary (R2) volume." Ex-1007; *see also* Ex-1007, 34:66-35:6 (it is very desirable to 'bundle' the write data for all write commands in the channel command word chain into a single write command transmitted over a link to the secondary storage system 246.").

177. A POSITA would have been motivated to combine the teachings of Lubbers and Sicola and analogous write bundling taught by Ofek for several reasons. First, Lubbers and Ofek teach effectively using bandwidth and minimizing latency. Ex-1005, [0041] (describing how "both latency and the variability in latency increases with more complex fabrics"), [0042] (explaining load balancing based on criteria such as "path availability, path performance, path load, controller load, and the like"), [0042]; Ex-1007, Abstract (explaining "[i]n order to more effectively use the capacity of the data link, the data storage system containing the

primary (R1) volume bundles the write data for all write commands in the chain into a single write command for transmission over the data link to the secondary data storage system containing the secondary (R2) volume”).

178. Second, both Lubbers and Ofek teach a goal of maintaining data integrity. Ex-1005, [0032] (explaining “[s]torage controllers 105 manage the tasks of allocating physical capacity to specified LUNs, moving data between physical storage devices, and other functions that maintain integrity and availability of the data stored therein”); Ex-1007, 3:10-18 (explaining the data storage system “achieves nearly 100 percent data integrity by assuring that all data is copied to a geographically remote site, and in those cases when a back-up copy is not made due to an error of any sort, an indication is stored that the data has not been copied, but instead must be updated at a future time.”), 35:60-36:2 (describing that “the entry in the link queue 504 includes a one-byte password for confirming the validity of the entry”), 37:17-24 (explaining “[e]ach link adapter scans the link queue 504 in an iterative loop, looking for unlocked entries to service, beginning at the head of the queue,” and then “locks the next entry to service, checks the password to determine if the entry is valid, and if so, gets the buffer pointer from the entry, reads the buffer, and builds a job to be executed for transferring data from cache across the link in a direct memory access (DMA) operation”). A POSITA would have been motivated to combine the teachings of Ofek and

Lubbers and Sicola to include data integrity assurances like object checking and verification for maintaining accurate data in a backup system.

179. A POSITA would recognize that bundling writes into a single write transaction for transmission, as taught by Ofek, would also use a cache memory, but would merely involve a buffer rather than a transaction log, which Lubbers and Sicola also already teach. Ex-1005, [0032] (explaining “how hosts 102 access physical storage capacity by addressing read and write operations to specified LUNs” (logical units) using the worldwide LUN ID (WWLID)).

180. Additionally, Lubbers, Sicola, and Ofek are analogous art to the ’182 Patent for at least two reasons. First, Lubbers, Sicola, and Ofek are directed toward the same field of endeavor of the ’182 Patent—techniques for redundant data storage. Ex-1001, 1:29-31; Ex-1005, [0001], [0007]; Ex-1006, 1:8-14; Ex-1007, 1:8-11.

181. Second, POSITAs knew that communication delays and component failures were a common data storage problem as of June 6, 2003. For example, Lubbers, Sicola, and Ofek each are reasonably pertinent to that problem and provide techniques for redundant data storage to overcome that problem. Ex-1005, [0038]; Ex-1006, 3:63-4:6; Ex-1007, 3:3-22. The ’182 Patent also confirms that communication delays and component failures were a problem in data storage as of June 6, 2003. Ex-1001, 1:45-51.

182. There are a variety of ways in which a POSITA would have been motivated to combine the teachings of Lubbers and Sicola with Ofek. For example, a POSITA would have been motivated to combine batching and sending write commands, as taught by Ofek, with the teachings of Lubbers and Sicola to minimize the effects of latency and improve data transmission path performance for backup data systems. A POSITA would have reasonably expected to succeed in having a first controller batch and send write commands in this manner because Lubbers and Sicola already contemplate including multiple write requests in logs before transmitting them to be merged into the alternate site (Ex-1005, [0060], Ex-1006, 10:22-53), and Ofek contemplates using chains of channel commands bundling multiple write commands (Ex-1007, Abstract, 34:66–35:6). Thus, sending multiple write commands in a batch to an alternate site would have been a routine matter well within a POSITA’s ability to achieve predictable results.

183. It also would have been obvious to bundle write data for write commands (“*batches*”) based on quantity of write commands, taught by Lubbers and Sicola, or based on time elapsed between batches, taught by Ofek because doing so would similarly achieve the goals of minimizing latency and improving data transmission and path performance. Further, for at least these same reasons, a POSITA would have found it obvious to adaptively size bundles of write data for write commands (“*batches*”) based on availability of communication bandwidth.

184. A POSITA would have been motivated to combine teachings for adaptively assessing batch size to balance the effects of latency, high availability, disaster tolerance, data transmission path performance, and system performance, speed, and cost. A POSITA would have reasonably expected to succeed in having a first controller determine the size of send batches in this manner because Lubbers and Sicola already contemplate different modes of operation to focus on certain of the competing considerations, and Ofek likewise includes a user-configurable skew parameter. Ex-1007, 17:39-41.

185. A POSITA would have found it obvious to replace a first write transaction in a bundle of write data (“*send batch*”) at a primary site (“*primary storage facility*”) by overwriting it with a second write transaction in the bundle of write data. Lubbers teaches ordering and sequencing write commands such that a write transaction can be overwritten. Ex-1005, [0060] (describing an ordering algorithm that “uses a ‘group sequence number’ and the remote groups 705 ensure that the data is written in order sequence”), [0072] (explaining “the controller determines whether the data under the supervision of the peer controllers is consistent” by “comparing GSNs of the various controllers” and that if “the GSNs differ, then the data states of various controllers differ”). Further, Lubbers describes a circumstance in which overwriting is beneficial. Ex-1005, [0072]-[0073] (if “the controller does not have the most current data, then the controller

goes offline to the host to receive a full copy from the controller that assumes the role of the source controller.”).

186. Sicola provides further details on overwriting write transactions in a batch in a multi-site storage system. Ex-1006, 12:16-30 (describing asynchronous operation involving a “micro-logging” process in which the controller “writes the data into its write-back cache”), 13:6-13 (explaining a “micro-merge” process during asynchronous operation in which “outstanding writes were logged to the initiator controller A1’s write-back cache, which is also mirrored in partner controller A2’s (mirrored) write-back cache, so that the cache data is available to controller A2 if controller A1 fails.”).

187. Ofek teaches a more specific overwriting procedure. Ex-1007, 39:28-31 (the channel adapter first checks whether the track is on disk, which is then copied from disk to the cache slot in cache); Ex-1007, 39:41-43 (if the track does reside in cache, then an “overwrite cache option” for the volume can be executed); Ex-1007, 39:43-52 (“if the cache contains a remote write pending record that has not yet been transmitted by a link adapter to the data storage system having the secondary volume, then a new version will overwrite this write pending in cache.”). Further, Ofek teaches that “[i]n the adaptive copy modes, however, the overwrite cache option substantially increases the performance of remote

mirroring of data, by reducing the number of remote pending records that are maintained in cache or on disk.” Ex-1007, 39:52-56.

188. A POSITA would have been motivated to combine the teachings of Lubbers, Sicola, and Ofek to enable overwriting write transactions in a cache before or after transmitting and storing the transactions to improve the performance of remote mirroring data, by reducing the number of remote pending records that are maintained in cache or on disk.

189. A POSITA would have reasonably expected to succeed in overwriting write transactions within a send batch or in combinations of send batches at either the primary site or secondary sites because Lubbers and Sicola already contemplate replacing controller data with the most current data, Ex-1005, [0073]; Ex-1006, 12:16–30, 13:6–13, and Ofek contemplates overwriting transactions within bundles of write transactions in a cache. Ex-1007, 37:64–67, 39:22–56. Thus, overwriting write transactions within a send batch at the primary site would have been a routine matter well within a POSITA’s ability to achieve predictable results. Additionally, a POSITA would have reasonably expected to succeed in combining receive batches and having write transactions overwrite other write transactions in a receive batch at the alternate site because Lubbers and Sicola already contemplate replacing controller data with the most current data, Ex-1005, [0073]; Ex-1006,

12:16–30, 13:6–13, and Ofek contemplates overwriting transactions within bundles of write transactions in a cache, Ex-1007, 35:13–24, 36:13–16, 40:21–28.

190. For the reasons explained above, the same technique for bundling writes to create “*batches*” and overwriting transactions would apply equally to batches being sent (“*send batches*”) from the primary site (“*primary storage facility*”) and being received (“*receive batches*”) at the alternate site (“*secondary storage facility*”). See e.g., Ex-1005, [0026]; Ex-1007, 35:13–16 (describing write and reads commands for the secondary volume (R2)) are “transmitted over the link 240”), 37:33–37:35 (explaining the channel adapter uses control logic to bundle remote write commands in a command word chain).

191. A POSITA would have reasonably expected to succeed in having the third controller (“*third redundancy appliance*”) determine data is safe to apply before writing it to storage at the alternate site (“*secondary storage facility*”) because Lubbers already teaches the importance of data reliability and integrity, and using LUNs, UUIDs, Ex-1005, [0032], and Ofek further teaches using a one-byte password (among other things) to confirm the validity of write requests before writing them to storage, Ex-1007, 35:60–36:2. Thus, determining write requests in a receive batch are safe to apply before writing those requests to storage at the alternate site (“*secondary storage facility*”) would have been a routine matter well within a POSITA’s ability to achieve predictable results.

A. Claim 9: “The data redundancy system according to claim 1, wherein the first redundancy appliance forwards redundant data to the secondary storage facility as send batches.”

192. Claim 9 would have been obvious based on the combined teachings of Lubbers, Sicola, and Ofek. A POSITA would have been motivated to combine the teachings of Lubbers, Sicola, and Ofek for the reasons described above in ¶¶174-191.

193. Lubbers teaches the first controller (“*first redundancy appliance*”) accesses and transfers a replica of the data (“*redundant data*”) to the alternate site (“*secondary storage facility*”) for the reasons discussed in element [1d].

194. It would have been obvious to combine Ofek’s teachings for bundling write commands and forwarding “*redundant data to the secondary storage facility as send batches*” with the teachings of Lubbers and Sicola.

195. The Lubbers system already discloses batching write requests in the form of transaction logs. Ex-1005, [0060] (explaining “the log is sized to store all write transactions until the transaction is committed to each member of a copy set” and describing that “I/O commands and the data associated with those commands are written to the data log, and after destination becomes available the information in the data log is merged into the destination volume to conform the data in the destination volume to the data in the source volume”).

196. A POSITA would recognize that bundling write requests for transmission would similarly involve using a cache memory, but would merely involve gathering the data in a buffer rather than a transaction log, which Lubbers also already teaches.

197. Like Lubbers, Ofek teaches “*send batches*.” Ex-1007, Abstract (explaining that “to more effectively use the capacity of the data link, the data storage system containing the primary (R1) volume bundles the write data for all write commands in the chain into a single write command for transmission over the data link to the secondary data storage system containing the secondary (R2) volume”), 34:66-35:6 (describing it is “desirable to ‘bundle’ the write data for all write commands in the channel command word chain into a single write command transmitted over a link to the secondary storage system 246”).

198. A POSITA would have been motivated to combine the teachings of Lubbers and Ofek based on balancing considerations like latency, path performance, and load capacity, taught by Lubbers.

199. A POSITA would have reasonably expected to succeed in combining the teachings of Lubbers and Ofek in this manner because Lubbers already teaches sending batched transactions via transaction logs. See ¶¶174-191.

200. Accordingly, the combined teachings of Lubbers, Sicola, and Ofek would have rendered obvious claim 9.

B. *Claim 10: “The data redundancy system according to claim 9, wherein a send batch size is determined according to either a time interval for collecting transactions into the send batch or a predetermined quantity of transactions to be included in the send batch.”*

201. Claim 10 would have been obvious based on the combined teachings of Lubbers, Sicola, and Ofek. A POSITA would have been motivated to combine the teachings of Lubbers, Sicola, and Ofek for the reasons described above in ¶¶174-191.

202. Lubbers teaches using “a time delay” or “a predetermined quantity of transactions” to determine “send batch size.” Ex-1005, [0011] (describing that “[d]uring the merge process the system maintains a ratio of merge data writes out of the data log to data writes from a host computer into the data log within a desired range until the data log reaches a predetermined size, [0059] (describing an “asynchronous ... operation mode in which a modification to one member of a copy set can be propagated to other members of the copy set after a time delay”), [0081].

203. Thus, based on the teachings of Lubbers, a POSITA would have reasonably expected to succeed in combining the teachings of Lubbers, Sicola, and Ofek to achieve predictable results.

204. Accordingly, the combined teachings of Lubbers, Sicola, and Ofek would have rendered obvious claim 10.

C. Claim 11: “The data redundancy system according to claim 9, wherein a send batch size is determined adaptively based on available communication bandwidth between the primary storage facility and the secondary storage facility.”

205. Claim 11 would have been obvious based on the combined teachings of Lubbers, Sicola, and Ofek. A POSITA would have been motivated to combine the teachings of Lubbers, Sicola, and Ofek for the reasons described above in ¶¶174-191.

206. It would have been obvious to include Ofek’s teachings to “*adaptively*” determine a “*send batch size*” based on “*available communication bandwidth*” with the teachings of Lubbers and Sicola.

207. Lubbers already teaches different modes of operations and balancing competing criteria for system performance. Ex-1005, [0042] (explaining “[m]ulti-pathing software enables a host 102 to identify multiple connection options that lead to desired storage and select amongst the available paths based on selected criteria,” including “path availability, path performance, path load, controller load, and the like”), [0059] (describing an “asynchronous ... operation mode in which a modification to one member of a copy set can be propagated to other members of the copy set after a time delay”).

208. Likewise, Ofek teaches balancing similar operational requirements, and adjusting certain parameters like “*send batch size*” based on “*available*

communication bandwidth” to achieve desired system performance. Ex-1007, 7:10-21 (explaining the data storage system can operate in a “point-in-time or asynchronous mode” to eliminate “any performance penalty if the communication link between the primary and secondary data storage systems is too slow), 12:22-24 (explaining “operational modes are selectable at the logical volume level based on the performance, distance, and speed of recovery requirements”), 17:39-42, 18:9-13 (explaining the adaptive copy—write pending mode for specified volumes “is convenient in situations where the write activity is caused by heavy batch loads or data reorganization can severely impact performance due to the data storage systems maintaining a full synchronous state”).

209. A POSITA would have been motivated to adaptively determining a “*send batch size*” in the combined system of Lubbers, Sicola, and Ofek in this manner to balance latency, high availability, disaster tolerance, data transmission path performance, and system performance, speed, and cost.

210. A POSITA would have recognized that quantity of data that could be reliably transmitted from the primary storage facility to the secondary storage facility would likely depend on the available bandwidth of the communication link between them, particularly given the teachings of Lubbers, Sicola, and Ofek discussed above. ¶¶174-191, 205-209. Thus, a POSITA would have reasonably expected to succeed in adaptively sizing send batches in this manner because

Lubbers and Ofek already teach adjusting system operation based on balancing these competing objectives. Ex-1007, 17:39-41.

211. Accordingly, the combined teachings of Lubbers, Sicola, and Ofek would have rendered obvious claim 11.

D. Claim 12: “The data redundancy system according to claim 1, wherein a first write transaction in a send batch at the primary storage facility is replaced with a second write transaction where the second write transaction overwrites data written by the first write transaction.”

212. Claim 12 would have been obvious based on the combined teachings of Lubbers, Sicola, and Ofek. A POSITA would have been motivated to combine the teachings of Lubbers, Sicola, and Ofek for the reasons described above in ¶¶174-191.

213. It would have been obvious to include Ofek’s teachings to “*overwrite*” a “*first write transaction in a send batch*” with “*a second write transaction*” within the system of Lubbers and Sicola.

214. Lubbers contemplates a circumstance in which an initial write transaction needs to be replaced with a second write transaction. Ex-1005, [0072]-[0073] (explaining that when “the controller does not have the most current data, then the controller goes offline to the host to receive a full copy from the controller that assumes the role of the source controller”).

215. Sicola provides further details on overwriting write transactions in a batch in a multi-site storage system. Ex-1006, 13:6-15 (explaining that asynchronous system operation involves a micro-logging and micro-merge process in which “‘outstanding’ writes were logged to the initiator controller A1’s write-back cache, which is also mirrored in partner controller A2’s write-back cache,” and that during failover, “the partner controller (A2) re-issues these remote copies from the micro-log.”).

216. Ofek teaches the use of a channel adapter that writes one or more records into a cache that can be “*overwritten*” with new writes that in operation, “the channel adapter writes one or more records into cache.” Ex-1007, 37:64-67, 39:22-24, 39:41-43 (explaining that if a “track does reside in cache,” then the “channel adapter is working in an ‘overwrite cache option’ for the volume,” in which “a new version will overwrite this write pending record in the cache”), 39:52-56.

217. A POSITA would have been motivated to replace a first write transaction in a bundle of write data (“*send batch*”) at a primary site (“*primary storage facility*”) by overwriting it with a second write transaction in the bundle in the combined system of Lubbers, Sicola, and Ofek in this manner to balance latency, high availability, disaster tolerance, data transmission path performance, and system performance, speed, and cost.


218. A POSITA would have reasonably expected to succeed in adaptively sizing send batches in this manner because Lubbers and Ofek already teach overwriting transactions stored in caches. Ex-1007, 17:39-41.

219. Accordingly, the combined teachings of Lubbers, Sicola, and Ofek would have rendered obvious claim 12.

E. Claim 13: “The data redundancy system according to claim 1, wherein the third redundancy appliance stores redundant data to the storage associated with the secondary facility according to receive batches.”

220. Claim 13 would have been obvious based on the combined teachings of Lubbers, Sicola, and Ofek. A POSITA would have been motivated to combine the teachings of Lubbers, Sicola, and Ofek for the reasons described above in ¶¶174-191.

221. Lubbers, Sicola, and Ofek teach “*the third redundancy appliance stores redundant data to the storage associated with the secondary facility*” for the reasons discussed above in element [1e], which is shown in annotated Figure 3 below:



Ex-1005, FIG. 3 (annotated).

222. Lubbers, Sicola, and Ofek also teach “receive batches,” and storing the write transactions “according to receive batches” for the same reasons discussed in claim 9 that the combined Lubbers, Sicola, and Ofek system teaches “send batches”). Ex-1007, 35:13-16, 37:33-37:35.

223. A POSITA would have been motivated to combine the teachings of Ofek to receive and decode “receive batches” with the teachings of Lubbers, Sicola, and Ofek in this manner to balance the effects of latency, high availability, disaster tolerance, data transmission path performance, and system performance, speed, and cost.

224. A POSITA would have reasonably expected to succeed in combining the teachings of Lubbers, Sicola, and Ofek such that the third controller (*“third redundancy appliance”*), at the alternate site (*“secondary storage facility”*), writes redundant data (*“stores redundant data”*) to the physical storage device 302 (*“the storage associated with the secondary facility”*) after decoding bundles of write commands in a channel command word chain (*“receive batches”*) to efficiently use resources during write processes.

225. Accordingly, the combined teachings of Lubbers, Sicola, and Ofek would have rendered obvious claim 13.

F. *Claim 14: “The data redundancy system according to claim 13, wherein the third redundancy appliance stores a receive batch after the receive batch is determined safe to apply.”*

226. Claim 14 would have been obvious based on the combined teachings of Lubbers, Sicola, and Ofek. A POSITA would have been motivated to combine the teachings of Lubbers, Sicola, and Ofek for the reasons described above in ¶¶174-191.

227. It would have been obvious to apply Ofek’s teachings to store *“a receive batch after the receive batch is determined safe to apply”* and the teachings of Lubbers and Sicola because a stated goal and functions already described in Lubbers and Sicola are for maintaining data integrity and availability. Ex-1005, [0032] (explaining “[s]torage controllers 105 manage the tasks of allocating

physical capacity to specified LUNs, moving data between physical storage devices, and other functions that maintain integrity and availability of the data stored therein”) Ex-1005, [0032].

228. Ofek also teaches the importance of ensuring data integrity before storing write requests, shown in annotated Figure 18:




FIG. 18

Ex-1007, FIG. 18 (annotated), 35:18-20, 35:60-36:2 (describing “the entry in the link queue 504 includes a one-byte password for confirming the validity of the entry”), 3:10-18 (explaining the data storage system “achieves nearly 100 percent data integrity by assuring that all data is copied to a geographically remote site, and in those cases when a back-up copy is not made due to an error of any sort, an

indication is stored that the data has not been copied, but instead must be updated at a future time.”).

229. A POSITA would have been motivated to store “*a receive batch after the receive batch is determined safe to apply*” in the combined system of Lubbers, Sicola, and Ofek in this manner to improve data availability and reliability, which is a shared goal.

230. A POSITA would have reasonably expected to succeed in determining whether write commands are safe to apply to a standby system in this manner as merely combining the system software teachings of Lubbers and Sicola with the teachings of Ofek to achieve a mutual goal. Ex-1007, 17:39-41.

231. Accordingly, the combined teachings of Lubbers, Sicola, and Ofek would have rendered obvious claim 14.

G. *Claim 15: “The data redundancy system according to claim 14, wherein a receive batch is determined safe to apply when a data integrity test is successfully performed.”*

232. Claim 15 would have been obvious based on Lubbers, Sicola, and Ofek for the same reasons discussed above in claim 14. A POSITA would have been motivated to combine the teachings of Lubbers, Sicola, and Ofek for the reasons described above in ¶¶174-191.

233. As stated in claim 14, Ofek expressly teaches a one-byte password for confirming the validity of the entry to verify data integrity, and how the system

achieves nearly 100% data integrity by ensuring no errant data is stored in the backup system. Ex-1007, 3:10-18, 35:60-36:2.

234. Further, a POSITA would have been motivated to combine Lubbers, Sicola, and Ofek for the reasons above.

235. Accordingly, the combined teachings of Lubbers, Sicola, and Ofek would have rendered obvious claim 15.

H. Claim 16: “The data redundancy system according to claim 13, wherein a first write transaction in a first receive batch at the secondary storage facility is replaced with a second write transaction in a second receive batch wherein the second write transaction overwrites data written by the first write transaction and the first and second receive batches are combined.”

236. Claim 16 would have been obvious based on the combined teachings of Lubbers, Sicola, and Ofek. A POSITA would have been motivated to combine the teachings of Lubbers, Sicola, and Ofek for the reasons described above in ¶¶174-191.

237. Lubbers, Sicola, and Ofek would have rendered obvious “receive batches” for the reasons discussed in claim 13. Additionally, Lubbers, Sicola, and Ofek would have rendered obvious a “second write transaction overwrites data written by a first write transaction” for the reasons discussed in claim 12.

238. It would have been obvious to combine the teachings of Ofek to “a first write transaction in a first receive batch” and “a second write transaction in a

second receive batch” when “the first and second receive batches are combined” and the teachings of Lubbers and Sicola. Ex-1007, 36:13-16 (explaining “it is very desirable to ‘bundle’ the write data ... into a single write command transmitted over a link to the secondary write system 246,” which is then decoded), 35:13-24 (describing write and reads commands for the secondary volume (R2)) are “transmitted over the link 240” in which “[e]ach entry is placed in a link queue 504 by a channel adapter”).

239. Thus, Ofek teaches a first write bundle (“first receive batch”) and a second write bundle (“second receive batch”) that are sequentially sent to the secondary storage volume (R2) (“secondary storage facility”).

240. Additionally, Ofek teaches write bundles can be combined, and transactions can be overwritten while in cache. Ex-1007, 40:21-28 (explaining that “keeping multiple versions of tracks in cache will tie up cache resources, so ... no more than a certain number of such old versions should be permitted in cache at any given time”) which also applies to the bundled writes received at the secondary storage volume. Ex-1007, 40:21-28. That teaching could be combined with the teachings of Lubbers and Sicola and applied to the primary site or the alternate site (“*secondary storage facility*”).

241. A POSITA would have been motivated to combine the teachings of Ofek to “*a first write transaction in a first receive batch*” and “*a second write*

transaction in a second receive batch” when “the first and second receive batches are combined” with the teachings of Lubbers and Sicola in this manner to free up cache resources and improve system performance, which is a shared goal.

242. A POSITA would have reasonably expected to succeed in using an “overwriting” technique in this manner as merely adjusting the system software of Lubbers and Sicola with the teachings of Ofek to achieve a mutual goal.


243. Accordingly, the combined teachings of Lubbers, Sicola, and Ofek would have rendered obvious Claim 16.

I. Claim 17

1. ***[17preamble]: “A data redundancy method for storing a primary copy of data at a primary storage facility and for storing data that is redundant of the primary copy at a secondary storage facility, the method comprising”***

244. For purposes of this analysis, I have been told to assume that the preamble is limiting for this claim.

245. As discussed for elements [1a] and [1d], Lubbers and Sicola teaches or suggests a multi-site storage area network that includes multiple sites where a primary site serves as the “*primary storage facility*” “*for storing a primary copy of data*” and an alternative site, the “*secondary storage facility*,” where the redundant data is stored.




Ex-1005, FIG. 3 (annotated).

246. Accordingly, Lubbers teaches or suggests this preamble.

2. [17a]: “receiving a sequence of write requests at a first redundancy appliance at the primary storage facility;”

247. As discussed for element [1b], Lubbers and Sicola teaches or suggests the first controller (“first redundancy appliance”), at the primary site (“primary storage facility”), handling data traffic (“sequence of write requests”) by packing as many host writes as possible into a single buffer teaches a “a first redundancy appliance” “receiving a sequence of write requests.” Further, as discussed for element [1a], Lubbers and Sicola teaches or suggests storage at a primary site (“primary storage facility”).




Ex-1005, FIG. 3 (annotated).

248. Accordingly, Lubbers and Sicola would have rendered obvious this element.

3. [17b]: “storing data for the sequence of write requests in primary storage at the primary storage facility;”

249. As discussed for element [1b], Lubbers and Sicola teaches or suggests the first controller logs data, implements and manages logical units, and stores the same (“storing data”) in physical storage devices 302 (“*primary storage*”) coupled to the first controller using the FCAL 303 at the primary site (“*primary storage facility*”).



Ex-1005, FIG. 3 (annotated).

250. Accordingly, Lubbers and Sicola would have rendered obvious this element.


4. ***[17c]: “logging at least a sequence of write requests for transfer to the secondary storage facility;”***

251. As discussed for elements [1b] and [1c], and claim 12, Lubbers, Sicola, and Ofek would have rendered obvious this element.

5. ***[17d]: “forwarding the logged write requests as send batches to a second redundancy appliance at the secondary storage facility disjoint from the primary storage facility, the second storage facility batching the received write request as receive batches;”***

252. As discussed for element [1d], Lubbers and Sicola would have rendered obvious “*the secondary storage facility* disjoint from the *primary storage facility*.”

253. As discussed for claim 9, Lubbers, Sicola, and Ofek would have rendered obvious “forwarding write requests as send batches to *a second redundancy appliance at the secondary storage facility . . . the second storage facility* batching the received write request as receive batches.”




Ex-1005, FIG. 3 (annotated).

254. As discussed for element [17c], Lubbers and Sicola would have rendered obvious logged sequences of write requests, rephrased in this claim as “*logged write requests.*”

255. Accordingly, Lubbers, Sicola, and Ofek would have rendered obvious this element.

6. [17e]: “storing data associated with the receive batches in secondary storage at the secondary storage facility”

256. As discussed for claim 13, Lubbers, Sicola, and Ofek would have rendered obvious this element.




Ex-1005, FIG. 3 (annotated).

7. [17f]: “replacing a first write transaction in a receive batch at the second storage facility with a second write transaction wherein the second write transaction overwrites data written by the first write transaction; and”

257. As discussed for claim 14, Lubbers, Sicola, and Ofek would have rendered obvious this element.

8. [17g]: “the second redundancy appliance acting as the first redundancy appliance in response to a failover event.”

258. As discussed for claims [1c], [1f] and 2, Lubbers and Sicola would have rendered obvious this element.



Ex-1005, FIG. 3 (annotated).

J. Claim 18

1. [18preamble]: “A data redundancy system, comprising:”

259. For purposes of this analysis, I have been told to assume that the preamble is limiting for this claim.

260. As discussed for element [1preamble], Lubbers teaches or suggests “a data redundancy system.”

2. ***[18a]: “a primary storage facility including at least one primary redundancy appliance for receiving a sequence of write requests and for storing data for the sequence of write requests in storage associated with the primary storage facility, the primary storage facility batching the write requests as send batches for transfer to a secondary storage facility; and”***

261. As discussed for elements [1a], [1b], 1[c], Lubbers and Sicola would have rendered obvious this element, except “*batching the write requests as send batches.*”

262. As discussed for claim 9, Lubbers, Sicola, and Ofek would have rendered obvious “the primary storage facility batching the write requests as send batches for transfer to a secondary storage facility.”

3. ***[18b]: “the secondary storage facility disjoint from the primary storage facility including at least one secondary redundancy appliance for receiving the sequence of write requests and for storing data for the sequence of write requests in storage associated with the secondary storage facility, the second storage facility batching the received write request as receive batches, the second redundancy appliance acting as the primary redundancy appliance in response to a failover event;”***

263. As discussed for elements [1a]-[1f], Lubbers and Sicola would have rendered obvious this element, except “*batching received write requests.*” As discussed for claim 13, Lubbers, Sicola, and Ofek would have rendered obvious “*the second storage facility batching the received write request as receive batches.*”

4. *[18c]: “wherein a first write transaction in a receive batch at the second storage facility is replaced with a second write transaction where the second write transaction over writes data written by the first write transaction.”*

264. As discussed for claim 14, Lubbers, Sicola, and Ofek would have rendered obvious this element.

XI. CONCLUSION

265. I declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true, and that these statements were made with 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.

January 13, 2025
Date: _____



Kevin C. Almeroth, Ph.D.