

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re patent of: Farber et al.	§	
	§	Petition for <i>Inter Partes</i> Review
U.S. Patent No. 6,928,442	§	
	§	Attorney Docket No.: 47015.137
Issued: Aug. 9, 2005	§	
	§	Customer No.: 116298
Title: ENFORCEMENT AND	§	
POLICING OF LICENSE	§	Real Parties
CONTENT USING	§	in Interest: Rackspace US, Inc. and
CONTENT-BASED	§	Rackspace Hosting, Inc.
IDENTIFIERS	§	

DECLARATION OF MELVIN RAY MERCER

I, Melvin Ray Mercer, do hereby declare as follows:

1. I am making this Declaration at the request of Petitioner, Rackspace US, Inc., in connection with a Petition for *Inter Partes* Review of U.S. Patent No. 6,928,442 to Farber and Lachman, entitled “ENFORCEMENT AND POLICING OF LICENSE CONTENT USING CONTENT-BASED IDENTIFIERS” (“the ‘442 Patent”).

2. I am being compensated at my normal consulting rate for my work. My compensation is not dependent on and in no way affects the substance of my statements in this Declaration.

3. I have no financial interest in Petitioner. I have been informed that PersonalWeb Technologies, LLC (“PersonalWeb”) and Level 3 Communications, LLC (“Level 3”) each purport to own 50% of the ‘442 Patent. I have no financial

interest in PersonalWeb or Level 3, and I have had no contact with either company.

I similarly have no financial interest in the '442 Patent, and have had no contact

with the named inventors of the '442 Patent: David A. Farber and Ronald D.

Lachman.

4. In the preparation of this Declaration, I have studied:
 - a. the '442 Patent, RACK-1001;
 - b. the prosecution history of the '442 Patent, RACK-1002;
 - c. Ex Parte Reexamination 90/010,260, RACK-1003;
 - d. U.S. Patent No. 5,649,196, ("Woodhill"), RACK-1004;
 - e. U.S. Patent No. 4,845,715, ("Francisco"), RACK-1005;
 - f. U.S. Patent No. 6,135,646, ("Kahn"), RACK-1006;
 - g. Albert Langer, "Re: dl/describe (File descriptions)," post to the "alt.sources" newsgroup on August 7, 1991 ("Langer"), RACK-1007;
 - h. IPR2013-00082 ('791 Patent) RACK-1008; and
 - i. IPR2013-00083 ('280 Patent) RACK-1009.
5. In forming the opinions expressed below, I have considered:
 - a. The documents listed above;
 - b. The relevant legal standards, including the standard for obviousness provided in *KSR International Co. v. Teleflex, Inc.*,

550 U.S. 398 (2007) and any additional authoritative documents as cited in the body of this declaration; and

- c. My knowledge and experience based upon my work in this area as described below.

My Qualifications and Professional Experience

6. My qualifications are set forth in my curriculum vitae, a copy of which is attached as Appendix A to this Report. A short synopsis of that material follows.

7. I have more than 45 years of dual industrial and academic experience in Electrical Engineering and Computer Engineering. I received a B.S. in Electrical Engineering from Texas Tech University in 1968. From 1968 to 1973, I was a Research/Development Engineer at General Telephone and Electronics Sylvania in Mountain View, California, and I received an M.S. in Electrical Engineering from Stanford University in 1971. During this period, I programmed minicomputer systems (predecessors to personal computers, smartphones, and modern servers) in machine language, assembly language and various higher-level languages. I wrote simple Operating Systems, and most of the applications involved real-time processing as a significant aspect of the systems design.

8. From 1973 to 1977, I was a Member of Technical Staff at Hewlett-Packard's Santa Clara Division and subsequently at Hewlett-Packard Laboratories

in Palo Alto, California. During this time, I continued to develop application programs - mostly in the area of real-time data acquisition and control of systems. In addition to designing the software associated with these systems, I also designed interface hardware to interact with the software of the computers and accomplish various tasks. Some of the applications I developed involved a significant number of data files associated with (to my knowledge) the first full scale study of the impact of environmental factors on the degradation of liquid crystal materials – such as those used in electronic clocks and watches.

9. From 1977 to 1980, I was a Lecturer in the Division of Mathematics, Statistics, and Computer Science at the University of Texas at San Antonio. As the director of a laboratory for teaching students to program and build hardware interfaces with computers, I purchased, built, and operated some of the earliest personal computers. I received a Ph.D. in Electrical Engineering from the University of Texas at Austin in 1980.

10. From 1980 to 1983, I was a Member of Technical Staff at Bell Laboratories in Murray Hill, New Jersey. My work involved the programming of computers and the hardware design of components for communication systems. Among other things, I was part of a three-person team that designed, tested, and directed the manufacture of an integrated circuit that was a key component in a digital telephone modem. This work involved significant amounts of data – mostly

produced on a Cray machine with issues of version control, data access etc. I also was involved with telephone switching system engineers, who at that time were developing the 5ESS Telephone Switch. Such telephone systems have very sophisticated data handling constraints with, for example the classes of service and charges to customers for those services.

11. In 1983, I was appointed Assistant Professor of Electrical and Computer Engineering at the University of Texas at Austin. In 1987, I was promoted to Associate Professor and Professor in 1991. During this period I taught Computer Engineering courses at the undergraduate and graduate level, and I directed the research of graduate students. I consulted with (and my research was funded by) numerous industrial organizations (including AT&T).

12. In 1995, I was appointed Professor of Electrical and Computer Engineering, Leader of the Computer Engineering Group, and Holder of the Computer Engineering Chair at Texas A&M University in College Station, Texas. My teaching, my research, my technical publications, and my supervision of graduate students during this period included the areas of the design and implementation of digital hardware and software systems, and my administrative duties - including the growth and enhancement of the Computer Engineering Group - directly involved Internet-based communication and control issues. As with my work at The University of Texas at Austin during this period, I taught

courses at the undergraduate and graduate level, I directed the research of graduate students, and I consulted and did research with numerous organizations on topics related to the issues in this case. For example I oversaw the collection of data in an experiment to compare the relative efficiency of testing methods for integrated circuits. Multiple companies were involved, and significant data protection was required to properly manage information proprietary to the various manufacturers.

13. In September 2005, I retired from my teaching position, and the Regents of the Texas A&M University System appointed me as Professor Emeritus of Electrical and Computer Engineering at Texas A&M University.

14. Since 1984, I have been an independent consultant and provided private consultation and advice in Electrical and Computer Engineering to numerous entities including IBM Corp., Rockwell International, Motorola Semiconductor, AT&T, Inc. and SigmaTel. Part of my consulting work at Rockwell International was directly related to the design of telephone modems.

15. In 1994, my wife and I formed a company called Conference Management Services. It organizes technical conferences around the world, and it was one of the earliest companies to utilize on line databases as the basis for all organizing activities (including paper submission, review, and scheduling, conference registration, event scheduling and planning, support for exhibits, and hotel registration). I, with the help of undergraduate and graduate students in

Computer Engineering, designed and implemented the early versions of this system. Numerous entities (technical contributors, submission reviewers, administrative personnel, hotels, conference centers, etc.), provided and acquired information specific to their part of the conference activities. This information was carefully controlled and compartmentalized as required by our system. The interval of time when I was involved in these activities included the priority date for the patents at issue in this case. Today, Conference Management Services continues to provide such database / electronic file system enabled activities.

16. I first served as an expert witness at the request of the Office of the State Attorney General of Texas in 1984. The case was about long distance telephone service. Since that time, I have been hired by numerous law firms to provide them and their clients with expert consultation and expert testimony - often in the areas of patent infringement litigation related to Electrical and Computer Engineering.

17. I was actively involved in numerous professional organizations including the Institute of Electrical and Electronics Engineers ("IEEE"), and I was recognized as an IEEE Fellow in 1994. I was the Program Chairman for the 1989 International Test Conference, which is an IEEE-sponsored annual conference with (at that time) more than one thousand attendees and over one hundred presented papers. I won the Best Paper Award at the 1982 International Test Conference.

18. I also won a Best Paper Award at the 1991 Design Automation Conference, an annual conference with (at that time) more than ten thousand attendees and five hundred submitted papers, many of which related to the design of integrated circuit based systems. The subject of that award-winning paper involved trade-offs between power consumption and processing speed in integrated circuits.

19. I also won a Best Paper Award at the 1999 VLSI Test Symposium. I am the inventor of two United States patents that relate to the design of integrated circuits. I was selected as a National Science Foundation Presidential Young Investigator in 1986. That award included \$ 500,000 in research funding and the document commemorating that award was signed by the President of the United States.

20. Based upon these and other technical activities, I am familiar with the knowledge and capabilities of one of ordinary skill in the areas of distributed information systems and associated access controls. Specifically, my work with students -- undergraduates as well as Masters and Ph.D. candidates, with colleagues in academia, with engineers practicing in industry, with customers of Conference Management Services, and with lawyers and technical experts in the Computer Engineering area allowed me to become personally familiar with the level of skill of individuals and the general state of this art. Specifically, my

experience (1) in the industry, (2) with undergraduate and post-graduate students, (3) with colleagues from academia, and (4) with engineers practicing in the industry allowed me to become personally familiar with the level of skill of individuals and the general state of the art. Unless otherwise stated, my testimony below refers to the knowledge of one of ordinary skill in the computer network and software fields during the time period around the priority date for the '442 Patent.

My Understanding of the Relevant Legal Standards

21. I have been asked to provide my opinions regarding whether the claims of the '442 Patent are anticipated or would have been obvious to a person having ordinary skill in the art at the time of the alleged invention, in light of the prior art. It is my understanding that, to anticipate a claim under 35 U.S.C. § 102, a reference must teach every element of the claim. Further, it is my understanding that a claimed invention is unpatentable under 35 U.S.C. § 103 if the differences between the invention and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which the subject matter pertains. I also understand that the obviousness analysis takes into account factual inquiries including the level of ordinary skill in the art, the scope and content of the prior art, and the differences between the prior art and the claimed subject matter.

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

23. It is my understanding that some claims can be interpreted as “means plus function” claims under 35 U.S.C. § 112, paragraph 6. I understand that determining the broadest reasonable interpretation of “means plus function” claims requires first, defining the particular function of the limitation and second, identifying the corresponding structure for that function in the specification. I also understand that structure disclosed in the specification is corresponding structure only if the specification or prosecution history clearly links or associates that structure to the function recited in the claim.

Subject Matter of the ‘442 Patent

24. I have reviewed and understand the specification, claims, and drawings (RACK-1001) and the file history (RACK-1002) of U.S. Patent No. 6,928,442 (“the ‘442 Patent”) along with the file history of the *ex parte* reexamination 90/010,260 filed on the ‘442 Patent (RACK-1003).

25. The ‘442 Patent is directed to “a system in which a set of data files are distributed across a plurality of computers.” The data files may include digital messages, images, videos or audio signals. Identifiers for the data files being determined, using a function of the data comprising the particular data file,” “where name is based, at least in part, on a given function of the data in a data item or file.” The ‘442 Patent further describes how the system includes a license table including a listing data files and a licensed user that is authorized to have access to the object. (RACK-1001, Title, Abstract, 11:66-12:10.)

26. According to the ‘442 Patent, prior art systems identified data items based on their location or address within the data processing system. (RACK-1001, 1:21-31.) For example, files were often identified by their context or “pathname,” *i.e.*, information specifying a path through the computer directories to the particular file (*e.g.*, C:\MyDocuments\LawSchool\2L\PatentsOutline.doc). (See RACK-1001, 1:39-43.) The ‘442 Patent contends that all prior art systems operated in this manner, stating that “[i]n *all* of the prior data processing systems[,]

the names or identifiers provided to identify data items ... are *always* defined relative to a specific context,” and “there is *no* direct relationship between the data names and the data item.” (RACK-1001, 1:66–2:4, 2:13-14 (emphasis added).)

27. According to the ‘442 Patent, this prior art practice of identifying a data item by its context or pathname had certain shortcomings. For example, with pathname identification, the same data name may refer to different data items, or conversely, two different data names may refer to the same data item.

(RACK-1001, 2:13-17.) Moreover, because there is no correlation between the contents of a data item and its pathname, there is no *a priori* way to confirm that the data item is in fact the one named by the pathname. (RACK-1001, 2:19- 22.) Furthermore, context or pathname identification may more easily result in the creation of unwanted duplicate data items, *e.g.*, multiple copies of a file on a file server. (RACK-1001, 2:48-59.)

28. The ‘442 Patent purports to address these shortcomings. It suggests that “it is therefore desirable to have a mechanism ... to determine a common and substantially unique identifier for a data item, using only the data in the data item and not relying on any sort of context.” (RACK-1001, 3:7-12.) Moreover, “[i]t is further desirable to have a mechanism for reducing multiple copies of data items ... and to have a mechanism which enables the identification of identical data items so as to reduce multiple copies.” (RACK-1001, 3:13-16.)

29. The '442 Patent uses the terms “True Name” and “data identifier” to refer to the substantially unique identifier for a particular data item (RACK-1001, 6:7-11) and explains that a True Name is computed using a message digest function (*see* RACK-1001, 12:57-13:17). Preferred embodiments use either of the MD5 or SHA message digest functions to calculate a substantially unique identifier from the contents of the data item. (RACK-1001, 12:55-13:17.)

30. In the preferred embodiments, the substantially unique identifiers are used to “augment” standard file management functions of an existing operating system. (*See* RACK-1001, 6:12-20.) For example, a local directory extensions (LDE) table¹ is indexed by a pathname or contextual name of a file and also includes True Names for most files. (*See* RACK-1001, 8:21-28.) A True File registry (TFR) lists True Names, and stores “location, dependency, and migration information about True Files.” (*See* RACK-1001, 8:29-30, 35-37.) True Files are identified in the True File registry by their True Names, and can be looked up in the registry by their True Names. (*See* RACK-1001, 8: 32-34, 23:61–63.)

31. When a data item is to be “assimilated” into the data processing system, its substantially unique identifier (True Name) is calculated and compared

¹ The patent describes an LDE table as a data structure which provides information about files and directories in the system and includes information in addition to that provided by the native file system. (*See* RACK-1001, 8:19-26.)

to the True File Registry to see if the True Name already exists in the Registry. (See RACK-1001, 14:44-54.) If the True Name already exists, this means that the data item already exists in the system and the to-be-assimilated data item (*i.e.*, the scratch file) need not be stored. (See RACK-1001, 14:55-64.) Conversely, if the True Name does not exist in the Registry, then a new entry is created in the Registry which is then set to the just-calculated True Name value, and the data items can be stored. (See RACK-1001, 14:65- 15:4.)

32. The '442 Patent describes a “mechanism that ensures that licensed files are not used by unauthorized parties.” (RACK-1001, 32:42-44.) The disclosed technique includes recording a file identifier in a license table along with the identity of each user or system which is authorized/licensed to have a copy of the file. (RACK-1001, 8:53-56 and 32:56-64). The license table information is used to compare the contents of a user's system with the license table data to determine whether the user has an authorized or unauthorized copy of a file. (RACK 1001, 32:65-33:7). The system can also deny access to the file without proper authorization. (RACK-1001, 32:48-49.)

33. I have been informed that the '442 Patent claims priority to U.S. Patent Application No. 08/960,079 filed Oct. 24, 1997 which itself claims priority to U.S. Patent Application No. 08/425,160, filed on April 11, 1995. I

understand this means the '442 Patent is considered to have been filed on April 11, 1995 for the purposes of determining whether a reference will qualify as prior art.

File History Overview

34. I understand that during the prosecution of the '442 Patent, when considering the present claims, the original patent Examiner did not rely upon a reference where the content of a data file is used to determine an identifier for the data file.

35. During an interview the examiner requested that the patent owner identify support for certain aspects of the claims.

In the interview the Examiner requested that applicants show support in the specification for the licensing aspects of the claims. During the interview applicants' representative directed the Examiner, *inter alia*, to pages 62-63 of the specification, to the section titled "Track for Licensing Purposes." In that section one possible way of using the invention to track data for licensing purposes is described. Note that the present application describes a number of so-called mechanisms which can be used together to build a system. The "Track for Licensing Purposes" mechanism is what the patent application refers to as an extended mechanism ("Extended mechanisms run within application programs over the operating system. These mechanisms provide solutions to specific problems and applications." P12L29-31) The Examiner is also directed to the description of the license table on pg. 15, lines 20-23 and pg. 18, line 18 to pg. 19, line 3. Applicants give these references as examples only, and they are not intended to in any way limit the scope of the claimed invention.

(RACK-1002, pp. 129-130.)

36. Portions of the specification relating to the license table are reproduced below.

Each record 150 of the license table 136 records a relationship between a licensable data item and the user licensed to have access to it. Each license table record 150 includes the information summarized in the following table, with reference to FIGURE 9:

Field	Description
True Name	True Name of a data item subject to license validation.

Field	Description
licensee	identity of a user authorized to have access to this object.

(RACK-1002, pp. 329-330.) As the specification above discloses, consistent with the ordinary understanding, to determine if a file is licensed or authorized, one must use a two- step process to first determine whether the content matches the list of items and second determine whether the user is licensed or authorized to have access to the file.

After the claims were amended to include limitations relating to “licensed or authorized” users and copies, a Notice of Allowance was issued. (RACK-1002, pp. 76-84.)

37. The applicant stated that none of the cited prior art provided the ability to restrict access of files to only licensed or authorized parties as required by the claims. (RACK-1002, pp. 130-133.)

38. During the reexamination proceedings, the patent examiner rejected all of the claims based on the Hellman reference. In seeking to maintain claims 1, 2, 4, 7, 23, 27, 28 and 30, the patent owner identified two main groups for the claims.

THE CLAIMED INVENTION

For the purposes of this discussion, the claims of the '442 Patent can be considered in two groups: *viz.*, those relating to *obtaining or delivering data items* [claims 1-9, 14-21, 37, 45, 52/4]; and those relating to *policing unauthorized or unlicensed copies of data items* [claims 10-13, 22-36, 38-44, 46-51, 52/51, 53, 54-56].²
(RACK-1003, p 222.)

39. The patent owner distinguished claim 1 in group 1 by suggesting that the Hellman reference failed to disclose providing the name in a request and causing a copy of the file to be provided to the requester. (RACK-1003, pp 230-231). With respect to claim 23 in group 2 claims, the patent owner suggested that Hellman failed to provide a list of file names to determine whether unauthorized or unlicensed copies of some of the plurality of data files are present on a particular computer. (RACK-1003, pp. 236-238.)

40. I further understand that the original Examiner in the '442 Patent, as well as the reexamination Examiner, were not made aware of U.S. Patent No. 4,845,715 to Francisco and U.S. Patent No. 6,135,646 to Kahn.

41. In light of the information in the Woodhill, Kahn, Francisco and Langer references, it is my opinion that the methods claimed in the '442 Patent were well known to persons skilled in that art at the time the application for the '442 Patent was filed. As will be described in extensive detail below, at the time the application for the '442 Patent was filed, there was nothing new about using a function of the contents of a data file to determine a unique identifier for the data file and then subsequently utilizing the unique identifier to store and/or retrieve the identified data file.

Claims of the '442 Patent

42. I understand that claims 1, 2, 4, 7, 23, 27, 28, and 30 of the '442 Patent are being challenged in the above-referenced *inter partes* review.

43. It is my understanding that in order to properly evaluate the '442 Patent, the terms of the claims must be construed. It is my understanding that the claims are to be given their broadest reasonable interpretation in light of the specification. It is my further understanding that claim terms are given their ordinary and accustomed meaning as would be understood by one of ordinary skill in the art, unless the inventor, as a lexicographer, has set forth a special meaning

for a term. In order to construe the claims, I have reviewed the entirety of the ‘442 Patent, as well as its prosecution history.

44. It is my understanding that a panel of the Patent Trial and Appeal Board (PTAB) has already construed certain claim terms of the ‘442 Patent appearing in the parent patents to the ‘442 Patent. Specifically, the PTAB construed the claim term “data file” in the ‘791 Patent and ‘280 Patent, to be “a named data item, such as a simple file that includes a single, fixed sequence of data bytes or a compound file that includes multiple, fixed sequences of data bytes.” Based on my review of the PTAB’s constructions (see Decision to Institute *Inter Partes* Review, IPR-2013-00082 (RACK-1008) and IPR2013-00083 (RACK-1009)), I believe this construction is consistent with the ordinary meaning and the specification of the ‘442 Patent. (*See* RACK-1001, 5:47-55.)

45. The term “data file” appears in claims 1, 2, 4, 7, 23 and 30. Except to the extent qualified in the paragraphs below, I have assumed and used the above construction of “data file” of the PTAB in forming my conclusions herein.

46. It is my understanding that features set forth in a preamble of a claim are construed as limitations if they recite essential structure or steps, or if they are necessary to give life, meaning, and vitality to the claim. The preamble of claim 7 includes the phrase “wherein some of the computers communicate with each other using a TCP/IP communication protocol.” The body of claim 7 makes no

reference to how the plurality of computers communicate with each other and certainly do not specify any features unique to using a TCP/IP communication protocol. Thus, it is my assessment that this feature set forth in the preamble does not add essential structure or steps to the claim and should not be entitled to patentable weight when considering the limitations within the body of claim 7.

Level of Skill in the Art

47. In my opinion, the level of ordinary skill in the art needed to have the capability of understanding the scientific and engineering principles applicable to the '442 Patent is a bachelor's degree in Electrical Engineering , Computer Engineering or Computer Science; or equivalent industry experience as one designing data storage systems and programming software applications. A strength in one of these areas can compensate for a weakness in another.

Technical Basis Underlying the Grounds of Rejections Set Forth in the Petition for *Inter Partes* Review

48. In the paragraphs that follow, I detail my opinions regarding various grounds of invalidity based on application of specific prior art references to specific challenged claims of the '442 Patent. In each case, I first summarize the applied reference(s) and then apply the teachings and/or suggestions thereof to specific claims of the '442 Patent, interpreted in accord with the foregoing constructions of specific claim terms. In the absence of an express construction of

terminology or language, I have concluded that the inventor has not set forth special meaning or lexicography, and that such terminology or language would be understood by one of ordinary skill in the art in accordance with ordinary and accustomed meanings.

Challenge #1: Claims 1, 2, 4, 23, 27 and 30 are obvious under 35 U.S.C. § 103 based on Woodhill in view of Francisco

49. U.S. Patent No. 5,649,196 to Woodhill et al., entitled “System and Method for Distributed Storage Management on Networked Computer Systems Using Binary Object Identifiers” (“Woodhill,” RACK-1004) describes use of context- and location-independent identifiers for purposes that are analogous to those disclosed in the ‘442 Patent. The ‘442 Patent granted on U.S. Application No. 09/28,160, filed April 1, 1999 claiming priority as a divisional application to U.S. Application No.: 08/960,079 filed Oct. 24, 1997, and also claiming priority to U.S. Application No.: 08/425,160 filed April 11, 1995. The Woodhill patent granted on U.S. Application No. 08/555,376 filed Nov. 9, 1995, which was itself a file wrapper continuation of application 08/085,596, filed July 1, 1993. I have been advised that the ‘442 Patent has an effective priority date of April 11, 1995, while the Woodhill patent has an earlier effective prior art date of July 1, 1993 under 35 U.S.C. §102(e).

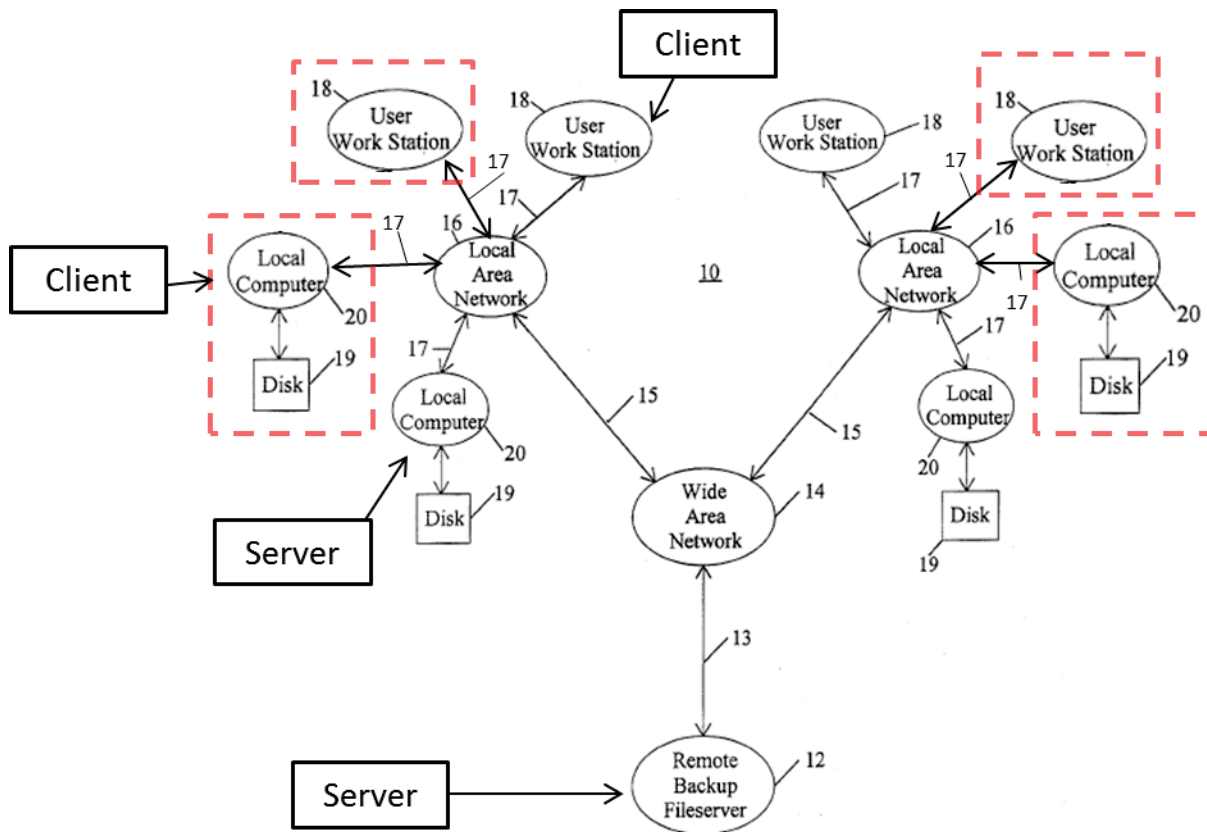
50. Woodhill discloses a distributed storage system that used “Binary Object Identifiers” to identify and access files, and to manage file back-ups, amongst other functions. (RACK-1004.) As Woodhill explains, a “Binary Object Identifier 74 [of Fig. 3] ... is a unique identifier for each binary object to be backed up.” (RACK-1004, 4:45-47.) Woodhill’s Binary Object Identifiers include three fields—a CRC value, a LRC value, and a hash value—each calculated from the contents of the binary object. (RACK-1004, 8:1- 33.) As Woodhill emphasized, “[t]he critical feature to be recognized in creating a Binary Object Identifier 74 is that the identifier should be based on the contents of the binary object so that the Binary Object Identifier 74 changes when the contents of the binary object changes.” (RACK-1004, 8:58-62.) Woodhill used these identifiers to identify binary objects that had changed since the most recent backup, so that “only those binary objects associated with the file that have changed must be backed up.” (RACK-1004, 9:7-14.) “[D]uplicate binary objects, even if resident on different types of computers in a heterogeneous network, can be recognized from their identical Binary Object Identifiers 74.” (RACK-1004, 8:62-65.)

51. Woodhill discloses that the method of distributed management of storage space and data using content based identifiers is implemented on a networked computer system having a plurality of local computers, work stations, local area networks, a wide area network and at least one remote backup file

server. “[T]he Distributed Storage Manager program 24 stores a compressed copy of every binary object it would need to restore every disk drive 19 on every local computer 20 somewhere on the local area network 16 other than on the local computer 20 on which it normally resides. At the same time, the Distributed Storage Manager program 24 transmits every new or changed binary object to the remote backup file server 12.” (RACK-1004, 9:30-38.)

52. The local computers can act as clients when making requests for files and can act as servers when responding to requests for providing files served on their local storage devices. Woodhill recognizes that “[s]ince most restores of files on a local area network 16 consist of requests to restore the most recent backup version of a file, the local copies of binary objects serve to handle very fast restores for most restore requests that occur on the local area network 16. If the local backup copy of a file does not exist or a prior version of a file is required, it must be restored from the remote backup file server 12.” (RACK-1004, 10:27-34.)

53. In order to better explain the description set forth in Woodhill, I have annotated Fig. 1 of the Woodhill to represent the description that “[e]ach local area network 16 includes multiple user workstations 18 and local computers 20 each in communication with their respective local area network 16 via data paths 17.” (RACK-1004, 3:24-26.) The annotated portions of the figure set forth below are noted by the red box surrounding the feature.



54. It is my understanding that U.S. Patent No. 4,845,715 to Francisco issued on July 4, 1989 and is therefore prior art to the '442 Patent under 35 U.S.C. §102(b). Francisco discloses:

“an improved method for maintaining the integrity of a data processing system through controlled authentication and subsequent authorization of both selected programs and potential users thereof. In its broader aspects,

the invention includes the generation and storage of a selective electronic identification indicia, based upon the nature and content of the program itself, for each software program in the system together with a separately stored correlation of such electronic identification indicia with user identity therewith in association with a regeneration of such electronic identification indicia each time the program is sought to be used and a checking of said regenerated electronic identification indicia against a stored catalog of such identification indicia and against a stored permitted user register therefore.” (RACK-1005, 1:38-53.)

55. Francisco further describes that the “[a]mong the advantages of the subject invention is a markedly improved system security to ensure only utilization of authenticated programs, the utilization of such programs only by authorized users thereof and the immediate detection of any modifications or changes introduced into a software program.” (RACK-1005, 1:54-59.)

56. The ‘442 Patent discloses that the license table 136 records both information about the data and the user licensed to have access to it. (RACK-1001, 11:66-12:10.) The ‘442 Patent discloses a license table that records the content based identifier of the data item subject to license validation and the identity of the user authorized to have access to this object (see table reproduced

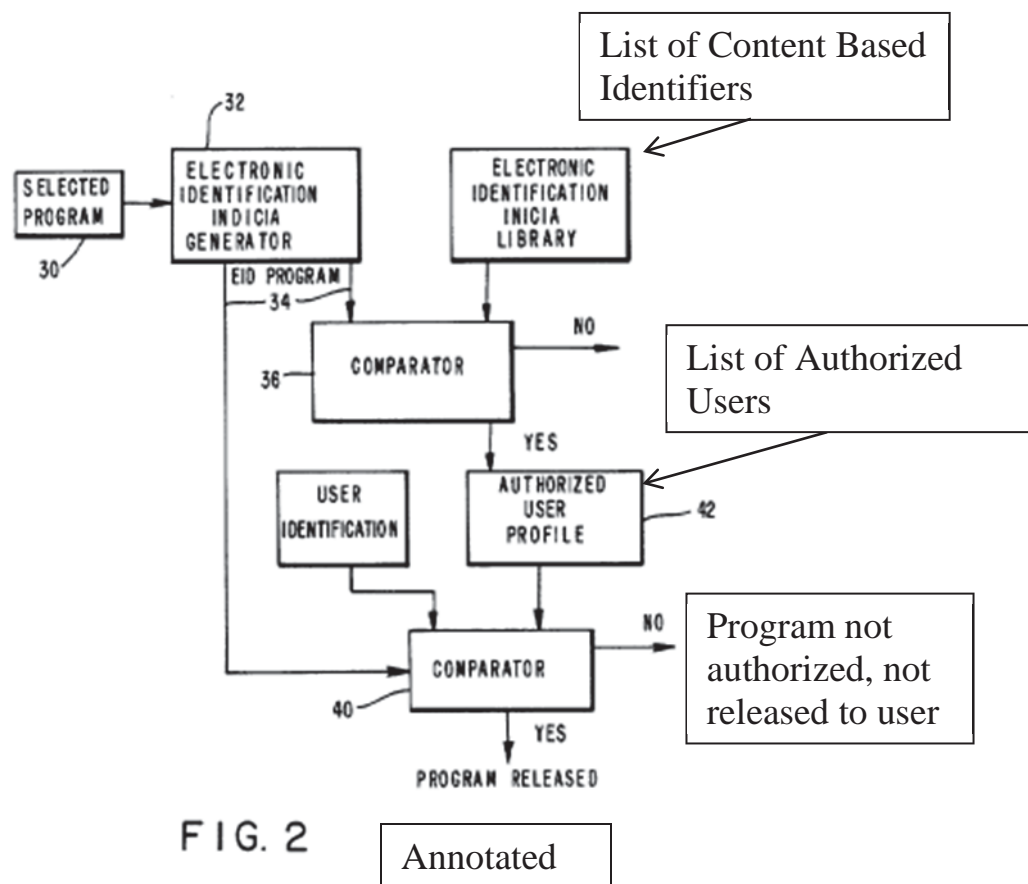
below). To determine if a file is licensed the file must be identified in the license table and the licensee must be identified.

Field	Description
True Name	True Name of a data item subject to license validation.

Field	Description
licensee	identity of a user authorized to have access to this object.

(RACK-1002, pp. 329-330.)

57. Francisco describes use of “electronic identification indicia” calculated over the content of a binary encoding of a program to generate a unique and selective identifier for the program. (See RACK-1005, 2:22-42.) Similar to the technique set forth above for the ‘442 Patent, *Francisco* uses the content based electronic identification indicia for an authenticity and/or integrity check of the content of a selected program as well as to correlate individual user profiles against programs that the user is authorized to use. (See RACK-1005, 1:38-59.) I have annotated FIG. 2 of *Francisco* to show the list of content based electronic identification indicia in the library that are used to determine whether selected program 30 will be considered an authorized program and released to the user.



58. I have reviewed the identification algorithm example (see RACK-1005, 2:31-42) described in *Francisco* and confirm that it constitutes a function over all the contents of a program encoding. In a manner analogous to the ‘442 patent’s license table, an authorized user library 16 (in *Francisco*) stores electronic identification indicia for particular programs in correlative relation with identifications of users that are authorized, or entitled, to use the particular programs. (See RACK-1005, 2:51-64.)

59. *Francisco* describes a system in which identifiers that are unique and selective for individual program files are used to determine those users that are authorized to use the program. *Francisco*’s “electronic identification indicia” are

calculated over the content of a binary encoding of a program to generate a unique and selective identifier for the program. (*See* RACK-1005, 2:22-42.) I have reviewed the identification algorithm example (*see* RACK-1005, 2:31-42.) described in *Francisco* and confirm that it constitutes a hash (albeit a simple hash) over all the contents of a program encoding.

60. It is my understanding that there are many legal bases under which two prior art references could be combined. It is my opinion that a person of ordinary skill in the art would have found it obvious to combine the method of Woodhill in view of *Francisco* because it would have been a use of known techniques to improve a similar methodology in the same way.

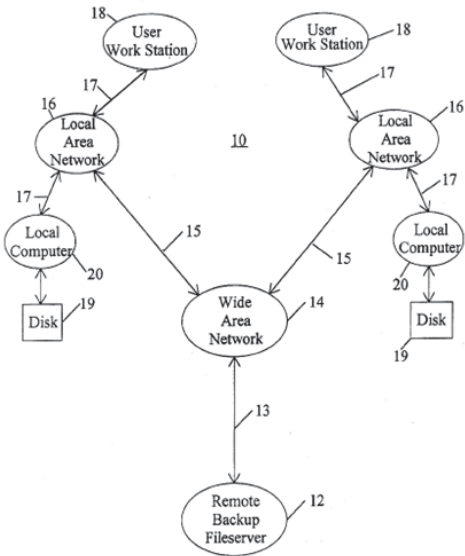
61. Woodhill and *Francisco* both disclose methods of managing data storage systems utilizing content based unique identifiers. Both references disclose using a function of the contents of the data file to determine a unique identifier. Woodhill discloses that file information includes access control data and therefore suggests that access control to the data files is already a part of the system set forth in the Woodhill patent. Although Woodhill does not expressly disclose how to use access control data to limit the provision of files to unauthorized/unlicensed parties, it was well known at the time to maintain a listing of authorized users and compare the requesting parties credentials against a listing of authorized users. As an example, *Francisco* discloses “a method for maintaining integrity through

selectively coded preauthorized software program and user identification and subsequent automatic authentication of both a selected program and permitted user thereof when system resources are to be utilized.” (RACK-1005, 1:10-15.)

62. It would have been obvious to combine Woodhill’s access control data with the user authentication techniques taught by Francisco to only provide copies of the requested file to an authorized party identified in a table as set forth in Francisco and not provide copies of files to unauthorized parties as set forth in Francisco. Combining the known authorization techniques in Francisco to improve the suggested access control features of Woodhill would yield predictable results and provide greater access control for the data of Woodhill.

63. It is my opinion that a person of ordinary skill in the art would find that Woodhill in view of Francisco renders obvious each and every element of at least claims 1, 2, 4, 23, 27 and 30. I address each of the elements of these claims in the charts set forth below.

65. The following chart describes in detail how Woodhill in combination with Francisco discloses each and every element of the claims 1, 2, 4, 23, 27 and 30.

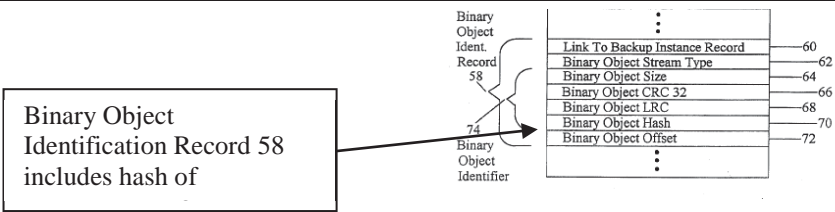
The '442 Patent Claims	U.S. Pat. No. 5,649,196 to Woodhill et al. in view of U.S. Pat. No. 4,845,715 to Francisco
<p>[1a] In a system in which a plurality of files are distributed across a plurality of computers, a method comprising:</p>	<p>Woodhill discloses a system for delivering content including distributing a set of data files across a network of servers. Woodhill is titled “System and Method for Distributed Storage Management on Networked Computer Systems Using Binary Object Identifiers.” Woodhill describes a simplified networked computer system including a remote backup file server and one or more local computers with data storage devices (clients and servers) connected to one or more local area networks.</p> <p>“The Distributed Storage Manager stores a compressed copy of every binary object it would need to restore every disk drive 19 on every local computer 20 somewhere on the local area network 15 other than on the local computer 20 on which it normally resides. At the same time, the Distributed Storage Manager program 24 transmits every new or changed binary object to the remote backup file server 12.”</p> <p>(RACK-1004, 9:31-38.)</p>  <pre> graph TD subgraph LAN1 [Local Area Network 16] U1[User Work Station 18] L1[Local Computer 20] D1[Disk 19] L1 <--> D1 L1 <--> 17 LAN1 end subgraph LAN2 [Local Area Network 16] U2[User Work Station 18] L2[Local Computer 20] D2[Disk 19] L2 <--> D2 L2 <--> 17 LAN2 end LAN1 <--> 15 WAN[Wide Area Network 14] LAN2 <--> 15 WAN WAN <--> 13 RBS[Remote Backup Fileserver 12] </pre>

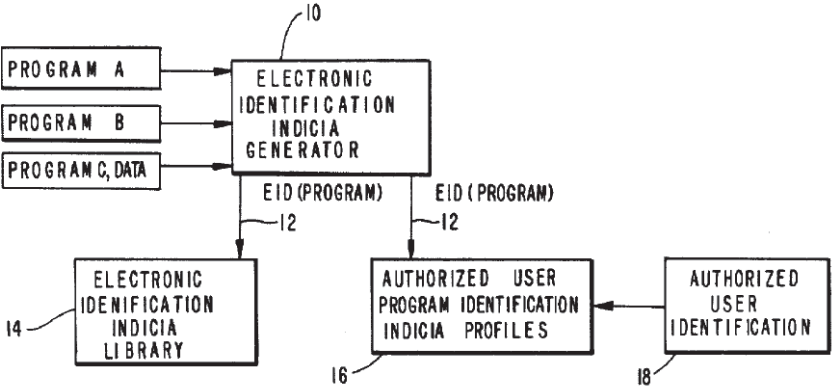
The '442 Patent Claims	U.S. Pat. No. 5,649,196 to Woodhill et al. in view of U.S. Pat. No. 4,845,715 to Francisco
<p>[1b] obtaining a name for a data file, the name being based at least in part on a given function of the data, wherein the data used by the given function to determine the name comprises the contents of the data file;</p>	<p>Just like the disclosure in the '280 patent, Woodhill discloses using content based data identifiers to identify (name) sequences of bytes, called “binary objects.” If the file is less than 1 megabyte, then a single binary object represents the file. (RACK-1004, 4:23-26.)</p> <p>“In step 138, a Binary Object Identification Record 58 is created in File Database 25 for each of the binary objects currently being processed. Each of these Binary Object Identification Records 58 are associated with the Backup Instance Record 42 created in step 130. The Binary Object Identifier 74 portion of each Binary Object Identification Record 58 is comprised of the Binary Object Size field 64, the Binary Object CRC32 field 66, the Binary Object LRC field 68 and the Binary Object Hash field 70. <u>Each of the fields of the Binary Object Identifier 74 may be four (4) bytes in length and is calculated from the contents of each binary object.</u></p> <p>The Binary Object Size field 64 may be set equal to the byte-size of the binary object. The Binary Object CRC32 field 66 may be set equal to the standard 32-bit Cyclical Redundancy Check number calculated against the contents of the binary object taken one (1) byte (8 bits) at a time. Those of ordinary skill in the art will readily recognize the manner in which the Cyclical Redundancy Check number is calculated. The Binary Object LRC field 68 may be set equal to the standard Longitudinal Redundancy Check number calculated against the contents of the binary object taken four (4) bytes (32 bits) at a time using the following algorithm:</p>

The '442 Patent Claims	U.S. Pat. No. 5,649,196 to Woodhill et al. in view of U.S. Pat. No. 4,845,715 to Francisco
	<pre> BINARY OBJECT LRC = (initialized value) for each double word (32 bits) of the binary object data: LRC = LRC (XOR) double word of binary object data end loop </pre> <p><u>The Binary Object Hash field 70 is calculated against the contents of the binary object taken one (1) word (16 bits) at a time using the following algorithm:</u></p> <pre> HASH = (initialized value) for each word (16 bits) of the binary object: rotate current HASH value by 5 bits HASH = HASH + 1 HASH = HASH + (current word (16 bits) of binary object) end loop </pre> <p>(RACK-1004, 7:60 – 8: 32 (emphasis added).)</p>
<p>[1c] and in response to a request for the a data file, the request including at least the name of the particular file, causing a copy of the file to be provided from a given one of the plurality of computers,</p>	<p>Woodhill discloses, responsive to a client request for the data file, the request including a hash of the contents of the data file, causing the data file to be provided to the client. For example, the data processing system of Woodhill allows restores (client request) of binary objects (data files) using their Binary Object Identifiers during the Backup/Restore Routine. The identifiers are sent as part of the request to restore a binary object. Also, the data processing system of Woodhill performs periodic self-audits by initiating a restore of a randomly selected binary object, identified by its Binary Object Identification Record, which includes its Binary Object Identifier. A Binary Object Identifier includes a hash value calculated by a hash function of the contents of the binary object that it identifies. For example, the technique of restoring files using the content based identifiers is described in detail below:</p> <p>“The technique of "granularizing" "large" files also becomes useful when a current version of a file (comprised of current versions of binary objects) must be</p>

The '442 Patent Claims	U.S. Pat. No. 5,649,196 to Woodhill et al. in view of U.S. Pat. No. 4,845,715 to Francisco
	<p>restored to a previous version of that file (comprised of previous versions of binary objects). Each binary object comprising the current version of the file can be restored to the binary object comprising the previous version of the file by restoring and updating only those "granules" of the current version of the binary objects that are different between the current and previous versions of the binary objects. This technique is illustrated in the flow chart depicted in FIG. 5i. Program control begins at step 442 where the Distributed Storage Manager program 24 obtains from the user the identities of the current and previous versions of the file (comprised of binary objects) which needs to be restored. Program control continues with step 443 where the Distributed Storage Manager program 24 compiles a list of all binary objects comprising the current version of the user-specified file. This information is obtained from File Database 25. Program control then continues with step 444 where the Distributed Storage Manager program 24 calculates "contents identifiers" for each "granule" within the current version of each binary object as it exists on the local computer 20. Program control then continues with step 446 where the Distributed Storage Manager program 448 transmits an "update request" to the remote backup file server 12 <u>which includes the Binary Object Identification Record 58</u> for the previous version of each binary object as well as the list of "contents identifiers" calculated in step 444. Program control continues with step 448 where the Distributed Storage Manager program 24 reconstitutes each previous version of the binary objects according to the technique illustrated in the flow chart depicted in FIG. 5h. Program control then continues with step 450</p>

The '442 Patent Claims	U.S. Pat. No. 5,649,196 to Woodhill et al. in view of U.S. Pat. No. 4,845,715 to Francisco
	<p>where the Distributed Storage Manager program 24, for each binary object, compares the "contents identifier" of the next "granule" in the work area of remote backup file server 12 against the corresponding "contents identifier" calculated in step 444. Program control continues with step 452 where the Distributed Storage Manager program 24 determines whether the "contents identifiers" match. If so, program control is returned to step 450 since this "granule" is the same on the local computer 20 and on the remote backup file server 12. If the Distributed Storage Manager program 24 determines, in step 452, that the "contents identifiers" do not match, program control continues with step 454 where the Distributed Storage Manager program 24 <u>transmits the "granule" to the local computer 20.</u> Program control then continues with step 456 where the "granule" received by the local computer 20 is written directly to the current version of the binary object at the appropriate location. Program control then continues with step 458 where the Distributed Storage Manager program 24 determines whether there are any more "granules" to be examined for the binary object currently being processed. If so, program control is returned to step 450; otherwise the file restore routine is terminated at step 460. After all "granules" are received from the remote backup file server 12, the binary object has been restored to the state of the previous version."</p> <p>(RACK-1004, 17:18-18:9 (emphasis added).)</p>

The '442 Patent Claims	U.S. Pat. No. 5,649,196 to Woodhill et al. in view of U.S. Pat. No. 4,845,715 to Francisco
	 <p>The diagram shows a box labeled "Binary Object Identification Record 58 includes hash of" with an arrow pointing to a table. The table has the following fields and reference numerals: Link To Backup Instance Record (60), Binary Object Stream Type (62), Binary Object Size (64), Binary Object CRC 32 (66), Binary Object LRC (68), Binary Object Hash (70), and Binary Object Offset (72). A bracket groups the first three fields (60-64) and is labeled "Binary Object Ident. Record 58". A bracket groups the last three fields (68-72) and is labeled "Binary Object Identifier".</p> <p>FIG. 3</p>
<p>[1d] wherein a copy of the requested file is only provided to licensed parties.</p>	<p>“Those of ordinary skill in the art will recognize that these data streams may represent regular data, extended attribute data, access control data, etc.” RACK-1004, 7:40.</p> <p>Access control data is understood by those skilled in the art to provide information concerning which users are entitled to access the particular data file.</p> <p>Although Woodhill does not expressly disclose how to use access control data to limit the provision of files to unauthorized/unlicensed parties, it was well known at the time to maintain a listing of authorized users and compare the requesting party’s credentials against a listing of authorized users before providing a copy of the requested file. As an example, U.S. Patent 4,845,714 to Francisco, discloses “a method for maintaining integrity through selectively coded preauthorized software program and user identification and subsequent automatic authentication of both a selected program and permitted user thereof when system resources are to be utilized.” (RACK-1005, 1:10-15.)</p> <p>It would have been obvious to combine Woodhill’s access control data with the user authentication techniques taught by Francisco to only provide copies of the requested file to an authorized party identified in a table as set forth in Francisco and not provide</p>

The '442 Patent Claims	U.S. Pat. No. 5,649,196 to Woodhill et al. in view of U.S. Pat. No. 4,845,715 to Francisco
	<p>copies of files to unauthorized parties as set forth in Francisco. Francisco discloses using a library of all authorized profiles correlating authorized user identifications with all programs which each user is entitled to use. (RACK-1005, 2:56-64.)</p> <p>“FIG. 1 is a schematic flow chart illustrative of library type storage of electronic identification indicia for both software programs and authorized user profiles therefore” (RACK- 1005, 2:3-6.)</p>  <p style="text-align: center;">FIG. 1</p> <p>Referring to Fig. 2 reproduced below:</p> <p>“the electronic identification of the user making the request for access is introduced into a second comparator 40. The comparator 40 may again suitably comprise the automatic logic unit of a general purpose digital computer. Also introduced into the second comparator 40 is the second electronic identification indicia 34 emanating from the generator 32 and the authorized user profile 42 obtained from the profile register 16.</p> <p><u>If the paired inputs to the second comparator do not</u></p>

The '442 Patent Claims	U.S. Pat. No. 5,649,196 to Woodhill et al. in view of U.S. Pat. No. 4,845,715 to Francisco
	<p><u>match, the requested program 30 will not be released for use</u> and an appropriate signal made to the system monitor to initiate appropriate investigative and corrective action. If, however, the paired signal inputs to the second comparator 40 provide a match, the requested program 30 may be released for use by the particular identified user.”</p> <p>(RACK-1005, 3:20-36 (emphasis added).)</p> <div data-bbox="600 795 1304 1495" data-label="Diagram"> <pre> graph TD SP[SELECTED PROGRAM 30] --> EIDG[ELECTRONIC IDENTIFICATION INDICIA GENERATOR 32] EIDG -- "EID PROGRAM 34" --> C1[COMPARATOR 36] EIL[ELECTRONIC IDENTIFICATION INDICIA LIBRARY] --> C1 C1 -- NO --> Exit1(()) C1 -- YES --> C2[COMPARATOR 40] UI[USER IDENTIFICATION] --> C2 AUP[AUTHORIZED USER PROFILE 42] --> C2 C2 -- NO --> Exit2(()) C2 -- YES --> PR[PROGRAM RELEASED] </pre> <p>The flowchart illustrates a two-step comparison process. First, a 'SELECTED PROGRAM 30' is input to an 'ELECTRONIC IDENTIFICATION INDICIA GENERATOR 32', which produces an 'EID PROGRAM 34'. This is compared in 'COMPARATOR 36' against an 'ELECTRONIC IDENTIFICATION INDICIA LIBRARY'. If the comparison fails (NO), the process ends. If it succeeds (YES), the 'EID PROGRAM 34' is then compared in 'COMPARATOR 40' against a 'USER IDENTIFICATION' and an 'AUTHORIZED USER PROFILE 42'. If this second comparison fails (NO), the process ends. If it succeeds (YES), the 'PROGRAM RELEASED'.</p> </div> <p>FIG. 2</p> <p>Francisco’s electronic identification indicia (EIDs) library includes the contents of the authorized user profile library 16 along with a correlation of the particular programs authorized for usage by each user. (See RACK-1005, 2:51-64.) The ‘442 Patent discloses that a</p>

The '442 Patent Claims	U.S. Pat. No. 5,649,196 to Woodhill et al. in view of U.S. Pat. No. 4,845,715 to Francisco
	<p>“licensee” is “a user authorized to have access to the object.” (RACK-1001, 12:9-10.) Thus, the user profile library 16 in Francisco is understood to provide a listing of each user authorized to have access to the program (a licensed party) in the same manner as disclosed in the '442 Patent. As explained above, and shown graphically in Fig 2 reproduced above, only if the “User Identification “ matches the “Authorized [Licensed] User Profile” is the requested program provided to the requester.</p>
<p>[2.] A method as in claim 1 further comprising: determining, using at least the name, whether a copy of the data file is present on a particular one of said computers.</p>	<p>Woodhill discloses a first computer using at least the file name (Binary Object Identification) to determine if the same binary object is present on a second computer.</p> <p>“Program control then continues with step 444 where the Distributed Storage Manager program 24 calculates "contents identifiers" for each "granule" within the current version of each binary object as it exists on the local computer 20. Program control then continues with step 446 where the Distributed Storage Manager program 448 transmits an "update request" to the remote backup file server 12 which includes the Binary Object Identification Record 58 for the previous version of each binary object as well as the list of "contents identifiers" calculated in step 444. Program control continues with step 452 where the Distributed Storage Manager program 24 <u>determines whether the "contents identifiers" match. If so, program control is returned to step 450 since this "granule" is the same on the local computer 20 and on the remote backup file server 12.</u>”</p> <p>(RACK-1004, 17:36-46 (emphasis added).)</p>

The '442 Patent Claims	U.S. Pat. No. 5,649,196 to Woodhill et al. in view of U.S. Pat. No. 4,845,715 to Francisco
	<p>In addition to the disclosure in Woodhill, the Francisco reference also utilizes an indicia based on the contents of the file to determine if a file is present on a particular computer.</p> <p>“As shown in Fig. 2, the practice of the herein disclosed method, a selected program 30 requested to be released for use is introduced into an electronic identification indicia generator 32 and to therein generate a second electronic identification signal 34 (EID-Program S). This second electronic identification indicia 34 is first introduced into a comparator 36 together with the first electronic identification 12, for such selected program (EID-Program S), the latter being retrieved from the library 14. . . . <u>If such first and second electronic identification indicia 12 and 34 for the selected Program S do not match, it is indicative of the fact that the requested program differs in some respect from the base or true program from which the first electronic identification indicia 12 (EID-Program S) was derived and such lack of match serves as a signal to take appropriate investigative and corrective action.</u>”</p> <p>(RACK-1005, 2:65-3: 16 (emphasis added).)</p> <p>Both references disclose utilizing a comparison of the content based names to determine is a file is present on a computer.</p>

The ‘442 Patent Claims	U.S. Pat. No. 5,649,196 to Woodhill et al. in view of U.S. Pat. No. 4,845,715 to Francisco
<p>[4.] A method as in claim 1, further comprising: maintaining accounting information relating to the data files.</p>	<p>The ‘442 Patent describes examples of accounting information relating to the data files to be the type of information associated with uses of the files, as well as the times when files are created, deleted, or copied. (RACK-1001, 31:57-32:24.)</p> <p>As set forth below, Woodhill expressly discloses tracking this same information.</p> <p>“For each File Identification Record 34 in File Database 25, one or more Backup Instance Records 42 are created that contain information about the file (identified by File Identification Record 34) at the time that the file is backed up. Each time that a file is backed up, a Backup Instance Record 42 is created for that file. Each Backup Instance Record 42 consists of the following elements: (1) Link to File Identification Record 44; (2) <u>Backup Cycle Identifier 46 (identifies the particular backup cycle during which the Backup Instance Record 42 is created)</u>; (3) File Size 48; (4) Last Modified Date/Time 50; (5) <u>Last Access Date/Time 52</u>; (6) File Attributes 54 (e.g., read-only, system, hidden); (7) <u>Delete Date 56 (date on which the file was deleted)</u>; and (8) <u>Insert Date 57 (date on which the Backup Instance Record 42 was created)</u>.”</p> <p>(RACK-1004, 3:64-4:11.)</p>
<p>[23a] A method comprising: obtaining a list of file names, at</p>	<p>As explained above in [1b], Woodhill discloses a system for determining unique identifiers based on the content of the file and then maintaining listings of</p>

The '442 Patent Claims	U.S. Pat. No. 5,649,196 to Woodhill et al. in view of U.S. Pat. No. 4,845,715 to Francisco
least one file name for each of a plurality of files, each of said file names having been determined, at least in part, by applying a function to the contents of the corresponding file; and	<p>those identifiers across a networked computer system. (RACK-1004, 3:7-5:11). For example:</p> <p>“The Distributed Storage Manager program 24 of the present invention builds and maintains the File Database 25 on one of the disk drives 19 on each local computer 20 in the networked computer system 10 according to the structure illustrated in FIG. 3. <u>The File Database 25 stores information relating to each file that has been backed up by the Distributed Storage Manager program 24 since the initialization of that program on each local computer 20.</u>” (RACK-1004, 3:45-54.)</p> <div><div></div><div></div></div> <p>Fig. 2, reproduced above, illustrates how the system allocates the storage on each computer in the network, including providing a “file database 25” which includes the information of Fig. 3, reproduced</p>

The '442 Patent Claims	U.S. Pat. No. 5,649,196 to Woodhill et al. in view of U.S. Pat. No. 4,845,715 to Francisco
	<p>above, for each file that has been backed up by the system. As explained above with respect to [1b], the Binary Object Identifier 74 is determined as a function of the contents of the file.</p> <p>As explained below, Francisco also discloses obtaining a list of file names by first determining a unique identifier for each file based on a function of the contents of the file and then storing a listing of the identifiers in a library.</p> <p>“This invention relates to data processing system security and more particularly to a method for maintaining integrity through selectively coded preauthorization software program and user identification and subsequent automatic authentication of both a selected program and permitted user thereof when system resources are to be utilized.”</p> <p>(RACK-1005, 1:9-15.)</p> <p>“In its broadest aspects, the invention includes the generation and storage of a selective electronic identification indicia, based upon the nature and content of the program itself, for each software program in the system together with a separately stored correlation of such electronic identification indicia with user identity therewith in association with a regeneration electronic indicia against a stored catalog of such identification indicia and against a stored permitted user register therefore.”</p> <p>(RACK-1005, 1:41-59.)</p>

The '442 Patent Claims	U.S. Pat. No. 5,649,196 to Woodhill et al. in view of U.S. Pat. No. 4,845,715 to Francisco
	<p>“The generator 10 . . . is adapted to generate a first electronic identification indicia 12 (EID-Program A) that uniquely and selectively identifies the submitted program. By way of example, in a relatively simple approach thereto such <u>generator 10 could generate a selective and unique electronic identification indicia by use of a preprogrammed algorithm in accord with which the total number of ones and zeroes in the binary coded input Program A could be totalled. . . . The resulting electronic numerical indicia would then, in all probability, be selectively unique for the particular program.</u> The algorithm would be periodically varied to enhance system security.</p> <p><u>This first electronic identification indicia 12 for a particular program, herein termed EID (Program A), is stored, together with similarly generated indicia for other programs B, C, D ..., in an EID library 14,</u> which could suitably be a read only memory (ROM) or a random access memory (RAM).”</p> <p>(RACK-1005, 2:22-50 (emphasis added).)</p>
<p>[23b] using at least said list to determine whether unauthorized or unlicensed copies of some of the plurality of data files are present on a particular computer.</p>	<p>As set forth above, Woodhill discloses a method for distributed storage management on networked computer systems utilizing content based identifiers. Woodhill suggests that access control can be a part of the system, but does not expressly disclose a technique for accomplishing access control. (See Rack-1004, 7:40.) Although Woodhill does not expressly disclose how to use access control data to limit the provision of files to unauthorized/unlicensed parties or to prevent unauthorized files to be used on the system, it was well known at the time to maintain a listing of authorized users and compare the requesting party’s</p>

The '442 Patent Claims	U.S. Pat. No. 5,649,196 to Woodhill et al. in view of U.S. Pat. No. 4,845,715 to Francisco
	<p>credentials against a listing of authorized users before providing a copy of the requested file. As an example, Francisco discloses “a method for maintaining integrity through selectively coded preauthorized software program and user identification and subsequent automatic authentication of both a selected program and permitted user thereof when system resources are to be utilized.” (RACK-1005, 1:10-15.)</p> <p>It would have been obvious to combine Woodhill’s access control data with the user authentication techniques taught by Francisco to only provide copies of the requested file if the file is authenticated and is provided to an authorized party identified in a table as set forth in Francisco. Conversely, copies of files are not provided to unauthorized parties as set forth in Francisco. Francisco discloses using a library of all authorized profiles correlating authorized user identifications with all programs which each user is entitled to use. (RACK-1004, 2:56-64.)</p> <p>“Among the advantages of the subject invention is a markedly improved system security to ensure only utilization of authenticated programs, the utilization of such programs only by authorized users thereof and the immediate detection of any modifications or changes introduced into a software program.”</p> <p>(RACK-1005, 1:54-59.)</p> <p>Thus, as an initial check utilizing the content based identifiers, Francisco discloses that a copy of a program being requested can be determined to be unauthorized if there have been any modifications or</p>

The '442 Patent Claims	U.S. Pat. No. 5,649,196 to Woodhill et al. in view of U.S. Pat. No. 4,845,715 to Francisco
	<p>changes in the software program.</p> <p>As a further step, Francisco operates by checking a list of file names determined at least in part by the content of the file and then checking to see if that file is authorized for use by the requesting user. As set forth below, even if Francisco determines that the content of the file is authorized, the copy of the file residing on a computer may be unlicensed or unauthorized if the user is not authorized for use of the designated file.</p> <p>“As shown in Fig. 2 [reproduced below], the practice of the herein disclosed method, a selected program 30 requested to be released for use is introduced into an electronic identification indicia generator 32 and to therein generate a second electronic identification signal 34 (EID-Program S). This second electronic identification indicia 34 is first introduced into a comparator 36 together with the first electronic identification 12, for such selected program (EID-Program S), the latter being retrieved from the library 14. . . . If such first and second electronic identification indicia 12 and 34 for the selected Program S do not match, it is indicative of the fact that the requested program differs in some respect from the base or true program from which the first electronic identification indicia 12 (EID-Program S) was derived and such lack of match serves as a signal to take appropriate investigative and corrective action.”</p> <p>(RACK-1005, 2:65-3:16.)</p>

**The '442 Patent
Claims**

**U.S. Pat. No. 5,649,196 to Woodhill et al. in view of
U.S. Pat. No. 4,845,715 to Francisco**

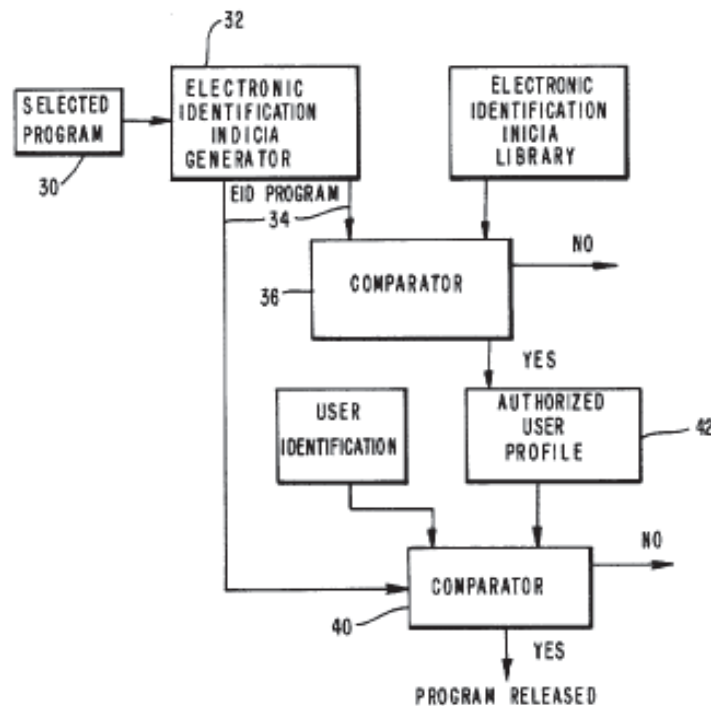


FIG. 2

“If the paired inputs to the second comparator do not match, the requested program 30 will not be released for use and an appropriate signal made to the system monitor to initiate appropriate investigative and corrective action.” (RACK-1005, 3:29-33.) Thus, Francisco determines that the requested copy of the file is not authorized.

It would have been obvious to combine the method determining content based file identifiers and distributed storage of Woodhill with Francisco’s method for maintaining a library of authorized files identifiers and comparing those identifiers to other files to determine if they are authorized files. Woodhill already suggests access control within the

The '442 Patent Claims	U.S. Pat. No. 5,649,196 to Woodhill et al. in view of U.S. Pat. No. 4,845,715 to Francisco
	system and Francisco provides a known technique of determining which files are authorized based on a comparison of content based identifiers such that the addition of Francisco's access control would utilize a known element to obtain a predictable result.
[27.] A method as in claim 23 wherein the function is a message digest function or a hash function.	As set forth above in [1b], Woodhill discloses that the function is a hash function.
[30.] A method as in claim 23 wherein the function produces a substantially unique value based on the data comprising the data file.	Woodhill discloses using four different functions, including a hash function calculated against the contents of the binary object to generate substantially unique values. (RACK-1004, 7:60-8:65.) These functions, including the hash function value, is included in the Binary Object Identifier, that "is used to uniquely identify a particular binary object." (RACK-1004, 8:33-36 and Fig. 3.)

Challenge #2: Claim 7 is obvious under 35 U.S.C. § 103 based on Woodhill in view of Kahn

66. The Woodhill reference is described above in greater detail and the relevant portions are set forth in the chart below. Woodhill creates unique file identifiers based on a hash of the contents of the data file and further discloses that file information includes access control data. Thus, Woodhill suggests that access control to the data files is already a part of the system set forth in the Woodhill

patent. Although Woodhill does not expressly disclose that the data file represents a digital image, it does not exclude the possibility that the data files being stored within the system would include digital images. Digital images were a common file type at the time of filing of the '442 Patent and thus, it is my opinion that a person of ordinary skill in the art would understand that the Woodhill reference contemplates that the data files would include digital images.

67. It is my understanding that U.S. Patent No. 6,135,646 to Kahn et al. issued on October 24, 2000 based on an application claiming priority to Oct. 22, 1993 which is before the earliest priority, April 11, 1995, claimed by the '442 Patent. Therefore, it is my understanding that Kahn qualifies as prior art to the '442 Patent under 35 U.S.C. §102(e).

68. Kahn provides an overview of the disclosed management of digital objects as follows:

In what follows, we provide examples of operations which may be conducted with assistance of the servers and services of the kind shown in FIG. 1. The operations include registration of rights, recordation of transfers of rights, deposit of objects in repositories, generation and use of handles, arranging compensation for use of objects and for licensing and transfer of rights (e.g., exclusive rights) in objects (under simple or non-simple terms), use of digital signature and other protective mechanisms to insure the

integrity of the transactions within the system, management of rights, and obtaining objects from repositories.

(RACK-1006, 7:59)

69. Kahn discloses an example of a system for storing and controlling delivery of digital object to authorized users of the system. (RACK-1006, 2:17-32.) Kahn describes that the holders of rights in digital objects are enabled to control the terms and conditions under which they are accessed by users in a network, or are granted to others. (RACK-1006, Abstract and 3:4-26.) The digital objects can include, among other things, digital image files. (RACK-1006, 1:17-21.) Still further, users can be authorized to receive delivery of the digital objects if their certificates match the information on file with the system. (RACK-1006, 22:26-26:38.) However, users requesting digital objects that do not have authorization to receive a digital object will not be sent a copy of the requested digital object. (RACK-1006, 25:1-3 and 54-57.) Digital objects can be requested between system components communication via direct connections such as TCP/IP. (RACK-1006, 10:39-48.)

70. It is my understanding that there are many legal bases under which two prior art references could be combined. It is my opinion that a person of ordinary skill in the art would have found it obvious to combine the method of

Woodhill in view of Kahn because it would have been a use of known techniques to improve a similar methodology in the same way.

71. Woodhill and Kahn both disclose methods of managing data storage systems having unique identifiers. Woodhill creates unique file identifiers based on a hash of the contents of the data file and further discloses that file information includes access control data. Thus, Woodhill suggests that access control to the data files is already a part of the system set forth in the Woodhill patent. Although Woodhill does not expressly disclose that the data file represents a digital image, it does not exclude the possibility that the data files being stored within the system would include digital images.

72. The Kahn reference expressly discloses that stored files within the system may be “digital objects” that can be “any set of sequences of bits or digits and an associated unique identifier.” (RACK-1006, 1:11-13.) “Digital objects may include conventional digital representations of works (books, papers, images, sounds, software). . . .” (RACK-1006, 1:18-20.) Thus, Kahn expressly discloses that a data file can represent a digital image and the use of a distributed storage system such as Woodhill to store such digital images would be an obvious use of a known system to achieve predictable results.

73. Woodhill recognizes that the data streams within its system include those which provide access control data. (RACK-1004, 7:40.) Access control data

is understood by those skilled in the art to provide information concerning which users are entitled to access the particular data file. For example, a person of ordinary skill in the art would understand that the access control data of Woodhill refers to controlling access to certain files to system administrators while restricting access to such files by regular users of the system. Although Woodhill does not expressly disclose how to use access control data to limit the provision of files to unauthorized/unlicensed parties, it was well known at the time to maintain a listing of authorized users and compare the requesting party's credentials against a listing of authorized users. Moreover, to a person of ordinary skill in the art, a data storage or archive system would not provide restricted files in response to client requests from unauthorized users.

74. The Kahn reference provides a system that restricts access to digital image files, among other digital representations of works, to those that are authorized to obtain copies of the files. As explained more fully in the chart below, the digital objects of Kahn are stored on a network that prevents unauthorized access.

75. Kahn discloses a procedure for retrieving a digital object from the system after a user has set-up an account. (RACK-1006, 23:48-26:38.) When requesting a copy of a digital file, the system checks whether the requester is an authorized party and if they are not, rather than sending the requested object, the

system sends “an invalid-random-value-tag response to the requesting system.” (RACK-1006, 25:1-8 and 25:54-58.) Only after verifying that the requesting party is authorized to receive the requested file, does the system “transmit [] to the requesting system: ...the [digital] object, signed by the repository.” (RACK-1006, 26:31-38.)

76. In this particular timeframe, I was involved in the evaluation of a system to control and distribute digital images via a network. The recent introduction of the World Wide Web and its promise of broad distribution of digital images was a “current hot topic.” Monetization of existing archives of digital images via their distribution using networked systems would have been a strong motivation to combine the exact file identification and distribution of Woodhill with the support for “business requirements” to safeguard and legitimately restrict the use of digital images to be disseminated as taught in Kahn. Thus, one of ordinary skill in the art would have been motivated to combine Kahn and Woodhill to utilize two existing capabilities in order to realize the benefits and insure against potential problems associated with their combination.

77. It would have been obvious to combine Woodhill’s access control data with the user authentication techniques taught by Kahn to only provide copies of the requested digital image file to an authorized party and not provide copies of files to unauthorized parties as set forth in Kahn.

78. It is my opinion that a person of ordinary skill in the art would find that Woodhill in view of Kahn renders obvious each and every element of at least claim 7. I address each of the elements of claim 7 in the chart set forth below.

The '442 Patent Claims	U.S. Pat. No. 5,649,196 to Woodhill et al. in view of U.S. Pat. No. 6,135,646 to Kahn et al.
[7a] A method, in a system in which a plurality of files are distributed across a plurality of computers, wherein some of the computers communicate with each other using a TCP/IP communication protocol, the method comprising:	As explained in detail above, Woodhill discloses a method for distributing files across of a plurality of computers. The statement in the preamble that “wherein some of the computers communicate with each other using a TCP/IP communication protocol” appears to be extra verbiage that should not be provided patentable weight as it is not referenced in the body of the claim. Since the claim is silent as to what type of communications need to occur, it is only necessary that some components communicate within the system using TCP/IP. To the extent that this well known communication protocol is determined to be important to the claimed method steps, it is respectfully submitted that the Kahn patent recognizes that “messages will be sent between system components via direct connections (e.g., “TCP/IP”).” (RACK-1006, 10:44-46.) It would be obvious to have some computers in the network of Woodhill communicate with each other using a TCP/IP communication protocol as taught by Kahn.
[7b] obtaining a name for a data file, the contents of said data file representing a digital image, the name having been determined using at least a given function of the data in the data file, wherein the data	See the discussion of claim element [1b] above concerning the features of Woodhill disclosing the claim elements relating to obtaining a name for a data file, the name being determined using at least a given function of the data in the data file. While Woodhill does not expressly disclose that the data file represents a digital image, it does not exclude the possibility that the data files being stored within the

The ‘442 Patent Claims	U.S. Pat. No. 5,649,196 to Woodhill et al. in view of U.S. Pat. No. 6,135,646 to Kahn et al.
used by the given function to determine the name comprises the contents of the data file; and	<p>system would include digital images. Digital images were a common file type at the time of filing of the ‘442 Patent and thus, it is my opinion that a person of ordinary skill in the art would understand that the Woodhill reference contemplates that the data files would include digital images.</p> <p>However, to the extent that Woodhill is deemed to not disclose such file types to a person of ordinary skill in the art at the time of filing, the Kahn reference expressly discloses that files may be “digital objects” that can be any set of sequences of bits or digits and an associated unique identifier.” “Digital objects may include conventional digital representations of works (books, papers, images, sounds, software). . . .” Thus, Kahn expressly discloses that a data file can represent a digital image.</p>
[7c] in response to a request for the data file, the request including at least the name of the data file, providing a copy of the file from a given one of the plurality of computers,	See the discussion of claim element [1c] above concerning the disclosure of Woodhill.
[7d] wherein a copy of the requested file is not provided to unlicensed parties or to unauthorized parties.	<p>“Those of ordinary skill in the art will recognize that these data streams may represent regular data, extended attribute data, <u>access control data</u>, etc.”</p> <p>(RACK-1004, 7:40 (emphasis added).)</p> <p>Access control data is understood by those skilled in the art to provide information concerning which users are entitled to access a particular data file.</p>

The '442 Patent Claims	U.S. Pat. No. 5,649,196 to Woodhill et al. in view of U.S. Pat. No. 6,135,646 to Kahn et al.
	<p>Moreover, to a person of ordinary skill in the art, a data storage or archive system would not provide restricted files in response to client requests from unauthorized or unlicensed users.</p> <p>An example of a system for controlling access to digital image files is provided in Kahn. The Kahn reference provides a system that restricts access to digital image files, among other digital representations of works, to those that are authorized to obtain copies of the files. As explained more fully in the passages below, the digital objects of Kahn are stored on a network that prevents unauthorized access.</p> <p>“In general, in one aspect, the invention features a method of managing digital objects in a network, the objects are stored at locations accessible in the network using a storage technique which renders the digital objects secure against unauthorized access.”</p> <p>(RACK-1006, 2:17-21.)</p> <p>“Another general aspect of the invention concerns maintaining a record of information concerning digital objects stored on a network. <u>The digital objects are stored on the network in a manner that restricts unauthorized access to and transactions associated with the digital objects.</u> A reference service is provided on the network, separate from the storage of the digital objects, for recording information about accesses to and transactions associated with the digital objects. Information about accesses to and transactions associated with the digital objects is recorded in the reference service.”</p>

The '442 Patent Claims	U.S. Pat. No. 5,649,196 to Woodhill et al. in view of U.S. Pat. No. 6,135,646 to Kahn et al.
	<p>(RACK-1006, 3:28-39.)</p> <p>“In this way, it is not necessary to store the objects at the same location as the validation information and <u>any authorized person on the network</u> (e.g., a court, or a government employee, or the rights holder, or a user) <u>may have access to</u> the validation and historical information and, <u>if authorized, the object itself</u>. When applied broadly to a large number and variety of rights holders and users, the system will produce a digital object infrastructure of enormous value to the conduct of business.”</p> <p>(RACK-1006, 4:63-5:5.)</p> <p>“In general, in one aspect, the invention features a method of managing digital objects in a network, the objects are stored at locations accessible in the network using a storage technique which renders the digital objects secure against unauthorized access.”</p> <p>(RACK-1006, 2:17-21.)</p> <p>Kahn discloses the procedure for retrieving a digital object from the system after a user has set-up an account. (RACK-106, 23:48-26:38.) When requesting a copy of a digital file, the system checks whether the requester is an authorized party and if they are not, rather than sending the requested object, the system sends “an invalid-random-value-tag response to the requesting system.” (RACK-1006, 25:1-8 and 25:54-58.) Only after verifying that the requesting party is authorized to receive the requested file, does the system “transmit [] to the requesting system: ...the [digital] object, signed by the repository.”</p>

The '442 Patent Claims	U.S. Pat. No. 5,649,196 to Woodhill et al. in view of U.S. Pat. No. 6,135,646 to Kahn et al.
	Woodhill discloses a method of distributing files across a network of servers and using content based identifiers to uniquely identify data stored on the network. Kahn discloses a method of restricting access to digital images to rights holders and users they authorize to obtain a copy of the digital image. It would have been obvious to combine the Kahn access controls with the distributed storage of Woodhill.

Challenge #3: Claim 28 is obvious under 35 U.S.C. § 103 based on Woodhill in view of Francisco and Langer

79. The Woodhill and Francisco and their application to claim 23 are discussed above in greater detail.

80. It is my understanding that the reference, Albert Langer, “Re: dl/describe (File descriptions),” post to the “alt.sources” newsgroup on August 7, 1991 (“Langer”) (RACK-1007), was published more than a year before the priority date of the ‘442 Patent and is therefore prior art to the ‘442 Patent under 35 U.S.C. §102(b). The Langer reference identifies a problem in “uniquely identifying files which may have different name and/or be in different directories on different systems (and also of being sure that files with the same name are identical ...). (RACK-1007, p. 3.) The Langer reference suggests that a solution to this problem would be to create unique identifiers based on the contents of the file. The Langer reference discloses “[a] simple method of defining a unique identifier that does

NOT include a particular site identifier would be to use a hash function on the entire contents of the file. This can be generated locally without requiring a registration system and if long enough the chances of collision are negligible. I would suggest using a cryptographic hash function such as MD5 which generates a 16 byte result.” (RACK-1007, p. 4.)

81. It is my understanding that there are many legal bases under which two prior art references could be combined. It is my opinion that a person of ordinary skill in the art would have found it obvious to combine the method of Francisco in view of Langer because at least it would have been a use of known techniques to improve a similar methodology in the same way.

82. Woodhill and Langer both disclose methods of managing data storage in distributed storage systems utilizing content based unique identifiers. Both references disclose using a hash function of the contents of the data file to determine a unique identifier. Woodhill recognizes “[t]hose of ordinary skill in the art will recognize that there exist many different ways to generate the Binary Object Identifier 74 (e.g. . . . using different calculations) and that the procedures set forth above is only one way of establishing the Binary Object Identifier 74.” (RACK-1004, 8:52-58.) Woodhill recognizes that it is important that “the possibility of two different binary objects being assigned the same Binary Object

Identifier 74 be very small.” (RACK-1004, 8:33-36.) Thus, there is a motivation to generate a content based identifier with limited chances of collision.

83. Langer expressly identifies the MD5 hash function as the hash function of choice to generate a unique identifier with negligible chances of collision. Accordingly, a person of ordinary skill in the art would be motivated to utilize the “cryptographic hash function such as MD5” disclosed in Langer as the hash function in determining the Binary Object Identifier of Woodhill because, at a minimum, the combination would have used known techniques to improve a similar system in the same way.

84. It is my opinion that a person of ordinary skill in the art would find that Woodhill, in view of Francisco and Langer renders obvious each and every element of claim 28. I address each of the elements of these claims in the charts set forth below.

The ‘442 Patent Claims	U.S. Pat. No. 5,649 to Woodhill in view of U.S. Pat. No. 4,845,715 to Francisco in view of Langer
[28.] A method as in claim 23 wherein the function is selected from the functions: MD4, MD5, and SHA.	As set forth above in [23], the combination of Woodhill and Francisco disclose all of the features of claim 23. Woodhill expressly discloses performing a hash of the contents of the data file. “The Binary Object Hash field 70 is calculated against the contents of the binary object taken one (1) word (16 bits) at a time using the following algorithm:

The '442 Patent Claims	U.S. Pat. No. 5,649 to Woodhill in view of U.S. Pat. No. 4,845,715 to Francisco in view of Langer
	<pre data-bbox="721 296 1287 443"> HASH = (initialized value) for each word (16 bits) of the binary object: rotate current HASH value by 5 bits HASH = HASH + 1 HASH = HASH + (current word (16 bits) of binary object) end loop </pre> <p data-bbox="597 499 956 537">(RACK-1004, 8:22-32.)</p> <p data-bbox="597 564 1430 1192"> Woodhill does not expressly disclose that the type of hash function contemplated for use in determining the identifier. Woodhill does disclose that “[t]hose of ordinary skill in the art will recognize that there exist many different ways to generate the Binary Object Identifier 74 (e.g. . . . using different calculations) and that the procedures set forth above is only one way of establishing the Binary Object Identifier 74.” (RACK-1004, 8:52-58.) Woodhill recognizes that it is important that “the possibility of two different binary objects being assigned the same Binary Object Identifier 74 be very small.” (RACK-1004, 8:33-36.) </p> <p data-bbox="597 1220 1414 1551"> However, it was well known to persons of ordinary skill in the art that MD5 hash functions were useful in performing hashes of the contents of data files. Thus, it would have been obvious to a person of ordinary skill in the art to utilize an MD5 hash function to perform the described hash of the contents of the data file according to Woodhill. </p> <p data-bbox="597 1579 1393 1761"> In the alternative, a person of ordinary skill in the art would have looked to the disclosure of the Langer reference for the specific suggestion to utilize an MD5 hash function: </p> <p data-bbox="597 1789 1403 1871"> “A simple method of defining a unique identifier that does NOT include a particular site identifier would be </p>

The '442 Patent Claims	U.S. Pat. No. 5,649 to Woodhill in view of U.S. Pat. No. 4,845,715 to Francisco in view of Langer
	<p>to use a hash function of the entire contents of the file. This can be generated locally without requiring a registration system and if long enough the chances of collision are negligible. ... I would suggest using a cryptographic hash function such as MD5 which generates a 16 byte result.”</p> <p>(RACK-1007, p. 4.)</p> <p>A person of ordinary skill in the art would be motivated to utilize the “cryptographic hash function such as MD5” disclosed in Langer as the hash function in determining the Binary Object Identifier of Woodhill because, at a minimum, the combination would have used known techniques to improve a similar system in the same way.</p>

Availability for Cross-Examination

85. In signing this declaration, I recognize that the declaration will be filed as evidence in a contested case before the Patent Trial and Appeal Board of the United States Patent and Trademark Office. I also recognize that I may be subject to cross examination in the case and that cross examination will take place within the United States. If cross examination is required of me, I will appear for cross examination within the United States during the time allotted for cross examination.

Right to Supplement

86. I reserve the right to supplement my opinions in the future to respond to any arguments that Patentee raises and to take into account new information as it becomes available to me.

Jurat

87. 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 further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Dated: Oct 10, 2013

By: Dr. M. Ray Mercer
Dr. M. Ray Mercer
College Station, TX