


**BEFORE THE PATENT TRIAL AND APPEAL BOARD IN THE UNITED
STATES PATENT AND TRADEMARK OFFICE**

Trial Nos.: IPR 2014-00066
Re: U.S. Patent No. 6,928,442
Patent Owners: PersonalWeb Technologies, LLC & Level 3 Communications
Petitioner: Rackspace US, Inc. & Rackspace Hosting, Inc.

* * * * *

UPDATED DECLARATION OF ROBERT B. K. DEWAR

1. My name is Robert B. K. Dewar. I reside at 1591 Carpenter Hill Road, Bennington, VT 05201, and am working for the software company Ada Core Technologies in New York, New York, that I founded. I have been retained by PersonalWeb Technologies, LLC. I understand that PersonalWeb is an owner of patents involved in several *Inter Parte Reviews* (IPRs). This declaration relates to U.S. Patent No. 6,928,442 (the '442 patent), and the IPR that it is involved in which I understand to be IPR No. 2014-00066.
2. I am currently Professor Emeritus at New York University in the Department of Computer Science. My curriculum vitae (CV) is attached



at the rear of this declaration. My educational and professional background is summarized below.

3. To summarize my educational and academic background, I was educated at a preparatory school in England, and attended one year of public school there. I then immigrated to the United States where I attended the University of Chicago Lab School. I then attended the University of Chicago where I earned BS and PhD degrees in the field of Chemistry. My PhD work in Chemistry involved substantial use of computers. After receiving my PhD degree, I took a position as Assistant Professor of Computer Science at the Illinois Institute of Technology in 1968.
4. I thereafter received tenure as an Associate Professor of Computer Science at the Illinois Institute of Technology. I was Full Professor of Computer Science, CIMS, NYU, since 1975, and am currently Professor Emeritus at New York University in the Department of Computer Science. I was Chairman of the Department of Computer Science from 1977-1980 at CIMS, NYU. I was Director of Undergraduate Studies, Computer Science from 1981-1987, CIMS, NYU. I was Associate Director of Courant Institute, from 1994-1997.
5. My teaching experience as a Professor includes (i) undergraduate courses relating to Introduction to programming for non-majors, first course for

majors (CS1/CS2 equivalents, taught in Fortran, Pascal, SNOBOL, Ada 95), Assembly language and Computer Architecture, Operating Systems, Microprocessors, Text processing systems, Programming Languages, Formal language theory, Numerical Analysis, Data processing, and Introduction to Algol-68, and (ii) graduate courses relating to Formal language theory, Theory of Computation, Algorithms, Programming Languages, Compilers, Data Communications, Microprocessor Architectures, Real Time Systems, Information Theory, and Programming by Transformation.

6. I have authored or co-authored numerous articles and books relating to computer science, programming, and data processing systems, as shown in my CV.
7. My research has centered on compilers, programming languages, operating systems, and microprocessor architectures. I have engaged in extensive consulting (apart from serving as an expert witness) in the area of operating systems and distributed systems, including writing several operating systems for Honeywell used in distributed applications such as airline reservations and freight management.
8. In 1994 I founded Ada Core Technologies to commercially exploit the work on Ada (programming language) that I was involved with at New

York University. I became CEO and President of Ada Core Technologies.

9. I retired from New York University in 2005.
10. I have received numerous awards including the SIGAda (Special Interest Group on Ada) award for outstanding Ada Community Contributions, as well as several other awards, which are listed in my attached CV.
11. I am familiar with USENETs and Bulletin Board Systems, and used the same in the 1980s and 1990s.
12. I have served as an expert witness on a number of occasions in cases involving patent and copyright issues. A list of several cases in which I served as an expert witness can be found in my attached CV.
13. I am being compensated for my time at the rate of \$500/hour for service that I am providing in connection with these Inter Partes Reviews (IPRs). My compensation is not contingent upon my performance, the outcome of this review, or any issues involved or related to this review.
14. I have reviewed and am familiar with at least the following documents, and any other document mentioned herein: the '442 patent and the prosecution history thereof; U.S. Patent No. 5,649,196 to Woodhill (Woodhill); U.S. Patent No. 4,845,715 to Francisco; U.S. Patent No. 6,135,646 to Kahn, Langer; the Mercer declaration; and the Mercer

deposition transcript. It is my understanding that the '442 patent has an effective filing date of April 11, 1995.

My Understanding of the Law Regarding Patent Validity

15. The following is what I have been told about the law regarding validity of a patent, and it represents my understanding of the same. It is my understanding that Rackspace (who I understand is a “petitioner” in this proceeding) has the burden of proving invalidity in an IPR proceeding by a preponderance of the evidence. As explained below, it is my opinion that the challenged claims of the '442 patent are not invalid under this standard.

16. Anticipation. It is my understanding that a claim of a patent is “anticipated” only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference. If even a single element is missing, the claim is not anticipated. It is my understanding that a feature is “inherent” in a reference only if that feature is necessarily present in the reference - not merely probably or possibly present. Furthermore, it is my understanding that in order to anticipate, a prior art reference must not only disclose all elements of the claim, but must also disclose those elements arranged as in the claim.

17. Obviousness. It is my understanding that a patent claim is invalid for “obviousness” if the claimed invention as a whole would have been obvious to one of ordinary skill in the art at the time of the invention in view of a single prior art reference or a combination of prior art references. I understand that a determination of whether a claimed invention would have been obvious requires taking into consideration factors which include assessing the scope and content of the prior art, the differences between the prior art and the claimed invention, and the level of ordinary skill in the art. It is my understanding that the level of ordinary skill in the art is determined by considering the type of problems encountered in the art, the prior art solutions to those problems, the rapidity with which innovations are made, the sophistication of the technology involved, and the educational level of active workers in the field. It is my understanding that a prior art reference may be considered to teach away from the invention when a person of ordinary skill, upon reading the reference, would be led in a direction divergent from the path that was taken by the claimed invention; and that the general rule is that references that teach away cannot serve to create a prima facie case of obviousness.

Level of Ordinary Skill in the Art

18. As of 1995, it is my opinion that a person of ordinary skill in the art to which the '442 patent pertains would likely be a person with a bachelors degree in computer science with ten to fifteen years of teaching or work experience in the field. Of course, one of ordinary skill in the art might have either more formal education with less experience, or less formal education with more experience.

Issues Involving the '442 Patent

19. It is my understanding that the Board, in a decision dated April 15, 2014, instituted a trial regarding the '442 patent for only the following: (a) whether claims 1, 2, 4, 23, 27 and 30 are unpatentable as obvious under 35 U.S.C. §103(a) over Woodhill (Ex. 1004 – U.S. Patent No. 5,649,196) and Francisco (Ex. 1005); (b) whether claim 7 is unpatentable as obvious under 35 U.S.C. §103(a) over Woodhill and Kahn (Ex. 1006 - U.S. Patent No. 6,135,646); and (c) whether claim 28 is unpatentable under §103(a) over Woodhill, Francisco, and Langer (Ex. 1007).

Claim Constructions

20. I have used the following interpretations of claim terms in my declaration. I have been asked to use these constructions for purposes of this proceeding.

21. *Preamble of claim 7 is limiting.* I have been told to assume that the preamble of claim 7 of the ‘442 patent is entitled to patentable weight in this proceeding.

22. *“file” and “data file” (claims 1, 2, 4, 7, 23, and 30).* The ordinary meaning of “file” in this art at the time of the invention was the following definition given in the 1994 Microsoft Computer Dictionary: **“file** A complete, named collection of information, such as a program, a set of data used by a program, or a user-created document. A file is the basic unit of storage that enables a computer to distinguish one set of information from another. A file might or might not be stored in human-readable form, but it is still the ‘glue’ that binds a conglomeration of instructions, numbers, words or images into a coherent unit that a user can retrieve, change, delete, save, or send to an output device.”

(Microsoft Press Computer Dictionary, Second Edition (1994) – Ex.

2011.) This is consistent with a “named” data item referred to at col.

5:48-49 of the ‘442 patent. A “named” data item as described in the ‘442 patent means that the data item is a coherent unit that is individually manipulatable by a user so that a user can selectively retrieve, change, delete, save, or send it to an output device. Thus, one of skill in the art would have understood that a “named” data item is one that is necessarily

a complete coherent unit that a user can selectively retrieve, change, delete, save, or send to an output device. Thus, the ordinary meaning of “file” as described in the 1994 Microsoft Computer Dictionary is both proper and consistent with the description of “file” in the ‘442 patent.

23. A mere *segment* of a basic file is NOT a “file.” Both the specification of the ‘442 patent and the 1994 Microsoft Computer Dictionary confirm this. The specification of the ‘442 patent *differentiates* between a “file” and a mere “segment” of a file, making clear that they are different and that a mere “segment” is not a “file.” As illustrated in Fig. 2 of the ‘442 patent, above, the ‘442 patent differentiates between a “*file*” 120 and a mere “*segment*” 122 of a basic file. (‘442 patent, Fig. 2; col 5:25-35.) The specification of the ‘442 patent explains, as illustrated in Fig. 2, that a “file” is at a different position in the hierarchy than a mere “segment” of a file. (‘442 patent, col. 5:25-30.) The ‘442 patent also explains that while a “file” can be named by a user, a mere “segment” of a file cannot be so named. (‘442 patent, col. 5:44-45; col. 8:25.) Regardless of how “file” is construed, the specification (which includes the drawings) of the ‘442 patent makes clear that a mere “segment” cannot be a “file.”
24. Woodhill’s binary objects and granules are not “files” – they are mere segments of basic files. According to the 1994 Microsoft Computer

Dictionary, a mere segment of a file cannot be a “file” at least because:

(a) a mere segment of a file is not a “complete” collection of information,
(b) a mere segment of a file is not a coherent unit that a user can retrieve,
change, delete, save or send to an output device, (c) a segment typically
has no file header, and (d) a mere segment of a file is not named.

(Microsoft Press Computer Dictionary, Second Edition (1994) – Ex.
2011.)

25. I note that this construction would NOT “preclude a data file from
including smaller data files” in all cases. For example, a *package* such as
a ZIP file may include a plurality of inner files – these would ALL be
“files” under this construction because each of these files would satisfy
the limitations of a “file” discussed above. Each inner file would be
individually named and each would satisfy the definition of a “file” in the
1994 Microsoft Computer Dictionary discussed above. Both before
formation of the ZIP file and after unpacking the ZIP file, each of the
inner files thereof would have a file header and each would be
individually manipulable by users, with each being a coherent unit that a
user can retrieve, change, delete, save or send to an output device. In
contrast, that is *not* the case with the binary objects and granules in
Woodhill. Binary objects and granules in Woodhill do not have headers,

are not individually manipulable as files by users, are not inner files of a package file (e.g., ZIP file), and do not satisfy the definition of a “file” in the 1994 Microsoft Computer Dictionary. A binary object or granule in Woodhill: (a) is determined solely based on size (number of bits) and is NOT based on the extent of any coherent unit that is named or individually manipulable; (b) has no header because it is based solely on size, (c) is NOT a coherent unit that a user can retrieve, change, delete, save or send to an output device, (d) is NOT a “complete” collection of information, and (e) is not named. The binary objects and granules in Woodhill, for example, are mere portions of a basic file, do not have headers, cannot be named, and are not complete coherent units that a user can selectively change, delete, save or send to an output device.

26. One of ordinary skill in the art at the time of the invention would also have understood that a “file” has a file header. The 1994 Microsoft Computer Dictionary explains the ordinary meaning that “*a file header identifies a data file by name, size, and time and date of creation or revision*”; it might also identify the program with which the file was created.” (Microsoft Press Computer Dictionary, Second Edition (1994) – Ex. 2011.) Accordingly, a “file” at the time of the invention, and in the ‘442 patent, must also be construed as including a file header. In contrast,

binary objects and granules in Woodhill do not have file headers, and thus cannot be “files” or “data files” for this additional reason.

27. The proper construction for a “data file” is the same as the construction for a “file” discussed above, except that a “data file” includes data.

28. ***“determining, using at least the name, whether a copy of the data file is present on a particular one of said computers” (claim 2).*** The ordinary meaning of “determine” is “to decide or settle conclusively and authoritatively.” *The American Heritage Dictionary* (1975). (Ex. 2020.) And the ordinary meaning of “whether” is “if it is so that; if the case is that.” *The American Heritage Dictionary* (1975). (Ex. 2020.)

Accordingly, this clause in claim 2 requires *deciding or settling conclusively and authoritatively, using at least the name, if it is so that a copy of the data file is present on a particular one of said computers.*

This is the proper construction for claim 2.

29. ***“determine whether unauthorized or unlicensed copies of some of the plurality of data files are present on a particular computer” (claim 23).***

As mentioned above, the ordinary meaning of “determine” is “to decide or settle conclusively and authoritatively”, and the ordinary meaning of “whether” is “if it is so that; if the case is that.” *The American Heritage Dictionary* (1975). (Ex. 2020.) Accordingly, this clause in claim 23

requires *decide or settle conclusively and authoritatively [using the list of file names]* if it is so that *unauthorized or unlicensed copies of some of the plurality of data files are present on a particular computer*. In order to be capable of *determining whether* a particular file, binary object, or granule is present on a given computer, there must be both (a) a capability to determine that it is present on that computer if that is the case, and (b) a capability to determine that it is not present on that computer if that is the case.

30. ***“file name” & “name for a data file” (claims 1, 2, 7, 23)***. A “file name” and “name for a data file” in this art are more than simply an identifier. It cannot reasonably be said that all identifiers are file names. For example, while a fingerprint is an identifier for a person, it is not a “name” for that person. Something is not a “file name” simply because it identifies a file. Any attempt to simply replace “file name” with “file identifier” (or “name” with “identifier”) in a construction would be incorrect and contrary to the claim language itself. Even Woodhill contrasts the difference between a file “name” and a binary object “identifier” (*see e.g.*, file name 40 and file name 80 in Figs. 3-4 of Woodhill, compared to binary object identifier 74 in Fig. 3 of Woodhill), evidencing that they are not necessarily the same thing in this art. One of

ordinary skill in the art at the time of the invention would have understood that there can be a significant difference between a “file name” (or “name for a data file”) and a mere “identifier.” “Name” should not be simply replaced with “identifier.” One of ordinary skill in the art would have recognized that a “file name” (or “name for a data file”) is used to both (a) identify a particular file that the name can be used to differentiate the file from other files, and (b) refer to that file, access that file, search for that file, and address that file so that a user can use the name to selectively retrieve, change, delete, save or send the file to an output device. (E.g., ‘442 patent, col. 16:13-16; col. 17:15-50; col. 19:1-27; col. 23:61-67; col. 24:56-67; col. 34:35-49; and Ex. 2011.) In contrast, identifiers relied upon in the cited art, such as Woodhill’s binary object identifiers 74 and “contents identifiers,” are not “file names” or “names for data files.”

The Challenged Claims Would Not Have Been Obvious

31. It is my opinion that the challenged claims of the ‘442 patent would not have been obvious for at least the following reasons.

Woodhill and Francisco both fail to disclose using a list of content-based file names to determine whether unauthorized

or unlicensed copies of some files are present on a particular computer as in claim 23.

32. Claim 23 of the '442 patent requires: *“using at least said list [of content-based file names] to determine whether unauthorized or unlicensed copies of some of the plurality of data files are present on a particular computer.”* Both Woodhill and Francisco fail to disclose this subject matter. Thus, even the alleged combination (which I do not agree with in any event) fails to meet the claimed subject matter of claim 23.

33. Woodhill certainly does not disclose this subject matter, and it is my understanding that Francisco is relied upon by petitioner in this respect. Indeed, Woodhill fails to disclose deciding or settling conclusively and authoritatively, using a list of binary object identifiers or contents identifiers (alleged file names), if it is so that unauthorized or unlicensed copies of a plurality of data files are present on a particular computer.¹

¹ With respect to Woodhill's granularization restore procedure for example, Woodhill is only analyzing whether a given granule is present in one particular large database file. Woodhill does not and cannot figure out if that granule is present in any of the other files on the local computer 20 or any other computer. The same applies to binary objects in Woodhill.

34. It is my understanding that petitioner thus argues that implementing Francisco's authorization check mechanism in Woodhill's distribution computer network system would have been an obvious improvement. However, Francisco also fails to disclose this claimed subject matter. Therefore, even the alleged combination fails to meet claim 23 for at least the following reasons.

35. **First**, petitioner's allegation on page 31 of the Petition that "Francisco discloses that a copy of a program residing on a computer being requested *can be determined to be unauthorized if there have been any modifications or changes in the software program*" is incorrect. Francisco's *authenticity* check at comparator 36 determines whether there have been any changes to the program – this is not an authorization check. Francisco's *authorization* check is performed later at comparator 40 *after* the authenticity check, and does not use a list of content based names. Francisco makes clear that determining whether changes have been made in Francisco is *not* an "authorization" check – it is an "authenticity check.

36. Francisco first performs an "authentication" check at comparator 36 to determine whether the requested program has changed relative to the base version. Then, "subsequent" to the authentication check, Francisco determines at comparator 40 whether the requesting user is authorized to

access the program. (Francisco at Fig. 2; col. 1:40-42; and col. 2:65 to col. 3:35.) The program identifier is compared with only a single value in comparator 36 during the “*authentication*” check, and is not compared with anything in comparator 40 during the “*authorization*” check.

37. Francisco’s *authentication* check at comparator 36 merely compares a program identifier (electronic identification indicia) 34 from generator 32 with a single program identifier 12. (Francisco at Fig. 2; col. 1:40-42; col. 2:65 to col. 3:20; and col. 4:12-28; and Mercer Dep. 60 [Ex. 2015].) Thus, there is no comparison of an identifier (12 or 34) to a list of content-based names in the authentication check at comparator 36. Moreover, comparator 36 merely performs an authentication check to make sure that the requested program has not changed – this is *not* an authorization check and does *not* determine whether the requested program is unauthorized or unlicensed.

38. If there is a match between identifiers 12 and 34 in comparator 36, then the requested program has not changed and it is only then that Francisco moves on to comparator 40 for determining whether the requesting user is authorized to access the program. (Francisco at Fig. 2; col. 3:17-36; and col. 4:30-51.) Francisco explains that the *authorization* check is performed “subsequent” to the *authentication* check. (Francisco at Fig.

2; and col. 1:40-42.) Because the authorization check at comparator 40 is “subsequent” to the authentication check, they are not the same and are different checks. In making the authorization check at comparator 40, Francisco compares an identification indicia of the requesting user with a user profile 42. (Francisco at Fig. 2; col. 3:20-36; col. 4:30-33; and col. 4:40-51.) Comparator 40 compares these two user indicia with each other in order to determine whether access is authorized. *Id.* See also Francisco at col. 4:48-51. Francisco’s comparator 40 does *not* compare program identifier 34 with anything – it merely forwards identifier 34 in a “program release” signal if the two user indicia match each other. No list of content-based file names is used at comparator 40 for any reason.

39. Accordingly, Francisco compares the program identifier (12 or 34) with a single value in comparator 36 to make sure that the file is *authentic* (*i.e.*, that it has not changed), and if the file is found to be authentic then Francisco later compares *user information* in another comparator 40 to determine if access is authorized. Thus, Francisco’s program identifier 34 (or 12): (i) is only compared to a single value for determining whether the program has changed, and (ii) is not compared to anything for determining whether access by a particular user is authorized. No list

of content-based file names is used during Francisco's authorization check at comparator 40.

40. The only possible alleged "list" of content-based file names in Francisco could be argued to be in the "electronic identification inicia library" in the upper right corner of Fig. 2 of Francisco. However, *nothing from this library is used during Francisco's authorization check at comparator 40*. Accordingly, even providing Francisco's system in Woodhill fails to meet the subject matter of claim 23. If Francisco's system were provided in Woodhill (with Woodhill's binary object identifiers 74 or records 58 possibly taking the place of Francisco's program identifiers), the invention of claim 23 still would not be met because the combination would NOT use a list of content-based file names to determine whether unauthorized or unlicensed copies of a plurality of data files are present on a particular computer.

41. **Second**, as explained above in the claim construction section, claim 23 requires deciding or settling conclusively and authoritatively, using a list of content-based file names, if it is so that unauthorized or unlicensed copies of a plurality of data files are present on a particular computer. It is my understanding that petitioner relies upon Francisco for this subject matter. In order to be capable of *determining whether* a particular file,

binary object, or granule is present on a given computer, there must be both (a) a capability to determine that it is present on that computer if that is the case, and (b) a capability to determine that it is NOT PRESENT on that computer if that is the case. Woodhill and Francisco are both incapable of at least the latter, because for a given file, binary object or granule they are only capable of analyzing information relating to *one file* on a given computer. They are blind regarding other files on that computer and elsewhere in the system. In Francisco (relied upon in this respect by petitioner), a single program (presumably on some server) is “selected” and requested to be released. (Francisco at col. 2:66-69 and Fig. 2.) Francisco then generates an identification signal 34 for that particular program and sends it to comparator 36. (Francisco at col. 2:66 to col. 3:3, and Fig. 2.) *Francisco does not analyze or generate any identification signal for any other programs on the server on which the selected program is presumably located.* Therefore, Francisco cannot determine whether a duplicate of that program, or a different and unlicensed version of that program, is located at a different location on that server. Because Francisco’s Fig. 2 procedure is blind to all programs on a given computer except for the one that is selected, Francisco is incapable of determining that a given program is *not present* on that

particular server if that is indeed the case. Unlike the '442 patent, Francisco (and Woodhill) does not search a list of content-based names for determining whether unauthorized or unlicensed files are present at a particular computer.

42. Assume a situation in Francisco where two versions of a given program are on a given server, where version A is licensed but version B is unlicensed. If a user selects version A of the program, then as explained above Francisco would do no analysis of version B and is incapable of figuring out that unlicensed version B is present on the server. Again, Francisco does not use a list of content-based names to determine whether unlicensed or unauthorized programs are on a particular server/computer.

One skilled in the art would not have modified Woodhill to add Francisco's or Kahn's access control system because Woodhill already has access control that is unrelated to binary object identifiers.

43. Given that Woodhill already has access control, there would have been no logical reason to have modified Woodhill to have provided further access control with Francisco or Kahn. One skill in the art, as a matter of common sense and logic which a person of ordinary skill in the art is presumed to have, would not think to layer additional and unnecessary

access control onto the existing access control of Woodhill. Thus, there would have been no motivation or logical reason for the alleged modification. Claims 1, 2, 4, 7, 23, 27, 28 and 30 define over the cited art, and are nonobvious, in this respect.

44. During the backup procedure, Woodhill separates each file being backed up into multiple data streams. (Woodhill, col. 7:41-43.) “These data streams may represent regular data, extended attribute data, *access control list data*, etc.” (Woodhill, col. 7:44-47.) Petitioner’s expert acknowledges that this “access control list data” relates to authorization checking in Woodhill. (Mercer Dep. 113, 115 [Ex. 2015].) Referring to Fig. 5A of Woodhill, it can be seen that this “access control list data” exists in step 132 *before* binary object identifiers 74 are even calculated in step 138. (Woodhill, Fig. 5A, col. 7:40-67.) Indeed, binary objects in Woodhill include such extended attribute data and access control list data, and Woodhill creates binary object identifiers 74 in part by hashing such extended attribute data and access control list data. (Woodhill, col. 7:41-67 and Fig. 5A.) Given that Woodhill’s “access control list data” exists *prior* to binary object identifiers 74 being created, it would be illogical and nonsensical to modify Woodhill as alleged by petitioner to use the binary object identifiers 74 for checking authorization/unauthorization.

Moreover, as Woodhill already had an authorization/deauthorization system based on this “access control list data” (which is unrelated to binary object identifiers 74), there would have been no logical reason to have modified Woodhill as alleged by petitioner to add another authorization/unauthorization checking feature to Woodhill.

One would not have modified Woodhill to check whether access to a file by a computer was authorized when the computer already had the current version of that file.

45. In Woodhill as allegedly modified by Francisco or Kahn, petitioner contends that the “request” of claims 1 and 7 is the “update request” described in Woodhill at column 17:42-46. In Woodhill’s restore procedure, the local computer 20 *already has* the most recent (current) version of the large database file at issue. (Woodhill at col. 17:39-40, col. 17:59-60.) Because the local computer 20 already has the most recent (current) version of the large database file at issue, there is no logical reason to have modified Woodhill to implement a system for checking whether that same local computer 20 is authorized to access a previous version of that same file. One would have assumed that the local computer 20 is permitted to access a prior version of a file if that computer already has the current version of that file. Given that the local computer 20 already has the most recent version of the file, as a matter of

common sense there would have been no logical reason to have checked whether that same computer is authorized to access the previous version of that same file.

46. Petitioner's alleged reason for making the alleged modification is also fundamentally flawed. It is my understanding that petitioner's rationale for the alleged modification is "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." However, the local computer in Woodhill's granularization restore procedure already has the most recent version of this file, so there would be no concern of that computer/user not having authorization. Moreover, Woodhill already has an access control system. (Woodhill, col. 7:44-47.) If Woodhill already has an access control system, there would have been no logical reason to have added another one. Indeed, the fact that Woodhill already had an access control system would NOT have motivated one of ordinary skill in the art to have added another one to Woodhill.

47. Woodhill also describes a "self-audit" procedure at col. 18:10-38. There also would have been no logical reason to have modified that procedure. Woodhill's auditing procedure is for performing "self-audits" of the

system. (Woodhill at col. 18:10-12.) Given that the system already has the binary object, as a matter of common sense there would have been no logical reason in a “self” auditing system to have checked whether the system is authorized to access that binary object which it already has. Moreover, there are no potential unauthorized users involved in Woodhill’s self-audit procedure. Still further, there is no “request” in the self-auditing procedure that includes a binary object identifier 74 or a contents identifier – so claims 1 and 7 would not be met in any event.

Francisco and Woodhill teach away from claim 23.

48. It is my understanding that a prior art reference may be considered to *teach away* when a person of ordinary skill, upon reading the reference, would be led in a direction divergent from the path that was taken by the patent. Here, Francisco and Woodhill both would have led one of ordinary skill in the art away from the path taken by the inventors. While both Woodhill and Francisco have content-based identifiers and both have access control systems, neither Woodhill nor Francisco compares their content-based identifiers to anything for determining whether access is authorized or for searching for unauthorized copies. And neither Woodhill nor Francisco uses a list of content-based names for

determining whether unauthorized or unlicensed copies of files exist on a particular computer.

49. Francisco's identification "indicia" 12 and 34 are content-based.

(Francisco at Fig. 2; col. 2:29-42; and col. 3:49 to col. 4:17.) Francisco's *authentication* check at comparator 36 compares a program identifier (electronic identification indicia) 34 from generator 32 with a single program identifier 12. (Francisco at Fig. 2; col. 1:40-42; col. 2:65 to col. 3:20; and col. 4:12-28.) Comparator 36 merely performs an authentication check to make sure that the requested program has not changed – this is not a determination of whether a particular user is authorized to access the program. If there is a match between 12 and 34 in comparator 36, then the requested program has not changed and it is only then that Francisco moves on to comparator 40 for determining whether the requesting user is authorized to access the program.

(Francisco at Fig. 2; col. 3:17-36; and col. 4:30-51.) Francisco clearly explains that the *authorization* check is performed "subsequent" to the *authentication* check. (Francisco at Fig. 2; and col. 1:40-42.) In making the authorization check at comparator 40, Francisco compares an identification indicia of the requesting user with a user profile 42.

(Francisco at Fig. 2; col. 3:20-36; col. 4:30-33; and col. 4:40-51.)

Comparator 40 compares these two user indicia with each other in order to determine whether the requesting user is authorized to access the program. *Id.* See also Francisco at col. 4:48-51. Francisco's comparator 40 does *not* compare program identifier 34 (or 12) with anything in making the authorization check. There is no list of alleged content-based names used in Francisco's authorization check at comparator 40 – contrary to claim 23.

50. Woodhill's binary object identifiers 74 are also content based.

(Woodhill at col. 8:22-31.) But Woodhill does not use the binary object identifiers 74 in its access control system. During the backup procedure, Woodhill separates each file being backed up into multiple data streams. (Woodhill, col. 7:41-43.) "These data streams may represent regular data, extended attribute data, *access control list data*, etc." (Woodhill, col. 7:44-47.) Referring to Fig. 5A of Woodhill, it can be seen that this "access control list data" exists in step 132 *before* binary object identifiers 74 are even calculated in step 138. (Woodhill, Fig. 5A, col. 7:40-67.) Indeed, binary objects in Woodhill include such extended attribute data and access control list data, and Woodhill creates binary object identifiers 74 in part by hashing such extended attribute data and access control list data. (Woodhill, col. 7:41-67 and Fig. 5A.) Given that

Woodhill's "access control list data" exists *prior* to binary object identifiers 74 even being created, Woodhill cannot use the binary object identifiers 74 in its access control system and teaches away from doing so. Again, there is no list of alleged content-based names used in any authorization check in Woodhill.

51. Therefore, while both Woodhill and Francisco have content-based identifiers and both have access control features, neither Woodhill nor Francisco compares their content-based identifiers to anything for determining whether access is authorized, or for determining whether unauthorized or unlicensed copies of files exist on a particular computer. These references teach away from claim 23.

**Prior art elements are not used for their intended purpose
regarding claim 23.**

52. Woodhill's binary object identifiers 74 and "contents identifiers", and Francisco's identifiers 12 and 34, are NOT used for comparisons to determine whether access is authorized. And, unlike claim 23, there is no list of such content-based identifiers in either Woodhill or Francisco that is used to determine whether unauthorized or unlicensed copies of files exist on a particular computer. Thus, any modification to meet the challenged claims would result in these identifiers NOT performing the same functions that they are known to perform in Woodhill and

Francisco. This is strong evidence of nonobviousness in my opinion. While both Woodhill and Francisco use content-based identifiers for various purposes, neither compares such identifiers to determine whether access to a file is authorized. And neither uses a list of such content-based identifiers to determine whether unauthorized or unlicensed copies of files exist on a particular computer. The inventors of claim 23 conceived of a new and nonobvious use for substantially unique content-based identifiers which is not disclosed or suggested in the relied upon art.

Granules corresponding to “contents identifiers” in the alleged request in Woodhill are never provided to the requesting computer in response to the request (claims 1 and 7).

53. Claims 1 and 7 each require a “request” for a data file, which request includes the name of that particular data file. Claims 1 and 7 further require that, in response to the request, a copy of the requested file is provided from a given one of the computers and (a) only to licensed parties (claim 1), or (b) is not provided to unlicensed or unauthorized parties (claim 7). This subject matter cannot be met if it is argued that “files” are granules and “names” are “contents identifiers” with respect to Woodhill’s granularization restore procedure (regardless of whether it is modified).

54. In the alleged Woodhill combination, with Francisco or Kahn, granules corresponding to “contents identifiers” in the alleged request (“update request”) are NEVER provided from the remote backup server 12 to the requesting computer in response to the “update request.” (Woodhill at col. 17:18-67.) Thus, Woodhill/Francisco and Woodhill/Kahn cannot meet claims 1 and 7 if the claimed data files are argued to be granules and the names are considered “contents identifiers.” Moreover, no name of the large database file at issue is included in the “update request,” so that the claimed data file in claims 1 and 7 cannot be considered to be the large database file.

55. In Woodhill, and in any modification of Woodhill/Francisco or Woodhill/Kahn, granules corresponding to “contents identifiers” in the alleged request (“update request”) are never provided to the requesting computer in response to the update request. (Woodhill at Fig. 5I and col. 17:18-67.) In the relied upon restore procedure, no granule having an identifier in the “update request” is ever provided to the requesting local computer 20 (alleged licensed party). The “contents identifiers” in the alleged “request” were calculated in step 444 and thus are NOT for granules that are being requested – the requesting local computer 20 already has the granules corresponding to the “contents identifiers”

calculated in step 444 and therefore they are never provided to the local computer even in the allegedly modified system. And because granules identified in the request are never sent from remote server 12 to the local computer in the restore procedure (the local computer already has them), it would not make sense to use identifiers for such granules to determine whether the system is authorized to provide such granules to the local computer. Thus, in addition to all other reasons set forth above, such combinations (which I disagree with in any event) fail to meet claims 1 and 7.

56. Turning to the “contents identifiers” for the reconstituted granules in Woodhill, the “contents identifiers” for the granules reconstituted at col. 17:48-50 are only used for comparison in box 450 of Fig. 5i *after* the Fig. 5h reconstitution of those granules, and are *never* included in any “request” regarding a data item.

Binary objects corresponding to Record 58 allegedly in the “update request” are not provided to the requesting computer in response to the “update request” (claims 1 and 7).

57. In the alleged Woodhill/Francisco and Woodhill/Kahn combinations binary objects corresponding to the binary object identification records 58 allegedly in the “update request” are NOT provided to the requesting

computer in response to the update request. (Woodhill at col. 17:18 to col. 18:9.) The granularization restore procedure in Woodhill as allegedly modified by Francisco or Kahn transmits select granules – not binary objects. And record 58 and binary object identifier 74 do not identify particular granules. Thus, Woodhill/Francisco and Woodhill/Kahn cannot meet claims 1 and 7 if, as alleged by petitioner, the claimed data items are binary objects and the identifiers are considered binary object identifiers 74 or binary object identification records 58.

58. As explained above, claims 1 and 7 each require a “request” for a data file, which request includes the name of that particular data file. Claims 1 and 7 further require that, in response to the request, a copy of the requested file is provided from a given one of the computers and (a) only to licensed parties (claim 1), or (b) is not provided to unlicensed or unauthorized parties (claim 7). This subject matter cannot be met if the claimed files are considered to be Woodhill’s binary objects and the identifiers to be Woodhill’s binary object identifiers 74 or records 58 as argued by petitioner, because the binary objects corresponding to the identifier allegedly in the “request” are not disclosed as being transmitted

to the local computer 20 in response to the request in Woodhill's granularization restore procedure.

59. In Woodhill, and in any modification of Woodhill, binary objects corresponding to binary object identification records 58 allegedly in the request ("update request") are *not* provided to the requesting computer in response to the request. (Woodhill at Fig. 5I and col. 17:18-67.)

Woodhill never discloses such an occurrence, and instead indicates to the contrary. The restore procedure transmits only particular granules based on whether they have changed – not binary objects. I also note that the granularization restore procedure in Woodhill is different than the self-audit procedure in that the different procedures transmit different items, use different comparisons, and so forth. The two procedures cannot be combined in any logical manner.

60. Thus, in the granularization process upon which petitioner relies, Woodhill is not designed to provide, and *never* provides, any binary object to the local computer 20 in response to the "update request." No binary object sequence of bits is ever sent from the backup file server 12 to a local computer 20 in the restore procedure of Woodhill's granularization feature (regardless of whether modified by Francisco/Kahn). In Woodhill's granularization process, the backup file

server 12 sporadically sends only select “granules” of a large database file to the local computer. (Woodhill, col. 17:60-65.) Binary object identifiers 74 and records 58 do not identify granules, and are not names for granules. Woodhill’s granularization embodiment does *not* send large database files or binary objects from server 12 to the local computer – just select granules.

**Binary objects and granules are not “data files” as required
by the challenged claims.**

61. The binary objects and granules in Woodhill are not “files” or “data files,” for the reasons explained above in the claim construction section.
62. It is my understanding that petitioner relies on a “granule” or “binary object” from Woodhill as the alleged “data file” with respect to all challenged claims. Because these are not data files as explained above in the claim construction section, petitioner’s allegations regarding all challenged claims are incorrect.
63. Claim 1 for example states as follows, *inter alia*: “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*; 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 . . .” There is no “request” for a “data file” in Woodhill’s granularization restore procedure relied upon by petitioner. Neither a binary object nor a granule is a “data file.” Woodhill’s “update request,” relied upon by petitioner, does not contain a name for a requested “data file” as required by claims 1 and 7.

64. Woodhill’s granularization embodiment relates to *large* database files each of which contains *many* binary objects and *many* granules (Woodhill, col. 14:52-67; col. 15:11-17; col. 17:18-22; col. 17:31-35.) Petitioner relies upon the “update request” described by Woodhill at col. 17:42. However, that “update request” does not request a large database file (the only type of file handled by Woodhill’s granularization embodiment) or any other “data file,” and does not contain any name for a large database file. An alleged request for a binary object is not the same as a request for a large database file. Accordingly, there is not even a “request” for a “data file” in Woodhill’s granularization restore procedure, and there is no such request that includes a name for a requested “data file.” Woodhill’s granularization restore at column 17 sends particular granules (not binary objects, and not an entire large database file) to a local computer. (Woodhill, col. 17:50-64.) Indeed, a purpose of Woodhill’s granularization embodiment is to avoid

transmitting an entire large database file, which teaches away from petitioner's allegations regarding claims 1 and 7. Accordingly, even the alleged combinations fail to meet claims 1 and 7.

65. Furthermore, while Woodhill mentions Binary Object Identification Record 58 at col. 17:44, Woodhill indicates that portions (62 and 72) of the Record 58 other than the binary object identifier 74 are used to identify a given binary object in a message. (Woodhill, col. 11:65 to col. 12:5.) Woodhill indicates in this respect that binary object identifiers 74 are *not* used in any alleged "request" for a binary object. (Woodhill, col. 11:65 to col. 12:5.) Woodhill indicates that binary object identifiers 74 are *not* used for accessing binary objects. *Id.*

66. Petitioner briefly mentions Woodhill's "self-audit" embodiment described at col. 18:10-38. (Ex. 1010, pg. 32.) However, Woodhill's self-audit embodiment also fails to disclose or suggest the subject matter of claims 1 and 7, regardless of whether modified by Francisco or Kahn. (Woodhill, col. 18:10-38.) There is no "request" in Woodhill's self-audit procedure that includes a binary object identifier 74 (let alone a name for a "data file"). Woodhill explains in the description of the self-auditing procedure that the Record 58 is "stored in File Database 25." (Woodhill, col. 18:16-19.) While Woodhill states that the Record 58 is "stored in

File Database 25”, the Record 58 is never described as part of a “request” for the binary object. (Woodhill, col. 18:16-19.) Moreover, Woodhill elsewhere indicates that, in a message, portions of the Record 58 other than the binary object identifier 74 are used to identify a binary object. (Woodhill, col. 11:65 to col. 12:5.) Accordingly, in Woodhill’s self-auditing procedure there is no description of a “request” for a binary object or anything else that includes a binary object identifier 74. To the contrary, Woodhill indicates that any such “request” would not include a binary object identifier 74. (E.g., Woodhill, col. 11:65 to col. 12:5.) Additionally, there is no request for a “data file” in Woodhill’s self-audit procedure as explained above. Woodhill restores a randomly selected binary object – not a “data file.” (Woodhill, col. 18:25.) A binary object is not a “data file.” And neither a binary object identifier 74 nor a binary object identification record 58 is a “name” for a data file. Moreover, because Woodhill’s “self-audit” embodiment is an automatic procedure performed by the system on a periodic basis to check itself, it is entirely unrelated to the final lines of claims 1 and 7 (e.g., “a copy of the requested file is only provided to licensed parties”). Since Woodhill’s “self-audit” embodiment is an automatic procedure for checking itself,

there would have been no logical reason to have implemented an authorization checking system in Woodhill's "self-audit" procedure.

67. Furthermore, it would be incorrect to switch back and forth between Woodhill's granularization embodiment, and Woodhill's "self-audit" embodiment, in an attempt to mix and match to meet all elements of a given claim. These are different embodiments that function separately on different bases. It is my understanding that in order to meet a claim a reference must not only disclose the elements of the claim, but must also disclose those elements "arranged or combined in the same way as in the claim."

Binary objects and granules cannot represent a "digital image" as required by claim 7.

68. Claim 7 requires a "request" for a data file, which request includes the name of that particular data file. Claim 7 further requires that, in response to the request, a copy of the requested file is provided from a given one of the computers. Petitioner relies on the "update request" in Woodhill's granularization restore procedure for the alleged "request," and contends that the granules represent "data files" and that the "contents identifiers" represent names for those alleged "files." (Ex. 1010 at 32-33 and 54.)

69. However, claim 7 expressly recites “the contents of said data file representing a *digital image*.” This cannot occur with a granule (or a binary object) in Woodhill, because a granule (or binary object) in Woodhill is only a small portion of a large database file and cannot itself represent a digital image.
70. Woodhill’s granularization restore procedure which petitioner seeks to modify is used for processing *large database files* each of which includes many binary objects and many granules. (‘442 patent at col. 14:52-65, col. 17:18-35.) In 1995, one of ordinary skill in the art would have expected a “large database file” of the type handled by Woodhill’s granularization procedure to be at least 1 to 2 gigabytes in size. Woodhill explains that a binary object is 1 megabyte in size, and that a granule can be 1 kilobyte. (Woodhill, col. 7:54-55, col. 14:58-67.) Given such sizes, a single 1 gigabyte large database file to be processed by Woodhill’s granularization restore procedure would include at least 1,024 binary objects and 1,048,576 granules. Thus, each granule is a small portion of a binary object, and each binary object is a small portion of the large database file. While a large database file in Woodhill could possibly itself represent a “digital image,” a small segment thereof such as a granule or binary object could not. Each granule and binary object could

possibly represent a small portion of the data for a digital image, but could not itself represent a digital image as required by claim 7.

Accordingly, a granule or binary object in Woodhill cannot be the type of “file” (a “file” that represents a digital image) called for in claim 7.

Accordingly, even the alleged modification to Woodhill fails to meet claim 7 because a granule or binary object cannot represent a “digital image” as called for in claim 7.

71. Additionally, any alleged modification to Woodhill’s self-audit procedure also cannot meet claim 7, because again a binary object (a portion of an overall file in Woodhill) cannot itself represent a “digital image.” Moreover, there is also no “request” for a data file in Woodhill’s self-audit procedure that includes a binary object identifier 74 of a particular binary object. And, as explained above, there are no unlicensed or unauthorized parties involved in Woodhill’s “self” audit procedure, so the claim also cannot be met in this respect and there would have been no logical reason to have provided an authorization check in a “self” audit procedure.

72. Furthermore, no “name” of the large database file is included in Woodhill’s “update request.” Therefore, the claimed data file called for in claim 7 cannot be considered the large database file in Woodhill/Kahn

because no name therefor is provided in any alleged “request for the data file” as required by claim 7.

The alleged identifier in the argued combinations is not a “file name” and is not used for accessing.

73. As explained above in the claim construction section, one of ordinary skill in the art would have recognized that a “file name” and “name for a data file” are used to both (a) identify a particular data item so that the name can be used to differentiate the data item from other data items, and (b) refer to that data item, access that data item, search for that data item, and address that data item so that a user can use the name to selectively retrieve, change, delete, save or send the item to an output device. (E.g., ‘442 patent, col. 16:13-16; col. 17:15-50; col. 19:1-27; col. 23:61-67; col. 24:56-67; col. 34:35-49; and Ex. 2011.)

74. The alleged identifiers in the Woodhill/Francisco and Woodhill/Kahn combinations are NOT “file names” or “names for data files.” Woodhill’s binary object identifiers 74 are not “file names” for binary objects, and Woodhill’s contents identifiers are not “file names” for granules. And Francisco’s program identifiers are also not “file names” for data files. Because these are not “file names” or “names for data files,” even the alleged modification fails to meet at least claims 1, 11, 24 and 32.

75. Woodhill uses “names” for files (see name 40 in Fig. 3 of Woodhill), but does *not* provide names for binary objects or granules. Woodhill merely employs a hash to compare same file components to determine whether there has been a change – Woodhill does not use a hash as a file name, or as a substitute for a name, in a computer network. Indeed, Woodhill’s granularization process upon which petitioner relies relates to a large database file that has many binary objects and many granules – these granules and binary objects of such a large database file are not provided with file names in Woodhill’s granularization process and there would have been no logical reason to have done so.

76. Something is not a “file name” simply because it is an “identifier.” For example, while a fingerprint is an identifier for a person, it is not a “name” for that person. One of ordinary skill in the art at the time of the invention would have understood that there is a significant difference between a “file name” and a mere “identifier” as explained above.

77. Woodhill’s “binary object identifiers 74” are not used to access binary objects, search for binary objects, or address binary objects. Binary object identifiers 74 are not even mentioned in columns 16-17 of Woodhill. When a Binary Object Identification Record 58 is provided such as at col. 17:44, and a binary object is accessed, Woodhill expressly

explains how this is done and it does NOT involve any binary object identifier 74. (Woodhill at col. 11:65 to col. 12:7, and col. 9:18-20.) Thus, Woodhill itself describes that fields 62 and 72 of Binary Object Identification Record 58 are used to access binary objects (i.e., the field for binary object identifier 74 is NOT used). Thus, binary object identifiers 74 are not compared to anything or otherwise used in any respect when accessing, searching for, or addressing binary objects in the allegedly modified system. They are not “file names” or “names for data files.”

78. Likewise, Woodhill’s “contents identifiers” are not used to access granules, search for granules, or address granules. For instance, in Woodhill’s procedure for reconstituting a binary object on a granule-by-granule basis, “contents identifiers” are not used or even mentioned. (Woodhill at Fig. 5H and col. 16:16 to col. 17:17.) This same procedure is used in the restore procedure. (Woodhill at col. 17:46-50.)

79. Therefore, the “file name” or “name for a data file” aspect of the challenged claims is not disclosed by Woodhill as allegedly modified by Francisco or Kahn.

80. Additionally, while Woodhill mentions names in its claims (*e.g.*, Woodhill at col. 22:3-9; col. 22:35-50; col. 22:67 to col. 23:2; col. 23:10-

19; and col. 24:24-39), this is not “prior art” to the challenged claims.

The “name” language in Woodhill’s issued claims is not “prior art” to the challenged claims. The subject matter added to Woodhill on January 5, 1996 (including the “name” language in Woodhill’s claims) was added after the April 11, 1995 effective filing date of the ‘442 patent, was not in Woodhill’s originally-filed specification, and is not “prior art” to the challenged claims of the ‘442 patent (see prosecution history of Woodhill at Ex. 2007). Woodhill’s binary object identifiers and “contents identifiers” are not names for binary objects or granules. Binary objects and granules in Woodhill are not files and thus do not have file names.

One of ordinary skill in the art reviewing Woodhill’s specification as originally filed would have understood that binary object identifiers 74 were not “names” or “file names” for binary objects, and that “contents identifiers” were not “names” or “file names” for granules. Many identifiers are not names. For example, while a fingerprint is an identifier for a person, it is not a “name” for that person. It is my understanding that subject matter added to Woodhill after the effective filing date of the challenged claims is not “prior art” and cannot be used against the challenged claims in this proceeding.

MD4, MD5 and SHA of Claim 28 Would not have been obvious.

81. Claim 28 requires an MD4, MD5 or SHA. It is my understanding that petitioner does not contend that MD4 or SHA would have been obvious to have used in Woodhill or Francisco. However, I understand that petitioner does contend that it would have been obvious to have modified Woodhill/Francisco to use an MD5 based on Langer.
82. It is my opinion that it would *not* have been obvious to have calculated Woodhill's Binary Object Identifier 74 (or "contents identifier") using an MD5 hash function. There were thousands and thousands of hashes known in the art at the time of the invention. There would have been no logical reason to have selected an "MD5", out of the many thousands of possible hashes, for use in Woodhill. An MD5 is a non-reversible cryptographic hash – Woodhill's alleged hashes in column 8 are not. Moreover, an MD5 produces a 16 byte hash value, whereas Woodhill desires a 4 byte hash value. (Woodhill, col. 8:1-3.) Thus, an MD5 would not be compatible with Woodhill's system (whereas many other available hashes would have been), and there would have been no logical reason to have selected it from the large number of possible other hash functions that were available at the time of the invention.

MISCELLANEOUS

83. In signing this declaration, I recognize that the declaration will be filed as evidence in a contested proceeding 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 proceeding 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.

84. I reserve the right to supplement my opinions in the future to respond to any new arguments that petitioner may raise and take into account new information as it becomes available to me.

85. The prior declaration I signed in this IPR proceeding was signed/executed on July 15, 2014. 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, and may jeopardize the validity of any patent from this proceeding.

(signature on next page)

Date: 8/3/2014

Robert B.K. Dewar

Robert B.K. Dewar