

**AREE WITHIN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
SHERMAN DIVISION**

**INTELLECTUAL VENTURES I LLC AND
INTELLECTUAL VENTURES II LLC,**

Plaintiff,

vs.

AMERICAN AIRLINES, INC.,

Defendant.

Civil Action No. 4:24-cv-00980

JURY TRIAL

DEFENDANT AMERICAN AIRLINES, INC.'S P.R. 3-3 INVALIDITY CONTENTIONS

Pursuant to P.R. 3-3, P.R. 3-4, and the Court's Scheduling Order in the above-captioned case (Dkt. 44, 49), Defendant American Airlines, Inc. ("American") serves its Preliminary Invalidity Contentions with respect to the claims identified by Plaintiff Intellectual Ventures I LLC and Intellectual Ventures II LLC (collectively "IV") in its Disclosure of Asserted Claims and Infringement Contentions, served May 23, 2025 ("Infringement Contentions").

These Preliminary Invalidity Contentions are limited to the claims asserted from the six patents identified in IV's Original Complaint (Dkt. 1). The Court has not yet ruled on IV's Opposed Motion for Leave to File an Amended Complaint. *See* Dkt. 48. If the Court grants IV's Motion, American reserves the right to supplement these contentions to raise invalidity and/or unenforceability contentions directed to any new or additional asserted claims.

I. PRELIMINARY STATEMENT AND RESERVATION OF RIGHTS

A. The Asserted Claims

IV asserts the following patents, claims, and priority dates in its Infringement Contentions, served May 23, 2025.

Asserted Patent	Asserted Claims	Asserted Priority Date(s)
'844 patent	7 and 11	December 30, 2004
'722 patent	14, 16, 17, and 18	December 18, 2000
'785 patent	30, 35, and 37	March 31, 2003
'326 patent	1, 4, and 18	January 12, 2004
'469 patent	24, 25, 26, 28, and 32	September 29, 2003
'582 patent	1, 2, 5, and 7	March 13, 2002

The patents identified in the foregoing table are collectively referred to as the “Asserted Patents,” and the claims identified in the foregoing table are collectively referred to as the “Asserted Claims.” Any reference to an “asserted priority date” in these Preliminary Invalidity Contentions refers to the “Asserted Priority Dates” identified in the foregoing table, and further defined below.

American contends that each of the Asserted Claims is invalid under at least one or more of 35 U.S.C. §§ 101, 102, 103, or 112. Pursuant to the Local Patent Rules, American does not provide any contentions regarding any claims not asserted by IV. To the extent that the Court permits IV to assert additional claims (or patents) against American in the future, American reserves all rights to amend or supplement these Preliminary Invalidity Contentions or to otherwise disclose new or supplemental invalidity contentions regarding such claims. Furthermore, because discovery is ongoing,¹ American reserves the right to revise, amend,

¹ Defendant’s ongoing efforts include but are not limited to: serving subpoenas on prior artists and inventors regarding prior art and claim scope, seeking additional information related to the references and systems disclosed in these Preliminary Invalidity Contentions, and seeking additional information related to available prior art systems, as well as IV’s Infringement

and/or supplement the information provided herein, including identifying, charting, and relying on additional references, should discovery yield additional information or references.²

American further reserves the right to amend these contentions in response to any claim construction rulings, as permitted by the Local Patent Rules or with permission of the Court.

The Infringement Contentions are deficient in multiple respects and do not provide American with sufficient information to understand the specific accused features and components and the alleged factual and evidentiary bases for IV's infringement allegations. Among other things, the Infringement Contentions lack the specificity required by P.R. 3-1, fail to properly identify accused instrumentalities, and fail to explain adequately IV's infringement theories for numerous limitations. IV has prejudiced American's ability to understand, for purposes of preparing these Preliminary Invalidity Contentions, what IV alleges to be the scope of the Asserted Claims. If IV modifies any assertion or contention in its Infringement Contentions, or presents any new assertion or contention relevant to these Preliminary Invalidity Contentions to the extent allowed by the Local Patent Rules or the Court, American reserves the right to supplement or otherwise amend these initial Invalidity Contentions.

B. Priority Date of the Patents-in-Suit

IV's Infringement Contentions contain allegations regarding the priority date to which IV alleges it is entitled for each of the Asserted Claims. American does not agree that IV is entitled

Contentions and the products accused of infringing therein. No depositions have been taken as of this time, including, without limitation, depositions of any inventors, authors, or entities listed on any references or systems identified in these Preliminary Invalidity Contentions. Further, Defendant reserves the right to review and supplement these Preliminary Invalidity Contentions with respect to any additional prior art that becomes apparent as discovery proceeds.

² IV has failed to respond to American's Interrogatory No. 5, seeking the identification of all allegations by any Person or Entity that any claim of the Patents-in-Suit or any Related Patent is invalid or unenforceable, and Interrogatory No. 6, seeking all prior art of which IV is aware or that anyone has ever contended or asserted to be prior art to the Asserted Claims. American incorporates by reference any responses to those interrogatories IV may provide in the future.

to the Asserted Priority Dates for each of the Asserted Claims, as IV has failed to prove it is entitled to its Asserted Priority Dates. Furthermore, IV has failed to meaningfully respond to American's Interrogatory No. 7, seeking details regarding the conception and reduction to practice of the Asserted Claims.

Any reference to an "asserted priority date" in these Preliminary Invalidity Contentions refers to the priority dates identified in IV's Infringement Contentions. Reference to a "priority date" or an "asserted priority date" should not be construed to mean that American agrees that any of the Asserted Patents are in fact entitled to such priority date, or that IV has provided proper notice as to its contentions for a priority date.³ To the extent IV alleges that any prior art relied on in these Preliminary Invalidity Contentions does not actually qualify as prior art to an Asserted Patent, American reserves the right to rebut those allegations (e.g., by demonstrating an earlier priority date for the challenged prior art and/or a later priority date for a particular Asserted Patent and/or Asserted Claim).⁴ Likewise, to the extent IV successfully establishes an invention date before any of the prior-art references relied on by American, then those references serve as evidence of obviousness, particularly, contemporaneous invention by others.

C. Claim Construction

American's Preliminary Invalidity Contentions are based on (1) American's present understanding of the Asserted Claims, (2) the claim constructions IV appears to be proposing based on the Infringement Contentions, all without regard to whether American agrees with IV's apparent or expressed claim constructions. American reserves the right to supplement or otherwise amend these Preliminary Invalidity Contentions in response to any proposed claim

³ American reserves the right to rely on additional documents and evidence, including without limitation the documents cited in IV's Infringement Contentions, in the event that IV fails to establish that any Asserted Claim of any Asserted Patent is entitled to its Asserted Priority Date.

⁴ American reserves the right to rely on additional documents and evidence to rebut any efforts by IV to allege any reference was not publicly available or otherwise available as prior art.

constructions or alleged supporting evidence offered by IV, any report from any expert witness for IV regarding claim construction issues, any claim construction briefing filed by IV, and any position taken by IV concerning claim construction, infringement, or invalidity.

American takes no position on any matter of claim construction in these Preliminary Invalidity Contentions. If American's apparent claim constructions herein are consistent with any explicit, apparent, or implied claim constructions in the Infringement Contentions, no inference is intended and no inference should be drawn that American agrees with any of IV's claim constructions. Any statement herein describing or tending to describe any claim element is provided solely for the purpose of understanding and/or applying the cited prior art. American expressly reserves the right (1) to propose any claim construction American considers appropriate, (2) to contest any claim construction proposed by IV that American considers inappropriate or inaccurate, and/or (3) to take positions with respect to claim construction issues that are inconsistent with, or even contradictory to, claim construction positions expressed or implied in these Preliminary Invalidity Contentions.

Prior art not included in these Preliminary Invalidity Contentions, whether now known to American, might become relevant depending on the claim constructions proposed by American and/or the Court's claim construction rulings. American reserves all rights to supplement or modify the positions and information in these Preliminary Invalidity Contentions, including without limitation the prior art and grounds of invalidity set forth herein, pursuant to P.R. 3-6 after the Court has construed the Asserted Claims.

D. Ongoing Discovery and Supplementation

American's investigation, including its investigation of prior art and grounds for invalidity, is ongoing. Furthermore, American's invalidity positions will be the subject of expert testimony. American bases these Preliminary Invalidity Contentions on its current knowledge

and understanding of the Asserted Claims, IV's Infringement Contentions, the prior art, systems, and other facts and information available as of the date of these contentions.

American reserves the right to supplement these Preliminary Invalidity Contentions, including, without limitation, by adding additional prior art and grounds of invalidity in accordance with the Federal Rules of Civil Procedure, the Local Rules, the Local Patent Rules, the Docket Control Order, any Order issued by this Court, or otherwise.

E. Prior Art Identification and Citations Thereto

In these Preliminary Invalidity Contentions, American identifies specific portions of prior art references that disclose the elements of the Asserted Claims. While American has identified exemplary prior art references for each element, they do not necessarily identify every disclosure of the same element in each prior art reference. A person of ordinary skill in the art would read a prior art reference as a whole and in the context of other publications, literature, and general knowledge in the field and would rely upon other information including other publications and general scientific or engineering knowledge. American therefore reserves the right to rely upon other unidentified portions of the prior art references and on other publications and prior art products and expert testimony to provide context and to aid understanding and interpretation of the identified portions of the prior art.

American also reserves the right to rely upon (1) other portions of the cited prior art references, other publications, prior art products, and the testimony of experts to establish that the alleged inventions would have been obvious to a person of ordinary skill in the art, including on the basis of modifying or combining certain cited references; (2) all versions of a cited prior art publication; (3) admissions relating to prior art in the Asserted Patents or related patents, the prosecution history of the Asserted Patents or related patents, or other admissions obtained

during discovery; and (4) foreign counterparts of any U.S. patents identified in American's Preliminary Invalidity Contentions.

Where American identifies a particular figure in a prior art reference, the identification should be understood to encompass the caption and description of the figure as well as any text relating to the figure in the remainder of the prior art reference (e.g., for patent references, text in the specification and prosecution history). Similarly, where an identified portion of text refers to a figure or other material, the identification of the text should be understood to include the referenced figure or other material.

All related prior art references that are themselves subparts or related documents of a broader set of documents should be considered a single prior art reference, as that is how a person of ordinary skill in the art would consider such related references.

F. Invalidity, Unenforceability and/or Ineligibility Based On Non-Required Disclosure(s)

American reserves the right to prove the invalidity, unenforceability and/or ineligibility of one or more of the Asserted Claims on bases other than those required to be disclosed in these disclosures pursuant to P.R. 3-3.

G. No Patentable Weight

American reserves the right to argue that various portions of the Asserted Claims, such as an intended use or result, non-functional descriptive material, and certain preamble language, are entitled to no patentable weight. Mapping of a portion of an Asserted Claim to a prior art reference does not represent that such portion of the claim is entitled to patentable weight when comparing the claimed subject matter to the prior art.

II. INVALIDITY CONTENTIONS

As explained herein and in Exhibits A-F, American contends that each of the Asserted Claims is invalid under 35 U.S.C. §§ 101, 102, 103, and/or 112.

A. P.R. 3-3(a) – Identification of Prior Art

Pursuant to P. R. 3-3(a), and subject to American's reservation of rights as stated herein, American identifies the prior art that anticipates or renders obvious the Asserted Claims in the tables set forth below.⁵ On information and belief, each listed reference qualifies as prior art to the Asserted Patents.

To the extent that any of the following are prior art, American reserves the right to rely upon foreign counterparts of the U.S. patents identified herein; U.S. counterparts of foreign patents and foreign patent applications identified herein; and U.S. and foreign patents and patent applications corresponding to articles and publications identified herein. American also reserves the right to rely upon parent or provisional patents or ancestor patents or patent applications from which any of the patents or patent applications identified herein claim priority to as continuation, divisional, or continuation-in-part applications. Identifications of dates of publication are made based on currently available information and American reserves the right to rely upon an earlier date should evidence supporting an earlier date be discovered.

⁵ American also hereby identifies any systems or products that embody the technology described in any patent or publication identified in these Preliminary Invalidity Contentions. American reserves the right to rely on any documents or other evidence regarding any such systems.

1. Identification of Prior Art

a) '582 patent⁶

Patent/Publication No./Document Name	Date of Filing/Issue/Publication ⁷
U.S. Patent No. 6,304,866 (“Chow”)	Filed – June 27, 1997 Issued – October 16, 2001
U.S. Patent Application Publication No. 2002/0019844 (“Kurowski”)	Filed – January 12, 2001 Published – February 14, 2002
U.S. Patent No. 6,330,583 (“Reiffin”)	Filed – September 9, 1994 Issued – December 11, 2001
U.S. Patent No. 6,212,617 (“Hardwick”)	Filed – June 30, 1998 Issued – April 3, 2001
U.S. Patent Application Publication No. 2002/0103889 (“Markson”)	Filed – June 19, 2001 Published – August 1, 2002
U.S. Patent Application Publication No. 2002/0091786 (“Yamaguchi”)	Filed – November 1, 2001 Published – July 11, 2002
U.S. Patent No. 5,765,146 (“Wolf”)	Filed – November 4, 1993 Issued – June 9, 1998
EP Patent Application Publication No. EP0377993 (“Lorie”)	Filed – December 20, 1989 Published – July 18, 1990
P. Maheshwari, “Improving task scheduling for larger grain execution of parallel functional programs,” [1993] Proceedings of the Twenty-sixth Hawaii International Conference on System Sciences, Wailea, HI, USA, 1993, pp. 594-602 vol.2, doi: 10.1109/HICSS.1993.284066. (“Maheshwari”)	January 8, 1993
“Analysis of First-Come-First-Serve Parallel Job Scheduling,” Uwe Schwiegelshohn and Ramin Yahyapour. Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms, pages 629-638 (1998). Available at https://dl.acm.org/doi/10.5555/314613.315031 (“Schwiegelshohn”)	January 1, 1998
Gamma System - University of Wisconsin-Madison	At least as early as 1990

⁶ Defendant notes that there is a pending IPR (IPR2025-00785) asserting that all claims of the '582 patent are invalid. Defendant hereby incorporates by reference, in full, the Petition and the Exhibits in IPR2025-00785. American further incorporates by reference any post grant review proceeding (including any IPRs or EPRs) that have been or may be filed challenging the patentability of one or more Asserted Claim.

⁷ For any patent that claims the benefit of a provisional application(s), American may rely on the provisional application(s) for purposes of establishing invalidity. American may also rely upon the date of filing of any issued patent or patent application to show the priority date of that reference.

Patent/Publication No./Document Name	Date of Filing/Issue/Publication ⁷
Volcano Parallel Database System - University of Wisconsin-Madison	At least as early as 1990
U.S. 6,334,137 (“Iida”)	Filed – August 24, 1999 Issued – December 25, 2001
U.S. 2002/0099716 (“Thompson”)	Filed – January 25, 2001 Published – July 25, 2002
JP 20003/30959 (“Aiba”)	Filed – May 18, 1999 Issued – May 24, 2006
U.S. 2002/0040381 (“Steiger”)	Filed – September 28, 2001 Published – April 4, 2002
U.S. 5,325,493 (“Herrell”)	Filed – March 30, 1993 Issued – June 28, 1994
U.S. 5,535,393 (“Reeve”)	Filed – June 5, 1995 Issued – July 9, 1996
U.S. 5,765,166 (“Gotfried”)	Filed – April 23, 1996 Issued – June 9, 1998
S. Kuo, M. Winslett, Y. Chen, Y. Cho, M. Subramaniam and K. Seamons, “Parallel input/output with heterogeneous disks,” Proceedings. Ninth International Conference on Scientific and Statistical Database Management (Cat. No.97TB100150), Olympia, WA, USA, 1997, pp. 79-90, doi: 10.1109/SSDM.1997.621154.	Published Aug. 11-13, 1997 Retrieved from https://ieeexplore.ieee.org/document/621154
U.S. 5,945,990 (“Morrison”)	Filed – November 19, 1996 Issued – August 31, 1999
WO 03/012696 (“Smith”)	Filed – July 29, 2002 Published – February 13, 2003
U.S. 6,480,876 (“Rehg”)	Filed – May 28, 1998 Issued – November 12, 2002
“Some Thoughts on Parallel Processing,” Lynn D. Yarbough, Communications of the ACM, Vol. 3, Issue 10, page 539, October 1, 1960 (available at: https://dl.acm.org/doi/10.1145/367415.367426).	Oct. 1, 1960
“A Survey of Some Theoretical Aspects of Multiprocessing,” J. L. Baer, ACM Computing Surveys (CSUR), Volume 5, Issue 1 Pages 31 – 80, March 1, 1973 (available at: https://dl.acm.org/doi/10.1145/356612.356615)	March 1, 1973
“Parallel Sorting Algorithms for Tightly Coupled Multiprocessors,” Michael J. Quinn, Parallel Computing 6 (1988) 349-357.	1988
“Removing Skew Effect in Join Operation on Parallel Processors,” Ron-Chung Hu, Richard R. Muntz, Computer Science Technical Report, University of	June 1989

Patent/Publication No./Document Name	Date of Filing/Issue/Publication ⁷
California, Los Angeles, June 1989.	
"Parallel Architectures for Database Systems," A.R. Hurson, L.L. Miller, S.H. Pakzad, M.H. Eich, B. Shirazi, Advances in Computers, Vol. 28, pages 107-151, 1989.	1989
"Parallel Algorithms for Merging and Sorting," Narsingh Deo, Dilip SarKar, Information Sciences, Vol. 56, pages 151-161, 1991.	1991
"Parallel Algorithms for least Median of Squares Regression," Chong-Wei Xu and Wei-Kei Shiue, Computational Statistics & Data Analysis, Vol. 16, pages 349-362, 1993.	1993
"On the Average Running Time of Odd-Even Merge Sort," Christine Rub, Max Plank-Institute Fur Informtik, Research Report, April 1995.	April 1995
"Parallel Algorithms," Guy E. Blelloch and Bruce M. Maggs, a chapter in "The Computer Science and Engineering Handbook." CRC Press, 1997, ISBN: 0-8493-2909-4. Available at https://dl.acm.org/doi/pdf/10.5555/1882723.1882748 .	1997
"An Approach to Parallelizing Isotonic Regression," Anthony J. Kearsley, Richard A. Tapia, Michael W.Trosset, Applied Mathematics and Parallel Computing, H. Fisher et. al. (eds.), Physica-Verlag Heidelberg (1996), pages 141-147.	1996
"Part 1: The Parallel Computing Environment, Alice E. Koniges, Morris A. Jette, and David C. Eder, in "Industrial Strength Parallel Computing," Morgan Kaufmann, (2000). ISBN 9781558605404. Available at: http://wayback.cecm.sfu.ca/PSG/book/intro.html and at https://books.google.com/books?id=mWalawBciCQC&pg=PA1&source=gbs_toc_r&cad=2#v=onepage&q&f=false .	2000
GPFS: A Parallel File System, Jason Barkes, Marcelo R. Barrios, Francis Cougard, Paul G. Crumley, Didac Marin, Hari Reddy, Theerapong Thitayanun, IBM International Technical Support Organization, April 1998.	April 1998
"Analysis of First-Come-First-Serve Parallel Job Scheduling," Uwe Schwiegelshohn and Ramin Yahyapour. Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (1998). Available at http://dx.doi.org/10.1145/314613.315031 .	1998
"Improving First-Come-First-Serve-Job Scheduling	March 30, 1998

Patent/Publication No./Document Name	Date of Filing/Issue/Publication ⁷
by Gang Scheduling,” Uwe Schwiegelshohn and Ramin Yahyapour. Job Scheduling Strategies for Parallel Processing, IPPS/SPDP '98 Workshop, Orlando, Florida, USA, March 30, 1998, pp. 180-198, available at https://www.academia.edu/14780669/Improving_first_come_first_serve_job_scheduling_by_gang_scheduling	
Bricker, et al., “Condor Technical Summary,” University of Wisconsin – Madison Computer Sciences Department	Jan. 1992
W. Gentzsch, “Sun Grid Engine: Towards Creating a Compute Power Grid,” Sun Microsystems, Inc.	2001
O. Sievert, “A Gentle Introduction to Globus,” Prepared for Parallel Computation, CSE160/CSE260, Spring 1999	1999
Anderson, et al., “SETI@home Internet Distributed Computing for SETI,” ASP Conference Series, Vol. 213	2000
Bamford, et al., “Architecture of Oracle Parallel Server,” Proceedings of the 24 th Very Large Database (VLDB) Conference New York, USA, 1998	1998
J. Howard, “An Overview of the Andrew File System,” Carnegie Mellon University	1988
S. Hotovy, “Workload Evolution on the Cornell Theory Center IBM SP2” Cornell Theory Center	1996
Foster et al., “Wide-Area Implementation of the Message Passing Interface,” Parallel Computing 24(12),1735-1749	1998
Heymann et al., “Adaptive Scheduling for Master-Worker Applications on the Computational Grid,” Department of Computer Sciences University of Wisconsin – Madison	2000
D. Gelernter and N. Carriero, “How to Write Parallel Programs: A Guide to the Perplexed,” <i>ACM Computing Surveys</i> , Vol. 21, No. 3	Sept. 1989

b) '722 patent⁸

Patent/Publication No./Document Name	Date of Filing/Issue/Publication
U.S. Patent No. 5,721,825 (“Lawson”)	Filed – October 3, 1996 Issued – February 24, 1998
U.S. Patent No. 6,480,883 (“Tsutsumitake”)	Filed – June 29, 1999 Issued – November 12, 2002
PCT Patent Application Publication No. WO 99/26127 (“Kostes”)	Filed – November 13, 1998 Published – May 27, 1999
U.S. Patent No. 6,999,991 (“Ikeda”)	Filed – July 24, 2000 Issued – February 14, 2006
U.S. Patent No. 6,990,591 (“Pearson”)	Filed – December 22, 1999 Issued – January 24, 2006
European Patent Application Publication No. EP1043671 (“Bird”)	Filed – March 17, 2000 Published – October 11, 2000
U.S. Patent No. 6,721,740 (“Skinner”)	Filed – June 5, 1998 Issued – April 13, 2004
U.S. Patent No. 6,710,702 (“Averbuch”)	Filed – November 22, 1999 Issued – March 23, 2004
U.S. Patent No. 6,763,384 (“Gupta”)	Filed – July 10, 2000 Issued – July 13, 2004
U.S. Patent No. 6,859,829 (“Parupudi”)	Filed – February 23, 1999 Issued – February 22, 2005
U.S. Patent No. 6,216,132 (“Chandra”)	Filed – November 20, 1997 Issued – April 10, 2001
U.S. Patent No. 6,910,070 (“Mishra”)	Filed – January 24, 2000 Issued – June 21, 2005
U.S. Patent No. 5,754,939 (“Herz”)	Filed – October 31, 1995 Issued – May 19, 1998
U.S. Patent No. 6,823,359 (“Heidingsfeld”)	Filed – November 21, 2000 Issued – November 23, 2004
U.S. Patent No. 6,757,283 (“Yamanaka”)	Filed – January 24, 2000 Issued – June 29, 2004
U.S. Patent No. 7,051,084 (“Hayton”)	Filed – November 2, 2000 Issued – May 23, 2006
U.S. Patent No. 6,633,910 (“Rajan”)	Filed – December 14, 1999 Issued – October 14, 2003
U.S. Patent No. 9,300,560 (“Leighton”)	Filed – March 4, 2013

⁸ American notes that there is a pending IPR asserting that all claims of the '722 patent are invalid, IPR2025-00987, and a settled IPR, IPR2025-00200. American hereby incorporates by reference, in full, the Petition and the Exhibits in both IPR2025-00987 and IPR2025-00200. American further incorporates by reference any post grant review proceeding (including any IPRs or EPRs) that has been or may be filed challenging the patentability of one or more Asserted Claim.

Patent/Publication No./Document Name	Date of Filing/Issue/Publication
	Issued – March 29, 2016
U.S. Patent No. 5,951,694 (“Choquier”)	Filed – February 3, 1997 Issued – September 14, 1999
PointCast System – PointCast Incorporated	At least as early as 1997
Castanet System – Marimba, Inc.	At least as early as 1996
Active Desktop System – Microsoft Corporation	At least as early as 1998
U.S. 2002/0032722 (“Baynes”)	Filed – September 12, 2001 Published – March 14, 2002
U.S. 7,103,680 (“Holdsworth”)	Filed – April 25, 2000 Issued – September 5, 2006
U.S. 6,886,044 (“Miles”)	Filed – June 30, 1999 Issued – April 26, 2005
U.S. 6,654,786 (“Fox”)	Filed – August 30, 2000 Issued – November 25, 2003
U.S. 5,913,032 (“Schwartz”)	Filed – September 30, 1996 Issued – June 15, 1999
U.S. 7,516,196 (“Madan”)	Filed – March 21, 2000 Issued – April 7, 2009
U.S. 2010/0325562 (“Andrews”)	Filed – August 27, 2010 Published – December 23, 2010
U.S. 6,643,682 (“Todd”)	Filed – February 22, 2000 Issued – November 4, 2003
JP JPH11353216A (“Masaru”)	Filed – June 8, 1998 Published – December 24, 1999
U.S. 6,694,352 (“Omoigui”)	Filed – December 16, 1999 Issued – February 17, 2004
U.S. 6,336,119 (“Banavar”)	Filed – March 30, 2000 Issued – January 1, 2002
U.S. 5,978,842 (“Noble”)	Filed – July 18, 1997 Issued – November 2, 1999
U.S. 6,760,759 (“Chan”)	Filed – May 1, 2000 Issued – July 6, 2004
GB2348025 (“Bird GB”)	Filed – March 19, 1999 Published – May 12, 1999
M. Hauswirth and M. Jazayeri, “A Component and Communication Model for Push Systems,” Proceedings of the ESEC/FSE 99 - Joint 7th European Software Engineering Conference (“Hauswirth”)	Pub. at least as early as September 6, 1999.
G. S. Barnes Nelson, “Messaging Systems: SAS® Tools for Internet-Based Communications,” Proceedings of the 25th Annual SAS Users Group International Conference (SUGI 25), Paper 288-25 (“Barnes Nelson”)	Pub. at least as early as March 19, 1999.
Segall, et al., “Content Based Routing with Elvin4,”	Pub. at least as early as June 2000.

Patent/Publication No./Document Name	Date of Filing/Issue/Publication
Proceedings of AUUG2K (No. 39) (“Segall”)	

c) ’785 patent⁹

Patent/Publication No./Document Name	Date of Filing/Issue/Publication
U.S. Patent No. 6,970,941 (“Caronni I”)	Filed – December 10, 1999 Issued – November 29, 2005
U.S. Patent No. 7,814,228 (“Caronni II”)	Filed – February 13, 2003 Issued – October 12, 2010
U.S. Patent No. 6,766,371 (“Hipp”)	Filed – October 5, 2000 Issued – July 20, 2004
U.S. Patent Application Publication No. 2003/0028671 (“Mehta”)	Filed – June 10, 2002 Published – February 6, 2003
Huitema, C., Network Working Group Request for Comment (RFC): 1383, entitled An Experiment in DNS Based IP Routing (“RFC 1383”)	Published December 1992
U.S. Patent No. 7,133,404 (“Alkhatib”)	Filed – August 10, 2001 Issued – November 7, 2006
U.S. Patent Application Publication No. 2002/0184529 (“Foster I”)	Filed – April 19, 2002 Published – December 5, 2002
U.S. Patent Application Publication No. 2003/0208602 (“Bhalla”)	Filed – April 8, 2002 Published – November 6, 2003
U.S. Patent Application Publication No. 2002/0181395 (“Foster II”)	Filed – April 19, 2002 Published – December 5, 2002
U.S. Patent Application Publication No. 2002/0091859 (“Tuomenoksa”)	Filed – April 11, 2001 Published – July 11, 2002
U.S. Patent Application Publication No. 2004/0177136 (“Chen”)	Filed – March 3, 2003 Published – September 9, 2004
European Patent Publication No. EP1098496 (“Engmann”)	Published – October 19, 2000
U.S. Patent Application Publication No. 2002/0069278 (“Forslow”)	Filed – December 5, 2000 Published – June 6, 2002
U.S. Patent No. 6,701,437 (“Hoke”)	Filed – November 9, 1998 Issued – March 2, 2004
U.S. Patent Application Publication No. 2002/0116502 (“Iyer”)	Filed – February 22, 2001 Published – August 22, 2002
U.S. Patent Application Publication No.	Filed – September 28, 2001

⁹ Defendant notes that there is a pending IPR (IPR2025-00786) asserting that all claims of the ’785 patent are invalid. Defendant hereby incorporates by reference, in full, the Petition and the Exhibits in IPR2025-00786. American further incorporates by reference any post grant review proceeding (including any IPRs or EPRs) that has been or may be filed challenging the patentability of one or more Asserted Claim.

Patent/Publication No./Document Name	Date of Filing/Issue/Publication
2003/0065785 (“Jain”)	Published – April 3, 2003
U.S. Patent Application Publication No. 2002/0035624 (“Kim”)	Filed – July 9, 2001 Published – March 21, 2002
U.S. Patent Application Publication No. 20020154635 (“Liu”)	Filed – April 23, 2001 Published – October 24, 2002
U.S. Patent No. 8,068,817 (“Viswanath”)	Filed – August 27, 2002 Issued – November 29, 2011
U.S. Patent Application Publication No. 2002/0116523 (“Warrier”)	Filed – February 22, 2001 Published – August 22, 2002
U.S. Patent Application Publication No. 2003/0229697 (“Borella”)	Filed – June 10, 2002 Published – December 11, 2003
U.S. Patent Application Publication No. 2003/0084162 (“Johnson”)	Filed – October 31, 2001 Published – May 1, 2003
PCT Patent Application Publication No. WO2003088625 (“Volz”)	Filed – March 28, 2003 Priority – April 8, 2002 Published – October 23, 2003
U.S. Patent Application Publication No. 2004/0098507 (“Thubert”)	Filed – November 20, 2002 Published – May 20, 2004
U.S. Patent Application Publication No. 2003/0016636 (“Tari”)	Filed – July 15, 2002 Published – January 23, 2003
U.S. Patent No. 8,107,483 (“Dunk”)	Filed – March 12, 2010 Priority – November 27, 2002 Issued – January 31, 2012
PCT Patent Application Publication No. WO2002015014 (“Wootton”)	Filed – July 31, 2001 Published – February 21, 2002
U.S. Patent Application Publication No. 2002/0038382 (“Ryu”)	Filed – August 30, 2001 Published – March 28, 2002
U.S. Patent Application Publication No. 2003/0172184 (“Kong”)	Filed – March 6, 2003 Published – September 11, 2003
U.S. Patent Application Publication No. 2004/0037260 (“Kakemizu”)	Filed – August 7, 2003 Published – February 26, 2004
Cisco VPN Routers and Software Systems – Cisco Systems	At least as early as 2002
VMware Systems - VMware	At least as early as 2001
U.S. Patent Application Publication No. 2002/0029276 (“Bendinelli”)	Filed – April 11, 2001 Published – October 7, 2003
“Virtual Enterprise Networks: The Next Generation of Secure Enterprise Networking,” Proceedings of the Sixteenth Annual Computer Security Applications Conference (ACSAC'00) (“NextGen”)	Date of Conference: December 11-15, 2000 Published on IEEE Xplore: August 6, 2002

d) '844 patent¹⁰

Patent/Publication No./Document Name	Date of Filing/Issue/Publication
“Optimizing the Migration of Virtual Computers,” Association for Computing Machinery, Special Interest Group on Operating Systems, Operating Systems Review, Proceedings of the 5th Symposium on Operating Systems Design and Implementation, Volume 36, Issue S1 (“Sapuntzakis”)	Published December 31, 2002
U.S. Patent No. 7,089,300 (“Birse”)	Filed – October 19, 1999 Issued – August 8, 2006
U.S. Patent Application Publication No. 2005/0283597 (“Holzmann”)	Filed – June 22, 2004 Published – December 22, 2005
U.S. Patent Application Publication No. 2002/0034105 (“Kulkarni”)	Filed – January 8, 2001 Published – March 21, 2002
U.S. Patent Application Publication No. 2005/0004925 (“Stahl”)	Filed – May 7, 2004 Published – January 6, 2005
U.S. Patent Application Publication No. 2006/0168154 (“Zhang”)	Filed – November 19, 2004 Published – July 27, 2006
U.S. Patent No. 8,001,085 (“Kiselev”)	Filed – November 25, 2003 Issued – August 16, 2011
U.S. Patent Application Publication No. 2004/0172509 (“Takeda”)	Filed – June 23, 2003 Published – September 2, 2004
U.S. Patent No. 6,671,782 (“Menon”)	Filed – March 29, 2000 Issued – December 30, 2003
U.S. Patent Application Publication No. 2003/0236850 (“Kodama”)	Filed – March 21, 2002 Published – December 25, 2003
EP Patent No. EP0604013 (“Nelson”)	Filed – November 15, 1993 Issued – February 20, 2002
PCT Application Publication No. WO 2006/074869 (“Bayer”)	Filed – January 4, 2006 Published – July 20, 2006
Shrira, Liuba et al., “Thresher: an efficient storage manager for copy-on-write snapshots,” ATEC '06: Proceedings of the annual conference on USENIX '06 Annual Technical Conference (Shrira)	Published at least by May 30, 2006
Laros, James H. et al., “Implementing Scalable Disk-Less Clusters Using The Network File	Published at least by January 2003

¹⁰ Defendant notes that there is a pending IPR (IPR2025-00931) asserting that all claims of the '844 patent are invalid. Defendant hereby incorporates by reference, in full, the Petition and the Exhibits in IPR2025-00931. American further incorporates by reference any post grant review proceeding (including any IPRs or EPRs) that has been or may be filed challenging the patentability of one or more Asserted Claim.

Patent/Publication No./Document Name	Date of Filing/Issue/Publication
System (NFS),” (“Laros”) Retrieved from: https://www.researchgate.net/publication/255223070_Implementing_scalable_disk-less_clusters_using_the_Network_File_System_NFS	
U.S. Patent Application Publication No. 2004/0153639 (“Cherian”)	Filed – February 5, 2003 Published – August 5, 2004
U.S. Patent Application Publication No. 2007/0057958 (“Bucher”)	Filed – August 9, 2006 Published – March 15, 2007
U.S. Patent No. 7,197,608 (“Mikuma”)	Filed – July 12, 2004 Issued – March 27, 2007
U.S. Patent Application Publication No. 2005/0283575 (“Kobayashi”)	Filed – August 13, 2004 Published – December 22, 2005
U.S. Patent No. 6,453,334 (“Vinson”)	Filed – June 16, 1998 Issued – September 17, 2002
U.S. Patent No. 5,930,513 (“Taylor”)	Filed – June 6, 1996 Issued – July 27, 1999
U.S. Patent Application Publication No. 2002/0002660 (“Malcolm”)	Filed – July 31, 1998 Published – January 3, 2002
U.S. Patent Application Publication No. 2002/0091763 (“Shah”)	Filed – May 15, 2001 Published – July 11, 2002
U.S. Patent No. 7,398,382 (“Rothman”)	Filed – December 29, 2004 Issued – July 8, 2008
U.S. Patent No. 6,317,826 (“McCall”)	Filed – October 19, 1998 Issued – November 13, 2001
Southwest Airlines NetApp System – NetApp/Southwest	At least as early as 2002
VMware Systems - VMware	At least as early as 2001
U.S. Patent No. 6,618,736 (“Menage”)	Filed – March 9, 2001 Issued – September 9, 2003
Apple NetBoot	At least as early as 2002
U.S. Patent Application Publication No. 2002/0186698 (“Ceniza”)	Filed – June 12, 2001 Published – December 12, 2002

e) '469 patent¹¹

Patent/Publication No./Document Name	Date of Filing/Issue/Publication
U.S. Patent No. 6,894,990 ("Agarwal")	Filed – October 13, 2000 Issued – May 17, 2005
U.S. Patent Application Publication No. 2002/0138625 ("Bruner")	Filed – March 21, 2001 Published – September 26, 2002
U.S. Patent No. 6,445,777 ("Clark")	Filed – December 12, 1998 Issued – September 3, 2002
DeSanctis, et al., "Aeronautical Communications for Personal and Multimedia Services via Satellite," AM. INST. OF AERONAUTICAL. AND ASTRONAUTICS ("DeSanctis")	Published at least by April 2003
U.S. Patent No. 6,968,394 ("El-Rafie")	Filed – March 21, 2000 Published – November 22, 2005
U.S. Patent Application No. 2002/0075844 ("Hagen")	Filed – April 10, 2001 Published -- June 20, 2002
Jahn, et al., "Dimensioning of Aeronautical Satellite Services," AM. INST. OF AERONAUTICAL. AND ASTRONAUTICS ("Jahn I")	Published at least by October 2002
Jahn, et al., "Passenger Multimedia Service Concept via Future Satellite Systems" DGLR JAHRBUCH ("Jahn II")	2002
U.S. Patent Application No. 2003/0051041 ("Kalavade")	Filed – August 6, 2002 Published – March 13, 2003
Korean Patent No. 2003/0029267 ("Lee")	Filed – October 5, 2001 Published – April 14, 2003
U.S. Patent Application No. 2007/0008937 ("Mody")	Filed – July 21, 2004 Published – January 11, 2007
U.S. Patent Application No. 2003/0055975 ("Nelson")	Filed – June 19, 2001 Published – March 28, 2003
U.S. Patent No. 6,377,981 ("Ollikainen")	Filed November 20, 1997 Published – April 23, 2002
Czech Patent No. 2,000,578 ("Rothblatt I")	Filed – August 19, 1998 Published – August 16, 2000
U.S. Patent No. 6,105,060 ("Rothblatt II")	Filed – September 5, 1997

¹¹ American notes that there is a pending IPR (IPR2025-00782) asserting that the asserted claims of the '469 patent are invalid. American hereby incorporates by reference, in full, the Petition and the Exhibits in IPR2025-00782. American further incorporates by reference any post grant review proceeding (including any IPRs or EPRs) that has been or may be filed challenging the patentability of one or more Asserted Claim.

Patent/Publication No./Document Name	Date of Filing/Issue/Publication
	Published – August 15, 2000
U.S. Patent Application No. 2004/0109449 (“Seo”)	Filed – August 15, 2003 Published – June 10, 2004
Suwanateep, Manasuda, “Inflight Internet Connectivity Technology and Its Application on E-Commerce,” U. OF THAILAND (2002) (“Suwanateep”)	Published at least by July 2002
U.S. Patent No. 5,987,430 (“Van Horne”)	Filed – August 28, 1997 Published – November 16, 1999
U.S. Patent Application Publication No. 2003/0050041 (“Wu”)	Filed – September 7, 2001 Published – March 13, 2003

f) '326 patent¹²

Patent/Publication No./Document Name	Date of Issue/Publication/Priority
PCT Patent App. Pub. No. 2000/038387 (“Baldemair”)	Filed – December 21, 1999 Published – June 29, 2000
U.S. Patent No. 7,418,050 (“Gardner”)	Filed – February 26, 2003 Issued – August 26, 2008
U.S. Patent App. Pub. No. 2003/0022473 (“Yamaura”)	Filed – June 17, 2002 Published – January 30, 2003
U.S. Patent No. 6,510,133 (“Uesugi”)	Filed – May 28, 1998 Issued – January 21, 2003
JP Patent App. Pub. No. JP2002319917 (“Mori”)	Filed – April 24, 2001 Published – October 31, 2002
JP Patent App. Pub. No. JP2003309533 (“Miyoshi”)	Filed – April 17, 2002 Published – October 31, 2003
JP Patent App. Pub. No. JPH11113049 (“Koga”)	Filed – September 30, 1997 Published – April 23, 1999
Thompson, D. E., “Modelling Adjacent Channel Interference in 3G Networks,” THE INSTITUTE OF ELECTRICAL ENGINEERS (2003) (“Thompson”)	2003
Weste, Neil et al., “VLSI for OFDM,” IEEE COMMUNICATIONS MAGAZINE (Oct. 1998) (“Weste”)	1998
“WLAN Channel Bonding: Causing Greater Problems Than It Solves” (“Texas Instruments White Paper”)	Sep. 2003

¹² American notes that there is a pending IPR (IPR2025-01055) asserting that the asserted claims of the '326 patent are invalid. American hereby incorporates by reference, in full, the Petition and the Exhibits in IPR2025-01055. American further incorporates by reference any post grant review proceeding (including any IPRs or EPRs) that has been or may be filed challenging the patentability of one or more Asserted Claim.

Patent/Publication No./Document Name	Date of Issue/Publication/Priority
U.S. Patent No. 7,095,708 (“Alamouti”)	Filed – June 14, 2000 Issued – August 22, 2006
U.S. Patent No. 6,975,585 (“Olafsson”)	Filed – July 27, 2000 Issued – December 13, 2005

1. Identification Of Prior Art Sales/Public Uses

Item Offered for Sale and/or Publicly Used	Date of Offer/Public Use	Person or Entity Who Made and Received Offer, Made Public Use, or Made Information Known
NetApp’s WAFL File System		This architecture was thoroughly documented in both U.S. Patents 5,819,292 and 7,334,095, as well as in technical literature such as “File System Design for an NFS File Server Appliance” by Dave Hitz et al.
VMware’s GSX and ESX server platforms		
Sun Microsystems’ diskless workstations		
The Episode File System	Early 1990s	
Petal		HP Labs
“WLAN Solutions: TNETW1130 Converged Single-Chip MAC and Baseband Processor for IEEE 802.11 a/b/g” (“TI Baseband Processor for IEEE 802.11”)	2003	Texas Instruments, Inc.
Airpath Wireless Hot Spot	No later than June 2003	Airpath Wireless, Inc., Mainstream Data, Inc., Todd Meyers
Connexion	No later than June 2003	The Boeing Co., Connexion by Boeing, Scott E. Carson, Cisco Systems, Inc., American Airlines, Inc., Delta Air Lines, Inc., United Airlines, Inc., Lufthansa Group
Maritime Telecommunications Network Cruise Ship Internet Services	No later than June 2003	MTN Satellite Communications (MTN), formerly known as Maritime Telecommunications Network, PCTEL, Inc., Digital Seas International (DSI), Norwegian Cruise Line (NCL), Semester at Sea (SaS), and other cruise line customers of MTN
Stratos	No later than	Stratos Global Corporation; Inmarsat

Item Offered for Sale and/or Publicly Used	Date of Offer/Public Use	Person or Entity Who Made and Received Offer, Made Public Use, or Made Information Known
	September 2003	Global Ltd.
Tenzing Global	No later than April 2001	Tenzing Communications, Inc.; Tenzing Technologies, LLC; Airbus SE; Telia company AB, SITA, OnAir company, Cathay Pacific Airways Ltd., Singapore Airlines Ltd., Virgin Atlantic Airways Ltd.
Wireless Cabin	No later than February 2003	Inmarsat Global Ltd., Telefonaktiebolaget LM Ericsson, Siemens AG, Kid-Systeme GmbH, ESYS Consulting Ltd., Airbus SE, University of Bradford

B. P.R. 3-3(b) – Anticipation and Obviousness

Pursuant to P.R. 3-3(b), and subject to American's reservation of rights, American attaches claim charts hereto that are directed to the prior art references that anticipate each of the Asserted Claims under 35 U.S.C. §§ 102(a), (b), (e), and/or (g), either expressly or inherently, and/or the prior art references that, in the alternative, would have rendered the Asserted Claims obvious under 35 U.S.C. § 103. *See* Exhibits A-F. Some claim charts also contain explanations about the motivation to combine the references. The combinations contained in American's claim charts are exemplary. Any prior art reference cited herein may be combined with any other reference to demonstrate the invalidity of any of the Asserted Claims, as set forth below and in Exhibits A-F.

To the extent any claim limitation is construed to have a similar meaning, or to encompass similar feature(s) and/or function(s), as any other claim limitation, the citations to prior art references for each of those claim limitations in American's claim charts are incorporated by reference with respect to each other.

American's claim charts provide exemplary citations to the prior art references that teach or suggest every element of the Asserted Claims. To the extent that an element of an Asserted

Claim is not shown in a chart, the Asserted Claims would have been obvious based on a combination of one or more other prior art references, as set forth below and in Exhibits A-F.

Much of the art cited in these Preliminary Invalidity Contentions reflects common knowledge and the state of the art at the time of the earliest filing date of the Asserted Patents. American may rely on additional citations, references, expert testimony, and other material to provide context or to aid in understanding the cited portions of the references and/or cited features of the systems. American also may rely on expert testimony explaining relevant portions of references, relevant hardware or software products or systems, and other discovery regarding these subject matters. Additionally, American may rely on other portions of any prior art reference or other references relied on by the same authors or describing the same systems for purposes of explaining the background and general technical subject area of the reference.

Where an individual reference is cited with respect to all elements of an Asserted Claim, American contends that the reference anticipates the claim under 35 U.S.C. §§ 102(a), (b), (e), and/or (g) and also renders obvious the claim under 35 U.S.C. § 103, both by itself in view of the knowledge of a person of ordinary skill in the art and in combination with the other cited references to the extent the reference is not found to disclose one or more claim elements. A single prior art reference, for example, can establish obviousness where the differences between the disclosures within the reference and the claimed invention would have been obvious to one of ordinary skill in the art. For example, “[c]ombining two embodiments disclosed adjacent to each other in a prior art patent does not require a leap of inventiveness.” *Boston Scientific Scimed, Inc. v. Cordis Corp.*, 554 F.3d 982, 991 (Fed. Cir. 2009). To the extent IV contends that an embodiment within a particular item of prior art does not fully disclose all limitations of a

claim, American accordingly reserves its rights to rely on other embodiments in that prior art reference, or other information, to show single reference obviousness under 35 U.S.C. § 103(a).

Where an individual reference is cited with respect to fewer than all elements of an Asserted Claim, American contends that the reference renders obvious the claim under 35 U.S.C. § 103(a) in view of each other reference and combination of references that discloses the remaining claim element(s), as indicated in the claim charts submitted herewith. “Under § 103, the scope and content of the prior art are to be determined; differences between the prior art and the claims at issue are to be ascertained; and the level of ordinary skill in the pertinent art resolved. Against this background, the obviousness or nonobviousness of the subject matter is determined.” *KSR Int’l Co. v. Teleflex Inc.*, 550 U.S. 398, 406 (2007) (quoting *Graham v. John Deere Co.*, 383 U.S. 1, 17 (1966)).

Exemplary motivations to combine references are discussed in the accompanying charts. American reserves the right to rely upon any references or assertions identified herein in connection with American’s contentions that each Asserted Claim is invalid under 35 U.S.C. § 103 and to rely upon expert testimony addressing such references and assertions. The fact that prior art is identified to anticipate the Asserted Claims presents no obstacle in also relying on that reference as a basis for invalidity based on obviousness. It is established that “a rejection for obviousness under § 103 can be based on a reference which happens to anticipate the claimed subject matter.” *In re Meyer*, 599 F.2d 1026, 1031 (C.C.P.A. 1979). To the extent any cited prior art item may not fully disclose a limitation of an Asserted Claim or is alleged by IV to lack disclosure of the limitation, such limitation is present and identified in another prior art item as shown in the attached claim charts.

Many of the cited references cite or relate to additional references and/or products, services, or projects. Many of the cited references also cite software, hardware, or systems. American might rely upon such cited additional references and/or products and copies or exemplars of such software, hardware, or systems. American will produce or make available for inspection any such cited references, products, software, hardware, or systems that it intends to rely upon. American may also rely upon the disclosures of the references cited and/or discussed during the prosecution of the Asserted Patents and/or the assertions presented regarding those references.

American reserves the right to further streamline and reduce the number of anticipation or obviousness references relied upon with respect to a given Asserted Claim and to exchange or otherwise modify the specific references relied upon for anticipation and within each obviousness combination for each Asserted Claim.

C. P.R. 3-3(c) – Claim Charts

American attaches the following claim charts pursuant to P.R. 3-3(c):

Asserted Patent	Claim Chart Exhibits
'582 patent	A1-A6
'722 patent	B1-B12
'785 patent	C1-C33
'844 patent	D1–D26; OTDP-1–OTDP-2
'469 patent	E1-E18
'326 patent	F1-F13

The attached claim charts are based, in whole or in part, on IV's asserted theories of infringement in this case, to the extent discernible from IV's Infringement Contentions. As an initial matter, all portions of each prior art reference cited in each of the attached claim charts are relied upon to support the disclosure of each patent claim limitation, as all portions provide general support. Representative descriptions and supporting citations are nevertheless provided but are merely exemplary; they do not necessarily reflect every instance where a particular claim

term or claim limitation may be disclosed in or taught by the prior art reference. References to figures or drawings refer to the figures/drawings themselves, as well as to any accompanying text or text necessary to understand the figures or drawings. References to text refers to the text itself, as well as the accompanying figures or drawings that accompany the text. American reserves the right to rely on additional, or different, portions of the prior art references, other publications and expert testimony to establish what these references would have taught one of ordinary skill in the art, or in what manner they would have motivated a particular combination of references. Moreover, in certain instances, representative documents for certain prior art systems are cited, but, again, they are merely exemplary; they do not necessarily reflect or include every document relating to the prior art system that exists and that discloses or teaches a particular claim term or claim limitation. American reserves the right to rely on any and all documents that describe or relate to prior art systems, including relying on the system itself. American also reserves the right to rely on the testimony of the authors, named inventors, or anyone else with knowledge of the prior art references and systems identified herein, as well as expert testimony regarding any such references or systems.

D. Obviousness and Motivation to Combine

The primary references identified above, and as further described in Exhibits A-F, each discloses, either expressly or inherently, every element of the Asserted Claims, thereby anticipating those claims. To the extent IV contends that any primary reference does not anticipate the Asserted Claims, it would have been obvious to combine or modify the primary references with concepts from other prior art, as explained herein and in Exhibits A-F.

In particular, for each limitation of the Asserted Claims that IV contends is not met by a particular primary reference, American contends that the limitation (and claim as a whole) is obvious based on a combination of that particular primary reference with (1) any other primary

reference disclosing that limitation, (2) any admitted prior art, as explained in the background of each patent or discussed in the file history, (3) any reference identified in Exhibits A-F as disclosing that limitation, and/or (4) the knowledge of a person of ordinary skill in the art and/or any of the references and concepts discussed herein regarding the relevant background and state of the art. The specific combinations of prior art that American contends render the Asserted Claims obvious are readily determinable as described herein, which is the most efficient manner of identifying the combinations in light of the fact that IV has asserted 21 claims across 6 patents. American's obviousness grounds for each dependent claim incorporate the obviousness grounds for the claim(s) from which the dependent claim depends in addition to any obviousness grounds identified in the charts for the dependent claim.

American does not yet have the benefit of IV's positions on the prior art, including what (if any) elements it contends are missing in each prior art reference, whether IV agrees that a reference is in fact prior art, and whether IV agrees that a person of ordinary skill in the art would be motivated to combine specific references. American reserves the right to supplement these obviousness positions (including identifying additional prior art combinations and the associated reasons to combine) as discovery in the case progresses, including expert discovery.

Each prior art reference may be combined with one or more other prior art references to render obvious the Asserted Claims in combination, as explained in more detail below. The disclosures of these references also may be combined with information known to persons skilled in the art at the time of the alleged invention and understood and supplemented in view of the common sense of persons skilled in the art at the time of the alleged invention, including any statements in the intrinsic record of the Asserted Patents and related applications.

A person of ordinary skill would have been motivated to combine the prior art cited in the attached claim charts based on the nature of the problem to be solved, the teachings of the prior art, and the knowledge of persons of ordinary skill in the art. The identified prior art references, including portions cited in the claim charts, address the same or similar technical issues and suggest the same or similar solutions to those issues as the Asserted Claims. On such bases, on an element-by-element basis, American intends to combine one or more prior art references identified in the claim charts attached as Exhibits A-F with each other to address any contention from IV that a particular prior art item supposedly lacks one or more elements of an Asserted Claim. In other words, American contends that each charted prior art reference can be combined with other charted prior art references when a particular prior art item lacks or does not explicitly disclose an element or feature of an Asserted Claim.

A motivation, teaching or suggestion to combine the prior art identified exists for each obviousness combination identified in Exhibits A-F. In *KSR International Co. v. Teleflex Inc., et al.*, 550 U.S. 398 (2007), the Supreme Court held that a claimed invention can be obvious even if there is no teaching, suggestion, or motivation for combining the references to produce the claimed invention. *KSR* holds that patents that are based on new combinations of elements or components already known in a technical field may be found obvious. Specifically, the Court in *KSR* rejected a rigid application of the “teaching, suggestion, or motivation [to combine] test.” *Id.* at 418-419. “In determining whether the subject matter of a patent claim is obvious, neither the particular motivation nor avowed purpose of the patentee controls. What matters is the objective reach of the claim.” *Id.* at 419. “Under the correct analysis, any need or problem known in the field of endeavor at the time of the invention and addressed by the patent can provide a reason for combining the elements in the manner claimed.” *Id.* at 420. The Court noted that

“[c]ommon sense teaches, however, that familiar items may have obvious uses beyond their primary purposes, and in many cases a person of ordinary skill will be able to fit the teachings of multiple patents together like pieces of a puzzle.” *Id.* at 420. The Court further stated that “[t]he combination of familiar elements according to known methods is likely to be obvious when it does no more than yield predictable results” *Id.* at 416. “When there is a design need or market pressure to solve a problem and there are a finite number of identified, predictable solutions, a person of ordinary skill has good reason to pursue the known options within his or her technical grasp. If this leads to the anticipated success, it is likely the product not of innovation but of ordinary skill and common sense.” *Id.* at 421.

Based on *KSR*, the USPTO issued a set of Examination Guidelines. *See* Examination Guidelines for Determining Obviousness Under 35 U.S.C. §103 in view of the Supreme Court Decision in *KSR International Co. v. Teleflex, Inc.*, 72 Fed. Reg. 57526 (October 10, 2007). The Guidelines provide a variety of rationales on which to base a finding of obviousness. Those rationales are based on *KSR* and other precedent. The rationales include:

- (A) Combining prior art elements according to known methods to yield predictable results;
- (B) Simple substitution of one known element for another to obtain predictable results;
- (C) Use of known technique to improve similar devices (methods, or products) in the same way;
- (D) Applying a known technique to a known device (method, or product) ready for improvement to yield predictable results;
- (E) “Obvious to try”— choosing from a finite number of identified, predictable solutions, with a reasonable expectation of success;
- (F) Known work in one field of endeavor may prompt variations of it for use in either the same field or a different one based on design incentives or

other market forces if the variations would have been predictable to one of ordinary skill in the art;

- (G) 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.

Id. at 57529. American contends that one or more of these rationales apply in considering the obviousness of the Asserted Claims. A person of ordinary skill in the art at the time of the claimed invention(s) had reason to combine or modify one or more of the references listed and charted in Exhibits A-F in light of the knowledge of a person of ordinary skill at the time and information in the prior art references cited herein.

Additional evidence establishing that there would have been a motivation to combine the prior art references identified above pursuant to the Scheduling Order includes the interrelated teachings of multiple prior art references; the effects of demands known to the design community or present in the marketplace; the existence of a known problem for which there was an obvious solution encompassed by the Asserted Claims; the existence of a known need or problem in the field of the endeavor at the time of the invention(s); and the background knowledge that would have been possessed by a person having ordinary skill in the art.

Further, the prior art references explicitly or implicitly reference other prior art references, share common authors or inventors, were published in the same journals, presented at the same conferences, were presented as proposals to industry groups, and/or were developed at common companies, schools, or organizations, all of which would motivate one of skill in the art to combine them. These references are also within the field of the Asserted Patents and are directed to similar subject matter within that field. Additionally, any products, devices, or processes described in the references existed and/or were invented before or during the period in which the claimed inventions were developed, providing further motivation to combine them.

Thus, the motivation to combine the teachings of the prior art references disclosed herein is found in the references themselves and: (1) the nature of the problem being solved, (2) the express, implied and inherent teachings of the prior art, (3) the knowledge of persons of ordinary skill in the art, (4) the fact that the prior art is generally directed toward methods and systems for processing wireless signals, and (5) the predictable results obtained in combining the different elements of the prior art. Additionally, one would be motivated to address at least the alleged problems or achieve the purported objectives identified in the description of the Asserted Patents. *See, e.g.,* the exemplary modifications to combine detailed below. Moreover, the references charted in the Invalidity Contentions recognized and solved these problems.

Any reference or combination of references that anticipates or makes obvious an asserted independent claim also makes obvious any asserted claim dependent on that independent claim because every element of each dependent claim was known by a person of ordinary skill at the time of the alleged invention, and it would have been obvious to combine those known elements with the independent claims at least as a matter of common sense and routine innovation. Accordingly, American contends that each asserted claim would have been obvious not only by the combinations explicitly defined in these contentions, but also by any combination of references that renders obvious an asserted claim.

Numerous prior art references, including those identified above and in the attached claim charts reflect common knowledge and the state of the prior art prior to the priority dates of the Asserted Patents. As it would be unduly burdensome to create detailed claim charts for the thousands of invalidating combinations, American has provided illustrative examples of such invalidating combinations in the attached claim charts. For at least the reasons described above and below in the examples provided, as well as in the attached claim charts, it would have been

obvious to one of ordinary skill in the art to combine any of a number of prior art references, including any combination of those identified in the attached claim charts, to meet the limitations of the Asserted Claims. As such, American's inclusion of exemplary combinations does not preclude them from identifying other invalidating combinations as appropriate.

Secondary considerations of nonobviousness "simply cannot overcome a strong prima facie case of obviousness." *Wyers v. Master Lock Co.*, 616 F.3d 1231, 1246 (Fed. Cir. 2010). For any secondary consideration to be relevant, IV must establish a nexus between the secondary consideration and the claimed inventions. *Ormco Corp. v. Align Tech., Inc.*, 463 F.3d 1299, 1312 (Fed. Cir. 2006).

IV has not provided any evidence of any secondary considerations, much less any nexus between such secondary considerations and the claimed inventions of the Asserted Patents. American contends that no secondary considerations evidence exists that supports the validity of any Asserted Claim and reserve all rights to respond to any argument or evidence presented by IV regarding any alleged secondary considerations of non-obviousness.

1. Statement Regarding Exemplary Obviousness Combinations and Motivation to Combine

The suggested obviousness combinations discussed in the attached claim charts are not to be construed to suggest that any reference included in the combinations is not anticipatory. Further, to the extent that IV contends that any of the anticipatory prior art fails to disclose one or more limitations of the Asserted Claims, American reserves the right to identify other prior art references that, when combined with the anticipatory prior art, would render the claims obvious despite an allegedly missing limitation. American will further specify the motivations to combine the prior art, including through reliance on expert testimony, at the appropriate later stage of this lawsuit.

Further, American notes that there are pending IPRs asserting that at least the asserted claims of the patents are invalid. American hereby incorporates by reference, in full the Petition and Exhibits in the following IPRs: IPR2025-00785; IPR2025-00782; IPR2025-00786; IPR2025-01055; IPR2025-00987; IPR2025-00840; and IPR2025-00931.

2. Secondary Considerations

A patentee bears the burden of production with respect to evidence of secondary considerations of non-obviousness. *ZUP, LLC v. Nash Mfg., Inc.*, 896 F.3d 1365, 1373 (Fed. Cir. 2018). As of the date of these Preliminary Invalidity Contentions, IV has not yet identified any evidence of secondary considerations. American reserves all rights to further respond to any secondary considerations of non-obviousness raised by IV, including by updating, modifying, and/or adding to these Preliminary Invalidity Contentions. American is not aware of any unexpected results (none is mentioned in the Asserted Patents or their file histories), long felt need, commercial success (or any nexus to any allegedly successful commercial embodiment), or awards for the alleged inventions of the Asserted Patents.

E. P.R. 3-3(d): Invalidity Under 35 U.S.C. § 112

As set forth below, American contends that the Asserted Claims are invalid under 35 U.S.C. § 112 because the claims (1) are indefinite; (2) fail to satisfy the enablement requirement; and/or (3) fail to satisfy the written description requirement.

American's invalidity contentions for the Asserted Claims identified in this section are made in the alternative and do not constitute, and should not be interpreted as, admissions regarding the construction or scope of the Asserted Claims identified herein or that any of the Asserted Claims identified herein would not have been anticipated and/or obvious in light of the prior art.

IV has not yet provided a claim construction for any of the terms or phrases that American anticipates will be in dispute. American, therefore, cannot provide a complete list of § 112 defenses because American does not know whether IV will proffer a construction for certain terms and phrases that is broader than, or inconsistent with, the construction that would be supportable by the disclosure set forth in the specification. American offers these contentions without prejudice to any position they may ultimately take as to any claim construction issues.

Accordingly, American reserves the right to supplement, amend, and/or modify these § 112 invalidity contentions as discovery progresses.

1. Invalidity Under 35 U.S.C. § 112, ¶ 1

35 U.S.C. § 112, ¶ 1 requires that the specification contain a written description of the invention. “[T]he hallmark of written description is disclosure.” *Boston Scientific Corp. v. Johnson & Johnson*, 647 F.3d 1353, 1361–62 (Fed. Cir. 2011) (citation omitted). The test for whether a specification adequately describes an invention is “whether the disclosure of the application relied upon reasonably conveys to those skilled in the art that the inventor had possession of the claimed subject matter as of the filing date [T]he test requires an objective inquiry into the four corners of the specification from the perspective of a person of ordinary skill in the art [It] is a question of fact.” *Ariad Pharms., Inc. v. Eli Lilly and Co.*, 598 F.3d 1336, 1351 (Fed. Cir. 2010) (*en banc*); *Boston Scientific*, 647 F.3d at 1362.

The enablement requirement of Section 112 demands that the patent specification enable “those skilled in the art to make and use the full scope of the claimed invention without ‘undue experimentation.’” *Genentech, Inc. v. Novo Nordisk A/S*, 108 F.3d 1361, 1365 (Fed. Cir. 1997) (quoting *In re Wright*, 999 F.2d 1557, 1561 (Fed. Cir. 1993)). “[T]he scope of the claims must be less than or equal to the scope of the enablement.” *Nat’l Recovery Tech., Inc. v. Magnetic Separation Sys., Inc.*, 166 F.3d 1190, 1196 (Fed. Cir. 1999).

American contends that the following Asserted Claims are invalid under 35 U.S.C. § 112, ¶ 1. Each Asserted Claim identified below (and each Asserted Claim that depends therefrom) is invalid under ¶ 1 because the specification of the Asserted Patent fails to provide a sufficient written description, enabling disclosure, and/or fails to inform, with reasonable certainty, those skilled in the art about the scope of the invention the meaning of the term(s) and/or phrase(s). For each listed term or phrase, American believes the term or phrase is invalid under § 112, ¶ 1, as is any limitation including such terms or phrases, for the same reason(s).

Asserted Patent	Asserted Claim(s)	Term(s) and/or Phrase(s) Lacking Written Description and/or Enablement
'582 patent	1	• “distributing descriptions of all of said partitions to each of a plurality of subtask processors” • “generating at least one output combining all of the subtask outputs and reflecting the processing of all of said subtasks”
'722 patent	14	• “live object recognizable by the client device” • “determine a node type to which the identified category maps” • “cohesive respective application environments”
'785 patent	30	• “virtual network manager” • “DNS server for the virtual network ... return a network address associated with a network route director”
'844 patent	7, 11	7: • “only additional data blocks not previously contained in said root image” 11: • “to create cohesive respective application environments” • “at an operational level between file systems”
'469 patent	24, 25, 26, 28, 32	Claim 24: • “Internet Hotspot” • “operatively coupled” • “subscriber access unit” • “authenticating”/“authenticate” • “a remote location experiencing a relatively high volume of transient traffic” • “web-ready device” Claim 25:

Asserted Patent	Asserted Claim(s)	Term(s) and/or Phrase(s) Lacking Written Description and/or Enablement
		“the data connection is one of a wired data connection and a wireless data connection” Claim 32: “the wireless connection is one of an 802.11a wireless area network, an 802.11b wireless area network, an 802.11g wireless area network, and an 802.11n wireless area network”
'326 patent	1, 4, 18	Claims 1 & 18: “plurality of data subcarriers” “partially filling”/ “partially fill” “one or more guard bands” Claim 4: “transmitting to multiple radios”

2. Invalidity Under 35 U.S.C. § 112, ¶ 2

Claims are indefinite under 35 U.S.C. § 112, ¶ 2 when they “fail to inform, with reasonable certainty, those skilled in the art about the scope of the invention.” *Nautilus, Inc. v. Biosig Instruments, Inc.*, 34 S. Ct. 2120, 2124 (2014). To the extent an asserted apparatus claim includes both apparatus and method limitations, that claim is invalid for indefiniteness under § 112, ¶ 2 because it fails to identify or notify the public of what constitutes direct infringement. *See IPXL Holdings, L.L.C. v. Amazon.com, Inc.*, 430 F.3d 1377 (Fed. Cir. 2005).

In addition to American’s reservation of rights stated above, American will present arguments as to indefiniteness will be presented at that time, i.e., through the Markman meet and confer and briefing process.

American contends that the following Asserted Claims are invalid under 35 U.S.C. § 112, ¶ 2. Each Asserted Claim identified below (and each Asserted Claim that depends therefrom) is invalid under ¶ 2 because they fail to inform, with reasonable certainty, those skilled in the art about the scope of the claimed invention. For each listed term or phrase, American believes the

term or phrase (as well as terms within any such phrase) is invalid under § 112, ¶ 2, and any limitation including such term or phrase is also indefinite.

Asserted Patent	Asserted Claim(s)	Indefinite Term(s) or Phrase(s)
'582 patent	1, 5	1: “automatically determining file allocation and logically subdividing records of said input file into a plurality of partitions” “simultaneously executing at least a respective one of the subtasks of the computer-executable process in each of at least some of said processors on a respective one of the partitions with each subtask reading and processing the respective partition so as to process the respective partition and produce respective subtask output” “repeating step (c) in at least some of the subtask processors each with another unprocessed partition on a first-come/first-served basis” 5: “wherein the output in step (e) is an accumulation of output records from said subtasks in an arbitrary order”
'722 patent	14	“live object recognizable by the client device” “gateway device ... configured to identify a category of the update message” “determine a node type to which the identified category maps” “cohesive respective application environments”
'785 patent	30, 37	30: “virtual network manager” “receive a registration request from an agent associated with a device” “DNS server for the virtual network ... return a network address associated with a network route director” 37: “maintain data to associate a virtual network address with a device in the virtual network”
'844 patent	7, 11	7: “only additional data blocks not previously contained in said root image” “accessed by at least one of said compute nodes”

Asserted Patent	Asserted Claim(s)	Indefinite Term(s) or Phrase(s)
		11: “merging the blocks of said root image with the blocks of respective leaf images” “to create cohesive respective application environments” “at an operational level between file systems”
'469 patent	24, 25, 26, 28, 32	Claim 24: “Internet Hotspot” “operatively coupled” “subscriber access unit” “authenticating”/“authenticate” “a remote location experiencing a relatively high volume of transient traffic” “web-ready device” Claim 25: “the data connection is one of a wired data connection and a wireless data connection” Claim 32: “the wireless connection is one of an 802.11a wireless area network, an 802.11b wireless area network, an 802.11g wireless area network, and an 802.11n wireless area network”
'326 patent	1, 4, 18	Claims 1 & 18: “full spectral synthesis capability” “combining”/“combine” “plurality of data subcarriers” “partially filling”/ “partially fill” “one or more guard bands” Claim 4 “transmitting to multiple radios”

F. Invalidity Under 35 U.S.C. § 101

The following Asserted Claims do not qualify as patent-eligible subject matter and are therefore invalid under 35 U.S.C. § 101.

Asserted Patent	Asserted Claim(s)
'582 patent	1, 2, 5, and 7
'722 patent	14, 16, 17, and 18

Asserted Patent	Asserted Claim(s)
'785 patent	30, 35, and 37
'844 patent	7, 11
'469 patent	24, 25, 26, 28, 32
'326 patent	1, 4 ,18

The Court has not yet construed the Challenged Claims. And IV has not yet provided a claim construction for any of the terms or phrases that American anticipates will be in dispute. American, therefore, cannot provide a complete list of § 101 defenses because American does not know whether IV will proffer a construction for certain terms and phrases that is broader than, or inconsistent with, the construction that would be supportable by the disclosure set forth in the specification. American offers these contentions without prejudice to any position they may ultimately take as to any claim construction issue. Accordingly, American reserves the right to supplement, amend, and/or modify these § 101 contentions as discovery progresses.

IV has not yet provided any subject matter eligibility contentions or otherwise disclosed its complete factual and legal basis in opposition to these Preliminary Invalidity Contentions. Further, to the extent related to subject matter eligibility, including describing the industry at the relevant time and each element of each Challenged Claim as found in the prior art, IV also has not provided any validity contentions under 35 U.S.C. §§ 102 or 103. IV also has not provided any expert report on the issue of subject matter eligibility, the inventors have not yet been deposed, and fact and expert discovery is at an early stage. American therefore reserves the right to further amend or supplement these Preliminary Invalidity Contentions in response to any further contentions, factual basis, legal basis, fact discovery, or expert discovery by or of IV asserting that the Challenged Claims are not ineligible under § 101, or that the prior art (as disclosed in these Preliminary Invalidity Contentions and otherwise) does not describe the industry, at the relevant time, or that any element of each Challenged Claim, both individually

and in combination with other elements, was not (i) well understood; (ii) routine; and (iii) conventional. This may include, but is not limited to, identifying additional prior art and other evidence regarding the state of the industry, providing expert testimony, and providing further factual and legal bases under ¶¶ (a)(2)(A) and (a)(2)(B) for why the Challenged Claims are invalid under 35 U.S.C. § 101.

American's § 101 contentions are made in the alternative, and they should not be interpreted, relied upon, or in any way affect the non-infringement or other invalidity arguments American has asserted. American does not infringe any of the Challenged Claims, which are also invalid under 35 U.S.C. §§ 102, 103, and/or 112.

American reserve the right to present expert evidence and testimony in further support of its contention that the Challenged Claims are ineligible under 35 U.S.C. § 101. Such testimony may include evidence and information regarding identifying each exception to eligibility, including why the Challenged Claims are directed to an abstract idea, as well as a description of the industry at the relevant time and what was well understood routine, and conventional in the industry. American specifically reserves the right to present expert testimony on any and all bases, and to do so in its opening expert report on invalidity in the time and manner required by the Court's Scheduling Order.

American incorporates by reference herein its Motion to Dismiss, including its supporting Reply. Additionally, American contends as follows:

1. The '582 patent

The Asserted Claims (claims 1, 2, 5, and 7) of the '582 patent are invalid under 35 U.S.C. § 101 because they are directed to an abstract idea—specifically, the idea of dividing a large task into subtasks and processing them using available resources on a first-come/first-served basis—and lack any inventive concept that transforms this abstract idea into patent-eligible subject

matter. The Asserted Claims of the '582 patent therefore fail at both steps of the Alice test and are therefore invalid. *See e.g. Alice Corp. v. CLS Bank Int'l*, 573 U.S. 208 (2014).

The '582 patent, titled “Method for Managing Distributed Processing of Tasks in a Network,” claims a method for subtask processing in distributed computing systems. Claim 1—the only independent claim of the '582 patent—recites:

1. A method of effecting on a **preexisting input file** a computer-executable process **comprised of a plurality of subtasks**, the method comprising the steps of:

(a) automatically determining file allocation and **logically subdividing records of said input file into a plurality of partitions**;

(b) **distributing descriptions of all of said partitions to each of a plurality of subtask processors**[:]

(c) **simultaneously executing** at least a respective **one of the subtasks** of the computer-executable process in each of at least some of said processors on a respective one of the partitions with each subtask reading and processing the respective partition so as to process the respective partition and produce respective subtask output and;

(d) thereafter **repeating step (c)** in at least some of the subtask processors each **with another unprocessed partition on a first-come/first-served basis**; and

(e) generating at least one output **combining all of the subtask outputs** and reflecting the processing of all of said subtasks.

The patent admits that these steps mostly were well known in the art, including partitioning data, distributing tasks, parallel processing, and aggregating results. '582 patent at 3:35-47, 3:67-4:15, 4:59-63. Partitioning (step a) relies on conventional techniques, such as byte ranges or track addresses, while subtask distribution (step b) employs generic mechanisms using control files; processing subtasks (step c) involves standard read-process-write operations, and aggregation (step e) is “of course” a conventional merging operation. *Id.* at 3:35-4:15, 4:59-63. The dependent Asserted Claims of the '582 patent fare no better. Claim 2 merely states that a

separate processor may perform the partitioning and distribution—an obvious and conventional division of labor that lacks any inventive concept. Claim 5 recites that subtask outputs can be accumulated in arbitrary order but offers no technical detail or mechanism for doing so—just a result. Claim 7 limits the input file to being stored on a network-attached storage device, which the patent itself identifies as a well-known technology.

The claimed first-come/first-served scheduling (step d) was argued during prosecution as the “primary difference” over the prior art. *See* ’582 patent, Prosecution History, Mar. 2, 2007 Remarks. According to the applicant, the prior art used “load information to distribute the load between processors,” whereas:

With the instant invention ... load information is not created Instead, the load sharing is done as a byproduct of the fact that the load-sharing process take parts of the load on a first-come/first-served basis. A comparison would be to a road intersection where, according to the prior art, there is a traffic light that determines who can go when. The instant invention is more like such an intersection with a four-way stop so that the individual drivers determine who can go and when.

Id. The patentee’s admission about the “primary difference” between the claimed invention and the prior art is critical in two ways. First, by pinpointing the “first available” assignment technique as the point of novelty, the inventors have defined what the claimed invention is “directed to” for purposes of the *Alice* Step 1 inquiry. *See Affinity Labs of Tex., LLC v. DIRECTV, LLC*, 838 F.3d 1253, 1257–58 (Fed. Cir. 2016) (“The ‘abstract idea’ step of the inquiry calls upon us to look at the ‘focus of the claimed advance over the prior art’ to determine if the claim’s ‘character as a whole’ is directed to excluded subject matter.”). Second, by the simple comparison between the invention and the managing of traffic loads at a street intersection, the inventor acknowledged that his “invention” merely applies a longstanding method of organizing human activities to computers.

The methods described in the '582 patent implement that abstract concept using generic computing components, such as processors and memory, which are standard elements in distributed systems. The claims use simple functional language to describe desired outcomes—partitioning, distributing, processing, scheduling, and aggregating—without detailing how these operations are performed or improved. Consequently, the Asserted Claims of the '582 patent fail at both steps of the Alice test and are therefore invalid.

a) The Asserted Claims of the '582 patent recite an abstract scheduling technique.

The Asserted Claims of the '582 patent are directed to the abstract idea of dividing a large task into subtasks and processing them using available resources on a first-come/first-served basis. Claim 1 of the '582 patent—the only independent claim of the '582 patent—comprises five steps: (a) “logically subdividing records . . . into a plurality of partitions,” (b) distributing descriptions of those partitions to multiple “subtask processors,” (c) “simultaneously executing” at least one “subtask” on each of the subtask processors, (d) thereafter allocating any remaining “unprocessed partition” to the subtask processors “on a first-come/first-served basis,” and (e) aggregating “all of the subtask outputs” into a final combined output. Stripped of jargon, Claim 1 covers the basic concept of splitting up a task and distributing the work, then collecting the results.

Claim 1’s distributed task management approach is an abstract idea that humans have long practiced without computers; it is akin to breaking a large job into smaller tasks, handing those tasks out to a group of workers, giving each new work as they free up, and then assembling the final product. The practice of breaking up a large job into discrete tasks, passing out a first round of assignments to team members, instructing them to ask for a new assignment when finished, and then combining the results has been employed by humans since the building of the

Egyptian pyramids. The Federal Circuit has consistently held that processes based on such fundamental organizational practices are abstract. *See, e.g., Intellectual Ventures I LLC v. Capital One Bank (USA)*, 792 F.3d 1363, 1367 (Fed. Cir. 2015) (organizing, storing, and retrieving information held abstract); *Intellectual Ventures I LLC v. Symantec Corp.*, 838 F.3d 1307 (Fed. Cir. 2016) (automatically organizing information into wanted and unwanted categories and discarding the unwanted information held abstract).

“[T]here is a critical difference between patenting a particular concrete solution to a problem and attempting to patent the abstract idea of a solution to the problem in general.” *Electric Power Group, LLC v. Alstom S.A.*, 830 F.3d 1350, 1356 (Fed. Cir. 2016). The Asserted Claims of the ’582 patent fall on the latter side of that line: they seek to monopolize the general idea of a solution (parallel task processing with on-the-fly scheduling) rather than a concrete technological solution. The Asserted Claims of the ’582 patent do not improve computer technology or solve a specific technical problem in how computers operate but instead apply conventional and common-sense workload management ideas to distributed computing. Partitioning data, distributing tasks, processing in parallel, and aggregating results are standard operations in distributed computing systems. ’582 patent at 3:35–4:15, 4:59–63. Simply executing a known business/workflow practice on generic computers is not a patentable improvement to those computers.

The ’582 patent’s recitation of assigning work using first-come/first-served scheduling—a ubiquitous method of task processing akin to queue management (e.g., a line at a concession stand or a four-way stop)—demonstrates its abstract nature. Claim 1 applies this principle to parallel computer processing without offering any new technological innovation. Moreover, The Asserted Claims of the ’582 patent use abstract functional language to describe desired

outcomes—such as “logically subdividing,” “distributing,” and “aggregating”—without detailing how these steps are performed. This kind of result-oriented claiming that specifies an outcome (efficient parallel processing) without limiting the claim to any innovative technique for achieving it is a hallmark of an abstract idea. Thus, the ’582 patent “merely applies a well-known idea using generic computers ‘to the particular technological environment of the Internet.’” *Intellectual Ventures I LLC v. Symantec Corp.*, 838 F.3d 1307, 1314 (Fed. Cir. 2016). In *Two-Way Media*, the Federal Circuit held that reciting a series of steps to achieve a desired result, without specifying how those steps are performed, renders the claim abstract. *See Two-Way Media Ltd. v. Comcast Cable Commc’ns, LLC*, 874 F.3d 1329, 1337–38 (Fed. Cir. 2017) (“routing” and “accumulating records” held abstract and invalid). Even though the specification describes conventional methods for these operations, the Asserted Claims of the ’582 patent do not specify how data is partitioned, how partitioned subtasks are distributed, or how results are aggregated. ’582 patent at 3:35–4:15, 4:59–63. Such high-level, results-oriented language underscores the claim’s abstract nature. As the Federal Circuit noted in *Electric Power*, claims that focus on the outcome of data processing without specifying the means to achieve it are abstract and invalid. 830 F.3d at 1356.

The ’582 patent are not directed to an improvement in data processing because using multiple computers in parallel may increase the speed of execution of processes and balances the load on resources. Dkt. 28 at 9-10. To the extent that the claimed method achieves these benefits, it only does so by applying the abstract idea of task partitioning and parallel execution on generic computers, rather than by reciting a new technology. Improved speed and load balancing are expected results of running a workload in parallel; they are performance outcomes, not inventive technical means. The ’582 patent claims do not improve the internal operation of the computer

processors or change how the network functions—they just use more computing power in a straightforward way. Merely accelerating an existing process via parallelization is akin to the scenario in *SAP America v. InvestPic*, where claims “selecting certain information, analyzing it using mathematical techniques, and reporting or displaying the results of the analysis” were held abstract. 898 F.3d 1161, 1167-1168 (Fed. Cir. 2018). Indeed, the claims of the ’582 patent “fit into the familiar class of claims that do not ‘focus ... on [] an improvement in computers as tools, but on certain independently abstract ideas that use computers as tools.’” *Id.* at 1168 (quoting *Electric Power*, 830 F.3d at 1354). Furthermore, IV’s reliance on *Enfish, LLC v. Microsoft Corp.*, 822 F.3d 1327 (Fed. Cir. 2016), is misplaced. Unlike *Enfish*, where the claims improved database memory structure, the ’582 patent does not specify any new hardware or improvement to system architecture—it merely applies a standard method of parallelization to a computing process.

Because the Asserted Claims of the ’582 patent of the ’582 patent are directed to a fundamental method for organizing and processing tasks, it falls squarely within the category of claims deemed ineligible under Step 1 of the *Alice* framework.

b) The ’582 patent lacks an inventive concept.

As to *Alice* Step 2, The Asserted Claims of the ’582 patent recite only conventional steps and generic components without introducing any meaningful technical improvements or innovations. The claim elements, individually and in combination, reflect routine operations in distributed computing, as explicitly acknowledged in the specification. Partitioning input data into subsets, distributing tasks to processors, and processing tasks in parallel were standard techniques for optimizing workload distribution at the time. ’582 patent at 3:35-4:15, 4:59-63.

The first-come/first-served scheduling technique, likewise, is a well-established method for processing tasks based on availability, like the applicant’s example of the four-way stop sign.

The Asserted Claims of the '582 patent “merely appl[y] [this] well-known idea using generic computers.” *Symantec*, 838 F.3d at 1314. The specification does not describe any modification, improvement, or adaptation of this scheduling technique to make it unique or unconventional in the context of distributed computing. The Asserted Claims of the '582 patent simply apply this known concept predictably to allocate unprocessed subtasks, a principle embedded in all multi-threaded and distributed computing architectures and offers no specific mechanisms or technical enhancements. Thus, the Asserted Claims of the '582 patent are “directed to a result or effect that itself is the abstract idea and merely invoke[s] generic processes and machinery” rather than “a specific means or method that improves the relevant technology.” *Yu v. Apple Inc.*, 1 F.4th 1040, 1043 (Fed. Cir. 2021). “What is claimed is simply a generic environment in which to carry out the abstract idea.” *Id.*

Confirming the routine nature of the claimed methods, both the specification and the Asserted Claims of the '582 patent rely entirely on generic computing components, such as processors, memory, and standard network architecture. '582 patent at 4:27–36; *Trinity Info Media, LLC v. Covalent, Inc.*, 72 F.4th 1355, 1364, 1366–67 (Fed. Cir. 2023) (generic computer components such as “processors” and “memory” insufficient to add an inventive concept). The patent does not propose any novel configurations or functionalities for these components but instead recites conventional “functions in general terms, without limiting them to technical means for performing the functions that are arguably an advance over conventional computer and network technology.” *Elec. Power*, 830 F.3d at 1351.

For example, Claim 1’s reliance on functional language—such as “automatically determining,” “logically subdividing,” “distributing,” “executing,” and “combining”—describes desired outcomes without specifying how those outcomes are achieved. It does not explain how

data is partitioned into subsets, how subtasks are allocated, how unallocated tasks are managed, or how results are aggregated into a final output. '582 patent at 6:47–4:15, 4:59–63. This lack of specificity leaves the claim framed at a high level of abstraction, with no concrete implementation details that could transform the method into a patent-eligible application. The dependent Asserted Claims of the '582 patent fare no better. Claim 2 merely delegates partitioning and distribution to another processor, without explaining how that processor performs the task. Claim 5 generically allows output in arbitrary order but lacks any technical detail on how that's achieved. Claim 7 limits the input file to network-attached storage—a conventional environment the patent admits was already well known.

Courts have consistently held analogous claims ineligible under Step 2 of the Alice framework. In *Content Extraction*, 776 F.3d at 1348, the court found “no ‘inventive concept’” in the use of generic computer components to perform routine activities. Similarly, in *Symantec Corp.*, 838 F.3d at 1320, the Federal Circuit emphasized that for a computer-implemented invention to be patent-eligible, it must involve more than well-understood, routine, and conventional activities. Like the claims in these cases, the Asserted Claims of the '582 patent merely apply abstract principles of task management to distributed computing using conventional techniques, which does not suffice to render it patent-eligible.

The patent's description of using heterogeneous computers and balancing load is not evidence that the claims capture a technical innovation. As detailed above, the combination of steps in the Asserted Claims of the '582 patent are entirely conventional—they are the standard recipe for parallel processing (split, distribute, process, repeat, merge). There is no inventive coordination or unconventional synergy arising from the arrangement of these steps; each step does what one would normally expect, and together they achieve the expected outcome of

parallel execution. The Asserted Claims of the '582 patent simply assemble known techniques to carry out an abstract goal. Combining conventional steps in a computer environment, without more, does not create an inventive concept. IV emphasizes that the patent describes using multiple heterogeneous computers and optimizing execution, but those aspects were already familiar in distributed computing.

Viewed as a whole, the Asserted Claims of the '582 patent combine conventional steps in a predictable manner to achieve the abstract goal of managing distributed tasks. They offer no technological advancements or inventive concepts beyond the abstract idea itself. Accordingly, the '582 patent fails to meet the requirements for patent eligibility under 35 U.S.C. § 101.

2. The '722 patent

The Asserted Claims (claims 14, 16, 17, and 18) of the '722 patent recite the abstract idea of receiving information, identifying relevant content, registering interest, and disseminating updates to interested parties. This is a long-standing and fundamental practice in organizing and distributing information—not a technological improvement or specific technical solution. Because the claims implement an abstract idea using only generic networking and data-handling functions without reciting any technological improvement or unconventional hardware, the claims fail both steps of the *Alice* framework.

a) The Asserted Claims of the '722 patent are Directed to an Abstract Idea

The Asserted Claims of the '722 patent are directed to the abstract idea of receiving information, identifying relevant content, registering interest, and disseminating updates to interested parties. Claim 14—the only representative independent claim—describes a method for enabling dynamic updates to “live objects” displayed on a client device through a structured communication framework involving a source device, a routing network, and intermediary

nodes. An input source sends a data representation to a client device via a routing network; this data includes at least one live object—an identifiable element capable of receiving updates. The client device recognizes the live object, determines its identifier, and registers with the routing network to receive updates about it. This registration includes connection information, which is stored at a node in the routing network. When the input source sends an update message for the live object, the message is routed through a gateway that determines its category and selects the appropriate node type to handle it. The node then identifies the registered client and forwards the message. Upon receipt, the client updates the relevant property of the live object.

Claim 16 adds to claim 14 that the client stores metadata related to the live object, and that the update message modifies this metadata. Claim 17 adds to claim 14 that the update message includes a message type, which the client uses to determine how to apply the update. Claim 18 expands on claim 14 by allowing the client to register for multiple live objects and receive updates for each.

What the Asserted Claims of the '722 patent effectively describe is a rules-based information forwarding mechanism, based on identifiers and interest registration—a framework that has existed in mailrooms, subscription services, and newsrooms long before the advent of digital networks. Its implementation using computers and a “routing network” merely automates this longstanding human practice.

Stripped of jargon, the Asserted Claims of the '722 patent describe a familiar, abstract pattern: sending information, having a recipient identify what is relevant, and distributing future updates to those who have expressed interest. This pattern is not new or technological. It is, for example, identical in substance to how a traditional newsroom operates. A journalist (the “input source”) provides a summary of several developing stories (the “data representation”) to a

central bulletin board in the newsroom. Each story on the board is a “live object” in the sense that it may be updated. Editors (the “client devices”) review the bulletin board and mark their interest in particular stories. When updates come in from the field, a receptionist (the “gateway device”) categorizes the update and determines which editor (the “node”) should receive it. The update is then delivered to the relevant editor, who revises the draft story accordingly.

This real-world newsroom example illustrates that the Asserted Claims of the ’722 patent merely recite an abstract method of organizing and delivering information—a practice that has been implemented manually for decades. The fact that the ’722 patent describes this process in the context of a computing environment does not render it any less abstract. Claims directed to abstract frameworks for information delivery—like updating content in response to user interaction or preference—have repeatedly been held patent-ineligible. *See, e.g., Broadband iTV, Inc. v. Amazon.com, Inc.*, No. 6:20-CV-00921-ADA, 2022 WL 4703425, at *11 (W.D. Tex. Sept. 30, 2022) (J. Albright), *aff’d*, 113 F.4th 1359 (Fed. Cir. 2024); *see also Broadband iTV, Inc. v. Oceanic Time Warner Cable, LLC*, 135 F. Supp. 3d 1175, 1178 (D. Haw. 2015) (claims directed to abstract idea of “using the same hierarchical ordering based on metadata to facilitate the display and locating of video content” ineligible under § 101). Because the Asserted Claims of the ’722 patent focus on the result of an implementing an abstract idea—delivering updated content to interested recipients—rather than a specific technological means of achieving that result, they fail at *Alice* step one.

The specification attempts to present these steps as an improvement over prior art by claiming to enhance network efficiency through dynamic client registration and message routing. It asserts that the use of node specialization and dynamic mapping reduces unnecessary message traffic. The routing network purportedly eliminates the need for the input source to track which

clients are displaying the live objects by maintaining a registry of registered clients. The system is described as being useful for applications where content on client devices must be updated dynamically based on changes in data, such as stock prices or sports scores. However, the specification itself reveals that the claimed methods are implemented using conventional networking techniques. The routing network, which consists of nodes and gateways, is built using generic computing components, including single-processor computer systems running common operating systems like Linux. The specification indicates that the system utilizes existing messaging protocols, such as HTTP, and known client-server architectures. Furthermore, the patent acknowledges that the basic concepts of client registration and update handling were already present in existing methods that included client-driven approaches, applets, and server-driven updates using TCP/IP connections.

b) The Asserted Claims of the '722 Patent Do Not Recite an Inventive Concept

Step 2 of *Alice* cannot save the Asserted Claims of the '722 patent from ineligibility under 35 U.S.C. § 101. The Asserted Claims of the '722 patent do no more than implement the abstract idea of selectively distributing content updates to interested recipients using generic computing components performing routine, conventional functions. The purported improvement—routing update messages based on client registration—does not introduce any new computing technology or method of enhancing network efficiency. Instead, it merely applies well-known client-server communication techniques to the problem of distributing data updates. The specification itself discloses that the system utilizes conventional networking components and known data transmission protocols without any significant modification.

Claim 14 recites basic components: an “input source,” a “client device,” a “routing network,” “nodes,” and “gateways.” But none of these are claimed in a specific or inventive

configuration. The specification describes the input source as a web server or dynamic content provider, the client device as a browser-enabled device such as a PC or mobile phone, and the routing network as an internet-based message delivery structure. The “gateway” and “node” elements, while given roles in the message flow, are described in general terms that mirror conventional content delivery mechanisms—such as proxies, relays, load balancers, or intermediate servers.

There is no technical detail in the claim about how messages are classified, how registration is implemented, or how update routing is performed. The claim merely recites the expected result: that a client receives an update for a live object it previously registered for. Courts have routinely held that such claims, which rely on routine network behavior and leave implementation to the practitioner, do not supply an inventive concept. *See, e.g., Two-Way Media Ltd. v. Comcast Cable Commc’ns, LLC*, 874 F.3d 1329, 1339 (Fed. Cir. 2017) (“The claim uses a conventional ordering of steps—first processing the data, then routing it, controlling it, and monitoring its reception—with conventional technology to achieve its desired result.”).

The specification itself confirms that the invention is designed to operate on existing infrastructure and leverages known technologies. Client devices are said to use standard web browsers with existing support for Java, JavaScript, or plug-ins. Update messages are transmitted over HTTP and are parsed using client-side scripts. The routing network is implemented using conventional clusters, nodes, and gateways, and messages are routed using category/type mappings or registry lookups—well-understood design patterns in networked systems. The so-called “live objects” are simply HTML elements or other page components that can be identified by an ID and updated via DOM manipulation or script execution, which common techniques in web development since at least the late 1990s. The use of an “activation module” to register

interest in updates is just another name for client-side code that subscribes to a push service—an architecture long practiced in content syndication and notification systems. Nothing in Claim 14 requires, or even enables, a novel technical implementation of these mechanisms.

None of the limitations in claims 16, 17, or 18 inject any inventive concept into the '722 patent:

- **Claim 16's** mention of metadata storage and update is a routine feature of client-side state tracking. Web applications and software clients have long-maintained associated metadata (e.g., timestamps, tags, formatting) alongside core content.
- **Claim 17** introduces a message type field used to determine how to process an update. But conditional logic based on message types is a staple of basic programming and messaging protocols. This is no more inventive than a “switch” statement tied to a message header.
- **Claim 18** expands registration to multiple live objects, which is a straightforward extension of the core idea. There is no novel mechanism for batch registration, grouping, or prioritization; the claim merely allows more than one subscription, a feature inherent in any scalable publish-subscribe system.

Thus, the dependent claims likewise apply well-known programming constructs and user-experience features without claiming how any of them are implemented in a novel or unconventional manner.

Furthermore, the combination of client registration and dynamic message routing was not an inventive concept at the time of the alleged invention. Technologies such as IP Multicast, Java Message Service (JMS), and CORBA Event Services were already implementing message routing based on client registration well before the priority date. The combination of these

known elements does not produce unexpected results or provide a technological improvement. Here, the claimed method merely takes existing client registration and message routing techniques and applies them in a straightforward, predictable way.

The Asserted Claims of the '722 patent therefore do not improve any specific computer functionality or solve any technical problem in a new way—they merely implement a conventional information delivery model using abstract functional language and general-purpose computing components. As such, the Asserted Claims of the '722 patent fail *Alice* step two and are invalid under 35 U.S.C. § 101.

* * *

The Asserted Claims of the '722 patent are invalid under 35 U.S.C. § 101 because they fail both prongs of the *Alice* framework. At step one, these claims are directed to the abstract idea of distributing updates about selected content to interested recipients—an organizing principle that has long been practiced by humans and implemented in conventional information systems. These claims describe this idea at a high level of generality, using functional language untethered to any specific improvement in technology or network design. At step two, the claims do not recite any inventive concept sufficient to transform this abstract idea into a patent-eligible application. The components are generic, the operations are conventional, and the implementation relies entirely on known computing techniques. These claims simply automate a familiar content delivery paradigm using routine architecture, offering no technical advancement over the prior art.

3. The '844 patent

The Asserted Claims (claims 7, 11) of the '844 patent recite the abstract idea of organizing and storing information using a “root-leaf” structure. This is a long-standing and fundamental practice in organizing and distributing information—not a technological

improvement or specific technical solution. Because the claims implement an abstract idea using only generic networking and data-handling functions without reciting any technological improvement or unconventional hardware, the claims fail both steps of the *Alice* framework.

a) The Asserted Claims of the '844 Patent are Directed to an Abstract Idea

The Asserted Claims of the '844 patent are directed to the abstract idea of organizing and storing information using a “root-leaf” structure. Although applied to a computer network, the method steps of independent claim 7—storing root and leaf images and caching accessed blocks—mirror basic data collection and organization practices, analogous to longstanding human activities. The patent attempts to claim the concept of storing some information in a central repository for permanent storage and common use (a root image), storing other non-identical data locally for faster retrieval and personal use (leaf images), and keeping recently used items temporarily on hand (cached). Such abstract ideas for how to organize and store data on computer systems are one of “the ‘basic tools of scientific and technological work’ that are free to all men and reserved exclusively to none.” *Synopsys, Inc. v. Mentor Graphics Corp.*, 839 F.3d 1138, 1146 (Fed. Cir. 2016) (quoting *Alice*, 573 U.S. at 216). Courts have consistently held such practices abstract. *See Intellectual Ventures I LLC v. Erie Indem. Co.*, 850 F.3d 1315, 1327 (Fed. Cir. 2017) (organizing and accessing data held abstract); *Content Extraction & Transmission LLC v. Wells Fargo Bank, Nat’l Ass’n*, 776 F.3d 1343, 1347 (Fed. Cir. 2014) (data collection and storage held abstract); *Kaavo Inc. v. Amazon.com Inc.*, 323 F. Supp. 3d 630, 641 (D. Del. 2018) (“setting up and managing a cloud computing environment” held abstract); *Versata Software, Inc. v. NetBrain Techs., Inc.*, No. 13-676-LPS-CJB, 2015 WL 5768938, at *7 (D. Del. Sept. 30, 2015) (“representing information in a hierarchy amounts to an abstract idea.”).

The specification confirms that the described components and techniques, including root-leaf storage systems and caching mechanisms, were well understood at the time of the invention. Dkt. 1-1 at 1:31–45, 6:37–7:5. The Asserted Claims of the '844 patent provide no specifics on storage structures/hardware, caching algorithms/techniques, or other technical improvements, relying instead on functional language that describes high-level results without specific implementation details. *See ChargePoint, Inc. v. SemaConnect, Inc.*, 920 F.3d 759, 769 (Fed. Cir. 2019) (“Even a specification full of technical details about a physical invention may nonetheless conclude with claims that claim nothing more than the broad law or abstract idea underlying the claims.”).

Claim 7’s scant details about what data gets saved where are inherently abstract. *See Content Extraction*, 776 F.3d at 1347 (“The concept of data collection, recognition, and storage is undisputedly well-known. Indeed, humans have always performed these functions.”). “This type of activity, *i.e.*, organizing and accessing [data], includes longstanding conduct that existed well before the advent of computers and the Internet.” *Erie Indem.*, 850 F.3d at 1327. Claim 7 includes no purportedly novel technical details for cluster computing—just an abstract idea for how to organize and store data. In the end, the claim is little different from a librarian storing different types of books and papers in different sections of a library. Claim 7 therefore fails *Alice* Step 1 because it recites a conventional root-leaf storage system using a high-level abstract concept to organize and store data without any specific technological improvement recited in the claim. Claim 11 fares no better; it merely adds a generic “merging” step—joining the blocks of the root and leaf images to “create cohesive respective application environments”—without describing *how* that merging is performed or any technical rules guiding it. Claim 11 therefore provides no new or inventive mechanism for merging, just the result.

b) The Asserted Claims of the '844 Patent Do Not Recite an Inventive Concept

The Asserted Claims of the '844 patent likewise fail Step 2, in part, because “these claims use generic computers to perform generic computer functions.” *Symantec*, 838 F.3d at 1315. The specification concedes that the claimed methods are standard, including storing root images on a first storage unit, leaf images on second storage units, and caching accessed blocks. Dkt. 1-1 at 2:14–21, 6:38–67. Thus, even if the specification discloses ideas that could be applied to useful ends, claims 7 and 11 “contain[] no restriction on how the result is accomplished.” *Symantec*, at 1316.

The claimed steps are described at a high level of generality, with no new algorithms, configurations, or technological improvements. Generic terms like “providing,” “storing,” and “caching” describe desired outcomes without specifying any “inventive concept” for how these tasks are performed. *See BSG Tech*, 899 F.3d at 1290-91 (“If a claim’s only ‘inventive concept’ is the application of an abstract idea using conventional and well-understood techniques, the claim has not been transformed into a patent-eligible application of an abstract idea”); *Elec. Power Grp., LLC v. Alstom S.A.*, 830 F.3d 1350, 1356 (Fed. Cir. 2016) (drawing the “distinction between ends sought and particular means of achieving them, between desired results (functions) and particular ways of achieving (performing) them”).

Breaking down each element of the Asserted Claims of the '844 patent reveals how each element is conventional:

Claim 7:

- **“storing blocks of a root image of said compute nodes on a first storage unit”**

This is conventional. Shared read-only root images were routinely stored on

central storage devices, as used in Sun's diskless workstations and LTSP environments. (See Hitz et al.; NetApp WAFL).

- **“storing leaf images for respective compute nodes on respective second storage units...”** This element is likewise conventional. Systems such as VMware and Episode stored deltas or overlays for individual machines using CoW. These deltas excluded unchanged base data, which was referenced from the shared root image.
- **“caching blocks of said root image that have been accessed...”** This element is likewise conventional. Caching mechanisms, including block-level caching and LRU algorithms, were widely used by 2004 in both operating systems and storage appliances (see NetApp; LTSP).

Claim 11:

- **“merging said blocks of the root image and said leaf images...”**

The element recited by Claim 11 is similarly conventional and routine. The merging recited by Claim 11 had been widely implemented in systems such as VMware's and NetApp's platforms.

Courts have invalidated analogous claims. In *Symantec Corp.*, 838 F.3d 1307, the Federal Circuit invalidated claims for managing email messages using conventional computer components. Similarly, the court in *Content Extraction*, 776 F.3d 1343, held claims for data recognition and storage invalid because they relied on routine computing methods. The Asserted Claims of the '844 patent fall into the same category, as they merely recite routine computing techniques and hardware.

The Asserted Claims of the '844 patent offer no inventive step, combining conventional techniques predictably to achieve abstract goals. They rely on generic computing elements and routine data management practices, failing *Alice* Step 2. Thus, The Asserted Claims of the '844 patent are ineligible under §101.

* * *

Nonetheless, at their core, claims 7 and 11 are directed to the abstract idea of organizing and storing data using a hierarchical structure of a shared base and individualized overlays. The '844 patent relies entirely on general-purpose computing devices with generic storage media (Dkt. 1-1 at 4:27–56), and unclaimed components, such as “Union Block Devices” (UBDs), described as low-level drivers with no disclosed novel functionality (*id.* at 5:19–64). The Asserted Claims of the '844 patent likewise employ broad functional language like “providing,” “storing,” and “caching,” without specifying any improved algorithms, hardware, or software. Thus, the Asserted Claims of the '844 patent merely recite an abstract idea for organizing and storing data without any inventive concept, thereby failing at both steps of the *Alice* test.

4. The '785 patent

The Asserted Claims (claims 30, 35, and 37) of the '785 patent recite the abstract idea of membership management and address look-up in a communications network. This is a long-standing and fundamental practice in administrative tasks—not a technological improvement or specific technical solution. Because the claims implement an abstract idea using only generic administrative functions without reciting any technological improvement or unconventional hardware, the claims fail both steps of the *Alice* framework.

a) The Asserted Claims of the '785 Patent are Directed to an Abstract Idea

The Asserted Claims of the '785 patent are directed to the abstract idea of membership management and address look-up in a communications network—mere administrative tasks performed every day by receptionists, security personnel, and bouncers. Stripped of jargon, the claimed elements in independent claim 30 are basic administrative functions involved in setting up and maintaining a private communication network—operations that map directly onto longstanding abstract concepts such as maintaining a directory of users, assigning identifiers, and responding to queries with directing information. Claim 35 builds on claim 30 by adding a “join module” that receives join and leave requests—analogueous to a sign-in/sign-out log for network members. Claim 37 further adds that the join module maintains a mapping between virtual addresses and devices—a conventional association table akin to a user directory or ARP (address resolution protocol) table.

The Asserted Claims of the '785 patent all reflect the abstract idea of membership management and address look-up in a communications network—akin to a corporate office receptionist managing an internal directory. The receptionist receives visitors (devices) requesting access (registration), checks whether they're authorized to enter (authentication), assigns them a visitor badge with an internal extension number (virtual address), and, when asked, tells other employees where to find the visitor (DNS resolution and routing). The receptionist also tracks when the visitor checks in and out (join and leave requests) and keeps a temporary log mapping badge numbers to visitors (address-device association). These are routine administrative functions that merely facilitate communication within a defined space—just as the Asserted Claims do within a virtual network. There is no innovation in how the receptionist performs these tasks, and these tasks have been performed by humans for centuries. Likewise,

the Asserted Claims do not improve any underlying networking technology or propose any inventive technical implementation. The Asserted Claims of the '785 patent are simply automating a familiar organizational role using conventional computing tools, rendering them abstract under *Alice* step one.

Tellingly, the method of originally filed claim 64 (which issued as claim 48) did not include any requirement that the steps of the claim be performed by a computer. On February 4, 2009 (5 years before the *Alice* decision), the Examiner rejected claim 64 under 35 U.S.C. § 101, applying the *Bilski* “machine-or-transformation” test for patentable subject matter that existed in 2009. In other words, without the recitation of a “machine” (*i.e.*, a computer), the claims were not patentable under § 101. In response to the Section 101 rejection, the Applicants amended claim 64 to be “computer-implemented,” as none of the originally-filed claims required the use of a computer. Although an amendment requiring that the method be performed by a computer was sufficient in 2009-2010 to overcome the Examiner’s Section 101 rejections, this is not the case today. Under current Supreme Court and Federal Circuit law, the *Alice* steps apply.

The '785 patent applicants’ computer-implemented amendment to independent claim 64 crystallizes two points. First, the amendment demonstrates that the Applicants understood that the method described in the '785 patent and claimed in the originally-filed claim 64 could be practiced without the use of a computer, *i.e.*, the method could be performed in the human mind, or by a human using a pen and paper. *See CyberSource Corp. v. Retail Decisions, Inc.*, 654 F.3d 1366, 1372–73 (Fed. Cir. 2011) (holding that a claim whose “steps can be performed in the human mind, or by a human using a pen and paper” is directed to an “unpatentable mental process[]”); *Elec. Power Grp.*, 830 F.3d at 1355 (“[C]laims focusing on organizing, characterizing, and/or sorting information are patent-ineligible because they are not significantly

differentiable from ordinary mental processes.”). Second, this amendment demonstrates that the focus of the claims of the ’785 patent is not on any specific asserted improvement in computer capabilities. *In re TLI Commc’ns L.L.C. Pat. Litig.*, 823 F.3d at 612 (“[T]he claims were not directed to specific improvement to computer functionality. Rather, they are directed to the use of conventional or generic technology in a nascent but well-known environment, without any claim that the invention reflects an inventive solution to any problem presented by combining the two.”).

Under *Alice*, a claim that is unpatentable without the recitation of a computer does not become patentable by reciting a computer. That the claims of the ’785 patent merely implement the non-statutory subject matter on a generic computing device is plain from the fact that the originally-filed claims did not require that the claimed systems and methods include or be performed by a computer, respectively. Thus, the language of the Asserted Claims, their prosecution history, and the change in the law demonstrate that the Asserted Claims of the ’785 patent are directed to an abstract administrative idea applied to a communications network.

b) The Asserted Claims of the ’785 Patent Do Not Recite an Inventive Concept

The Asserted Claims of the ’785 patent fail at *Alice* step 2 because they do not recite anything beyond generic computer functions performed using conventional components. The specification confirms that the system is built using well-known networking elements (*e.g.*, processors, memory, interfaces, DNS servers, NAT traversal techniques, virtual IP addressing) without any specialized circuitry, protocol, or configuration.

The ’785 patent itself acknowledges that the components and methods it employs were conventional at the time of invention. For example, the specification discusses the use of Network Address Translation (NAT) and virtual addresses to facilitate secure communication,

but these techniques were already well-known in the field of networking. The patent does not propose any novel way of utilizing NAT or managing virtual addresses beyond their ordinary and conventional uses.

Claim 30, for instance, merely describes using a DNS server to resolve device addresses within a virtual network, which is a standard function of DNS servers generally. The idea of assigning virtual addresses to devices and resolving these addresses through DNS lookups is not unique or inventive. Instead, the patent recites high-level results—such as secure communication—without specifying how these results are technologically achieved. There is no disclosed protocol, no new addressing scheme, no algorithm for conflict resolution, and no structural innovation in how devices are registered or managed. The system instead relies on existing IP, UDP, NAT, and IPsec technologies. The claims are results-oriented: “register the device,” “assign an address,” “respond to a DNS query”—without any accompanying inventive technical solution.

The claim elements—processor, memory, network interface, modules—are standard hardware or logical software components performing their expected functions. There is no claim that the invention requires new hardware, reconfigured components, or even custom protocols. The Asserted Claims of the ’785 patent are therefore equivalent to the claims in *BSG Tech LLC v. BuySeasons, Inc.*, 899 F.3d 1281, 1291 (Fed. Cir. 2018), where the claims at issue merely limited an abstract idea to a particular technological environment and were held ineligible. Similarly, in *ChargePoint, Inc. v. SemaConnect, Inc.*, the court invalidated claims that combined generic networking capabilities with existing concepts without adding any inventive technology. 920 F.3d 759, 774 (Fed. Cir. 2019) (“But network control is the abstract idea itself, and ‘a claimed invention’s use of the ineligible concept to which it is directed cannot supply the

inventive concept that renders the invention ‘significantly more’ than that ineligible concept.”) (quoting *BSG Tech*, 899 F.3d at 1290). The Asserted Claims of the ’785 patent follow the same pattern by applying well-understood networking techniques without advancing the state of the art.

* * *

The Asserted Claims of the ’785 patent are invalid under 35 U.S.C. § 101. They are directed to abstract administrative operations for managing virtual networks and fail to recite any inventive concept in either their individual components or overall arrangement. They do not improve any underlying technology, nor do they disclose a novel technical mechanism for achieving the stated goals. As such, they fail both steps of the *Alice* test and are ineligible for patent protection.

5. The ’326 patent

a) The Asserted Claims of the ’326 Patent are Directed to an Abstract Idea

The Asserted Claims of the ’326 patent recite the abstract idea of combining channels to increase data throughput (claims 1 and 18) and transmitting data to multiple radios (claim 4). These claims fail Step 1 of *Alice*. See *Alice*, 573 U.S. at 216-17. They recite long-standing and fundamental practices in sending and receiving information over a network—not a technological improvement or specific technical solution. There is nothing to the patent beyond the concept of combining two data transmission channels together, including the space between them, to yield a larger channel.

Channel bonding and transmitting data to multiple radios is an abstract idea as claims directed to sending and receiving data are abstract in nature. See, e.g., *Affinity Labs of Texas, LLC. v. DIRECTTV, LLC*, 838 F.3d 1253, 1261 (Fed. Cir. 2016) (explaining that patents directed

to “the conveyance and manipulation of information using wireless communication and computer technology” do not disclose eligible subject matter). Claims directed to formatting data are also abstract. *See Dropbox, Inc. v. Synchronoss Techs., Inc.*, 815 F. App’x 529, 537 (Fed. Cir. 2020) (finding claims ineligible that were directed, in part, to “formatting data”). Additionally, the Asserted Claims of the ’326 patent are directed towards manipulating data subcarriers, filling frequency gaps, and operations such as Fast Fourier Transforms, all of which are fundamentally mathematical operations. The Asserted Claims are thus invalid, as “[p]urely conventional or obvious pre-solution activity is normally not sufficient to transform an unpatentable law of nature into a patent-eligible application of such a law.” *Mayo Collaborative Services v. Prometheus Labs., Inc.*, 566 U.S. 66, 79, (2012) (cleaned up); *see also Digitech Image Techs., LLC v. Elecs. for Imaging, Inc.*, 758 F.3d 1344, 1351 (Fed. Cir. 2014) (“Without additional limitations, a process that employs mathematical algorithms to manipulate existing information to generate additional information is not patent eligible.”).

b) The ’326 Patent Claims Lack an Inventive Concept.

Step 2 of *Alice* does save the Asserted Claims of the ’326 patent from ineligibility under 35 U.S.C. § 101. The Asserted Claims of the ’326 patent do no more than implement the abstract idea of increasing data throughput by using generic computing components performing routine, conventional functions. Namely, channel bonding (claims 1 and 18) and transmitting data to multiple radios (claim 4) lack an inventive concept because such claims amount only to generic components and fail to demonstrate unconventionality. *See Dropbox*, 815 F. App’x at 537 (finding that “data transfer connection[s] are generic computer-related concepts as they literally are just data transfer modalities” and are thus not inventive concepts).

The Asserted Claims of the ’326 patent are invalid under 35 U.S.C. § 101. They are directed to the abstract idea of transferring data over a network and fail to recite an inventive

concept. Because the claims implement an abstract idea using only generic networking and data-handling functions without reciting any technological improvement or unconventional mechanisms, the claims fail both steps of the Alice framework.

6. '469 patent

a) The Asserted Claims of the '469 Patent recite the business practice of placing a paid hotspot in a “transient location”

Each of the Asserted Claims of the '469 patent are invalid under 35 U.S.C. § 101 because the claims are directed to an abstract idea and fail to describe an inventive concept. *Alice Corp. v. CLS Bank Int'l*, 573 U.S. 208, 216-17 (2014). For example, all of the Asserted Claims of the '469 patent are generally directed to an abstract idea for Section 101 purposes and under Intellectual Venture's apparent constructions. The sole asserted independent claim, claim 24, concerns placing a known composition of commercial components in a “remote location” to establish a subscribers-only internet hotspot. The specification also makes clear that the invention of the Asserted patent is the idea of placing paid hotspots where users would otherwise not have Internet access:

The present invention provides rural “Hotspots” (such as Wi-Fi access, for example) to enable wireless and hardwired, satellite distributed Internet access for anyone with a PC or other web-ready device (wireless ready or cabled) and a valid credit card.

'469 patent at 1:34-38. The concept of charging people money for internet access in a “transient location” like a restaurant or motel is merely a business idea. Paying for services is a “fundamental economic practice long prevalent in our system of commerce” and, “as such, is a patent ineligible abstract idea.” *W. Express Bancshares, LLC v. Green Dot Corp.*, 816 F. App'x 485, 487 (Fed. Cir. 2020).

Applying this fundamental economic practice to tangible components such as commercially available network equipment does not alter the fact that the claims are directed to

“an art or principle in the abstract, and not for any particular method or machinery.” *ChargePoint, Inc. v. SemaConnect, Inc.*, 920 F.3d 759, 769-70 (Fed. Cir. 2019) (quoting *Wyeth v. Stone*, 30 F. Cas. 723 (C.C.D. Mass. 1840)) (“In short, the inventors here had the good idea to add networking capabilities to existing charging stations to facilitate various business interactions. But that is where they stopped, and that is all they patented. We therefore hold that claim 1 is ‘directed to’ an abstract idea”). The patentee’s implementation rests wholly on the existing knowledge of a person of skill in the art: “The Hotspot 10 comprises all the equipment installed at the rural location that is necessary to provide Internet access via satellite.” ’469 patent at 4:18-21 (emphasis added). “The Ethernet wireless access point 40 can employ a security protocol of a type that is known to persons skilled in the art, such as the 802.11b security protocol, or any subsequent or future versions of the 802.11 standard.” ’469 patent at 5:19-25 (emphasis added). The Federal Circuit has “repeatedly found the concept of controlling access to resources via software to be an abstract idea.” *Ericsson Inc. v. TCL Commc’n Tech. Holdings Ltd.*, 955 F.3d 1317, 1327 (Fed. Cir. 2020); see also *Smartflash LLC v. Apple Inc.*, 680 F. App’x 977 (Fed. Cir. 2017).

b) The Asserted Claims of the ’469 patent lack an inventive concept.

The specification provides no solution to overcome technical challenges associated with “remote”-ness, or a “high volume of transient traffic” or “subscriber access.” The claims of the Asserted Patent are results-driven and fail to overcome ineligibility. Each limitation recites generic components which are “capable of” performing connectivity, routing, and authorization using any method whatsoever. As set forth above and in the attached claim charts, paid Internet hotspots were very well known and widely used at the time of the Asserted Patent. Placing a paid Internet hotspot in “a remote location” is insufficient to render the claims eligible where no

particular improvement is recited to, for example, teach how a subscriber access unit is “capable of” authorizing payments in “a remote location.” *Universal Secure Registry LLC v. Apple Inc.*, 10 F.4th 1342, 1346 (Fed. Cir. 2021); *Elec. Power Grp., LLC v. Alstom S.A.*, 830 F.3d 1350, 1354 (Fed. Cir. 2016) (“limiting the claims to the particular technological environment of power-grid monitoring is, without more, insufficient to transform them into patent-eligible applications of the abstract idea at their core.”). The “essentially result-focused, functional character of claim language has been a frequent feature of claims held ineligible under § 101, especially in the area of using generic computer and network technology to carry out economic transactions.”

Each of the asserted dependent claims likewise recite the use of conventional components or processes to achieve results for which those components are designed. ’469 patent, claim 25 (“a wired data connection” “a wireless data connection”); claim 26 (“a plurality of users may access the Internet simultaneously,” “at the remote location,” “establishing data connections with the router via their web-ready devices”); claim 28 (“the data connections include wireless data connections”); claim 32 (“the wireless connection,” “802.11a wireless area network,” “802.11b wireless area network,” “802.11g wireless area network,” “802.11n wireless area network”). See also *CyberSource Corp. v. Retail Decisions, Inc.*, 654 F.3d 1366, 1370-73 (Fed. Cir. 2011) (a method of verifying the validity of credit card transactions over the Internet held patent-ineligible as directed to an abstract idea); *buySAFE, Inc. v. Google, Inc.*, 765 F.3d 1350, 1354-55 (Fed. Cir. 2014) (guaranteeing performance in an Internet transaction was patent-ineligible); *TLI Communications LLC Patent Litigation*, 823 F.3d 607, 615 (Fed. Cir. 2016) (“vague, functional descriptions of server components are insufficient to transform the abstract idea into a patent-eligible invention.”). The claimed steps are described at a high level of generality, using generic terms, confirming the ’469 is ineligible under 35 U.S.C. §101.

G. Invalidity Based on Obviousness-Type Double Patenting

1. Obviousness-Type Double Patenting of the '844 Patent in View of U.S. Patent No. 7,721,282

American provides further notice to Plaintiffs that the asserted claims of the '844 patent are invalid under the doctrine of obviousness-type double patenting ("OTDP") in view of claim 15 of U.S. Patent No. 7,721,282 (the "'282 patent"). Because the asserted claims of the '844 patent are directed to the same invention recited by claims presented in the '282 patent, and because any differences in the asserted claims of the '844 patent are, at most, obvious modifications of claims in the '282 patent, the asserted '844 patent claims are invalid. *See AbbVie Inc. v. Mathilda & Terrence Kennedy Inst. of Rheumatology Tr.*, 764 F.3d 1366, 1373 (Fed. Cir. 2014).

The '282 patent is available as an OTDP reference against the '844 patent. The '282 patent was filed on March 30, 2006, issued on May 18, 2010, and is set to expire on December 19, 2027. The '844 patent was filed on February 21, 2007, issued on December 11, 2012, and is set to expire on April 6, 2028. Thus, the '282 patent was filed earlier than the '844 patent, issued earlier than the '844 patent, and is set to expire earlier than the '844 patent.

As is the case here, where "the applicant chooses to file separate applications for overlapping subject matter . . . the doctrine of obviousness-type double patenting ensures that a particular invention (and obvious variants thereof) does not receive an undue patent term extension." *AbbVie*, 764 F.3d at 1373. Therefore, the earlier-filed, earlier-issued, and earlier-expiring '282 patent qualifies as a double patenting reference.

Given that the '844 patent and the reference '282 patent share at least one common inventor and common ownership, determining whether a patent claim is invalid for OTDP entails two steps. *See id.* at 1374. The first step determines the differences between the respective claims

of the two patents. The second step determines whether those differences render the later-expiring set of claims sufficiently different to be patentably distinct. *Id.* “The second part of this analysis is analogous to the obviousness inquiry under 35 U.S.C. § 103 in the sense that if an earlier claim renders obvious or anticipates a later claim, the later claim is not patentably distinct and is thus invalid for obviousness-type double patenting.” *UCB, Inc. v. Accord Healthcare, Inc.*, 890 F.3d 1313, 1323 (Fed. Cir. 2018) (citation omitted). The disclosure of the reference patent may be used to determine whether the claims merely define an obvious variation of what is disclosed and claimed in the reference patent. *Abbvie*, 764 F.3d at 1381.

The attached claim chart in Exhibit 844-OTDP-1 demonstrates that, to the extent there are any differences between the asserted claims of the ’844 patent and claims of the ’282 patent, these differences are not patentably distinct. The charts provide reasons why a person of ordinary skill in the art would conclude that the alleged invention defined in the asserted claims of the ’844 patent would have been an obvious variation of the alleged invention defined in the corresponding claim or claims of the ’282 patent.

2. Obviousness-Type Double Patenting of the ’844 Patent in View of U.S. Patent No. 7,870,106

American provides further notice to Plaintiffs that the asserted claims of the ’844 patent are invalid under the doctrine of obviousness-type double patenting (“OTDP”) in view of claim 1 of U.S. Patent No. 7,870,106 (the “’106 patent”). Because the asserted claims of the ’844 patent are directed to the same invention recited by claims presented in the ’106 patent, and because any differences in the asserted claims of the ’844 patent are, at most, obvious modifications of claims in the ’106 patent, the asserted ’844 patent claims are invalid. *See AbbVie Inc. v. Mathilda & Terrence Kennedy Inst. of Rheumatology Tr.*, 764 F.3d 1366, 1373 (Fed. Cir. 2014).

The '106 patent is available as an OTDP reference against the '844 patent. The '106 patent was filed on February 2, 2006, issued on January 11, 2011, and is set to expire on May 18, 2027. The '844 patent was filed on February 21, 2007, issued on December 11, 2012, and is set to expire on April 6, 2028. Thus, the '106 patent was filed earlier than the '844 patent, issued earlier than the '844 patent, and is set to expire earlier than the '844 patent.

As is the case here, where “the applicant chooses to file separate applications for overlapping subject matter . . . the doctrine of obviousness-type double patenting ensures that a particular invention (and obvious variants thereof) does not receive an undue patent term extension.” *AbbVie*, 764 F.3d at 1373. Therefore, the earlier-filed, earlier-issued, and earlier-expiring '106 patent qualifies as a double patenting reference.

Given that the '844 patent and the reference '106 patent share at least one common inventor and common ownership, determining whether a patent claim is invalid for OTDP entails two steps. *See id.* at 1374. The first step determines the differences between the respective claims of the two patents. The second step determines whether those differences render the later-expiring set of claims sufficiently different to be patentably distinct. *Id.* “The second part of this analysis is analogous to the obviousness inquiry under 35 U.S.C. § 103 in the sense that if an earlier claim renders obvious or anticipates a later claim, the later claim is not patentably distinct and is thus invalid for obviousness-type double patenting.” *UCB, Inc. v. Accord Healthcare, Inc.*, 890 F.3d 1313, 1323 (Fed. Cir. 2018) (citation omitted). The disclosure of the reference patent may be used to determine whether the claims merely define an obvious variation of what is disclosed and claimed in the reference patent. *Abbvie*, 764 F.3d at 1381.

The attached claim chart in Exhibit 844-OTDP-2 demonstrates that, to the extent there are any differences between the asserted claims of the '844 patent and claims of the '106 patent,

these differences are not patentably distinct. The charts provide reasons why a person of ordinary skill in the art would conclude that the alleged invention defined in the asserted claims of the '844 patent would have been an obvious variation of the alleged invention defined in the corresponding claim or claims of the '106 patent.

American reserves the right to assert any other double patenting challenges against any of the Asserted Patents.

III. ACCOMPANYING DOCUMENT PRODUCTION

Pursuant to P.R. 3-4(a), American is producing documentation in its possession that shows the operation of any aspects or elements of each Accused Instrumentality identified by IV in its P.R. 3-1(c) charts. *See, e.g.*, AA_IV_0002168 – AA_IV_0074538. However, American does not make any of the technologies accused of infringement. Moreover, for certain allegations, American has not identified any allegedly infringing use of the technologies based on IV's infringement allegations. American has collected certain source code in its possession relevant to the accused technologies and will make that code available pursuant to the source code provisions of the Protective Order, once entered by the Court. However, American does not represent that any of the code in its possession is relevant to IV's infringement allegations (to the extent that those allegations can be understood to identify an accused product).

Additionally, pursuant to P.R. 3-4(b), American is producing herewith a copy of each prior art reference identified above pursuant to P.R. 3-3(a) which does not appear in the file histories of the Asserted Patents, as well as additional prior art references. To the extent any such prior art reference is not in English, American has produced to IV or will produce to IV an English translation of the portion(s) of the non-English prior art reference(s) relied upon by American.

Dated: June 20, 2025

/s/ John B. Campbell

John B. Campbell
Texas State Bar No. 24036314
jcampbell@McKoolSmith.com
Kenneth M. Scott
Texas State Bar No. 24137497
kscott@McKoolSmith.com
McKool Smith, P.C.
303 Colorado Street Suite 2100
Austin, TX 78701
Telephone: (512) 692-8700
Telecopier: (512) 692-8744

Emily Tannenbaum
New York State Bar No. 5928130
etannenbaum@mckoolsmith.com
McKool Smith, P.C.
1301 Avenue of the Americas, 32nd Floor
New York, NY 10019
Telephone: (212) 402-9400
Telecopier: (212) 402-9444

Casey L. Shomaker
Texas State Bar No. 24110359
cshomaker@mckoolsmith.com
McKool Smith, P.C.
300 Crescent Court, Suite 1500
Dallas, TX 75201
Telephone: (214) 978-4000
Facsimile: (214) 978-4044

Alan P. Block
California State Bar No. 143783
ablock@mckoolsmith.com
McKool Smith, P.C.
300 South Grand Avenue, Suite 2900
Los Angeles, CA 90071
Telephone: (213) 694-1054
Telecopier: (213) 694-1234

ATTORNEYS FOR DEFENDANT
AMERICAN AIRLINES, INC.

CERTIFICATE OF SERVICE

I hereby certify that a true and correct copy of the above and foregoing document has been served on all counsel of record via email on June 20, 2025.

/s/ John B. Campbell
John B. Campbell

Exhibit D-1

Invalidity of U.S. Patent No. 8,332,844 in View of “Optimizing the Migration of Virtual Computers,” Association for Computing Machinery, Special Interest Group on Operating Systems, Operating Systems Review, Proceedings of the 5th Symposium on Operating Systems Design and Implementation, Volume 36, Issue S1 (“Sapuntzakis”)

Sapuntzakis was published on December 31, 2002.¹ Sapuntzakis is prior art to the '844 Patent under at least Pre-AIA 35 U.S.C. § 102(b) as to claims 7 and 11 (the “Asserted Claims”) of the '844 Patent.

To the extent Plaintiff alleges Sapuntzakis does not disclose any particular limitation of the Asserted Claims of the '844 Patent, either expressly or inherently, any purported differences are such that the claimed subject matter as a whole would have been obvious in view of the knowledge of one skilled in the art. It would have further been obvious to a person of ordinary skill in the art as of the priority date of the '844 Patent to modify Sapuntzakis or combine the teachings of Sapuntzakis with other prior art in a manner that would have rendered the Asserted Claims invalid as obvious, including but not limited to:

- U.S. Patent Application Publication No. 2005/0283597 (“Holzmann”);
- EP Patent No. EP0604013 (“Nelson”);
- U.S. Patent No. 7,089,300 (“Birse”);
- U.S. Patent Application Publication No. 2005/0004925 (“Stahl”); and/or
- U.S. Patent Application Publication No. 2002/0034105 (“Kulkarni”).

As the United States Supreme Court held in *KSR Int'l Co. v. Teleflex, Inc.*, “[t]he combination of familiar elements according to known methods is likely to be obvious when it does no more than yield predictable results.” 550 U.S. 398, 416 (2007). The Supreme Court further held: “[w]hen a work is available in one field of endeavor, design incentives and other market forces can prompt variations of it, either in the same field or a different one. If a person of ordinary skill can implement a predictable variation, § 103 likely bars its patentability. For the same reason, if a technique has been used to improve one device, and a person of ordinary skill in the art would recognize that it would improve similar devices in the same way, using the technique is obvious unless its actual application is beyond his or her skill.” *Id.* at 417. The Supreme Court further held: “[w]hen there is a design need or market pressure to solve a problem and there are a finite number of identified, predictable solutions, a person of ordinary skill has good reason to

¹ Defendant’s invalidity charts, in some instances, rely at least in part on Plaintiffs’ apparent positions regarding the scope of its claims for purposes of asserting infringement. Nothing in these claim charts should be understood as an admission relating to infringement, either literally or under the doctrine of equivalents, or as an admission relating to Defendant’s understanding of the proper interpretation or scope of the Asserted Claims. Defendant reserves the right to rely on additional citations or sources of evidence that may also be applicable, or that may become applicable in light of any Court Order on claim construction, changes in Plaintiffs’ infringement contentions, or information obtained during discovery as the case proceeds.

Exhibit D-1

pursue the known options within his or her technical grasp. If this leads to the anticipated success, it is likely the product not of innovation but of ordinary skill and common sense.” *Id.* at 421.

Defendant reserves the right to amend or supplement this claim chart at a later date as more fully set forth in the Invalidity Contentions.

Exhibit D-1

Claims	Optimizing the Migration of Virtual Computers (“Sapuntzakis”)
[7.pre] A method for providing data to a plurality of compute nodes, comprising:	To the extent the preamble is limiting, Sapuntzakis discloses or renders obvious a method for providing data to a plurality of compute nodes. For example: “Computers and storage in the Collective system act as caches of capsules. As users travel, the Collective can move their capsules to computers close to them, giving users a consistent environment. Capsules could be moved with users as they commute between home and work. Capsules can be duplicated, distributed to many different machines, and updated like any other data; this can form the basis for administering a group of computers.” Sapuntzakis at 2. To the extent that Sapuntzakis does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Sapuntzakis, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[7.1] storing blocks of a root image of said compute nodes on a first storage unit;	Sapuntzakis discloses or renders obvious storing blocks of a root image of said compute nodes on a first storage unit. For example: “At the root of the hierarchy is a root disk, which is a complete capsule disk.” Sapuntzakis at 5 “Figure 1 shows an example of a capsule hierarchy illustrating how it may be used in a university. The root capsule contains all the software available to all students.” Sapuntzakis at 5. <pre> graph TD Root[University Capsule Disk] --> D1[Department1 Latest Disk] Root --> D2[Department2 Latest Disk] D1 --> S1[Student1 capsule snapshot Latest Disk] D1 --> S2[Student2 capsule Latest Disk] D2 --> S3[Student3 capsule Latest Disk] D2 --> D2U[Department2 updated capsule Latest Disk] S1 --> S1L[Student1 capsule Latest Disk] D2U --> S4[Student4 capsule Latest Disk] </pre> Figure 1: An example capsule hierarchy. Sapuntzakis at 5.

Exhibit D-1

Claims	Optimizing the Migration of Virtual Computers (“Sapuntzakis”)
	<pre> graph TD subgraph VMware_Process [VMware Process (one per capsule)] VM[Virtual Machine] Shim[Shim C-library] end Shim -- Local RPC --> LDSP[Local Disk Server Process (one per compute server)] subgraph LDSP_Box [Local Disk Server Process] L_COW[Local COW Disks] S_COW[Shadow COW Disks] end L_COW --- S_COW S_COW -- Network RPC --> HCP[HCP Server] subgraph HCP_Box [HCP Server] R_COW[Remote COW Disks] end </pre> Figure 3: Implementation of capsule disks and demand paging. Sapuntzakis at 6. Further, the evidence cited for claim limitation [7.pre] applies to this claim limitation. To the extent that Sapuntzakis does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Sapuntzakis, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.

Exhibit D-1

Claims	Optimizing the Migration of Virtual Computers (“Sapuntzakis”)
[7.2] storing leaf images for respective compute nodes on respective second storage units, said leaf images including only additional data blocks not previously contained in said root image and changes made by respective compute nodes to the blocks of the root image, wherein said leaf images of respective compute nodes do not include blocks of said root image that are unchanged by respective compute nodes; and	Sapuntzakis discloses or renders obvious storing leaf images for respective compute nodes on respective second storage units, said leaf images including only additional data blocks not previously contained in said root image and changes made by respective compute nodes to the blocks of the root image, wherein said leaf images of respective compute nodes do not include blocks of said root image that are unchanged by respective compute nodes. For example: “We can store the disks in these capsules efficiently by creating a hierarchy, where each child capsule could be viewed as inheriting from the parent capsule with the differences in disk state between parent and child captured in a separate copy-on-write (COW) virtual disk.” Sapuntzakis at 5. “Writes to a capsule disk are performed by writing the data to the latest COW disk and updating its bitmap file. Reads involve searching the latest COW disk and its ancestor disks in turn until the required block is found. Since the root COW disk contains a copy of all the blocks, the search is guaranteed to terminate.” Sapuntzakis at 5–6. “Since the latest COW disk is always local, all writes are local.” Sapuntzakis at 6.

Exhibit D-1

Claims	Optimizing the Migration of Virtual Computers (“Sapuntzakis”)
	Bitmap 11111111 DiskBlocks Root COW Disk 10010101 COW Disk 1 11000001 COW Disk 2 01011000 Latest COW Disk Flat Disk (as seen by the capsule) Block present in COW disk Figure 2: An example capsule disk. Sapuntzakis at 6. Further, the evidence cited for claim limitations [7.pre]-[7.1] apply to this claim limitation. To the extent this reference does not expressly or inherently disclose this limitation, it would have been obvious in view of the background knowledge of a person of ordinary skill in the art. This reference further renders the limitation obvious in light of other references, such as Birse and/or Kulkarni. Birse and Kulkarni concern similar subject matter as this reference, namely, distributed computing systems. As set forth in the accompanying contentions, storing data blocks not previously contained in a root image and/or changes made by compute nodes to the blocks of the root image was a well-known and commonly-used technique. Modifying this reference with Birse and/or Kulkarni’s teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. A person of ordinary skill in the art would have further understood how to modify the disclosures of

Exhibit D-1

Claims	Optimizing the Migration of Virtual Computers (“Sapuntzakis”)
	this reference with those of Birse or Kulkarni for example, to incorporate this limitation. <i>See, e.g.,</i> “The block device driver 154 services file system read and write requests to the system volumes 160. In one embodiment, the block device driver 154 appears to the file system 152 to be a standard hard drive device driver. However, in reality, the block device driver 154 may be configured to service the read and write requests of the file system 152 by accessing the NC server 170. As will be discussed further below, in one embodiment, one or more shadow volumes, may be created for the NC client 150 and stored with the non-persistent client-specific data 184 for the purpose of storing user preferences and/or changes to the operating system. Therefore, some mechanism is needed to direct reads and writes to the appropriate volume, e.g., the shadow volumes or the system volumes 174. According to one embodiment, the block device driver 154 has knowledge of the separate shadow volumes and system volumes 174 and redirects read and write requests to the appropriate volume. Alternatively, server-side logic may perform the redirection locally.” Birse at 7:35-53; <i>see also</i> Birse at 11:30-32. “In the embodiment depicted, it is assumed that the shadow volume contains only those portions of the operating system that have been modified by the user rather than a complete copy of the operating system as modified. The various options regarding the granularity of the portions written to the shadow volume are discussed below.” Birse 11:49-52. <i>See also</i> Birse at 10:63-65, 11:6-17, 12:54–59, 14:49-61, FIGS. 4, 6, 10. “FIG. 2 is a flow diagram outlining the acts performed by the server 102 to create the differences file 120 from the old image 114 and the new image 116. In the present example, the updated image is an updated device operating system for the PDA 104.” Kulkarni at [0036].

Exhibit D-1

Claims	Optimizing the Migration of Virtual Computers (“Sapuntzakis”)
	“Creating the differences file 120 is accomplished by applying a compression algorithm to the new image 116, using the old image 114 as a reference. Since the old image 114 is already stored on the PDA 104, data from the new image 116 that is also contained in the old image 114 does not need to be downloaded from the server 102 to the PDA 104. When the new image 116 is created at the PDA 104, data that is contained in both images is copied from the image 139 already on the PDA 104. Data that is only contained in the new image 116 is downloaded to the PDA 104 in the modified differences file 124 and is copied to the flash memory image 139.” Kulkarni at [0037]. “The differences file 120 and the modified difference file 124 are composed of two alternating types of sections: data sections contain data from the new image 116 that is not found in the old image 114; copy sections contain data that is found in both the new image 116 and the old image 114. To distinguish between the two types of sections, a header is placed at the beginning of each section, the header giving specific information about the section that follows the header.” Kulkarni at [0038]. <i>See also</i> Kulkarni at [0026]-[0041], [0068], FIG. 1, FIGS. 2a-2d, FIG. 6. To the extent that Sapuntzakis does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Sapuntzakis, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[7.3] caching blocks of said root image that have been accessed by at least one of said compute nodes in a cache memory.	Sapuntzakis discloses or renders obvious caching blocks of said root image that have been accessed by at least one of said compute nodes in a cache memory. For example: “Each COW disk can either be local or remote. Each remote COW disk has a corresponding local shadow COW disk which contains all the locally cached blocks of the remote COW disk.” Sapuntzakis at 6. “Reads, on the other hand, could either be local or remote. In the case of a remote read, the disk server requests the block from the shadow COW disk. If the block is

Exhibit D-1

Claims	Optimizing the Migration of Virtual Computers (“Sapuntzakis”)
	not cached locally, it is fetched remotely and added to the shadow COW.” Sapuntzakis at 6. “First, the memory in most systems caches disk blocks.” Sapuntzakis at 7. Further, the evidence cited for claim limitations [7.pre]-[7.2] apply to this claim limitation. To the extent this reference does not expressly or inherently disclose this limitation, it would have been obvious in view of the background knowledge of a person of ordinary skill in the art. This reference further renders the limitation obvious in light of other references, such as Holzmann, Nelson, and/or Stahl. Holzmann, Nelson, and Stahl concern similar subject matter as this reference, namely, distributed computing systems. As set forth in the accompanying contentions, caching data from comparatively slow storage devices in faster devices (including solid-state memory) was a well-known and commonly-used technique. Modifying this reference with Holzmann’s, Nelson’s, and/or Stahl’s caching teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. A person of ordinary skill in the art would have further understood how to modify the disclosures of this reference with those of Holzmann, Nelson, or Stahl for example, to incorporate this limitation. <i>See, e.g.,</i> “The invention is directed generally to computer storage systems. More specifically, the invention is directed to a system and method for booting multiple servers using a single operating system image located on a solid state disk in a storage area network.” Holzmann at [0002].

Exhibit D-1

Claims	Optimizing the Migration of Virtual Computers (“Sapuntzakis”)
	<pre> graph TD 107[INTERNAL POWER SOURCE] --- 104((MEMORY)) 107 --- 103[CONTROL MODULE] 107 --- 105[INTERFACE MODULE] 104 <--> Data 103 103 <--> Data 102[NON-VOLATILE STORAGE MEANS] 103 <--> 105 105 <--> 106([EXTERNAL DEVICES]) </pre> Fig. 1 Holzmann at FIG. 1 “Referring now to the figures, the present invention is directed to a system and method for booting multiple servers or other external devices from a storage area network using a single OS image. FIG. 1 is a block diagram illustrating the component parts of the invented system and the functions of each component part. The invented system comprises a solid state disk system 101 having a storage means 102, a control module 103, a memory module 104, an interface module 105 that communicates with external devices 106, and an internal power source 107.” Holzmann at [0017]. “The memory module 104 comprises at least one direct-access memory module for holding data currently or recently used by external devices 106. The memory module 104 is more quickly accessible, and performs read and write processing functions more quickly, than non-volatile or disk memory, such as storage means 102. The memory module 104 preferably comprises at least one random-access memory (RAM) module.” Holzmann at [0020].

Exhibit D-1

Claims	Optimizing the Migration of Virtual Computers (“Sapuntzakis”)
	“This invention relates in general to the fields of Computer Operating Systems, Multi-processor Hardware Systems, Object Oriented Programming, and Virtual Memory Systems. In particular, the invention relates to improved techniques for establishing and efficiently handling relationships between a client program, a file system server and a virtual memory manager ("VMM").” Nelson at [0001]. “This servicing of multiple client programs is now illustrated by means of Figure 6. A first client program 102 asks the CFS 218 to cache remote file 200. The CFS 218 provides a first CFS file object 108 for the first client 102 to use in processing requests on the file. The CFS 218 establishes a communications link 244 to the file server 206 located on the remote node 204 to obtain the file attributes which the CFS 218 caches in the CFS attribute cache 239. The CFS 218 also establishes the communications between the CFS file object manager 221 and the VMM 230 which caches the file data in the data cache 231 by means of the pager_object 234/fs_cache_object 232 pipeline. This pipeline results from the map and bind operations previously described. Pager object 234 is connected to fs_cache object 232 in the pager object manager 247, and also connected via the link between the fs_cache object 232 and the cache object 223 in the fs_cache object manager 243 in the CFS 218. When a second client program 118 asks the CFS 218 to cache the same remote file 200, the CFS 218 implements a second CFS file object 122. However, since the CFS 218 has already set up a CFS attribute cache 239 to handle file attributes for the remote file 200, the second CFS file object 122 is merely directed to the existing CFS attribute cache 239. Similarly, when the second CFS file object 122 wishes to read or write data to the file 200, the file object manager 221 again invokes an operation to map the file in the CFS address_space using its address_space object 225. Once again this causes the VMM's address_space object manager 227 to invoke a bind () operation on the file object 122 which is again directed to the CFS 218. This time the CFS 218 checks its cache of bind results 237 to see if that file 200 is already being cached by the VMM 230. When the CFS 218 finds that it is already

Exhibit D-1

Claims	Optimizing the Migration of Virtual Computers (“Sapuntzakis”)
	being cached, the CFS 218 returns a pointer to the fs_cache_object/pager_object pair being used, and the bind operation is completed. Thus, duplicate caches are not established and the redundant network communications overhead is avoided by the process.” Nelson at [0038]-[0039]. “In an alternate embodiment, if the first client program 102 had set up the first CFS file object 108 with the CFS 218, the first client program 102 may pass the first CFS file object 108 to the second client program 118. In that case, both client programs 102 & 118 would use the same CFS file object 108, and as before both client programs would be serviced by the CFS 218 with only one set of caches.” Nelson at [0040].

Exhibit D-1

Claims	Optimizing the Migration of Virtual Computers (“Sapuntzakis”)
	FIGURE 6 Nelson at FIG. 6. See also Nelson at [0010], [0034]-[0037], [0052]-[0119], FIG. 4. “Repository FS Implementation 222 is an implementation of a file system responsible for maintaining data on the persistent read-only file system 226 that is to have its data shared. Repository FS Implementation 222 receives and responds to file system requests from MapFS Implementation 212. Repository FS Instance in memory 224 represents a transient in-memory status of a particular Repository file system in persistent storage. Status includes, for example, information about what

Exhibit D-1

Claims	Optimizing the Migration of Virtual Computers (“Sapuntzakis”)
	files exist (or do not) and have been successfully or unsuccessfully looked up, file size, permission data, actual cached data blocks, etc.” Stahl at [0028]. “The mapping module 502 looks 308 in the View FS instance 208 (via file system communication module 506) to see whether the file is cached. Preferably, the view caches filenames in a manner similar to MapFS—that is, in a table of names and look-up results. Again, since this is the first time the file is being looked up, it will not exist in the View's cache either, unless it has been looked up through some non-MapFS process. Thus, the mapping module 502 asks the 310 View FS implementation 206 to look up the name in the View in a manner preferably similar to the way in which the VFS layer examines MapFS instance data described above. The View FS Implementation 206 looks up 312 the file in the persistent view FS instance on disk 210 and updates 314 the View FS instance in memory 208 by populating an in-memory description of the file, and inserting an entry in its name-to-inode table. The View FS implementation 206 returns 316 a reference to the file in the View FS Instance 208 to the mapping module 502.” Stahl at [0037]. <i>See also</i> Stahl at paragraphs [0017], [0044]. To the extent that Sapuntzakis does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Sapuntzakis, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[11.0] The method as recited in claim 7 further comprising: merging the blocks of said root image with the blocks of respective leaf images to create cohesive respective application environments.	Sapuntzakis discloses or renders obvious the method as recited in claim 7 further comprising merging the blocks of said root image with the blocks of respective leaf images to create cohesive respective application environments. For example: “To reduce the start-up time of a capsule, the COW disks corresponding to a capsule disk are read page-by- page on demand, rather than being pre-fetched.” Sapuntzakis at 6.

Exhibit D-1

Claims	Optimizing the Migration of Virtual Computers (“Sapuntzakis”)
	“Reads involve searching the latest COW disk and its ancestor disks in turn until the required block is found.” Sapuntzakis at 5-6. “Note that the COW disk hierarchy is not visible to the VMM, or to the OS and applications inside the capsule; all of them see a normal flat disk as illustrated in the figure.” Sapuntzakis at 6. Further, the evidence cited for claim limitations [7.pre]-[7.3] apply to this claim limitation. To the extent this reference does not expressly or inherently disclose this limitation, it would have been obvious in view of the background knowledge of a person of ordinary skill in the art. This reference further renders the limitation obvious in light of other references, such as Birse, Stahl, and/or Kulkarni. Birse, Stahl, and Kulkarni concern similar subject matter as this reference, namely, distributed computing systems. As set forth in the accompanying contentions, merging data blocks for multiple computer nodes in distributed environments was a well-known and commonly-used technique. Modifying this reference with Birse, Stahl, and/or Kulkarni’s teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. A person of ordinary skill in the art would have further understood how to modify the disclosures of this reference with those of Birse, Stahl, and/or Kulkarni for example, to incorporate this limitation. <i>See, e.g.,</i> “In one embodiment, the shadow system volume 930 is not visible to the user. Consequently, the user system volume 924 will appear to be the location where all writes are directed. For example, if a user modifies a preference, that preference will appear to be written to the preference file located in the preferences folder on the user system volume 924.” Birse at 14:61-67. “When an NC client boots from the network and accesses the operating system from an NC server, the user's preferences are dynamically merged with the system environment provided to the NC client. Advantageously, since, the user's desktop

Exhibit D-1

Claims	Optimizing the Migration of Virtual Computers (“Sapuntzakis”)
	preferences and other customized settings are all preserved from session to session and supplied to the NC client as it boots from the network, the user may login to any NC client on the network and have the same user experience.” Birse at 9:2-10. “Referring now to FIG. 5, NC client read request processing will now be described according to one embodiment of the present invention. ... At step 510, the file system 152 generates a read request directed to the system volumes 160. At step 520, the read request is received by the block device driver 154 and a determination is made whether the read request specifies a portion of the operating system that has been modified by the user or whether the read request specifies a portion of the operating system that remains unchanged. If the read request corresponds to a portion of the operating system that has not been modified by the user, the processing continues with step 530. However, if the read request corresponds to a portion of the operating system that has been previously modified by the user, then processing continues with step 540. At step 530, the block device driver directs the NC server 170 to retrieve the information associated with the read request from the system volumes 174. At step 540, the block device driver directs the NC server 170 to retrieve the information associated with the read request from the user’s shadow volume. In this manner, the file system 152 need not be aware of the particular mechanism used to coordinate user modifications with the original operating system as contained in the system volumes 174.” Birse at 11:39-41, 11:55-12:10. “The implementations described herein create and utilize a differences file to create a new image in flash memory of an electronic device from sections of an old image in the flash memory and sections of a new image, the new image being an updated version of the old image. The new image is created incrementally in flash memory so that an entire image does not have to first be constructed in device RAM, thereby saving significant overhead. A “flash manager” is described which allows the new image to be built incrementally until a memory block is created that is of sufficient size to be written to the flash memory.” Kulkarni at [0013].

Exhibit D-1

Claims	Optimizing the Migration of Virtual Computers (“Sapuntzakis”)
	“Repository FS Implementation 222 is an implementation of a file system responsible for maintaining data on the persistent read-only file system 226 that is to have its data shared. Repository FS Implementation 222 receives and responds to file system requests from MapFS Implementation 212. Repository FS Instance in memory 224 represents a transient in-memory status of a particular Repository file system in persistent storage. Status includes, for example, information about what files exist (or do not) and have been successfully or unsuccessfully looked up, file size, permission data, actual cached data blocks, etc.” Stahl at [0028]. “Overlay FS Implementation 214 is an implementation of a file system responsible for providing writeable storage on the persistent overlay file system instance 220. Overlay FS Implementation 214 receives and responds to file system requests—including write requests—from MapFS Implementation 212. Overlay FS Instance in memory 218 behaves analogously to Repository FS Instance 224. In a preferred embodiment, the Overlay FS is a directory tree rooted at a location specified when the MapFS is instantiated.” Stahl at [0029]. “The flash manager 136 uses the modified differences file 124' and the image 139 in flash memory 132 (the old image) to start building the new image. At step 404, a portion of the new image is constructed and temporarily stored in the new memory block 138. When the new memory block 138 contains a portion of the new image that constitutes at least one-half of a write block size for the flash memory 132 (“Yes” branch, step 406), then the portion of the new image stored in the new memory block 138 is added to the image 139 in the flash memory 132. As previously discussed, the size of a write block depends on the type of flash memory, i.e., the brand of the flash memory. It is common for flash memory to have a write block of 64 Kb. Also as previously discussed, the flash manager 136 may manage several new memory blocks 138. When any of the new memory blocks 138 reaches a size of one-half of a write block size, that new memory block 138 is written to the flash memory.” Kulkarni at [0069]. <i>See also</i> Kulkarni at [0026]-[0041], [0068], FIG. 1, FIGS. 2a-2d, FIG. 6.

Exhibit D-1

Claims	Optimizing the Migration of Virtual Computers (“Sapuntzakis”)
	To the extent that Sapuntzakis does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Sapuntzakis, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.

Exhibit D-2

Invalidity of U.S. Patent No. 8,332,844 in View of United States Patent No. 7,089,300 (“Birse”)

Birse titled “Method and apparatus for administering the operating system of a net-booted environment” to Cameron Stuart Birse et al. was filed on October 18, 1999 and issued on August 8, 2006.¹ Birse is prior art to the ’844 Patent under at least Pre-AIA 35 U.S.C. §§ 102(e) as to claims 7 and 11 (the “Asserted Claims”) of the ’844 Patent.

To the extent Plaintiff alleges Birse does not disclose any particular limitation of the Asserted Claims of the ’844 Patent, either expressly or inherently, any purported differences are such that the claimed subject matter as a whole would have been obvious in view of the knowledge of one skilled in the art. It would have further been obvious to a person of ordinary skill in the art as of the priority date of the ’844 Patent to modify Birse or combine the teachings of Birse with other prior art in a manner that would have rendered the Asserted Claims invalid as obvious, including but not limited to:

- U.S. Patent Application Publication No. 2005/0283597 (“Holzmann”);
- EP Patent No. EP0604013 (“Nelson”);
- Optimizing the Migration of Virtual Computers” (“Sapuntzakis”);
- U.S. Patent Application Publication No. 2005/0004925 (“Stahl”); and/or
- U.S. Patent Application Publication No. 2002/0034105 (“Kulkarni”).

As the United States Supreme Court held in *KSR Int’l Co. v. Teleflex, Inc.*, “[t]he combination of familiar elements according to known methods is likely to be obvious when it does no more than yield predictable results.” 550 U.S. 398, 416 (2007). The Supreme Court further held: “[w]hen a work is available in one field of endeavor, design incentives and other market forces can prompt variations of it, either in the same field or a different one. If a person of ordinary skill can implement a predictable variation, § 103 likely bars its patentability. For the same reason, if a technique has been used to improve one device, and a person of ordinary skill in the art would recognize that it would improve similar devices in the same way, using the technique is obvious unless its actual application is beyond his or her skill.” *Id.* at 417. The Supreme Court further held: “[w]hen there is a design need or market pressure to solve a problem and there are a finite number of identified, predictable solutions, a person of ordinary skill has good reason to pursue the known options within his or her technical grasp. If this leads to the anticipated success, it is likely the product not of innovation but of ordinary skill and common sense.” *Id.* at 421.

¹ Defendant’s invalidity charts, in some instances, rely at least in part on Plaintiff’s apparent positions regarding the scope of its claims for purposes of asserting infringement. Nothing in these claim charts should be understood as an admission relating to infringement, either literally or under the doctrine of equivalents, or as an admission relating to Defendant’s understanding of the proper interpretation or scope of the Asserted Claims. Defendant reserves the right to rely on additional citations or sources of evidence that may also be applicable, or that may become applicable in light of any Court Order on claim construction, changes in Plaintiff’s infringement contentions, or information obtained during discovery as the case proceeds.

Exhibit D-2

Defendant reserves the right to amend or supplement this claim chart at a later date as more fully set forth in the Invalidity Contentions.

Exhibit D-2

Claims	United States Patent No. 7,089,300 (“Birse”)
[7.0] A method for providing data to a plurality of compute nodes, comprising:	To the extent the preamble is limiting, Birse discloses or renders obvious a method for providing data to a plurality of compute nodes. For example: “A method and apparatus are provided for supplying a reliable and maintainable operating system in a net-booted environment. According to one embodiment, a network computer (NC) system including an NC server and multiple NC clients is managed by an NC client causing the remainder of the NC clients that are subsequently booted to receive operating system software that is configured differently than that currently in effect by replacing one or more system volumes on the NC server containing the operating system software with one or more different system volumes.” Birse at Abstract. “In order to obtain the benefits of low-cost administration, all storage of application programs and operating system software for an NC system is preferably on the NC server. This doesn't mean the NC clients have no mass storage, but rather that the NC clients boot over the network by accessing a stored copy of the operating system from the NC server. Additionally, when NC clients want to run an application, they access it from the NC server. Advantageously, in this manner, an application may be upgraded by simply replacing the old version of the application with a new version on the NC server. Then, the next time an NC client requests the application, the NC client will receive the new version.” Birse at 4:53-64. <i>See also</i> Birse at 2:51-3:14.

Exhibit D-2

Claims	United States Patent No. 7,089,300 ("Birse")
	FIG. 1 Birse at FIG. 1 "FIG. 8 is a simplified block diagram illustrating two NC systems 810 and 820 coupled via the Internet 805." Birse at 14:8-9. FIG. 8 Birse at FIG. 8.

Exhibit D-2

Claims	United States Patent No. 7,089,300 (“Birse”)
	“FIG. 2 is a block diagram of a digital processing system which may be used in accordance with one embodiment of the present invention. For example, the digital processing system 200 shown in FIG. 2 may be used as an NC client, an NC server, or other server computer system, such as an NT server. The digital processing system 200 may be interfaced to external systems through a network interface 268. It will be appreciated that the network interface 268 may be considered as part of the digital processing system 200. The network interface 268 may be an analog modem, an ISDN modem, a cable modem, a token ring interface, a satellite transmission interface, a wireless interface, or other interface(s) for providing a data communication link between two or more digital processing systems. The digital processing system 200 includes a processor 252, which may represent one or more processors and may include one or more conventional types of such processors, such as Motorola PowerPC processor, an Intel Pentium (or x86) processor, etc. A memory 254 is coupled to the processor 252 by a bus 256. The memory 254 may be a dynamic random access memory (DRAM) and/or may include static RAM (SRAM). The processor may also be coupled to other types of storage areas/memories (e.g., cache, Flash memory, disk, etc.), which could be considered as part of the memory 254 or separate from the memory 254. The bus 256 further couples the processor 252 to a display controller 258, an optional mass memory 262, the network interface 268, and an input/output (I/O) controller 264. The mass memory 262 may represent a magnetic, optical, magneto-optical, tape, and/or other type of machine-readable medium/device for storing information. For example, the mass memory 262 may represent a hard disk, a read-only or writeable optical CD, etc. The display controller 258 controls in a conventional manner a display 260, which may represent a cathode ray tube (CRT) display, a liquid crystal display (LCD), a plasma display, or other type of display device. The I/O controller 264 controls I/O device(s) 266, which may include one or more keyboards, mouse/trackball or other pointing devices, magnetic and/or

Exhibit D-2

Claims	United States Patent No. 7,089,300 (“Birse”)
	optical disk drives, printers, scanners, digital cameras, microphones, etc.” Birse at 7:63-8:35. FIG. 2 Birse at FIG. 2. To the extent that Birse does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Birse, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[7.1] storing blocks of a root image of said compute nodes on a first storage unit;	Birse discloses or renders obvious storing blocks of a root image of said compute nodes on a first storage unit. For example: “The NC server 170 includes one or more system volumes 174, one or more application volumes 176, a user registry 178, non-persistent client-specific data 184, and persistent user data 186. The system volumes 174 include a protected, read-only, master copy of the operating system software.” Birse at 5:10-14. “FIG. 9 is a block diagram that conceptually illustrates NC client interaction with a SplitOS 920 according to one embodiment of the present invention. According to

Exhibit D-2

Claims	United States Patent No. 7,089,300 (“Birse”)
	the example depicted, the SplitOS 920 of the NC server 170 contains a read-only core system volume image 922 and a read-write user system volume image 924. The core system volume 922 preferably contains those parts of the system that do not need to be written back to during system operation. One goal of the core system volume 922 is to provide all the system components that are mandatory for system operation.” Birse at 14:27-34. “The user system volume 924 contains all the user-configurable system components, including preferences and all the system additions installed by application software, such as application-installed extensions and libraries. Additionally, the user system volume 924 may contain applications that cannot be run from an AppleShare server and/or other system components that do not function on a read-only volume.” Birse at 14:41-44. “While for purposes of illustration, the boot server process 172, the server software management process 180, the user environment process 182, and the administration tool 188 are shown as part of a single NC server 170, in alternative embodiments, they may be distributed in whole or in part across multiple server computer systems. Additionally, the system volumes 174, the application volumes 176, the user registry 178, the non-persistent client-specific data 184, and the persistent user data 186 may be independently distributed across multiple server computer systems.” Birse 6:59-63. “FIG. 6 is a simplified block diagram illustrating the layout of an exemplary hierarchical directory structure that may be used by an NC server according to one embodiment of the present invention. According to the embodiment depicted, a hard drive 610 of the NC server 170 or other mass storage device associated with the NC server 170 includes an NC folder (directory) 620. Within the NC folder 620 are an NC admin folder 630 and an NC shared data folder 640.” Birse 12:30-38.

Exhibit D-2

Claims	United States Patent No. 7,089,300 (“Birse”)
	<pre> graph TD HD[(HARD DRIVE 610)] --- NC[NC 620] NC --- NCADMIN[NC ADMIN 630] NC --- NCSD[NC SHARED DATA 640] NCADMIN --- MOSI[MASTER OS IMAGE 650] NCSD --- SOSI[SHARED OS IMAGE 660] NCSD --- BI[BOOT IMAGE 670] NCSD --- CL[CLIENTS 680] CL --- NC1[NC #1 690] NC1 --- SI[SHADOW IMAGE 695] </pre> FIG. 6 Birse at FIG. 6. “The shared operating system image 660 is a read- write version of the read-only master operating system image 650 and may be mounted by the NC client 150 as part of step 395.” Birse 12:44–47.

Exhibit D-2


Claims	United States Patent No. 7,089,300 (“Birse”)
	 FIG. 10 Birse at FIG. 10. Further, the evidence cited for claim limitation [7.pre] applies to this claim limitation. To the extent that Birse does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Birse, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[7.2] storing leaf images for respective compute nodes on respective second storage units, said leaf images including only additional data blocks not previously contained in said root image and changes made by respective compute nodes to the blocks of the root image, wherein said leaf images of respective compute nodes do not include blocks of said root image that are unchanged by respective compute nodes; and	Birse discloses or renders obvious storing leaf images for respective compute nodes on respective second storage units, said leaf images including only additional data blocks not previously contained in said root image and changes made by respective compute nodes to the blocks of the root image, wherein said leaf images of respective compute nodes do not include blocks of said root image that are unchanged by respective compute nodes. For example: “In this example, the NC clients are numbered from 1 to N and the clients folder 680 includes a folder for each NC client. For example, an NC #1 folder 690 is an area for storing client-specific information, such as a shadow image 695, corresponding to NC client #1.” Birse 12:54–59.

Exhibit D-2

Claims	United States Patent No. 7,089,300 (“Birse”)
	“The block device driver 154 services file system read and write requests to the system volumes 160. In one embodiment, the block device driver 154 appears to the file system 152 to be a standard hard drive device driver. However, in reality, the block device driver 154 may be configured to service the read and write requests of the file system 152 by accessing the NC server 170. As will be discussed further below, in one embodiment, one or more shadow volumes, may be created for the NC client 150 and stored with the non-persistent client-specific data 184 for the purpose of storing user preferences and/or changes to the operating system. Therefore, some mechanism is needed to direct reads and writes to the appropriate volume, e.g., the shadow volumes or the system volumes 174. According to one embodiment, the block device driver 154 has knowledge of the separate shadow volumes and system volumes 174 and redirects read and write requests to the appropriate volume. Alternatively, server-side logic may perform the redirection locally.” Birse at 7:35-53; <i>see also</i> Birse at 11:30-32. “FIG. 4 is a flow diagram illustrating how NC client write requests directed to the system volumes 160 are handled according to one embodiment of the present invention.” Birse at 10:63-65. “According to the embodiment depicted, at step 410, the file system 152 generates a write request directed to the system volumes 160. The write request is received by the block device driver 154 and redirected by the block device driver 154 at step 420 by translating it into a write request directed to the user's shadow volume on the NC server 170. At step 430, the NC server 170 stores information associated with the write request in the shadow volume associated with the NC client 150. In this manner, the file system 152 need not be aware that attempted modifications to the system volumes 160 are recorded instead in a client-specific shadow volume residing on the NC server 170.” Birse at 11:6-17.

Exhibit D-2

Claims	United States Patent No. 7,089,300 (“Birse”)
	<pre> graph TD A([OS WRITE PROCESSING]) --> B[FILE SYSTEM GENERATES A WRITE REQUEST 410] B --> C[BLOCK DEVICE DRIVER REDIRECTS THE WRITE REQUEST TO THE USER'S SHADOW VOLUME ON THE NC SERVER 420] C --> D[NC SERVER STORES INFORMATION ASSOCIATED WITH THE WRITE REQUEST IN THE USER'S SHADOW VOLUME 430] D --> E([END]) </pre> FIG. 4 Birse at FIG. 4. “In the embodiment depicted, it is assumed that the shadow volume contains only those portions of the operating system that have been modified by the user rather than a complete copy of the operating system as modified. The various options regarding the granularity of the portions written to the shadow volume are discussed below.” Birse at 11:49-52. “Preferably, the NC server 170 also creates a shadow system volume 930 for each connected NC client. The shadow system volume 930 shadows the user system

Exhibit D-2

Claims	United States Patent No. 7,089,300 (“Birse”)
	volume 924 by storing modifications that are made to the user system volume 924.” Birse 14:49-53. “As described above, according to one embodiment, the shadow system volume 930 is used by the block device driver 154 of the NC client 150 to implement a “copy-on-write” storage scheme in which modifications (writes) directed to the user system volume 930 are instead copied to the shadow system volume 930.” Birse at 14:56-61. Further, the evidence cited for claim limitations [7.pre]-[7.1] apply to this claim limitation. To the extent this reference does not expressly or inherently disclose this limitation, it would have been obvious in view of the background knowledge of a person of ordinary skill in the art. This reference further renders the limitation obvious in light of other references, such as Sapuntzakis and/or Kulkarni. Sapuntzakis and Kulkarni concern similar subject matter as this reference, namely, distributed computing systems. As set forth in the accompanying contentions, storing data blocks not previously contained in a root image and/or changes made by compute nodes to the blocks of the root image was a well-known and commonly-used technique. Modifying this reference with Sapuntzakis and/or Kulkarni’s teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. A person of ordinary skill in the art would have further understood how to modify the disclosures of this reference with those of Sapuntzakis or Kulkarni for example, to incorporate this limitation. <i>See, e.g.,</i> “We can store the disks in these capsules efficiently by creating a hierarchy, where each child capsule could be viewed as inheriting from the parent capsule with the differences in disk state between parent and child captured in a separate copy-on-write (COW) virtual disk.” Sapuntzakis at 5. “Writes to a capsule disk are performed by writing the data to the latest COW disk

Exhibit D-2

Claims	United States Patent No. 7,089,300 (“Birse”)
	and updating its bitmap file. Reads involve searching the latest COW disk and its ancestor disks in turn until the required block is found. Since the root COW disk contains a copy of all the blocks, the search is guaranteed to terminate.” Sapuntzakis at 5–6. “Since the latest COW disk is always local, all writes are local.” Sapuntzakis at 6. The diagram illustrates a capsule disk structure. At the top, a 'Bitmap' is shown with four rows of 8 bits each: 11111111, 10010101, 11000001, and 01011000. To the right of the bitmap is a 'DiskBlocks' section with four rows of 8 blocks each. The top row is labeled 'Root COW Disk' and all its blocks are shaded gray. The second row is labeled 'COW Disk 1' and has blocks 1, 3, 5, and 7 shaded gray. The third row is labeled 'COW Disk 2' and has blocks 1, 2, and 3 shaded gray. The fourth row is labeled 'Latest COW Disk' and has blocks 1, 2, 3, 4, 5, and 6 shaded gray. Below these is a 'Flat Disk (as seen by the capsule)' with 8 blocks, all of which are white. Dashed vertical lines connect the shaded blocks in the COW disks to their corresponding positions in the Flat Disk. A legend at the bottom indicates that a gray square represents a 'Block present in COW disk'. Figure 2: An example capsule disk. Sapuntzakis at 6. <i>See also</i> Sapuntzakis at 1-12. “FIG. 2 is a flow diagram outlining the acts performed by the server 102 to create the differences file 120 from the old image 114 and the new image 116. In the present example, the updated image is an updated device operating system for the PDA 104.” Kulkarni at [0036].

Exhibit D-2

Claims	United States Patent No. 7,089,300 (“Birse”)
	“Creating the differences file 120 is accomplished by applying a compression algorithm to the new image 116, using the old image 114 as a reference. Since the old image 114 is already stored on the PDA 104, data from the new image 116 that is also contained in the old image 114 does not need to be downloaded from the server 102 to the PDA 104. When the new image 116 is created at the PDA 104, data that is contained in both images is copied from the image 139 already on the PDA 104. Data that is only contained in the new image 116 is downloaded to the PDA 104 in the modified differences file 124 and is copied to the flash memory image 139.” Kulkarni at [0037]. “The differences file 120 and the modified difference file 124 are composed of two alternating types of sections: data sections contain data from the new image 116 that is not found in the old image 114; copy sections contain data that is found in both the new image 116 and the old image 114. To distinguish between the two types of sections, a header is placed at the beginning of each section, the header giving specific information about the section that follows the header.” Kulkarni at [0038]. <i>See also</i> Kulkarni at [0026]-[0041], [0068], FIG. 1, FIGS. 2a-2d, FIG. 6. To the extent that Birse does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Birse, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[7.3] caching blocks of said root image that have been accessed by at least one of said compute nodes in a cache memory.	Birse discloses or renders obvious caching blocks of said root image that have been accessed by at least one of said compute nodes in a cache memory. For example: “The processor may also be coupled to other types of storage areas/memories (e.g., cache, Flash memory, disk, etc.), which could be considered as part of the memory 254 or separate from the memory 254.” Birse at 8:17-19. “It will be appreciated that the digital processing system 200 represents only one example of a system, which may have many different configurations and

Exhibit D-2

Claims	United States Patent No. 7,089,300 (“Birse”)
	architectures, and which may be employed with the present invention. For example, Macintosh and Intel systems often have multiple busses, such as a peripheral bus, a dedicated cache bus, etc.” Birse at 8:36-41. Further, the evidence cited for claim limitations [7.pre]-[7.2] apply to this claim limitation. To the extent this reference does not expressly or inherently disclose this limitation, it would have been obvious in view of the background knowledge of a person of ordinary skill in the art. This reference further renders the limitation obvious in light of other references, such as Holzmann, Nelson, and/or Stahl. Holzmann, Nelson, and Stahl concern similar subject matter as this reference, namely, distributed computing systems. As set forth in the accompanying contentions, caching data from comparatively slow storage devices in faster devices (including solid-state memory) was a well-known and commonly-used technique. Modifying this reference with Holzmann’s, Nelson’s, and/or Stahl’s caching teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. A person of ordinary skill in the art would have further understood how to modify the disclosures of this reference with those of Holzmann, Nelson, or Stahl for example, to incorporate this limitation. <i>See, e.g.,</i> “The invention is directed generally to computer storage systems. More specifically, the invention is directed to a system and method for booting multiple servers using a single operating system image located on a solid state disk in a storage area network.” Holzmann at [0002].

Exhibit D-2

Claims	United States Patent No. 7,089,300 (“Birser”)
	<pre> graph TD 107[INTERNAL POWER SOURCE] --- 104[MEMORY] 107 --- 103[CONTROL MODULE] 107 --- 105[INTERFACE MODULE] 104 <--> Data 103 103 <--> Data 102[NON-VOLATILE STORAGE MEANS] 103 <--> 105 105 <--> 106[EXTERNAL DEVICES] subgraph 101 [Solid State Disk System] 107 104 103 102 105 end </pre> Fig. 1 Holzmann at FIG. 1 “Referring now to the figures, the present invention is directed to a system and method for booting multiple servers or other external devices from a storage area network using a single OS image. FIG. 1 is a block diagram illustrating the component parts of the invented system and the functions of each component part. The invented system comprises a solid state disk system 101 having a storage means 102, a control module 103, a memory module 104, an interface module 105 that communicates with external devices 106, and an internal power source 107.” Holzmann at [0017]. “The memory module 104 comprises at least one direct-access memory module for holding data currently or recently used by external devices 106. The memory module 104 is more quickly accessible, and performs read and write processing functions more quickly, than non-volatile or disk memory, such as storage means 102. The memory module 104 preferably comprises at least one random-access memory (RAM) module.” Holzmann at [0020].

Exhibit D-2

Claims	United States Patent No. 7,089,300 ("Birse")
	"This invention relates in general to the fields of Computer Operating Systems, Multi-processor Hardware Systems, Object Oriented Programming, and Virtual Memory Systems. In particular, the invention relates to improved techniques for establishing and efficiently handling relationships between a client program, a file system server and a virtual memory manager ("VMM")." Nelson at [0001]. "This servicing of multiple client programs is now illustrated by means of Figure 6. A first client program 102 asks the CFS 218 to cache remote file 200. The CFS 218 provides a first CFS file object 108 for the first client 102 to use in processing requests on the file. The CFS 218 establishes a communications link 244 to the file server 206 located on the remote node 204 to obtain the file attributes which the CFS 218 caches in the CFS attribute cache 239. The CFS 218 also establishes the communications between the CFS file object manager 221 and the VMM 230 which caches the file data in the data cache 231 by means of the pager_object 234/fs_cache_object 232 pipeline. This pipeline results from the map and bind operations previously described. Pager object 234 is connected to fs_cache object 232 in the pager object manager 247, and also connected via the link between the fs_cache object 232 and the cache object 223 in the fs_cache object manager 243 in the CFS 218. When a second client program 118 asks the CFS 218 to cache the same remote file 200, the CFS 218 implements a second CFS file object 122. However, since the CFS 218 has already set up a CFS attribute cache 239 to handle file attributes for the remote file 200, the second CFS file object 122 is merely directed to the existing CFS attribute cache 239. Similarly, when the second CFS file object 122 wishes to read or write data to the file 200, the file object manager 221 again invokes an operation to map the file in the CFS address_space using its address_space object 225. Once again this causes the VMM's address_space object manager 227 to invoke a bind () operation on the file object 122 which is again directed to the CFS 218. This time the CFS 218 checks its cache of bind results 237 to see if that file 200 is already being cached by the VMM 230. When the CFS 218 finds that it is already

Exhibit D-2

Claims	United States Patent No. 7,089,300 (“Birse”)
	being cached, the CFS 218 returns a pointer to the fs_cache_object/pager_object pair being used, and the bind operation is completed. Thus, duplicate caches are not established and the redundant network communications overhead is avoided by the process.” Nelson at [0038]-[0039]. “In an alternate embodiment, if the first client program 102 had set up the first CFS file object 108 with the CFS 218, the first client program 102 may pass the first CFS file object 108 to the second client program 118. In that case, both client programs 102 & 118 would use the same CFS file object 108, and as before both client programs would be serviced by the CFS 218 with only one set of caches.” Nelson at [0040].

Exhibit D-2

Claims	United States Patent No. 7,089,300 ("Birse")
	FIGURE 6 Nelson at FIG. 6. See also Nelson at [0010], [0034]-[0037], [0052]-[0119], FIG. 4. “Repository FS Implementation 222 is an implementation of a file system responsible for maintaining data on the persistent read-only file system 226 that is to have its data shared. Repository FS Implementation 222 receives and responds to file system requests from MapFS Implementation 212. Repository FS Instance in memory 224 represents a transient in-memory status of a particular Repository file system in persistent storage. Status includes, for example, information about what

Exhibit D-2

Claims	United States Patent No. 7,089,300 (“Birse”)
	files exist (or do not) and have been successfully or unsuccessfully looked up, file size, permission data, actual cached data blocks, etc.” Stahl at [0028]. “The mapping module 502 looks 308 in the View FS instance 208 (via file system communication module 506) to see whether the file is cached. Preferably, the view caches filenames in a manner similar to MapFS—that is, in a table of names and look-up results. Again, since this is the first time the file is being looked up, it will not exist in the View's cache either, unless it has been looked up through some non-MapFS process. Thus, the mapping module 502 asks the 310 View FS implementation 206 to look up the name in the View in a manner preferably similar to the way in which the VFS layer examines MapFS instance data described above. The View FS Implementation 206 looks up 312 the file in the persistent view FS instance on disk 210 and updates 314 the View FS instance in memory 208 by populating an in-memory description of the file, and inserting an entry in its name-to-inode table. The View FS implementation 206 returns 316 a reference to the file in the View FS Instance 208 to the mapping module 502.” Stahl at [0037]. <i>See also</i> Stahl at paragraphs [0017], [0044]. To the extent that Birse does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Birse, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[11.0] The method as recited in claim 7 further comprising: merging the blocks of said root image with the blocks of respective leaf images to create cohesive respective application environments.	Birse discloses or renders obvious the method as recited in claim 7 further comprising merging the blocks of said root image with the blocks of respective leaf images to create cohesive respective application environments. For example: “In one embodiment, the shadow system volume 930 is not visible to the user. Consequently, the user system volume 924 will appear to be the location where all writes are directed. For example, if a user modifies a preference, that preference will

Exhibit D-2

Claims	United States Patent No. 7,089,300 (“Birse”)
	appear to be written to the preference file located in the preferences folder on the user system volume 924.” Birse at 14:61-67. “When an NC client boots from the network and accesses the operating system from an NC server, the user's preferences are dynamically merged with the system environment provided to the NC client. Advantageously, since, the user's desktop preferences and other customized settings are all preserved from session to session and supplied to the NC client as it boots from the network, the user may login to any NC client on the network and have the same user experience.” Birse at 9:2-10. “Referring now to FIG. 5, NC client read request processing will now be described according to one embodiment of the present invention. ... At step 510, the file system 152 generates a read request directed to the system volumes 160. At step 520, the read request is received by the block device driver 154 and a determination is made whether the read request specifies a portion of the operating system that has been modified by the user or whether the read request specifies a portion of the operating system that remains unchanged. If the read request corresponds to a portion of the operating system that has not been modified by the user, the processing continues with step 530. However, if the read request corresponds to a portion of the operating system that has been previously modified by the user, then processing continues with step 540. At step 530, the block device driver directs the NC server 170 to retrieve the information associated with the read request from the system volumes 174. At step 540, the block device driver directs the NC server 170 to retrieve the information associated with the read request from the user’s shadow volume. In this manner, the file system 152 need not be aware of the particular mechanism used to coordinate user modifications with the original operating system as contained in the system volumes 174.” Birse at 11:39-41, 11:55-12:10. Further, the evidence cited for claim limitations [7.pre]-[7.3] apply to this claim limitation.

Exhibit D-2

Claims	United States Patent No. 7,089,300 (“Birser”)
	To the extent this reference does not expressly or inherently disclose this limitation, it would have been obvious in view of the background knowledge of a person of ordinary skill in the art. This reference further renders the limitation obvious in light of other references, such as Sapuntzakis, Stahl, and/or Kulkarni. Sapuntzakis, Stahl, and Kulkarni concern similar subject matter as this reference, namely, distributed computing systems. As set forth in the accompanying contentions, merging data blocks for multiple computer nodes in distributed environments was a well-known and commonly-used technique. Modifying this reference with Sapuntzakis, Stahl, and/or Kulkarni’s teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. A person of ordinary skill in the art would have further understood how to modify the disclosures of this reference with those of Sapuntzakis, Stahl, and/or Kulkarni for example, to incorporate this limitation. <i>See, e.g.,</i> “To reduce the start-up time of a capsule, the COW disks corresponding to a capsule disk are read page-by- page on demand, rather than being pre-fetched.” Sapuntzakis at 6. “Reads involve searching the latest COW disk and its ancestor disks in turn until the required block is found.” Sapuntzakis at 5-6. “Note that the COW disk hierarchy is not visible to the VMM, or to the OS and applications inside the capsule; all of them see a normal flat disk as illustrated in the figure.” Sapuntzakis at 6. <i>See also</i> Sapuntzakis at 1-12. “The implementations described herein create and utilize a differences file to create a new image in flash memory of an electronic device from sections of an old image in the flash memory and sections of a new image, the new image being an updated version of the old image. The new image is created incrementally in flash memory so that an entire image does not have to first be constructed in device RAM, thereby

Exhibit D-2

Claims	United States Patent No. 7,089,300 (“Birser”)
	saving significant overhead. A “flash manager” is described which allows the new image to be built incrementally until a memory block is created that is of sufficient size to be written to the flash memory.” Kulkarni at [0013]. “Repository FS Implementation 222 is an implementation of a file system responsible for maintaining data on the persistent read-only file system 226 that is to have its data shared. Repository FS Implementation 222 receives and responds to file system requests from MapFS Implementation 212. Repository FS Instance in memory 224 represents a transient in-memory status of a particular Repository file system in persistent storage. Status includes, for example, information about what files exist (or do not) and have been successfully or unsuccessfully looked up, file size, permission data, actual cached data blocks, etc.” Stahl at [0028]. “Overlay FS Implementation 214 is an implementation of a file system responsible for providing writeable storage on the persistent overlay file system instance 220. Overlay FS Implementation 214 receives and responds to file system requests—including write requests—from MapFS Implementation 212. Overlay FS Instance in memory 218 behaves analogously to Repository FS Instance 224. In a preferred embodiment, the Overlay FS is a directory tree rooted at a location specified when the MapFS is instantiated.” Stahl at [0029]. “The flash manager 136 uses the modified differences file 124’ and the image 139 in flash memory 132 (the old image) to start building the new image. At step 404, a portion of the new image is constructed and temporarily stored in the new memory block 138. When the new memory block 138 contains a portion of the new image that constitutes at least one-half of a write block size for the flash memory 132 (“Yes” branch, step 406), then the portion of the new image stored in the new memory block 138 is added to the image 139 in the flash memory 132. As previously discussed, the size of a write block depends on the type of flash memory, i.e., the brand of the flash memory. It is common for flash memory to have a write block of 64 Kb. Also as previously discussed, the flash manager 136 may manage several new memory blocks 138. When any of the new memory blocks 138 reaches a size

Exhibit D-2

Claims	United States Patent No. 7,089,300 (“Birse”)
	of one-half of a write block size, that new memory block 138 is written to the flash memory.” Kulkarni at [0069]. <i>See also</i> Kulkarni at [0026]-[0041], [0068], FIG. 1, FIGS. 2a-2d, FIG. 6. To the extent that Birse does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Birse, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.

Exhibit D-21

Invalidity of U.S. Patent No. 8,332,844 in View of United States Patent No. 7,398,382 (“Rothman”)

Rothman titled “Method and apparatus to enhance platform boot efficiency” to Michael A. Rothman was filed on December 29, 2004 and issued on July 8, 2008.¹ Rothman is prior art to the ’844 Patent under at least Pre-AIA 35 U.S.C. § 102(e) as to claims 7 and 11 (the “Asserted Claims”) of the ’844 Patent, as discussed in detail in the chart that follows.

To the extent Plaintiff alleges Rothman does not disclose any particular limitation of the Asserted Claims of the ’844 Patent, either expressly or inherently, any purported differences are such that the claimed subject matter as a whole would have been obvious in view of the knowledge of one skilled in the art. It would have further been obvious to a person of ordinary skill in the art as of the priority date of the ’844 Patent to modify Rothman or combine the teachings of Rothman with other prior art in a manner that would have rendered the Asserted Claims invalid as obvious, including but not limited to:

- U.S. Patent Application Publication No. 2005/0283597 (“Holzmann”);
- EP Patent No. EP0604013 (“Nelson”);
- Optimizing the Migration of Virtual Computers” (“Sapuntzakis”);
- U.S. Patent No. 7,089,300 (“Birse”);
- U.S. Patent Application Publication No. 2005/0004925 (“Stahl”); and/or
- U.S. Patent Application Publication No. 2002/0034105 (“Kulkarni”).

As the United States Supreme Court held in *KSR Int’l Co. v. Teleflex, Inc.*, “[t]he combination of familiar elements according to known methods is likely to be obvious when it does no more than yield predictable results.” 550 U.S. 398, 416 (2007). The Supreme Court further held: “[w]hen a work is available in one field of endeavor, design incentives and other market forces can prompt variations of it, either in the same field or a different one. If a person of ordinary skill can implement a predictable variation, § 103 likely bars its patentability. For the same reason, if a technique has been used to improve one device, and a person of ordinary skill in the art would recognize that it would improve similar devices in the same way, using the technique is obvious unless its actual application is beyond his or her skill.” *Id.* at 417. The Supreme Court further held: “[w]hen there is a design need or market pressure to solve a problem and there are a finite number of identified, predictable solutions, a person of ordinary skill has good reason to

¹ Defendant’s invalidity charts, in some instances, rely at least in part on Plaintiff’s apparent positions regarding the scope of its claims for purposes of asserting infringement. Nothing in these claim charts should be understood as an admission relating to infringement, either literally or under the doctrine of equivalents, or as an admission relating to Defendant’s understanding of the proper interpretation or scope of the Asserted Claims. Defendant reserves the right to rely on additional citations or sources of evidence that may also be applicable, or that may become applicable in light of any Court Order on claim construction, changes in Plaintiff’s infringement contentions, or information obtained during discovery as the case proceeds.

Exhibit D-21

pursue the known options within his or her technical grasp. If this leads to the anticipated success, it is likely the product not of innovation but of ordinary skill and common sense.” *Id.* at 421.

Defendant reserves the right to amend or supplement this claim chart at a later date as more fully set forth in the Invalidity Contentions.

Exhibit D-21

Claims	United States Patent No. 7,398,382 (“Rothman”)
[7.pre] A method for providing data to a plurality of compute nodes, comprising:	To the extent the preamble is limiting, Rothman discloses or renders obvious a method for providing data to a plurality of compute nodes. For example: “Methods and apparatus for improving network boot efficiency are disclosed. Under embodiments of the method, boot images that are initially sent from a boot server to various clients are cached at network devices along communication paths between the boot server and the clients. In response to subsequent boot image requests from the clients, boot images cached at the network devices are downloaded directly from the network devices rather than from the boot server, reducing network traffic to and from the boot server and domains in between. In addition, boot program bootstrap files may also be cached and downloaded in a similar manner. Techniques are also disclosed for intercepting boot image download and network bootstrap program requests at the network devices, and for maintaining valid boot image cache configurations across the network. The network devices generally include switches, routers, bridges, and gateway servers.” Rothman at Abstract. “Each of domains 102, 104, 106 and 110 are depicted as hosting multiple clients, including exemplary desktop clients 128 and laptop clients 130. Other types of clients may also be hosted by a given domain, including workstations and various types of servers (both not shown) in addition to the various domain and DHCP servers shown in the figures herein. Domain 108 may also host similar types of clients (not shown for clarity) to those shown in domains 102,104,106 and 110. It will be understood that each domain may include 10’s, 100’s or even 1000’s of clients, only a fraction of which are depicted in the figures herein for simplicity and clarity. Some of the clients may be connected to their respective domain or DHCP server via a wired link (e.g., Ethernet), while others may be connected via a wireless link (e.g., 802.11a-g) facilitated by an appropriate wireless access point or router. For simplicity, wireless access points and routers are not depicted in the figures herein. At least a portion of the clients in each domain are configured to boot their respective operating systems using a network boot process.” Rothman at 3:44-62. “In a block 204, the client obtains an operating system boot image over the network

Exhibit D-21

Claims	United States Patent No. 7,398,382 (“Rothman”)
	from a boot server. In conjunction with delivering the boot image to the client, the boot image is cached at a network device. Such as a switch, router, bridge, etc., or a gateway server that is proximate to the client in a block 206. The image is then booted in a block 208 to enable run-time use, which is continued in a continuation block 210.” Rothman at 4:6-12. To the extent that Rothman does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Rothman, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[7.1] storing blocks of a root image of said compute nodes on a first storage unit;	Rothman discloses or renders obvious storing blocks of a root image of said compute nodes on a first storage unit. For example: “In one embodiment, NBP file 322 contains information that identifies the IP address for switch 114 as the storage location for boot image file(s) 322. Accordingly, rather than send a boot image download request message to PXE server 126, the message is sent directly to switch 114 at step 9 using, e.g., a management port, as shown in FIG. 3<i>b</i>.” Rothman at 9:8-14. “In response to receiving the boot image download request message, one or more corresponding boot image files, as depicted by boot image 6 in FIG. 1<i>b</i>, is/are retrieved from a local storage device hosted by switch 114 and downloaded to client 132 at step 10.” Rothman at 9:22-26. “For example, a machine-readable medium can include Such as a read only memory (ROM); a random access memory (RAM); a magnetic disk storage media; an optical storage media; and a flash memory device, etc. In addition, a machine-readable medium can include propagated signals such as electrical, optical, acoustical or other form of propagated signal (e.g., carrier waves, infrared signals, digital signals, etc.).” Rothman at 15:35-43. <i>See also</i> Rothman at cl. 1.

Exhibit D-21

Claims	United States Patent No. 7,398,382 (“Rothman”)
	Further, the evidence cited for claim limitation [7.pre] applies to this claim limitation. To the extent that Rothman does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Rothman, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[7.2] storing leaf images for respective compute nodes on respective second storage units, said leaf images including only additional data blocks not previously contained in said root image and changes made by respective compute nodes to the blocks of the root image, wherein said leaf images of respective compute nodes do not include blocks of said root image that are unchanged by respective compute nodes; and	Rothman discloses or renders obvious storing leaf images for respective compute nodes on respective second storage units, said leaf images including only additional data blocks not previously contained in said root image and changes made by respective compute nodes to the blocks of the root image, wherein said leaf images of respective compute nodes do not include blocks of said root image that are unchanged by respective compute nodes. For example: “In a decision block 404, a determination is made to whether any filter indicia is marked in the header data. If filter indicia is found, the packet processing is redirected to slow path processing in a block 406. In a block 408, a copy of each packet is then cached prior to forwarding the packet, the cached packets are then reassembled, if necessary to form a file or message, and the file is stored or the message is handled. In general, a cached file will be stored on a local non-volatile store, such as a local hard disk built into the network device chassis or otherwise communicatively-coupled to the network device. For example, the network device could actually store the cached files on a local file server that it accesses via a network link. Message handling operations are typically performed in the manner described above for handling PXE download request and boot image download request messages. If a boot image file or NBP file is cached, a local image map stored by the network device is updated in a block 410 to reflect the addition of a new file that is now cached by the network device. The logic then returns to block 400 to process the next packet. If no filter indicia is found, fast path processing is simply continued to forward the packet, as depicted by a block 412.” Rothman at 10:61-11:16.

Exhibit D-21

Claims	United States Patent No. 7,398,382 (“Rothman”)
	Further, the evidence cited for claim limitations [7.pre]-[7.1] apply to this claim limitation. To the extent this reference does not expressly or inherently disclose this limitation, it would have been obvious in view of the background knowledge of a person of ordinary skill in the art. This reference further renders the limitation obvious in light of other references, such as Sapuntzakis, Birse, and/or Kulkarni. Sapuntzakis, Birse, and Kulkarni concern similar subject matter as this reference, namely, distributed computing systems. As set forth in the accompanying contentions, storing data blocks not previously contained in a root image and/or changes made by compute nodes to the blocks of the root image was a well-known and commonly-used technique. Modifying this reference with Sapuntzakis, Birse, and/or Kulkarni’s teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. A person of ordinary skill in the art would have further understood how to modify the disclosures of this reference with those of Sapuntzakis, Birse, or Kulkarni for example, to incorporate this limitation. <i>See, e.g.,</i> “We can store the disks in these capsules efficiently by creating a hierarchy, where each child capsule could be viewed as inheriting from the parent capsule with the differences in disk state between parent and child captured in a separate copy-on-write (COW) virtual disk.” Sapuntzakis at 5. “Writes to a capsule disk are performed by writing the data to the latest COW disk and updating its bitmap file. Reads involve searching the latest COW disk and its ancestor disks in turn until the required block is found. Since the root COW disk contains a copy of all the blocks, the search is guaranteed to terminate.” Sapuntzakis at 5–6. “Since the latest COW disk is always local, all writes are local.” Sapuntzakis at 6.

Exhibit D-21

Claims	United States Patent No. 7,398,382 (“Rothman”)
	The diagram illustrates a capsule disk structure. It consists of a Bitmap and DiskBlocks. The Bitmap is a sequence of bits: 11111111, 10010101, 11000001, and 01011000. The DiskBlocks are represented by a grid of blocks. The Root COW Disk is the first disk, followed by COW Disk 1, COW Disk 2, and the Latest COW Disk. The Flat Disk (as seen by the capsule) is the final disk. A legend indicates that a shaded block represents a block present in a COW disk. Figure 2: An example capsule disk. Sapuntzakis at 6. See also Sapuntzakis at 1-12. “The block device driver 154 services file system read and write requests to the system volumes 160. In one embodiment, the block device driver 154 appears to the file system 152 to be a standard hard drive device driver. However, in reality, the block device driver 154 may be configured to service the read and write requests of the file system 152 by accessing the NC server 170. As will be discussed further below, in one embodiment, one or more shadow volumes, may be created for the NC client 150 and stored with the non-persistent client-specific data 184 for the purpose of storing user preferences and/or changes to the operating system. Therefore, some mechanism is needed to direct reads and writes to the appropriate volume, e.g., the shadow volumes or the system volumes 174. According to one embodiment, the block device driver 154 has knowledge of the separate shadow volumes and system volumes 174 and redirects read and write requests to the

Exhibit D-21

Claims	United States Patent No. 7,398,382 (“Rothman”)
	appropriate volume. Alternatively, server-side logic may perform the redirection locally.” Birse at 7:35-53; <i>see also</i> Birse at 11:30-32. “In the embodiment depicted, it is assumed that the shadow volume contains only those portions of the operating system that have been modified by the user rather than a complete copy of the operating system as modified. The various options regarding the granularity of the portions written to the shadow volume are discussed below.” Birse 11:49-52. <i>See also</i> Birse at 10:63-65, 11:6-17, 12:54–59, 14:49-61, FIGS. 4, 6, 10. “FIG. 2 is a flow diagram outlining the acts performed by the server 102 to create the differences file 120 from the old image 114 and the new image 116. In the present example, the updated image is an updated device operating system for the PDA 104.” Kulkarni at [0036]. “Creating the differences file 120 is accomplished by applying a compression algorithm to the new image 116, using the old image 114 as a reference. Since the old image 114 is already stored on the PDA 104, data from the new image 116 that is also contained in the old image 114 does not need to be downloaded from the server 102 to the PDA 104. When the new image 116 is created at the PDA 104, data that is contained in both images is copied from the image 139 already on the PDA 104. Data that is only contained in the new image 116 is downloaded to the PDA 104 in the modified differences file 124 and is copied to the flash memory image 139.” Kulkarni at [0037]. “The differences file 120 and the modified difference file 124 are composed of two alternating types of sections: data sections contain data from the new image 116 that is not found in the old image 114; copy sections contain data that is found in both the new image 116 and the old image 114. To distinguish between the two types of sections, a header is placed at the beginning of each section, the header giving specific information about the section that follows the header.” Kulkarni at [0038].

Exhibit D-21

Claims	United States Patent No. 7,398,382 (“Rothman”)
	<i>See also</i> Kulkarni at [0026]-[0041], [0068], FIG. 1, FIGS. 2a-2d, FIG. 6. To the extent that Rothman does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Rothman, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[7.3] caching blocks of said root image that have been accessed by at least one of said compute nodes in a cache memory.	Rothman discloses or renders obvious caching blocks of said root image that have been accessed by at least one of said compute nodes in a cache memory. For example: “Methods and apparatus for improving network boot efficiency are disclosed. Under embodiments of the method, boot images that are initially sent from a boot server to various clients are cached at network devices along communication paths between the boot server and the clients. In response to subsequent boot image requests from the clients, boot images cached at the network devices are downloaded directly from the network devices rather than from the boot server, reducing network traffic to and from the boot server and domains in between. In addition, boot program bootstrap files may also be cached and downloaded in a similar manner. Techniques are also disclosed for intercepting boot image download and network bootstrap program requests at the network devices, and for maintaining valid boot image cache configurations across the network. The network devices generally include switches, routers, bridges, and gateway servers.” Rothman at Abstract. “In a block 204, the client obtains an operating system boot image over the network from a boot server. In conjunction with delivering the boot image to the client, the boot image is cached at a network device, such as a switch, router, bridge, etc., or a gateway server that is proximate to the client in a block 206.” Rothman at 4:6-11. “In general, a cached file will be stored on a local non-volatile store, such as a local hard disk built into the network device chassis or otherwise communicatively-coupled to the network device. For example, the network device could actually store the cached files on a local file server that it accesses via a network link.” Rothman

Exhibit D-21

Claims	United States Patent No. 7,398,382 (“Rothman”)
	at 11:1-6. “Architecture blocks corresponding to an exemplary network device 900 that supports the network device caching operations described above are shown in FIG. 9. The architecture blocks include a media interface 902, an SRAM store 904, a DRAM store 906, a fast path processing block 908, a slow path processing block 910, network services 911, a non-volatile (NV) store 912, and a local file store 914. The fast path processing block includes a packet header filter 916, while the slow path processing block includes a file system 918, cache management logic 920, and message handler 922.” Rothman at 14:16-25. Further, the evidence cited for claim limitations [7.pre]-[7.2] apply to this claim limitation. To the extent this reference does not expressly or inherently disclose this limitation, it would have been obvious in view of the background knowledge of a person of ordinary skill in the art. This reference further renders the limitation obvious in light of other references, such as Holzmann, Nelson, and/or Stahl. Holzmann, Nelson, and Stahl concern similar subject matter as this reference, namely, distributed computing systems. As set forth in the accompanying contentions, caching data from comparatively slow storage devices in faster devices (including solid-state memory) was a well-known and commonly-used technique. Modifying this reference with Holzmann’s, Nelson’s, and/or Stahl’s caching teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. A person of ordinary skill in the art would have further understood how to modify the disclosures of this reference with those of Holzmann, Nelson, or Stahl for example, to incorporate this limitation. <i>See, e.g.,</i> “The invention is directed generally to computer storage systems. More specifically, the invention is directed to a system and method for booting multiple servers using a single operating system image located on a solid state disk in a storage area network.” Holzmann at [0002].

Exhibit D-21

Claims	United States Patent No. 7,398,382 (“Rothman”)
	<pre> graph TD 107[INTERNAL POWER SOURCE] --- 104((MEMORY)) 107 --- 103[CONTROL MODULE] 107 --- 105[INTERFACE MODULE] 104 <--> Data 103 103 <--> Data 102[NON-VOLATILE STORAGE MEANS] 104 <--> 105 103 <--> 105 105 <--> 106([EXTERNAL DEVICES]) subgraph 101 [Solid State Disk System] 107 104 103 102 105 end </pre> Fig. 1 Holzmann at FIG. 1 “Referring now to the figures, the present invention is directed to a system and method for booting multiple servers or other external devices from a storage area network using a single OS image. FIG. 1 is a block diagram illustrating the component parts of the invented system and the functions of each component part. The invented system comprises a solid state disk system 101 having a storage means 102, a control module 103, a memory module 104, an interface module 105 that communicates with external devices 106, and an internal power source 107.” Holzmann at [0017]. “The memory module 104 comprises at least one direct-access memory module for holding data currently or recently used by external devices 106. The memory module 104 is more quickly accessible, and performs read and write processing functions more quickly, than non-volatile or disk memory, such as storage means

Exhibit D-21

Claims	United States Patent No. 7,398,382 (“Rothman”)
	102. The memory module 104 preferably comprises at least one random-access memory (RAM) module.” Holzmanna at [0020]. “This invention relates in general to the fields of Computer Operating Systems, Multi-processor Hardware Systems, Object Oriented Programming, and Virtual Memory Systems. In particular, the invention relates to improved techniques for establishing and efficiently handling relationships between a client program, a file system server and a virtual memory manager (‘VMM’).” Nelson at [0001]. “This servicing of multiple client programs is now illustrated by means of Figure 6. A first client program 102 asks the CFS 218 to cache remote file 200. The CFS 218 provides a first CFS file object 108 for the first client 102 to use in processing requests on the file. The CFS 218 establishes a communications link 244 to the file server 206 located on the remote node 204 to obtain the file attributes which the CFS 218 caches in the CFS attribute cache 239. The CFS 218 also establishes the communications between the CFS file object manager 221 and the VMM 230 which caches the file data in the data cache 231 by means of the pager_object 234/fs_cache_object 232 pipeline. This pipeline results from the map and bind operations previously described. Pager object 234 is connected to fs_cache object 232 in the pager object manager 247, and also connected via the link between the fs_cache object 232 and the cache object 223 in the fs_cache object manager 243 in the CFS 218. When a second client program 118 asks the CFS 218 to cache the same remote file 200, the CFS 218 implements a second CFS file object 122. However, since the CFS 218 has already set up a CFS attribute cache 239 to handle file attributes for the remote file 200, the second CFS file object 122 is merely directed to the existing CFS attribute cache 239. Similarly, when the second CFS file object 122 wishes to read or write data to the file 200, the file object manager 221 again invokes an operation to map the file in the CFS address_space using its address_space object 225. Once again this causes the VMM's address_space object manager 227 to invoke a bind () operation on the file object 122 which is again directed to the CFS 218.

Exhibit D-21

Claims	United States Patent No. 7,398,382 (“Rothman”)
	This time the CFS 218 checks its cache of bind results 237 to see if that file 200 is already being cached by the VMM 230. When the CFS 218 finds that it is already being cached, the CFS 218 returns a pointer to the fs_cache_object/pager_object pair being used, and the bind operation is completed. Thus, duplicate caches are not established and the redundant network communications overhead is avoided by the process.” Nelson at [0038]-[0039]. “In an alternate embodiment, if the first client program 102 had set up the first CFS file object 108 with the CFS 218, the first client program 102 may pass the first CFS file object 108 to the second client program 118. In that case, both client programs 102 & 118 would use the same CFS file object 108, and as before both client programs would be serviced by the CFS 218 with only one set of caches.” Nelson at [0040].

Exhibit D-21

Claims	United States Patent No. 7,398,382 (“Rothman”)
	Nelson at FIG. 6. See also Nelson at [0010], [0034]-[0037], [0052]-[0119], FIG. 4. “Repository FS Implementation 222 is an implementation of a file system responsible for maintaining data on the persistent read-only file system 226 that is to have its data shared. Repository FS Implementation 222 receives and responds to file system requests from MapFS Implementation 212. Repository FS Instance in memory 224 represents a transient in-memory status of a particular Repository file system in persistent storage. Status includes, for example, information about what

Exhibit D-21

Claims	United States Patent No. 7,398,382 (“Rothman”)
	files exist (or do not) and have been successfully or unsuccessfully looked up, file size, permission data, actual cached data blocks, etc.” Stahl at [0028]. “The mapping module 502 looks 308 in the View FS instance 208 (via file system communication module 506) to see whether the file is cached. Preferably, the view caches filenames in a manner similar to MapFS—that is, in a table of names and look-up results. Again, since this is the first time the file is being looked up, it will not exist in the View's cache either, unless it has been looked up through some non-MapFS process. Thus, the mapping module 502 asks the 310 View FS implementation 206 to look up the name in the View in a manner preferably similar to the way in which the VFS layer examines MapFS instance data described above. The View FS Implementation 206 looks up 312 the file in the persistent view FS instance on disk 210 and updates 314 the View FS instance in memory 208 by populating an in-memory description of the file, and inserting an entry in its name-to-inode table. The View FS implementation 206 returns 316 a reference to the file in the View FS Instance 208 to the mapping module 502.” Stahl at [0037]. <i>See also</i> Stahl at paragraphs [0017], [0044]. To the extent that Rothman does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Rothman, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[11.0] The method as recited in claim 7 further comprising: merging the blocks of said root image with the blocks of respective leaf images to create cohesive respective application environments.	Rothman discloses or renders obvious the method as recited in claim 7 further comprising merging the blocks of said root image with the blocks of respective leaf images to create cohesive respective application environments. For example: “If a boot image file or NBP file is cached, a local image map stored by the network device is updated in a block 510 to reflect the addition of a new file that is now cached by the network device.” Rothman at 11:64-67.

Exhibit D-21

Claims	United States Patent No. 7,398,382 (“Rothman”)
	Further, the evidence cited for claim limitations [7.pre]-[7.3] apply to this claim limitation. To the extent this reference does not expressly or inherently disclose this limitation, it would have been obvious in view of the background knowledge of a person of ordinary skill in the art. This reference further renders the limitation obvious in light of other references, such as Sapuntzakis, Birse, Stahl, and/or Kulkarni. Sapuntzakis, Birse, Stahl, and Kulkarni concern similar subject matter as this reference, namely, distributed computing systems. As set forth in the accompanying contentions, merging data blocks for multiple computer nodes in distributed environments was a well-known and commonly-used technique. Modifying this reference with Sapuntzakis, Birse, Stahl, and/or Kulkarni’s teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. A person of ordinary skill in the art would have further understood how to modify the disclosures of this reference with those of Sapuntzakis, Birse, Stahl, and/or Kulkarni for example, to incorporate this limitation. <i>See, e.g.,</i> “To reduce the start-up time of a capsule, the COW disks corresponding to a capsule disk are read page-by- page on demand, rather than being pre-fetched.” Sapuntzakis at 6. “Reads involve searching the latest COW disk and its ancestor disks in turn until the required block is found.” Sapuntzakis at 5-6. “Note that the COW disk hierarchy is not visible to the VMM, or to the OS and applications inside the capsule; all of them see a normal flat disk as illustrated in the figure.” Sapuntzakis at 6. <i>See also</i> Sapuntzakis at 1-12. “In one embodiment, the shadow system volume 930 is not visible to the user. Consequently, the user system volume 924 will appear to be the location where all

Exhibit D-21

Claims	United States Patent No. 7,398,382 (“Rothman”)
	writes are directed. For example, if a user modifies a preference, that preference will appear to be written to the preference file located in the preferences folder on the user system volume 924.” Birse at 14:61-67. “When an NC client boots from the network and accesses the operating system from an NC server, the user's preferences are dynamically merged with the system environment provided to the NC client. Advantageously, since, the user's desktop preferences and other customized settings are all preserved from session to session and supplied to the NC client as it boots from the network, the user may login to any NC client on the network and have the same user experience.” Birse at 9:2-10. “Referring now to FIG. 5, NC client read request processing will now be described according to one embodiment of the present invention. ... At step 510, the file system 152 generates a read request directed to the system volumes 160. At step 520, the read request is received by the block device driver 154 and a determination is made whether the read request specifies a portion of the operating system that has been modified by the user or whether the read request specifies a portion of the operating system that remains unchanged. If the read request corresponds to a portion of the operating system that has not been modified by the user, the processing continues with step 530. However, if the read request corresponds to a portion of the operating system that has been previously modified by the user, then processing continues with step 540. At step 530, the block device driver directs the NC server 170 to retrieve the information associated with the read request from the system volumes 174. At step 540, the block device driver directs the NC server 170 to retrieve the information associated with the read request from the user's shadow volume. In this manner, the file system 152 need not be aware of the particular mechanism used to coordinate user modifications with the original operating system as contained in the system volumes 174.” Birse at 11:39-41, 11:55-12:10. “The implementations described herein create and utilize a differences file to create a new image in flash memory of an electronic device from sections of an old image in the flash memory and sections of a new image, the new image being an updated

Exhibit D-21

Claims	United States Patent No. 7,398,382 (“Rothman”)
	version of the old image. The new image is created incrementally in flash memory so that an entire image does not have to first be constructed in device RAM, thereby saving significant overhead. A “flash manager” is described which allows the new image to be built incrementally until a memory block is created that is of sufficient size to be written to the flash memory.” Kulkarni at [0013]. “Repository FS Implementation 222 is an implementation of a file system responsible for maintaining data on the persistent read-only file system 226 that is to have its data shared. Repository FS Implementation 222 receives and responds to file system requests from MapFS Implementation 212. Repository FS Instance in memory 224 represents a transient in-memory status of a particular Repository file system in persistent storage. Status includes, for example, information about what files exist (or do not) and have been successfully or unsuccessfully looked up, file size, permission data, actual cached data blocks, etc.” Stahl at [0028]. “Overlay FS Implementation 214 is an implementation of a file system responsible for providing writeable storage on the persistent overlay file system instance 220. Overlay FS Implementation 214 receives and responds to file system requests—including write requests—from MapFS Implementation 212. Overlay FS Instance in memory 218 behaves analogously to Repository FS Instance 224. In a preferred embodiment, the Overlay FS is a directory tree rooted at a location specified when the MapFS is instantiated.” Stahl at [0029]. “The flash manager 136 uses the modified differences file 124’ and the image 139 in flash memory 132 (the old image) to start building the new image. At step 404, a portion of the new image is constructed and temporarily stored in the new memory block 138. When the new memory block 138 contains a portion of the new image that constitutes at least one-half of a write block size for the flash memory 132 (“Yes” branch, step 406), then the portion of the new image stored in the new memory block 138 is added to the image 139 in the flash memory 132. As previously discussed, the size of a write block depends on the type of flash memory, i.e., the brand of the flash memory. It is common for flash memory to have a write block of

Exhibit D-21

Claims	United States Patent No. 7,398,382 (“Rothman”)
	64 Kb. Also as previously discussed, the flash manager 136 may manage several new memory blocks 138. When any of the new memory blocks 138 reaches a size of one-half of a write block size, that new memory block 138 is written to the flash memory.” Kulkarni at [0069]. <i>See also</i> Kulkarni at [0026]-[0041], [0068], FIG. 1, FIGS. 2a-2d, FIG. 6. To the extent that Rothman does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Rothman, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.

Exhibit D-23

Invalidity of U.S. Patent No. 8,332,844 in View of Southwest Airlines NetApp System

Southwest Airlines NetApp System was known in this country, used in public, sold or offered for sale no later than 2002, as demonstrated by at least the materials cited herein or the testimony of knowledgeable witnesses and corroborating documents.¹ Specifically, based on information currently available, Southwest Airlines NetApp System was (1) known or used in this country before the alleged invention of the claimed subject matter of the asserted claims, (2) in public use or on sale in this country more than one year before the filing date of the patent, or (3) invented by another who did not abandon, suppress, or conceal, before the alleged invention of the claimed subject matter of the asserted claim. Thus, Southwest Airlines NetApp System is prior art to the '844 Patent under at least Pre-AIA 35 U.S.C. §§ 102(a) or (g) as to claims 7 and 11 (the "Asserted Claims") of the '844 Patent, as discussed in detail in the chart that follows.

At least the following documents describe the relevant functionality disclosed by Southwest Airlines NetApp System:

- <https://web.archive.org/web/20040708152749/http://www.netapp.com/news/feature/2002/southwest.pdf>;
- <https://web.archive.org/web/20040719184445/http://www.netapp.com:80/ftp/snapshot-brochure.pdf>;
- <https://web.archive.org/web/20040805003209/http://www.netapp.com/ftp/reliability-brochure.pdf>;
- <https://web.archive.org/web/20040717035040/http://www.netapp.com/ftp/product-overview.pdf>;
- <https://web.archive.org/web/20040719153814/http://www.netapp.com:80/ftp/snapshot-brief.pdf>; and
- <https://web.archive.org/web/20030322200556/http://www.netapp.com:80/ftp/snapstore.pdf>.

Based on information currently available, the persons or entities involved in the conception and diligent reduction to practice of the ideas in Southwest Airlines NetApp System include at least persons involved in the Southwest Airlines NetApp System and their testing and commercial use, including at least Southwest Airlines Co. The conception of Southwest Airlines NetApp System occurred at least as early as 2002.

American offers this invalidity chart based on its current understanding of Southwest Airlines NetApp System. American's investigation of this system is ongoing. American may rely on materials produced by Plaintiffs or by third parties regarding Southwest

¹ Defendant's invalidity charts, in some instances, rely at least in part on Plaintiff's apparent positions regarding the scope of its claims for purposes of asserting infringement. Nothing in these claim charts should be understood as an admission relating to infringement, either literally or under the doctrine of equivalents, or as an admission relating to Defendant's understanding of the proper interpretation or scope of the Asserted Claims. Defendant reserves the right to rely on additional citations or sources of evidence that may also be applicable, or that may become applicable in light of any Court Order on claim construction, changes in Plaintiff's infringement contentions, or information obtained during discovery as the case proceeds.

Exhibit D-23

Airlines NetApp System and related software to demonstrate invalidity of the asserted claims. American has not had an opportunity to review source code for Southwest Airlines NetApp System. American may rely on source code, when available, implementing the functionality described herein, or other source code implementing substantially similar functionality, to demonstrate the invalidity of the asserted claims.

To the extent Plaintiff alleges Southwest Airlines NetApp System does not disclose any particular limitation of the Asserted Claims of the '844 Patent, either expressly or inherently, any purported differences are such that the claimed subject matter as a whole would have been obvious in view of the knowledge of one skilled in the art. It would have further been obvious to a person of ordinary skill in the art as of the priority date of the '844 Patent to modify Southwest Airlines NetApp System or combine the teachings of Southwest Airlines NetApp System with other prior art in a manner that would have rendered the Asserted Claims invalid as obvious, including but not limited to:

- U.S. Patent Application Publication No. 2005/0283597 (“Holzmann”);
- EP Patent No. EP0604013 (“Nelson”);
- Optimizing the Migration of Virtual Computers” (“Sapuntzakis”);
- U.S. Patent No. 7,089,300 (“Birse”);
- U.S. Patent Application Publication No. 2005/0004925 (“Stahl”); and/or
- U.S. Patent Application Publication No. 2002/0034105 (“Kulkarni”).

As the United States Supreme Court held in *KSR Int’l Co. v. Teleflex, Inc.*, “[t]he combination of familiar elements according to known methods is likely to be obvious when it does no more than yield predictable results.” 550 U.S. 398, 416 (2007). The Supreme Court further held: “[w]hen a work is available in one field of endeavor, design incentives and other market forces can prompt variations of it, either in the same field or a different one. If a person of ordinary skill can implement a predictable variation, § 103 likely bars its patentability. For the same reason, if a technique has been used to improve one device, and a person of ordinary skill in the art would recognize that it would improve similar devices in the same way, using the technique is obvious unless its actual application is beyond his or her skill.” *Id.* at 417. The Supreme Court further held: “[w]hen there is a design need or market pressure to solve a problem and there are a finite number of identified, predictable solutions, a person of ordinary skill has good reason to pursue the known options within his or her technical grasp. If this leads to the anticipated success, it is likely the product not of innovation but of ordinary skill and common sense.” *Id.* at 421.

Defendant reserves the right to amend or supplement this claim chart at a later date as more fully set forth in the Invalidity Contentions.

Exhibit D-23

Claims	Southwest Airlines NetApp System
[7.pre] A method for providing data to a plurality of compute nodes, comprising:	To the extent the preamble is limiting, Southwest Airlines NetApp System discloses or renders obvious a method for providing data to a plurality of compute nodes. For example: The NetApp system implemented at Southwest Airlines in 2002 exemplifies a robust and efficient method for managing root and leaf image data. The system architecture at Southwest leveraged clustered NetApp filers integrated with Oracle9i RAC to achieve high availability and seamless data management. The key technologies used in this setup included WAFL (Write Anywhere File Layout), SnapMirror, and SnapRestore. These technologies facilitated the management of root and leaf image data, maintained consistent application environments, and supported rapid recovery after system disruptions. To the extent that Southwest Airlines NetApp System does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Southwest Airlines NetApp System, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[7.1] storing blocks of a root image of said compute nodes on a first storage unit;	Southwest Airlines NetApp System discloses or renders obvious storing blocks of a root image of said compute nodes on a first storage unit. For example: The root image, which contains the foundational system data and configuration, was maintained on a centralized NetApp Filer within the Southwest Airlines system. WAFL technology, a core component of the NetApp system, enabled efficient storage and handling of snapshots by managing data without duplicating unchanged blocks. This approach ensured that the root image could be consistently accessed and maintained across multiple compute nodes. The root image is stored on a NetApp Filer, utilizing WAFL (Write Anywhere File Layout) for efficient data management. The root image is maintained as a snapshot. As highlighted in the Southwest Airlines case study, “[Southwest Airlines] switched from a single system to a clustered approach based on Oracle9i™ Real Application Clusters (RAC), which makes it possible to run an Oracle® database across a cluster of multiple servers. If one server in a cluster goes down, another can pick up the workload instantly, and the system can be easily expanded by adding more servers. . . . NetApp filers are able to handle storage tasks with speed and efficiency, and they feature storage-specific intelligence and tools. . . . ‘The tight integration between Oracle9i RAC and NetApp filers gives [Southwest Airlines] the highly available total package we need for these key applications.’” See https://web.archive.org/web/20040708152749/http://www.netapp.com/news/feature/2002/southwest.pdf.

Exhibit D-23

Claims	Southwest Airlines NetApp System
	To the extent that Southwest Airlines NetApp System does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Southwest Airlines NetApp System, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant's Invalidity Contentions.
[7.2] storing leaf images for respective compute nodes on respective second storage units, said leaf images including only additional data blocks not previously contained in said root image and changes made by respective compute nodes to the blocks of the root image, wherein said leaf images of respective compute nodes do not include blocks of said root image that are unchanged by respective compute nodes; and	Southwest Airlines NetApp System discloses or renders obvious storing leaf images for respective compute nodes on respective second storage units, said leaf images including only additional data blocks not previously contained in said root image and changes made by respective compute nodes to the blocks of the root image, wherein said leaf images of respective compute nodes do not include blocks of said root image that are unchanged by respective compute nodes. For example: When data modifications occurred, WAFL's snapshot technology used Copy-on-Write (CoW) to efficiently record changes without duplicating the entire root image. Instead, WAFL captured only the data blocks that had been altered, creating leaf images that contained only additional data not previously present in the root image. This method significantly conserved storage space while maintaining the integrity of the original root image. NetApp documentation notes that "A Snapshot copy is a 'frozen,' read-only view of a WAFL volume that provides easy access to old versions of files, directory hierarchies, and/or LUNs (logical unit numbers). The system creates a Snapshot copy by preserving all the currently used disk blocks. A Snapshot copy is typically created in less than one second, regardless of the size of the volume or the level of activity on the NetApp storage system. After a Snapshot copy has been taken, changes to data objects are reflected in updates to the current version of the objects, as if the Snapshot copy did not exist. Meanwhile, the Snapshot versions of the data remain completely stable." See https://web.archive.org/web/20040719184445/http://www.netapp.com:80/ftp/snapshot-brochure.pdf. "A core capability of NetApp technology, Snapshot software delivers near-instantaneous, transparent, online incremental images of your data. It requires minimal disk space and operates invisibly in the background." See https://web.archive.org/web/20040805003209/http://www.netapp.com/ftp/reliability-brochure.pdf. Snapshots in WAFL capture only the changes since the last snapshot, keeping the base root image intact. Leaf images contain only additional data blocks and changes made by the compute nodes. The snapshot process in WAFL ensures that unchanged data remains part of the root image, while only the modified data is stored separately as a leaf image. This efficient use of space aligns with the patent claim

Exhibit D-23

Claims	Southwest Airlines NetApp System
	requirement that the leaf images exclude unchanged data. According to NetApp documentation, “Snapshot copies write only changed 4kB blocks to the active file system. Even the most efficient competitive implementations consume twice the storage volume of the active file system.” See https://web.archive.org/web/20040719184445/http://www.netapp.com:80/ftp/snapshot-brochure.pdf. “A core capability of NetApp technology, Snapshot software delivers near-instantaneous, transparent, online incremental images of your data. It requires minimal disk space and operates invisibly in the background.” See https://web.archive.org/web/20040805003209/http://www.netapp.com/ftp/reliability-brochure.pdf. “[Snapshot] Enables online backups, providing near-instantaneous access to previous versions of data without requiring complete, separate copies.” See https://web.archive.org/web/20040717035040/http://www.netapp.com/ftp/product-overview.pdf. WAFL’s CoW-based snapshot technology ensures that only changes made by the compute nodes are stored as leaf images, while unchanged data remains part of the root image. “<i>After a Snapshot copy has been created, changes to data objects are reflected in updates to the current version of the objects, as if Snapshot copies did not exist. Meanwhile, the Snapshot versions of the data remain completely stable. Therefore, Snapshot copies incur no performance overhead; users can comfortably store up to 255 Snapshot copies per WAFL volume, all of which are accessible as read-only and online versions of the data.</i>” See https://web.archive.org/web/20040719153814/http://www.netapp.com:80/ftp/snapshot-brief.pdf. To the extent this reference does not expressly or inherently disclose this limitation, it would have been obvious in view of the background knowledge of a person of ordinary skill in the art. This reference further renders the limitation obvious in light of other references, such as Sapuntzakis, Birse, and/or Kulkarni. Sapuntzakis, Birse, and Kulkarni concern similar subject matter as this reference, namely, distributed computing systems. As set forth in the accompanying contentions, storing data blocks not previously contained in a root image and/or changes made by compute nodes to the blocks of the root image was a well-known and commonly-used technique. Modifying this reference with Sapuntzakis, Birse, and/or Kulkarni’s teachings would have only required simple substitution, would not have required

Exhibit D-23

Claims	Southwest Airlines NetApp System
	extensive experimentation, and would have yielded predictable results. A person of ordinary skill in the art would have further understood how to modify the disclosures of this reference with those of Sapuntzakis, Birse, or Kulkarni for example, to incorporate this limitation. <i>See, e.g.</i>, “We can store the disks in these capsules efficiently by creating a hierarchy, where each child capsule could be viewed as inheriting from the parent capsule with the differences in disk state between parent and child captured in a separate copy-on-write (COW) virtual disk.” Sapuntzakis at 5. “Writes to a capsule disk are performed by writing the data to the latest COW disk and updating its bitmap file. Reads involve searching the latest COW disk and its ancestor disks in turn until the required block is found. Since the root COW disk contains a copy of all the blocks, the search is guaranteed to terminate.” Sapuntzakis at 5–6. “Since the latest COW disk is always local, all writes are local.” Sapuntzakis at 6. Bitmap 11111111 10010101 11000001 01011000 DiskBlocks Root COW Disk COW Disk 1 COW Disk 2 Latest COW Disk Flat Disk (as seen by the capsule) Block present in COW disk Figure 2: An example capsule disk. Sapuntzakis at 6.

Exhibit D-23

Claims	Southwest Airlines NetApp System
	<i>See also</i> Sapuntzakis at 1-12. “The block device driver 154 services file system read and write requests to the system volumes 160. In one embodiment, the block device driver 154 appears to the file system 152 to be a standard hard drive device driver. However, in reality, the block device driver 154 may be configured to service the read and write requests of the file system 152 by accessing the NC server 170. As will be discussed further below, in one embodiment, one or more shadow volumes, may be created for the NC client 150 and stored with the non-persistent client-specific data 184 for the purpose of storing user preferences and/or changes to the operating system. Therefore, some mechanism is needed to direct reads and writes to the appropriate volume, e.g., the shadow volumes or the system volumes 174. According to one embodiment, the block device driver 154 has knowledge of the separate shadow volumes and system volumes 174 and redirects read and write requests to the appropriate volume. Alternatively, server-side logic may perform the redirection locally.” Birse at 7:35-53; <i>see also</i> Birse at 11:30-32. “In the embodiment depicted, it is assumed that the shadow volume contains only those portions of the operating system that have been modified by the user rather than a complete copy of the operating system as modified. The various options regarding the granularity of the portions written to the shadow volume are discussed below.” Birse 11:49-52. <i>See also</i> Birse at 10:63-65, 11:6-17, 12:54-59, 14:49-61, FIGS. 4, 6, 10. “FIG. 2 is a flow diagram outlining the acts performed by the server 102 to create the differences file 120 from the old image 114 and the new image 116. In the present example, the updated image is an updated device operating system for the PDA 104.” Kulkarni at [0036]. “Creating the differences file 120 is accomplished by applying a compression algorithm to the new image 116, using the old image 114 as a reference. Since the old image 114 is already stored on the PDA 104, data from the new image 116 that is also contained in the old image 114 does not need to be downloaded from the server 102 to the PDA 104. When the new image 116 is created at the PDA 104, data that is contained in both images is copied from the image 139 already on the PDA 104. Data

Exhibit D-23

Claims	Southwest Airlines NetApp System
	that is only contained in the new image 116 is downloaded to the PDA 104 in the modified differences file 124 and is copied to the flash memory image 139.” Kulkarni at [0037]. “The differences file 120 and the modified difference file 124 are composed of two alternating types of sections: data sections contain data from the new image 116 that is not found in the old image 114; copy sections contain data that is found in both the new image 116 and the old image 114. To distinguish between the two types of sections, a header is placed at the beginning of each section, the header giving specific information about the section that follows the header.” Kulkarni at [0038]. <i>See also</i> Kulkarni at [0026]-[0041], [0068], FIG. 1, FIGS. 2a-2d, FIG. 6. To the extent that Southwest Airlines NetApp System does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Southwest Airlines NetApp System, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[7.3] caching blocks of said root image that have been accessed by at least one of said compute nodes in a cache memory.	Southwest Airlines NetApp System discloses or renders obvious caching blocks of said root image that have been accessed by at least one of said compute nodes in a cache memory. For example: To enhance performance and reduce latency, the Southwest system utilized Data ONTAP’s caching features. Frequently accessed root image data was temporarily stored in cache memory, which minimized disk I/O and facilitated faster data retrieval. This capability was crucial for maintaining high performance during peak system demands. As explained in the NetApp reliability brochure, “Our Data ONTAP™ microkernel ensures reliable operation through high-performance, highly optimized code.” <i>See</i> https://web.archive.org/web/20040805003209/http://www.netapp.com/ftp/reliability-brochure.pdf. “[Data ONTAP is] Operating system software that optimizes data serving by combining patented file system technology and a microkernel design dedicated to multiprotocol data access.” <i>See</i> https://web.archive.org/web/20040717035040/http://www.netapp.com/ftp/product-overview.pdf. Data ONTAP employs in-memory caching to optimize data retrieval and minimize disk I/O. “<i>Southwest Airlines discovered the value of Snapshot technologies early on, says Kerry Schwab. During implementation, IT professionals could take a Snapshot at key points, and then try a number of options. If we decided against going</i>

Exhibit D-23

Claims	Southwest Airlines NetApp System
	<i>down a certain path, we could start over very quickly,’ he says. ‘We could run through some what-if scenarios—what if I increase parallelism on this, or what if I spread the data out this way instead of that way? If we had been using traditional storage, we would have had to take several hours to make a backup, and several more hours to restore. The Snapshot feature let us try things without the incremental costs of dumping and restoring all the data.’”</i> <i>See</i> https://web.archive.org/web/20040708152749/http://www.netapp.com/news/feature/2002/southwest.pdf. To the extent this reference does not expressly or inherently disclose this limitation, it would have been obvious in view of the background knowledge of a person of ordinary skill in the art. This reference further renders the limitation obvious in light of other references, such as Holzmann, Nelson, and/or Stahl. Holzmann, Nelson, and Stahl concern similar subject matter as this reference, namely, distributed computing systems. As set forth in the accompanying contentions, caching data from comparatively slow storage devices in faster devices (including solid-state memory) was a well-known and commonly-used technique. Modifying this reference with Holzmann’s, Nelson’s, and/or Stahl’s caching teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. A person of ordinary skill in the art would have further understood how to modify the disclosures of this reference with those of Holzmann, Nelson, or Stahl for example, to incorporate this limitation. <i>See, e.g.,</i> “The invention is directed generally to computer storage systems. More specifically, the invention is directed to a system and method for booting multiple servers using a single operating system image located on a solid state disk in a storage area network.” Holzmann at [0002].

Exhibit D-23

Claims	Southwest Airlines NetApp System
	<pre> graph TD subgraph 101 [Solid State Disk System] 107[INTERNAL POWER SOURCE] 104((MEMORY)) 103[CONTROL MODULE] 102[NON-VOLATILE STORAGE MEANS] 105[INTERFACE MODULE] 104 <--> Data 103 103 <--> Data 102 105 <--> 104 105 <--> 103 105 <--> 107 end 105 <--> Data 106[EXTERNAL DEVICES] </pre> Fig. 1 Holzmann at FIG. 1 “Referring now to the figures, the present invention is directed to a system and method for booting multiple servers or other external devices from a storage area network using a single OS image. FIG. 1 is a block diagram illustrating the component parts of the invented system and the functions of each component part. The invented system comprises a solid state disk system 101 having a storage means 102, a control module 103, a memory module 104, an interface module 105 that communicates with external devices 106, and an internal power source 107.” Holzmann at [0017]. “The memory module 104 comprises at least one direct-access memory module for holding data currently or recently used by external devices 106. The memory module 104 is more quickly accessible, and performs read and write processing functions more quickly, than non-volatile or disk memory, such as storage means 102. The memory module 104 preferably comprises at least one random-access memory (RAM) module.” Holzmann at [0020]. “This invention relates in general to the fields of Computer Operating Systems, Multi-processor Hardware Systems, Object Oriented Programming, and Virtual Memory Systems. In particular, the

Exhibit D-23

Claims	Southwest Airlines NetApp System
	invention relates to improved techniques for establishing and efficiently handling relationships between a client program, a file system server and a virtual memory manager ('VMM')." Nelson at [0001]. "This servicing of multiple client programs is now illustrated by means of Figure 6. A first client program 102 asks the CFS 218 to cache remote file 200. The CFS 218 provides a first CFS file object 108 for the first client 102 to use in processing requests on the file. The CFS 218 establishes a communications link 244 to the file server 206 located on the remote node 204 to obtain the file attributes which the CFS 218 caches in the CFS attribute cache 239. The CFS 218 also establishes the communications between the CFS file object manager 221 and the VMM 230 which caches the file data in the data cache 231 by means of the pager_object 234/fs_cache_object 232 pipeline. This pipeline results from the map and bind operations previously described. Pager object 234 is connected to fs_cache object 232 in the pager object manager 247, and also connected via the link between the fs_cache object 232 and the cache object 223 in the fs_cache object manager 243 in the CFS 218. When a second client program 118 asks the CFS 218 to cache the same remote file 200, the CFS 218 implements a second CFS file object 122. However, since the CFS 218 has already set up a CFS attribute cache 239 to handle file attributes for the remote file 200, the second CFS file object 122 is merely directed to the existing CFS attribute cache 239. Similarly, when the second CFS file object 122 wishes to read or write data to the file 200, the file object manager 221 again invokes an operation to map the file in the CFS address_space using its address_space object 225. Once again this causes the VMM's address_space object manager 227 to invoke a bind () operation on the file object 122 which is again directed to the CFS 218. This time the CFS 218 checks its cache of bind results 237 to see if that file 200 is already being cached by the VMM 230. When the CFS 218 finds that it is already being cached, the CFS 218 returns a pointer to the fs_cache_object/pager_object pair being used, and the bind operation is completed. Thus, duplicate caches are not established and the redundant network communications overhead is avoided by the process." Nelson at [0038]-[0039]. "In an alternate embodiment, if the first client program 102 had set up the first CFS file object 108 with the CFS 218, the first client program 102 may pass the first CFS file object 108 to the second client program 118. In that case, both client programs 102 & 118 would use the same CFS file object

Exhibit D-23

Claims	Southwest Airlines NetApp System
	108, and as before both client programs would be serviced by the CFS 218 with only one set of caches.” Nelson at [0040]. FIGURE 6 Nelson at FIG. 6. See also Nelson at [0010], [0034]-[0037], [0052]-[0119], FIG. 4. “Repository FS Implementation 222 is an implementation of a file system responsible for maintaining data on the persistent read-only file system 226 that is to have its data shared. Repository FS Implementation 222 receives and responds to file system requests from MapFS Implementation 212.

Exhibit D-23

Claims	Southwest Airlines NetApp System
	Repository FS Instance in memory 224 represents a transient in-memory status of a particular Repository file system in persistent storage. Status includes, for example, information about what files exist (or do not) and have been successfully or unsuccessfully looked up, file size, permission data, actual cached data blocks, etc.” Stahl at [0028]. “The mapping module 502 looks 308 in the View FS instance 208 (via file system communication module 506) to see whether the file is cached. Preferably, the view caches filenames in a manner similar to MapFS—that is, in a table of names and look-up results. Again, since this is the first time the file is being looked up, it will not exist in the View's cache either, unless it has been looked up through some non-MapFS process. Thus, the mapping module 502 asks the 310 View FS implementation 206 to look up the name in the View in a manner preferably similar to the way in which the VFS layer examines MapFS instance data described above. The View FS Implementation 206 looks up 312 the file in the persistent view FS instance on disk 210 and updates 314 the View FS instance in memory 208 by populating an in-memory description of the file, and inserting an entry in its name-to-inode table. The View FS implementation 206 returns 316 a reference to the file in the View FS Instance 208 to the mapping module 502.” Stahl at [0037]. <i>See also</i> Stahl at paragraphs [0017], [0044]. To the extent that Southwest Airlines NetApp System does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Southwest Airlines NetApp System, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant's Invalidity Contentions.
[11.0] The method as recited in claim 7 further comprising: merging the blocks of said root image with the blocks of respective leaf images to create cohesive	Southwest Airlines NetApp System discloses or renders obvious the method as recited in claim 7 further comprising merging the blocks of said root image with the blocks of respective leaf images to create cohesive respective application environments. For example: An essential feature of the Southwest system was its ability to rapidly restore a consistent application environment after data loss or system changes. NetApp's SnapRestore technology played a pivotal role in this process by merging the root image with the modifications stored in snapshots. This rapid restoration ensured that the system could quickly revert to a previous stable state. The documentation emphasizes

Exhibit D-23

Claims	Southwest Airlines NetApp System
respective application environments.	that “[SnapRestore] Allows rapid restoration of a file system to an earlier point in time, typically in only a few seconds.” See https://web.archive.org/web/20040717035040/http://www.netapp.com/ftp/product-overview.pdf. SnapRestore technology from NetApp merges the root image with the changes made by compute nodes, creating a cohesive application environment after recovery. “SnapRestore software uses Snapshot technology to perform near-instantaneous data restoration. . . . With SnapRestore, data can be restored from any one of the Snapshots stored on the file system. This allows an application development team, for example, to revert to Snapshots from various stages of their design, or test engineers to quickly and easily return data to a baseline state. Restoring to the base environment takes only seconds, and the restored environment is identical to the point at which the Snapshot copy was created.” See https://web.archive.org/web/20030322200556/http://www.netapp.com:80/ftp/snaprestore.pdf. “While Snapshot software is perfect for recovering single files, SnapRestore is made for recovering multiterabyte volumes. SnapRestore uses Snapshot technology to perform near-instantaneous data restoration at the volume level, recovering an entire file system back to a known good state in seconds without interrupting users and applications on other volumes.” See https://web.archive.org/web/20040805003209/http://www.netapp.com/ftp/reliability-brochure.pdf. Further, the evidence cited for claim limitations [7.pre]-[7.3] apply to this claim limitation. To the extent this reference does not expressly or inherently disclose this limitation, it would have been obvious in view of the background knowledge of a person of ordinary skill in the art. This reference further renders the limitation obvious in light of other references, such as Sapuntzakis, Birse, Stahl, and/or Kulkarni. Sapuntzakis, Birse, Stahl, and Kulkarni concern similar subject matter as this reference, namely, distributed computing systems. As set forth in the accompanying contentions, merging data blocks for multiple computer nodes in distributed environments was a well-known and commonly-used technique. Modifying this reference with Sapuntzakis, Birse, Stahl, and/or Kulkarni’s teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. A person of ordinary skill in the art would have further understood how to modify the

Exhibit D-23

Claims	Southwest Airlines NetApp System
	disclosures of this reference with those of Sapuntzakis, Birse, Stahl, and/or Kulkarni for example, to incorporate this limitation. <i>See, e.g.,</i> “To reduce the start-up time of a capsule, the COW disks corresponding to a capsule disk are read page-by- page on demand, rather than being pre-fetched.” Sapuntzakis at 6. “Reads involve searching the latest COW disk and its ancestor disks in turn until the required block is found.” Sapuntzakis at 5-6. “Note that the COW disk hierarchy is not visible to the VMM, or to the OS and applications inside the capsule; all of them see a normal flat disk as illustrated in the figure.” Sapuntzakis at 6. <i>See also</i> Sapuntzakis at 1-12. “In one embodiment, the shadow system volume 930 is not visible to the user. Consequently, the user system volume 924 will appear to be the location where all writes are directed. For example, if a user modifies a preference, that preference will appear to be written to the preference file located in the preferences folder on the user system volume 924.” Birse at 14:61-67. “When an NC client boots from the network and accesses the operating system from an NC server, the user's preferences are dynamically merged with the system environment provided to the NC client. Advantageously, since, the user's desktop preferences and other customized settings are all preserved from session to session and supplied to the NC client as it boots from the network, the user may login to any NC client on the network and have the same user experience.” Birse at 9:2-10. “Referring now to FIG. 5, NC client read request processing will now be described according to one embodiment of the present invention. ... At step 510, the file system 152 generates a read request directed to the system volumes 160. At step 520, the read request is received by the block device driver 154 and a determination is made whether the read request specifies a portion of the operating system that has been modified by the user or whether the read request specifies a portion of the operating system that remains unchanged. If the read request corresponds to a portion of the operating system that has not been modified by the user, the processing continues with step 530. However, if

Exhibit D-23

Claims	Southwest Airlines NetApp System
	the read request corresponds to a portion of the operating system that has been previously modified by the user, then processing continues with step 540. At step 530, the block device driver directs the NC server 170 to retrieve the information associated with the read request from the system volumes 174. At step 540, the block device driver directs the NC server 170 to retrieve the information associated with the read request from the user's shadow volume. In this manner, the file system 152 need not be aware of the particular mechanism used to coordinate user modifications with the original operating system as contained in the system volumes 174." Birse at 11:39-41, 11:55-12:10. "The implementations described herein create and utilize a differences file to create a new image in flash memory of an electronic device from sections of an old image in the flash memory and sections of a new image, the new image being an updated version of the old image. The new image is created incrementally in flash memory so that an entire image does not have to first be constructed in device RAM, thereby saving significant overhead. A "flash manager" is described which allows the new image to be built incrementally until a memory block is created that is of sufficient size to be written to the flash memory." Kulkarni at [0013]. "Repository FS Implementation 222 is an implementation of a file system responsible for maintaining data on the persistent read-only file system 226 that is to have its data shared. Repository FS Implementation 222 receives and responds to file system requests from MapFS Implementation 212. Repository FS Instance in memory 224 represents a transient in-memory status of a particular Repository file system in persistent storage. Status includes, for example, information about what files exist (or do not) and have been successfully or unsuccessfully looked up, file size, permission data, actual cached data blocks, etc." Stahl at [0028]. "Overlay FS Implementation 214 is an implementation of a file system responsible for providing writeable storage on the persistent overlay file system instance 220. Overlay FS Implementation 214 receives and responds to file system requests—including write requests—from MapFS Implementation 212. Overlay FS Instance in memory 218 behaves analogously to Repository FS Instance 224. In a preferred embodiment, the Overlay FS is a directory tree rooted at a location specified when the MapFS is instantiated." Stahl at [0029].

Exhibit D-23

Claims	Southwest Airlines NetApp System
	“The flash manager 136 uses the modified differences file 124' and the image 139 in flash memory 132 (the old image) to start building the new image. At step 404, a portion of the new image is constructed and temporarily stored in the new memory block 138. When the new memory block 138 contains a portion of the new image that constitutes at least one-half of a write block size for the flash memory 132 (“Yes” branch, step 406), then the portion of the new image stored in the new memory block 138 is added to the image 139 in the flash memory 132. As previously discussed, the size of a write block depends on the type of flash memory, i.e., the brand of the flash memory. It is common for flash memory to have a write block of 64 Kb. Also as previously discussed, the flash manager 136 may manage several new memory blocks 138. When any of the new memory blocks 138 reaches a size of one-half of a write block size, that new memory block 138 is written to the flash memory.” Kulkarni at [0069]. <i>See also</i> Kulkarni at [0026]-[0041], [0068], FIG. 1, FIGS. 2a-2d, FIG. 6. To the extent that Southwest Airlines NetApp System does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Southwest Airlines NetApp System, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.

Exhibit D-24

Invalidity of U.S. Patent No. 8,332,844 in View of NetBoot

Apple NetBoot was known in this country, used in public, sold or offered for sale no later than 2002, as demonstrated by at least the materials cited herein or the testimony of knowledgeable witnesses and corroborating documents.¹ Specifically, based on information currently available, Apple NetBoot was (1) known or used in this country before the alleged invention of the claimed subject matter of the asserted claims, (2) in public use or on sale in this country more than one year before the filing date of the patent, or (3) invented by another who did not abandon, suppress, or conceal, before the alleged invention of the claimed subject matter of the asserted claim. Thus, Apple NetBoot is prior art to the '844 Patent under at least Pre-AIA 35 U.S.C. §§ 102(a) or (g) as to claims 7 and 11 (the "Asserted Claims") of the '844 Patent, as discussed in detail in the chart that follows.

At least the following documents describe the relevant functionality disclosed by Apple NetBoot:

- https://ftpmirror.your.org/pub/misc/apple/Apple_Software_Updates/English-North_American/Mac_OS_X_Server/NetBoot_Desktop_Admin/Inside_NetBooting.pdf;
- https://ftpmirror.your.org/pub/misc/apple/Apple_Software_Updates/English-North_American/Mac_OS_X_Server/NetBoot_Desktop_Admin/;
- <https://www.macintoshrepository.org/15912-mac-os-x-v10-2-netboot-for-mac-os-9-cd-5/7/>; and
- https://web.archive.org/web/20201120105030/https://support.apple.com/kb/DL1192?locale=en_GB.

Based on information currently available, the persons or entities involved in the conception and diligent reduction to practice of the ideas in Apple NetBoot include at least persons involved in the Apple NetBoot and their testing and commercial use, including at least Apple Inc. The conception of Apple NetBook occurred at least as early as 2002.

American offers this invalidity chart based on its current understanding of Apple NetBoot. American's investigation of this system is ongoing. American may rely on materials produced by Plaintiffs or by third parties regarding Apple NetBoot and related software to demonstrate invalidity of the asserted claims. American has not had an opportunity to review source code for Apple NetBoot.

¹ Defendant's invalidity charts, in some instances, rely at least in part on Plaintiff's apparent positions regarding the scope of its claims for purposes of asserting infringement. Nothing in these claim charts should be understood as an admission relating to infringement, either literally or under the doctrine of equivalents, or as an admission relating to Defendant's understanding of the proper interpretation or scope of the Asserted Claims. Defendant reserves the right to rely on additional citations or sources of evidence that may also be applicable, or that may become applicable in light of any Court Order on claim construction, changes in Plaintiff's infringement contentions, or information obtained during discovery as the case proceeds.

Exhibit D-24

American may rely on source code, when available, implementing the functionality described herein, or other source code implementing substantially similar functionality, to demonstrate the invalidity of the asserted claims.

To the extent Plaintiff alleges Apple NetBoot does not disclose any particular limitation of the Asserted Claims of the '844 Patent, either expressly or inherently, any purported differences are such that the claimed subject matter as a whole would have been obvious in view of the knowledge of one skilled in the art. It would have further been obvious to a person of ordinary skill in the art as of the priority date of the '844 Patent to modify Apple NetBoot or combine the teachings of Apple NetBoot with other prior art in a manner that would have rendered the Asserted Claims invalid as obvious, including but not limited to:

- U.S. Patent Application Publication No. 2005/0283597 (“Holzmann”);
- EP Patent No. EP0604013 (“Nelson”);
- “Optimizing the Migration of Virtual Computers” (“Sapuntzakis”);
- U.S. Patent No. 7,089,300 (“Birse”);
- U.S. Patent Application Publication No. 2005/0004925 (“Stahl”); and/or
- U.S. Patent Application Publication No. 2002/0034105 (“Kulkarni”).

As the United States Supreme Court held in *KSR Int’l Co. v. Teleflex, Inc.*, “[t]he combination of familiar elements according to known methods is likely to be obvious when it does no more than yield predictable results.” 550 U.S. 398, 416 (2007). The Supreme Court further held: “[w]hen a work is available in one field of endeavor, design incentives and other market forces can prompt variations of it, either in the same field or a different one. If a person of ordinary skill can implement a predictable variation, § 103 likely bars its patentability. For the same reason, if a technique has been used to improve one device, and a person of ordinary skill in the art would recognize that it would improve similar devices in the same way, using the technique is obvious unless its actual application is beyond his or her skill.” *Id.* at 417. The Supreme Court further held: “[w]hen there is a design need or market pressure to solve a problem and there are a finite number of identified, predictable solutions, a person of ordinary skill has good reason to pursue the known options within his or her technical grasp. If this leads to the anticipated success, it is likely the product not of innovation but of ordinary skill and common sense.” *Id.* at 421.

Defendant reserves the right to amend or supplement this claim chart at a later date as more fully set forth in the Invalidity Contentions.

Exhibit D-24

Claims	Apple NetBoot
[7.pre] A method for providing data to a plurality of compute nodes, comprising:	To the extent the preamble is limiting, Apple NetBoot discloses or renders obvious a method for providing data to a plurality of compute nodes. For example: “NetBoot is a suite of software that supports these features: • Single point of administration-- As all computers that start up from a NetBoot server use the same client Mac OS image, you need only to manage this single image. • Personal desktop--Users can log in to any computer connected to a NetBoot server and obtain their personalized Macintosh user experience.” See https://ftpmirror.your.org/pub/misc/apple/Apple_Software_Updates/English-North_American/Mac_OS_X_Server/NetBoot_Desktop_Admin/Inside_NetBooting.pdf. “The Mac OS used by your client computers is downloaded from the server when a client computer starts up. The client Mac OS stored as a single image on the server. Every client computer uses the same Mac OS image. You need only to manage this image to administer all NetBoot clients.” See https://ftpmirror.your.org/pub/misc/apple/Apple_Software_Updates/English-North_American/Mac_OS_X_Server/NetBoot_Desktop_Admin/Inside_NetBooting.pdf. “You can use a NetBoot server and network-startup computers in a variety of computing environments. The simplest environment is a single lab that has one NetBoot server supporting network-startup computers. A more complex environment is an existing network of computers that spans buildings and crosses routers and that has a NetBoot server and network-startup computers integrated into it.” See https://ftpmirror.your.org/pub/misc/apple/Apple_Software_Updates/English-North_American/Mac_OS_X_Server/NetBoot_Desktop_Admin/Inside_NetBooting.pdf. To the extent that Apple NetBoot does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple NetBoot, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.

Exhibit D-24

Claims	Apple NetBoot
[7.1] storing blocks of a root image of said compute nodes on a first storage unit;	Apple NetBoot discloses or renders obvious storing blocks of a root image of said compute nodes on a first storage unit. For example: “Macintosh Management server--This server program stores and looks up information about user accounts and workgroups, provides authentication services, and can assign a default location on the server for users to store documents. The Macintosh Management server can be used to administer client computers that start up from a local hard drive.” See https://ftpmirror.your.org/pub/misc/apple/Apple_Software_Updates/English-North_American/Mac_OS_X_Server/NetBoot_Desktop_Admin/Inside_NetBooting.pdf. “If the user has permission to save files and folders, the items are saved to the disk images on the server, in a location you set up in Macintosh Manager. Users may also save items to another AFP server a local hard drive or other storage device if you set up permission in Macintosh Manager.” See https://ftpmirror.your.org/pub/misc/apple/Apple_Software_Updates/English-North_American/Mac_OS_X_Server/NetBoot_Desktop_Admin/Inside_NetBooting.pdf. To the extent that Apple NetBoot does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple NetBoot, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[7.2] storing leaf images for respective compute nodes on respective second storage units, said leaf images including only additional data blocks not previously contained in said root image and changes made by respective compute nodes to the blocks of the root image, wherein said leaf images of respective compute nodes do not include blocks of said root image	Apple NetBoot discloses or renders obvious storing leaf images for respective compute nodes on respective second storage units, said leaf images including only additional data blocks not previously contained in said root image and changes made by respective compute nodes to the blocks of the root image, wherein said leaf images of respective compute nodes do not include blocks of said root image that are unchanged by respective compute nodes. For example: “Once a connection is established between a client computer and the server, the Mac OS ROM image is copied from the server to the random access memory (RAM) in the client computer. This is done using Trivial File Transfer Protocol (TFTP).” See https://ftpmirror.your.org/pub/misc/apple/Apple_Software_Updates/English-North_American/Mac_OS_X_Server/NetBoot_Desktop_Admin/Inside_NetBooting.pdf.

Exhibit D-24

Claims	Apple NetBoot
that are unchanged by respective compute nodes; and	“For NetBoot client computers, you use Macintosh Manager to set options that define which preferences and settings can be saved. When user preferences are changed, the changes are stored in a location on the server that's referred to as a ‘shadow volume.’ The shadow volume contains the differences between the Mac OS image and what the user has changed. For example, if the user changes the desktop pattern to show Bubble Poppy instead of the Mac OS default, the changes to the Mac OS image to store this user preference are written to the shadow volume instead.” See https://ftpmirror.your.org/pub/misc/apple/Apple_Software_Updates/English-North_American/Mac_OS_X_Server/NetBoot_Desktop_Admin/Inside_NetBooting.pdf. To the extent this reference does not expressly or inherently disclose this limitation, it would have been obvious in view of the background knowledge of a person of ordinary skill in the art. This reference further renders the limitation obvious in light of other references, such as Sapuntzakis, Birse, and/or Kulkarni. Sapuntzakis, Birse, and Kulkarni concern similar subject matter as this reference, namely, distributed computing systems. As set forth in the accompanying contentions, storing data blocks not previously contained in a root image and/or changes made by compute nodes to the blocks of the root image was a well-known and commonly-used technique. Modifying this reference with Sapuntzakis, Birse, and/or Kulkarni’s teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. A person of ordinary skill in the art would have further understood how to modify the disclosures of this reference with those of Sapuntzakis, Birse, or Kulkarni for example, to incorporate this limitation. See, e.g., “We can store the disks in these capsules efficiently by creating a hierarchy, where each child capsule could be viewed as inheriting from the parent capsule with the differences in disk state between parent and child captured in a separate copy-on-write (COW) virtual disk.” Sapuntzakis at 5. “Writes to a capsule disk are performed by writing the data to the latest COW disk and updating its bitmap file. Reads involve searching the latest COW disk and its ancestor

Exhibit D-24

Claims	Apple NetBoot
	disks in turn until the required block is found. Since the root COW disk contains a copy of all the blocks, the search is guaranteed to terminate.” Sapuntzakis at 5–6. “Since the latest COW disk is always local, all writes are local.” Sapuntzakis at 6. The diagram illustrates the structure of a capsule disk. It consists of a Bitmap, a set of DiskBlocks, and a Flat Disk. The Bitmap is a sequence of bits (11111111, 10010101, 11000001, 01011000) representing the presence of blocks in the COW disks. The DiskBlocks are a set of blocks (represented by shaded and unshaded squares) that are mapped to the Flat Disk. The Flat Disk is a sequence of blocks (represented by squares) that are seen by the capsule. A legend indicates that a shaded square represents a block present in a COW disk. Figure 2: An example capsule disk. Sapuntzakis at 6. <i>See also</i> Sapuntzakis at 1-12. “The block device driver 154 services file system read and write requests to the system volumes 160. In one embodiment, the block device driver 154 appears to the file system 152 to be a standard hard drive device driver. However, in reality, the block device driver 154 may be configured to service the read and write requests of the file system 152 by accessing the NC server 170. As will be discussed further below, in one embodiment, one or more shadow volumes, may be created for the NC client 150 and stored with the non-persistent client-specific data 184 for the purpose of storing user preferences and/or changes to the operating system. Therefore, some mechanism is needed to direct reads

Exhibit D-24

Claims	Apple NetBoot
	and writes to the appropriate volume, e.g., the shadow volumes or the system volumes 174. According to one embodiment, the block device driver 154 has knowledge of the separate shadow volumes and system volumes 174 and redirects read and write requests to the appropriate volume. Alternatively, server-side logic may perform the redirection locally.” Birse at 7:35-53; <i>see also</i> Birse at 11:30-32. “In the embodiment depicted, it is assumed that the shadow volume contains only those portions of the operating system that have been modified by the user rather than a complete copy of the operating system as modified. The various options regarding the granularity of the portions written to the shadow volume are discussed below.” Birse 11:49-52. <i>See also</i> Birse at 10:63-65, 11:6-17, 12:54–59, 14:49-61, FIGS. 4, 6, 10. “FIG. 2 is a flow diagram outlining the acts performed by the server 102 to create the differences file 120 from the old image 114 and the new image 116. In the present example, the updated image is an updated device operating system for the PDA 104.” Kulkarni at [0036]. “Creating the differences file 120 is accomplished by applying a compression algorithm to the new image 116, using the old image 114 as a reference. Since the old image 114 is already stored on the PDA 104, data from the new image 116 that is also contained in the old image 114 does not need to be downloaded from the server 102 to the PDA 104. When the new image 116 is created at the PDA 104, data that is contained in both images is copies from the image 139 already on the PDA 104. Data that is only contained in the new image 116 is downloaded to the PDA 104 in the modified differences file 124 and is copied to the flash memory image 139.” Kulkarni at [0037]. “The differences file 120 and the modified difference file 124 are composed of two alternating types of sections: data sections contain data from the new image 116 that is not found in the old image 114; copy sections contain data that is found in both the new image 116 and the old image 114. To distinguish between the two types of sections, a

Exhibit D-24

Claims	Apple NetBoot
	header is placed at the beginning of each section, the header giving specific information about the section that follows the header.” Kulkarni at [0038]. <i>See also</i> Kulkarni at [0026]-[0041], [0068], FIG. 1, FIGS. 2a-2d, FIG. 6. To the extent that Apple NetBoot does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple NetBoot, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[7.3] caching blocks of said root image that have been accessed by at least one of said compute nodes in a cache memory.	Apple NetBoot discloses or renders obvious caching blocks of said root image that have been accessed by at least one of said compute nodes in a cache memory. For example: “Perhaps the number one variable that can have a positive impact on performance is increasing the physical RAM in each client. This cuts down on the need to “page” (the technique which implements virtual memory) and can provide crisper performance by allowing more caching in memory. For best results, it is highly recommended that each client has a least 64 MB of physical RAM, and that the cache size be increased to at least 2 MB.” <i>See</i> https://ftpmirror.your.org/pub/misc/apple/Applesoftwareupdates/english-north-american/mac-os-x-server/netboot/desktop-admin/inside-netbooting.pdf. To the extent this reference does not expressly or inherently disclose this limitation, it would have been obvious in view of the background knowledge of a person of ordinary skill in the art. This reference further renders the limitation obvious in light of other references, such as Holzmann, Nelson, and/or Stahl. Holzmann, Nelson, and Stahl concern similar subject matter as this reference, namely, distributed computing systems. As set forth in the accompanying contentions, caching data from comparatively slow storage devices in faster devices (including solid-state memory) was a well-known and commonly-used technique. Modifying this reference with Holzmann’s, Nelson’s, and/or Stahl’s caching teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. A person of ordinary skill in the art would have further understood how to modify the disclosures of this reference with those of Holzmann, Nelson, or Stahl for example, to incorporate this limitation. <i>See, e.g.,</i>

Exhibit D-24

Claims	Apple NetBoot
	“The invention is directed generally to computer storage systems. More specifically, the invention is directed to a system and method for booting multiple servers using a single operating system image located on a solid state disk in a storage area network.” Holzmann at [0002]. <pre> graph TD 107[INTERNAL POWER SOURCE] --- 104[MEMORY] 107 --- 103[CONTROL MODULE] 107 --- 105[INTERFACE MODULE] 104 <--> Data 103 103 <--> Data 102[NON-VOLATILE STORAGE MEANS] 103 <--> Data 105 105 <--> Data 106[EXTERNAL DEVICES] subgraph 101 [Solid State Disk System] 104 103 105 102 end </pre> Fig. 1 Holzmann at FIG. 1 “Referring now to the figures, the present invention is directed to a system and method for booting multiple servers or other external devices from a storage area network using a single OS image. FIG. 1 is a block diagram illustrating the component parts of the invented system and the functions of each component part. The invented system comprises a solid state disk system 101 having a storage means 102, a control module 103, a memory module 104, an interface module 105 that communicates with external devices 106, and an internal power source 107.” Holzmann at [0017].

Exhibit D-24

Claims	Apple NetBoot
	“The memory module 104 comprises at least one direct-access memory module for holding data currently or recently used by external devices 106. The memory module 104 is more quickly accessible, and performs read and write processing functions more quickly, than non-volatile or disk memory, such as storage means 102. The memory module 104 preferably comprises at least one random-access memory (RAM) module.” Holzmann at [0020]. “This invention relates in general to the fields of Computer Operating Systems, Multi-processor Hardware Systems, Object Oriented Programming, and Virtual Memory Systems. In particular, the invention relates to improved techniques for establishing and efficiently handling relationships between a client program, a file system server and a virtual memory manager (‘VMM’).” Nelson at [0001]. “This servicing of multiple client programs is now illustrated by means of Figure 6. A first client program 102 asks the CFS 218 to cache remote file 200. The CFS 218 provides a first CFS file object 108 for the first client 102 to use in processing requests on the file. The CFS 218 establishes a communications link 244 to the file server 206 located on the remote node 204 to obtain the file attributes which the CFS 218 caches in the CFS attribute cache 239. The CFS 218 also establishes the communications between the CFS file object manager 221 and the VMM 230 which caches the file data in the data cache 231 by means of the pager_object 234/fs_cache_object 232 pipeline. This pipeline results from the map and bind operations previously described. Pager object 234 is connected to fs_cache object 232 in the pager object manager 247, and also connected via the link between the fs_cache object 232 and the cache object 223 in the fs_cache object manager 243 in the CFS 218. When a second client program 118 asks the CFS 218 to cache the same remote file 200, the CFS 218 implements a second CFS file object 122. However, since the CFS 218 has already set up a CFS attribute cache 239 to handle file attributes for the remote file 200, the second CFS file object 122 is merely directed to the existing CFS attribute cache 239. Similarly, when the second CFS file object 122 wishes to read or write data to the file 200, the file object manager 221 again invokes an operation to map the file in the CFS

Exhibit D-24

Claims	Apple NetBoot
	address_space using its address_space object 225. Once again this causes the VMM's address_space object manager 227 to invoke a bind () operation on the file object 122 which is again directed to the CFS 218. This time the CFS 218 checks its cache of bind results 237 to see if that file 200 is already being cached by the VMM 230. When the CFS 218 finds that it is already being cached, the CFS 218 returns a pointer to the fs_cache_object/pager_object pair being used, and the bind operation is completed. Thus, duplicate caches are not established and the redundant network communications overhead is avoided by the process.” Nelson at [0038]-[0039]. “In an alternate embodiment, if the first client program 102 had set up the first CFS file object 108 with the CFS 218, the first client program 102 may pass the first CFS file object 108 to the second client program 118. In that case, both client programs 102 & 118 would use the same CFS file object 108, and as before both client programs would be serviced by the CFS 218 with only one set of caches.” Nelson at [0040].

Exhibit D-24

Claims	Apple NetBoot
	Nelson at FIG. 6. See also Nelson at [0010], [0034]-[0037], [0052]-[0119], FIG. 4. “Repository FS Implementation 222 is an implementation of a file system responsible for maintaining data on the persistent read-only file system 226 that is to have its data shared. Repository FS Implementation 222 receives and responds to file system requests from MapFS Implementation 212. Repository FS Instance in memory 224 represents a transient in-memory status of a particular Repository file system in persistent storage. Status includes, for example, information about what files exist (or do not) and have been

Exhibit D-24

Claims	Apple NetBoot
	successfully or unsuccessfully looked up, file size, permission data, actual cached data blocks, etc.” Stahl at [0028]. “The mapping module 502 looks 308 in the View FS instance 208 (via file system communication module 506) to see whether the file is cached. Preferably, the view caches filenames in a manner similar to MapFS—that is, in a table of names and look-up results. Again, since this is the first time the file is being looked up, it will not exist in the View's cache either, unless it has been looked up through some non-MapFS process. Thus, the mapping module 502 asks the 310 View FS implementation 206 to look up the name in the View in a manner preferably similar to the way in which the VFS layer examines MapFS instance data described above. The View FS Implementation 206 looks up 312 the file in the persistent view FS instance on disk 210 and updates 314 the View FS instance in memory 208 by populating an in-memory description of the file, and inserting an entry in its name-to-inode table. The View FS implementation 206 returns 316 a reference to the file in the View FS Instance 208 to the mapping module 502.” Stahl at [0037]. <i>See also</i> Stahl at paragraphs [0017], [0044]. To the extent that Apple NetBoot does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple NetBoot, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[11.0] The method as recited in claim 7 further comprising: merging the blocks of said root image with the blocks of respective leaf images to create cohesive respective application environments.	Apple NetBoot discloses or renders obvious the method as recited in claim 7 further comprising merging the blocks of said root image with the blocks of respective leaf images to create cohesive respective application environments. For example: “When the user logs out, changes to preferences that are being saved are stored in a file in the Macintosh Manager Items folder located in the Macintosh Manager sharepoint.

Exhibit D-24

Claims	Apple NetBoot
	The next time the user logs in, his or her saved preferences are retrieved from the Macintosh Manager Items folder. Using the desktop pattern example discussed previously, information about the user's desktop pattern is saved when the user logs out and restored when the user logs in. The shared Mac OS image will not be modified. The shadow volume will contain all of the changes made during this procedure. The environment the user sees is a combination of the shared Mac OS image and the changes to it that are stored in the shadow volume.” See https://ftpmirror.your.org/pub/misc/apple/Applesoftwareupdates/english-north-american/mac-os-x-server/netboot/desktop-admin/inside-netbooting.pdf. Further, the evidence cited for claim limitations [7.pre]-[7.3] apply to this claim limitation. To the extent this reference does not expressly or inherently disclose this limitation, it would have been obvious in view of the background knowledge of a person of ordinary skill in the art. This reference further renders the limitation obvious in light of other references, such as Sapuntzakis, Birse, Stahl, and/or Kulkarni. Sapuntzakis, Birse, Stahl, and Kulkarni concern similar subject matter as this reference, namely, distributed computing systems. As set forth in the accompanying contentions, merging data blocks for multiple computer nodes in distributed environments was a well-known and commonly-used technique. Modifying this reference with Sapuntzakis, Birse, Stahl, and/or Kulkarni’s teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. A person of ordinary skill in the art would have further understood how to modify the disclosures of this reference with those of Sapuntzakis, Birse, Stahl, and/or Kulkarni for example, to incorporate this limitation. See, e.g., “To reduce the start-up time of a capsule, the COW disks corresponding to a capsule disk are read page-by- page on demand, rather than being pre-fetched.” Sapuntzakis at 6. “Reads involve searching the latest COW disk and its ancestor disks in turn until the required block is found.” Sapuntzakis at 5-6.

Exhibit D-24

Claims	Apple NetBoot
	“Note that the COW disk hierarchy is not visible to the VMM, or to the OS and applications inside the capsule; all of them see a normal flat disk as illustrated in the figure.” Sapuntzakis at 6. <i>See also</i> Sapuntzakis at 1-12. “In one embodiment, the shadow system volume 930 is not visible to the user. Consequently, the user system volume 924 will appear to be the location where all writes are directed. For example, if a user modifies a preference, that preference will appear to be written to the preference file located in the preferences folder on the user system volume 924.” Birse at 14:61-67. “When an NC client boots from the network and accesses the operating system from an NC server, the user's preferences are dynamically merged with the system environment provided to the NC client. Advantageously, since, the user's desktop preferences and other customized settings are all preserved from session to session and supplied to the NC client as it boots from the network, the user may login to any NC client on the network and have the same user experience.” Birse at 9:2-10. “Referring now to FIG. 5, NC client read request processing will now be described according to one embodiment of the present invention. ... At step 510, the file system 152 generates a read request directed to the system volumes 160. At step 520, the read request is received by the block device driver 154 and a determination is made whether the read request specifies a portion of the operating system that has been modified by the user or whether the read request specifies a portion of the operating system that remains unchanged. If the read request corresponds to a portion of the operating system that has not been modified by the user, the processing continues with step 530. However, if the read request corresponds to a portion of the operating system that has been previously modified by the user, then processing continues with step 540. At step 530, the block device driver directs the NC server 170 to retrieve the information associated with the read request from the system volumes 174. At step 540, the block device driver directs the NC server 170 to retrieve the information associated with the read request from the

Exhibit D-24

Claims	Apple NetBoot
	user's shadow volume. In this manner, the file system 152 need not be aware of the particular mechanism used to coordinate user modifications with the original operating system as contained in the system volumes 174." Birse at 11:39-41, 11:55-12:10. "The implementations described herein create and utilize a differences file to create a new image in flash memory of an electronic device from sections of an old image in the flash memory and sections of a new image, the new image being an updated version of the old image. The new image is created incrementally in flash memory so that an entire image does not have to first be constructed in device RAM, thereby saving significant overhead. A "flash manager" is described which allows the new image to be built incrementally until a memory block is created that is of sufficient size to be written to the flash memory." Kulkarni at [0013]. "Repository FS Implementation 222 is an implementation of a file system responsible for maintaining data on the persistent read-only file system 226 that is to have its data shared. Repository FS Implementation 222 receives and responds to file system requests from MapFS Implementation 212. Repository FS Instance in memory 224 represents a transient in-memory status of a particular Repository file system in persistent storage. Status includes, for example, information about what files exist (or do not) and have been successfully or unsuccessfully looked up, file size, permission data, actual cached data blocks, etc." Stahl at [0028]. "Overlay FS Implementation 214 is an implementation of a file system responsible for providing writeable storage on the persistent overlay file system instance 220. Overlay FS Implementation 214 receives and responds to file system requests—including write requests—from MapFS Implementation 212. Overlay FS Instance in memory 218 behaves analogously to Repository FS Instance 224. In a preferred embodiment, the Overlay FS is a directory tree rooted at a location specified when the MapFS is instantiated." Stahl at [0029]. "The flash manager 136 uses the modified differences file 124' and the image 139 in flash memory 132 (the old image) to start building the new image. At step 404, a portion of the

Exhibit D-24

Claims	Apple NetBoot
	new image is constructed and temporarily stored in the new memory block 138. When the new memory block 138 contains a portion of the new image that constitutes at least one-half of a write block size for the flash memory 132 (“Yes” branch, step 406), then the portion of the new image stored in the new memory block 138 is added to the image 139 in the flash memory 132. As previously discussed, the size of a write block depends on the type of flash memory, i.e., the brand of the flash memory. It is common for flash memory to have a write block of 64 Kb. Also as previously discussed, the flash manager 136 may manage several new memory blocks 138. When any of the new memory blocks 138 reaches a size of one-half of a write block size, that new memory block 138 is written to the flash memory.” Kulkarni at [0069]. <i>See also</i> Kulkarni at [0026]-[0041], [0068], FIG. 1, FIGS. 2a-2d, FIG. 6. To the extent that Apple NetBoot does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple NetBoot, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.

Exhibit D-25

Invalidity of U.S. Patent No. 8,332,844 in View of VMware Systems

VMware Systems was known in this country, used in public, sold or offered for sale no later than 2001, as demonstrated by at least the materials cited herein or the testimony of knowledgeable witnesses and corroborating documents.¹ Specifically, based on information currently available, VMware Systems was (1) known or used in this country before the alleged invention of the claimed subject matter of the asserted claims, (2) in public use or on sale in this country more than one year before the filing date of the patent, or (3) invented by another who did not abandon, suppress, or conceal, before the alleged invention of the claimed subject matter of the asserted claim. Thus, VMware Systems is prior art to the '844 Patent under at least Pre-AIA 35 U.S.C. §§ 102(a) or (g) as to claims 7 and 11 (the "Asserted Claims") of the '844 Patent, as discussed in detail in the chart that follows.

At least the following documents describe the relevant functionality disclosed by VMware Systems:

- VMware GSX Server 3 Virtual Machine Guide;
- VMware GSX Server 3 Specifications Datasheet; and
- VMware ESX Server 2.1 Specifications Datasheet.

Based on information currently available, the persons or entities involved in the conception and diligent reduction to practice of the ideas in VMware Systems include at least persons involved in the VMware Systems and their testing and commercial use, including at least VMware. The conception of VMware Systems occurred at least as early as 2001.

American offers this invalidity chart based on its current understanding of VMware Systems. American's investigation of this system is ongoing. American may rely on materials produced by Plaintiffs or by third parties regarding VMware Systems and related software to demonstrate invalidity of the asserted claims. American has not had an opportunity to review source code for VMware Systems. American may rely on source code, when available, implementing the functionality described herein, or other source code implementing substantially similar functionality, to demonstrate the invalidity of the asserted claims.

To the extent Plaintiff alleges VMware Systems does not disclose any particular limitation of the Asserted Claims of the '844 Patent, either expressly or inherently, any purported differences are such that the claimed subject matter as a whole would have been

¹ Defendant's invalidity charts, in some instances, rely at least in part on Plaintiff's apparent positions regarding the scope of its claims for purposes of asserting infringement. Nothing in these claim charts should be understood as an admission relating to infringement, either literally or under the doctrine of equivalents, or as an admission relating to Defendant's understanding of the proper interpretation or scope of the Asserted Claims. Defendant reserves the right to rely on additional citations or sources of evidence that may also be applicable, or that may become applicable in light of any Court Order on claim construction, changes in Plaintiff's infringement contentions, or information obtained during discovery as the case proceeds.

Exhibit D-25

obvious in view of the knowledge of one skilled in the art. It would have further been obvious to a person of ordinary skill in the art as of the priority date of the '844 Patent to modify VMware Systems or combine the teachings of VMware Systems with other prior art in a manner that would have rendered the Asserted Claims invalid as obvious, including but not limited to:

- U.S. Patent Application Publication No. 2005/0283597 (“Holzmann”);
- EP Patent No. EP0604013 (“Nelson”);
- “Optimizing the Migration of Virtual Computers” (“Sapuntzakis”);
- U.S. Patent No. 7,089,300 (“Birse”);
- U.S. Patent Application Publication No. 2005/0004925 (“Stahl”); and/or
- U.S. Patent Application Publication No. 2002/0034105 (“Kulkarni”).

As the United States Supreme Court held in *KSR Int’l Co. v. Teleflex, Inc.*, “[t]he combination of familiar elements according to known methods is likely to be obvious when it does no more than yield predictable results.” 550 U.S. 398, 416 (2007). The Supreme Court further held: “[w]hen a work is available in one field of endeavor, design incentives and other market forces can prompt variations of it, either in the same field or a different one. If a person of ordinary skill can implement a predictable variation, § 103 likely bars its patentability. For the same reason, if a technique has been used to improve one device, and a person of ordinary skill in the art would recognize that it would improve similar devices in the same way, using the technique is obvious unless its actual application is beyond his or her skill.” *Id.* at 417. The Supreme Court further held: “[w]hen there is a design need or market pressure to solve a problem and there are a finite number of identified, predictable solutions, a person of ordinary skill has good reason to pursue the known options within his or her technical grasp. If this leads to the anticipated success, it is likely the product not of innovation but of ordinary skill and common sense.” *Id.* at 421.

Defendant reserves the right to amend or supplement this claim chart at a later date as more fully set forth in the Invalidity Contentions.

Exhibit D-25

Claims	VMware Systems
[7.pre] A method for providing data to a plurality of compute nodes, comprising:	To the extent the preamble is limiting, VMware Systems discloses or renders obvious a method for providing data to a plurality of compute nodes. For example: VMware GSX Server was introduced in January 2001, and VMware ESX Server was introduced in March 2001. These systems enable multiple virtual machines (compute nodes) to run concurrently on a single physical server. The virtual machines share data through root image files stored on centralized storage units and manage data changes via snapshots and redo logs. The systems efficiently provide data to a plurality of compute nodes through virtualization and dynamic resource allocation. To the extent that VMware Systems does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of VMware Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant's Invalidity Contentions.
[7.1] storing blocks of a root image of said compute nodes on a first storage unit;	VMware Systems discloses or renders obvious storing blocks of a root image of said compute nodes on a first storage unit. For example: The VMware GSX Server stores virtual machine configurations (including root images) in a dedicated directory, consistent with storing root image blocks on a first storage unit. "The virtual machine typically is stored on the host computer in a set of files, all of which are in a directory set aside for that particular virtual machine." <i>See</i> VMware GSX Server 3 Virtual Machine Guide. VMware GSX Server stores virtual machine data, including root images, as encapsulated virtual disk files on a storage unit. "Virtual disk files encapsulate all virtual machine data and can run on any system with VMware

Exhibit D-25

Claims	VMware Systems
	GSX Server, ESX Server, or Workstation installed.” <i>See</i> VMware GSX Server 3 Specifications Datasheet. VMware ESX Server stores root images as isolated entities within a physical system, effectively using a storage unit for each VM’s root image. “VMware ESX Server transforms physical systems into a pool of logical computing resources... Operating systems and applications are isolated in multiple virtual machines that reside on a single piece of hardware.” <i>See</i> VMware ESX Server 2.1 Specifications Datasheet. To the extent that VMware Systems does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of VMware Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[7.2] storing leaf images for respective compute nodes on respective second storage units, said leaf images including only additional data blocks not previously contained in said root image and changes made by respective compute nodes to the blocks of the root image, wherein said leaf images of respective compute nodes do not include blocks of said root image that are unchanged by respective compute nodes; and	VMware Systems discloses or renders obvious storing leaf images for respective compute nodes on respective second storage units, said leaf images including only additional data blocks not previously contained in said root image and changes made by respective compute nodes to the blocks of the root image, wherein said leaf images of respective compute nodes do not include blocks of said root image that are unchanged by respective compute nodes. For example: After a snapshot is taken, changes are stored in redo-log files, which act as leaf images containing only the modified data blocks. “Once you take a snapshot, the virtual machine starts saving any changes to the virtual machine to one or more redo-log files.” <i>See</i> VMware GSX Server 3 Virtual Machine Guide.

Exhibit D-25

Claims	VMware Systems
	The snapshot mechanism captures only the changes made during runtime, storing modified data separately from the unchanged root image. “Snapshots save a point-in-time copy of the running virtual machine state.” <i>See</i> VMware ESX Server 2.1 Specifications Datasheet. To the extent this reference does not expressly or inherently disclose this limitation, it would have been obvious in view of the background knowledge of a person of ordinary skill in the art. This reference further renders the limitation obvious in light of other references, such as Sapuntzakis, Birse, and/or Kulkarni. Sapuntzakis, Birse, and Kulkarni concern similar subject matter as this reference, namely, distributed computing systems. As set forth in the accompanying contentions, storing data blocks not previously contained in a root image and/or changes made by compute nodes to the blocks of the root image was a well-known and commonly-used technique. Modifying this reference with Sapuntzakis, Birse, and/or Kulkarni’s teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. A person of ordinary skill in the art would have further understood how to modify the disclosures of this reference with those of Sapuntzakis, Birse, or Kulkarni for example, to incorporate this limitation. <i>See, e.g.,</i> “We can store the disks in these capsules efficiently by creating a hierarchy, where each child capsule could be viewed as inheriting from the parent capsule with the differences in disk state between parent and child captured in a separate copy-on-write (COW) virtual disk.” Sapuntzakis at 5. “Writes to a capsule disk are performed by writing the data to the latest COW disk and updating its bitmap file. Reads involve searching the latest COW disk and its ancestor disks in turn until the required block is found. Since the root COW disk contains a copy of all the blocks, the search is guaranteed to terminate.” Sapuntzakis at 5–6.

Exhibit D-25

Claims	VMware Systems
	“Since the latest COW disk is always local, all writes are local.” Sapuntzakis at 6. The diagram illustrates a capsule disk structure. At the top, a 'Bitmap' row shows binary values: 11111111, 10010101, 11000001, and 01011000. Below these are 'DiskBlocks' represented as rows of shaded and unshaded squares. The rows are labeled: 'Root COW Disk', 'COW Disk 1', 'COW Disk 2', and 'Latest COW Disk'. Below these is a 'Flat Disk (as seen by the capsule)' row. Dashed arrows point from the DiskBlocks of the COW disks to the Flat Disk. A legend at the bottom indicates that a shaded square represents a 'Block present in COW disk'. Figure 2: An example capsule disk. Sapuntzakis at 6. <i>See also</i> Sapuntzakis at 1-12. “The block device driver 154 services file system read and write requests to the system volumes 160. In one embodiment, the block device driver 154 appears to the file system 152 to be a standard hard drive device driver. However, in reality, the block device driver 154 may be configured to service the read and write requests of the file system 152 by accessing the NC server 170. As will be discussed further below, in one embodiment, one or more shadow volumes, may be created for the NC client 150 and stored with the non-persistent client-specific data 184 for the purpose of storing user preferences and/or changes to the operating system. Therefore, some mechanism is needed to direct reads and writes to the appropriate volume, e.g., the shadow volumes or the system volumes 174. According to one embodiment, the block device driver 154 has knowledge of the separate shadow

Exhibit D-25

Claims	VMware Systems
	volumes and system volumes 174 and redirects read and write requests to the appropriate volume. Alternatively, server-side logic may perform the redirection locally.” Birse at 7:35-53; <i>see also</i> Birse at 11:30-32. “In the embodiment depicted, it is assumed that the shadow volume contains only those portions of the operating system that have been modified by the user rather than a complete copy of the operating system as modified. The various options regarding the granularity of the portions written to the shadow volume are discussed below.” Birse 11:49-52. <i>See also</i> Birse at 10:63-65, 11:6-17, 12:54-59, 14:49-61, FIGS. 4, 6, 10. “FIG. 2 is a flow diagram outlining the acts performed by the server 102 to create the differences file 120 from the old image 114 and the new image 116. In the present example, the updated image is an updated device operating system for the PDA 104.” Kulkarni at [0036]. “Creating the differences file 120 is accomplished by applying a compression algorithm to the new image 116, using the old image 114 as a reference. Since the old image 114 is already stored on the PDA 104, data from the new image 116 that is also contained in the old image 114 does not need to be downloaded from the server 102 to the PDA 104. When the new image 116 is created at the PDA 104, data that is contained in both images is copied from the image 139 already on the PDA 104. Data that is only contained in the new image 116 is downloaded to the PDA 104 in the modified differences file 124 and is copied to the flash memory image 139.” Kulkarni at [0037]. “The differences file 120 and the modified difference file 124 are composed of two alternating types of sections: data sections contain data from the new image 116 that is not found in the old image 114; copy sections contain data that is found in both the new image 116 and the old image 114. To distinguish between the two types of

Exhibit D-25

Claims	VMware Systems
	sections, a header is placed at the beginning of each section, the header giving specific information about the section that follows the header.” Kulkarni at [0038]. <i>See also</i> Kulkarni at [0026]-[0041], [0068], FIG. 1, FIGS. 2a-2d, FIG. 6. To the extent that VMware Systems does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of VMware Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[7.3] caching blocks of said root image that have been accessed by at least one of said compute nodes in a cache memory.	To the extent this reference does not expressly or inherently disclose this limitation, it would have been obvious in view of the background knowledge of a person of ordinary skill in the art. This reference further renders the limitation obvious in light of other references, such as Holzmann, Nelson, and/or Stahl. Holzmann, Nelson, and Stahl concern similar subject matter as this reference, namely, distributed computing systems. As set forth in the accompanying contentions, caching data from comparatively slow storage devices in faster devices (including solid-state memory) was a well-known and commonly-used technique. Modifying this reference with Holzmann’s, Nelson’s, and/or Stahl’s caching teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. A person of ordinary skill in the art would have further understood how to modify the disclosures of this reference with those of Holzmann, Nelson, or Stahl for example, to incorporate this limitation. <i>See, e.g.,</i> “The invention is directed generally to computer storage systems. More specifically, the invention is directed to a system and method for booting multiple servers using a single operating system image located on a solid state disk in a storage area network.” Holzmann at [0002].

Exhibit D-25

Claims	VMware Systems
	<pre> graph TD 107[INTERNAL POWER SOURCE] --- 104((MEMORY)) 107 --- 103[CONTROL MODULE] 107 --- 105[INTERFACE MODULE] 104 <--> Data 103 103 <--> Data 102[NON-VOLATILE STORAGE MEANS] 103 <--> 105 105 <--> 106([EXTERNAL DEVICES]) </pre> Fig. 1 Holzmann at FIG. 1 “Referring now to the figures, the present invention is directed to a system and method for booting multiple servers or other external devices from a storage area network using a single OS image. FIG. 1 is a block diagram illustrating the component parts of the invented system and the functions of each component part. The invented system comprises a solid state disk system 101 having a storage means 102, a control module 103, a memory module 104, an interface module 105 that communicates with external devices 106, and an internal power source 107.” Holzmann at [0017]. “The memory module 104 comprises at least one direct-access memory module for holding data currently or recently used by external devices 106. The memory module 104 is more quickly accessible, and performs read and write processing functions more quickly, than non-volatile or disk memory, such as storage means 102. The memory module 104 preferably comprises at least one random-access memory (RAM) module.” Holzmann at [0020].

Exhibit D-25

Claims	VMware Systems
	“This invention relates in general to the fields of Computer Operating Systems, Multi-processor Hardware Systems, Object Oriented Programming, and Virtual Memory Systems. In particular, the invention relates to improved techniques for establishing and efficiently handling relationships between a client program, a file system server and a virtual memory manager (‘VMM’).” Nelson at [0001]. “This servicing of multiple client programs is now illustrated by means of Figure 6. A first client program 102 asks the CFS 218 to cache remote file 200. The CFS 218 provides a first CFS file object 108 for the first client 102 to use in processing requests on the file. The CFS 218 establishes a communications link 244 to the file server 206 located on the remote node 204 to obtain the file attributes which the CFS 218 caches in the CFS attribute cache 239. The CFS 218 also establishes the communications between the CFS file object manager 221 and the VMM 230 which caches the file data in the data cache 231 by means of the pager_object 234/fs_cache_object 232 pipeline. This pipeline results from the map and bind operations previously described. Pager object 234 is connected to fs_cache object 232 in the pager object manager 247, and also connected via the link between the fs_cache object 232 and the cache object 223 in the fs_cache object manager 243 in the CFS 218. When a second client program 118 asks the CFS 218 to cache the same remote file 200, the CFS 218 implements a second CFS file object 122. However, since the CFS 218 has already set up a CFS attribute cache 239 to handle file attributes for the remote file 200, the second CFS file object 122 is merely directed to the existing CFS attribute cache 239. Similarly, when the second CFS file object 122 wishes to read or write data to the file 200, the file object manager 221 again invokes an operation to map the file in the CFS address_space using its address_space object 225. Once again this causes the VMM's address_space object manager 227 to invoke a bind () operation on the file object 122 which is again directed to the CFS 218. This time the CFS 218 checks its cache of bind results 237 to see if that file 200 is already being cached by the VMM 230. When the CFS 218 finds that it is already

Exhibit D-25

Claims	VMware Systems
	being cached, the CFS 218 returns a pointer to the fs_cache_object/pager_object pair being used, and the bind operation is completed. Thus, duplicate caches are not established and the redundant network communications overhead is avoided by the process.” Nelson at [0038]-[0039]. “In an alternate embodiment, if the first client program 102 had set up the first CFS file object 108 with the CFS 218, the first client program 102 may pass the first CFS file object 108 to the second client program 118. In that case, both client programs 102 & 118 would use the same CFS file object 108, and as before both client programs would be serviced by the CFS 218 with only one set of caches.” Nelson at [0040].

Exhibit D-25

Claims	VMware Systems
	FIGURE 6 Nelson at FIG. 6. See also Nelson at [0010], [0034]-[0037], [0052]-[0119], FIG. 4. “Repository FS Implementation 222 is an implementation of a file system responsible for maintaining data on the persistent read-only file system 226 that is to have its data shared. Repository FS Implementation 222 receives and responds to file system requests from MapFS Implementation 212. Repository FS Instance in memory 224 represents a transient in-memory status of a particular Repository file system in persistent storage. Status includes, for example, information about what

Exhibit D-25

Claims	VMware Systems
	files exist (or do not) and have been successfully or unsuccessfully looked up, file size, permission data, actual cached data blocks, etc.” Stahl at [0028]. “The mapping module 502 looks 308 in the View FS instance 208 (via file system communication module 506) to see whether the file is cached. Preferably, the view caches filenames in a manner similar to MapFS—that is, in a table of names and look-up results. Again, since this is the first time the file is being looked up, it will not exist in the View's cache either, unless it has been looked up through some non-MapFS process. Thus, the mapping module 502 asks the 310 View FS implementation 206 to look up the name in the View in a manner preferably similar to the way in which the VFS layer examines MapFS instance data described above. The View FS Implementation 206 looks up 312 the file in the persistent view FS instance on disk 210 and updates 314 the View FS instance in memory 208 by populating an in-memory description of the file, and inserting an entry in its name-to-inode table. The View FS implementation 206 returns 316 a reference to the file in the View FS Instance 208 to the mapping module 502.” Stahl at [0037]. <i>See also</i> Stahl at paragraphs [0017], [0044]. To the extent that VMware Systems does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of VMware Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[11.0] The method as recited in claim 7 further comprising: merging the blocks of said root image with the blocks of respective leaf images to create cohesive respective application environments.	The evidence cited for claim limitations [7.pre]-[7.3] apply to this claim limitation. To the extent this reference does not expressly or inherently disclose this limitation, it would have been obvious in view of the background knowledge of a person of ordinary skill in the art. This reference further renders the limitation obvious in light of other references, such as Sapuntzakis, Birse, Stahl, and/or Kulkarni. Sapuntzakis, Birse, Stahl, and Kulkarni concern similar subject matter as this reference, namely, distributed

Exhibit D-25

Claims	VMware Systems
	computing systems. As set forth in the accompanying contentions, merging data blocks for multiple computer nodes in distributed environments was a well-known and commonly-used technique. Modifying this reference with Sapuntzakis, Birse, Stahl, and/or Kulkarni's teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. A person of ordinary skill in the art would have further understood how to modify the disclosures of this reference with those of Sapuntzakis, Birse, Stahl, and/or Kulkarni for example, to incorporate this limitation. <i>See, e.g.,</i> “To reduce the start-up time of a capsule, the COW disks corresponding to a capsule disk are read page-by- page on demand, rather than being pre-fetched.” Sapuntzakis at 6. “Reads involve searching the latest COW disk and its ancestor disks in turn until the required block is found.” Sapuntzakis at 5-6. “Note that the COW disk hierarchy is not visible to the VMM, or to the OS and applications inside the capsule; all of them see a normal flat disk as illustrated in the figure.” Sapuntzakis at 6. <i>See also</i> Sapuntzakis at 1-12. “In one embodiment, the shadow system volume 930 is not visible to the user. Consequently, the user system volume 924 will appear to be the location where all writes are directed. For example, if a user modifies a preference, that preference will appear to be written to the preference file located in the preferences folder on the user system volume 924.” Birse at 14:61-67. “When an NC client boots from the network and accesses the operating system from an NC server, the user's preferences are dynamically merged with the system environment provided to the NC client. Advantageously, since, the user's desktop preferences and other customized settings are all preserved from session to session

Exhibit D-25

Claims	VMware Systems
	and supplied to the NC client as it boots from the network, the user may login to any NC client on the network and have the same user experience.” Birse at 9:2-10. “Referring now to FIG. 5, NC client read request processing will now be described according to one embodiment of the present invention. ... At step 510, the file system 152 generates a read request directed to the system volumes 160. At step 520, the read request is received by the block device driver 154 and a determination is made whether the read request specifies a portion of the operating system that has been modified by the user or whether the read request specifies a portion of the operating system that remains unchanged. If the read request corresponds to a portion of the operating system that has not been modified by the user, the processing continues with step 530. However, if the read request corresponds to a portion of the operating system that has been previously modified by the user, then processing continues with step 540. At step 530, the block device driver directs the NC server 170 to retrieve the information associated with the read request from the system volumes 174. At step 540, the block device driver directs the NC server 170 to retrieve the information associated with the read request from the user’s shadow volume. In this manner, the file system 152 need not be aware of the particular mechanism used to coordinate user modifications with the original operating system as contained in the system volumes 174.” Birse at 11:39-41, 11:55-12:10. “The implementations described herein create and utilize a differences file to create a new image in flash memory of an electronic device from sections of an old image in the flash memory and sections of a new image, the new image being an updated version of the old image. The new image is created incrementally in flash memory so that an entire image does not have to first be constructed in device RAM, thereby saving significant overhead. A “flash manager” is described which allows the new image to be built incrementally until a memory block is created that is of sufficient size to be written to the flash memory.” Kulkarni at [0013]. “Repository FS Implementation 222 is an implementation of a file system responsible for maintaining data on the persistent read-only file system 226 that is

Exhibit D-25

Claims	VMware Systems
	to have its data shared. Repository FS Implementation 222 receives and responds to file system requests from MapFS Implementation 212. Repository FS Instance in memory 224 represents a transient in-memory status of a particular Repository file system in persistent storage. Status includes, for example, information about what files exist (or do not) and have been successfully or unsuccessfully looked up, file size, permission data, actual cached data blocks, etc.” Stahl at [0028]. “Overlay FS Implementation 214 is an implementation of a file system responsible for providing writeable storage on the persistent overlay file system instance 220. Overlay FS Implementation 214 receives and responds to file system requests—including write requests—from MapFS Implementation 212. Overlay FS Instance in memory 218 behaves analogously to Repository FS Instance 224. In a preferred embodiment, the Overlay FS is a directory tree rooted at a location specified when the MapFS is instantiated.” Stahl at [0029]. “The flash manager 136 uses the modified differences file 124’ and the image 139 in flash memory 132 (the old image) to start building the new image. At step 404, a portion of the new image is constructed and temporarily stored in the new memory block 138. When the new memory block 138 contains a portion of the new image that constitutes at least one-half of a write block size for the flash memory 132 (“Yes” branch, step 406), then the portion of the new image stored in the new memory block 138 is added to the image 139 in the flash memory 132. As previously discussed, the size of a write block depends on the type of flash memory, i.e., the brand of the flash memory. It is common for flash memory to have a write block of 64 Kb. Also as previously discussed, the flash manager 136 may manage several new memory blocks 138. When any of the new memory blocks 138 reaches a size of one-half of a write block size, that new memory block 138 is written to the flash memory.” Kulkarni at [0069]. <i>See also</i> Kulkarni at [0026]-[0041], [0068], FIG. 1, FIGS. 2a-2d, FIG. 6. To the extent that VMware Systems does not disclose any elements of this claim

Exhibit D-25

Claims	VMware Systems
	limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of VMware Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant's Invalidity Contentions.

Exhibit D-26

Invalidity of U.S. Patent No. 8,332,844 in View of United States Patent No. 6,618,736 (“Menage”)

Menage titled “Template-Based Creation and Archival of File Systems” to Paul Menage was filed on March 9, 2001 and issued on September 9, 2003.¹ Menage is prior art to the ’844 Patent under at least Pre-AIA 35 U.S.C. §§ 102(a) and (b) as to claims 7 and 11 (the “Asserted Claims”) of the ’844 Patent, as discussed in detail in the chart that follows.

To the extent Plaintiff alleges Menage does not disclose any particular limitation of the Asserted Claims of the ’844 Patent, either expressly or inherently, any purported differences are such that the claimed subject matter as a whole would have been obvious in view of the knowledge of one skilled in the art. It would have further been obvious to a person of ordinary skill in the art as of the priority date of the ’844 Patent to modify Menage or combine the teachings of Menage with other prior art in a manner that would have rendered the Asserted Claims invalid as obvious, including but not limited to:

- U.S. Patent Application Publication No. 2005/0283597 (“Holzmann”);
- EP Patent No. EP0604013 (“Nelson”);
- Optimizing the Migration of Virtual Computers” (“Sapuntzakis”);
- U.S. Patent No. 7,089,300 (“Birse”);
- U.S. Patent Application Publication No. 2005/0004925 (“Stahl”); and/or
- U.S. Patent Application Publication No. 2002/0034105 (“Kulkarni”).

As the United States Supreme Court held in *KSR Int’l Co. v. Teleflex, Inc.*, “[t]he combination of familiar elements according to known methods is likely to be obvious when it does no more than yield predictable results.” 550 U.S. 398, 416 (2007). The Supreme Court further held: “[w]hen a work is available in one field of endeavor, design incentives and other market forces can prompt variations of it, either in the same field or a different one. If a person of ordinary skill can implement a predictable variation, § 103 likely bars its patentability. For the same reason, if a technique has been used to improve one device, and a person of ordinary skill in the art would recognize that it would improve similar devices in the same way, using the technique is obvious unless its actual application is beyond his or her skill.” *Id.* at 417. The Supreme Court further held: “[w]hen there is a design need or market pressure to solve a problem and there are a finite number of identified, predictable solutions, a person of ordinary skill has good reason to

¹ Defendant’s invalidity charts, in some instances, rely at least in part on Plaintiff’s apparent positions regarding the scope of its claims for purposes of asserting infringement. Nothing in these claim charts should be understood as an admission relating to infringement, either literally or under the doctrine of equivalents, or as an admission relating to Defendant’s understanding of the proper interpretation or scope of the Asserted Claims. Defendant reserves the right to rely on additional citations or sources of evidence that may also be applicable, or that may become applicable in light of any Court Order on claim construction, changes in Plaintiff’s infringement contentions, or information obtained during discovery as the case proceeds.

Exhibit D-26

pursue the known options within his or her technical grasp. If this leads to the anticipated success, it is likely the product not of innovation but of ordinary skill and common sense.” *Id.* at 421.

Defendant reserves the right to amend or supplement this claim chart at a later date as more fully set forth in the Invalidity Contentions.

Exhibit D-26

Claims	United States Patent No. 6,618,736 (“Menage”)
[7.pre] A method for providing data to a plurality of compute nodes, comprising:	To the extent the preamble is limiting, Menage discloses or renders obvious a method for providing data to a plurality of compute nodes. For example: “Accordingly, it would be desirable to store the file systems of multiple virtual private servers within the same physical device or comparatively small set of devices.” Menage at 2:3-6; <i>see also</i> Menage at 7:43-46. “The present invention relates to file system creation and archival, particularly in the context of multiple virtual private servers running on a single physical host machine.” Menage at 3:56–58. “Thus, as shown in FIG. 7, the file system 200 of each virtual process 101 may be conceptualized as a combination of the set of shared storage units 302 and the set of private storage units 304 from the corresponding template 310.” Menage at 9:54-57.

Exhibit D-26

Claims	United States Patent No. 6,618,736 (“Menage”)
	FIG. 7 Menage at FIG. 7. To the extent that Menage does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Menage, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[7.1] storing blocks of a root image of said compute nodes on a first storage unit;	Menage discloses or renders obvious storing blocks of a root image of said compute nodes on a first storage unit. For example: “As illustrated in FIG. 3, the foregoing problems are solved by providing a set of shared storage units 302 and a set of private storage units 304. In one embodiment, the set of shared storage units 302 is used to store data that is common to the files systems 200 of multiple virtual processes 101, such as standard application

Exhibit D-26

Claims	United States Patent No. 6,618,736 (“Menage”)
	programs, utilities, databases, etc. The set of shared storage units 302 forms an initial file system 200 for each virtual process 101.” Menage at 8:8-15. The diagram illustrates a storage architecture. A 'Storage Device' (204) contains 'Shared Storage Units' (302) and 'Private Storage Units' (304). The 'Shared Storage Units' (302) are numbered 1, 2, 3, and so on. The 'Private Storage Units' (304) are numbered 1001 (1), 1002 (2), 1003 (3), and so on. A 'Virtual Process #1' (101) is connected to the 'Private Storage Units' (304) and a 'Usage Map' (306). The 'Usage Map' (306) contains entries for 1 (1001), 0 (1002), and 1 (1003), and so on. A bracket (200) groups the 'Shared Storage Units' (302) and 'Private Storage Units' (304). A bracket (310) groups the 'Usage Map' (306). FIG. 3 Menage at FIG. 3. “The data is typically organized into ‘files,’ such as word processing documents, spreadsheets, executable programs, and the like. The files are stored within a plurality of ‘storage units’ of the storage device, sometimes referred to as ‘disk blocks’ or allocation units.” Menage at 7:36-40. Further, the evidence cited for claim limitation [7.pre] applies to this claim limitation.

Exhibit D-26

Claims	United States Patent No. 6,618,736 (“Menage”)
	To the extent that Menage does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Menage, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[7.2] storing leaf images for respective compute nodes on respective second storage units, said leaf images including only additional data blocks not previously contained in said root image and changes made by respective compute nodes to the blocks of the root image, wherein said leaf images of respective compute nodes do not include blocks of said root image that are unchanged by respective compute nodes; and	Menage discloses or renders obvious storing leaf images for respective compute nodes on respective second storage units, said leaf images including only additional data blocks not previously contained in said root image and changes made by respective compute nodes to the blocks of the root image, wherein said leaf images of respective compute nodes do not include blocks of said root image that are unchanged by respective compute nodes. For example: “Thus, as shown in FIG. 7, the file system 200 of each virtual process 101 may be conceptualized as a combination of the set of shared storage units 302 and the set of private storage units 304 from the corresponding template 310.” Menage at 9:54-57.

Exhibit D-26

Claims	United States Patent No. 6,618,736 (“Menage”)
	<pre> graph TD VP1[Virtual Process #1] --> T1[Template #1] VP2[Virtual Process #2] --> T2[Template #2] VP3[Virtual Process #3] --> T3[Template #3] subgraph 200 T1 T2 T3 end T1 --> SSU[Shared Storage Units] T2 --> SSU T3 --> SSU subgraph 302 SSU end </pre> FIG. 7 Menage at FIG. 7. “The set of private storage units 304, on the other hand, are used to store changes to the initial file system 200 that occur subsequent to the creation of the file system 200. Each private storage unit 304 preferably corresponds to one of the shared storage units 302. For example, as shown in FIG. 3, the number in parentheses within each private storage unit 304 refers to the corresponding shared storage unit 302.” Menage at 8:15-21.

Exhibit D-26

Claims	United States Patent No. 6,618,736 (“Menage”)
	The diagram illustrates a storage device architecture. A central box labeled 'Storage Device' (204) contains two main sections: 'Shared Storage Units' (302) and 'Private Storage Units' (304). The 'Shared Storage Units' (302) are represented by a vertical stack of boxes labeled 1, 2, 3, and an ellipsis. The 'Private Storage Units' (304) are also a vertical stack of boxes labeled 1001 (1), 1002 (2), 1003 (3), and an ellipsis. To the right of the 'Private Storage Units' (304) is a box labeled 'Usage Map' (306). The 'Usage Map' (306) contains a vertical stack of boxes labeled 1 (1001), 0 (1002), 1 (1003), and an ellipsis. A dashed box (310) encloses the 'Private Storage Units' (304) and the 'Usage Map' (306). To the right of the dashed box (310) is a box labeled 'Virtual Process #1' (101). Arrows indicate connections: an arrow from the 'Virtual Process #1' (101) to the 'Private Storage Units' (304), and an arrow from the 'Usage Map' (306) to the 'Private Storage Units' (304). A bracket (200) is positioned to the left of the 'Shared Storage Units' (302) and the 'Private Storage Units' (304). A bracket (308) is positioned to the left of the 'Usage Map' (306). FIG. 3 Menage at FIG. 3. “Throughout the following description, a set of private storage units 304 for a virtual process 101 and a corresponding usage map 306 are collectively referred to as a ‘template’ 310. While FIG. 3 illustrates only a single template 310, it should be recognized that at least one template 310 is preferably provided for each virtual process 101. As detailed below, the file system 200 of each virtual process 101 comprises a combination of the shared storage units 302 and the private storage units 304 from one or more templates 310.” Menage at 8:58-67.

Exhibit D-26

Claims	United States Patent No. 6,618,736 (“Menage”)
	“In one implementation, a usage map 306 is also provided for each set of private storage units 304. The usage map 306 includes a plurality of indicators 308 for indicating whether a corresponding private Storage unit 304 contains valid data. As used herein, the phrase “contains valid data” means that the private storage unit 304 stores data representing a change to the initial file system 200.” Menage at 8:28-34. “In one embodiment, the usage map 306 is initially reset or ‘initialized’ to indicate that none of the private storage units 304 contain valid data. In the case of a bitmap embodiment, for example, all of the bits could be reset to 0.” Menage at 8:46-49. “Referring now to FIGS. 4-6, there is shown a system 400 for creating and managing the file systems 200 of a plurality of virtual processes 101 according to an embodiment of the invention. In one implementation, a system call wrapper 111, as described with reference to FIG. 1, intercepts a system call 115 for writing a data item to a shared storage unit 202. Various system calls 115 may be intercepted, such as the UNIX(R) write() system call 115. Of course, the invention is not limited to any particular operating system.” Menage at 9:1-9.

Exhibit D-26

Claims	United States Patent No. 6,618,736 (“Menage”)
	<pre> graph TD VP[Virtual Process 101] -- "Write() data to SU 2 (115)" --> SCW[System Call Wrapper 111] SCW --> SWM[Storage Unit Writing Module 402] SCW --> UMU[Usage Map Updating Module 404] SWM --> PSU[Private Storage Units 304] UMU --> UM[Usage Map 306] subgraph PSU [Private Storage Units 304] direction TB P1[1001] P2["data 1002"] P3[1003] P4[...] end subgraph UM [Usage Map 306] direction TB U1["1 (1001)"] U2["1 (1002)"] U3["1 (1003)"] U4[...] end U1 --> 308 U2 --> 308 U3 --> 308 </pre> FIG. 4 Menage at FIG. 4. “Rather than writing the data item to the shared storage unit 302, however, a storage unit writing module 402 writes the data item to the corresponding private storage unit 304. In various embodiments, the storage unit writing module 402 may be implemented as a component of the system call wrapper 111 or, as depicted, a separate software module executing within either the user address space 103 or the operating system address space 105. Additionally, a usage map updating module 404 stores an indication 308 in the usage

Exhibit D-26

Claims	United States Patent No. 6,618,736 (“Menage”)
	map 306 that the private storage unit 304 (to which the data item is written) contains valid data. In the depicted embodiment, the usage map updating module 404 sets a bit in the usage map 306 that corresponds to the private storage unit 304. In various embodiments, the usage map updating module 404 may be implemented as a component of the system call wrapper 111 or, as depicted, a separate software module executing within either the user address space 103 or the operating system address space 105.” Menage at 9:10-28. “As illustrated in FIG. 5, the system call wrapper 111 is also configured, in one embodiment, to intercept a system call 115 for reading a data item from a shared storage unit 302. Various system calls 115 may be intercepted, such as the UNIX(R) read() system call 115. Of course, the invention is not limited to any particular operating system. After the system call 115 for reading a data item is intercepted, a usage map checking module 502 checks for an indication 308 in the usage map 306 that the corresponding private storage unit 304 contains valid data. If it does, a storage unit reading module 504 reads the data item from the private Storage unit 304 rather than the shared storage unit 302. In various embodiments, the usage map checking module 404 and the storage unit reading module 504 may be implemented as components of the system call wrapper 111 or, as illustrated, as separate software modules executing within either the user address space 103 or the operating system address space 105.” Menage at 9:29-46.

Exhibit D-26

Claims	United States Patent No. 6,618,736 (“Menage”)
	<pre> graph TD VP[Virtual Process 101] -- "115: Read() data from SU 2" --> SCW[System Call Wrapper 111] SCW --> UMC[Usage Map Checking Module 502] UMC <--> UM[Usage Map 306] UMC --> SURM[Storage Unit Reading Module 504] SURM --> PSU[Private Storage Units 304] PSU --> data[data] </pre> FIG. 5 Menage at FIG. 5. “As shown in FIG. 6, if the usage map checking module 502 determines that the usage map 306 does not include an indication 308 of valid data for the corresponding private storage unit 304, then the storage unit reading module 504 reads the data item from the shared storage unit 302 specified in the intercepted system call 115.” Menage at 9:48-53.

Exhibit D-26

Claims	United States Patent No. 6,618,736 ("Menage")
	<pre> graph TD VP[Virtual Process 101] -- "Read() data from SU 2 115" --> SCW[System Call Wrapper 111] SCW --> SUCM[Usage Map Checking Module 502] SCW --> SUM[Storage Unit Reading Module 504] SUCM --> UM[Usage Map 306] SUCM --> SSU[Shared Storage Units 302] UM --> SUM SSU --> SUM SUM --> data[data] </pre> FIG. 6 Menage at FIG. 6. Further, the evidence cited for claim limitations [7.pre]-[7.1] apply to this claim limitation. To the extent this reference does not expressly or inherently disclose this limitation, it would have been obvious in view of the background knowledge of a person of ordinary skill in the art. This reference further renders the limitation obvious in light of other references, such as Sapuntzakis, Birse, and/or Kulkarni. Sapuntzakis, Birse, and Kulkarni concern similar subject matter as this reference, namely, distributed computing systems. As set forth in the accompanying contentions, storing data blocks not previously contained in a root image and/or changes made by compute nodes to the blocks of the root image was a well-known and commonly-used technique. Modifying this reference

Exhibit D-26

Claims	United States Patent No. 6,618,736 (“Menage”)
	with Sapuntzakis, Birse, and/or Kulkarni’s teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. A person of ordinary skill in the art would have further understood how to modify the disclosures of this reference with those of Sapuntzakis, Birse, or Kulkarni for example, to incorporate this limitation. <i>See, e.g.,</i> “We can store the disks in these capsules efficiently by creating a hierarchy, where each child capsule could be viewed as inheriting from the parent capsule with the differences in disk state between parent and child captured in a separate copy-on-write (COW) virtual disk.” Sapuntzakis at 5. “Writes to a capsule disk are performed by writing the data to the latest COW disk and updating its bitmap file. Reads involve searching the latest COW disk and its ancestor disks in turn until the required block is found. Since the root COW disk contains a copy of all the blocks, the search is guaranteed to terminate.” Sapuntzakis at 5–6. “Since the latest COW disk is always local, all writes are local.” Sapuntzakis at 6.

Exhibit D-26

Claims	United States Patent No. 6,618,736 (“Menage”)
	The diagram illustrates a capsule disk structure. It consists of a Bitmap and DiskBlocks. The Bitmap is a sequence of bits: 11111111, 10010101, 11000001, and 01011000. The DiskBlocks are represented by a grid of blocks. The Root COW Disk is the first row of blocks. The COW Disk 1 is the second row. The COW Disk 2 is the third row. The Latest COW Disk is the fourth row. The Flat Disk (as seen by the capsule) is the bottom row of blocks. A legend indicates that a shaded block represents a block present in a COW disk. Figure 2: An example capsule disk. Sapuntzakis at 6. See also Sapuntzakis at 1-12. “The block device driver 154 services file system read and write requests to the system volumes 160. In one embodiment, the block device driver 154 appears to the file system 152 to be a standard hard drive device driver. However, in reality, the block device driver 154 may be configured to service the read and write requests of the file system 152 by accessing the NC server 170. As will be discussed further below, in one embodiment, one or more shadow volumes, may be created for the NC client 150 and stored with the non-persistent client-specific data 184 for the purpose of storing user preferences and/or changes to the operating system. Therefore, some mechanism is needed to direct reads and writes to the appropriate volume, e.g., the shadow volumes or the system volumes 174. According to one embodiment, the block device driver 154 has knowledge of the separate shadow volumes and system volumes 174 and redirects read and write requests to the

Exhibit D-26

Claims	United States Patent No. 6,618,736 (“Menage”)
	appropriate volume. Alternatively, server-side logic may perform the redirection locally.” Birse at 7:35-53; <i>see also</i> Birse at 11:30-32. “In the embodiment depicted, it is assumed that the shadow volume contains only those portions of the operating system that have been modified by the user rather than a complete copy of the operating system as modified. The various options regarding the granularity of the portions written to the shadow volume are discussed below.” Birse 11:49-52. <i>See also</i> Birse at 10:63-65, 11:6-17, 12:54–59, 14:49-61, FIGS. 4, 6, 10. “FIG. 2 is a flow diagram outlining the acts performed by the server 102 to create the differences file 120 from the old image 114 and the new image 116. In the present example, the updated image is an updated device operating system for the PDA 104.” Kulkarni at [0036]. “Creating the differences file 120 is accomplished by applying a compression algorithm to the new image 116, using the old image 114 as a reference. Since the old image 114 is already stored on the PDA 104, data from the new image 116 that is also contained in the old image 114 does not need to be downloaded from the server 102 to the PDA 104. When the new image 116 is created at the PDA 104, data that is contained in both images is copied from the image 139 already on the PDA 104. Data that is only contained in the new image 116 is downloaded to the PDA 104 in the modified differences file 124 and is copied to the flash memory image 139.” Kulkarni at [0037]. “The differences file 120 and the modified difference file 124 are composed of two alternating types of sections: data sections contain data from the new image 116 that is not found in the old image 114; copy sections contain data that is found in both the new image 116 and the old image 114. To distinguish between the two types of sections, a header is placed at the beginning of each section, the header giving specific information about the section that follows the header.” Kulkarni at [0038].

Exhibit D-26

Claims	United States Patent No. 6,618,736 (“Menage”)
	<i>See also</i> Kulkarni at [0026]-[0041], [0068], FIG. 1, FIGS. 2a-2d, FIG. 6. To the extent that Menage does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Menage, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[7.3] caching blocks of said root image that have been accessed by at least one of said compute nodes in a cache memory.	Menage discloses or renders obvious caching blocks of said root image that have been accessed by at least one of said compute nodes in a cache memory. For example: “A template-based approach to file system 200 creation and management is advantageous in a number of respects. For example, there is no need to duplicate an initial file system 200 for multiple virtual processes 101. Moreover, since all the virtual processes 101 use the same set of shared storage units 302, the shared storage units 302 may be easily cached for high speed access in the computer memory 102 or in a memory integrated with the storage device 204.” Menage at 9:66-10:6. Further, the evidence cited for claim limitations [7.pre]-[7.2] apply to this claim limitation. To the extent this reference does not expressly or inherently disclose this limitation, it would have been obvious in view of the background knowledge of a person of ordinary skill in the art. This reference further renders the limitation obvious in light of other references, such as Holzmann, Nelson, and/or Stahl. Holzmann, Nelson, and Stahl concern similar subject matter as this reference, namely, distributed computing systems. As set forth in the accompanying contentions, caching data from comparatively slow storage devices in faster devices (including solid-state memory) was a well-known and commonly-used technique. Modifying this reference with Holzmann’s, Nelson’s, and/or Stahl’s caching teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. A person of ordinary skill in the art would have further understood how to modify the disclosures

Exhibit D-26

Claims	United States Patent No. 6,618,736 (“Menage”)
	of this reference with those of Holzmann, Nelson, or Stahl for example, to incorporate this limitation. <i>See, e.g.,</i> “The invention is directed generally to computer storage systems. More specifically, the invention is directed to a system and method for booting multiple servers using a single operating system image located on a solid state disk in a storage area network.” Holzmann at [0002]. <pre> graph TD 107[INTERNAL POWER SOURCE] --> 104((MEMORY)) 107 --> 103[CONTROL MODULE] 104 <--> Data 103 103 <--> Data 102[NON-VOLATILE STORAGE MEANS] 103 <--> 105[INTERFACE MODULE] 105 <--> 106([EXTERNAL DEVICES]) </pre> Fig. 1 Holzmann at FIG. 1 “Referring now to the figures, the present invention is directed to a system and method for booting multiple servers or other external devices from a storage area network using a single OS image. FIG. 1 is a block diagram illustrating the component parts of the invented system and the functions of each component part. The invented system comprises a solid state disk system 101 having a storage means 102, a control module 103, a memory module 104, an interface module 105 that

Exhibit D-26

Claims	United States Patent No. 6,618,736 (“Menage”)
	communicates with external devices 106, and an internal power source 107.” Holzmann at [0017]. “The memory module 104 comprises at least one direct-access memory module for holding data currently or recently used by external devices 106. The memory module 104 is more quickly accessible, and performs read and write processing functions more quickly, than non-volatile or disk memory, such as storage means 102. The memory module 104 preferably comprises at least one random-access memory (RAM) module.” Holzmann at [0020]. “This invention relates in general to the fields of Computer Operating Systems, Multi-processor Hardware Systems, Object Oriented Programming, and Virtual Memory Systems. In particular, the invention relates to improved techniques for establishing and efficiently handling relationships between a client program, a file system server and a virtual memory manager (‘VMM’).” Nelson at [0001]. “This servicing of multiple client programs is now illustrated by means of Figure 6. A first client program 102 asks the CFS 218 to cache remote file 200. The CFS 218 provides a first CFS file object 108 for the first client 102 to use in processing requests on the file. The CFS 218 establishes a communications link 244 to the file server 206 located on the remote node 204 to obtain the file attributes which the CFS 218 caches in the CFS attribute cache 239. The CFS 218 also establishes the communications between the CFS file object manager 221 and the VMM 230 which caches the file data in the data cache 231 by means of the pager_object 234/fs_cache_object 232 pipeline. This pipeline results from the map and bind operations previously described. Pager object 234 is connected to fs_cache object 232 in the pager object manager 247, and also connected via the link between the fs_cache object 232 and the cache object 223 in the fs_cache object manager 243 in the CFS 218. When a second client program 118 asks the CFS 218 to cache the same remote file 200, the CFS 218 implements a second CFS file object 122. However, since the CFS

Exhibit D-26

Claims	United States Patent No. 6,618,736 (“Menage”)
	218 has already set up a CFS attribute cache 239 to handle file attributes for the remote file 200, the second CFS file object 122 is merely directed to the existing CFS attribute cache 239. Similarly, when the second CFS file object 122 wishes to read or write data to the file 200, the file object manager 221 again invokes an operation to map the file in the CFS address_space using its address_space object 225. Once again this causes the VMM's address_space object manager 227 to invoke a bind () operation on the file object 122 which is again directed to the CFS 218. This time the CFS 218 checks its cache of bind results 237 to see if that file 200 is already being cached by the VMM 230. When the CFS 218 finds that it is already being cached, the CFS 218 returns a pointer to the fs_cache_object/pager_object pair being used, and the bind operation is completed. Thus, duplicate caches are not established and the redundant network communications overhead is avoided by the process.” Nelson at [0038]-[0039]. “In an alternate embodiment, if the first client program 102 had set up the first CFS file object 108 with the CFS 218, the first client program 102 may pass the first CFS file object 108 to the second client program 118. In that case, both client programs 102 & 118 would use the same CFS file object 108, and as before both client programs would be serviced by the CFS 218 with only one set of caches.” Nelson at [0040].

Exhibit D-26

Claims	United States Patent No. 6,618,736 ("Menage")
	Nelson at FIG. 6. See also Nelson at [0010], [0034]-[0037], [0052]-[0119], FIG. 4. “Repository FS Implementation 222 is an implementation of a file system responsible for maintaining data on the persistent read-only file system 226 that is to have its data shared. Repository FS Implementation 222 receives and responds to file system requests from MapFS Implementation 212. Repository FS Instance in memory 224 represents a transient in-memory status of a particular Repository file system in persistent storage. Status includes, for example, information about what

Exhibit D-26

Claims	United States Patent No. 6,618,736 (“Menage”)
	files exist (or do not) and have been successfully or unsuccessfully looked up, file size, permission data, actual cached data blocks, etc.” Stahl at [0028]. “The mapping module 502 looks 308 in the View FS instance 208 (via file system communication module 506) to see whether the file is cached. Preferably, the view caches filenames in a manner similar to MapFS—that is, in a table of names and look-up results. Again, since this is the first time the file is being looked up, it will not exist in the View's cache either, unless it has been looked up through some non-MapFS process. Thus, the mapping module 502 asks the 310 View FS implementation 206 to look up the name in the View in a manner preferably similar to the way in which the VFS layer examines MapFS instance data described above. The View FS Implementation 206 looks up 312 the file in the persistent view FS instance on disk 210 and updates 314 the View FS instance in memory 208 by populating an in-memory description of the file, and inserting an entry in its name-to-inode table. The View FS implementation 206 returns 316 a reference to the file in the View FS Instance 208 to the mapping module 502.” Stahl at [0037]. <i>See also</i> Stahl at paragraphs [0017], [0044]. To the extent that Menage does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Menage, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[11.0] The method as recited in claim 7 further comprising: merging the blocks of said root image with the blocks of respective leaf images to create cohesive respective application environments.	Menage discloses or renders obvious the method as recited in claim 7 further comprising merging the blocks of said root image with the blocks of respective leaf images to create cohesive respective application environments. For example: “As illustrated in FIG. 3, the foregoing problems are solved by providing a set of shared storage units 302 and a set of private storage units 304. In one embodiment, the set of shared storage units 302 is used to store data that is common to the files systems 200 of multiple virtual processes 101, such as standard application

Exhibit D-26

Claims	United States Patent No. 6,618,736 (“Menage”)
	programs, utilities, databases, etc. The set of shared storage units 302 forms an initial file system 200 for each virtual process 101.” Menage at 8:8-15. “The set of private storage units 304, on the other hand, are used to store changes to the initial file system 200 that occur subsequent to the creation of the file system 200. Each private storage unit 304 preferably corresponds to one of the shared storage units 302. For example, as shown in FIG. 3, the number in parentheses within each private storage unit 304 refers to the corresponding shared storage unit 302.” Menage at 8:15-21. <i>See also</i> Menage at cl. 51. Further, the evidence cited for claim limitations [7.pre]-[7.3] apply to this claim limitation. To the extent this reference does not expressly or inherently disclose this limitation, it would have been obvious in view of the background knowledge of a person of ordinary skill in the art. This reference further renders the limitation obvious in light of other references, such as Sapuntzakis, Birse, Stahl, and/or Kulkarni. Sapuntzakis, Birse, Stahl, and Kulkarni concern similar subject matter as this reference, namely, distributed computing systems. As set forth in the accompanying contentions, merging data blocks for multiple computer nodes in distributed environments was a well-known and commonly-used technique. Modifying this reference with Sapuntzakis, Birse, Stahl, and/or Kulkarni’s teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. A person of ordinary skill in the art would have further understood how to modify the disclosures of this reference with those of Sapuntzakis, Birse, Stahl, and/or Kulkarni for example, to incorporate this limitation. <i>See, e.g.,</i> “To reduce the start-up time of a capsule, the COW disks corresponding to a capsule disk are read page-by- page on demand, rather than being pre-fetched.” Sapuntzakis at 6.

Exhibit D-26

Claims	United States Patent No. 6,618,736 (“Menage”)
	“Reads involve searching the latest COW disk and its ancestor disks in turn until the required block is found.” Sapuntzakis at 5-6. “Note that the COW disk hierarchy is not visible to the VMM, or to the OS and applications inside the capsule; all of them see a normal flat disk as illustrated in the figure.” Sapuntzakis at 6. <i>See also</i> Sapuntzakis at 1-12. “In one embodiment, the shadow system volume 930 is not visible to the user. Consequently, the user system volume 924 will appear to be the location where all writes are directed. For example, if a user modifies a preference, that preference will appear to be written to the preference file located in the preferences folder on the user system volume 924.” Birse at 14:61-67. “When an NC client boots from the network and accesses the operating system from an NC server, the user's preferences are dynamically merged with the system environment provided to the NC client. Advantageously, since, the user's desktop preferences and other customized settings are all preserved from session to session and supplied to the NC client as it boots from the network, the user may login to any NC client on the network and have the same user experience.” Birse at 9:2-10. “Referring now to FIG. 5, NC client read request processing will now be described according to one embodiment of the present invention. ... At step 510, the file system 152 generates a read request directed to the system volumes 160. At step 520, the read request is received by the block device driver 154 and a determination is made whether the read request specifies a portion of the operating system that has been modified by the user or whether the read request specifies a portion of the operating system that remains unchanged. If the read request corresponds to a portion of the operating system that has not been modified by the user, the processing continues with step 530. However, if the read request corresponds to a

Exhibit D-26

Claims	United States Patent No. 6,618,736 (“Menage”)
	portion of the operating system that has been previously modified by the user, then processing continues with step 540. At step 530, the block device driver directs the NC server 170 to retrieve the information associated with the read request from the system volumes 174. At step 540, the block device driver directs the NC server 170 to retrieve the information associated with the read request from the user’s shadow volume. In this manner, the file system 152 need not be aware of the particular mechanism used to coordinate user modifications with the original operating system as contained in the system volumes 174.” Birse at 11:39-41, 11:55-12:10. “The implementations described herein create and utilize a differences file to create a new image in flash memory of an electronic device from sections of an old image in the flash memory and sections of a new image, the new image being an updated version of the old image. The new image is created incrementally in flash memory so that an entire image does not have to first be constructed in device RAM, thereby saving significant overhead. A “flash manager” is described which allows the new image to be built incrementally until a memory block is created that is of sufficient size to be written to the flash memory.” Kulkarni at [0013]. “Repository FS Implementation 222 is an implementation of a file system responsible for maintaining data on the persistent read-only file system 226 that is to have its data shared. Repository FS Implementation 222 receives and responds to file system requests from MapFS Implementation 212. Repository FS Instance in memory 224 represents a transient in-memory status of a particular Repository file system in persistent storage. Status includes, for example, information about what files exist (or do not) and have been successfully or unsuccessfully looked up, file size, permission data, actual cached data blocks, etc.” Stahl at [0028]. “Overlay FS Implementation 214 is an implementation of a file system responsible for providing writeable storage on the persistent overlay file system instance 220. Overlay FS Implementation 214 receives and responds to file system requests—including write requests—from MapFS Implementation 212. Overlay FS Instance in memory 218 behaves analogously to Repository FS Instance 224. In a preferred

Exhibit D-26

Claims	United States Patent No. 6,618,736 (“Menage”)
	embodiment, the Overlay FS is a directory tree rooted at a location specified when the MapFS is instantiated.” Stahl at [0029]. “The flash manager 136 uses the modified differences file 124’ and the image 139 in flash memory 132 (the old image) to start building the new image. At step 404, a portion of the new image is constructed and temporarily stored in the new memory block 138. When the new memory block 138 contains a portion of the new image that constitutes at least one-half of a write block size for the flash memory 132 (“Yes” branch, step 406), then the portion of the new image stored in the new memory block 138 is added to the image 139 in the flash memory 132. As previously discussed, the size of a write block depends on the type of flash memory, i.e., the brand of the flash memory. It is common for flash memory to have a write block of 64 Kb. Also as previously discussed, the flash manager 136 may manage several new memory blocks 138. When any of the new memory blocks 138 reaches a size of one-half of a write block size, that new memory block 138 is written to the flash memory.” Kulkarni at [0069]. <i>See also</i> Kulkarni at [0026]-[0041], [0068], FIG. 1, FIGS. 2a-2d, FIG. 6. To the extent that Menage does not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Menage, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.