

**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
Volcano Parallel Database 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.

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
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 California, Los Angeles, June 1989.	June 1989
“Parallel Architectures for Database Systems,” A.R.	1989

Patent/Publication No./Document Name	Date of Filing/Issue/Publication ⁷
Hurson, L.L. Miller, S.H. Pakzad, M.H. Eich, B. Shirazi, Advances in Computers, Vol. 28, pages 107-151, 1989.	
“Parallel Algorithms for Merging and Sorting,” Narsingh Deo, Dilip SarKar, Information Sciences, Vol. 56, pages 151-161, 1991.	1991
“Parallel Algorithms for Ileast 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=gbv&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 by Gang Scheduling,” Uwe Schwiegelshohn and Ramin Yahyapour. Job Scheduling Strategies for	March 30, 1998

Patent/Publication No./Document Name	Date of Filing/Issue/Publication ⁷
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 patent

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. 2003/0065785 (“Jain”)	Filed – September 28, 2001 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. 2002/0154635 (“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 patent

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 System (NFS),” (“Laros”) Retrieved from: https://www.researchgate.net/publication/255223070_Implementing_scalable_disk-less_clusters_using_the_Network_File_System_NFS	Published at least by January 2003
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 patent⁹

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

⁹ 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.

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 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 patent¹⁰

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

¹⁰ 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.

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
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

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
		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 September 2003	Stratos Global Corporation; Inmarsat 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”

Asserted Patent	Asserted Claim(s)	Term(s) and/or Phrase(s) Lacking Written Description and/or Enablement
		• “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: • “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”

Asserted Patent	Asserted Claim(s)	Indefinite Term(s) or Phrase(s)
'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” 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
'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—analogue 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 A-5

Exhibit A-5 to Defendant's Preliminary Invalidity Contentions

Gamma System

Gamma System was known in this country, used in public, sold or offered for sale no later than 1990, 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, Gamma 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, Gamma System anticipates the asserted claims under at least 35 U.S.C. §§ 102(a) or (g).

At least the following documents describe the relevant functionality disclosed by Gamma System:

- <https://apps.dtic.mil/sti/pdfs/ADA223436.pdf>;
- <https://research.cs.wisc.edu/techreports/1990/TR921.pdf>;
- https://minds.wisconsin.edu/bitstream/handle/1793/11036/file_1.pdf;
- <https://www.cs.ubc.ca/~rap/teaching/504/2022W2/slides/parallel.pdf>;
- <https://wiki.epfl.ch/edicpublic/documents/Candidacy%20exam/gamma.pdf>;
- <https://web.eecs.umich.edu/~mozafari/fall2015/eecs584/papers/gamma.pdf>;
- <https://cs.brown.edu/courses/cs227/archives/2012/papers/newsq/p85-dewitt.pdf>;
- <https://minds.wisconsin.edu/bitstream/handle/1793/58708/TR635.pdf?sequence=1>; and
- <https://www.vldb.org/conf/1986/P228.PDF>.

Based on information currently available, the persons or entities involved in the conception and diligent reduction to practice of the ideas in Gamma System include at least persons involved in the Gamma System and their testing and commercial use, including at least University of Wisconsin-Madison. The conception of Gamma System occurred at least as early as 1990.

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 A-5

Defendant offers this invalidity chart based on its current understanding of Gamma System. Defendant's investigation of this system is ongoing. Defendant may rely on materials produced by Plaintiffs or by third parties regarding Gamma System and related software to demonstrate invalidity of the asserted claims. Defendant has not had an opportunity to review source code for Gamma System. Defendant 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 Gamma System does not disclose any particular limitation of the Asserted Claims of the '582 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 '582 Patent to modify Gamma System or combine the teachings of Gamma 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 No. 6,304,866 ("Chow");
- U.S. Patent No. 6,330,583 ("Reiffin");
- U.S. Patent Application Publication No. 2002/0019844 ("Kurowski");
- U.S. Patent Application Publication No. 2002/0091786 ("Yamaguchi");
- U.S. Patent No. 5,765,146 ("Wolf");
- EP Patent Application Publication No. EP0377993 ("Lorie");
- "Analysis of First-Come-First-Served Parallel Job Scheduling," Uwe Schwiegelshohn and Ramin Yahyapour. Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (1998) (<http://dx.doi.org/10.1145/314613.315031>) ("Schwiegelshohn 1");
- "Improving First-Come-First-Serve-Job Scheduling 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 ("Schwiegelshohn 2"); and/or
- U.S. Patent Application Publication No. 2002/0103889 ("Markson").

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:

Exhibit A-5

“[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.


Claims	Gamma System
[1.0] 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:	To the extent the preamble is limiting, Gamma System discloses a method of effecting on a preexisting input file a computer-executable process comprised of a plurality of subtasks. For example: Gamma was a parallel relational database system developed at the University of Wisconsin-Madison in the late 1980s. It embodied many concepts that later influenced systems like MapReduce, particularly in how it handled partitioned parallelism for query execution. https://apps.dtic.mil/sti/pdfs/ADA223436.pdf  Partitioning of R into N logical buckets for Hybrid hash-join. Figure 8 https://research.cs.wisc.edu/techreports/1990/TR921.pdf

Exhibit A-5

Claims	Gamma System
	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 Reiffin, Chow, and/or Kurowski. Reiffin, Chow, and Kurowski concern similar subject matter as this reference, namely, distributed computing systems and parallel processing. As set forth in the accompanying contentions, parallel processing of data and tasks was a well-known and commonly-used technique. Modifying this reference with Reiffin, Chow, and/or Kurowski'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 Reiffin, Chow, or Kurowski for example, to incorporate this limitation. <i>See, e.g.,</i> “This invention relates to computer networks, and more particularly, to a distributed parallel processing network system wherein a large compute-intensive task may be partitioned into subtasks which are then distributed among a plurality of personal computer or workstations for parallel execution of the subtasks in the background concurrently with the execution in the foreground of the respective local tasks of the individual workstations.” Reiffin at 1:7-14. “It is a further object of the invention to provide a system wherein a network of workstations or personal computers may operate in parallel to form a powerful multicomputer system for parallel processing of large time-consuming compute-intensive applications in a fraction of the time that it would take any of the individual computers to do the job alone.” Reiffin at 2:1-7; <i>see also</i> Reiffin at 3:35-6:23, FIGS. 1-5. “Broadly, the present invention concerns the performance of an aggregate database processing task in a multiprocessing computer database system. The aggregate task is executed with multiple concurrently operating task execution units, using incremental and on-demand sub-task allocation. A command is received to perform a machine-executed task divisible into multiple sub-tasks, i.e., an ‘aggregate task’. Then, multiple task execution units

Exhibit A-5

Claims	Gamma System
	are independently operated on-demand to sequentially self-allocate and execute sub-tasks of the aggregate task. Each sub-task is sized according to a specific ‘granularity’, used to divide the aggregate task into sub-tasks. The granularity may be predetermined for all task execution units or repeatedly dynamically determined by each task execution unit at the time of self-allocation. Preferably, the granularity is not too large as to cause potentially uneven sub-task allocation, and not so small as to incur excessive overhead in allocating sub-tasks.” Chow at 1:61-2:11. “The invention affords its users with a number of distinct advantages. Chiefly, the invention guides a multiprocessing system to perform an aggregate task as quickly as possible, with minimum workload skew among independently operating processing elements.”Chow at 2:42-46. “Generally, the method sequence 400 is used to perform an aggregate database processing task, by breaking the task down into smaller units, and executing the units with multiple concurrently operating task execution units.” Chow at 5:1-5. “independently operating each one of multiple task execution units to sequentially and independently self-allocate and execute sub-tasks of the aggregate task on-demand.” Chow at 10:10-13 (Claim 1); <i>see also</i> Claim 8; <i>see also</i> Chow at 1:10-58, 1:60-2:51, 2:54-3:14, 3:18-4:24, 4:26-5:27, 5:29-6:47, 6:49-7:17, 7:19-8:51, 8:53-9:62, FIGS. 1-8. “A distributed computing system achieves a highly distributed environment where very large computation intensive tasks are broken down into thousands of sub-tasks and then distributed to thousands of clients running on a variety of computers across the Internet. The idle CPU time of each of these thousands of client computers is used to perform these computations by running custom application modules in a low priority. A task server keeps track of information associated with each of the clients and uses the information to assign one or more tasks associated with a computing problem to each client computer. A file server provides the

Exhibit A-5

Claims	Gamma System
	application modules to the client computers for executing their assigned tasks. An application server provides input data for the application modules and receives output data from the application modules. Status and performance information for machines, accounts and teams is collected by the task server and displayed on a background page of each client. Incentives for commitments of computing time are provided to users of potential client computers.” Kurowski at Abstract. “The client 200 is one of the pieces in the larger picture of how the distributed computing system 100 described herein provides a highly distributed environment for large scale computing problems. One goal of all these different pieces is to enable the system to take on a very large computing problem, break it down into thousands of tasks and distribute it to tens of thousands of computers across the internet so this large computation could be performed by tens of thousands of parallel sub-computations and the results brought back from all these different clients 200.” Kurowski at [0192].

Exhibit A-5


Claims	Gamma System
	 FIG. 3 Kurowski at FIG. 3. See also Kurowski at [0035]-[0059], [0061]-[0077], [0096]-[0100], [0183]-[0200], FIGS. 1-7, 19, 24-26B. To the extent that Gamma 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 Gamma System, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant's Invalidity Contentions.
[1.1] (a) automatically determining file allocation and logically subdividing records	Gamma System discloses or renders obvious automatically determining file allocation and logically subdividing records of said input file into a plurality of partitions. For example:

Exhibit A-5

Claims	Gamma System
of said input file into a plurality of partitions;	Gamma sub-divides data using horizontal partitioning. “Gamma employs the concept of <i>horizontal partitioning</i> [33] (also termed <i>declustering</i> [29]) to distribute the tuples of a relation among multiple disk drives. This design enables large relations to be processed by multiple processors concurrently without incurring any communications overhead.” https://minds.wisconsin.edu/bitstream/handle/1793/11036/file_1.pdf The following also discloses the logical subdividing of records into a plurality of partitions: <pre> graph TD A["name age salary Bob 20 10K Shideh 18 35K Ted 50 60K Kevin 62 120K Angela 55 140K Mike 45 90K"] B["name age salary Ted 50 60K Kevin 62 120K"] C["name age salary Bob 20 10K Mike 45 90K"] D["name age salary Shideh 18 35K Angela 55 140K"] A --> B A --> C A --> D </pre> https://www.cs.ubc.ca/~rap/teaching/504/2022W2/slides/parallel.pdf As can be seen in the preceding image, the input data is subdivided into partitions.

Exhibit A-5

Claims	Gamma System
	Partitioning of R into N logical buckets for Hybrid hash-join. Figure 8 https://research.cs.wisc.edu/techreports/1990/TR921.pdf The following image is a diagram showing the process of subdividing into partitions. Fig. 3.

Exhibit A-5

Claims	Gamma System
	https://wiki.epfl.ch/edicpublic/documents/Candidacy%20exam/gamma.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 Reiffin, Chow, Kurowski, and/or Yamaguchi. Reiffin, Chow, Kurowski, and Yamaguchi concern similar subject matter as this reference, namely, distributed computing systems and parallel processing. As set forth in the accompanying contentions, parallel processing of data and tasks was a well-known and commonly-used technique. Modifying this reference with Reiffin, Chow, Kurowski, and/or Yamaguchi'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 Reiffin, Chow, Kurowski, Yamaguchi for example, to incorporate this limitation. <i>See, e.g.,</i> “This invention relates to computer networks, and more particularly, to a distributed parallel processing network system wherein a large compute-intensive task may be partitioned into subtasks which are then distributed among a plurality of personal computer or workstations for parallel execution of the subtasks in the background concurrently with the execution in the foreground of the respective local tasks of the individual workstations.” Reiffin at 1:7-14. “If a computer has a task which is both large and compute-intensive, the task is partitioned into a plurality of smaller subtasks to be executed simultaneously in parallel by the other computers of the network. Each of these subtasks is assigned its respective identifier. One of the network disk drives has a directory constituting a pool for storing all of the subtask identifiers, preferably in the form of a queue. Each of the other computers of the network selects a subtask identifier from the pool, copies the corresponding subtask program from the originating computer, executes the selected subtask in the background while concurrently executing its own local task in the foreground, returns the computed results of the subtask

Exhibit A-5

Claims	Gamma System
	computation to the originating computer, and then selects another subtask from the pool to repeat the cycle.” Reiffin at 2:19-33. “This background execution of network tasks is particularly suited for the parallel multiprocessing of large compute-intensive programs which frequently arise in science, engineering, financial analysis, graphics and other fields. The workstations or personal computer originally given such a task may partition the data and program code into a plurality of subtasks. A plurality of subtask identifiers corresponding to the partitioned subtasks are stored in a network directory of a disk drive of one of the computers to constitute a pool or queue identifying remote subtasks to be executed. As explained in more detail below, each computer of the network may access this pool, select a subtask identifier, copy the corresponding remote subtask from the originating computer, execute the subtask in the background without substantial delay of its local processing, and return the computed subtask results to the originating computer.” Reiffin at 4:60-5:8; <i>see also</i> Reiffin at 3:35-6:23, FIGS. 1-5. “Reducing task completion time is especially challenging in ‘multiprocessing environments’, which refer to computer systems where multiple computers, digital data processing units, threads, or other processes perform concurrent tasks. For ease of discussion, the term ‘program execution threads’ is used to refer to these and other arrangements, where some or all of hardware device performs a specific task by executing an instruction stream. ‘Threads’ is a known shorthand term. In multi-thread systems, one important consideration is how to apportion the work among the various concurrently operating threads. Typically, workload is apportioned by somehow logically dividing the data to be processed. For example, a block of data might be divided evenly into a number of parts equal to the number of available threads, so that each thread can independently process a separate portion of the data. This is called ‘static’ or ‘a priori’ partitioning.” Chow at 1:26-42.

Exhibit A-5

Claims	Gamma System
	“A command is received to perform a machine-executed task divisible into multiple sub-tasks, i.e., an ‘aggregate task’. Then, multiple task execution units are independently operated on-demand to sequentially self-allocate and execute sub-tasks of the aggregate task. Each sub-task is sized according to a specific ‘granularity’, used to divide the aggregate task into sub-tasks. The granularity may be predetermined for all task execution units or repeatedly dynamically determined by each task execution unit at the time of self-allocation. Preferably, the granularity is not too large as to cause potentially uneven sub-task allocation, and not so small as to incur excessive overhead in allocating sub-tasks. Operating ‘on-demand’, each task execution unit sequentially allocates and executes one sub-task at a time, then proceeding to allocate the next unallocated pice of work to itself according to the appropriate granularity.” Chow at 1:66-2:15. “The method of claim 1, the execution of sub-task forming intermediate data, wherein: the method further comprises establishing a predetermined number of sort partitions accessible to each of the task execution units; the execution of sub-tasks of the aggregate task rotationally distributes input data into the sort partitions; and the method further comprises performing further processing of data in each partition by assigning each partition to a task execution unit and processing data in each partition exclusively by the task execution unit assigned to that partition.” Chow at 11:5-17 (claim 9); <i>see also</i> Chow at 1:10-58, 1:60-2:51, 2:54-3:14, 3:18-4:24, 4:26-5:27, 5:29-6:47, 6:49-7:17, 7:19-8:51, 8:53-9:62, FIGS. 1-8 “The client 200 is one of the pieces in the larger picture of how the distributed computing system 100 described herein provides a highly distributed environment for large scale computing problems. One goal of all these different pieces is to enable the system to take on a very large computing problem, break it down into thousands of tasks and distribute it to tens of thousands of computers across the internet so this large computation could be performed

Exhibit A-5


Claims	Gamma System
	by tens of thousands of parallel sub-computations and the results brought back from all these different clients 200.” Kurowski at [0192].  FIG. 3 Kurowski at FIG. 3. See also Kurowski at [0035]-[0059], [0061]-[0077], [0096]-[0100], [0183]-[0200], FIGS. 1-7, 19, 24-26B. “In a logical partitioning system, a Web server program is executed on each of several logical partitions, the load on each of the logical partitions is monitored, and when a logical partition is overloaded, data being distributed by the logical partition is automatically switched for

Exhibit A-5

Claims	Gamma System
	distribution by another logical partition that has spare capacity to accommodate the load, whereby load balancing is carried out.” Yamaguchi at Abstract. “Logical partitions LPAR 0 180 b to LPAR5 185 b execute Web server programs that distribute home page data constituted of data written according to a page-description language, such as HTML data, and image data such as GIF data to the Web clients 104 to 106. The Web server programs are stored in the local memories 180 e to 185 e used respectively by logical partitions LPAR0 180 b to LPAR5 185 b, and the home page data is stored in the local disks 120 b to 125 b used respectively by logical partitions LPAR0 180 b to LPAR5 185 b. Home page data is stored in a local disk 142 of the service processor (SVP) 141; at system start-up, the service processor (SVP) 141 transfers data from its local disk 142 to the local disks 120 b to 125 b of respective LPARs based on information entered from the console device 111 by the system administrator giving LPAR numbers and the types of home page data distributed by the LPARs with those numbers. Storage of the home page data is not limited to the local disk 142 in the service processor (SVP) 141; a storage device external to the system can be used instead. In this case, the service processor (SVP) 141, or each of the logical partitions LPAR0 180 b to LPAR5 185 b, or a single LPAR, references the LPAR load table 170 b described below in the shared memory 170 and requests data transfer from the storage device.” Yamaguchi at [0042]. “A load balancing unit 107 is provided between the Web clients 104 to 106 and LPARs. The load balancing unit 107 comprises a URL table (FIG. 14) including URLs (Uniform Resource Locators) giving the addresses of the home page data distributed by each LPAR. Each entry 1100 in the URL table is generated in correspondence with an LPAR number, and stores an LPAR valid flag 1111 that indicates whether the LPAR is active or not, as well as the URL data 1110 of home page data to be distributed by the LPAR. The load balancing unit 107 compares the URLs of home page data in distribution requests made by the Web clients 104 to 106 with each of the URLs of home page data held by the LPARs; the distribution request is sent to the LPAR holding the matching URL. If there are a plurality of LPARs holding the

Exhibit A-5

Claims	Gamma System
	matching URL, these LPARs are selected sequentially or randomly, for example, to prevent load from concentrating on a single LPAR.” Yamaguchi at [0043]. “The load balancing unit 107 has to provide a function of uniquely associating URLs requested by Web clients 104 to 106 with the IP addresses of the logical partitions LPAR0 180 b to LPAR8 188 b. Therefore, the load-balancing unit 107 may function in association with a domain name system (DNS) server, for example, but the description given herein will be limited to noting only that information associating the URLs with the LPARs is required, as shown in the URL table in FIG. 14.” Yamaguchi at [0044]. FIG. 1 Yamaguchi at FIG. 1.

Exhibit A-5

Claims	Gamma System
	<i>See also</i> Yamguchi at [0016]-[0029], [0031]-[0040], [0045]-[0092], FIGS. 2-14. To the extent that Gamma 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 Gamma System, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[1.2] (b) distributing descriptions of all of said partitions to each of a plurality of subtask processors	Gamma System discloses or renders obvious distributing descriptions of all of said partitions to each of a plurality of subtask processors. For example: The data is distributed to partitions to be processed. “The multiprocessor join algorithms provided by Gamma are based on the concept of partitioning the two relations to be joined into disjoint subsets called <i>buckets</i> [21], [28], [6] by applying a hash function to the join attribute of each tuple.” https://minds.wisconsin.edu/bitstream/handle/1793/11036/file_1.pdf “Thus, doubling the cluster size in order to halve (approximately) the increase in the load on the remaining nodes when a failure occurs has the (quite negative) side effect of doubling the probability that data will actually be unavailable. For this reason, Teradata <i>recommends a cluster size of 4 or 8</i> processors” (emphasis added) https://web.eecs.umich.edu/~mozafari/fall2015/eecs584/papers/gamma.pdf “To evaluate the performance of the hypercube version of Gamma three different metrics were used. First, the set of Wisconsin [BITT83] benchmark queries were run on a <i>30 processor configuration using three different sizes of relations: 100,000, 1 million, and 10 million tuples.</i>” (emphasis added) https://web.eecs.umich.edu/~mozafari/fall2015/eecs584/papers/gamma.pdf

Exhibit A-5


Claims	Gamma System
	Gamma used horizontal partitioning:  The diagram illustrates horizontal partitioning. At the top, a table lists six records: Bob (20, 10K), Shideh (18, 35K), Ted (50, 60K), Kevin (62, 120K), Angela (55, 140K), and Mike (45, 90K). Below this, three cylinders represent data partitions. Arrows show the distribution: the first partition contains Ted and Kevin; the second contains Bob and Mike; the third contains Shideh and Angela. https://www.cs.ubc.ca/~rap/teaching/504/2022W2/slides/parallel.pdf The following diagram demonstrates the partitioning of data on different processors:

Exhibit A-5

Claims	Gamma System										
	<table border="1"> <thead> <tr> <th>Value</th><th>Destination Process</th></tr> </thead> <tbody> <tr> <td>0</td><td>(Processor #3, Port #5)</td></tr> <tr> <td>1</td><td>(Processor #2, Port #13)</td></tr> <tr> <td>2</td><td>(Processor #7, Port #6)</td></tr> <tr> <td>3</td><td>(Processor #9, Port #15)</td></tr> </tbody> </table> Fig. 4. An example split table. https://wiki.epfl.ch/edicpublic/documents/Candidacy%20exam/gamma.pdf “The Gamma system runs on a 32-node Intel iPSC/2 Hypercube with a disk attached to each node (the system could use processors without disks for parallel processing). In addition to round-robin, range and hash partitioning, Gamma also provides hybrid-range partitioning that combines the best features of the hash and range partitioning strategies. Once a relation has been partitioned, Gamma provides both clustered and nonclustered indices on either the partitioning or nonpartitioning attributes. The indices are implemented as B-trees or hash tables.” https://cs.brown.edu/courses/cs227/archives/2012/papers/newsq/p85-dewitt.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 Reiffin, Chow, and/or Kurowski. As noted above regarding element [1.0] Reiffin, Chow, Kurowski, and this reference concern similar subject matter, and modifying this reference with Reiffin’s, Chow’s, or Kurowski’s teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. <i>See, e.g.,</i>	Value	Destination Process	0	(Processor #3, Port #5)	1	(Processor #2, Port #13)	2	(Processor #7, Port #6)	3	(Processor #9, Port #15)
Value	Destination Process										
0	(Processor #3, Port #5)										
1	(Processor #2, Port #13)										
2	(Processor #7, Port #6)										
3	(Processor #9, Port #15)										

Exhibit A-5

Claims	Gamma System
	“If a computer has a task which is both large and compute-intensive, the task is partitioned into a plurality of smaller subtasks to be executed simultaneously in parallel by the other computers of the network. Each of these subtasks is assigned its respective identifier. One of the network disk drives has a directory constituting a pool for storing all of the subtask identifiers, preferably in the form of a queue. Each of the other computers of the network selects a subtask identifier from the pool, copies the corresponding subtask program from the originating computer, executes the selected subtask in the background while concurrently executing its own local task in the foreground, returns the computed results of the subtask computation to the originating computer, and then selects another subtask from the pool to repeat the cycle.” Reiffin at 2:19-33; <i>see also</i> Reiffin at 3:35-6:23, FIGS. 1-5. “A command is received to perform a machine-executed task divisible into multiple sub-tasks, i.e., an ‘aggregate task’. Then, multiple task execution units are independently operated on-demand to sequentially self-allocate and execute sub-tasks of the aggregate task. Each sub-task is sized according to a specific ‘granularity’, used to divide the aggregate task into sub-tasks. The granularity may be predetermined for all task execution units or repeatedly dynamically determined by each task execution unit at the time of self-allocation. Preferably, the granularity is not too large as to cause potentially uneven sub-task allocation, and not so small as to incur excessive overhead in allocating sub-tasks. Operating ‘on-demand’, each task execution unit sequentially allocates and executes one sub-task at a time, then proceeding to allocate the next unallocated pice of work to itself according to the appropriate granularity.” Chow at 1:66-2:15. “FIG. 2 depicts an exemplary task execution unit 200, representing one particular hardware embodiment of the task execution units 101-104. The task execution unit 200 includes a processing unit 202, such as a microprocessor or other processing machine. As a specific example, the processing unit 202 may comprise a 64-bit IBM Power-PC microprocessor chip.” Chow at 3:64-4:3.

Exhibit A-5

Claims	Gamma System
	“After the steps 400 are initiated in task 402, step 404 receives the aggregate task. This may involve, for example, receipt of the task-to-be-performed by the task execution units 101-104, or by a central host (not shown) coupled to the task execution units 101-104.” Chow at 5:7-11. <pre> graph LR 210((210)) --- PU[PROCESSING UNIT] subgraph 204 [204] PU --- FAS[FAST ACCESS STORAGE 206] PU --- NVS[NVS 208] FAS --- DSU[DATA STORAGE UNIT] NVS --- DSU end subgraph 200 [200] PU DSU end </pre> The diagram illustrates a task execution unit (200) enclosed in a dashed box. A processing unit (210) is connected to a data storage unit (204). The data storage unit (204) contains fast access storage (206) and non-volatile storage (NVS) (208). The processing unit (210) is also connected to a central host (not shown) via a line labeled 210. FIG. 2 Chow at FIG. 2.

Exhibit A-5


Claims	Gamma System
	 FIG. 4 Chow at FIG. 4. “Still another mechanism to pool and redistribute data during multi-level processing is the Shared Sort technique. Broadly, in the Shared Sort technique, data is pooled into and distributed from a common partition using the straw model. FIG. 8 shows this process in greater detail. After the routine 800 starts in step 802, step 804 establishes the common partition. The common partition is accessible by all task execution units 101-104. Preferably, the common partition is structured to include multiple sort trees, thus avoiding concurrency conflicts from multiple task execution units simultaneously attempting to store data in the common partition.” Chow at 9:31-42.

Exhibit A-5


Claims	Gamma System
	 <pre> graph TD 600 --> 802([START]) 802 --> 804[ESTABLISH COMMON PARTITION] 804 --> 806[AGGREGATE TASK: PROCESS DATA USING STRAW MODEL; POOL RESULTS IN COMMON PARTITION] 806 --> 808[REDISTRIBUTE DATA USING STRAW MODEL; PROCESS AGAIN POOLING RESULTS IN COMMON PARTITION] 808 --> 810{MORE PROCESSING REMAINING ?} 810 -- YES --> 808 810 -- NO --> 614([END]) </pre> FIG. 8 Chow at FIG. 8. See also Chow at 1:10-58, 1:60-2:51, 2:54-3:14, 3:18-4:24, 4:26-5:27, 5:29-6:47, 6:49-7:17, 7:19-8:51, 8:53-9:62, FIGS. 1-8. “It was mentioned above that a user can monitor the progress of his or her account, select which projects to participate in, and watch on the system web site the standing of any team he or she might create or join. Specifically, all machine, user, team, and task information,

Exhibit A-5

Claims	Gamma System
	summaries of such information, or subsets of such information, may be made available through reports and interactive pages on an ‘integrated’ web site. This ‘integrated’ web site serves as a central account management site for all client computers. By way of example, such ‘integrated’ web site may be accessed via the system web site for the distributed computing system 100. The ‘integrated’ web site may be accessed by clicking on a ‘members’ link on the system web site and then logging in by entering a member ID and password.” Kurowski at [0188]; <i>see also</i> Kurowski at [0035]-[0059], [0061]-[0077], [0096]-[0100], [0183]-[0200], FIGS. 1-7, 19, 24-26B. To the extent that Gamma 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 Gamma System, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[1.3] 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;	Gamma System discloses or renders obvious 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. For example: In Gamma, the subdivision allows subtasks to be simultaneously executed. “Since all relations are declustered over multiple disk drives, <i>parallelizing the selection operation</i> involves simply initiating a selection operator on the set of relevant nodes with disks.” (emphasis added) https://web.eecs.umich.edu/~mozafari/fall2015/eecs584/papers/gamma.pdf The following describes specific methods of distributing tasks:

Exhibit A-5

Claims	Gamma System
	“All relations in Gamma are horizontally partitioned [RIES78] across all disk drives in the system. The Gamma query language (gdl - a extension of QUEL [STON76]) provides the user with four alternative ways of distributing the tuples of a relation: • round robin • hashed • range partitioned with user-specified placement by key value • range partitioned with uniform distribution” https://minds.wisconsin.edu/bitstream/handle/1793/58708/TR635.pdf?sequence=1

Exhibit A-5




Exhibit A-5


Claims	Gamma System
	 The diagram illustrates the Gamma System architecture, divided into a HOST layer and GAMMA PROCESSORS layer. In the HOST layer, there are two QUERY MANAGERS (one for each active Gamma user), a SCHEMA, and a CATALOG MANAGER. The CATALOG MANAGER is connected to both QUERY MANAGERS. In the GAMMA PROCESSORS layer, there are RECOVERY PROCESSES, SCHEDULER PROCESSES, and DEADLOCK DETECTION PROCESSES. The SCHEDULER PROCESSES are connected to the CATALOG MANAGER and the RECOVERY PROCESSES. The DEADLOCK DETECTION PROCESSES are connected to the SCHEDULER PROCESSES. Below these are four OPERATOR PROCESSES, each connected to a DATABASE. The OPERATOR PROCESSES are connected to the SCHEDULER PROCESSES. The DATABASES are connected to the OPERATOR PROCESSES. Annotations include: 'One for each active Gamma user' pointing to the QUERY MANAGERS; 'One for each multisite query' pointing to the SCHEDULER PROCESSES; and 'Horizontally partitioned data: round robin; hashed; range partitioned' pointing to the DATABASES. https://www.vldb.org/conf/1986/P228.PDF “In addition to demonstrating that parallelism can really be made to work in a database machine context, the <i>Gamma prototype shows how parallelism can be controlled with minimal control overhead</i> through a combination of the use of algorithms based on hashing and the pipelining of data between processes. Except for 2 messages to initiate each operator of a query tree and 1 message when the operator terminates, the execution of a query is entirely self-scheduling.” (emphasis added) https://minds.wisconsin.edu/bitstream/handle/1793/58708/TR635.pdf?sequence=1 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 Reiffin, Chow, and/or Kurowski. As noted above regarding element [1.0] Reiffin, Chow, Kurowski, and this reference concern similar subject matter, and modifying this reference with Reiffin’s,

Exhibit A-5

Claims	Gamma System
	Chow's, or Kurowski's teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. <i>See, e.g.,</i> “This invention relates to computer networks, and more particularly, to a distributed parallel processing network system wherein a large compute-intensive task may be partitioned into subtasks which are then distributed among a plurality of personal computer or workstations for parallel execution of the subtasks in the background concurrently with the execution in the foreground of the respective local tasks of the individual workstations.” Reiffin at 1:7-14. “If a computer has a task which is both large and compute-intensive, the task is partitioned into a plurality of smaller subtasks to be executed simultaneously in parallel by the other computers of the network. Each of these subtasks is assigned its respective identifier. One of the network disk drives has a directory constituting a pool for storing all of the subtask identifiers, preferably in the form of a queue. Each of the other computers of the network selects a subtask identifier from the pool, copies the corresponding subtask program from the originating computer, executes the selected subtask in the background while concurrently executing its own local task in the foreground, returns the computed results of the subtask computation to the originating computer, and then selects another subtask from the pool to repeat the cycle.” Reiffin at 2:19-33; <i>see also</i> Reiffin at 3:35-6:23, FIGS. 1-5. “Each task execution unit may also comprise a hardware device such as a computer, microprocessor, programmed memory, discrete circuitry, application-specific integrated circuitry, programmable circuit components, logic circuitry, a combination of the foregoing, or another suitable digital data processing machine.” Chow at 3:58-63. “FIG. 4 depicts a sequence of method steps 400 to illustrate one example of the method aspect of the present invention. For ease of explanation, but without any limitation intended thereby, the example of FIG. 4 is described in the context of the multiprocessing system 100 described above, using task execution units 200. Generally, the method sequence 400 is used

Exhibit A-5

Claims	Gamma System
	to perform an aggregate database processing task, by breaking the task down into smaller units, and executing the units with multiple concurrently operating task execution units.” Chow at 4:63-5:5. “After step 404, step 408 concurrently begins processing the aggregate task by concurrently starting each of the task execution units 101-104 on different sub-tasks. When a task execution unit finishes its sub-task, it embarks on another sub-task, if any remain. Thus, the multiprocessing system 100 functions like a ‘straw model’, with several people drinking water from a common glass through individual straws. As a result, all task execution units work together, and finish at substantially the same time.” Chow at 5:13-21. FIG. 4 Chow at FIG. 4. “Thus, as shown above, the aggregate task is not allocated among the task execution units in advance. Instead, the aggregate task is broken into increments (i.e., sub-tasks), which are incrementally allocated to the task execution units ‘on demand’. In other words, the task

Exhibit A-5

Claims	Gamma System
	execution unit 200 operates ‘on-demand’ by sequentially allocating and executing one sub-task at a time, then proceeding to the next unexecuted sub-task.” Chow at 6:34-41. “‘Multi-level’ processing involves the coupling of pipelines of local processing steps via various pooling and redistribution techniques, as described below. Generally, the multi-level processing paradigm of the invention uses the straw model to allocate a first level of work among the task execution units. Then, results of the first-level work are pooled and redistributed among the task execution units. Results from each previous processing level are thus sequentially ‘pooled’ and re-allocated for pipelined local processing by task execution units as a new aggregate task. Although not required, redistribution of work may also be performed according to the straw model. After all work is performed, results are pooled for possible final processing steps and to present as a final output.” Chow at 7:35-48. “FIG. 5 takes a more detailed look at multi-level processing according to the invention. After the sequence 500 begins in step 502, step 504 initially processes the table, index, or other data structure. Preferably, allocation of work in step 504 may be performed according to the straw model of the invention, using the routine 400 (FIG. 4) until all remaining processing depends upon the results of the initial processing. This completes the first ‘level’ of processing.” Chow at 7:49-56.

Exhibit A-5


Claims	Gamma System
	 <pre> graph TD 500 --> 502([START]) 502 --> 504[COMPLETE FIRST LEVEL PROCESSING USING STRAW MODEL] 504 --> 506[POOL RESULTS] 506 --> 508[PERFORM SUBSEQUENCE LEVEL PROCESSING USING STRAW MODEL] 508 --> 510{FURTHER PROCESSING?} 510 -- YES --> 506 510 -- NO --> 511[FINAL POOLING OF RESULTS] 511 --> 512([END]) </pre> FIG. 5 Chow at FIG. 5. <i>See also</i> Chow at 1:10-58, 1:60-2:51, 2:54-3:14, 3:18-4:24, 4:26-5:27, 5:29-6:47, 6:49-7:17, 7:19-8:51, 8:53-9:62, FIGS. 1-8. “With this system, one client 200 can work on one part of a calculation while a different client 200 works on another part of the same calculation. The distributed network of ‘servers’ make the calculation all come together transparently.” Kurowski at [0043]. “The client 200 is one of the pieces in the larger picture of how the distributed computing system 100 described herein provides a highly distributed environment for large scale

Exhibit A-5

Claims	Gamma System
	computing problems. One goal of all these different pieces is to enable the system to take on a very large computing problem, break it down into thousands of tasks and distribute it to tens of thousands of computers across the internet so this large computation could be performed by tens of thousands of parallel sub-computations and the results brought back from all these different clients 200.” Kurowski at [0192]; <i>see also</i> Kurowski at [0035]-[0059], [0061]-[0077], [0096]-[0100], [0183]-[0200], FIGS. 1-7, 19, 24-26B. To the extent that Gamma 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 Gamma System, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[1.4] 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	Gamma System discloses or renders obvious repeating step (c) in at least some of the subtask processors each with another unprocessed partition on a first-come/first-served basis. For example: Gamma uses horizontal partitioning to allocate subtask processors. Specific methods are disclosed in the following paragraph. “All relations in Gamma are horizontally partitioned [RIES78] across all disk drives in the system. The Gamma query language (gdl - a extension of QUEL [STON76]) provides the user with four alternative ways of distributing the tuples of a relation: • round robin • hashed • range partitioned with user-specified placement by key value • range partitioned with uniform distribution” https://minds.wisconsin.edu/bitstream/handle/1793/58708/TR635.pdf?sequence=1

Exhibit A-5

Claims	Gamma System
	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 Reiffin, Chow, Schwiegelshohn 1, Schwiegelshohn 2, and/or Kurowski. As noted above regarding element [1.0] Reiffin, Chow, Kurowski, Schwiegelshohn 1, Schwiegelshohn 2, and this reference concern similar subject matter, and modifying this reference with Reiffin's, Chow's, Schwiegelshohn 1's, Schwiegelshohn 2's, or Kurowski's teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. <i>See, e.g.,</i> “If a computer has a task which is both large and compute-intensive, the task is partitioned into a plurality of smaller subtasks to be executed simultaneously in parallel by the other computers of the network. Each of these subtasks is assigned its respective identifier. One of the network disk drives has a directory constituting a pool for storing all of the subtask identifiers, preferably in the form of a queue. Each of the other computers of the network selects a subtask identifier from the pool, copies the corresponding subtask program from the originating computer, executes the selected subtask in the background while concurrently executing its own local task in the foreground, returns the computed results of the subtask computation to the originating computer, and then selects another subtask from the pool to repeat the cycle.” Reiffin at 2:19-33; <i>see also</i> Reiffin at 3:35-6:23, FIGS. 1-5. “After step 404, step 408 concurrently begins processing the aggregate task by concurrently starting each of the task execution units 101-104 on different sub-tasks. When a task execution unit finishes its sub-task, it embarks on another sub-task, if any remain. Thus, the multiprocessing system 100 functions like a ‘straw model’, with several people drinking water from a common glass through individual straws. As a result, all task execution units work together, and finish at substantially the same time.” Chow at 5:13-21.

Exhibit A-5

Claims	Gamma System
	“The work of an individual task execution unit 200 is described as follows, with reference to steps 410-420. After the task execution unit 200 starts in step 410, step 416 asks whether any work from the aggregate task remains to be processed. If not, the task execution unit 200 quits in step 420.” Chow at 5:22-27. “Otherwise, if step 416 finds work remaining, the task execution unit 200 in step 418 determines a unit of granularity for the next sub-task to be processed, and then, using the determined granularity, obtains the sub-task from the aggregate task. By obtaining the sub-task from the aggregate task, the task execution unit 200 self-allocates the sub-task. The size of the next sub-task to be processed is dictated by a corresponding granularity unit, also determined in step 418.” Chow at 5:29-37. “Thus, as shown above, the aggregate task is not allocated among the task execution units in advance. Instead, the aggregate task is broken into increments (i.e., sub-tasks), which are incrementally allocated to the task execution units ‘on demand’. In other words, the task execution unit 200 operates ‘on-demand’ by sequentially allocating and executing one sub-task at a time, then proceeding to the next unexecuted sub-task.” Chow at 6:34-41.

Exhibit A-5


Claims	Gamma System
	 <pre> graph TD START([START 402]) --> RECEIVE[RECEIVE TASK 404] RECEIVE --> S1[START TASK EXECUTION UNIT 410] RECEIVE --> S2[START TASK EXECUTION UNIT] RECEIVE --> S3[START TASK EXECUTION UNIT 408] S1 --> D1[DETERMINE GRANULARITY; GET NEXT SUB-TASK OF DETERMINED GRANULARITY 418] S2 --> D2[DETERMINE GRANULARITY; GET NEXT SUB-TASK OF DETERMINED GRANULARITY] S3 --> D3[DETERMINE GRANULARITY; GET NEXT SUB-TASK OF DETERMINED GRANULARITY] D1 --> R1{SUB-TASKS REMAINING? 416} D2 --> R2{SUB-TASKS REMAINING?} D3 --> R3{SUB-TASKS REMAINING?} R1 -- YES --> D1 R1 -- NO --> D1 R2 -- YES --> D2 R2 -- NO --> D2 R3 -- YES --> D3 R3 -- NO --> D3 D1 --> D1_2{DONE? 412} D2 --> D2_2{DONE?} D3 --> D3_2{DONE?} D1_2 -- YES --> CONT1[CONTINUE 414] D1_2 -- NO --> D1 D2_2 -- YES --> CONT2[CONTINUE] D2_2 -- NO --> D2 D3_2 -- YES --> CONT3[CONTINUE] D3_2 -- NO --> D3 CONT1 --> END([END 420]) CONT2 --> END CONT3 --> END </pre> FIG. 4 Chow at FIG. 4. See also Chow at 1:10-58, 1:60-2:51, 2:54-3:14, 3:18-4:24, 4:26-5:27, 5:29-6:47, 6:49-7:17, 7:19-8:51, 8:53-9:62, FIGS. 1-8. “Preferably, the client 200 requests work rather than being given a task. In other words, the client 200 preferably requests or asks for work rather than receiving a start message from a server.” Kurowski at [0069]. “As mentioned above, the client 200 preferably requests tasks from the Task Server 1200. At appropriate times in the client's execution, it may decide that its current task is finished and it needs to obtain a new task from the Task Server 1200. In step 220 (FIG. 7), the client 200 issues a request to the Task Server 1200 asking for a new task. This task may be further computation on one of the already existing application modules that the client has cached locally on the disk or it may require the download of a completely new application module. In either case, when the request is sent to the Task Server, the following information is preferably provided which assists the Task Server 1200 in assigning the most appropriate task

Exhibit A-5

Claims	Gamma System
	to the client: Unique machine GUID; CPU number (optional: default is zero); User Id. Regarding the CPU number, if there are more than one processors then the client preferably needs to inform the Task Server which CPU needs the next task. This is because there might be different kinds of tasks running on different CPUs and this information would help the Task Server determine what is the next appropriate task assignment for the client. As mentioned above, application modules are also referred to herein as computation modules, computational modules, task modules, and/or simply modules. In return, in step 222 the Task Server 1200 assigns a task to the client 200 based on information received from the client 200. In step 224 the Task Server 1200 assembles module information relating to the assigned task and sends this module information to the client 200 in step 226. Specifically, the Task Server 1200 sends back the following information to the client 200 which the client needs to determine how it can run the next task: Unique Computational Module ID; Computational Module version number; URL to get the Computation Module if it is not already cached locally; and Checksum for computation module binary files to make sure they are still valid after being downloaded.” Kurowski at [0154].

Exhibit A-5




Exhibit A-5

Claims	Gamma System
	<i>See also</i> Kurowski at [0035]-[0059], [0061]-[0077], [0096]-[0100], [0183]-[0200], FIGS. 1-6, 19, 24-26B. Today’s massively parallel computers are built to execute a large number of different and independent jobs with a varying degree of parallelism. For reasons of efficiency most of these architectures allow space sharing, i.e. the concurrent execution of jobs with little parallelism on disjoint node sets. This produces a complex management task with the scheduling problem, i.e. the assignment of jobs to nodes and time slots being a central part. ... In our paper we want to demonstrate that there are algorithmic issues in job scheduling where theoretical and applied research can both contribute to a solution. To this end we discuss First-come-first-serve (FCFS) scheduling, a simple job scheduling method which can often be found in real systems but is usually considered inadequate by many researchers. Schwiegelshohn, at 1 (p.629). “If the amount of requested nodes for a job exceeds the nodes available on the parallel computer, a scheduling policy is used to settle those resource conflicts. This policy is supposed to reflect the objectives of the owner of the parallel computer as well as the wishes of the users. As jobs may be submitted at any time, often some kind of First-come-first-serve (FCFS) strategy is used. Especially for parallel computers which do not support preemption and if no knowledge of future job arrival is available, various versions of FCFS strategy are frequently used. This also includes modifications like backfilling [17]. Also, an FCFS strategy reflects the notion of fairness which is probably the reason for being more or less accepted by most users. If several queues are used, FCFS may be applied separately to each queue or it may span several or even all queues.” Schwiegelshohn 2, at 184-185. To the extent that Gamma 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 Gamma System, those elements, as known to them or as disclosed by other prior art

Exhibit A-5

Claims	Gamma System
	identified in the cover pleading to Defendant’s Invalidity Contentions.
[1.5] (e) generating at least one output combining all of the subtask outputs and reflecting the processing of all of said subtasks.	Gamma System discloses or renders obvious generating at least one output combining all of the subtask outputs and reflecting the processing of all of said subtasks. For example: Gamma uses merge operators to combine all of the subtask outputs. “Gamma uses <i>split and merge operators</i> to execute relational algebra operators using both parallelism and pipelining. Sort-merge and three different hash join algorithms are supported. Near-linear speedup and scaleup for relational queries has been measured on this architecture. In Gamma, the control node (called scheduler) sends a control message to start a process on a node. The process replies with a message to identify itself. Once this process starts to execute, it reads tuples from its input stream, operates on the tuples and sends result tuples to other processes. When a process detects the end of its input stream, a control message is sent to the control node indicating that it has completed the execution. With the exception of these three control messages, the execution of a process is completely self-scheduling.” (emphasis added) https://cs.brown.edu/courses/cs227/archives/2012/papers/newsq1/p85-dewitt.pdf "We have implemented parallel versions of four join algorithms on the Gamma prototype: sort-merge, Grace, Simple, and Hybrid." https://wiki.epfl.ch/edicpublic/documents/Candidacy%20exam/gamma.pdf “Partition relations into disjoint subsets (buckets) by hashing on the join attribute. • Four types of parallel joins: sort-merge, Grace, Simple, Hybrid. • The Hybrid hash join almost always provides the best performance” https://www.cs.ubc.ca/~rap/teaching/504/2022W2/slides/parallel.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

Exhibit A-5

Claims	Gamma System
	reference further renders the limitation obvious in light of other references, such as Reiffin, Chow, and/or Kurowski. As noted above regarding element [1.0] Reiffin, Chow, Kurowski, and this reference concern similar subject matter, and modifying this reference with Reiffin's, Chow's, or Kurowski's teachings would have only required simple substitution, would not have required extensive experimentation, and would have yielded predictable results. <i>See, e.g.,</i> “If the network is a client-server type then both the subtask identifier pool and the suprograms constituting the subtasks are stored on a server computer and the subtasks are fetched from the latter. The fetched remote subtask or other network task is then executed. The computed results from the execution of each subtask are placed into files which are stored in a network directory which may be accessed by the workstations which originally partitioned and distributed the compute-intensive task.” Reiffin at 6:5-13; <i>see also</i> Reiffin at 3:35-6:23, FIGS. 1-5. “If step 510 finds that further processing still remains, the intermediate results are again pooled in step 506, and step 508 repeated. Steps 506 and 508 are repeated until processing of the data structure is completed. At this point, task 511 pools the results for final output, then, the routine 500 ends in step 512.” Chow at 7:66-8:4.

Exhibit A-5


Claims	Gamma System
	 <pre> graph TD 500 --> 502([START]) 502 --> 504[COMPLETE FIRST LEVEL PROCESSING USING STRAW MODEL] 504 --> 506[POOL RESULTS] 506 --> 508[PERFORM SUBSEQUENCE LEVEL PROCESSING USING STRAW MODEL] 508 --> 510{FURTHER PROCESSING?} 510 -- YES --> 506 510 -- NO --> 511[FINAL POOLING OF RESULTS] 511 --> 512([END]) </pre> FIG. 5 Chow at FIG. 5. <i>See also</i> Chow at 1:10-58, 1:60-2:51, 2:54-3:14, 3:18-4:24, 4:26-5:27, 5:29-6:47, 6:49-7:17, 7:19-8:51, 8:53-9:62, FIGS. 1-8. “With this system, one client 200 can work on one part of a calculation while a different client 200 works on another part of the same calculation. The distributed network of ‘servers’ make the calculation all come together transparently.” Kurowski at [0043]. “The client 200 is one of the pieces in the larger picture of how the distributed computing system 100 described herein provides a highly distributed environment for large scale

Exhibit A-5

Claims	Gamma System
	computing problems. One goal of all these different pieces is to enable the system to take on a very large computing problem, break it down into thousands of tasks and distribute it to tens of thousands of computers across the internet so this large computation could be performed by tens of thousands of parallel sub-computations and the results brought back from all these different clients 200.” Kurowski at [0192]. FIG. 3 Kurowski at FIG. 3. See also Kurowski at [0035]-[0059], [0061]-[0077], [0096]-[0100], [0183]-[0200], FIGS. 1-7, 19, 24-26B.

Exhibit A-5

Claims	Gamma System
	To the extent that Gamma 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 Gamma System, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant's Invalidity Contentions.
[2.0] The method defined in claim 1 wherein the automatic determination of file allocation and logical subdivision of records of said input file into said plurality of partitions in step (a) and the distribution of the description of all of said partitions in step (b) is carried out with at least one further processor in addition to the subtask processors.	Gamma System discloses or renders obvious the method defined in claim 1 wherein the automatic determination of file allocation and logical subdivision of records of said input file into said plurality of partitions in step (a) and the distribution of the description of all of said partitions in step (b) is carried out with at least one further processor in addition to the subtask processors. For example: Gamma uses a range of algorithms for distribution of partitions and has control node that is in addition to the subtask processors. “All relations in Gamma are horizontally partitioned [RIES78] across all disk drives in the system. The Gamma query language (gdl - a extension of QUEL [STON76]) provides the user with four alternative ways of distributing the tuples of a relation: • round robin • hashed • range partitioned with user-specified placement by key value • range partitioned with uniform distribution” https://minds.wisconsin.edu/bitstream/handle/1793/58708/TR635.pdf?sequence=1 “Gamma uses <i>split and merge operators</i> to execute relational algebra operators using both parallelism and pipelining. Sort-merge and three different hash join algorithms are supported. Near-linear speedup and scaleup for relational queries has been measured on this architecture. In Gamma, the control node (called scheduler) sends a control message to start a process on a node. The process replies with a message to identify itself. Once this

Exhibit A-5


Claims	Gamma System
	process starts to execute, it reads tuples from its input stream, operates on the tuples and sends result tuples to other processes. When a process detects the end of its input stream, a control message is sent to the control node indicating that it has completed the execution. With the exception of these three control messages, the execution of a process is completely self-scheduling.” (emphasis added) https://cs.brown.edu/courses/cs227/archives/2012/papers/newsq/p85-dewitt.pdf  Partitioning of R into N logical buckets for Hybrid hash-join. Figure 8 https://research.cs.wisc.edu/techreports/1990/TR921.pdf “Gamma uses split and merge operators to execute relational algebra operators using both parallelism and pipelining. Sort-merge and three different hash join algorithms are supported. Near-linear speedup and scaleup for relational queries has been measured on this architecture. In Gamma, the control node (called scheduler) sends a control message to start a process on a node. The process replies with a message to identify itself. Once this process starts to execute, it reads tuples from its input stream, operates on the tuples and sends result

Exhibit A-5

Claims	Gamma System
	tuples to other processes. When a process detects the end of its input stream, a control message is sent to the control node indicating that it has completed the execution. With the exception of these three control messages, the execution of a process is completely self-scheduling.” https://cs.brown.edu/courses/cs227/archives/2012/papers/newsq/p85-dewitt.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 Reiffin, Chow, Kurowski, and/or Wolf. Reiffin, Chow, Kurowski, and Wolf concern similar subject matter as this reference, namely, distributed computing systems and parallel processing. As set forth in the accompanying contentions, parallel processing of data and tasks was a well-known and commonly-used technique. Modifying this reference with Reiffin, Chow, Kurowski, and/or Wolf’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 Reiffin, Chow, Kurowski, or Wolf for example, to incorporate this limitation. <i>See, e.g.,</i> “If a computer has a task which is both large and compute-intensive, the task is partitioned into a plurality of smaller subtasks to be executed simultaneously in parallel by the other computers of the network. Each of these subtasks is assigned its respective identifier. One of the network disk drives has a directory constituting a pool for storing all of the subtask identifiers, preferably in the form of a queue. Each of the other computers of the network selects a subtask identifier from the pool, copies the corresponding subtask program from the originating computer, executes the selected subtask in the background while concurrently executing its own local task in the foreground, returns the computed results of the subtask computation to the originating computer, and then selects another subtask from the pool to repeat the cycle.” Reiffin at 2:19-33; <i>see also</i> Reiffin at 3:35-6:23, FIGS. 1-5.

Exhibit A-5

Claims	Gamma System
	“The present invention relates to multiprocessing computer database systems. More particularly, the invention concerns a method, apparatus, and article of manufacture for performing an aggregate database processing task with multiple concurrently operating task execution units, using incremental and on-demand sub-task allocation. This is also referred to as ‘straw model’ data processing.” Chow at 1:10-15. “Reducing task completion time is especially challenging in ‘multiprocessing environments’, which refer to computer systems where multiple computers, digital data processing units, threads, or other processes perform concurrent tasks. For ease of discussion, the term ‘program execution threads’ is used to refer to these and other arrangements, where some or all of hardware device performs a specific task by executing an instruction stream. ‘Threads’ is a known shorthand term. In multi-thread systems, one important consideration is how to apportion the work among the various concurrently operating threads. Typically, workload is apportioned by somehow logically dividing the data to be processed. For example, a block of data might be divided evenly into a number of parts equal to the number of available threads, so that each thread can independently process a separate portion of the data. This is called ‘static’ or ‘a priori’ partitioning.” Chow at 1:26-42. “Each task execution unit may also comprise a hardware device such as a computer, microprocessor, programmed memory, discrete circuitry, application-specific integrated circuitry, programmable circuit components, logic circuitry, a combination of the foregoing, or another suitable digital data processing machine.” Chow at 3:58-63. “FIG. 4 depicts a sequence of method steps 400 to illustrate one example of the method aspect of the present invention. For ease of explanation, but without any limitation intended thereby, the example of FIG. 4 is described in the context of the multiprocessing system 100 described above, using task execution units 200. Generally, the method sequence 400 is used to perform an aggregate database processing task, by breaking the task down into smaller units, and executing the units with multiple concurrently operating task execution units.”

Exhibit A-5

Claims	Gamma System
	Chow at 4:63-5:5; <i>see also</i> Chow at 1:10-58, 1:60-2:51, 2:54-3:14, 3:18-4:24, 4:26-5:27, 5:29-6:47, 6:49-7:17, 7:19-8:51, 8:53-9:62, FIGS. 1-8. “The present invention advantageously addresses the needs above as well as other needs by providing a system for use in distributed computing. The system includes a task server, a file server, and an application server. The task server is configured to keep track of information associated with each of a multiplicity of client computers and to use the information to assign one or more tasks associated with a computing problem to each client computer. The file server is configured to provide application modules to the client computers for executing their assigned tasks. The application server is configured to provide input data for the application modules to the client computers and to receive output data of the application modules from the client computers.”Kurowski at [0008]; <i>see also</i> Kurowski at [0035]-[0059], [0061]-[0077], [0096]-[0100], [0183]-[0200], FIGS. 1-7, 19, 24-26B. “Processors 104 are used for the concurrent parallel execution of tasks making up database queries, as described below. A query may originate either from one of the processors 104 or from a separate front-end query processor as described in the concurrently filed application of T. Borden et al., now U.S. Pat. No. 5,495,606. As further described in that application, within each cluster 102 the query splitting and scheduling steps described below may be performed by an additional processor or processors (not shown) similar to processors 104; such additional processors would not be counted among the P/Q processors 104 per complex 102 to which tasks are assigned.” Wolf at 4:63-5:7.

Exhibit A-5


Claims	Gamma System
	 FIG.1 Wolf at FIG. 1. See also Wolf at 4:40-6:12, 6:15-8:27, 8:30-9:58, 9:60-13:18, FIGS. 2-12. To the extent that Gamma 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 Gamma System, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant's Invalidity Contentions.
[5.0] The method defined in claim 1 wherein the output in step (e) is an accumulation of output records from said subtasks in an arbitrary order.	Gamma System discloses or renders obvious the method defined in claim 1 wherein the output in step (e) is an accumulation of output records from said subtasks in an arbitrary order. For example: See [1.0]-[1.5] above. The output is in arbitrary order unless there are specific steps taken to

Exhibit A-5

Claims	Gamma System
	sort it into a particular order. 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 Reiffin, Chow, and/or Lorie. Reiffin, Chow, and Lorie concern similar subject matter as this reference, namely, distributed computing systems and parallel processing. As set forth in the accompanying contentions, parallel processing of data and tasks was a well-known and commonly-used technique. Modifying this reference with Reiffin, Chow, and/or Lorie'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 Reiffin, Chow, or Lorie for example, to incorporate this limitation. <i>See, e.g.,</i> "If the network is a client-server type then both the subtask identifier pool and the suprograms constituting the subtasks are stored on a server computer and the subtasks are fetched from the latter. The fetched remote subtask or other network task is then executed. The computed results from the execution of each subtask are placed into files which are stored in a network directory which may be accessed by the workstations which originally partitioned and distributed the compute-intensive task." Reiffin at 6:5-13; <i>see also</i> Reiffin at 3:35-6:23, FIGS. 1-5. "Such apportionment of work is rarely completely balanced. One reason is that the extent of the data may be difficult to ascertain, or the data's processing time may be otherwise difficult to predict. Furthermore, even equisized data subparts may require comparatively different processing times. As a result of this lack of balance, some threads finish early, and other threads labor for longer periods. Still, the overall job is not complete until all threads have finished, including the last to finish its respective task. Thus, even though some threads complete their work early, other threads linger on, delaying completion of the ultimate task.

Exhibit A-5

Claims	Gamma System
	One name for this phenomenon is called “workload skew”. For users seeking faster execution of their multi-thread application programs, workload skew is a significant problem area.” Chow at 1:43-57. “If step 510 finds that further processing still remains, the intermediate results are again pooled in step 506, and step 508 repeated. Steps 506 and 508 are repeated until processing of the data structure is completed. At this point, task 511 pools the results for final output, then, the routine 500 ends in step 512.” Chow at 7:66-8:4. “After step 704, the task execution units in step 706 cooperatively perform the aggregate task of storing the input data into the sort partitions. Namely, utilizing the technique of FIG. 4, each of the task execution units 101-104 retrieves a data entry from the input source, performs pipeline query processing operations upon the data entry, and deposits the resultant data item(s) into the sort partitions. Within that sort partition, the data entry is placed according to a sort criteria. The sort criteria, however, does not determine which sort partition to place the data entry into. Each time a task execution unit performs pipeline query processing operations upon a data entry from the input source, it deposits the resultant data item(s) into different ones of the sort partitions according to a predetermined “round robin” order. Thus, each task execution unit evenly distributes data entries into all of the sort partitions.” Chow at 9:6-21.

Exhibit A-5


Claims	Gamma System
	 <pre> graph TD 702([START]) --> 704[ESTABLISH PARTITIONS] 704 --> 706[AGGREGATE TASK: DISTRIBUTE INPUT DATA AMONG PARTITIONS IN ROUND ROBIN FASHION] 706 --> 708[EACH TASK EXECUTION UNIT PERFORMS SUBSEQUENT PROCESSING OF ONE PARTITION] 708 --> 710([END]) </pre> FIG. 7 Chow at FIG. 7. See also Chow at 1:10-58, 1:60-2:51, 2:54-3:14, 3:18-4:24, 4:26-5:27, 5:29-6:47, 6:49-7:17, 7:19-8:51, 8:53-9:62, FIGS. 1-8. “In line C01, the coordinator receives the request from the host 16, and determines the processors to which the requests needs to be passed. Lines C04-C06 comprise a loop to read all the local distribution tables DT_i 28_i from all the processors P_i 12. These tables are received on a first come, first received fashion, i.e. in random order. When all DT's have been received, line C07 merges them and computes the global distribution table GDT 30. Lines C08-C10 are a loop that sends the global distribution table GDT 30 to all the processors P_i

Exhibit A-5

Claims	Gamma System
	12. Once this is done, the coordinator task 22 terminates.” Lorie at 6:50-56; <i>see also</i> Lorie at 2:1-5:9, 5:10-6:15, 6:21-7:3, 7:5-24, 7:27-38, 7:44-8:19, 8:25-32, 8:35-12:15, 12:17-18:23, FIGS. 1-6. To the extent that Gamma 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 Gamma System, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant’s Invalidity Contentions.
[7.0] The method defined in claim 1 wherein said input file resides on a network-attached storage and is derived therefrom.	Gamma System discloses or renders obvious the method defined in claim 1 wherein said input file resides on a network-attached storage and is derived therefrom. For example: Gamma uses network attached storage:

Exhibit A-5


Claims	Gamma System
	 The diagram illustrates the Gamma System architecture. At the top, a cloud-like shape represents the 'INTERCONNECTION NETWORK'. Inside this cloud, the words 'INTERCONNECTION' and 'NETWORK' are written, with 'NETWORK' highlighted in blue. Three circular nodes, labeled P_1, P_2, and P_{100}, are connected to the network by lines. Ellipses between P_2 and P_{100} indicate a range of processors. Below the network, a rectangular box is labeled 'MASS STORAGE SYSTEM'. A double-headed vertical arrow connects the bottom of the network cloud to the top of the mass storage system box. Figure 1 https://minds.wisconsin.edu/bitstream/handle/1793/58708/TR635.pdf?sequence=1


Exhibit A-5

Claims	Gamma System
	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 Reiffin, Chow, and/or Markson. Reiffin, Chow, and Markson concern similar subject matter as this reference, namely, distributed computing systems and parallel processing. As set forth in the accompanying contentions, storing data on network-attached storage locations was a well-known and commonly-used technique. Modifying this reference with Reiffin, Chow, and/or Markson'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 Reiffin, Chow, or Markson for example, to incorporate this limitation. <i>See, e.g.,</i> “If a computer has a task which is both large and compute-intensive, the task is partitioned into a plurality of smaller subtasks to be executed simultaneously in parallel by the other computers of the network. Each of these subtasks is assigned its respective identifier. One of the network disk drives has a directory constituting a pool for storing all of the subtask identifiers, preferably in the form of a queue. Each of the other computers of the network selects a subtask identifier from the pool, copies the corresponding subtask program from the originating computer, executes the selected subtask in the background while concurrently executing its own local task in the foreground, returns the computed results of the subtask computation to the originating computer, and then selects another subtask from the pool to repeat the cycle.” Reiffin at 2:19-33; <i>see also</i> Reiffin at 3:35-6:23, FIGS. 1-5. “The multiprocessing system 100 includes multiple task execution units 101-104, each coupled to a data storage unit 106. The task execution units 101-104 comprise electronic devices or processes for executing a sequence of data processing tasks. This data is stored on the data storage unit 106, and may be stored in the form of database tables, for instance. The data storage unit 106 comprises one or more devices for electronically storing data, such as an integrated circuit memory, magnetic hard drive, optical disk drive, magnetic or optical tape,

Exhibit A-5

Claims	Gamma System
	etc.” Chow at 3:23-49; <i>see also</i> Chow at 1:10-58, 1:60-2:51, 2:54-3:14, 3:18-4:24, 4:26-5:27, 5:29-6:47, 6:49-7:17, 7:19-8:51, 8:53-9:62, FIGS. 1-8. “FIG. 4A is a block diagram of one embodiment of storage area network 308A. In this embodiment, storage area network 308A is implemented as two networks that respectively provide network attached storage (NAS) and direct attached storage (DAS).” Markson at [0118]. “One or more control databases 322A, 322B are coupled to a control network 401. One or more storage managers 324A, 324B also are coupled to the control network 401. The control network is further communicatively coupled to one or more disk arrays 404A, 404B that participate respectively in network attached storage network 408 and direct attached storage network 402.” Markson at [0119]. “In one embodiment, network attached storage network 408 comprises a plurality of data movement servers that can receive network requests for information stored in storage units 404B and respond with requested data. A disk array controller 406B is communicatively coupled to the disk arrays 404B for controlling data transfer among them and the NAS network 408. In one specific embodiment, EMC Celerra disk arrays are used” Markson at [0120].

Exhibit A-5

Claims	Gamma System
	FIG. 4A  Markson at FIG. 4A. See also Markson at [0003]-[0015], [0062]-[0073], [0075]-[0079], [0082]-[0102], [0104]-[0117], [0121]-[0129], [0131]-[0174], [0176]-[0182], [0184]-[0193], [0195]-[0201], [0203]-[0212], FIGS. 1A-3C, 4B-9. To the extent that Gamma 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 Gamma System, those elements, as known to them or as disclosed by other prior art identified in the cover pleading to Defendant's Invalidity Contentions.