

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CHOLLA ENERGY LLC AND OBM, INC.,
Petitioners

IPR2024-01361
U.S. Patent No. 11,283,261

**PETITION FOR *INTER PARTES* REVIEW
UNDER 35 U.S.C. § 312 AND 37 C.F.R. § 42.104**

TABLE OF CONTENTS

Petitioners’ Exhibit List	5
I. Introduction.....	7
II. Grounds for Standing.....	7
III. Note.....	7
IV. Overview of the Technology	8
A. Background	8
1. Grid Stability with Renewables	8
2. Harnessing Curtailed Power with Green Datacenters	9
3. Geographically Distributing Compute Load.....	10
B. The ’261 Patent	11
C. Prosecution History	14
D. Effective Priority Date of the ’261 Patent.....	14
V. Level of Ordinary Skill in the Art	14
VI. Claim Construction.....	15
VII. Relief Requested and Reasons Therefor.....	15
VIII. Identification of How the Claims are Unpatentable.....	16
A. Grounds	16
B. Ground 1: Claims 1-2, 9-12 Obvious Over Kiani-Pelio	18
1. Kiani.....	18
2. Pelio.....	20

3.	Reasons to Combine Kiani and Pelio	21
4.	Claim 1	26
5.	Claim 2	50
6.	Claim 9	51
7.	Claim 10	51
8.	Claim 11	52
9.	Claim 12	53
C.	Ground 2: Claims 4-8 Obvious Over Kiani-Pelio-Chapel	53
1.	Chapel	53
2.	Reasons to combine Kiani, Pelio, and Chapel	54
3.	Claim 4	58
4.	Claim 5	64
5.	Claim 6	64
6.	Claim 7	67
7.	Claim 8	68
D.	Ground 3: Claim 3 Obvious Over Kiani-Pelio-Belady	69
1.	Belady	69
2.	Reasons to combine Pelio, Chapel, and Belady	70
3.	Claim 3	74
E.	Ground 4: Claims 13-16 Obvious Over Kiani-Pelio-Chapel- Forestiero	76

1.	Forestiero	76
2.	Reasons to Combine Kiani, Pelio, Chapel, and Forestiero	77
3.	Claim 13	82
4.	Claim 14	89
5.	Claim 15	90
6.	Claim 16	90
IX.	Discretionary Denial is Inappropriate.....	91
A.	No Basis for 35 U.S.C. § 325(d) Denial	91
B.	Discretionary Denial Under the <i>Fintiv</i> Factors is Not Appropriate ...	91
X.	Conclusion	91
XI.	Mandatory notices	92
A.	Real Party-in-Interest	92
B.	Related Matters.....	92
C.	Lead and back-up counsel and service information.....	92
XII.	Appendix of Challenged Claims.....	94
	Certificate of Word Count	100
	Certificate of Service	101

PETITIONERS' EXHIBIT LIST

Ex.1001	U.S. 11,283,261 by McNamara et al.
Ex.1002	Prosecution History of U.S. 11,283,261
Ex.1003	Declaration of Tajana S. Rosing, Ph.D. under 37 C.F.R. § 1.68
Ex.1004	<i>Curriculum Vitae</i> of Tajana S. Rosing, Ph.D.
Ex.1005	Kiani et al., “On the Fundamental Energy Trade-offs of Geographical Load Balancing,” Green Communications and Computing, IEEE Communications Magazine, May 2017
Ex.1006	U.S. Pub. 2011/0316337 by Pelio et al.
Ex.1007	U.S. Pub. 2010/0145542 by Chapel et al.
Ex.1008	U.S. 8,849,469 by Belady et al.
Ex.1009	Forestiero et al., “Hierarchical Approach for Efficient Workload Management in Geo-Distributed Data Centers,” IEEE Transactions on Green Communications and Networking, Vol.1, No. 1, March 2017
Ex.1010	U.S. Pub. 2010/0088261 by Montalvo
Ex.1011	Martinot, E., “Grid Integration of Renewable Energy: Flexibility, Innovation, and Experience,” Annu. Rev. Environ. Resour., vol. 41:223-251, 2016
Ex.1012	Yang et al., “ZCCloud: Exploring Wasted Green Power for High-Performance Computing,” 2016 IEEE International Parallel and Distributed Processing Symposium, 23-27 May 2016
Ex.1013	Bird et al., “Distributed (green) data centers: A new concept for energy, computing, and telecommunications,” Energy for Sustainable Development, Vol. 19, 2014, pp. 83-91
Ex.1014	Yang et al., “Large Scale and Extreme Scale Computing with Stranded Green Power: Opportunities and Costs,” IEEE

	Transactions on parallel and Distributed Systems, December 12, 2017
Ex.1015	Chien et al., “The Zero-Carbon Cloud: High Value, Dispatchable Demand for Renewable Power Generators,” The Electricity Journal, Vol. 28, Issue 8, 2015, pp. 110-118
Ex.1016	Brocanelli et al., “Reducing the expenses of geo-distributed datacenters with portable containerized modules,” Performance Evaluation, Vol. 79, 2014, pp. 104-119
Ex.1017	Li et al., “SolarCore: Solar Energy Driven Multi-Core Architecture Power Management,” 2011 IEEE 17 th International Symposium on High Performance Computer Architecture, February 12-16, 2011
Ex.1018	U.S. Pub. 2016/0109916 by Li et al.
Ex.1019	Zhang et al., “GreenWare: Greening Cloud-Scale Data Centers to Maximize the Use of Renewable Energy,” 12 th International Middleware Conference (MIDDLEWARE), Dec. 2011, Lisbon Portugal (pp. 143-164)
Ex.1020	Blouin, “Exploring Dual-Power Feeds in Data Centers,” https://www.belden.com/blogs/data-center/exploring-dual-power-feeds-in-data-centers , May 17, 2018; accessed February 8, 2024
Ex.1021	IEEE Xplore, https://ieeexplore.ieee.org/abstract/document/7926937 (Kiani)
Ex.1022	IEEE Xplore, https://ieeexplore.ieee.org/abstract/document/7557052 (Forestiero)
Ex.1023	U.S. 9,003,216 by Sankar et al.

I. INTRODUCTION

Cholla Energy LLC and OBM, Inc. (“Petitioners”) respectfully request that the Board review and cancel as unpatentable claims 1-16 (the “Challenged Claims”) of U.S. 11,283,261 (the “’261 patent,” Ex.1001).

The ’261 patent describes a flexible datacenter located “behind-the-meter” (BTM) that receives power before it undergoes step-up transformation for a power grid. This includes a “queue system” and “control system” that determine a computational operation and if enough BTM power is available at the datacenter, and assign the computational operation to the datacenter if there is. But placing a flexible datacenter BTM and using queues and control systems to distribute computational operations to datacenters with sufficient BTM power availability, and datacenter power management strategies based on fluctuating power availability, were already well-known. *See, e.g.*, Ex.1005-Ex.1009 and *infra*.

Because the Challenged Claims merely recite well-known concepts, Petitioners ask the Board to institute trial and find them unpatentable.

II. GROUNDS FOR STANDING

Petitioners certify that the ’261 patent is eligible for IPR and Petitioners are not barred or estopped from requesting IPR. 37 C.F.R. § 42.104(a).

III. NOTE

Emphasis in quotes has been added. Claim terms are in *italics*.

IV. OVERVIEW OF THE TECHNOLOGY

A. Background

1. Grid Stability with Renewables

Maintaining balance between energy supply and demand is a necessary but ongoing challenge in electricity grids. Electricity demand fluctuates for several reasons, as does supply (particularly as more renewable power sources are added to grids). Ex.1010, ¶[0069]; Ex.1011, pp. 2-3; Ex.1003, ¶27.

Some renewable power sources have large output swings which must be considered for stability and economic purposes. Ex.1011, p. 6. During periods of excess power supply, Independent System Operators (“ISOs”)/Regional Transmission Organizations (“RTOs”) initiate curtailment actions to manage power amounts from generation sources for a balanced system. Ex.1011, p. 15. Curtailment typically occurs due to overgeneration, transmission and/or distribution line limitations, or other grid conditions. Ex.1011, p. 15; Ex.1003, ¶28.

During curtailment, some generation sources are temporarily disconnected or operated at reduced capacities. Ex.1010, ¶¶[0045], [0055]–[0056]. Involuntary curtailment is “ordered ... in response to system-reliability conditions,” also referred to as reliability curtailment. Ex.1011, p. 15. Voluntary curtailment relies upon the wholesale market to cause generation sources to voluntarily reduce output due to market conditions, such as a negative wholesale price. Ex.1011, p. 15. This

is also known as economic curtailment. *See* Ex.1012, p. 3. But using curtailment as a tool to maintain grid stability results in wasted electricity (e.g., discarded supply). Ex.1003, ¶¶29-30.

2. Harnessing Curtailed Power with Green Datacenters

Wasted electricity from renewable power sources subject to curtailment is a well-known problem, and many have proposed co-locating different loads (e.g., datacenter’s compute load) to harness wasted energy. One approach co-located datacenters to harness stranded energy and eliminate costly transmission infrastructure and associated losses. Ex.1013, pp. 1, 5 (e.g., transportable containers). Such containers were a well-known strategy for flexible datacenter power source co-location well before the effective date of the ’261 patent. *See, e.g.,* Ex.1006, Abstract, [0024], [0048]-[0050], [0066]; Ex.1012, pp. 1, 4; Ex.1014, p. 7; Ex.1013, p. 5; Ex.1015, pp. 3, 4; Ex.1016, p. 2; Ex.1003, ¶¶31-33.

Datacenters implemented different strategies to manage the variability of stranded energy, including switching between the co-located renewable source and a stable grid source. *See, e.g.,* Ex.1017, pp. 1-2, 9. Another strategy relied solely upon the renewable power source, accepting the volatility of stranded power. Ex.1015, pp. 3, 4; *see also* Ex.1014, pp. 1, 2; Ex.1003, ¶¶34-35. Datacenters handled the volatility by ramping-up/ramping-down server power consumption based on stranded power availability. Ex.1012, p. 4. Such “power-aware

computing” controlled computing power consumption based on the power available. Ex.1018, [0024]. It was already known in the prior art to control power delivery within a datacenter when actual or potential “over capacity conditions” exist, such as the compute load exceeding the amount of power available. Ex.1007, [0020], [0037], [0049], [0050] (over capacity conditions); [0019], [0027]-[0028] (controlling power delivery); Ex.1003, ¶36.


3. Geographically Distributing Compute Load

To further improve power metrics, geographically distributed datacenters have achieved price and energy availability diversity by dynamically distributing workload requests to different datacenter sites. Ex.1016, p. 1. This included using co-located renewable energy for variable green energy (e.g., due to local weather conditions). Ex.1016, pp. 1, 3; Ex.1003, ¶¶37-38.

Known approaches “dynamically dispatched” workload requests “based on green energy availability” to different datacenters. Ex.1016, p. 1; Ex.1019, pp. 3, 4. Known approaches also deployed “portable containerized modules ... for geo-distributed data centers” to exploit diversity, price variability, green energy availability, and ambient temperature at different locations. Ex.1016, p. 2.

Known approaches also utilized queues to manage incoming workload requests for distribution to the different datacenters. Workload requests were “submitted to [a] shared queue” to dispatch the requests to datacenter sites with

green energy when available. *See, e.g.*, Ex.1014, p. 5. This included, e.g., queues at each datacenter location to handle scheduling of workload to green or brown energy. *See, e.g.*, Ex.1005, pp. 2-3; Figure 2:



Ex.1005, Fig. 2; Ex.1003, ¶¶39-40.

B. The '261 Patent

The '261 patent describes “a system for controlling the use of ‘behind-the-meter’ power.” Ex.1001, 1:15-16. Computation operations from different sources can have distinct latency, time sensitivity, or high reliability requirements.

Ex.1001, 5:66-6:9. Computation operations without time sensitivity may be batch processed or distributed across locations, including among critical and flexible, BTM datacenters. Ex.1001, 6:9-14, 23:14-17; Figure 10:





FIG. 10

Ex.1001, Fig. 10 (annotated).

A flexible datacenter includes a mobile container, like a storage trailer on wheels or the ground, or a stackable, inflatable, or floatable container. Ex.1001, 11:4-6. The flexible datacenter might not be a container at all (e.g., “situated within a building”). Ex.1001, 11:6-17. The flexible datacenter includes a behind-the-meter power input system, a power distribution system, climate control, a datacenter control system, and servers. Ex.1001, 10:56-11:3; Figure 2:



A queue system, centralized or decentralized, manages computational operation distribution amongst datacenters. Ex.1001, 6:62-66; 7:1-4, 7:21-24, 7:40-42, 26:51-27:1. Datacenters either receive computational operations distribution or “access and obtain computational operations” from the queue system. Ex.1001, 7:30-32 (distribute), 7:36-39 (access), 7:55-64 (either); Ex.1003, ¶¶41-45.

C. Prosecution History

To overcome an art rejection of the originally-filed claims, the Applicant amended the independent claims to require that the BTM power input system receive power from a power generator before the power underwent step-up transformation for a grid. Ex.1002, p. 449-452. After another rejection, the Applicant further amended the independent claims to require that the flexible datacenter be positioned to receive power directly from the power generation system without the power flowing through a grid first. Ex.1002, p. 165-168. After the Applicant filed a terminal disclaimer over the parent patent 10,367,353, the Examiner allowed the claims and quoted as allowable the claim language about the BTM power input system receiving power from the power generation system prior to the power undergoing step-up transformation for transmission on the grid, instead of sending the power to the grid. Ex.1002, 13-17, 150-51.

D. Effective Priority Date of the '261 Patent

The '261 patent was filed on July 29, 2019 as a continuation of U.S. 10,367,353 filed October 30, 2018. For purposes of this IPR only, Petitioners apply October 30, 2018 as the effective priority date.

V. LEVEL OF ORDINARY SKILL IN THE ART

A Person of Ordinary Skill in The Art (“POSITA”) in October 2018 would have had a working knowledge of the art that is pertinent to the '261 patent,

including green datacenter efforts, renewable power co-location techniques, and geographic workload distribution techniques available in the late 2010's. A POSITA would have had a bachelor's degree in computer science, computer engineering, electrical engineering, mechanical engineering, or an equivalent, and two years of professional experience relating to green datacenters, including the use of renewable energy in powering datacenters, workload distribution amongst geographically distributed datacenters, and power management techniques for datacenters. Lack of professional experience can be remedied by additional education, and vice versa. Ex.1003, ¶¶20-22.

VI. CLAIM CONSTRUCTION

Claim terms in IPR are construed according to their “ordinary and customary meaning” to those of skill in the art. 37 C.F.R. § 42.100(b). Petitioners submit that, for the purposes of this proceeding and the grounds presented herein, no claim term requires express construction. *Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co.*, 868 F.3d 1013, 1017 (Fed. Cir. 2017); Ex.1003, ¶46.

VII. RELIEF REQUESTED AND REASONS THEREFOR

Petitioners ask that the Board institute a trial for *inter partes* review and cancel the Challenged Claims in view of the analysis below.

VIII. IDENTIFICATION OF HOW THE CLAIMS ARE UNPATENTABLE

A. Grounds

Ground	Claims	Basis
#1	1-2, 9-12	35 U.S.C. § 103 over Kiani and Pelio
#2	4-8	35 U.S.C. § 103 over Kiani, Pelio, and Chapel
#3	3	35 U.S.C. § 103 over Kiani, Pelio, and Belady
#4	13-16	35 U.S.C. § 103 over Kiani, Pelio, Chapel, and Forestiero

Kiani et al., “On the Fundamental Energy Trade-offs of Geographical Load Balancing,” (Ex.1005, “**Kiani**”) published in May 2017 as part of the May 2017 issue of the IEEE Communications Magazine, volume 55, no. 5. Ex.1005, p. 1. Kiani includes a copyright date of 2017. Ex.1005, p. 1 (bottom). Kiani was added to IEEE Xplore on May 12, 2017, the same date as which the IEEE Communications Magazine published, and was publicly accessible as of that date. Ex.1021; Ex.1003, ¶¶50-51. Thus, Kiani was publicly accessible well before the earliest priority date of the ’261 patent. *See Hulu, LLC, v. Sound View Innovations, LLC*, IPR2018-01039, Paper 29 at 3, 21 (PTAB Dec. 20, 2019) (precedential)(determination is based on the totality of the evidence).

U.S. 2011/0316337 (Ex.1006, “**Pelio**”) published December 29, 2011.

U.S. 2010/0145542 (Ex.1007, “**Chapel**”) published June 10, 2010.

U.S. 8,849,469 (Ex.1008, “**Belady**”) issued September 30, 2014.

Forestiero, “Hierarchical Approach for Efficient Workload Management in Geo-Distributed Data Centers,” (Ex.1009, “**Forestiero**”) published in 2016. Forestiero includes a copyright date of 2016. Ex.1009, p. 1 (bottom). Forestiero was added to IEEE Xplore on August 31, 2016, and was publicly accessible as of that date. Ex.1022; Ex.1003, ¶¶280-81. Forestiero was part of the March 2017 issue of the IEEE Transactions on Green Communications and Networking. Ex.1022. Thus, because the journal in which Forestiero was included was disseminated to recipients well before the end of the issuing month, the IEEE Transactions journal including Forestiero would have been publicly disseminated by at least March 31, 2017. Ex.1003, ¶282. Thus, a POSITA would have found Forestiero to be publicly accessible as of August 31, 2016, or at the very latest March 31, 2017. *See Hulu*, Paper 29 at 3, 21.

Based on their public accessibility dates, Kiani, Pelio, Chapel, Belady, and Forestiero are prior art under 35 U.S.C. § 102(a)(1).

Petitioners’ grounds rely on the combined teachings of the references and not on a physical incorporation of elements. *See In re Mouttet*, 686 F.3d 1322, 1332 (Fed. Cir. 2012); Ex.1003, ¶¶79, 209, 270, 298.


Petitioners also cite additional prior art below as evidence of the background knowledge of a POSITA and to provide contemporaneous context to support Petitioners’ assertions regarding what a POSITA would have understood from the

prior art in the grounds. *See Yeda Research v. Mylan Pharm. Inc.*, 906 F.3d 1031, 1041-42 (Fed. Cir. 2018) (affirming the use of “supporting evidence relied upon to support the challenge”); 37 C.F.R. § 42.104(b); *see also K/S HIMPP v. Hear-Wear Techs., LLC*, 751 F.3d 1362, 1365-66 (Fed. Cir. 2014); *Arendi S.A.R.L. v. Apple Inc.*, 832 F.3d 1355, 1363 (Fed. Cir. 2016).

B. Ground 1: Claims 1-2, 9-12 Obvious Over Kiani-Pelio

1. Kiani

Kiani optimizes green energy (renewable) utilization and cost for geographically dispersed datacenter systems, including locating datacenters next to renewable power generators, “behind the meter.” Ex.1005, pp. 1-3; Fig. 1:




Ex.1005, Fig. 1.

Excess green energy beyond the datacenter’s need “can be injected into the power grid” or stored. Ex.1005, pp. 3, 4. When there isn’t enough green energy,

the co-located datacenter can use grid-supplied “brown energy.” Ex.1005, p. 3.

“[F]ront-end (FE) servers” near users field computation requests to the datacenter system. Ex.1005, p. 2. The FE servers distribute the requests to the “geo-dispersed data centers,” thus functioning as workload distribution centers. Ex.1005, p. 2; *see also* Fig. 2:




Ex.1005, Fig. 2.

Each datacenter has a queue. Ex.1005, p. 5. An FE server distributes a request by placing an allocated request into such queue. Ex.1005, p. 5. The FE server’s distribution relies on various criteria including datacenter green energy availability. Ex.1005, p. 3. Each datacenter controls its queue length to enforce quality of service, organizing its workload to “utilize[] green energy as much as possible.” Ex.1005, p. 4; Ex.1003, ¶¶52-57.

2. Pelio

Pelio co-locates datacenter containers with a generating station to receive electricity “upstream of any step-up transformer for transmission lines or grids,” before the voltages are “stepped up through transformers to the transmission and distribution utility grids.” Ex.1005, Abstract, [0002]. Generation sources include renewables with time-varying outputs like solar, wind, wave, and tidal. Ex.1005, [0033]; *see, e.g.*, Figure 1A (wind generator example (Ex.1005, [0044])):



Modular datacenter 160 includes containers “home to 1,000 or more servers each.” Ex.1005, [0050], [0071]. Each container receives power from a “modular power supply system[] ... directly from the utility.” Ex.1005, [0051]. Each container distributes the behind-the-meter power via one or both of “‘A-side’ and ‘B-side’ power feeds.” Ex.1005, [0079]. Due to the cyclical nature of renewable sources like solar and wind, when the servers no longer have renewable power the datacenter can either receive power via a backfeed from the grid, or drop pending jobs in favor of another datacenter. Ex.1005, [0037], [0106]; Ex.1003, ¶¶58-61.

3. Reasons to Combine Kiani and Pelio

A POSITA would have been motivated to combine the teachings of Kiani and Pelio for multiple reasons.

First, Kiani and Pelio are analogous art to the '261 patent. The Challenged Claims are in the field of “a system for controlling the use of ‘behind-the-meter’ power.” Ex.1001, 1:15-16. The Challenged Claims recite a flexible datacenter receiving power behind-the-meter, a queue system with subqueues for distributing computational operations, and a control system for determining computational operation assignment based on BTM power availability. *See* Ex.1001, 36:5-36, 37:19-38:17; Ex.1003, ¶¶62-63.

Kiani is in the same field, utilizing “[g]eographical load balancing” to optimize green energy utilization and price diversities at geographically dispersed

datacenters. Ex.1005, p. 1. Kiani greens datacenters “behind the meter” by locating the datacenter with a renewable power generator location before the energy is placed onto the grid. Ex.1005, pp. 2-3. Kiani uses queues to handle distributing computational operations amongst the datacenters based on green energy availability. Ex.1005, pp. 2, 3; Ex.1003, ¶64.

Pelio is also in the same field. Pelio co-locates datacenter containers at a generating station before step-up transformations for transmission and distribution (behind-the-meter). Ex.1006, [0002]. Pelio’s generator examples include renewable types with variable output. Ex.1005, [0037], [0106]; Ex.1003, ¶65.

Kiani is also reasonably pertinent. Per the ’261 patent, datacenter problems included increased power requirements and wasted energy from renewables. Ex.1001, 35:11-15. Kiani teaches solutions that optimize green energy utilization and electricity cost. Ex.1005, p. 1. Kiani further harnesses otherwise wasted green energy by routing computing operations to green datacenters when energy is available. Ex.1005, pp. 1-3. Kiani’s approach considers price diversity, locational diversity, and network delay. Ex.1005, p. 6; Ex.1003, ¶¶66-67.

Pelio’s teachings are also reasonably pertinent. Pelio provides solutions to the ’261 patent-identified problems including placing datacenter containers behind-the-meter to avoid transmission/distribution costs and transmission power loss. Ex.1006, Abstract, [0002], [0019]. Pelio also teaches using renewable power only

when it is available at the datacenter. Ex.1006, [0106]; Ex.1003, ¶68.

Second, a POSITA considering Kiani’s teachings for green load balancing would have been motivated to turn to Pelio for further details about implementing behind-the-meter operation. Kiani focuses on optimizing green energy utilization at distributed datacenters, leaving implementation details for behind-the-meter operation up to a POSITA. *See, e.g.*, Ex.1005, p. 1 (Abstract). Pelio provides behind-the-meter datacenter implementation details, including flexible structural implementations and power delivery teachings. *See, e.g.*, Ex.1006, [0095] (flexible), [0099] (power distribution), [0079] (power feeds); Ex.1003, ¶69.

Kiani discloses datacenters that are “generally dispersed in different geographical regions” for “redundancy, scalability, and high availability.” Ex.1005, p. 2. Pelio, in turn, provides behind-the-meter datacenter implementation details, e.g. a “flexible, just-in-time, modular data center” with a building shell for multiple containers, each of which can be modularly added/removed, supported, cooled, and powered. *See, e.g.*, Ex.1006, [0048]-[0052]. Each container can be “factory-built offsite” and “provide tremendous equipment configuration flexibility,” housing 1,000 or more servers each. Ex.1006, [0066], [0070]-[0071]. Thus, the building easily provides the type of power and scale that Kiani envisioned at geographically dispersed locations, with Pelio’s added benefit of “equipment configuration flexibility.” Ex.1003, ¶¶70-72.

Pelio also provides power provisioning implementation details for Kiani's behind-the-meter datacenter teachings. *See* Ex.1005, pp. 2-3. Pelio provides power to the datacenter by a “number of modular power supply systems 240-[2]43.” Ex.1006, [0051]. Pelio also teaches how to distribute the power to the different containers within the building shell via raceways, and A/B power feeds to servers within each containerized datacenter. Ex.1006, [0099], [0079]; Ex.1003, ¶73.

Third, a POSITA would have had a reasonable expectation of success. Pelio's teachings are compatible with Kiani's teachings; both teach datacenter placement behind-the-meter to harness green energy before grid transmission and distribution. *See, e.g.*, Ex.1005, pp. 2-3, 4 (behind-the-meter, inject excess power into the grid); Ex.1006, [0002], [0047] (datacenters at generator before stepping-up voltages). Further, Kiani and Pelio both contemplate using networks for task/job distribution between different datacenter locations. Ex.1005, p. 2; Ex.1006, [0106]; Ex.1003, ¶¶74-75.

Accordingly, Pelio's teachings provide datacenter implementation details that support Kiani's stated purposes, including behind-the-meter datacenter placement, supporting workload via multiple datacenters, and using networking infrastructure to manage the dispersed workload. These factors lead to a reasonable expectation of success combining Pelio's datacenter implementation teachings with Kiani's geographically distributed datacenter load balancing teachings.

Ex.1003, ¶75.

Fourth, combining Pelio’s teachings with Kiani’s teachings would have been predictable and within a POSITA’s skillset. It was already well-known to locate datacenters behind-the-meter to harness unutilized renewable energy. Others taught locating datacenters “in containers,” co-located with “renewable generators.” Ex.1012, pp. 1, 4; see also Ex.1014, p. 7 (containers co-located at generation sites); Ex.1013, p. 5 (modular containers at distributed locations); Ex.1015, pp. 3, 4 (harnessing curtailed power with co-located containers); Ex.1016, p. 1 (“portable containerized modules ... for geo-distributed data centers”); Ex.1003, ¶76.

A POSITA would have expected Pelio’s flexible datacenter teachings to predictably work with Kiani’s distributed datacenter system teachings, just as has occurred in other known implementations. *E.g.*, Ex.1013, p. 6; Ex.1016, p. 1. In combination, Kiani’s geographic load distribution would be to Pelio’s flexible datacenters with supporting power distribution. *See* Ex.1005, pp. 2-5 (workload distribution to geographically dispersed datacenters to harness green energy, including behind-the-meter); Ex.1006, [0002], [0047] (flexible datacenters behind-the-meter), [0051] (modular power supplies for flexible datacenters), [0079] (A/B side power), [0099] (power distribution within shell); Ex.1003, ¶77.

Implementing Pelio’s flexible datacenter teachings would beneficially avoid

power losses from transformers and transmission lines, thanks to behind-the-meter implementation. Ex.1005, pp. 2-3; Ex.1006, [0019], [0039]. Further, Pelio’s teachings enable Kiani’s system to better “evolve with developing technology,” with cheaper implementation. Ex.1006, [0060]; Ex.1003, ¶77.


The combination would have therefore been nothing more than the combination of known elements (Kiani’s distributed, behind-the-meter datacenters and workload distribution, and Pelio’s behind-the-meter, flexible datacenter implementation details), according to known methods (using Pelio’s flexible datacenter and power distribution teachings for Kiani’s datacenters), to yield predictable results (Kiani’s datacenters having added flexibility while still being behind-the-meter, with power distribution as taught by Pelio). Ex.1003, ¶78.

4. Claim 1

[1.0]

Kiani discloses [1.0], to the extent it is limiting.

Kiani uses “geographical load balancing” to “optimize the utilization of green energy and the cost of electricity.” Ex.1005, p. 1. “[G]reen data centers” “integrate[] renewable energy sources” into their operation. Ex.1005, pp.1-2. Kiani implements “front-end (FE) servers [] col-located with users” to receive service requests and distribute those requests to servers at the dispersed datacenters. Ex.1005, p. 2; Figure 1:




Ex.1005, Fig. 1 (annotated).

Each FE server bases its distribution of incoming workload to different data centers based on green energy availability. Ex.1005, p. 3. Generally, “each data center utilizes green energy as much as possible, and purchases brown energy only when the green energy generation is not adequate to serve all incoming workloads.” Ex.1005, p. 4; Ex.1003, ¶¶80-83.

[1.1]


First, Kiani uses “green data centers” (using renewable energy sources) to service computational requests. Ex.1005, pp. 1-2. A well-known way to “green a data center” was to co-locate the datacenter with a renewable power generator, the generator either owned by the datacenter operator or a third party that sells the power. Ex.1005, pp. 2-3. This co-location is “behind the meter.” Ex.1005, p. 2; Ex.1003, ¶¶84-85.

“[E]ach data center utilizes green energy as much as possible,” purchasing brown energy only when necessary to process all incoming workloads. Ex.1005, p. 4. Kiani obtains additional grid energy when the renewable energy is insufficient, just like the '261 patent. Ex.1001, 24:17-20. Kiani’s teachings demonstrate there are times where the green energy is sufficient for the workloads. Indeed, it was well known for “green” datacenters to operate only using green energy. *See, e.g.*, Ex.1012, p. 4; Ex.1014, pp. 3-4. Such co-location results in the datacenters receiving the generated power before flowing through a grid:



Ex.1005, Fig. 2 (annotated).


The overlapping circles icon in Kiani’s Figure 2 represents a step-up and/or step-down transformer. *See, e.g.*, Ex.1003, ¶¶86-89; Ex.1006, [0044] (transformer 164 stepping up voltage to grid, stepping down voltage from the grid); FIG. 1A:



Ex.1006, Fig. 1A (annotated).


Second, a POSITA would have recognized that Kiani’s datacenters are within building structures. See Ex.1005, p. 2; Ex.1003, ¶90. Each datacenter within a building discloses the “*flexible datacenter*,” since the ’261 patent understood that to encompass not only a container, but also a datacenter “situated within a building.” Ex.1001, 11:4-17. Thus, Kiani’s disclosure alone demonstrates the “*flexible datacenter*.” Ex.1003, ¶90.

Third, it would have been obvious to further implement Kiani’s datacenters within containers for increased operational flexibility, per Pelio. Pelio illustrates multiple “*flexible datacenter[s]*” with each of the containers 224A, 225B and third data center container 225. Ex.1006, Abstract, [0049]; FIG. 2:



Ex.1006, Fig. 2 (annotated).

Figure 4 illustrates “[a] number of data center containers 410-412 on a floor 414.” Ex.1006, [0097], Fig. 4:



Ex.1006, Fig. 4 (annotated).

The containers can flexibly house 1,000 or more servers each. Ex.1006, [0050], [0066], [0070]-[0071]. The containers are placed “in **flexible, just-in-time, modular data centers**.” Ex.1006, [0095]. Pelio’s containers are the very type of flexible datacenters envisioned in the ’261 patent. Ex.1001, 6:30-38. It would have been obvious to implement Pelio’s teachings with Kiani’s teachings to increase datacenter provisioning flexibility. *See, e.g.*, Ex.1006, [0095]; Ex.1003, ¶¶91-95.


Fourth, a POSITA would have understood that Kiani’s behind-the-meter green datacenters are “*positioned to receive electrical power directly from a power generation system*.” *See* Ex.1005, pp. 2-4 (extra green energy the datacenter doesn’t need injected into the grid). The claim explains that receiving electrical power “*directly*” from generator refers to receiving it “*without the received power flowing through a grid*.” Kiani’s behind-the-meter power does not “*flow[] through a grid*” before reaching the datacenter. Ex.1005, pp. 2-3; Ex.1003, ¶¶96-97.

Pelio confirms that a POSITA would have understood Kiani’s behind-the-meter teaching in this manner, disclosing that the datacenter containers receive the power from one or more power generation stations prior to stepping up the power to a grid. Ex.1006, [0002]. The modular datacenters therefore “**take their operating power directly from the electrical generating units** before any corresponding step-up transformer,” expressly avoiding the grid. Ex.1006, [0047]; Ex.1003, ¶¶98-99.

Thus, Pelio's co-located datacenters, examples of Kiani's co-located green datacenters, are able to avoid power losses from "stepping up and stepping down of voltages with transformers" and the long transmission line distances. Ex.1006, [0039]; [0019]; Ex.1003, ¶¶100-01.

[1.2]

First, Kiani teaches locating datacenters behind-the-meter with a generating station. Ex.1005, pp. 2-3; *see* [1.1]. It would have been obvious that Kiani's behind-the-meter (BTM) datacenter teaching would have a corresponding BTM power input system that receives the power from the generating station. Kiani's datacenters receive the power from the renewable source before step-up by a transformer, instead of the received power going to the grid:




Ex.1005, Fig. 2 (annotated).

Thus, Kiani alone renders obvious this limitation. Ex.1003, ¶¶102-04.

Second, Pelio confirms that a co-located datacenter has a BTM power input system. The datacenter receives power from a co-located power generation station before stepping up the voltages for transmission/distribution. Ex.1006, [0002]. As a result, the datacenter receives the “operating power directly from the electrical generating units.” Ex.1006, [0047]. This discloses BTM power input. *See* Ex.1001, 4:62-67; Ex.1003, ¶105.

Third, Pelio expressly discloses a “*power input system*” that receives power “*behind-the-meter*.” The “modular power supply systems 240-[2]43” are “utility sub-panels and transformers,” and/or generators or UPS systems, and are therefore each an example of a BTM “*power input system*,” as is the utility sub-panel 245. This is because each example receives the power prior to step-up transformation:



Ex.1006, Fig. 2 (annotated).



So, too, is the power input system at each datacenter container “*behind-the-*

meter” in Pelio. Indeed, because the power is provided to the modular datacenter in Pelio, the power so-provided is, by definition, provided to the “*flexible datacenter instead of sending the received power to the grid.*” In combination with Kiani’s teachings, the result would be a flexible datacenter that receive BTM power via a BTM power input system. Ex.1003, ¶¶106-08.

[1.3]

First, Kiani explains that servers in datacenters impose a base load (the power consumption from servers even when idle) and a proportional load (the extra power consumed proportional to the workload when not idle). Ex.1005, p. 1. Servers at each datacenter are “powered by the available green energy.” Ex.1005, p. 4. It would have been obvious that the servers, by being powered by the green power source, would have received that power via a “*power distribution system.*” This is a simple concept – any way in which the servers receive the green operating power constitutes a “*power distribution system.*” See, e.g., Ex.1007, [0027]; [0082]. Thus, Kiani alone renders obvious this limitation. Ex.1003, ¶¶109-10.

Second, the combination with Pelio further renders this limitation obvious (1) via raceways that distribute power to different containers, and (2) via power feeds within the containers themselves. The “overhead raceway 440 provides ... electrical power connections” for the building, distributing received BTM power to the different flexible datacenter containers. Ex.1006, [0099]; Figs. 2, 4:




Pelio further renders obvious a power distribution system in each datacenter container, such as would be found in Kiani's datacenters, when explaining the obvious fact that "electrical power [is] consumed by the servers and data storage units" within a data center container. Ex.1006, [0071], [0075]; *see also, e.g.*, Ex.1007, [0027]; [0082]. Pelio teaches power distribution to servers in a datacenter via one or both of "'A-side' and 'B-side' power feeds." Ex.1006, [0079]. Such

power feeds are well-known to be “independent electrical feeds coming into a data center from the utility company.” Ex.1020, p. 1. Each feed includes at least one power distribution unit. Ex.1020, p. 1; *see also* Ex.1007, [0082] (power distribution unit associated with an A source and a B source); Ex.1003, ¶¶111-14.

[1.4]

First, Kiani teaches “hundreds of thousands of servers spread across server farms.” Ex.1005, p. 1. A “server farm” is an exemplary datacenter. Ex.1023, 3:18-20; Ex.1003, ¶¶115-16. Kiani teaches sending computational requests to “the back-end servers at geo-dispersed data centers.” Ex.1005, p. 2. Accordingly, “servers at each data center” are powered by green energy for desired computational requests. Ex.1005, p. 4; Figure 1:



Ex.1005, Fig. 1 (annotated).

As further explained at [1.2] and [1.3], servers at each datacenter are

powered by the BTM power via a power distribution system and BTM power input. Thus, Kiani alone renders obvious [1.4]. Ex.1003, ¶¶116-17.


Second, the combination with Pelio further renders this obvious. Pelio’s datacenter containers contain “network servers” mounted in racks. Ex.1006, [0071]. “[S]tandard-size shipping containers” are typically “home to 1,000 or more servers each,” and can be ordered with servers already installed. Ex.1006, [0071], [0048]. These servers are “computer systems.” Ex.1006, [0008]; Ex.1003, ¶118.

Pelio’s network servers are powered by the power input system via the power distribution system. The “modular power supply systems 240-[2]43,” and utility sub-panel 245, are examples of a “*power input system*” because they receive the power prior to step-up transformation. Ex.1006, [0051]. The power is provided from the utility/power input system to the containers via “electrical cable raceways.” Ex.1006, [0049]. The raceways “provide[] for ... electrical power connections” between the containers and the utility power via the power input system, and the A/B side power feeds to the servers in the racks. Ex.1006, [0099]; [0079]; Ex.1003, ¶¶119-21.

[1.5]

First, Kiani discloses using a queue system to organize the assignment of computational operations across multiple datacenters. Kiani uses “geographical load balancing” to “optimize the utilization of green energy and the cost of

electricity” at the datacenters. Ex.1005, pp.1-2. To accomplish geographical load balancing, datacenters are deployed in various locations (typically farther away from users), with “front-end (FE) servers” closer to users. Ex.1005, p. 2. “Each FE server receives requests from its nearby users” and distributes them to back-end servers at geo-dispersed datacenters. Ex.1005, p. 2; Figure 1 (each circle):




Ex.1005, Fig. 1.

“[E]ach FE server functions as a workload distribution center” by distributing user requests to the datacenters. Ex.1005, p. 2. The FE servers select the datacenters based on parameters like server or content availability, network distance, datacenter efficiency, electricity cost, and renewable energy availability. Ex.1005, p. 2. It would have been obvious that the FE server has a queue system to handle receipt of incoming requests and distribution of outbound assignments to Kiani’s datacenters. Ex.1003, ¶124. This, alone, renders obvious the claimed

“queue system.” Further, for multiple FE servers, as seen in Figure 1, each has its own queue for distributing workload to datacenters. Ex.1003, ¶¶122-24.

Second, Kiani teaches queues at the different datacenters, which further teaches the “queue system.” The FE server distributes allocated requests to the assigned datacenter, where the requests “are first **placed in a queue**” before processing by an available server. Ex.1005, p. 5. Each datacenter has such a queue. Ex.1005, p. 5. In an example, Kiani teaches an “FE server” managing incoming workload distribution to datacenters “based on the availability of green energy” at the datacenters. Ex.1005, p. 3. Generally, “each data center utilizes green energy as much as possible,” turning to brown energy only when necessary. Ex.1005, p. 4. Kiani illustrates the datacenters’ queues in Figure 2:




Ex.1005, Fig. 2 (annotated).

In Figure 2, Kiani illustrates N datacenters, each with a plurality of servers.

All of the queues together, one at each datacenter, constitutes another example of a “queue system.” The combined teachings of Kiani and Pelio result in each flexible datacenter having a queue to receive incoming workload requests from an FE server, as taught by Kiani. Ex.1003, ¶¶125-27.

Third, Kiani further teaches the queue system comprising multiple queue subsystems, with the “queue at each data center” to receive workload allocation from FE servers. Ex.1005, pp. 3, 5. From an FE server, the “allocated requests to a data center are first placed in a queue before they can be processed by any available server” at the datacenter. Ex.1005, p. 5; Figure 2:



Ex.1005, Fig. 2 (annotated).

This is consistent with the '261 patent. According to dependent claims 10 and 11, “a first queue subsystem of the queue system is located at the first flexible datacenter,” and “a second queue subsystem of the queue system is located at the

second datacenter.” Thus, the ’261 patent distributes parts of the queue system amongst the datacenters, just like Kiani. Moreover, a POSITA would have appreciated that Kiani’s FE server operating in cooperation with the subqueues at each datacenter together demonstrates a “*queue system.*” Ex.1003, ¶¶128-31.

[1.6]

First, each datacenter has a queue subsystem that receives workload allocations from the FE server. *See* Ex.1005, pp. 3, 5; [1.5]. Each datacenter has a “QoS ... enforced **by maintaining an upper bound**” on that datacenter’s queue length. Ex.1005, p. 4. Thus, the upper bound is a rule associated with the queue subsystem, and because it impacts the queue itself, the upper bound is a rule that causes the queue to organize the allocated workload (“*a set of computational operations*”) at the datacenter (in the queue) for execution. Ex.1003, ¶¶132-33.

Kiani further teaches another rule at each datacenter: to “maximize the utilization of green energy,” while limiting “the allocated green workload by the available green resources.” Ex.1005, p. 4. If green energy is insufficient to satisfy QoS requirements, brown energy is used. Ex.1005, p. 4. For the brown workload, the rule is “to minimize the cost of electricity” while enforcing the datacenter’s QoS requirements. Ex.1005, p. 4. These rules cause each datacenter’s queue to organize its respective workload with green and/or brown energy, subject to the QoS requirements. Thus, Kiani discloses a plurality of rules associated with the

queue subsystem at each datacenter (Ex.1005, p. 4), including:

- An “upper bound” on the queue length;
- A QoS that determines the “upper bound”;
- Maximize green energy utilization for workload;
- Apply brown energy if green energy unavailable/insufficient;
- Minimize brown energy cost; and
- Maintain the QoS for brown energy operations. Ex.1003, ¶134.

Second, each FE server distributes incoming workload “based on the availability of green energy.” Ex.1005, p. 3. Additional distribution-determining parameters include server or content availability, network distance, datacenter efficiency, and/or electricity cost. Ex.1005, p. 2. Thus, the workload allocation to different datacenters is done according to “rules” (the multiple parameters), which is another act of “*organiz[ing] a set of computational operations.*” The organizing is “*according to rules associated with the queue subsystem*” since each datacenter can have different limitations “due to different service efficiency parameters.” Ex.1005, p. 5; Ex.1003, ¶135.


Thus, because the FE server sends a workload allocation to a queue subsystem at a datacenter due to rules dictating that particular allocation as the best option, the organizing is done “*according to rules*” that are “*associated with*” that queue subsystem. These additional rules, alone or in addition to the execution rules

at the datacenter itself, also disclose this claim limitation. Ex.1003, ¶¶136-37.

[1.7]

The '261 patent specification does not mention a “*routing control system*.” According to the dependent claims, the “*routing control system*” “*is a remote master control system*” positioned “*remotely from the flexible datacenter*” (claim 2) or “*at the flexible datacenter*” (claim 3). Ex.1003, ¶¶138-39.

Kiani discloses a routing control system, like the '261 patent's remote master control system, is configured to perform claim 1's operations. In Kiani, the FE server manages “the distribution of its incoming workload to different data centers based on the availability of green energy.” Ex.1005, p. 3. The FE server distributes allocated requests to the assigned datacenter. Ex.1005, p. 5. Thus, the FE server is an example of a “*routing control system*” (Ex.1003, ¶¶140-41):




Ex.1005, Fig. 2 (annotated).

[1.8]

Kiani discloses a “*routing control system*” with its “FE server” teachings.

See Ex.1005, pp. 3, 5; [1.7]:




Ex.1005, Fig. 2 (annotated).

The FE server receives user requests and distributes them to the geo-dispersed datacenters. Ex.1005, p. 2. The incoming workload (user requests) to the FE server represents a “*computational operation to be performed,*” because the workload constitutes operation requests from users. Ex.1005, pp. 2-3. By knowing what workload the FE server has to distribute, the FE server has “*determine[d], using the queue system,*” the computational operation to be performed (and, thus, distributed to a datacenter). Ex.1003, ¶¶142-45.

[1.9]

First, Kiani’s FE server teachings disclose a “*routing control system.*” See

Ex.1005, pp. 3, 5; [1.7]; Ex.1003, ¶¶46-47:



Ex.1005, Fig. 2 (annotated).

Second, Kiani’s FE server determines how much “green energy,” which is “behind the meter” (“*an amount of BTM power*”), is available to a datacenter (“*available to the plurality of computing systems*”). For example, the FE server distributes its incoming workload to different datacenters “based on the availability of green energy” from the BTM energy sources (e.g., wind, solar) taught in Kiani. Ex.1005, pp. 2-3; Ex.1003, ¶148.

Since each datacenter has multiple servers (“*computing systems*”), of which at least several are powered by “green energy” (BTM), The FE server’s determination of the amount of green energy available to a datacenter teaches a determination of BTM power available to the computing systems at the datacenter. Indeed, Kiani teaches that green workload is defined as “the workload served by a

green server” in a datacenter, with “some servers [] turned on and powered by the available green energy (green servers).” Ex.1005, p. 4; Ex.1003, ¶149.

Kiani demonstrates this concept with an example based on wind power simulation data at three different datacenters. *See* Ex.1005, p. 4. In one time slot, the FE server determines that “before hour 8 and from hours 15 to 24, the available wind power is adequate to serve all the service requests” at the datacenters, resulting in no brown workload allocation. Ex.1005, p. 4. Thus, in this example “the available wind power is the key decision factor to allocate workloads among the data centers” by the FE server. Ex.1005, p. 4; Ex.1003, ¶150.


Third, this limitation is further rendered obvious combined with Pelio’s teaching of receiving power BTM by a co-located datacenter. *See* [1.1] and [1.2]. In combination, Kiani’s FE server takes into account the BTM power available according to both Kiani’s and Pelio’s BTM teachings. Ex.1003, ¶¶151-52.

[1.10]

First, Kiani’s FE server distributes workload to datacenters based on the datacenters’ respective “availability of green energy.” Ex.1005, p. 3. And, as noted at [1.1] and [1.2], Kiani discloses providing green energy “behind the meter.” Ex.1005, pp. 2-3. The “availability of green energy” determination, together with Kiani’s “behind the meter” teaching for that green energy, demonstrates “*determine whether to route the computational operation*” to a datacenter “*using*

the amount of BTM power that is available.” Indeed, in an example “the available wind power is **the key decision factor** to allocate workloads among the data centers.” Ex. 1005, p. 4. As explained at [1.11] below, Kiani further teaches allocating workload to a datacenter, which demonstrates routing a computational operation to that datacenter. *See* Ex.1005, pp. 2, 4; Ex.1003, ¶¶153-55.

Second, Kiani-Pelio renders obvious a flexible datacenter placed BTM. *See* [1.1]; Ex.1005, pp. 2-3; Ex.1006, Abstract, [0049], [0071]; Figure 2:




Ex.1006, Fig. 2 (annotated).


Accordingly, Pelio’s flexible datacenter teachings are used combined with Kiani’s BTM teachings, as well as Kiani’s distributing (routing) workload (computational operations) to the BTM datacenter based on availability of BTM green energy (amount of BTM power available). Ex.1003, ¶¶156-58.

[1.11]

Kiani teaches each FE server allocating the workload to a determined




Kiani’s workload allocation discloses “*assign[ing] the computational operation,*” since Kiani’s distribution of the workload (an example of computational operations) to a determined datacenter is an act of assigning for performance at the datacenter. Ex.1005, p. 2; Ex.1003, ¶161.



Ex.1005, Fig. 4 (annotated).

Kiani illustrates the workload allocation in Figure 3; e.g., at hour 14:




Ex.1005, Fig. 3 (annotated).

Kiani's distribution/allocation of workload is expressly an act of "assign[ing]" to a given datacenter (e.g., datacenter 1 in the above example, which further in combination with Pelio is a flexible datacenter). Kiani's allocation of the workload is based on determining to route the workload based on the amount of BTM green energy at the datacenter. *See* [1.10]; Ex.1005, pp. 2-4; Ex.1006, Abstract, [0049], [0071]; Ex.1003, ¶¶162-67.

5. Claim 2

First, Kiani teaches locating the FE server (a "*routing control system*") close to users, and thus remotely from the datacenters that are typically farther away from those same users. Ex.1005, p. 2. Kiani illustrates this in Figure 1:



Ex.1005, Fig. 1 (annotated).

Each FE server is a "*remote master control system*." Ex.1003, ¶¶168-71.

Second, Kiani's teachings alone, and combined with Pelio, render obvious a


flexible datacenter placed BTM. *See* [1.1]; Ex.1005, pp. 2-3 (within a building, like the '261 patent); Ex.1006, Abstract, [0049], [0071]; Figure 2; Ex.1003, ¶¶172-73.

6. Claim 9

Kiani's FE server discloses a routing control system that distributes user requests to determined datacenters' queues. *See* [1.7]; Ex.1005, p. 2. This demonstrates that the FE server, a part of the "queue system," controls distribution of computational operations from the FE server. *See* [1.5]; Ex.1003, ¶¶174-77.

7. Claim 10

First, Kiani discloses a queue at each datacenter. *See* [1.5]. When an FE server allocates requests (workload), they are "placed in a queue" at the assigned datacenter. Ex.1005, p. 5. Thus, Kiani discloses a first queue at a first datacenter. *See also* Ex.1005, Figure 2; Ex.1003, ¶¶178-80:




Ex.1005, Fig. 2 (annotated).

Second, Kiani’s and Pelio’s combined teachings render obvious a flexible datacenter placed BTM. *See* [1.1]; Ex.1005, pp. 2-3; Ex.1006, Abstract, [0049], [0071]; Figure 2. Accordingly, such a flexible datacenter in combination with Kiani’s queue teachings results in a “*flexible datacenter*” with a queue (a “*first queue subsystem*”). Ex.1003, ¶¶181-82.

8. Claim 11

Kiani discloses multiple datacenters, with a “queue at each data center” for workload requests. Ex.1005, p. 5. Thus, with multiple datacenters, Kiani disclose a second queue at a second datacenter. *See also* Ex.1005, Figure 2:



Ex.1005, Fig. 2 (annotated).

Claim 11 does not recite that the second datacenter is a “*flexible*” datacenter. Thus, Kiani alone discloses this limitation with its datacenter disclosures. Even if

the second datacenter is a flexible datacenter as well, this would have also been obvious from Kiani alone and in combination with Pelio's flexible datacenter teachings, since Pelio teaches that it was known to have multiple flexible datacenters. *See* Ex.1006, [0049]; FIG. 2; Ex.1003, ¶¶183-87.


9. Claim 12

Kiani teaches multiple datacenters, each having its own queue. *See* [11.1]; Ex.1005, p. 5. Further, it would have been obvious for Kiani's additional datacenters to likewise be flexible datacenters. *See* [11.1]. This would have been nothing than the mere duplication of Pelio's flexible datacenter teaching for additional datacenters taught by Kiani. Moreover, Pelio teaches that it was known to have multiple flexible datacenters. *See* Ex.1006, [0049]; Ex.1003, ¶¶188-91.

C. Ground 2: Claims 4-8 Obvious Over Kiani-Pelio-Chapel

1. Chapel

Chapel relates to “electrical power distribution and management,” including intelligently monitoring and “control[ling] delivery of electricity” to devices including datacenter servers. Ex.1006, [0002], [0047], [0080]. The servers are plugged into receptacles, which are in communication with a local controller. Ex.1006, [0082]. The system “communicate[s] with a local controller” at the customer premises, which communicates “with a remote controller” to control power to the servers. Ex.1006, [0056], [0083]; *see, e.g.*, FIG. 6:



Ex.1006, Fig. 6.

When it is determined to “power down or power up a device or set of devices,” the remote controller sends a message to the local controller, which conveys it to the receptacle/server. Ex.1006, [0090]; Ex.1003, ¶¶192-95.

2. Reasons to combine Kiani, Pelio, and Chapel

First, Chapel describes approaches for “electrical power distribution and management” of electrical devices, including controlling power delivery to them within a datacenter. Ex.1007, [0002], [0047]. Chapel teaches controlling the power delivery for multiple reasons, including when power supply is insufficient (either actual or projected). Ex.1007, [0048]. For example, Chapel teaches solutions to the need of modulating power delivery in response to power supply changes, including intelligently monitoring and controlling electricity delivery over a single circuit or an entire datacenter. Ex.1007, [0002], [0019], [0056], [0079]. Chapel’s solution

includes control for powering up and down multiple devices at the datacenter based on predicted or actual power supply status. Ex.1007, [0090]. Thus, a POSITA would have found Chapel's teachings pertinent to problems identified in the '261 patent. Ex.1003, ¶¶196-98.

Second, Kiani and Pelio both describe datacenters receiving behind-the-meter power, thereby avoiding transmission/distribution costs and associated energy losses. *See* Ex.1005, pp. 2-3; Ex.1006, Abstract; [0002], [0019]. Pelio further teaches distributing the behind-the-meter power to the containers, and to servers within a datacenter container via “power feeds,” but leaves the details up to POSITAs. Ex.1006, [0051], [0079], [0099]. A POSITA would have been motivated to turn to Chapel for additional power distribution system implementation and control details. Chapel describes a datacenter including “a power distribution unit” associated with one or more power sources. Ex.1007, [0027]; [0082]. Power distribution units were well-known in datacenter design, including Kiani's and Pelio's datacenters. Indeed, behind-the-meter renewable power sources were already well-known. *See, e.g.*, Ex.1006, Abstract; Ex.1012, pp. 1, 4; Ex.1014, p. 7; Ex.1013, p. 6; Ex.1015, pp. 3, 4; Ex.1003, ¶¶199-200.

Third, Kiani and Pelio both recognize the variability of renewable sources, including “the weather dependency of green energy” (Ex.1005, pp. 2, 3) and the cycling “between full operation” and as far as “zero” output due to lack of the

renewable resource (Ex.1006, [0037], [0106]). Kiani's and Pelio's teachings place datacenters behind-the-meter to harness the power output before transmission costs and losses, and for equipment configuration flexibility, leaving up to a POSITA the details for handling power source variability. *See* Ex.1005, pp. 2-3; Ex.1006, [0002]; [0070]. Chapel teaches monitoring and controlling delivery of electricity to servers within the datacenter based on the available power supply. Ex.1007, [0002]; [0047]; [0080]-[0083]. Chapel accomplishes power management via the receptacles that servers plug into, one or more local controllers, and a remote controller. Ex.1007, [0080]-[0083]; [0090]. The remote controller sends a command to the local controller, which conveys the command within the datacenter to control power delivery. Ex.1007, [0090]. Thus, a POSITA would have been further motivated to turn to Chapel for Chapel's datacenter power management teachings. Ex.1003, ¶¶201-03.

Fourth, a POSITA would have had a reasonable expectation of success in combining the teachings of Chapel with Pelio and Kiani, because Chapel's teachings are compatible with Kiani's and Pelio's teachings – e.g., Kiani and Pelio both teach providing power to servers in datacenters from a variable power source behind-the-meter. Ex.1005, p. 3; Ex.1006, Abstract; [0002]; [0033]; [0037]; [0050]; [0106]. Chapel teaches how to modulate the power delivery to servers within a datacenter. Ex.1007, [0002]; [0047]; [0080]-[0083]. Kiani's and Pelio's

variable power supply would be monitored according to Chapel's monitoring teachings and proper control implemented per Chapel's power management teachings. *See id.*; Ex.1003, ¶204.

Combining Chapel's, Kiani's, and Pelio's teachings would have been predictable and within a POSITA's skillset. It was already known to locate datacenters behind-the-meter with renewable power sources; others also taught locating datacenters with "renewable generators." Ex.1012, pp. 1, 4; Ex.1014, p. 7 (containers co-located at generation sites); Ex.1013, p. 5 (modular containers at distributed site locations); Ex.1015, pp. 3, 4 (harnessing curtailed power with co-located containers); Ex.1016, p. 1 ("portable containerized modules ... for geo-distributed data centers"); Ex.1003, ¶205.

It was also well-known to control the datacenter's load in response to a renewable power source's varying power supply. A known approach was to have the servers "only work when stranded power is available, and otherwise [have them] fully powered off." Ex.1012, p. 4. Such "power-aware computing" controlled "the amount of power consumed by the computing based on the amount of power available." Ex.1018, [0024]. Chapel teaches, for example, powering down when supply is lower than demand to avoid an over-capacity condition, and powering up when supply is sufficient for more load. *See* Ex.1007, [0020], [0037], [0049], [0050]. Implementing Chapel's power control teachings with Kiani's and

Pelio's teachings would beneficially enable behind-the-meter datacenters to follow a varying power supply. Ex.1003, ¶206.

The Kiani-Pelio-Chapel combination would predictably result in Kiani-Pelio's datacenters including Chapel's teachings for controlling power delivered to servers. *See* Ex.1005, pp. 2-3; Ex.1006, Abstract; [0002]; [0033]; [0037]; [0050]; [0106]; Ex.1007, [0002]; [0047]; [0080]-[0083] (monitoring, controlling power delivery). Doing so would advantageously enable Kiani's datacenters to better follow the power supply and avoid grid electricity costs when the renewable supply dwindles. Ex.1003, ¶207.


Combining Kiani-Pelio-Chapel is nothing more than the combination of known elements (Kiani's, Pelio's behind-the-meter datacenter teachings and Chapel's datacenter power control teachings), according to known methods (Chapel's power monitoring and delivery control teachings with Kiani's datacenters), to yield predictable results (Kiani's power consumption following available supply, per Chapel's power delivery control teachings). Ex.1003, ¶208.

3. Claim 4

First, Kiani-Pelio renders obvious a power delivery system that delivers power to the plurality of computing systems. *See* [1.3]; Ex.1003, ¶¶210-12.

Second, both Kiani and Pelio teach placing datacenters near a power generating source behind-the-meter so a higher percentage of the produced power

actually reaches the datacenter load Ex.1005, pp. 2-3; *see also* Ex.1006, [0002], [0019]. Renewable power sources can have varying power output, e.g. due to “the weather dependency of green energy.” Ex.1005, p. 3. Kiani illustrates this variability with “the trends of wind power” at three different datacenters over 24 hours (each horizontal line showing power output from BTM wind turbines):



Ex.1005, Fig. 4.

Pelio similarly recognizes that there are times when power is less available from renewables than other times, and teaches some action to account for that shortfall. *See* Ex.1006, [0037]. Both Kiani and Pelio’s solutions include drawing inexpensive power from a grid (“brown energy,” Ex.1005, p. 4), or passing jobs to other datacenters/power down. Ex.1006, [0037]; [0043]; [0105]-[0106]. The

amount of green energy available, wind speed, sunlight conditions, cost of brown electricity, are all examples of “*power system conditions*.” Ex.1003, ¶¶213-15.

Kiani and Pelio both leave details up to a POSITA for how to control power at a datacenter due to varying green power availability. A POSITA would have been motivated to turn to Chapel, which teaches ways to control power within a datacenter to address actual or potential power shortfalls, such as those arising from variable sources of power like solar, wind, etc. Ex.1003, ¶216.

Third, Chapel teaches how to control power within a datacenter for actual or potential “over capacity condition[s].” Ex.1007, [0048]. Such a condition exists when the available power supply (e.g., from a renewable station) falls below the power demand of a datacenter. *See* Ex.1005, p. 3; Ex.1006, [0035]. Chapel teaches controlling power in a datacenter by modulating power delivery to one or more of the servers to address such conditions. Ex.1003, ¶217.

Chapel teaches “electrical power distribution and management,” including “intelligently monitor[ing] at least a portion of the circuit and control[ing] delivery ... over the circuit.” Ex.1007, [0002]. Chapel’s teachings apply to “a data center” with “data center equipment control and power distribution systems.” Ex.1007, [0019], [0047]. Chapel’s teachings enable “remote monitoring and/or control of multiple receptacles” providing power to servers. Ex.1007, [0056]. “[A] local controller 304” at “the customer premises” communicates with the receptacles, and

also “with a remote controller 308” operated by the premises owner, a controller of the power grid, etc. Ex.1007, [0056]; FIG. 3:





FIG.3

Ex.1007, Fig. 3 (annotated).

The “local controller” discloses a “*datacenter control system*,” since it is a control system at the datacenter. Ex.1007, [0056]. The local controller “control[s] delivery of power to the electrical device,” which includes turning off devices or operating the devices in reduced power mode – e.g., by interrupting power distribution. Ex.1007, [0018], [0025]-[0026]; Ex.1003, ¶¶218-19.

Chapel’s Figure 6 illustrates “control in a data center environment,” with servers plugged into controllable receptacles. Ex.1007, [0080]-[0081]. “[A] local controller 650” communicates with the receptacles and a remote controller 670. Ex.1007, [0082]. The local and remote controllers collectively provide “remote

and/or policy based operation” of the servers. Ex.1007, [0083]. For example, “an operator of a remote controller unit may choose to power down or power up a ... set of devices,” so the remote controller transmits a message to the local controller to implement the action. Ex.1007, [0090], FIG. 6; Ex.1003, ¶220.



Ex.1007, Fig. 6 (annotated).

Fourth, Chapel teaches the local controller’s (“*database control system*”) modulating power delivery is based on monitored conditions of the power system. Kiani and Pelio both teach datacenters receiving BTM power from variable renewable power sources. Ex.1005, pp. 2-3; Ex.1006, [0002], [0019], [0033], [0039]. It would have been obvious that to know how much workload to allocate as green workload (Ex.1005, p. 4) and/or whether there is need to throttle up/down another power source (per Kiani and Pelio, e.g. Ex.1006, [0033]), or throttle the load itself (Ex.1006, [0106], Ex.1007, [0089]), power sources’ varying output

would itself be monitored in Kiani and Pelio. Ex.1003, ¶¶221-22.

Chapel confirms this, teaching monitoring power availability conditions from power sources. Ex.1007, [0004], [0082]. This includes “**remotely monitoring** and/or controlling” datacenter servers. Ex.1007, [0047]. For example, Figure 3’s system 300 achieves “remote monitoring and/or control of multiple receptacles” using local controller 304, which communicates monitoring information to remote controller 308 (a “*control system*”). Ex.1007, [0056]. Chapel’s Figure 6 example has local controller 650 send power availability and load information to remote controller 670. Ex.1007, [0083]; Ex.1003, ¶223.

In combination, it would have been obvious for a remote controller (per Chapel) to receive power information (a “*monitored power system condition*”) about BTM power availability (e.g., BTM power sources per Kiani/Pelio, and recognized by Chapel). *See* Ex.1005, pp. 2-4; Ex.1006, [0033], [0106]; Ex.1007, [0004], [0089]. Other monitored conditions include wind speed, solar cover, cost of brown electricity, etc. With the power conditions monitoring taught by Chapel, Kiani-Pelio’s flexible datacenter would thus monitor power availability of the BTM power sources at the datacenter. Ex.1003, ¶224.

Further, the local controller at the datacenter per Kiani-Pelio-Chapel would modulate power delivery to the servers, as Chapel teaches, based on monitored conditions (e.g., monitored at a remote controller, like Kiani’s FE server with

instructions to the local controller for implementation). Ex.1003, ¶¶225-26.

4. Claim 5

First, Kiani-Pelio-Chapel teaches “*one or more monitored power system conditions.*” See [4.1]; Ex.1003, ¶¶227-29.

Second, Kiani particularly teaches that the availability of BTM power at the datacenter – and, thus, at the datacenter’s “back-end servers” – is a monitored power system condition. See Ex.1005, p. 3 (“based on the availability of green energy”); p. 2 (“back-end servers”). To know how much workload to allocate as green workload (per Kiani), and/or whether to throttle up/down another power source (per Kiani and Pelio) or the level of load, a POSITA would recognize that the current status of green energy availability BTM at the datacenter (and, thus, to servers in the datacenter) would need to be monitored. See [4.1]. The combination with Chapel renders obvious such monitoring of Kiani’s BTM green energy availability. Ex.1007, [0004], [0047], [0082]; Ex.1003, ¶¶230-32.

5. Claim 6

Kiani combined with Chapel alone, and further combined with Pelio, renders obvious claim 6 several ways, any one of which invalidates the claim. See claim 6 (“*one or more of ... or*”). Ex.1003, ¶¶233-35.

Low Local Power

First, it would have been obvious that in different situations, BTM co-

located power sources would generate excess power or not enough power. Kiani teaches that green energy can have a weather dependency, and thus “is not a constantly available resource to power” datacenters. Ex.1005, p. 3. Thus, there are times when the generation station generates less green energy than the datacenter needs for its workload – where “green energy generation is not adequate to serve all incoming workloads.” Ex.1005, p. 4. This renders obvious “*low local power generation*.” Thus, Kiani alone renders obvious claim 6. Ex.1003, ¶236.

Second, the combination with Pelio further renders this obvious. Pelio also recognizes the well-known fact that some types of renewable power generating stations, such as solar and wind, “cycle between full operation and zero.” Ex.1006, [0037]. Thus, there are times where the generation station generates less power than the datacenter needs, such as nighttime (solar), or low wind. Ex.1006, [0037]; *see also* Ex.1008, FIG. 3 (available capacity below needs); 10:14-21. Pelio thus further renders obvious “*low local power generation*.” Ex.1003, ¶237.

Excess Local Power Generation

First, it would have also been obvious that at other times, there would be excess local power generation. Kiani teaches using “green energy as much as possible,” only reverting to brown energy when there isn’t enough green energy generation. Ex.1005, p. 4. Indeed, “extra green energy generation at a data center can be injected into the power grid” when it isn’t all needed for incoming

workloads. Ex.1005, p. 4. Thus, Kiani teaches times of excess green energy beyond what the datacenter's workload needs, rendering obvious "*excess local power generation at a local station level.*" Ex.1003, ¶238.

Second, combining with Pelio further renders this obvious. Pelio recognizes that BTM renewable power generating stations, such as solar and wind, "cycle between full operation and zero." Ex.1005, [0037]. Thus, there are times where the power source generates more power than the datacenter needs, such as during peak sunshine or high wind. Ex.1006, [0037]; *see also* Ex.1008, FIG. 3 (capacity exceeds computational needs); 10:14-21. For example, Pelio teaches providing excess power, after the datacenter's needs, onto the grid. Ex.1006, [0044]-[0046]. Such extra generated power renders obvious "*excess local power generation at a local station level.*" Ex.1003, ¶¶239-40.

Transient Local Power

First, Kiani teaches that renewable local power generation is "*transient*," because "green energy is not a constantly available resource to power" datacenters due to, e.g., "weather dependency." Ex.1005, p. 3. This "weather dependency" describes a "*transient local power generation situation*["]. Ex.1003, ¶241.

Second, Pelio further teaches that the availability of renewable power is "*transient*," such as "a local data center's electric generating plant" going down "due to lack of fuel, sun, wind, gas," e.g. from "the daily cycles of a solar energy

generating station that depend on sunshine.” Ex.1006, [0106], [0037]. The cycling between full operation and zero also demonstrates a “*transient local power generating situation*[].” Ex.1003, ¶¶242-43.

6. Claim 7

First, Kiani-Pelio-Chapel renders obvious a “*datacenter control system ... configured to modulate power delivery to the plurality of computing systems*” based on monitored power system conditions. See [4.1]; Ex.1003, ¶¶244-46.

Second, Kiani-Pelio-Chapel further renders obvious the datacenter control system modulating power delivery “*based on an operational directive.*” Chapel teaches the local controller (“*datacenter control system*”) modulating power delivery by policies, including those communicated from a remote controller like “a message from the electrical utility provider specifying transition to an efficiency mode or a brownout mode.” Ex.1007, [0076]; *see also* [0058], [0066], [0090]. Such “messages” and policies communicated from remote controller(s) render obvious the claimed modulation of the power delivery “*based on an operational directive.*” Ex.1003, ¶¶247-48.

In combination, Kiani-Pelio-Chapel teaches a “local controller” (per Chapel) determining to control power delivery by interrupting power distribution from the power distribution system to one or more servers within Kiani-Pelio’s flexible datacenters, according to a message from a remote controller. Ex.1007, [0017],

[0025]-[0026], [0076]. A POSITA would have understood that Kiani-Pelio's flexible datacenters co-located BTM with a generation station would have times when there would be a reduction in power output (e.g., solar or wind conditions changing), causing an overload condition unless power is received from the grid or load is reduced. A POSITA would have thus implemented Chapel's teachings for reducing power delivery to one or more datacenter servers to reduce load and avoid an overload condition. Ex.1003, ¶¶249-50.

7. Claim 8

First, the Kiani-Pelio-Chapel combination renders obvious modulating power delivery based on an “*operational directive*.” A remote controller (e.g., Kiani's FE server) sends a message to the local controller to modulate the power distribution to servers. *See* [7.1]; Ex.1007, [0076]; *see also* [0058], [0066], [0090]; Ex.1005, p. 2; Ex.1007, [0056], [0083]; Ex.1003, ¶¶251-53.


Second, the combined Kiani-Pelio-Chapel teachings further render obvious monitoring BTM power availability at the servers within Kiani's datacenters (e.g., flexible datacenters per Pelio's teachings). *See* [5.1]; Ex.1005, p. 2; Ex.1006, Abstract, [0049], [0071]; Figure 2; Ex.1007, [0004], [0047], [0056], [0082]-[0083]. In combination, Kiani-Pelio-Chapel teaches Kiani's FE server monitoring BTM power availability according to Chapel's teachings, and providing a message to locally control a flexible datacenter depending on the amount of BTM power

available to the servers. Ex.1003, ¶¶254-56.

D. Ground 3: Claim 3 Obvious Over Kiani-Pelio-Belady

1. Belady

Belady describes “plural data centers powered by plural power sources,” including off-grid and/or on-grid renewables. Ex.1008, 2:6-14. Different datacenters are located at different geographic locations and connected with different power sources, including co-location to avoid transmission limitations. Ex.1008, 10:14-34; 6:53-7:17; FIG. 1:



Ex.1008, Fig. 1 (annotated).

Instead of moving power, Belady temporally and/or spatially moves computational tasks between datacenters based on power availability, and potentially also for cost-effectiveness. Ex.1008, 7:49-60 (temporally), 7:61-8:8 (spatially), 2:14-22 (cost). Temporal movement adjusts a time of task execution, while spatial movement transfers a task's execution from one datacenter to another. Ex.1008, 2:23-26, 2:36-41, 10:20-32. The management system uses “real-time information” to determine power availability and needs, and direct datacenter operations. Ex.1008, 12:41-47 (real-time), 12:48-53 (predicted); Ex.1008, 10:60-62 (manage execution of computational tasks); Ex.1003, ¶¶257-61.

2. Reasons to combine Pelio, Chapel, and Belady

First, Belady, like Kiani and Pelio, focuses on systems for controlling the use of BTM power. Belady powers different datacenters at different locations with different power sources, including renewable power that is behind-the-meter with a datacenter. Ex.1008, 2:6-14, 6:53-7:17. Just like Kiani, Pelio and the '261 patent, Belady teaches co-locating computing resources at power generating stations to harness behind-the-meter power. Thus, Belady is in the same field of endeavor. Ex.1003, ¶¶262-63.

Further, Belady is pertinent to problems such as “the exponential increase in power required” for datacenter operations, wasted energy generated from renewable energy sources, and grid transmission and distribution costs. *See*

Ex.1001, 35:11-15. Belady teaches co-locating datacenters with renewable power sources to reduce infrastructure costs and transmission losses. Ex.1008, 7:1-9. Belady also teaches moving computational tasks temporally, spatially, or both to harness power availability and cost-effectiveness. Ex.1008, 2:15-22. Thus, a POSITA would have also found Belady's teachings to be reasonably pertinent. Ex.1003, ¶¶264.

Second, a POSITA would have been motivated to turn from Kiani to Belady for the teachings about implementing management functionality (of which Kiani's FE server functions are examples) on top of the datacenter software at a datacenter itself. *See* Ex.1008, 8:23-26. The location of Kiani's FE server control functionality was a mere design choice, one of a limited number of options including remote placement (like in Kiani and Belady), local placement at a datacenter (like in Belady), or split between a remote location and a datacenter (like in Belady). *See* Ex.1008, 8:15-31. A POSITA would have been motivated to implement Belady's teaching of control at the datacenter to benefit from already-existing hardware and infrastructure. Since all of these options had already demonstrated success in implementation as Belady indicates, a POSITA would have therefore had a reasonable expectation of success in the combination of Belady's local implementation teachings with Kiani's FE server teachings. Ex.1003, ¶265.

Third, a POSITA would have further been motivated to turn from Kiani and Pelio for Belady’s teachings of moving computational tasks due to power availability, and how to detect and determine that it is appropriate to move. In addition to geographical load balancing to initially distribute computational operations, Kiani teaches that “load shifting algorithms” shift the load after initial distribution. Ex.1005, p. 2. Further, Pelio already recognized the desirability of moving pending compute jobs when there is insufficient renewable power. Ex.1006, [0106]; *see also* Ex.1013, p. 1 (shift computational load between datacenters due to lack of power). Belady provides implementation details regarding Kiani’s and Pelio’s suggestions. Ex.1003, ¶266.

For example, Belady’s management system (e.g., Kiani’s FE server teachings implemented at a datacenter per Belady’s teachings) manages task (like Kiani’s “requests,” Pelio’s “data center jobs”) execution amongst multiple datacenters based on variable power availability. Ex.1008, 2:15-22. This includes moving computation tasks from one datacenter powered by a solar generating source in one location during nighttime to another datacenter powered by another solar generating source geographically in daytime. Ex.1008, 10:14-34. A POSITA would have been motivated to turn to Belady’s teachings for implementation details load shifting/job movement. Ex.1003, ¶267.

Fourth, POSITA would have had a reasonable expectation of success applying Belady's teachings to the Kiani-Pelio teachings. Like Kiani and Pelio, Belady teaches collocating the datacenter with the generating source. Ex.1008, 6:53-7:17; *see also* Ex.1005, pp. 2-3, Ex.1006, [0002]. Kiani, Pelio, and Belady all recognize the utility of moving computing jobs/tasks between datacenters for cost advantage (Kiani) and/or in response to power availability issues (Pelio, Belady). Ex.1005, p. 2; Ex.1006, [0106]; Ex.1008, 2:14-22. A POSITA would have therefore recognized the compatibility of Kiani's, Pelio's and Belady's systems, with a reasonable expectation of success implementing Belady's teachings for moving tasks between datacenters – including by the management system (FE server functionality) at a datacenter instead of remote thereto. Ex.1003, ¶268.

In combination, Kiani's datacenters would implement Belady's real-time monitoring and computational task movement teachings by a control system (e.g., Kiani's FE server, and Belady's management system teachings). Doing so would have been predictable and within a POSITA's skillset. Kiani and Pelio recognized the desirability of moving jobs due to power availability issues, and Belady teaches ways to do so. And it was already well-known to use behind-the-meter renewable power sources for the datacenters' tasks. Ex.1005, pp. 2-3; Ex.1006, Abstract, [0002]; Ex.1012, pp. 1, 4; Ex.1014, p. 7; Ex.1013, p. 6; Ex.1015, pp. 3, 4; Ex.1003, ¶269.

3. Claim 3

First, Kiani teaches that the “*routing control system is a remote master control system.*” See [2.1] (Ground 1); Ex.1005, p. 2; Ex.1003, ¶¶271-73.

Second, the combination of Kiani and Pelio also renders obvious the “*flexible datacenter.*” See [1.1] (Ground 1); Ex.1003, ¶274.

Third, placement of the control system (e.g., Kiani’s FE server) is a mere design choice, with options including remote placement from the datacenter, at the datacenter, or split between remote and local systems. Ex.1003, ¶275. Kiani provides an example with the FE server positioned remote from the datacenters. See [2.1]. It would have been obvious to place the control system at the datacenter instead, as Belady demonstrates. Ex.1003, ¶275.

Belady teaches a “management system” that controls how “a plurality of computational tasks” are executed by multiple datacenters, including movement of the tasks between them based on power availability and/or cost. Ex.1008, 2:14-22. Belady teaches implementing the management system 120 – analogous to Kiani’s FE server – “by the processing equipment provided by the individual data centers.” Ex.1008, 12:20-30 (management system implements management functionality); 8:23-26; *see also* 8:15-20, 8:26-31 (alternatively placing the control system remotely, or split between remote and local systems); Ex.1003, ¶276.

In combination, it would have been obvious for Kiani’s FE server to be

implemented by processing equipment at the datacenter. Kiani’s original concept of datacenters was to have “powerful data centers” deployed away from users as just an example. Ex.1005, p. 2. But Kiani does not preclude the possibility of datacenters – “Internet-scale” or not – closer to users. Doing so was one of just a few options that were obvious to try. Accordingly, implementing an FE server at the datacenter itself (per Belady’s teachings) would result in the FE server being close to users while also being at the datacenter. Ex.1003, ¶277.

Moreover, Kiani-Pelio’s combined teachings provide flexible, highly-portable datacenters-on-wheels, resulting in mobile datacenters placeable in multiple locations. *See* Ex.1005, p. 2 (placing datacenters to achieve desired characteristics); Ex.1006, [0055]-[0060] (modular, containerized datacenters). With Pelio’s flexible datacenter teachings, there would beneficially be more flexibility in datacenter location options, including coincident with Kiani’s FE server locations. In the combination, Kiani’s FE server functions would remain the same, since even with FE server/datacenter co-location differences may still arise in “server or content availability ... efficiency of the data centers, the cost of the electricity, and availability of renewable energy” for different datacenters. Ex.1005, p. 2. As a result, the FE server would still perform its workload distribution function. Ex.1003, ¶¶278-79.

E. Ground 4: Claims 13-16 Obvious Over Kiani-Pelio-Chapel-Forestiero

1. Forestiero

Forestiero uses a “hierarchical approach for workload management in geographically distributed [datacenters, or DCs].” Ex.1009, p. 1 (Abstract). This includes a lower layer for site-specific internal workload management, and an upper layer implementing a shared algorithm amongst the sites “used to evaluate the behavior of single sites and distribute the workload among them.” Ex.1009, pp. 1-2; Figure 1:




Fig. 1. EcoMultiCloud scenario: upper and lower layer of two interconnected data centers.

Ex.1009, Fig. 1 (p. 4).

Forestiero’s approach decentralizes the intelligence and decision points. Ex.1009, p. 3. Each datacenter includes a “data center manager” (DCM)

implementing the upper layer, and a local manager (LM) implementing the lower layer, typically “deployed on the same host.” Ex.1009, p. 4. The DCM acquires its datacenter’s current state information, circulates that information amongst the DCMs at interconnected datacenters, with which the DCMs collectively decide how to distribute the workload. Ex.1009, p. 4. The information distributed between datacenters “is tiny, i.e., of the order of a few packets.” Ex.1009, p. 4; Ex.1003, ¶¶283-86.

2. Reasons to Combine Kiani, Pelio, Chapel, and Forestiero

First, Forestiero teaches how to handle workload in geo-distributed datacenter systems, with a DCM at each datacenter communicating with other datacenters’ DCMs to collectively choose a target datacenter for a workload, and a LM at each datacenter to manage assigned workload. Ex.1009, p. 4. Forestiero’s approaches accommodate multi-datacenter scenarios that balance multiple business objectives/constraints and adapt workload distribution to varying conditions. Ex.1009, p. 4. Forestiero’s teachings enable datacenters to manage their workload for better performance and cost. These issues are pertinent to the ’261 patent, which identifies the different costs of different locations and power types as a workload management problem. *See, e.g.*, Ex.1001, 6:14-17 (different costs of operations), 6:40-42 (reduce cost of datacenter use). A POSITA would have therefore found Forestiero reasonably pertinent. Ex.1003, ¶¶287-88.

Second, a POSITA would have been motivated to turn to Forestiero's geo-distributed workload management teachings to harness Forestiero's benefits of decentralized workload management. Implementing the decentralized decision intelligence via a DCM at each datacenter would beneficially remove the limitations of centralized approaches, including poor scalability, poor adaptability to changing conditions, and lack of autonomy. *See* Ex.1009, p. 3. Kiani's geographical load balancing teachings would result in the FE servers reducing their decision-making burden to fielding new requests, with the distribution decision logic implemented by DCMs at the different datacenters. *See* Ex.1005, p. 2 (FE servers receiving user requests); Ex.1009, p. 4 (DCMs collectively choosing the target datacenter for a user request); Ex.1003, ¶289.

Workload management is handled in a limited number of ways, including centralized, decentralized, or split between the two. Such approaches were already well-recognized. *See, e.g.*, Ex.1009, p. 3 (comparing centralized and decentralized approaches); Ex.1008, 8:23-26 (distributed management system functionality), 8:15-23, 8:26-31 (remote or split implementation). Thus, it would have been obvious to try implementing the load distribution decision-making in a decentralized manner, such Forestiero teaches. Forestiero's distributed decision-making teachings beneficially enhance these benefits by providing Kiani's system with a hierarchical architecture that considers both system-wide and local issues

when determining how to allocate/re-allocate workload. Ex.1009, pp. 3, 4. The combination further beneficially exploits “the centralized architecture and functionalities” of Kiani’s system while decentralizing the decision-making for autonomy and increased adaptability to changing conditions. Ex.1009, p. 3; Ex.1003, ¶¶290-91.

Third, a POSITA would have had a reasonable expectation of success combining Forestiero’s and Kiani’s teachings. Forestiero’s distribution of the workload distribution decision to a controller at each datacenter (e.g., in a collective decision-making process with the other datacenter controllers) is nothing more than a natural extension of Kiani’s recognition of the benefits of distributing decision-making (including the ability to adopt “different workload distribution strategies” at each FE server that consider different objectives at each site). Ex.1005, p. 2; *see* Ex.1009, p. 4 (distributed decision-making at each datacenter in a collective fashion); p. 3 (improved adaptability to changing conditions locally); Ex.1003, ¶292.

In combination, it would have been obvious for Kiani’s FE servers to receive user workload requests close to the users, and the decision-making of final distribution to occur in a distributed yet cooperative manner amongst the different, geographically-dispersed datacenters (per Forestiero’s teachings). The FE servers would contribute “high level information” as part of Forestiero’s disclosed

information exchange amongst DCMs, either with the FE servers operating as effectively peers to the DCMs, or in similar capacity as Forestiero's LMs. Ex.1003, ¶293.

Forestiero supports this understanding. Each DCM communicates with its LM to obtain state information of the local datacenter, which is then part of the information exchange and decision-making process collectively by the DCMs. *See* Ex.1009, p. 4. It would have been a trivial detail to implement Kiani's FE servers as information sources like Forestiero's LMs, since Forestiero's teachings are more about exchanging and using the information than about what the particular information sources are. Indeed, Forestiero's LM does not have to be deployed in the same host as the DCM, which demonstrates that other locations (and communication between the two) are also possible – like Kiani's FE server locations. *See* Ex.1009, p. 4 (DCM and LM deployment). Ex.1003, ¶294.

Forestiero's DCM information-sharing teachings further support the understanding. Each DCM shares "high level information" with all of the other DCMs. Ex.1009, p. 4. It would have been obvious for Kiani's FE servers to be a peer source of "high level information" for all of the DCMs, without the FE servers otherwise being part of the decision-making process. With the FE servers sharing information (such as green energy utilization, electricity cost minimization, profit maximization, etc. as taught by Kiani, p. 2), controllers at datacenters would be

able to make more efficient decisions according to Forestiero's distributed decision-making teachings. Ex.1003, ¶295.

A POSITA would have therefore found the combination of Kiani's and Forestiero's teachings predictable and within a POSITA's skillset. Specifically, a POSITA would have expected the combination of Kiani's system teachings and Forestiero's distributed decision-making teachings to result in Kiani's FE servers sharing their high-level information with the DCMs to contribute to Forestiero's collective decision-making teachings. This would provide Kiani's system with the benefits of a hierarchical architecture, multiple-objective and multiple-constraint balancing among locations, and increased adaptability between locations. *See* Ex.1009, pp. 3, 4; Ex.1003, ¶296.

The combination would have therefore been nothing more than the combination of known elements (Kiani's distributed FE server decision-making and Forestiero's decentralized information sharing and decision-making) according to known methods (using Forestiero's distributed decision-making at the different datacenters with Kiani's different datacenters and different FE servers), to yield predictable results (Kiani's FE servers fielding user requests with Forestiero's distributed decision-making at the datacenters taking into account information shared from all the datacenters and FE servers). Ex.1003, ¶297.

3. Claim 13

[13.0]

See [1.1].

[13.1]

See [1.2].

[13.2]

See [1.3].

[13.3]

See [1.7] and [4.1].

[13.4]

See [1.4].

[13.5]

First, Kiani discloses a queue subsystem at each datacenter that receives workload assignments from an FE server. *See* [1.5]; Ex.1005, p. 5; Ex.1003, ¶¶299-305.

Second, the queue at each datacenter, with or without the FE servers' queues, demonstrates that the “*queue subsystem*” at a given datacenter “*is part of a distributed queue system.*” *See* [1.5]; Ex.1003, ¶306.

Third, each datacenter's queue organizes computational operations at that datacenter for execution (“*performance*”) by servers at that datacenter. *See* [1.6];

Ex.1005, pp. 3-5 (managing workload according to green energy availability at a datacenter, maintaining the queue's upper bound, and QoS requirements); Ex.1003, ¶¶307-08.

[13.6]

First, Kiani discloses a control system that determines, using the queue system, a computational operation to be performed. *See* [1.8]. The FE server receives user requests and distributes the requests to servers at geo-dispersed datacenters. Ex.1005, p. 2. The incoming workload to the FE server teaches a “*computational operation to be performed.*” Ex.1005, pp. 2-3. Because the FE server knows about the workload to distribute, the FE server has “*determine[d], using the queue system,*” the computational operation to be performed after distribution to a datacenter. Ex.1003, ¶¶309-10.

Second, Chapel's local controller controls power distribution and management within a datacenter (and thus is a “*datacenter control system*”). *See* [13.3], [4.1] and [7.1]; Ex.1007, [0002]; Figure 2:




FIG.3

Ex.1007, Fig. 3 (annotated).

With a local controller at each datacenter in Kiani-Chapel’s combined teachings, it would have been obvious to implement workflow allocation control to the datacenters’ local controllers. This was just one of a handful of ways to implement the determination of a computational operation to be performed, including a centralized approach (e.g., Kiani’s FE server assigning workload) and a decentralized approach with local controllers. Ex.1003, ¶¶311-13.

Forestiero decentralizes allocation of workload (e.g., represented by virtual machine allocation per job request, Ex.1009, pp. 1, 2) in geo-distributed datacenter systems. A local manager (LM) handles workload management within that datacenter after the workload is assigned there by the upper layer. Ex.1009, p. 4. The datacenter manager (DCM) at each datacenter “runs the algorithms of the

upper layer.” Ex.1009, p. 4; *see also* Figure 1:




Fig. 1. EcoMultiCloud scenario: upper and lower layer of two interconnected data centers.

Ex.1009, Fig. 1 (annotated).

A DCM – e.g. at datacenter 1 above – communicates with the local LM for the datacenter’s current state (and extracts high-level information from it), transmits/receives high-level information to/from all the other DCMs at other datacenters, combines the collected information, and makes decisions about the distribution of the workload among the datacenters. Ex.1009, p. 4. Forestiero teaches that “all the DCMs execute the upper layer algorithms and ... choose the target [datacenter]” for a workload. Ex.1009, p. 4. The determination is based on the information received locally and from other datacenter sites (see Figure 2):




Fig. 2. EcoMultiCloud scenario: the DCMs of four data centers exchange high level information about the state of local data centers. Such information is used, for example, to decide which site should accommodate a new VM.

Ex.1009, Fig. 2 (annotated).

In combination, it would have been obvious for Kiani’s geographically distributed datacenters to include DCM functionality within a local controller (as taught by Chapel). See Ex.1009, p. 3; Ex.1003, ¶¶314-17. Kiani’s FE servers would receive incoming workload requests from users, with the decision intelligence decentralized to local controllers for workload distribution and management. Indeed, Forestiero teaches deploying the DCM “on the same host” as other local control entities at the datacenter, which in the combination includes Chapel’s local controller. See Ex.1009, p. 4. Thus, the combination teaches the “*datacenter control system*” (the DCM at a local controller, such as taught by the

Chapel-Forestiero combination) determining a computational operation to be performed (the workload). And, while Forestiero's Figure 2 exemplary embodiment provides a whole mesh example, Forestiero explains that other network architectures are also possible to accommodate a larger number of datacenters. *See* Ex.1009, pp. 4-5. Thus, the combination would work with Kiani's system. Ex.1003, ¶¶317-19.

[13.7]

First, Kiani discloses, and combined with Pelio renders obvious, a control system configured to determine an amount of BTM power that is available to the plurality of computing systems, via the FE server (a “*routing control system*”). *See* [1.9]. Further, the FE server determines how much “green energy” is available to a datacenter, where “green energy” includes energy produced with the datacenter “behind the meter.” *See* Ex.1005, pp. 2-3; Ex.1003, ¶¶320-21.

Second, it would have been obvious to implement the control in a decentralized manner by a datacenter control system (e.g., Chapel's local controller), according to Forestiero's decentralized DCM teachings. *See* [13.6]. In combination, Kiani's FE server would field user requests and distribute those requests according to the determination by Forestiero's DCMs (datacenter control systems, e.g. Chapel's local controllers). *See, e.g.*, Ex.1009, Figure 2; Ex.1003, ¶322.

In combination, the information collected by Forestiero's DCMs about datacenter conditions would include the availability of "green energy" per Kiani's teachings, and thus BTM power as taught by both Kiani and Pelio. Ex.1003, ¶¶323-24.

[13.8]

Kiani-Pelio renders obvious determining to route the computational operation to the flexible datacenter using the amount of BTM power available. *See* [1.10]. It would have been further obvious to implement the distribution functionality decentralized at the datacenters, instead of by centralized controller, per Forestiero. *See* [13.6]; Ex.1003, ¶¶325-28.


[13.9]

First, Kiani-Pelio renders obvious assigning the computational operation for performance at the flexible datacenter based on a determination to route the computational operation there. *See* [1.11]; Ex.1003, ¶¶329-30.

Second, it would have been further obvious to implement the distribution functionality, including assigning the workload, decentralized at the datacenters per Forestiero. *See* [13.6]. For example, Forestiero teaches that the DCMs (implemented by local controllers at Kiani's datacenters) make determinations about workload distribution, including the assignment of the workload, amongst the datacenters. Ex.1009, p. 4; Ex.1003, ¶¶331-32.

4. Claim 14

First, as demonstrated at [1.5], Kiani discloses at least two queue subsystems with a queue at each datacenter:



Second, Kiani discloses organizing computational operations at each datacenter in each datacenter's queue according to rules. *See* [1.6]; Ex.1005, pp. 3, 5. A “**QoS at each data center** is enforced by maintaining an upper bound” on the queue length of that datacenter. Ex.1005, p. 4. This is a rule associated with the queue subsystem at that datacenter. The upper bound is another rule for organizing each queue's allocated workload (“*computational operations*”) for execution, and “is **varied at different data centers** due to different service efficiency parameters.” Ex.1005, p. 5; *see also* p. 4. Accordingly, Kiani discloses the different datacenters having their own respective queues (“*at least two queue subsystems*”),

each with an “upper bound” (“*different rules for organizing computational operations*”) based on the unique parameters of each datacenter. Ex.1003, ¶¶333-37.

And, as further noted at [1.6], Kiani’s variety of rules are determined to facilitate accomplishing the targeted QoS for the workload at each datacenter, including how to maximize green energy utilization, how to apply brown energy (and minimize its cost), and maintain QoS with brown energy. *See* Ex.1005, p. 4. These other rules also influence the organizing of workload (“*computational operations*”) and are unique at each datacenter. Ex.1003, ¶¶338-40.

5. Claim 15

See [7.1]; Ex.1003, ¶¶341-42.

6. Claim 16

Kiani teaches a “*remote master control system*” (FE server) positioned remotely from the datacenters, which are flexible datacenters alone and combined with Pelio. *See* [2.1]. This alone renders obvious claim 16. Ex.1003, ¶¶343-45.

Chapel further teaches a local controller (a “*datacenter control system*”) receiving messages from a remote controller (a “*remote master control system*”) that identifies a power delivery modulation action, such as transitioning to an efficiency or brownout mode (an “*operational directive*”). *See* Ex.1007, [0076]; *see also* [0058], [0066], [0078]. Figures 3 and 6 of Chapel illustrate the local

controller receiving the one or more messages from a remote controller. *See* Ex.1007, [0056] (Figure 3); [0079]-[0082] (Figure 6); Ex.1003, ¶¶346-52.

IX. DISCRETIONARY DENIAL IS INAPPROPRIATE

A. No Basis for 35 U.S.C. § 325(d) Denial

None of Kiani, Pelio, Chapel, Belady, or Forestiero were cited or considered by the Examiner during prosecution of the '261 patent. *See* Ex.1002; Ex.1001, p. 1 (56). In contrast to the art considered by the Examiner, the prior art presented in this petition renders obvious all of the limitations from the claims. Accordingly, discretionary denial under 35 U.S.C. § 325(d) is not appropriate.

B. Discretionary Denial Under the *Fintiv* Factors is Not Appropriate

There is no copending litigation, and therefore the six factors considered for § 314 denial strongly favor institution. *See Apple Inc. v. Fintiv, Inc.*, IPR2020-00019, Paper 11 (Mar. 20, 2020).

X. CONCLUSION

Petitioners have established a reasonable likelihood that the Challenged Claims are unpatentable.

Respectfully submitted,

Dated: August 30, 2024
HAYNES AND BOONE, LLP
2801 N. Harwood St., Suite 2300
Dallas, Texas 75201
Customer No. 27683

/David L. McCombs/
David L. McCombs
Lead Counsel for Petitioners
Registration No. 32,271

XI. MANDATORY NOTICES

A. Real Party-in-Interest

Pursuant to 37 C.F.R. § 42.8(b)(1), Petitioners certify that the real parties-in-interest are Cholla Energy LLC and OBM, Inc.

B. Related Matters

Pursuant to 37 C.F.R. § 42.8(b)(2), to the best knowledge of Petitioners, the '261 patent is not involved in any pending cases, but was previously asserted in:

Case Heading	Number	Court	Filed
<i>Lancium LLC v. U.S. Data Mining Group, Inc. et al.</i>	6-23-cv-00344	W.D. Tex.	May 10, 2023

C. Lead and back-up counsel and service information

Lead Counsel

David L. McCombs
HAYNES AND BOONE, LLP
2801 N. Harwood St., Suite 2300
Dallas, TX 75201

Phone: 214-651-5533
Fax: 214-200-0853
david.mccombs.ipr@haynesboone.com
USPTO Reg. No. 32,271

Back-up Counsel

Adam C. Fowles
HAYNES AND BOONE, LLP
2801 N. Harwood St., Suite 2300
Dallas, TX 75201

Phone: 972-739-8674
Fax: 214-200-0853
adam.fowles.ipr@haynesboone.com
USPTO Reg. No. 65,005

Rodney R. Miller
WOMBLE BOND DICKINSON (US),
LLP
271 17th St. N.W., Suite 2400
Atlanta, GA 30363

Phone: 404-879-2435
Fax: 404-888-7490
Rodney.Miller@wbd-us.com
USPTO Reg. No. 66,073

Preston H. Heard
WOMBLE BOND DICKINSON (US),
LLP
271 17th St. N.W., Suite 2400
Atlanta, GA 30363

Phone: 404-888-7366
Fax: 404-888-7490
Preston.Heard@wbd-us.com
USPTO Reg. No. 64,675

James E. Dority
WOMBLE BOND DICKINSON (US),
LLP
550 S. Main St., Suite 400
Greenville, SC 29601

Phone: 864-255-5412
Fax: 404-888-7490
James.Dority@wbd-us.com
USPTO Reg. No. 77,313

Please address all correspondence in this proceeding to lead and back-up counsel. Petitioners consent to service in this proceeding by email at the addresses above.

XII. APPENDIX OF CHALLENGED CLAIMS

[1.0] A system comprising:

[1.1] a flexible datacenter that is positioned to receive electrical power directly from a power generation system without the received power flowing through a grid comprising:

[1.2] a behind-the-meter (“BTM”) power input system configured to receive power from the power generation system prior to the power undergoing step-up transformation for transmission to the grid and provide the received power to the flexible datacenter instead of sending the received power to the grid;

[1.3] a power distribution system; and

[1.4] a plurality of computing systems powered by the BTM power input system via the power distribution system;

[1.5] a queue system configured to organize a plurality of computational operations, wherein the queue system comprises multiple queue subsystems, and

[1.6] wherein each queue subsystem organizes a set of computational operations according to rules associated with the queue subsystem; and

[1.7] a routing control system configured to:

[1.8] determine, using the queue system, a computational operation to be performed;

[1.9] determine an amount of BTM power that is available to the plurality of

computing systems;

[1.10] [a routing control system configured to:] determine whether to route the computational operation to the flexible datacenter using the amount of BTM power that is available; and

[1.11] [a routing control system configured to:] based on a determination to route the computational operation to the flexible datacenter, assign the computational operation for performance at the flexible datacenter.

[2.0] The system of claim 1,

[2.1] wherein the routing control system is a remote master control system positioned remotely from the flexible datacenter.

[3.0] The system of claim 1,

[3.1] wherein the routing control system is a remote master control system positioned at the flexible datacenter.

[4.0] The system of claim 1, wherein the flexible datacenter further comprises:

[4.1] a datacenter control system, wherein the datacenter control system is configured to modulate power delivery to the plurality of computing systems based on one or more monitored power system conditions.

[5.0] The system of claim 4, wherein the one or more monitored power system conditions comprises:

[5.1] behind-the-meter (“BTM”) power availability at the plurality of computing systems.

[6.0] The system of claim 4, wherein the one or more monitored power system conditions comprises:

[6.1] one or more of excess local power generation at a local station level, excess local power generation that a grid cannot receive, local power generation subject to economic curtailment, local power generation subject to reliability curtailment, local power generation subject to power factor correction, low local power generation, start up local power generation situations, transient local power generation situations, or testing local power generation situations where there is an economic advantage to using local behind-the-meter power generation.

[7.0] The system of claim 1, wherein the flexible datacenter further comprises:

[7.1] a datacenter control system, wherein the datacenter control system is configured to modulate power delivery to the plurality of computing systems based on an operational directive.

[8.0] The system of claim 7,

[8.1] wherein the operational directive originates from the routing control system and depends on behind-the-meter (“BTM”) power availability at the plurality of computing systems.

[9.0] The system of claim 1,

[9.1] wherein the routing control system controls distribution of the plurality of computational operations from the queue system.

[10.0] The system of claim 1,

[10.1] wherein a first queue subsystem of the queue system is located at the flexible datacenter.

[11.0] The system of claim 10, further comprising:

[11.1] a second datacenter, wherein a second queue subsystem of the queue system is located at the second datacenter.

[12.0] The system of claim 11,

[12.1] wherein the second datacenter is a second flexible datacenter.

[13.0] A flexible datacenter comprising:

[13.1] a behind-the-meter (“BTM”) power input system configured to receive power from a power generation system prior to the power undergoing step-up transformation for transmission to a grid and provide the received power to the flexible datacenter instead of sending the received power to the grid;

[13.2] a power distribution system;

[13.3] a datacenter control system;

[13.4] a plurality of computing systems, wherein the plurality of computing systems are powered by the BTM power input system via the power distribution

system; and

[13.5] a queue subsystem configured to organize a set of computational operations for performance by the plurality of computing systems, wherein the queue subsystem is part of a distributed queue system;

[13.6] wherein the datacenter control system is configured to: determine, using the queue system, a computational operation to be performed;

[13.7] determine an amount of BTM power that is available to the plurality of computing systems;

[13.8] determine whether to route the computational operation to the flexible datacenter using the amount of BTM power that is available; and

[13.9] based on a determination to route the computational operation to the flexible datacenter, assign the computational operation for performance at the flexible datacenter.

[14.0] The flexible datacenter of claim 13,

[14.1] wherein the distributed queue system includes at least two queue subsystems having different rules for organizing computational operations.

[15.0] The flexible datacenter of claim 13,

[15.1] wherein the datacenter control system is configured to modulate power delivery to the plurality of computing systems based on an operational directive received from a remote master control system.

[16.0] The flexible datacenter of claim 15,

[16.1] wherein the remote master control system is positioned remotely from the flexible datacenter.

CERTIFICATE OF WORD COUNT

Pursuant to 37 C.F.R. § 42.24(d), Petitioners hereby certify, in accordance with and reliance on the word count provided by the word-processing system used to prepare this Petition, that the number of words in this paper is 13,903. Pursuant to 37 C.F.R. § 42.24(d), this word count excludes the table of contents, table of authorities, mandatory notices under § 42.8, certificate of service, certificate of word count, appendix of exhibits, and any claim listing.

Dated: August 30, 2024

/David L. McCombs/
David L. McCombs
Lead Counsel for Petitioners
Registration No. 32,271

CERTIFICATE OF SERVICE

The undersigned certifies that, in accordance with 37 C.F.R. § 42.6(e) and 37 C.F.R. § 42.105, service was made on Patent Owner as detailed below.

Date of service August 30, 2024

Manner of service Priority Mail Express

Documents served Petition for *Inter Partes* Review Under 35 U.S.C. § 312 and 37 C.F.R. § 42.104 of U.S. 11,283,261; Petitioners' Exhibit List; Exhibits 1001-1023.

Persons served Barnes & Thornburg LLP
P.O. Box 2786
Chicago, IL 60690-2786

/David L. McCombs/

David L. McCombs
Lead Counsel for Petitioners
Registration No. 32,271