

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CHOLLA ENERGY LLC AND OBM, INC.,
Petitioners

IPR2023-00914
U.S. Patent No. 11,016,456

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

TABLE OF CONTENTS

Petitioners’ Exhibit List	6
I. Introduction.....	8
II. Grounds for standing	8
III. Note.....	9
IV. OVERVIEW OF THE TECHNOLOGY.....	9
A. Background	9
B. Grid Stability with Renewables.....	9
C. Harnessing Curtailed Power with Green Datacenters.....	10
D. The ’456 Patent	11
E. Prosecution History	13
F. Effective Priority Date of the ’456 Patent.....	13
V. Level of ordinary skill in the art	14
VI. Claim construction.....	14
VII. Relief requested and reasons Therefor	14
VIII. Identification of how the claims are unpatentable.....	15
A. Grounds	15
B. Ground 1: Claims 1-4, 14, 17-18, 20, 25, 27 Obvious Over Pelio- Chapel.....	17
1. Pelio.....	17
2. Chapel	19

3.	Reasons to combine Pelio and Chapel.....	20
4.	Claim 1	26
5.	Claim 2.....	39
6.	Claim 3.....	41
7.	Claim 4.....	42
8.	Claim 14.....	44
9.	Claim 17.....	45
10.	Claim 18.....	46
11.	Claim 20.....	47
12.	Claim 25.....	48
13.	Claim 27.....	54
C.	Ground 2: Claims 5-8, 10, 12, 28-30 Obvious Over Pelio-Chapel- Yang	55
1.	Yang.....	55
2.	Reasons to combine Pelio and Yang.....	56
3.	Claim 5.....	59
4.	Claim 6.....	63
5.	Claim 7.....	63
6.	Claim 8.....	64
7.	Claim 10.....	64
8.	Claim 12.....	64

9.	Claim 28	64
10.	Claim 29	70
11.	Claim 30	72
D.	Ground 3: Claims 19, 21-24, and 26 Obvious Over Pelio-Chapel-Belady.....	72
1.	Belady	72
2.	Reasons to combine Pelio, Chapel, and Belady	74
3.	Claim 19	77
4.	Claim 21	79
5.	Claim 22	82
6.	Claim 23	85
7.	Claim 24	86
8.	Claim 26	87
E.	Ground 4: Claim 9 Obvious Over Pelio, Chapel, and Rasmussen.....	87
1.	Rasmussen.....	87
2.	Reasons to Combine Pelio, Chapel, and Rasmussen.....	88
3.	Claim 9	90
F.	Ground 5: Claim 9 Obvious Over Pelio, Chapel, and Rasmussen.....	91
1.	Corhodzic	91
2.	Reasons to Combine Pelio, Chapel, and Corhodzic	92
3.	Claim 15	94

4.	Claim 16	95
IX.	Discretionary denial is inappropriate.....	95
A.	No basis for 35 U.S.C. § 325(d) denial	95
B.	Discretionary denial under the <i>Fintiv</i> factors is not appropriate.....	96
X.	Conclusion	96
XI.	Mandatory notices	97
A.	Real party-in-interest.....	97
B.	Related matters	97
C.	Lead and back-up counsel and service information.....	97
XII.	Appendix of Challenged Claims.....	99
	Certificate of Word Count	107
	Certificate of Service	108

PETITIONERS' EXHIBIT LIST

Ex.1001	U.S. 11,016,456
Ex.1002	Prosecution History of U.S. 11,016,456
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	U.S. Pub. 2011/0316337 by Pelio et al.
Ex.1006	U.S. Pub. 2010/0145542 by Chapel et al.
Ex.1007	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.1008	U.S. 8,849,469 to Belady et al.
Ex.1009	RESERVED
Ex.1010	Rasmussen, "Impact of Leading Power Factor on Data Center Generator Systems," Schneider Electric, 2014
Ex.1011	U.S. 8,080,900 by Corhodzic et al.
Ex.1012	Martinot, E., "Grid Integration of Renewable Energy: Flexibility, Innovation, and Experience," Annu. Rev. Environ. Resour., vol. 41:223-251, 2016
Ex.1013	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 IEEE Xplore Date of Publication
Ex.1014	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.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	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.1020	U.S. Pub. 2010/0088261 by Montalvo
Ex.1021	IEEE Xplore, https://ieeexplore.ieee.org/document/7516102
Ex.1022	Wayback Machine capture of www.apc.com white papers index, https://web.archive.org/web/20141222215407/http://www.apc.com/prod_docs/results.cfm?DocType=White+Paper&Query_Type=10
Ex.1023	Zhabelova et al., “Data Center Power Dynamics Within the Settings of Regional Power Grid,” 2015 IEEE 20 th Conference on Emerging Technologies & Factory Automation (ETFA), September 8-11, 2015
Ex.1024	Metadata from Rasmussen (Ex.1010)
Ex.1025	Schneider Buys APC for \$6.1 Billion, wayback machine archive of October 30, 2006 report, https://web.archive.org/web/20090606055449/https://www.datacenterknowledge.com/archives/2006/10/30/schneider-buys-apc-for-61-billion
Ex.1026	Affidavit of Nathaniel E. Frank-White of the Internet Archive (including Exhibit A with exhibits 1025 and 1022, in that order)

I. INTRODUCTION

Cholla Energy LLC and OBM, Inc. (“Petitioners”) respectfully request that the Board review and cancel as unpatentable under 35 U.S.C. §103 claims 1-10, 12, and 14-30 (the “Challenged Claims”) of U.S. 11,016,456 (the “’456 patent,” Ex.1001).

The ’456 patent describes placing a flexible datacenter “behind-the-meter,” where the datacenter receives power before the power undergoes step-up transformation for use on a power grid. This includes a “datacenter control system” that is used to “modulate power delivery” to computing systems based on an operational directive that could be based on power availability, a command from a utility, etc. But placing a flexible datacenter behind-the-meter was already well-known and well-discussed long before the ’456 patent was filed, including strategies on how to manage power within the datacenter based on fluctuating power availability. *See, e.g.*, Ex.1005, Ex.1006, Ex.1007 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 ’456 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 quoted material has been added. Claim terms are presented in *italics*.

IV. OVERVIEW OF THE TECHNOLOGY

A. Background

B. 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.1020, ¶[0069]; Ex.1012, pp. 2-3; Ex.1003, ¶27.

Some renewable power sources have large output swings which must be taken into account for stability and economic purposes. Ex.1012, p. 6. For example, during periods of excess power supply, Independent System Operators (“ISOs”)/Regional Transmission Organizations (“RTOs”) initiate curtailment actions to manage the amount of power from various generation sources for a balanced system. Ex.1012, p. 15. Curtailment typically occurs due to overgeneration, a lack of transmission and/or distribution line capacity to handle generation output, or other grid conditions. Ex.1012, p. 15; Ex.1003, ¶28.

During curtailment, various generation sources are temporarily disconnected or operated at reduced capacities. Ex.1020, ¶¶[0045], [0055]–[0056]. Involuntary

curtailment is “ordered ... in response to system-reliability conditions,” also referred to as reliability curtailment. Ex.1012, p. 15. Voluntary curtailment relies upon the wholesale market to cause generation sources to voluntarily reduce output responsive to the market, such as responsive to a negative wholesale price.

Ex.1012, p. 15. Curtailment due to market conditions is also known as economic curtailment. *See* Ex.1007, p. 3. ISOs/RTOs use curtailment as one of their tools to optimize power transmission and avoid grid congestion/maintain grid stability, resulting in wasted electricity (e.g., discarded supply). Ex.1003, ¶¶29-30.

C. Harnessing Curtailed Power with Green Datacenters

The electricity from renewable power sources subject to curtailment is wasted if it is not put to any use. A well-known problem, many have proposed co-locating different loads to harness such potentially wasted energy. Ex.1003, ¶31.

One type of co-located load is a datacenter’s compute load. One approach co-located “small, cheap, and highly efficient mini data centers” to harness stranded energy. Ex.1014, p. 1. Doing so reduced/eliminated costly transmission infrastructure and associated losses. Ex.1014, p. 1. The co-located datacenters used transportable containers with rack servers. Ex.1014, p. 6. Such containers were a well-known strategy for flexible datacenter power source co-location. *See, e.g.*, Ex.1005, Abstract, [0024], [0048]-[0050], [0066]; Ex.1007, pp. 1, 4; Ex.1013, p. 7; Ex.1014, p. 6; Ex.1015, pp. 3, 4; Ex.1016, p. 2; Ex.1003, ¶¶32-33.

Datacenters implement different co-location strategies to manage the variability of stranded energy. One approach switches between the co-located renewable source and a “stable” grid power source for the datacenter, switching to grid power when the renewable output falls below a threshold and switching back to the renewable power when it rises above a threshold. Ex.1017, pp. 1-2 (“direct-coupled PV system design”); 9. Another approach relies solely upon the renewable power source, accepting the volatility of stranded power. Ex.1014, pp. 3, 4; *see also* Ex.1013, pp. 1, 2 (using only stranded power); Ex.1003, ¶¶34-35.


Co-located datacenters handle variability by ramping-up/ramping-down the server power consumption based on the stranded power availability. Ex.1007, p. 4. Such “power-aware computing” thus controlled computing power consumption “based on the amount of power available.” Ex.1018, [0024]. Indeed, it was already known for a controller 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.1006, [0020], [0037], [0049], [0050] (over capacity conditions); [0019], [0027]-[0028] (controlling power delivery); Ex.1003, ¶36.

D. The ’456 Patent

The ’456 patent focuses on “dynamic power delivery to a flexible datacenter us[ing] unutilized behind-the-meter power sources.” Ex.1001, 4:17-20. The flexible datacenter “modulate[s] power delivery” to servers based on either “the

availability of unutilized behind-the-meter power or an operational directive.”

Ex.1001, 4:21-24. The flexible datacenter is a mobile container, e.g. a “storage trailer ... on wheels.” Ex.1001, 6:15-22, 30-38. The container includes a behind-the-meter power input system, a power distribution system, climate control, a datacenter control system, and multiple servers. Ex.1001, 6:15-22; *see also* FIG. 2:



The datacenter control system is used to modulate the power delivery to the servers based on the power availability or operational directive, potentially in

cooperation with a remote master control system and/or local generator control system. Ex.1001, 6:22-29; 8:41-49. Modulating includes an action to “dynamically reduce or withdraw power delivery to one or more computing systems,” limit the number of powered servers to a smaller subset, or reduce the operating frequency or power mode of one or more of the computing systems, adjust which phases of three-phase voltage is used, or ramp-up servers to a desired consumption level. Ex.1001, 7:7-22; 9:19-42; 10:37-47; 12:17-27; Ex.1003, ¶¶37-42.

E. Prosecution History

The examiner allowed many of the claims in the first Office action, rejecting the remaining claims with only a nonstatutory double patenting rejection over claims 1-10 of U.S. 10,452,127. Ex.1002, 40-46. After the applicant filed a terminal disclaimer over the '127 patent, the examiner allowed the claims without providing a reason. Ex.1002, 32-36, 11.

F. Effective Priority Date of the '456 Patent

The '456 patent entered the U.S. national phase on July 31, 2019 from a PCT application filed February 13, 2018 which claims the benefit of Provisional application no. 62/616,348 filed January 11, 2018. For purposes of this IPR only, Petitioners do not dispute that the Challenged Claims are entitled to a January 11, 2018 priority date.

V. LEVEL OF ORDINARY SKILL IN THE ART

A Person of Ordinary Skill in The Art (“POSITA”) in January 2018 would have had a working knowledge of the networking art that is pertinent to the ’456 patent, including green datacenter efforts and renewable power co-location 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, or more particularly, the use of renewable energy in powering 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, ¶43.

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	1-4, 14, 17-18, 20, 25, 27	35 U.S.C. § 103 over Pelio and Chapel
#2	5-8, 10, 12, 28-30	35 U.S.C. § 103 over Pelio, Chapel, Yang
#3	19, 21-24, 26	35 U.S.C. § 103 over Pelio, Chapel, Belady
#4	9	35 U.S.C. § 103 over Pelio, Chapel, Rasmussen
#5	15-16	35 U.S.C. § 103 over Pelio, Chapel, Corhodzic

U.S. 2011/0316337 (Ex.1005, “**Pelio**”) published December 29, 2011.

U.S. 2010/0145542 (Ex.1006, “**Chapel**”) published June 10, 2010.

Yang et al., “ZCCloud: Exploring Wasted Green Power for High-Performance Computing” (Ex.1007, “**Yang**”) published May 2016. Yang was presented at the 2016 IEEE International Parallel and Distributed Processing Symposium held May 23-27, 2016. Yang includes a copyright date of 2016. Ex.1007, 1. Yang was added to IEEE Xplore on July 21, 2016 and was publicly accessible as of that date (and would have also been distributed to attendees without restriction). Ex.1021; Ex.1003, ¶¶194-95. Thus, Yang was publicly accessible well before the earliest priority date of the ’456 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. 8,849,469 (Ex.1008, “**Belady**”) issued September 30, 2014.

Rasmussen, “Impact of Leading Power Factor on Data Center Generator Systems,” by Schneider Electric (Ex.1010) is a white paper published in 2014. Rasmussen has a copyright date of 2014. Ex.1010, p. 19. The Internet Archive Wayback Machine captured the index of white papers available on Schneider Electric’s support website, www.apc.com, including Rasmussen, as of at least December 22, 2014. *See* Ex.1022, p. 3 (“WP-200 v0,” “Impact of Leading Power Factor on Data Center Generator Systems”); *see also* Ex.1026 (Nathaniel Frank-White Affidavit). “WP-200 v0” corresponds to the identification on Rasmussen’s front page as “White Paper 200, Revision 0.” Ex.1010, p. 1; Ex.1003, ¶¶371-72.

APC was a leading supplier of power control equipment for datacenters, leading to APC’s acquisition by Schneider Electric in 2007. Ex.1025, p. 1. A POSITA interested in seeking information about datacenter power control solutions would have known to access the APC support website and look at APC’s white papers search tool where Rasmussen was indexed, organized, and accessible. Ex.1003, ¶373; *see also* Ex.1022, p. 1 (search fields for searching all white papers). Indeed, others accessed Rasmussen, such as Zhabelova et al. *See* Ex.1023, pp. 2, 6 (citation [6]). Rasmussen’s metadata confirms the white paper was created November 12, 2014. Ex.1024. Thus, a POSITA would have found Rasmussen to be publicly accessible as of at least December 22, 2014. *See Hulu*, Paper 29 at 3,

21.

U.S. 8,080,900 (Ex.1011, “**Corhodzic**”) issued December 20, 2011.

Thus, Pelio, Chapel, Yang, Belady, Rasmussen, and Corhodzic are prior art under 35 U.S.C. § 102(a)(1). None of these references were before the Patent Office during prosecution.

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, ¶¶70, 206, 324, 382, 400.


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, 1041-1042 (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 1-4, 14, 17-18, 20, 25, 27 Obvious Over Pelio-Chapel

1. Pelio

Pelio describes collocating 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 photo-voltaics, wind, wave, and tidal.” Ex.1005, [0033]; *see, e.g.*, Figure 1A (datacenter 160 collocated with/powered by a wind generator farm 162 (Ex.1005, [0044])):



Modular datacenter 160 includes “data center containers” like standard

shipping containers length types “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, the datacenter can either receive power via a backfeed from the grid, or drop pending jobs in favor of another datacenter, when the servers no longer have power. Ex.1005, [0037], [0106]; Ex.1003, ¶¶47-50.

2. Chapel

Chapel “relates generally to electrical power distribution and management,” including intelligently monitoring and “control[ling] delivery of electricity over the circuit.” Ex.1006, [0002]. This includes controlling power delivery to electrical devices like servers in a datacenter. Ex.1006, [0047], [0080]. The servers are plugged into the receptacles, which are in communication with a local controller. Ex.1006, [0082]. A system with multiple smart receptacles (e.g., into which servers are connected) “communicate with a local controller” at the customer premises, which in turn communicates “with a remote controller” to control power to the computers. Ex.1006, [0056], [0083]; *see, e.g.*, FIG. 6:




FIG. 6
Ex.1006, Fig. 6.

When an operator of the remote controller desires to “power down or power up a device or set of devices,” the operator directs the remote controller to send a message to the local controller, which conveys it to the appropriate receptacle/server. Ex.1006, [0090]; Ex.1003, ¶¶51-54.

3. Reasons to combine Pelio and Chapel

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

First, Pelio and Chapel are analogous art, pertaining to the same field of endeavor as the '456 patent, power delivery to computers in a datacenter. The Challenged Claims focus on a “flexible datacenter” with behind-the-meter power input system, datacenter control system, and multiple computing systems, with the power delivery modulated based on unutilized behind-the-meter power availability

or operational directive. Ex.1001, 19:5-16 (claim 1). The operational directive can be based on a variety of conditions relating to behind-the-meter power availability. Ex.1001, 19:17-20:46 (dependent claims); Ex.1003, ¶¶55-56.

Pelio is in the same field of endeavor. Pelio co-locates datacenter containers at a generating station before step-up transformations for transmission and distribution (behind-the-meter). Ex.1005, Abstract, [0002]. Pelio's power generating station examples include various renewable types that have variable output. Ex.1005, [0037]. Pelio recognized that one alternative to deal with variable power output was for the co-located datacenter to process pending compute jobs only when sufficient power is available. Ex.1005, [0106]. Thus, Pelio is also in the field of power delivery to computers in a datacenter. Ex.1003, ¶57.

Pelio is also reasonably pertinent. The '456 patent explained that problems for datacenters included a large increase in power requirements and wasted energy from renewables. Ex.1001, 18:25-31. The '456 patent also identified the problem of controlling power at the datacenter based upon how much power was currently available. *See* Ex.1001, 4:24-27, 18:45-50. Another problem was the cost imposed by grid transmission and distribution. *See* Ex.1001, 18:51-57. Pelio's teachings provide multiple solutions to these problems, placing datacenter containers behind-the-meter to avoid transmission and distribution costs and transmission power loss. Ex.1005, Abstract, [0002], [0019]. Pelio also teaches using renewable power only

when it is available at the datacenter. Ex.1005, [0106]. Pelio’s teachings are therefore reasonably pertinent. Ex.1003, ¶58.

Chapel is also in the same field of endeavor. Chapel relates to “electrical power distribution and management” of electrical devices, including controlling power delivery to datacenter devices. Ex.1006, [0002], [0047]. Chapel teaches controlling the power delivery for multiple reasons, including when actual or projected power supply is insufficient for the datacenter. Ex.1006, [0048]; Ex.1003, ¶59.

Chapel is also reasonably pertinent. Chapel provides solutions to the need of modulating power delivery in response to power supply changes, by intelligently monitoring circuits in a datacenter and controlling electricity delivery. Ex.1006, [0002], [0019], [0056], [0079]. Chapel controls both powering up and down multiple devices at the premises based on power supply status. Ex.1006, [0090]. A POSITA would have found Chapel’s teachings reasonably pertinent to the problem of controlling power at the datacenter. Ex.1003, ¶60.

Second, Pelio teaches containers receiving behind-the-meter power, and distributing electrical power to servers within a container via “power feeds,” but leaves the power distribution details up to POSITAs. Ex.1005, [0002], [0019], [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.1006, [0027]; [0082]. Power distribution units were a well-known aspect of datacenter design, including Pelio’s datacenter containers. Indeed, Pelio’s behind-the-meter renewable sources were well-known for powering datacenters well before the ’456 patent. *See, e.g.*, Ex.1005, Abstract; Ex.1007, pp. 1, 4; Ex.1013, p. 7; Ex.1014, p. 6; Ex.1015, pp. 3, 4; Ex.1003, ¶¶61-62.

Third, Pelio recognizes the variability of renewable sources resulting in times where there is insufficient power for computing operations. Ex.1005, [0037], [0106]. Pelio’s teachings focus on siting datacenter containers behind-the-meter to harness the power output before transmission costs and losses arise, as well as providing equipment configuration flexibility, instead of details for handling power source variability. *See* Ex.1005, [0002]; [0070]. Chapel teaches monitoring and controlling power delivery to manage electrical power distribution within a datacenter based on the available power supply. Ex.1006, [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.1006, [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.1006, [0090]. Thus, a POSITA would have been further motivated to turn to Chapel for Chapel’s datacenter power management teachings. Ex.1003, ¶¶63-64.

Fourth, a POSITA would have had a reasonable expectation of success combining the teachings of Chapel with Pelio, because Chapel’s teachings are compatible with Pelio’s teachings that provide power to servers within a datacenter container from a variable power source. Ex.1005, Abstract; [0002]; [0033]; [0037]; [0050]; [0106]. Chapel teaches how to modulate the power delivery to devices within a datacenter. Ex.1006, [0002]; [0047]; [0080]-[0083]. 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, ¶65.

Combining Chapel’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 co-locating datacenter containers with “renewable generators.” Ex.1007, pp. 1, 4; *see also* Ex.1013, p. 7 (containers co-located at generation sites); Ex.1014, p. 6 (modular containers); Ex.1015, pp. 3, 4 (harnessing curtailed power with containers co-located); Ex.1016, p. 2 (“portable containerized modules”); Ex.1003, ¶66.

It was also well-known to control the datacenter’s load in response to a renewable 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.1007, 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.1006, [0020], [0037], [0049], [0050]. Implementing Chapel’s power control teachings with Pelio’s teachings would beneficially enable Pelio’s datacenter containers to follow a varying power supply. Ex.1003, ¶67.

The Pelio-Chapel combination would predictably result in Pelio’s datacenter containers including Chapel’s teachings for controlling power delivered to servers. *See* Ex.1005, Abstract; [0002]; [0033]; [0037]; [0050]; [0106] (containers behind-the-meter); Ex.1006, [0002]; [0047]; [0080]-[0083] (monitoring, controlling power delivery). Doing so would advantageously enable Pelio’s datacenters to better follow the power supply and avoid grid electricity costs when the renewable supply cyclically dwindles. Ex.1003, ¶68.


Indeed, the Pelio-Chapel combination is nothing more than the combination of known elements (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 Pelio’s datacenters), to yield predictable results (Pelio’s datacenter power consumption following available supply, according to Chapel’s power delivery control teachings). Ex.1003, ¶69.

4. Claim 1

[1.0]


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

First, Pelio discloses “data center containers,” each an example of a “*flexible datacenter*.” Ex.1005, Abstract. Pelio’s Figure 2 (below) illustrates multiple examples including “containers 224A and 225B ... and a third data center container 225.” Ex.1005, [0049]; FIG. 2:




Ex.1005, Fig. 2 (annotated); Ex.1003, ¶¶71-72.

These datacenter containers are mobile and home to thousands of servers each. Ex.1005, [0050], [0066], [0070]-[0071]. Using containers facilitates “**flexible, just-in-time, modular data centers**.” Ex.1005, [0095]; FIG. 4 (Ex.1005, [0097]):



These flexible datacenter containers, which “can be moved between trucks,” are the very type of datacenters the ’456 patent describes. *See* Ex.1005, [0071], [0089]; Ex.1001, 6:30-38; Ex.1003, ¶¶73-76.

Second, Pelio’s “modular datacenter”—which houses one or more datacenter containers—is another example of a “*flexible datacenter*.” Ex.1005, Abstract. Figure 2 illustrates “a flexible, just-in-time, modular data center” housing multiple containers. Ex.1005, [0024], [0048]; FIG. 2:



Ex.1005, Fig. 2 (annotated); Ex.1003, ¶¶77.

The containers can be moved in and out of the building shell, allowing changes in critical load “container-by-container” via the servers therein. Ex.1005, [0050], [0055], [0066]; [0067], [0068], [0071]. The modular datacenter itself is, therefore, also a “flexible datacenter.” Ex.1003, ¶¶78-80.

[1.1]

First, Pelio discloses the modular datacenters/containers receive power “*behind-the-meter*,” with “**data centers located at a power generation source** and connected to be powered ... **at a point before the voltages are stepped up through transformers to the transmission and distribution utility grids.**”


Ex.1005, [0002]; Ex.1003, ¶¶81-82.

Accordingly, power is provided directly to the collocated datacenter where it “does not have to be stepped-up to transmission level.” Ex.1005, [0019]. This

avoids power losses from “[e]ach stepping up and stepping down of voltages with transformers” and transmission line power losses. Ex.1005, [0039]; Ex.1003, ¶83.

Further, the power input to each datacenter container within the modular datacenter is also “*behind-the-meter*,” since the containers receive power via “modular power supply systems 240-[2]43” from the behind-the-meter power input to the modular datacenter. Ex.1005, [0051]. The power supply systems to each container can be “utility sub-panels and transformers,” or UPS systems or generators, or “[e]ach data center container may ... **obtain its power directly from the utility.**” Ex.1005, [0051]; Ex.1003, ¶84.

Pelio’s “modular power supply systems 240-[2]43” are examples of a “*power input system*,” as is the utility sub-panel 245:



Ex.1005, Fig. 2 (annotated); Ex.1003, ¶85.

Second, Pelio provides several “*behind-the-meter*” power generation station

examples. Ex.1005, [0005]. Figure 1A's example shows a datacenter "collocated with and powered by a wind generator farm 162," prior to step-up. Ex.1005, [0044]; Figure 1A:




Figure 1B's example shows a datacenter "collocated with and powered by solar electric generating station 150," prior to step-up. Ex.1005, [0045]; Figure 1B:





Figure 1C's example shows a datacenter "collocated with and powered by hydro electric generating station 106," prior to step-up. Ex.1005, [0046]; Figure 1C:



Common to each example is the datacenter taking its “power directly from the electrical generating units before any corresponding step-up transformer.” Ex.1005, [0047]. Since Pelio’s modular datacenter receives power before step-up, Pelio’s power input system at the modular datacenter, and power input system at each container, is likewise “*behind-the-meter.*” Ex.1003, ¶¶86-92.



[1.2]

Pelio renders obvious, alone and combined with Chapel, [1.2].

First, each datacenter container in Pelio includes a “*power distribution system*.” Pelio states the obvious: “electrical power [is] consumed by the servers and data storage units,” with up to “1,000 or more servers each” per container. Ex.1005, [0071], [0075]. It would have been obvious that the electrical power, input via the “modular power supply systems” and utility sub-panel (Ex.1005, [0051]), would be distributed to each server in a container via a “*power distribution system*.” See, e.g., Ex.1006, [0027], [0082]; Ex.1003, ¶¶93-94.

Pelio teaches distributing the power to servers in a container via one or both of “the ‘A-side’ and ‘B-side’ power feeds.” Ex.1005, [0079]. Power feeds are well-known to be “independent electrical feeds coming into a data center from the utility company” with corresponding power distribution units (PDUs) container-wide and rack-specific. Ex.1019, p. 1; see also Ex.1006, [0094] (datacenter using a PDU associated with an A source and a B source); Ex.1003, ¶95.

Second, Pelio’s modular datacenter itself includes a “*power distribution system*” with its “raceway 440” teaching, which “provides ... electrical power connections to uninterruptible power supply (UPS) unit 442 and utility power.” Ex.1005, [0099]; figures 2 and 4:



Third, Chapel confirms that a datacenter, like in Pelio, has a “*power distribution system*” (PDU) to distribute “*behind-the-meter*” power. Chapel discloses “a power distribution unit of a data center.” Ex.1006, [0027]; [0082]. The PDU outputs power received from one or more power sources to servers within a datacenter. Ex.1006, [0082]. The PDU “may be associated with ... an A source and a B source.” Ex.1006, [0082]; FIG. 6:

Some renewable generating stations “naturally cycle between full operation and zero,” such as solar with sunshine and others with wind, etc. Ex.1005, [0037]. Thus, Pelio teaches there are times when power is less available from renewables (if at all) than other times, and requires some action to account for that shortfall. Pelio’s solutions include drawing “inexpensive power” from a grid, or stopping servers from operating while power is not available. Ex.1005, [0037]; *see also* [0043]; [0105]-[0106]; Ex.1003, ¶105.

Pelio focuses on the design of a modular datacenter and collocation with a generating station (Ex.1005, Abstract), leaving details of how power can be controlled within a datacenter up to a POSITA. A POSITA would have been motivated to turn to Chapel, which teaches ways to control power in a datacenter to address potential shortfalls, such as those arising from variable sources of power including e.g., solar, wind, etc. (as Pelio taught). Ex.1003, ¶106.

Second, Chapel teaches how to control power within a datacenter for actual or potential “over capacity condition[s].” Ex.1006, [0048]. Such a condition exists when the available power supply (e.g., from a renewable station) falls below the power demand of Pelio’s datacenter. *See* Ex.1005, [0035]. Chapel teaches controlling power in a datacenter by modulating power delivery to one or more servers to address such over capacity conditions. Ex.1003, ¶107.

Chapel covers “electrical power distribution and management,” including

“intelligently monitor[ing] at least a portion of the circuit and control[ing] delivery ... over the circuit.” Ex.1006, [0002]. Chapel’s teachings apply to “a data center” with “data center equipment control and power distribution systems.” Ex.1006, [0019], [0047]. Chapel’s teachings enable “remote monitoring and/or control of multiple receptacles” providing power to servers. Ex.1006, [0056]. “[A] local controller 304” at the “customer premises” communicates with the receptacles, and also “with a remote controller 308” operated by an owner of the customer premises, a controller of the power grid, etc. Ex.1006, [0056]; FIG. 3:





FIG.3

Ex.1006, Fig. 3 (annotated).

The “local controller” discloses a “*datacenter control system*,” since it is a control system at the datacenter. Ex.1006, [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.1006, [0018], [0025]-[0026]; Ex.1003, ¶¶108-109.

Chapel’s Figure 6 illustrates “control in a data center environment,” with servers plugged into controllable receptacles. Ex.1006, [0080]-[0081]. “[A] local controller 650” communicates with the receptacles and a remote controller 670. Ex.1006, [0082]. The local and remote controllers collectively provide “remote and/or policy based operation” of the servers. Ex.1006, [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 requested action. Ex.1006, [0090]; FIG. 6:



Third, Chapel teaches the local controller (“*datacenter control system*”) modulating power delivery via policies, including those from a remote controller,

like “a message from the electrical utility provider specifying transition to an efficiency mode or a brownout mode.” Ex.1006, [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, ¶111.

In combination, 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 Pelio’s flexible datacenters, according to a message from a remote controller. Ex.1006, [0017], [0025]-[0026], [0076]. A POSITA would have understood that Pelio’s flexible datacenters collocated 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 delivery of power to one or more datacenter servers to reduce load and avoid an overload condition. Ex.1003, ¶¶112-113.

5. Claim 2

Chapel teaches a local controller (“*datacenter control system*”) receiving a message from a remote controller identifying a power delivery modulation action to take, such as transitioning “to an efficiency mode or a brownout mode”

(“operational directive”). See [1.4]; Ex.1006, [0076]; see also [0058], [0066], [0078]. In Chapel’s Figures 3 and 6, the local controller receives messages from a remote controller (“remote master control system”). Ex.1003, ¶¶114-116

In Figure 3, receptacles (which servers plug into) “communicate with a local controller 304,” and the “local controller 304 ... communicate[s] with a remote controller 308” from a remote location. Ex.1006, [0056], [0027]; FIG. 3:





FIG.3
Ex.1006, Fig. 3 (annotated); Ex.1003, ¶¶117-118.

Figure 6 illustrates remote “control in a data center environment.” Ex.1006, [0079]-[0080]. A “local controller” communicates with a remote controller to implement “remote and/or policy based operation” of the datacenter devices. Ex.1006, [0082]-[0083]; FIG. 6:



Ex.1006, Fig. 6 (annotated).

For example, in response to a remote controller operator choosing “to power down or power up ... a set of devices,” the remote controller transmits a message “to the local controller” for implementation. Ex.1006, [0090]; Ex.1003, ¶¶119-120.

In combination with Pelio’s teachings, a POSITA would implement Chapel’s local and remote controllers, and messaging control information teachings, to control datacenter power consumption. Ex.1003, ¶¶121-122.

6. Claim 3

Chapel teaches a remote controller sending messages (operational directives) to a datacenter’s local controller to control datacenter power usage. *See* claim 2; Ex.1006, [0027], [0056] (Figure 3), [0079]-[0083], [0090] (Figure 6). The messages are sent to deal with actual or potential “over-capacity conditions.” Ex.1006, [0026]. An over-capacity condition is a “condition potentially requiring

reduction of power provided to standard ... commercial customers” to avoid blackouts or brownouts. Ex.1006, [0026]. The conditions result from circumstances including “increasing industrial and residential power needs,” “aging power infrastructure,” and “practical limitations on new power generation.” Ex.1006, [0004]. Chapel’s “new power generation” includes Pelio’s renewable power sources. Ex.1005, [0005]; *see also* [0033] (additional generator types applicable too); Ex.1003, ¶¶123-126.


In combination, Chapel’s remote controller teachings would cause a local controller (via messaging, an “*operational directive*”) to modulate power delivery to datacenter equipment (including servers) with Pelio’s BTM-located generating source teachings, according to actual or potential over-capacity conditions from periodic/cyclical BTM-located power sources (e.g., cycling between full/low/zero output periods). *See* Ex.1005, [0019], [0033], [0037]; Ex.1006, [0028], [0088]. A POSITA would have also applied Chapel’s teachings to modulate power delivery where over-capacity may occur due to higher demand at the datacenter. Thus, the control relies on how much power is available relative to demand at the datacenter (“*depends on power availability at the flexible datacenter*”). Ex.1003, ¶¶127-129.

7. Claim 4

Pelio teaches a flexible datacenter. *See* [1.0]; Ex.1005, Abstract, [0024], [0048], [0049]; Ex.1003, ¶¶130-132.

Further, the Pelio-Chapel combination renders obvious a remote controller that provides messages of instruction to the datacenter for controlling power delivery. *See* claim 2; Ex.1006, [0027], [0076]; Ex.1003, ¶133.

Chapel teaches locating the remote controller (“*remote master control system*”) away from the datacenter. *See* claim 2; Ex.1006, [0056], [0083]. In combination, Pelio’s flexible datacenter is thus located remotely from Chapel’s remote controller. *See also* Ex.1006, FIGs. 3, 6:



Ex.1006, Fig. 3 (annotated).

implementing reduced power mode, and/or interrupting power distribution (Ex.1006, [0025]-[0026]), but also turning devices back on, ending reduced power mode, removing power interruption, etc. *See* Ex.1006, [0068]; Ex.1003, ¶140. And, since Chapel teaches performing the power delivery control when supply is expected to be less than demand (the load), it would have been obvious for the remote controller to increase power delivery to load when the supply increases again. *See also infra* [28.3]; Ex.1003, ¶¶140-141.

9. Claim 17

Chapel teaches using the local controller at the customer premises to “control[] delivery of power to the electrical device.” *See* [1.4]; Ex.1006, [0018]. Controlling power delivery includes turning off devices or operating the devices in reduced power mode – e.g., by interrupting power distribution. Ex.1006, [0025]-[0026]; Ex.1003, ¶¶142-144.

Chapel’s Figure 6 illustrates “control in a data center environment,” where datacenter devices (e.g., servers) are plugged into controllable receptacles. Ex.1006, [0080]-[0081]. These are connected to “a local controller” which also communicates with a “remote controller.” Ex.1006, [0082]-[0083]. The remote and local controllers operate together to achieve “remote and/or policy based operation of” datacenter devices. Ex.1006, [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 sends a message to the local controller which then implements the action. Ex.1006, [0090]; Ex.1003, ¶145.

In combination, the servers from Pelio’s teachings would be powered down or up by controlled power delivery to one or more of the servers, per Chapel’s teachings. Ex.1003, ¶¶146-147.

10. Claim 18

Receptacles, which servers (“*plurality of computing systems*”) plug into, “communicate with a local controller 304” that can be at the “customer premises.” *See* [1.4]; Ex.1006, [0056]. It would have been obvious that one option was for a controller to be implemented separately from the computing devices (and, thus, “*remotely*” therefrom). *See, e.g.*, Ex.1008, 8:15-20 (implementing controller distinct from the datacenter processing equipment). Thus, Chapel’s local controller alone renders this limitation obvious. Ex.1003, ¶¶148-150.

Further, claim 24 acknowledges there is no distinction between a datacenter control system and a remote master control system: “*the datacenter control system is a remote master control system.*” Thus, claim 1’s “*datacenter control system*” is broad enough to encompass local and remote scenarios, which Chapel teaches. Ex.1003, ¶152.

Chapel’s local controller 304 communicates with “**remote** controller 308” operated “by a private or public party.” Ex.1006, [0056]. Similarly, Chapel’s

Figure 6 “local controller 650” communicates with remote controller 670. Ex.1006, [0080]-[0082]. The remote controller controls power to the servers by transmitting a message to the local controller to implement the power-down/power-up action. Ex.1006, [0090]. Chapel’s “remote controller” are thus another example of a “*datacenter control system*” according to the claim, thus rendering [18.1] obvious. Ex.1003, ¶¶151-153.

11. Claim 20

Receptacles, which servers (the “*plurality of computing systems*”) plug into, “communicate with a local controller 304” that can be at the “customer premises.” Ex.1006, [0056]; FIGs. 3, 6:





FIG.3

Ex.1006, Fig. 3 (annotated).



Since the local controller is “at the customer premises,” the local controller is “collocated with the plurality of computing systems.” Ex.1003, ¶¶154-157.

12. Claim 25


[25.0]

Pelio discloses [25.0] to the extent it is limiting. Pelio discloses a “modular datacenter” comprising “a building shell,” a modular common power and cooling infrastructure, and multiple “data center containers.” See [1.0]; Ex.1005, Abstract, [0024], [0048]. The modular datacenter’s interior supports multiple containers and, therefore, distinct loads. Ex.1005, [0048]. The modular datacenter, with multiple containers (e.g., first and second flexible datacenters) as well as control elements, discloses a “system.” Ex.1003, ¶¶158-160.


[25.1]

Pelio discloses a first flexible datacenter with a first container. *See* [1.0];

Ex.1005, FIGs. 2, 4:



Ex.1005, Fig. 2 (annotated).



Ex.1005, Fig. 4 (annotated).

Ex.1003, ¶¶161-162.

[25.2]

See [1.1]. Ex.1003, ¶163.

[25.3]

See [1.2]. Ex.1003, ¶164.

[25.4]

Pelio, alone and combined with Chapel, renders obvious [25.4].

First, Pelio discloses a plurality of computing systems. *See* [1.3]; Ex.1003, ¶¶165-166.

Second, Pelio's plurality of computing systems receive power from the BTM power input system via the power distribution system. The datacenter servers receive power from one or both of "'A-side' and 'B-side' power feeds." *See* [1.2]; Ex.1005, [0079]. These feeds are well-known to include a power distribution unit and necessary hardware to power the servers. *See* Ex.1019, p. 1; Ex.1006, [0082]. The A-side and B-side power distribution system of a given flexible datacenter provides the BTM power to the servers, thus rendering obvious [25.4]. Ex.1003, ¶¶167-168.

Third, Pelio's modular datacenter (shell) itself includes a "raceway" providing "electrical power connections" for the datacenter containers. Ex.1005, [0099]. This further renders obvious the servers receiving power from the BTM power input system (e.g., utility power panel) via the raceway. Ex.1003, ¶169.

Fourth, Chapel confirms that a datacenter, like Pelio's, uses a "*power distribution system*" to distribute power from the power input. *See* Ex.1006,

[0027], [0082]; *and* [1.2]. In combination, Pelio's plurality of datacenter computing systems receive the power distributed via Chapel's PDU after receipt from Pelio's BTM power input system. Ex.1003, ¶¶170-171.

computing systems receive the power distributed via Chapel's PDU after receipt from Pelio's BTM power input system. Ex.1003, ¶¶170-171.


from Pelio's BTM power input system. Ex.1003, ¶¶170-171.

[25.5]

Each of Pelio’s multiple “data center containers” is an example of a different “flexible datacenter,” including first and second containers. *See* Ex.1005, Abstract, [0049]; [0097]; Figures 2, 4:

“flexible datacenter,” including first and second containers. *See* Ex.1005, Abstract,

[0049]; [0097]; Figures 2, 4:



Ex.1005, Fig. 2 (annotated).




FIG. 5

Ex.1001, Fig. 5; Ex.1003, ¶¶176-177.

First, Pelio-Chapel renders obvious a routing control system configured to modulate power delivery to computing devices in flexible datacenters. *See* [1.4]. Chapel teaches a “remote controller” and a “local controller” relative to the datacenter to control modulation of power delivery in each datacenter. *See* claim 2. Chapel’s Figure 3 demonstrates “remote monitoring and/or control of multiple receptacles” which provide power to servers via a remote controller. Ex.1006, [0056]. Figure 6 demonstrates “control in a data center environment” via remote control. Ex.1006, [0079]-[0080], [0082]-[0083]. For example, “to power down or power up ... a set of devices,” the remote controller transmits a message to the local controller to implement. Ex.1006, [0090]; Ex.1003, ¶¶178-181.

Second, in combination, a POSITA would implement Chapel’s local and remote controllers, and messaging control information, to modulate power delivery

at Pelio’s flexible datacenters. Indeed, Chapel teaches a remote controller (“grid controller”) communicating with multiple local controllers at different datacenter locations to operate in a requested mode of operation. Ex.1006, [0025]; [0077] (grid controller is a “remote grid controller”). Thus, a remote controller (“*routing control system*”) sends messages to each flexible datacenter to modulate power delivery at each datacenter, per Pelio-Chapel’s combined teachings. *See* Ex.1005, Abstract (flexible datacenters); Ex.1006, [0018] (modulating power delivery); Ex.1003, ¶¶182-184.

13. Claim 27

[27.0]

See [25.0]-[25.6]. Ex.1003, ¶185.

[27.1]

See claim 4, [25.6]. Ex.1003, ¶¶186-187.

[27.2]

First, Pelio-Chapel renders obvious modulating power delivery to a flexible datacenter based on power availability. *See* [14.1]. Chapel teaches monitoring power availability (e.g., Pelio’s BTM power availability as further shown at [28.1] *infra*). *See* Ex.1006, [0004], [0047], [0056], [0083]. In Chapel’s Figure 3 example, a remote controller communicates with a local controller to modulate power delivery. Ex.1006, [0056]. In Chapel’s Figure 6 example, servers are controlled via

messages from a remote controller to a local controller to implement policy based operation. Ex.1006, [0080]-[0083]. Modulating power delivery includes powering down or up a set of devices. Ex.1006, [0090]; Ex.1003, ¶¶188-191.

Second, it would have been obvious to repeat the Pelio-Chapel teachings with each flexible datacenter of Pelio. This is nothing more than the mere duplication of parts (Chapel’s modulating power delivery for each flexible datacenter taught by Pelio). Ex.1003, ¶¶192-193.

C. Ground 2: Claims 5-8, 10, 12, 28-30 Obvious Over Pelio-Chapel-Yang

1. Yang

Yang co-locates datacenters with renewable generating sources. Ex.1007, p. 1. In response to curtailing renewable power generation for economic or reliability concerns, Yang teaches powering datacenter computing “exclusively by wasted power generated from intermittent renewable sources.” Ex.1007, pp. 1, 2-3. Using containers eliminates building costs, and co-locating at with the generating source eliminates transmission. Ex.1007, pp. 1, 4. Yang activates computing resources “only when stranded power is available, and will be shutdown whenever there is insufficient power.” Ex.1007, p. 1; Ex.1003, ¶¶194-196.

2. Reasons to combine Pelio and Yang

A POSITA would have been motivated to combine Pelio and Chapel with Yang for several reasons.

First, Pelio and Yang pertain to the same field of endeavor as the '456 patent (power delivery to computers in a datacenter. Yang's "computing resources are deployed on generation sites and placed in containers" without transmission needs. Ex.1007, pp. 1, 4. Just like Pelio and the '456 patent, Yang co-locates container-based computing resources at generating stations for behind-the-meter power. Yang only powers container computing resources when there is "stranded power" (power that the grid does not want or cannot handle). Ex.1007, pp. 1, 2-3. Thus, Yang is in the same field of endeavor. Ex.1003, ¶¶197-198.

Yang is also reasonably pertinent to '456 patent-identified problems including what to do with "unutilized and typically wasted energy," how to control power at the datacenter based on generating station power availability, and grid transmission and distribution costs. *See* Ex.1001, 4:24-27; 18:25-31; 18:45-57. Yang teaches co-locating the datacenter to harness "unutilized and typically wasted energy" to exclusively power container datacenter computing. Ex.1007, pp. 1, 2-3; Ex.1001, 18:25-31. Yang describes controlling power at the datacenter based upon power availability from the energy source, activating computing resources "only when stranded power is available," and shutting down otherwise. Ex.1007, p. 1.

Thus, a POSITA would have also found Yang's teachings to be reasonably pertinent. Ex.1003, ¶199.

Second, a POSITA would have been motivated to combine Pelio-Chapel with Yang for Yang's teachings of harnessing "stranded power to enable scaling to higher levels of capability and performance" in datacenter computing. Ex.1007, p. 4. "[S]tranded power" is generated power that "cannot be used" due to economic or reliability constraints. Ex.1007, p. 3. Using only "stranded power" with co-located container-based computing beneficially provides "incremental scalability of infrastructure, low power cost, and negligible carbon footprint." Ex.1007, p. 9. A POSITA would have been motivated to combine Yang's teachings with Pelio-Chapel's container-based power delivery control to obtain these benefits. Ex.1003, ¶200.

Third, a POSITA would have had a reasonable expectation of success combining Yang's teachings with Pelio's and Chapel's teachings. Like Pelio, Yang provides mobile and scalable server computing in a container platform. *See* Ex.1005, [0055]-[0058] (scale as-needed, more efficient cooling management); Ex.1007, p. 4 (same). Further, Yang and Pelio teach locating container-based datacenters behind-the-meter to avoid transmission and distribution issues. *See* Ex.1005, Abstract; [0002]; [0033]; [0037]; [0050]; [0106] (containers behind-the-meter); Ex.1007, pp. 1, 4 (containers behind-the-meter). Since both Pelio and Yang

compatibly teach using container-based computing platforms, and siting those containers behind-the-meter, a POSITA would have a reasonable expectation of success implementing Yang’s “stranded power” teachings for powering Pelio’s datacenter. Ex.1003, ¶¶201-202.

Fourth, combining Yang’s teachings with Pelio-Chapel’s teachings would have been predictable and within a POSITA’s skillset. Locating datacenters behind-the-meter with renewable power sources was known. Ex.1005, Abstract, [0002]; Ex.1007, pp. 1, 4; *see also* Ex.1013, p. 7 (containers collocated at generation sites); Ex.1014, p. 6 (modular containers); Ex.1015, pp. 3, 4 (containers directly placed at generation sites for curtailed power); Ex.1016, p. 2 (portable containers); Ex.1003, ¶203.

It was also well-known to control the datacenter’s load in response to the varying supply from a renewable power source. For example, in Yang servers “only work when stranded power is available, and otherwise [have them] fully powered off.” Ex.1007, p. 4. Thereby, “power-aware computing” controlled “the amount of power consumed by the computing based on the amount of power available.” Ex.1018, [0024]. Yang’s power control teachings at a datacenter (e.g., only operating when stranded power available) would be a natural extension of Pelio’s understanding of the intermittent availability of behind-the-meter renewable power. Ex.1003, ¶204.

Implementing Yang’s “stranded power” teachings would beneficially reduce power consumption when power supply drops, which Pelio recognizes as a cyclical phenomenon for some renewable sources, and result beneficially in infrastructure scalability, “low power cost, and negligible carbon footprint.” Ex.1007, p. 9. Yang’s stranded power teachings further “exploit[] un-economic grid power” and “create[] dispatchable loads” for excess generation that “benefit grid stability, cost, and management.” Ex.1007, p. 9. A POSITA would have therefore been motivated to turn from the Pelio-Chapel combination to Yang’s teachings. Ex.1003, ¶205.

3. Claim 5


The Pelio-Chapel-Yang combination teaches controlling power delivery to computing systems based on an operational directive (per Pelio-Chapel, *see* [1.4]), based on availability of behind-the-meter generation (per Pelio, *see* [1.1]). *See* Ex.1005, [0032]-[0033]; Ex.1007, pp. 1, 3. Pelio-Chapel-Yang renders obvious claim 5 in several ways, any one of which invalidates the claim. *See* claim 5 (“*one or more of ... or*”). Ex.1003, ¶¶207-208, 212.

“excess local power generation at a local station level”

First, it would have been obvious that in different situations, BTM collocated power sources would generate excess power since certain renewable power sources cycle “between full operation and zero.” Ex.1005, [0037]. Thus, there would be times where the power source generates more power than the

datacenter may need, such as during peak sunshine or high wind. Ex.1005, [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.1005, [0044]-[0046]. Such extra generated power renders obvious “*excess local power generation at a local station level.*” Ex.1003, ¶¶209-211

Second, Yang further renders this obvious. Yang teaches ways in which to harness power that is curtailed or otherwise available from renewable power sources but not needed on the grid (stranded power). Ex.1007, p. 3. Yang places datacenters in containers that are co-located with renewable generators (“on a wind farm ... directly connected to the power backbone of wind turbines”). Ex.1007, pp. 1, 4; Fig. 3:



Ex.1007, Fig. 3 (annotated).

Intermittent power sources (e.g., solar/wind) have “time-varying productivity” which periodically results in “excess power [that] cannot be

economically distributed through [the] grid.” Ex.1007, pp. 1, 2. This volatility results occasionally in “excess generation” that Yang’s co-located datacenters harness. Ex.1007, p. 3. This also demonstrates “*excess local power generation at a local station level.*” Ex.1003, ¶¶213-214.

“local power generation subject to economic curtailment”

Yang teaches that at times of excess generation, “power markets will price some generator’s outputs with a low or negative price” resulting in uneconomic power, which is “called ‘economic curtailment’ or dispatch down.” Ex.1007, p. 3. In response to such economic pricing signals, “generators will reduce production.” Ex.1013, p. 3. POSITAs understood Yang’s “uneconomic” power as “*economic curtailment*” well before the ’456 patent. *See* Ex.1012, p. 15 (economic/voluntary curtailment happens through the market, reducing output responsive to pricing signals); Ex.1003, ¶¶215-217.

“local power generation subject to reliability[] curtailment”

The Pelio-Chapel-Yang combination teaches controlling power delivery when electricity markets are not able to achieve “merit order” for purchasing energy from the lowest-cost generators via pricing alone, resulting in power generation subject to reliability curtailment. Ex.1007, pp. 2-3. This includes transmission congestion (“generation and load vary at each point in the grid”). Ex.1007, pp. 1, 3; Ex.1005, [0032]-[0033]. Reliability curtailment is used to

resolve congestion issues “to ensure stability.” Ex.1013, p. 3. A POSITA would have understood Yang’s teaching of preventing power from entering the grid by curtailment mechanisms other than pricing signals as “*reliability curtailment*.” This was a well-known concept. Ex.1012, p. 15 (curtailment ordered by grid in response to system-reliability conditions); Ex.1003, ¶¶218-220.

“*low local power generation*”

First, it would have been obvious that in different situations, BTM collocated power sources would not generate enough power, as certain renewable power sources “cycle between full operation and zero.” Ex.1005, [0037]. Thus, there would be times where the power source generates less power than the datacenter needs, such as nighttime when there is no sunshine or low/no wind. Ex.1005, [0037]; *see also* Ex.1008, FIG. 3 (capacity below computational needs); 10:14-21. Such dearth of generated power renders obvious “*low local power generation*.” Ex.1003, ¶¶221-223.

Second, Yang further renders this obvious. Yang teaches that renewable power generation has “time-varying productivity.” Ex.1007, pp. 1, 2, 3. This includes shutting down the datacenter due to “insufficient power.” Ex.1007, p. 1. Such downtimes demonstrate another example of “*low local power generation*.” Ex.1003, ¶¶224-228.

“transient local power generation situations”

Pelio 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.1005, [0106], [0037]. The cycling between full operation and zero demonstrates “*transient local power generating situations*.” Ex.1003, ¶229.

Yang further renders this obvious. Yang’s datacenter computing is powered “exclusively by wasted power generated from intermittent renewable sources.” Ex.1007, pp. 1, 2-3. The computing resources “are activated **only when stranded power is available, and ... shutdown whenever there is insufficient power**.” Ex.1007, p. 1. This also describes “*transient local power generating situations*” (power available or not). Ex.1003, ¶¶230-231.

4. Claim 6

See claim 5 (“excess local power generation at a local station level”).
Ex.1003, ¶¶232-240.

5. Claim 7

See claim 5 (“local power generation subject to economic curtailment”).
Ex.1003, ¶¶241-250.

6. Claim 8

See claim 5 (“local power generation subject to reliability[] curtailment”).
Ex.1003, ¶¶251-257.

7. Claim 10

See claim 5 (“low local power generation”). Ex.1003, ¶¶258-266.

8. Claim 12

See claim 5 (“transient local power generation situations”). Ex.1003, ¶¶267-271.

9. Claim 28

[28.0]

Pelio, alone and combined with Chapel, renders obvious [28.0] to the extent it is limiting.

First, Pelio discloses a flexible datacenter that uses BTM power. *See* [1.0] (flexible datacenter), [1.1] (BTM power); Ex.1003, ¶¶272-273.

Second, Pelio teaches that power from renewable sources of energy “naturally cycle between full operation and zero” based on the renewable energy source (solar, wind, etc.). Ex.1005, [0037]. It would have been obvious that renewables like solar, wind, etc. result in a dynamic delivery of power to the collocated datacenters, thus rendering [28.0] obvious. Ex.1003, ¶274.

Third, Pelio combined with Chapel further renders [28.0] obvious. Chapel


teaches the local controller identifying a condition subject to a policy (e.g., a change in operating mode dictated by a grid supplier). Ex.1006, [0076], [0088], [0089] (e.g., remote grid controller); [0058], [0094]-[0095]. Responsive to instruction from a remote controller, the local controller modifies power delivery by controlling delivery of power to the servers. Ex.1006, [0027]-[0028], [0054]-[0056]; Ex.1003, ¶¶275-276.

[28.1]

Pelio teaches a datacenter receiving power via a BTM power source. *See* [1.1]; Ex.1005, [0002], [0019], [0039]. Some BTM power sources, like “solar photo-voltaics, wind, wave, and tidal,” add power to the grid independent of demand. Ex.1005, [0033]. It would have been obvious to monitor the varying output of such power sources to “throttle up and down” other power sources (Ex.1005, [0033]), or the load itself. Ex.1005, [0106] (dropping load when generating output drops); Ex.1006, [0089] (powering up/down depending upon over-capacity condition). For Pelio’s system to know whether to throttle up or down another power source, or the load level, the current status would need to be monitored. Ex.1003, ¶¶277-278.

This is confirmed by Chapel, which teaches monitoring power availability and distribution to deal with “actual or potential over capacity conditions.” Ex.1006, [0004], [0082]. Thus, Chapel “remotely monitor[s] and/or control[s]

electrical devices” in datacenters. Ex.1006, [0047]. For example, Figure 3’s system provides “remote monitoring and/or control of multiple receptacles” via a local controller communicating monitoring information to a remote controller (each a “control system”). Ex.1006, [0056]. Figure 6’s example also includes a local controller 650 sending information about power availability and load to remote controller 670. Ex.1006, [0083].



Ex.1006, Fig. 3 (annotated); Ex.1003, ¶279.

In combination, it would have been obvious for a remote controller (per Chapel) to receive power availability information from one or more local controllers (and other data sources) from the BTM power sources (per Pelio). *See* Ex.1005, [0033], [0106]; Ex.1006, [0004], [0089]. Pelio’s system would thus

monitor availability of BTM power at the datacenter. Ex.1003, ¶¶280-281.

[28.2]

See [1.1]. Ex.1003, ¶282.

[28.3]

First, Pelio teaches a flexible datacenter. *See* [1.0]. Pelio also teaches providing available BTM power to flexible datacenters for datacenter operation. *See* [1.1]; Ex.1005, [0002] (BTM), [0005] (BTM power availability dependent upon renewable source conditions), [0037] (cycles); Ex.1003, ¶¶283-284.

Second, the Pelio-Chapel combination renders obvious monitoring BTM power availability. *See* [28.1]; Ex.1006, [0004], [0047], [0056], [0083]. Moreover, Chapel's load control involves not only controlling power delivery by reducing power (Ex.1006, [0025]-[0026]), but also turning devices back on, ending reduced power mode, removing power interruption, etc. Ex.1003, ¶285. And, since Chapel teaches the remote controller controlling power delivery when supply is less than demand (load), it would have been obvious for the remote controller to also increase the demand when the supply becomes sufficient again (e.g., over-capacity condition threat removed). *Id.*

Third, Yang confirms it was known to increase load when supply is sufficient again, following power supply availability from a renewable source to power datacenter servers. This includes determining that a condition has been met

to increase the load: “ZCCloud [the datacenter] resources are **activated only when stranded power is available**.” Ex.1007, p. 1; *see also* p. 4 (“ZCCloud computing resources only work when stranded power is available”). The load following the availability of the renewable power supply, including ramping up when sufficient power is available, was well-known. *See also* Ex.1017, p. 9 (adjust server load to follow availability of renewable power). Yang thus teaches determining that “stranded power is available,” which demonstrates “*determining that a ramp-up condition is met*” to run the co-located datacenter. Ex.1003, ¶¶286-290.

In combination, Chapel’s remote controller, while monitoring power availability from the BTM power source at a flexible datacenter (per Pelio), would determine that a condition is met to increase the datacenter server load when there is sufficient power from the power source for operation (per Yang). Moreover, Chapel’s power modulation teachings would apply based on the variability of the renewable source’s power output (per Yang) to also operate on a per-computer or per-rack basis (per Chapel) to follow power output variations during operation (per Yang). Ex.1003, ¶¶290-291.

[28.4]

The Pelio-Chapel-Yang combination teaches determining that a ramp-up condition is met by a remote controller. *See* [28.3]. Further, Chapel teaches a remote controller sending a message to the local controller (a “*datacenter control*

system”) to modulate power delivery, including “specifying transition to an efficiency mode or a brownout mode.” Ex.1006, [0076]; *see also* [0058], [0066], [0090]. Such “messages” from remote controller(s), render obvious “*providing an operational directive to a datacenter control system*” (the local controller). *See* [1.4]; Ex.1003, ¶¶292-293.

Yang teaches starting to operate the datacenter once stranded power becomes available from the co-located renewable source. Ex.1007, pp. 1, 4. It would have been obvious for Yang’s starting operation to occur from the message from the remote controller per Chapel’s teachings. Ex.1003, ¶¶294-295.

[28.5]

As explained at [28.4] and [1.4], Pelio-Chapel renders obvious an “*operational directive*” instructing the “*datacenter control system*” (e.g., Chapel’s “local controller” at Pelio’s datacenter) to modulate power delivery to one or more computing systems at the datacenter. Ex.1006, [0025]-[0026], [0076]; *see also* [0058], [0066], [0090]; Ex.1003, ¶¶296-297.

Further, Yang teaches activating datacenter resources “only when stranded power is available.” Ex.1007, p. 1; *see also* p. 4. In combination, a POSITA would understand Chapel’s teaching of power modulation (“*power delivery*”) to enable delivery of Pelio’s BTM power to datacenter servers in response to the amount of renewable energy available as stranded power being greater than a threshold as

taught by Yang (i.e., when stranded power is available versus when it is not).

Ex.1003, ¶298. A POSITA would apply Chapel’s teaching of modulating power delivery (controlling power distribution) to remove the interruption once sufficient stranded power is available (per Yang’s teachings), thus resolving an over-capacity condition. Ex.1003, ¶¶298-299.

10. Claim 29

[29.0]

See [28.0]-[28.5]. Ex.1003, ¶300.

[29.1]

Pelio-Chapel, alone and combined with Yang, renders obvious [29.1].

Pelio-Chapel teaches modulating power delivery to reduce the load in response to the variable renewable power availability (obtained BTM). *See* [1.4]; Ex.1005, [0106] (throttling load when generating output drops); Ex.1006, [0025]-[0026], [0076]; *see also* [0058], [0066], [0090]. Thus, Pelio’s and Chapel’s load throttling is from lack of power, demonstrating a determination of insufficiency that results in the throttling. Ex.1003, ¶¶301-302.

Further, Yang teaches shutting down the datacenter “whenever there is insufficient power.” Ex.1007, pp. 1, 4 (powering off when stranded power not available). The combined teachings demonstrate determining a “*ramp-down condition is met*” (Pelio’s dropping generator output, Chapel’s over-voltage

condition, and/or Yang’s power supply from stranded renewable power unavailable) that depends on Pelio’s BTM power availability, such as of a renewable power source. Ex.1003, ¶¶303-304.

[29.2]

The Pelio-Chapel-Yang combination teaches a remote controller determining that a ramp-down condition is met by a remote controller. *See* [28.4], [29.1]. As further demonstrated at [28.3], Chapel teaches a local controller (“*datacenter control system*”) modulating power delivery from policies received via from a remote controller, such as to “transition to an efficiency mode or a brownout mode.” Ex.1006, [0076]; *see also* [0058], [0066], [0090]. These “messages” from remote controller(s) render obvious “*providing a second operational directive*.” *See* [1.4]; Ex.1003, ¶¶305-306.

Yang further teaches stopping datacenter operation if renewable power availability (e.g., as stranded power) falls below a sufficient amount. Ex.1007, pp. 1, 4; *see also* [29.1]. It would have been obvious for Yang’s “shutdown” to occur from another operational directive, such as a message from Chapel’s remote controller, which provides several examples of instructions to shut down individual or multiple servers. *See, e.g.*, Ex.1006, [0058], [0066], [0076], [0090]; Ex.1003, ¶¶307-308.

[29.3]

Pelio combined with Chapel renders obvious [29.3]. *See* [1.4] (modulating power delivery by disabling power delivery per Chapel, combined with Pelio’s BTM teachings); Ex.1003, ¶309.

11. Claim 30

See claim 4. Ex.1003, ¶¶310-311.

D. Ground 3: Claims 19, 21-24, and 26 Obvious Over Pelio-Chapel-Belady

1. Belady

Belady describes “[a] data center system” with “plural data centers powered by plural power sources.” Ex.1008, 2:6-7. The power sources include “any combination of renewable power sources,” off-grid or on-grid. Ex.1008, 2:7-14. 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:




FIG. 1

Ex.1008, Fig. 1 (annotated).

Instead of moving power, Belady moves computational tasks between datacenters temporally or spatially based on power availability. Ex.1008, 7:49-60 (temporally), 7:61-8:8 (spatially). A “management system ... mov[es] data components to the sites where power is available and cost-effective.” Ex.1008, 2:15-22. Temporal movement adjusts a time of computational task execution, while spatial movement transfers a computation task for execution from one datacenter to another in a different location. Ex.1008, 2:23-26, 2:36-41, 10:20-32. The management system uses “real-time information” from the datacenters and power sources to determine power availability and needs, and to 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, ¶¶312-316.

2. Reasons to combine Pelio, Chapel, and Belady

A POSITA would have been motivated to combine Belady with Pelio and Chapel for several reasons.

First, Belady, like Pelio, focuses on power delivery to computers in a datacenter. 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 Pelio and the '456 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, ¶¶317-318.

Further, Belady is reasonably pertinent to problems including “the exponential increase in power required” for datacenter operations, wasted energy from renewable energy sources, and grid transmission and distribution costs. *See* Ex.1001, 4:24-27; 18:25-31; 18:45-57. 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, ¶319.

Second, a POSITA would have been motivated to turn from Pelio to Belady for Belady's teachings of moving computational tasks due to power availability, and how to detect and determine that it is appropriate to move. Pelio already recognized the desirability of moving pending compute jobs to another datacenter when there is insufficient renewable power at a location. Ex.1005, [0106]. Pelio also teaches that "a cloud network architecture in the Internet could be relied upon to seamlessly assume the data center jobs" to move pending jobs. Ex.1005, [0106]; *see also* Ex.1014, p. 1 (known to shift computational load between datacenters due to non-availability of intermittent power). A POSITA would be motivated to turn to Belady for implementation details regarding Pelio's transfer suggestion. Ex.1003, ¶320.

For example, Belady's management system manages the execution of computational tasks (analogous to Pelio's "data center jobs") amongst multiple datacenters based on variable power availability. Ex.1008, 2:15-22. For example, the management system moves computation tasks from one datacenter powered by a solar generating source in one location during nighttime (a solar power lull) 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 details on how to implement Pelio's movement of datacenter jobs. Ex.1003, ¶321.

Third, a POSITA would have had a reasonable expectation of success applying Belady's computational task movement teachings to the Pelio-Chapel teachings. Like Pelio, Belady teaches co-locating the datacenter with the generating source. Ex.1008, 6:53-7:17; *see also* Ex.1005, [0002]. And both Pelio and Belady recognize the utility of moving computing jobs/tasks between datacenters in response to power availability from a renewable generating source at the current datacenter. Ex.1005, [0106]; Ex.1008, 2:15-22. A POSITA would have accordingly recognized the compatibility of Pelio's and Belady's systems, with a reasonable expectation of success implementing Belady's teachings for moving tasks between datacenters. Ex.1003, ¶322.

In combination, Pelio's datacenter containers would implement Belady's real-time monitoring and computational task movement teachings by a control system (e.g., a local controller and/or remote controller as Chapel teaches, and Belady's management system teachings). Doing so would have been predictable and within a POSITA's skillset. Pelio recognized the desirability of moving jobs due to power availability issues, and Belady teaches ways in which to do so. And it was already well-known to use behind-the-meter renewable power sources for the

datacenters' tasks. Ex.1005, Abstract, [0002]; Ex.1007, pp. 1, 4; Ex.1013, p. 7; Ex.1014, p. 6; Ex.1015, pp. 3, 4; Ex.1016, p. 2; Ex.1003, ¶323.


3. Claim 19

First, Pelio-Chapel renders obvious an “*operational directive*” (Chapel’s messages from a remote/grid controller). *See* [1.4]; Ex.1006, [0076]; *see also* [0058], [0066], [0090]. The local controller at the customer premises “control[s] delivery of power to the electrical device.” Ex.1006, [0018]; Ex.1003, ¶¶325-327.

Second, Chapel’s message is a “*workload directive*” based on power availability. Chapel teaches the local controller modulating the power delivery from policies communicated via message from a remote controller. Chapel’s power delivery control includes turning off devices or operating in reduced power mode by interrupting power distribution. Ex.1006, [0025]-[0026]. Since Chapel’s directive turns on or off, or puts into reduced power mode, datacenter computers, Chapel’s directive is related to the computers’ workload and is thus a “*workload directive*.” Ex.1003, ¶328.

Third, Pelio-Chapel renders obvious modulating power delivery to the plurality computing systems based on BTM power availability at the BTM power input system. *See* [14.1]. It would have been obvious to also base the directive on projected BTM power availability for the Pelio-Chapel combined teachings to also account for future power availability. Belady confirms this. Ex.1003, ¶329.

Belady uses “real-time information provided by the data centers” about power availability from “various power sources” as well as “predict[ed] power source availability.” Ex.1008, 12:41-47 (real-time), 12:48-53 (predicted). Belady uses this information to determine how to direct operations at one or more datacenters. *See* Ex.1008, 10:60-62 (management system uses the information). Belady teaches performing the determination by a management functionality 602 on the management side. *See* Figure 6:



Ex.1008, Fig. 6 (annotated); Ex.1003, ¶330.

In combination, a POSITA would have implemented Belady’s teachings of using real-time and predicted power availability when determining the message to

send to the local controller at the datacenter for implementation, as taught by Pelio-Chapel. Indeed, Belady teaches implementing the management system functions several ways including separately from the datacenters, such as would occur with Chapel's remote controllers. *See* Ex.1008, 12:20-30, 8:15-45; Ex.1003, ¶¶331-333.

4. Claim 21

[21.0]

See Ground 1, [1.0]-[1.4]. Ex.1003, ¶334.

[21.1]

First, Pelio generally teaches a plurality of flexible datacenters, each with its own plurality of computing systems. *See* [1.3]; Ex.1005, [0048], [0049] (multiple containers), [0071]; Ex.1003, ¶¶335-336.

Second, it would have been obvious to transfer workload between flexible datacenters on the same premises or different locations, and thus transfer between different computing systems of those datacenters. Pelio teaches different computing systems assuming the datacenter jobs of a container where there is insufficient renewable power available. Ex.1005, [0106]. This alone renders [21.1] obvious. Ex.1003, ¶337.

Belady provides implementation details for transferring workload. The “computational capacity [of a collocated datacenter] has an undulating shape **that reflects the availability of power throughout a 24 hour cycle**” from the

renewable power source. Ex.1008, 9:19-21. This is because computational capacity is available from a solar power source when sunlight is available (during the daytime), while at nighttime sunlight and, thus, power, is not available. Ex.1008, 9:21-24; Ex.1003, ¶¶338-339.

The “actual computational needs 204” does not always match the capacity (and, thus, power source power output). Ex.1008, 9:24-29. To address this, Belady teaches “location-based (i.e., spatial) movement of data components” (a data component being any part of a computational task). Ex.1008, 9:14-16, 6:37-38. In an example, a controller determines to shift computational tasks between datacenters at different locations (e.g., X and Y) based on the different periods of excess power generation at each location. Ex.1008, 10:14-34; Figure 3:




FIG. 3


Ex.1008, Fig. 2 (annotated); Ex.1003, ¶340.

Belady sends data components from location X to location Y when there is a lull at location X corresponding to a time of excess generation at location Y (“movement 310” in Figure 3). Ex.1008, 10:25-28. Similarly, during lulls in capacity at location Y data components are sent to location X (movement 312). Ex.1008, 10:29-32. In combination, a POSITA would have implemented Belady’s data movement teachings with Chapel’s datacenter control teachings. Ex.1003, ¶¶341-343.

[21.2]

First, Pelio teaches a plurality of flexible datacenters, each with its own plurality of computing systems. *See* [1.3]; Ex.1005, [0048], [0049] (multiple containers), [0071]. Pelio’s plurality of datacenters includes a “*first*” as well as a “*second*” datacenter, each with a respective plurality of computing systems. Ex.1005, [0071]; Ex.1003, ¶¶344-345.

Second, Belady renders obvious moving data components between two datacenters. *See* [21.1]; Ex.1008, 10:25-28; FIG. 1:



Ex.1008, Fig. 1 (annotated).

In combination, it would have been obvious for each flexible datacenter to move data components from a first flexible datacenter to a second flexible datacenter, as taught by Belady. Ex.1003, ¶¶346-348.

5. Claim 22


[22.0]

See Ground 1 [1.0]-[1.4]; and [21.0]-[21.2]. Ex.1003, ¶349.

[22.1]

In addition to Pelio's teaching of a datacenter coupled behind-the-meter to a generating station, Belady specifically teaches a first datacenter (a first flexible

datacenter in the Pelio-Belady combined teachings) coupled to a wind power source 112. Ex.1008, 6:43-67; 7:31-40; Ex.1003, ¶¶350-351.




Ex.1008, Fig. 1 (annotated).

Consistent with Pelio's behind-the-meter teaching, Belady also teaches co-locating the datacenter with the renewable power source, with examples of the power source also connected to the grid. Ex.1008, 6:53-7:17. For example, datacenter A in Belady's Figure 1 example receives power from renewable power source 112 (which includes operation as a BTM power source). Ex.1008, 7:31-40; Ex.1003, ¶¶352-353.

[22.2]

Belady teaches a second datacenter (a second flexible datacenter in the Pelio-Belady combined teachings) receiving power from the “solar power source 114,” which is different from wind power source 112. Ex.1008, 6:43-67; 7:31-40; Ex.1003, ¶¶354-355.



Ex.1008, Fig. 1 (annotated).


Thus, Belady teaches datacenter B receiving power from renewable power source 114 for datacenter B’s computing systems to operate using BTM power from the source 114. Ex.1008, 6:53-7:17, 7:31-40; Ex.1003, ¶¶356-357.

6. Claim 23

First, the Pelio-Chapel combination renders obvious the “*datacenter control system*” operating based on BTM power availability. *See* [1.4]; Ex.1003, ¶¶358-360.

Second, the Pelio-Chapel-Belady combination renders obvious transferring workload between pluralities of computing systems at different flexible datacenters. *See* [21.1]; Ex.1003, ¶361.

Third, Belady further teaches transferring workload between datacenters in response to detecting a change in BTM power availability at a datacenter (including a plurality of computing systems). This includes transferring data components from location X to location Y during a generation lull at location X, corresponding to excess generation at location Y, by movement 310. Ex.1008, 10:25-28. This also includes transferring data components from location Y to location X during a generation lull at location Y corresponding to excess generation at location X, by movement 312. Ex.1008, 10:29-32. Each point where the lull falls below actual computational need at a datacenter illustrates a time of “*change in BTM power availability*” at that datacenter. *See* Ex.1008, FIG. 3:



Ex.1008, Fig. 2 (annotated); Ex.1003, ¶¶362-363.


7. Claim 24

Pelio-Chapel alone, and further combined with Belady, renders obvious claim 24.

First, Pelio-Chapel renders obvious a remote master control system positioned remotely from flexible datacenters. *See* [1.4] (remote controller). It would have been obvious for Chapel's local controller and remote controller to be implemented together, since that is nothing more than the routine implementation of a controller at a remote location, a well-known alternative. Thus, Pelio-Chapel alone renders this limitation obvious. Ex.1003, ¶¶364-366.

Second, [24.1] is further obvious combined with Belady, which teaches that

management functionality 602 “can be implemented by the management system 120,” distinct from the individual datacenters by a separate system. Ex.1008, 12:20-30, 8:15-45.



Ex.1008, Fig. 1 (annotated); Ex.1003, ¶¶367-368.

8. Claim 26

See [23.1], [21.1]; Ex.1003, ¶¶369-370.

E. Ground 4: Claim 9 Obvious Over Pelio, Chapel, and Rasmussen

1. Rasmussen

Rasmussen describes the problem of “IT loads” of a datacenter exhibiting “leading power factor” that can become too large over time, since the generator

source can only “tolerate a small amount of leading power factor.” Ex.1010, p. 2.

Datacenters can cause the generator source to trip off line even when the load is “well within the generator ratings” when “IT load becomes directly supplied by the generator,” e.g. a power path to the load without intervening UPS protection.

Ex.1010, p. 2. While “IT equipment” (e.g., servers) power supplies can “eliminate harmonic currents [lagging power factor],” they were not intended to fix leading currents (leading power factor from capacitive loads), which is what affects the generators. Ex.1010, p. 4. To mitigate the capacitive loads’ impact on the associated generator, IT load (e.g., one or more servers) are removed from the generator so that “sources of out-of-phase current are removed.” Ex.1010, p. 14; Ex.1003, ¶¶371-375.

2. Reasons to Combine Pelio, Chapel, and Rasmussen

A POSITA would have been motivated to combine Pelio and Chapel with Rasmussen for several reasons.

First, Rasmussen is in the same field of power delivery to computers in a datacenter, describing issues related to power delivery from a generator to a server IT load in a datacenter. Ex.1010, p. 2. Thus, Rasmussen is in the same field of endeavor. Ex.1003, ¶¶376-377.

Rasmussen is also reasonably pertinent to problems including how to control power at the datacenter and ensure stability for the generator. *See* Ex.1001, 4:24-

27; 18:25-31; 18:45-57. Rasmussen teaches controlling power at the datacenter by removing IT devices as load (e.g., by turning them off). Ex.1010, p. 14. By implementing Rasmussen's teachings, a POSITA would solve the problems of controlling power and ensuring stability. Rasmussen is therefore reasonably pertinent. Ex.1003, ¶378.

Second, a POSITA would have been motivated to combine Rasmussen's teachings for correcting power factor issues with Pelio-Chapel's power delivery control teachings to ensure the generator's long-term stability. For example, the servers' power supplies create out-of-phase reactive currents, causing issues when the datacenter's server IT load increases occurs much later, potentially many years after the datacenter is finished. Ex.1010, p. 9. By implementing Rasmussen's mitigation teachings of selectively removing IT devices from the load, a POSITA would beneficially prevent the generator from exceeding its allowed power factor limit. *See* Ex.1010, pp. 9, 14; Ex.1003, ¶379.

Third, a POSITA would have had a reasonable expectation of success combining Rasmussen's teachings with the Pelio-Chapel teachings. The Pelio-Chapel combination results in controlling a datacenter container (per Pelio) by local and/or remote controllers to turn off/on one or more servers (per Chapel), consistent with Rasmussen's teachings of removing IT load when necessary. In fact, Rasmussen's teachings simply provide another reason for Pelio-Chapel's

control teachings to act upon (power factor correction for generator protection).

Ex.1003, ¶380.

Fourth, combining Rasmussen’s teachings with Pelio-Chapel’s teachings would have also been predictable and within a POSITA’s skillset. Chapel already recognizes the benefit of monitoring the power supply and changing the power status of one or more datacenter servers. *See* Ex.1006, [0002], [0047]. A POSITA would have recognized that removing load from a generator (per Rasmussen as well as Chapel) means controlling the power supply by reducing the load. It would have been a trivial matter to monitor the power supply, per Chapel’s teachings, and control power delivery to one or more servers (again per Chapel) to correct for power factor (per Rasmussen). Doing so would beneficially provide generator stability by preventing the power factor from exceeding an allowed limit. *See* Ex.1010, p. 14; Ex.1003, ¶381.

3. Claim 9

Datacenters can cause the generator source to trip off line even when the load is “well within the generator ratings” when “IT load becomes directly supplied by the generator,” e.g. a power path to the load without intervening UPS protection. Ex.1010, p. 2. This is due to the leading power factor that IT loads (e.g., servers) exhibit. Ex.1010, p. 2; Ex.1003, ¶¶383-386.

Rasmussen teaches mitigating the leading power factor problem on

generator operation by removing servers as IT loads from the generator. Ex.1010, p. 14. Doing so removes “sources of out-of-phase current.” Ex.1010, p. 14.

Rasmussen’s teaching of removing IT load from the generator to implement power factor correction is the very thing that the ’456 patent contemplated for achieving “the power factor correction required by the local station.” Ex.1001, 8:33-40; *see also* 9:3-12 (monitored power availability includes power generation subject to power factor correction); Ex.1003, ¶¶387-388.

In combination, Pelio-Chapel-Rasmussen renders obvious controlling power delivery to computing systems based on an operational directive (per Pelio-Chapel, *see* [1.4]) based on the monitoring of the BTM power for power factor correction needs, per Rasmussen. Ex.1003, ¶¶389-390.

F. Ground 5: Claim 9 Obvious Over Pelio, Chapel, and Rasmussen

1. Corhodzic

Corhodzic describes “electrical power distribution in data centers.” Ex.1011, 1:5-8. Power arrives at the datacenter from an electric utility. Ex.1011, 6:28-31. Racks of the datacenter are powered by three phase AC power from an electric utility generator. Ex.1011, 6:46-50. The datacenter distributes “one phase voltage and a neutral line” to each rack from the three-phase power. Ex.1011, 7:16-18. Each rack tray receives the same phase voltage line (e.g., phase A voltage), while

other racks receive other phases, such as a phase B voltage or phase C voltage.

Ex.1011, 7:21-36; Ex.1003, ¶¶391-393.

2. Reasons to Combine Pelio, Chapel, and Corhodzic

A POSITA would have been motivated to combine Pelio and Chapel with Corhodzic for several reasons.

First, Corhodzic is in the same field of power delivery to computers in a datacenter. Corhodzic relates to “electrical power distribution in data centers,” and particularly providing different phases of AC voltage to different racks in the datacenter. Ex.1011, 1:5-8, 6:46-50, 7:21-36. Thus, Corhodzic is in the same field of endeavor. Ex.1003, ¶¶394-395.

Corhodzic is also reasonably pertinent, providing “an integrated uninterruptible power system in which only one AC-to-DC rectification occurs between the utility AC grid” and the processing loads. Ex.1011, 1:54-58. Prior approaches did not use a “central UPS ... to condition power or provide back-up power to data center 104,” relying instead on a “local UPS” for each load. Ex.1011, 6:1-13. Corhodzic’s teachings integrate the UPS with the DC load so only one rectification occurs between the power source and the load, improving power conversion efficiency. Ex.1011, 2:20-36. Datacenters implementing Corhodzic’s teachings would benefit from the improved efficiency of power

conversion—the loss from more than one rectification can result in “millions of dollars of electricity annually.” Ex.1011, 1:40-46; Ex.1003, ¶396.

Second, it was well-known for datacenters to receive 3-phase power from a generating source, and for the datacenter to distribute one or more phases of the received power to different racks and associated computing systems, which Corhodzic confirms. A POSITA would have been motivated to turn to Corhodzic for details about receiving and distributing 3-phase power, including receiving 3-phase power from a generator and distributing the phases to different racks. Ex.1011, 6:28-31, 6:46-50, 7:16-18. 7:21-36. It would have been obvious to supplement Pelio-Chapel’s teachings with Corhodzic’s teachings. Ex.1003, ¶397.

Third, a POSITA would have had a reasonable expectation of success combining Corhodzic’s teachings with the Pelio-Chapel combined teachings. Per the Pelio-Chapel combination, a datacenter container (per Pelio) receives power from a generator power source, and the computing systems therein operate from the power distributed within the datacenter container from the source, consistent with Corhodzic’s teachings of receiving 3-phase power and distributing the phases of power to the different computing systems within the datacenter. Combining Corhodzic’s power distribution teachings would have thus been predictable and within a POSITA’s skillset, since Pelio and Chapel already taught distributing power to the computing systems. Ex.1003, ¶398.

Finally, a POSITA would have been motivated to turn to Corhodzic’s phase distribution teachings, and integrating the UPS with the load, for Corhodzic’s benefits of improving “electrical power conversion efficiency,” resulting in “substantial energy savings” extending to reduced conductor requirements and thermal load. Ex.1011, 2:20-28; Ex.1003, ¶399.

3. Claim 15

First, Pelio discloses a BTM power input system. *See* [1.1]; Ex.1003, ¶¶400-403.

Second, Pelio renders obvious the BTM power input system being configured to receive three-phase AC voltage from the collocated BTM power source with Pelio’s description of “phase conductors” of AC circuits in the datacenter. Ex.1005, [0016]. These “phase conductors” demonstrate plural phases, of which three phases is a well-known option. Ex.1003, ¶404.

Third, Corhodzic confirms how a POSITA would have understood Pelio’s teachings and further renders the limitation obvious. Corhodzic teaches that a datacenter receives “3-phase AC power” from the grid. Ex.1011, 10:38-40; 10:4-11. In combination, Pelio-Corhodzic teaches Pelio’s BTM power input system receiving three-phase voltage (per Corhodzic), which would be from Pelio’s BTM source. Ex.1003, ¶¶405-406.

4. Claim 16

First, Pelio discloses a BTM power input system, and distributing the power to the datacenter. *See* [1.1]; *see also* Ex.1005, [0079] (A-side/B-side power feeds). Per the Pelio-Chapel combination, this includes a local controller. *See* [1.4]; Ex.1006, [0056]; Ex.1003, ¶¶407-409.

Second, Corhodzic further renders obvious providing a single-phase AC voltage from the BTM power input system. In Corhodzic, “one phase voltage” is distributed to each rack in a datacenter. Ex.1011, 7:16-17. This includes different trays receiving different ones of the multiple phases of voltage. Ex.1011, 7:21-36. Corhodzic teaches accomplishing this with a “power distribution architecture.” Ex.1011, 10:55-67; Ex.1003, ¶410.

In combination, Pelio’s BTM power input system provides single-phase voltage (per Corhodzic), which would be from Pelio’s BTM source. It would have therefore been obvious that the datacenter control system (per Chapel, part of the datacenter) would receive the single-phase voltage per Corhodzic. Ex.1003, ¶¶411-412.

IX. DISCRETIONARY DENIAL IS INAPPROPRIATE

A. No basis for 35 U.S.C. § 325(d) denial

None of Pelio, Chapel, Yang, Belady, Rasmussen, or Corhodzic were cited or considered by the Examiner during prosecution of the ’456 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: May 13, 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 '456 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 flexible datacenter comprising:

[1.1] a behind-the-meter (BTM) power input system, wherein the BTM power input system is configured to receive power from a power generation station prior to the power undergoing step-up transformation for transmission to a power grid;

[1.2] a power distribution system;

[1.3] a plurality of computing systems; and

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

[2.0] The flexible datacenter of claim 1,

[2.1] wherein the datacenter control system is configured to receive the operational directive from a remote master control system.

[3.0] The flexible datacenter of claim 2,

[3.1] wherein the operational directive depends on BTM power availability at the flexible datacenter.

[4.0] The flexible datacenter of claim 3,

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

[5.0] The flexible datacenter of claim 3,

[5.1] wherein BTM power availability comprises: one or more of excess local power generation at a local station level, 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 BTM power generation to power the flexible datacenter.

[6.0] The flexible datacenter of claim 3, wherein BTM power availability comprises:

[6.1] excess local power generation at a local station level.

[7.0] The flexible datacenter of claim 3,

[7.1] wherein BTM power availability comprises: local power generation subject to economic curtailment.

[8.0] The method according to claim 3,

[8.1] wherein BTM power availability comprises: local power generation subject to reliability curtailment.

[9.0] The flexible datacenter of claim 3,

[9.1] wherein BTM power availability comprises: local power generation c to power factor correction.

[10.0] The flexible datacenter of claim 3,

[10.1] wherein BTM power availability comprises: low local power generation.

[12.0] The flexible datacenter of claim 3,

[12.1] wherein BTM power availability comprises: transient local power generation situations.

[14.0] The flexible datacenter of claim 1,

[14.1] wherein the datacenter control system is further configured to modulate power delivery to the plurality computing systems based on BTM power availability at the BTM power input system.

[15.0] The flexible datacenter of claim 1,

[15.1] wherein the BTM power input system comprises: an input configured to receive three-phase BTM AC voltage from the power generation station.

[16.0] The flexible datacenter of claim 15,

[16.1] wherein the BTM power input system is configured to provide a single-phase AC voltage to the datacenter control system.

[17.0] The flexible datacenter of claim 1,

[17.1] wherein the datacenter control system controllably enables or disables power delivery to respective computing systems of the plurality of computing systems.

[18.0] The flexible datacenter of claim 1,

[18.1] [w]herein the datacenter control system is positioned remotely from the plurality of computing systems.

[19.0] The flexible datacenter of claim 1,

[19.1] wherein the operational directive is a workload directive based on a current BTM power availability and a projected BTM power availability at the flexible datacenter.

[20.0] The flexible datacenter of claim 1,

[20.1] wherein the datacenter control system is collocated with the plurality of computing systems.

[21.0] The flexible datacenter of claim 1,

[21.1] wherein the datacenter control system is further configured to transfer a workload from the plurality of computing systems to a second plurality of computing systems, and

[21.2] wherein the second plurality of computing systems is positioned at a second flexible datacenter.

[22.0] The flexible datacenter of claim 21,

[22.1] wherein the flexible datacenter is coupled to a first power generation station such that the plurality of computing systems operates using BTM power from the first power generation station; and

[22.2] wherein the second flexible datacenter is coupled to a second power generation station such that the second plurality of computing systems operates using BTM power from the second power generation station.

[23.0] The flexible datacenter of claim 22,

[23.1] wherein the datacenter control system is configured to transfer the workload from the plurality of computing systems to the second plurality of computing systems in response to detecting a change in BTM power availability at the plurality of computing systems.

[24.0] The flexible datacenter of claim 23,

[24.1] wherein the datacenter control system is a remote master control system that is positioned remotely from the flexible datacenter and the second flexible datacenter.

[25.0] A system comprising:

[25.1] a first flexible datacenter that comprises:

[25.2] a behind-the-meter (BTM) power input system, wherein the BTM power input system is configured to receive power from a power generation station prior to the power undergoing step-up transformation for transmission to a power grid;

[25.3] a power distribution system; and

[25.4] a plurality of computing systems configured to receive power from

the BTM power input system via the power distribution system;

[25.5] a second flexible datacenter; and

[25.6] a routing control system configured to modulate power delivery to the first flexible datacenter and the second flexible datacenter.

[26.0] The system of claim 25,

[26.1] wherein the routing control system is configured to transfer a workload from the plurality of computing systems to a second plurality of computing systems at the second flexible datacenter in response to detecting a change in BTM power availability at the first flexible datacenter.

[27.0] The system of claim 25,

[27.1] wherein the routing control system is positioned remotely from the first flexible datacenter and the second flexible datacenter, and

[27.2] wherein the routing control system is configured to modulate power delivery to the first flexible datacenter and the second flexible datacenter based on BTM power availability at the first flexible datacenter and the second flexible datacenter.

[28.0] A method of dynamic power delivery to a flexible datacenter using behind-the-meter (BTM) power comprising:

[28.1] monitoring, by a control system, BTM power availability at the flexible datacenter,

[28.2] wherein the flexible datacenter includes a BTM power input system configured to receive power from a power generation station prior to the power undergoing step-up transformation for transmission to a power grid;

[28.3] determining that a ramp-up condition is met, wherein the ramp-up condition depends on BTM power availability at the flexible datacenter; and,

[28.4] based on determining that the ramp-up condition is met, providing an operational directive to a datacenter control system at the flexible datacenter,

[28.5] wherein the operational directive instructs the datacenter control system to enable BTM power delivery to one or more computing systems at the flexible datacenter.

[29.0] The method of claim 28, further comprising:

[29.1] determining that a ramp-down condition is met, wherein the ramp-down condition depends on BTM power availability at the flexible datacenter; and

[29.2] based on determining that the ramp-down condition is met, providing a second operational directive to the datacenter control system at the flexible datacenter,

[29.3] wherein the second operational directive instructs the datacenter control system to disable BTM power delivery to one or more computing systems at the flexible datacenter.

[30.0] The method of claim 28,

[30.1] wherein the 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 13,869. 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: May 13, 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 May 13, 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. 11,016,456; Petitioners' Exhibit List; Exhibits 1001-1008, 1010-1026.

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