

Dr. Andrew A. Chien is the William Eckhardt Professor in Computer Science and Director of the CERES Center for Unstoppable Computing at the University of Chicago and a Senior Computer Scientist at Argonne National Laboratory. From 2005 to 2010, Chien was Vice President of Research of Intel Corporation, and from 1998 to 2005, the SAIC Chair Professor in Computer Science and founder of the Center for Networked Systems at UCSD. His research has been recognized for excellence with numerous awards, and also supported by the NSF, DARPA, DOE, ONR, NASA, and industry. Dr. Chien currently serves as co-chair of the editorial board of the Communications of the ACM. From 1990 to 1998, he was a Professor of Computer Science at University of Illinois. Dr. Chien earned B.S., M.S., and Ph.D. degrees at the Massachusetts Institute of Technology, and is a Fellow of the ACM, IEEE, and AAAS.

Dr. Rich Wolski is a Professor of Computer Science at the University of California, Santa Barbara, and co-founder of Eucalyptus Systems Inc. Having received his M.S. and Ph.D. degrees from the University of California at Davis (while a research scientist at Lawrence Livermore National Laboratory) he has also held positions at the University of California, San Diego, and the University of Tennessee, the San Diego Supercomputer Center, and Lawrence Berkeley National Laboratory. Dr. Wolski has led several national-scale research efforts in the area of distributed systems and is the progenitor of the Eucalyptus open source cloud project.

Fan Yang is currently a third-year Ph.D. student in the Department of Computer Science, University of Chicago. She joined the Large-Scale System Group at UChicago in 2013. She has broad research interests in parallel and distributed systems, data-intensive computing, parallel programming and optimization, large-scale system software, and computer architecture. In the past year, she worked on graph computing and benchmarking of graph-processing systems. She currently focuses on green datacenter design and stranded power in renewable electricity generation. She received her B.Eng. in computer science from National University of Defense Technology, Changsha, Hunan Province, China, in 2012.

This work was supported in part by the Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under Award DE-SC0008603 and Contract DE-AC02-06CH11357, as well as the National Science Foundation under Awards CNS-1405959, STCI-0751315, CNS-0905237, and CNS-1218808. The authors also gratefully acknowledge generous support from Hewlett-Packard, Keysight, Huawei, Nvidia, and the Seymour Goodman Foundation.



The Zero-Carbon Cloud: High-Value, Dispatchable Demand for Renewable Power Generators

Variability is an ongoing challenge to growth of large-scale renewable power generation, posing challenges for the power grid and ambitious renewable portfolio standards. The authors propose Zero-Carbon Cloud (ZCCloud), a new high-value, dispatchable demand for renewables that improves their economic viability. Initial studies show that ZCCloud can create high-value computing resources with payback periods of just a few years.

Andrew A. Chien, Richard Wolski and Fan Yang

I. Rising Renewable Power Standards: The Stranded Power Opportunity

Generation costs for renewable power are an ongoing challenge to growing adoption of higher targets for renewable portfolio standards (RPS). Among leading RPS states, California produced 20 percent of its power from

renewable sources in 2010 and has an ambitious target of 33 percent by 2020, and even Midwestern states such as Illinois¹ have adopted RPS targets that increase renewables in the electrical utilities fuel mix from 10 percent in 2015 to 25 percent in 2025. The United States federal government has established a goal of 20 percent by 2020 for all federal agencies, a dramatic increase

from the current proportion of 10 percent in 2015. A number of European nations have established even more aggressive goals (Lew et al., 2013).

The variability and non-dispatchable nature of these renewable sources, combined with low incremental generation cost, creates significant challenges for power grid design and management (Power Grid, 2014). At present, when generation exceeds demand and the excess power exceeds the grid storage's limited abilities, it is simply discarded at the source as "stranded power." Power grids call this loss of excess power "curtailment" or "down dispatching."

Numerous power grids (independent system operators, or ISOs) around the world are grappling with this issue of stranded power. Figure 1 reports data from the Midcontinent Independent System Operator (MISO) showing total generation,

total wind power, and total power "curtailed/down-dispatched" for a recent two-and-a-half-year period. Despite improved grid connectivity and management, and MISO's economic dispatching market still suffers from a few percent waste, an extraordinary amount of power – 2.2 TWh in 2014. Comparable levels prevail for ERCOT (wind) and CAISO (solar and wind), and numerous regions in Europe (Denmark, Germany, Ireland, Italy) (Lew et al., 2013). In all of these power grids, the fraction of renewables is expected to increase by 100 percent or more in the next decade, creating even greater challenges to the maintenance of power-grid balance, and more stranded power (Johnson and Moyer, 2012; Bird et al., 2013a).

In the MISO region, wind power has significant penetration today with smaller states such as Iowa and Minnesota greater than 20 percent, and larger states such as

Illinois and Michigan at 5 percent, but all have renewable power standards goals to double this percentage by 2025 (Johnson and Moyer, 2012).

The fastest-growing sources of renewable energy are wind and solar (NREL, 2014; Renewable Energy, 2015), and these resources typically produce at 20–30 percent of peak capacity. This means that achieving these significant increases in renewable fraction will mean doubling or tripling the proportion of such renewable resources in the grid, and concomitantly increasing the dynamic variation in renewable power generation to approach or even exceed peak electrical demand. For example, consider the calculation in Table 1, where we use "renewable" to refer to variable wind and solar resources. It shows the dramatic increase dynamic range that the drive for higher RPS creates, reaching 39 to 72.5 percent in the MISO and CAISO examples specifically; European countries with more ambitious RPS goals fall at the top of this range, or even above it. For example, a recent feasibility study for a CAISO RPS goal of 50 percent by 2030 projected a dynamic range of 100 percent (Energy of Economics, 2014).

The growing dynamic range of these power grids is expected to lead to investment in growing grid infrastructure to share power across larger geographic regions, and also energy storage to ensure reliable power. While these




Figure 1: Stranded Power in the MISO Power Grid, June 2011–October 2013 (Bird et al., 2013b)

Table 1: Renewables Growth under RPS and Increasing Dynamic Range; CAISO Numbers Reflect California RPS Goals.

	Renewable Fraction	Peak Renewable @ 30% Productivity	Total Energy/Year	Peak Demand	Dynamic Range
CAISO 2010	10%	8.8 GW	230 TWH	40 GW	22%
CAISO 2020	33%	29 GW	230 TWH	40 GW	72%
MISO 2015	9%	13.26 GW	389 TWH	95 GW	14%
MISO 2025	25%	36.8 GW	389 TWH	95 GW	39%

techniques will be helpful, economic realities will limit the ability of these resources to fully absorb 10 GW variations, and consequently the significant quantities of stranded power are likely to increase (Energy of Economics, 2014).

II. ZCCloud: A High-Value Dispatchable Demand for Renewable Power Generators

A. Zero-Carbon Cloud services vision

Modern cloud computing technology based on containers can be rapidly deployed in remote locations with minimal infrastructure requirements—power and networking. These systems are flexible, and can dynamically adapt to the availability of power and computational demand. The basic idea behind Zero-Carbon Cloud is to use stranded power to create a volatile computational resource. The accelerating growth of cloud computing services creates new revenue opportunities for such computing services. Thus, ZCCloud will convert

intermittent renewable power into a high-value dispatchable demand that supplies useful cloud computing services. ZCCloud envisions a *pure* renewable-based computing services, a radical contrast to greening efforts (Greenpeace, 2014) that purchase a balancing average of renewable power. ZCCloud's approach accepts volatility in power as a fundamental operating constraint, and by doing so does not require support from conventional, carbon-generating sources.

B. Zero-Carbon Cloud deployment cost-efficiencies

ZCCloud's approach is to create a system with lower fixed

costs (physical plant and server) and lower variable costs (power). The lower fixed costs allow the computing plant to be operated competitively with only intermittent power. The exploitation of low-cost renewable power at the point of generation allows older computing hardware to be used cost-effectively. Our estimates suggest that nearly seven-fold cost reductions can be achieved based on industry-standard models of TCO (Barroso and Holzle, 2009; Patel and Shah, 2005), as shown in Figure 2.

Key strategies for cost reduction include: (1) eliminate buildings: use of containers and direct siting at renewable generation sites, (2) exploit old hardware, and (3)




Figure 2: Zero Carbon Cloud Reduced Costs: Containers (C), Used Servers (U) and Stranded Power




Figure 3: Performance, Cost and Perf/Cost. Depreciation (Left), Deprec. + End of Moore's Law (Right)

exploit low-cost excess renewable power.

C. Exploit end of Moore's Law: old hardware

The end of Dennard scaling (2005) has slowed computing improvement (single-thread). Unimproved software has experienced a modest 21 percent per year performance increase; three times slower than the 64 percent average for 1985–2005. Consequently, the performance/cost of old servers is now better than new hardware (Figure 3). Considering the growing technology costs for new processors (Intel, 2014) (see right), three-year old hardware has a five times advantage, and four-year old hardware a nearly 25 times advantage. Commercial drivers for new systems are energy efficiency and density, particularly because power can be as much as 40 percent of the server total cost of ownership (TCO). Thus, effective use of older hardware could significantly reduce capital costs associated with cloud computing.

Physical plant costs. Published data suggests data centers' physical plant costs contribute 25–50 percent of TCO costs (Barroso and Holzle, 2009; Patel and Shah, 2005). Using containers, directly and siting them at renewable generation sites, ZCCloud eliminates the need for purpose-built buildings to house the infrastructure and power distribution.

Server costs. Studies show that server costs make significant contributions to data center TCO (25–50 percent), and together with physical plant costs account for 75 percent of TCO (Barroso and Holzle, 2009; Patel and Shah, 2005). Exploiting the end of Moore's law, and the slowing improvement of computing technology enables ZCCloud to exploit older hardware, but deliver high-capability cloud services, cost-effectively. This benefit can be exploited aggressively by use low-cost renewable power enabling use of less energy-efficient older hardware.

Low-cost stranded renewable power. A primary element of the

remaining 25 percent is electricity costs. ZCCloud will exploit stranded power, reducing the cost of power well below retail power costs and even the wholesale prices paid by large data centers 10-fold. Further, the low-cost power will enable ZCCloud to employ older hardware that is less energy-efficient but computationally effective.

Together these improvements suggest a system that could provide computing as much as seven-fold cheaper than traditional methods, when stranded power is available.

III. Low Initial Cost and Incremental Scalability

The ZCCloud approach is both low-cost to initiate and incrementally scalable. ZCCloud can be initiated with a single container, populated with 10 racks of servers, each filled with 40 dual-processor Xeon servers, as detailed below.

Design of ZCCloud Building

Blocks: ZCCloud uses convention building blocks that achieve high

Table 2: ZCCloud Computing Container Sketch Design.

	Description	Performance	Power
Node	Dual sockets, Intel Xeon E5-2660 v3 CPUs	960 GFLOPS	0.125 kW
Rack	40 nodes per rack	38.4 TFLOPS	5 kW
Container	10 racks per container	0.384 PFLOPS	50 kW

Table 3: Initial and Scaled Scenarios for ZCCloud.

ZCCloud	Racks	Servers	CPUs	Total Teraflops	Total Memory	Power	Estimated Cost
Base (1-container)	10	400	800	384 Tflops	51 TBytes	50 kW	\$2 M to \$3 M
Medium (4-container)	40	1,600	3,200	1,536 Tflops	204 TBytes	200 kW	\$8 M to \$12 M
Large (32-container)	320	12,800	25,600	12,240 Tflops	1632 TBytes	1,600 kW	\$64 M to \$98 M

densities of computing per rack and per container. The computing nodes are connected with low-latency, high-speed 10-gigabit Ethernet switches. We assume the containers have a power usage effectiveness (PUE) of 2.0, which means the non-compute facilities, e.g. cooling system, has the same power consumption as compute nodes. This number is conservative with published commercial PUEs ranging from 2.0 down to 1.19. To enable real-time response even when stranded power is unavailable, ZCCloud also deploy an always-on front-end server for each container, the power consumption of which is nontrivial compared to the total container power. The resulting power and computing density for a single container is summarized in [Table 2](#).

We show cost estimates and scaling for ZCCloud in [Figure 4](#). A single-container ZCCloud could be initiated at a capital cost of several million dollars. Such a level of investment

is a modest hurdle compared to transmission line and generation projects that cost hundreds of millions of dollars.

ZCCloud can be scaled incrementally by adding containers to significant scale. For example, the medium-scale ZCCloud (replicating four times) has 1.6 teraflops and 204 terabyte of memory capacity and consumes only 200 kW of power. Pushing ZCCloud to a large scale yields a 32-container system that achieves 12.24 petaflops and 1.6 petabytes (on the scale of 2014 world-class supercomputer). Remarkably, the large ZCCloud system consumes only 1.6 megawatts, about the

generating capacity of the average wind turbine being installed today ([Tegen et al., 2013](#)).

Even a large ZCCloud can be installed in a moderate-size wind farm of 100 MW or less, and hardly tax its capacity ([Table 3](#)). For example, the Twin Groves farm ([Twin Groves, 2014](#)) includes 240 turbines, each 1.65 MW for peak generating capacity of 398 MW. It produces 1.3 terawatt-hours (1.3B kWh) annually. As shown in [Figure 4](#), a 7 petaflop, computing system could be powered by a single 1.6 MW turbine. One such would cover only 15% of a large wind farm such as Twin Groves, and there are dozens of such facilities



Figure 4: Scaling from **Small:** 1-container, 0.2 Petaflops, 0.05 MW, **Medium:** 4-container, 0.8 Petaflops both can be power by a single turbine. **Large:** 32-container, 6.4 PF, 1.6 MW, 1% of turbines in a wind farm

in the MISO region, so total capacity exceeding exaflops is possible.

To exploit renewable generation, an intermittent power source, ZCCloud's computing services will rapidly start up and shut down. This will enable them to respond on short time scales to the availability of stranded power. It is important to note that while renewables are intermittent generators, this term should be understood in the context of a power grid that has regular requirements, varying with human and industrial activity on a daily and weekly basis. Wind power generation varies with regional and local weather patterns (hours, days, weeks) and solar power generation varies with the daily cycle and, at a smaller scale, with weather patterns on a weekly scale. The ZCCloud building block offers intermittent and variable capacity based on the availability of stranded power. Volatility tied to wind power will have continuous available from hours (overnight) to days, due to the change in weather patterns. Volatility tied to solar stranded power appears to be likely tied to regular daily and weekly cycles, but involves shorter daylight periods.

ZCCloud will meet the technical challenges to create a dispatchable-load computing resource, demonstrating (1) a new capability to harvest at large-scale the intermittent power for computing (capability), and

subsequently (2) to harvest efficiently, deriving benefit from a large fraction of the available intermittent power for computing (efficient capability).

IV. Economic Viability

Commercial uses and markets for volatile computing resources



exist. Large-scale cloud providers, like Amazon AWS, offer "spot instances" – rentals at a bid "spot price" that are terminated when the market (or perhaps the provider) decides that they should be reclaimed. In short, they can be revoked at any time, yet still are deemed useful by a large user community (Ben-Yehuda et al., 2013).

To explore economic viability, we compare ZCCloud's volatile computing resources to Amazon's *Spot instances* model (Ben-Yehuda et al., 2013; Amazon, 2015a) that supports high-throughput computing jobs. While a comparison to supercomputing center operating

costs might be more direct match, the economics of most such centers are complicated by government support, and large users populations who do not directly pay for use of the resources. So we choose instead to compare to Amazon spot instances, an open and broadly used market with direct payment by most customers. Of course, more advanced service offerings are possible (Guo et al., 2015; He et al., 2015), and they would likely increase significantly the value of ZCCloud resources. We focus on simple volatile instances here in order to provide a lower bound.

A. An example of revenue and expenses

To illustrate the potential commercial viability of ZCCloud, we provide an example of the possible revenue and expenses associated with a 1 MW deployment. We assume that the computational infrastructure (servers, power distribution, and cooling) will be housed in self-contained "containerized" data centers (ASTModular, 2015).

We use the *m3.large* instance type from Amazon's AWS as the unit of computation (Amazon, 2015b). An *m3.large* instance contains two cores from an Intel Xeon E5-2670 v2 processor (2.5 GHz), and 7.5 GB of memory. The E5-2670 processor is almost three years old and no longer readily available in a server form factor. For expense calculation, we use the current Intel Xeon

E5-2660 v3 (2.6 GHz) in a two-CPU, 1U form factor. This CPU model supports 12 cores, making it possible to deploy 24 cores per 1U. We assume a rack capacity of 40U, and 10 racks per containerized data center. With these specifications, 1 rack of servers accounts for approximately 7 kW of power when fully committed. Thus 1 MW of power is equivalent to 142.85 racks worth of servers in terms of draw.

B. Revenue

One server, with two CPUs, can support 12 *m3.large* equivalents. One rack thus supports 480 *m3.large* equivalents and a 1 MW deployment (142.85 racks) is equivalent to 68,568 *m3.large* instances. Using a trace of AWS spot prices gathered between January and June of 2015, the average spot price in the *us-east-1* region for an *m3.large* instance was 0.018/hour of occupancy (rounded to the nearest hour). Thus a 1 MW deployment, at 100 percent utilization generates $0.018 \times 68568 = \$1234.22/\text{hour}$ or $\$10,794,247.20/\text{year}$ at the current rate charged by Amazon for an *m3.large* spot instance.

Current utilization rates for public clouds are held as trade secrets. However using rates observed for commercial private clouds (Wolski and Brevik, 2014), the estimated revenues for a 1 W system in this example configuration are likely between \$5.3 million (50 percent




Figure 5: Number of Years until Break-Even Based on Utilization and Server Cost

utilization) and \$7.5 million (70 percent utilization) annually.

C. Expenses

Retail cost for a single containerized data center (without servers) is approximately \$500,000. With 10 racks/container (dictated by power distribution capacity), a 143-rack system requires 15 containers or a total capital expense of \$7.5 million for the containerized data centers. The retail cost of a 1U server in this configuration is approximately \$5,500, making the total “capex” \$38.9 million if the 1 MW system were purchased at retail prices in the summer of 2015. Typically, however, servers purchased in this volume can be obtained at a substantial discount (we are unsure about discounts for the containers). Thus, a more likely server cost is approximately \$3,000 per server, making the total capex \$24.6M. Thus capex is

equivalent to revenue is 3.6 years (at full utilization) at retail prices, and 2.3 years if expected volume discounts are available for the servers.

Based on our experiences with private clouds, the personnel associated with the operation of this system is likely to be 2.5 full-time employees (FTE) at approximately \$200,000 apiece fully burdened. Thus the “opex” associated with this hypothetical 1 MW deployment is \$500,000/year.

Figure 5 depicts the time until revenues equal capex as a function of average utilization and server cost taking into account opex.

After this break-even time, the 1 MW system approximately generates between \$5 million and \$10 million in positive cash flow. This analysis does not take into account several financial factors, however.

It does not include insurance, depreciation, taxation, or

regulatory expenses. It also does not include the cost of conditioning power before it is consumed by the containerized data centers. Finally there is likely a set of expenses associated with connecting the facility to the common carrier Internet. The cost of network usage is neutral relative to the AWS pricing (we assume the cost will be similar) but size connectivity may incur an installation cost.

We have also used current equipment that is equivalent to what is offered by AWS for a specific utilitarian instance type. It is possible to reduce (drastically) the server cost by using older hardware. For example, surplus servers circa 2009 are available for \$300 per 1U. These servers will not produce the performance equivalent to an *m3.large* but might prove to be viable at some lower price point. Without a good price reference, however, it is difficult to estimate what the revenue generation from this lower-performing hardware might be in the marketplace.

On the revenue side, we have made no attempt to synthesize higher-value services, that either increase continuity (decrease volatility) (Guo et al., 2015; He et al., 2015) or capitalize on the ability to schedule parallel computations of significant scale. By comparison, Amazon's "on-demand" instances can command as much as 10 times higher prices per hour, potentially reducing these payback periods by a corresponding factor. We leave

study of these possibilities to future research.

V. Summary and Future Work

ZCCloud posits the creation of a new, high-value dispatchable demand for the power grid. Such demands are increasingly important in face of rising



renewable power standards (RPS) that increase the volatility of grid power generation mix to dynamic ranges to 40 percent and more. ZCCloud can be dispatched to exploit short (hours) and medium-term (days) power excess, and produce high value computational resources. ZCCloud can be deployed at low initial cost, scaled incrementally, and have payback periods as short as three years.

Many challenging questions remain. Does the geographic and temporal structure of stranded power make it practically accessible? How much of the benefit of lower-cost, older hardware can be captured

to reduce capital costs? What higher quality services can be synthesized with this base of volatility and compute capability? And, can this entire model become economically viable, not only self-sustaining but growing? We look forward to addressing these with the research community in the future.■

References

- Amazon elastic compute cloud (EC2). <http://aws.amazon.com/ec2/>.
- Amazon elastic compute cloud (EC2) instance types. <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instance-types.html> (accessed August 2015).
- ASTModular. <http://www.astmodular.com> (accessed August 2015).
- Barroso, L., Holze, U., 2009. *The Data Center as Computer: An Introduction to the Design of Warehouse-scale Machines*. Morgan-Claypool.
- Ben-Yehuda, O.A., Ben-Yehuda, M., Schuster, A., Tsafir, D., 2013. Deconstructing Amazon EC2 spot instance pricing. *ACM Trans. Econ. Comput.* 1 (September (3)), 16:1–16:20.
- Bird, et al., 2013a. Integrating variable renewable energy: challenges and solutions. Technical report. National Renewable Energy Laboratory (NREL).
- Bird, et al., 2013b. Wind and solar energy curtailment: experience and practices in the United States. Technical report. National Renewable Energy Laboratory.
- Energy of Economics, 2014. <http://ethree.com/>.
- Greenpeace, 2014, April. Clicking clean: how companies are creating a green internet.
- Guo, W., et al., 2015, June. Bidding for highly available services with low price in spot instance market. In: *Proceedings of HPDC 2015*.
- He, X., et al., 2015, June. Cutting the cost of hosting online services using cloud spot markets. In: *Proceedings of HPDC 2015*.

- Wall Street J., Intel server prices increased 47% from 2007 to 2014 In: <http://www.wsj.com/articles/>.
- Johnson, S., Moyer, E., 2012. Feasibility of U.S. renewable portfolio standards under cost caps and case study for Illinois, Technical report. The Center for Robust Decision Making on Climate and Energy Policy.
- Lew, et al., 2013, October. Wind and solar curtailment. In: Proceedings of the International Workshop on Large-Scale Integration of Wind Power into Power Systems as Well as on Transmission Networks for Offshore Wind Power Plants.
- NREL, 2014. <http://www.nrel.gov/docs/fy08osti/41869.pdf>.
- Patel, C., Shah, A., 2005. A cost model for planning, development, and operation of a data center. Technical Report HPL-2005-107R1. Hewlett-Packard Corporation: HP Labs.
- United States Department of Energy. <http://energy.gov/oe/downloads/future-grid-evolving-meet-america-s-needs-december-2014>.
- United States Energy Information Administration. http://www.eia.gov/forecasts/steo/report/renew_co2.cfm.
- Tegen, S., et al., 2013. 2011 cost of wind energy review. National Renewable Energy Laboratory Report, In: <http://www.nrel.gov/docs/fy08osti/41869.pdf>.
- Twin Groves, 2014. http://en.wikipedia.org/wiki/Twin_Groves_Wind_Farm.
- Wolski, R., Brevik, J., 2014. Using parametric models to represent private cloud workloads. IEEE Trans. Serv. Comput. 7 (4) 714–725.

Endnote:

1. Overall 33 of the nation's 50 states, including a long list of major industrial states across the nation.