2014 IEEE International Conference on High Performance Computing and Communications (HPCC), 2014 IEEE 6th International Symposium on Cyberspace Safety and Security (CSS) and 2014 IEEE 11th International Conference on Embedded Software and Systems (ICESS)

# Cloud Energy Broker: Towards SLA-driven Green Energy Planning for IaaS Providers

Md Sabbir Hasan, Yousri Kouki and Thomas Ledoux
*ASCOLA Research Group*
*EMN-INRIA, LINA*
*Ecole des Mines de Nantes, France*
*Firstname.Lastname@inria.fr*

Jean Louis Pazat
*MYRIADS Research Group*
*INSA, INRIA, IRISA*
*Rennes*
*jean-louis.pazat@irisa.fr*

*Abstract*—**Demand for Green services is increasing considerably as people are getting more environmental conscious to build a sustainable society. Therefore, enterprise and clients want to shift their workloads towards green Cloud environment offered by the Infrastructure-as-a-Service (IaaS) provider. The main challenge for an IaaS provider is to determine the best trade-off between its profit while using renewable energy and customers satisfaction. In order to address this issue, we provide a Cloud energy broker, which can adjust the availability and price combination to buy Green energy dynamically from the market to make datacenter green. We investigate a simplified power model from where we can formulate and predict power demand. Our energy broker tries to maximize of using renewable energy under strict budget constraint whereas it also tries to minimize the use of brown energy by capping the limit of overall energy consumption of datacenter. The energy broker was evaluated with a real workload traced by PlanetLab. Experimental results show that our energy broker successfully keeps the best trade-off.**

*Keywords*-**Cloud Computing; Green Computing; Renewable Energy; Service Level Agreement (SLA); Cross-layer SLA; Capacity Planning**

## I. Introduction

In response to the growing demand for Internet and Cloud computing services, large companies such as Amazon, IBM, Google, Yahoo!, Microsoft etc. responded greatly by making their own Cloud platforms and datacenters.It is obvious that datacenters consume enormous power that can lead to negative environmental implications (e.g., emission of several million tons of CO2 and global warming) in its life span, which is a serious concern for society and academia researchers in recent years [5]. Similar to other large consumers of power, datacenters find themselves increasingly pressured either by legislation or by public opinion to find options to reduce their carbon footprint. Therefore, demands for green products and services are ever increasing. In response, using renewable energy in the datacenter is one of the best ways to address this issue even though renewable sources are very intermittent in nature and generally incurs higher cost to produce energy.

While on-site and off-site renewable generation models are explicitly involve with datacenter to offset their carbon footprint reduction goal, some implicit model e.g., Renewable Energy Certificate (REC) and Power purchasing agreement (PPA), have created lots of attention to the datacenter owners or Cloud providers. REC, known as green certificate in Europe, is a tradable commodity proving that electricity generated using renewable sources. Therefore, purchasing of a green certificate equals to purchasing a claim that the certificate owner consumed energy from the renewable portion of the whole energy grid [12].

Our motivation is to exploit this REC market where multiple *Green Energy-as-a-Service* (GEaaS) providers will produce energy and feed to the Grid. Considering the Cloud infrastructure as a stack of XaaS layers, a Infrastructure-as-a-Service (IaaS) provider will buy a portion of green energy dynamically from those GEaaS providers to supply green computing services to the Software-as-a-Service (SaaS) providers or their clients. As renewable energy sources are very intermittent in nature, the renewable energy-feeding price would be very different from one to another provider depending on the location of site, availability of sources (Wind speed, solar irradiation etc.) and capacity factor of the plant. Committing to a single provider might result unavailability of required green energy requirement for certain time frame thus ensuring certain percentage of green energy availability in data center can not be met. On the other hand, when the generation of green energy is lowest due to weather or maintenance work in the plant, the price of energy also might go beyond the acceptable limit. Thus, providing contracted Green computing services to SaaS providers or end-users become extremely difficult for a IaaS provider.

In this paper, we propose a *Cloud energy broker*, which can adjust the availability and price combination to buy green energy dynamically from the market to

978-1-4799-6123-8/14 $31.00 © 2014 IEEE
DOI 10.1109/HPCC.2014.44

make datacenter green for a specific (as a example 30%) portion. We investigate a simplified power model from where we can formulate and predict power demand of a datacenter for next 24 hours by evaluating 7 days real data traced from PlanetLab [11]. Our monitoring window length is an hour that provides almost accurate predicted information. We also have taken a realistic consideration that Green provider can publish a day ahead green energy generation and price per hour, which is a common practice at European electricity and energy market along with smart-grid environment. Our energy broker tries to maximize of using renewable energy under strict budget constraint whereas it also tries to minimize the use of brown energy by capping the limit of overall energy consumption of datacenter. Moreover, our work is divided into two parts: planning and real time execution. We will only highlight the planning phase due to limitation of the page number. For providing Green services to the SaaS provider or client, strong Service Level Agreement (SLA) has to be addressed. Therefore, we explain different level of Service Level Objective (SLO) in each Cloud layers to realize how cross-layer SLA can be contracted in Cloud computing environment. Later, we validate our proposal using CSLA language [7] to justify our approach.

The rest of the paper is structured as follows. Section 2 briefs our proposed model and SLO in each layer. Section 3 describes the modeling details of our energy broker considering one IaaS and multiple energy providers available in the REC market. The results obtained from experimental evaluation are presented and discussed in Section 4. Section 5 provides a selection of relevant work related to this paper. Finally, Section 6 concludes this paper and provides some discussion on future work.

## II. Cross-layers SLA

The objective of this section is to present the SLA dependencies in a Cloud cross-layers architecture.

### A. Actors

The Cloud architecture is usually composed of several XaaS layers and SLAs are characterized at various levels in this stack to ensure the expected QoS for different stakeholders. As show in Figure 1, an End-User is a client of the SaaS provider, which is itself a client of the IaaS provider and as well as for Energy-as-a-Service (EaaS) provider.

In the REC market, Green Energy-as-a-Service (GEaaS) providers produce green energy and feed to the Power/Electric Grid but sell their renewable energy credits or renewable energy in a wholesale market to consumers (IaaS provider) for direct purchase. Even though IaaS provider consumes energy from the Power/Electric Grid, they have to pay directly to GEaaS providers for their consumption of certain portion of green energy which has been contracted with respective GEaaS provider. Hence, since the SLA has to be contracted between IaaS provider and GEaaS provider, the Grid monitoring infrastructure mentioned in Figure 1 is considered as supporting part (or third actor) to monitor/validate SLA between IaaS and GeaaS actors.

Except for the End-User, any Cloud layer plays a provider-consumer role: it is a provider for the upper layers and a consumer for the lower layers. Its main challenge is to maintain its consumer's satisfaction face to demand variations while minimizing the service costs due to resources fees and SLA penalties (in case of violation).

### B. SLAs

The Figure 1 presents Service Level Objectives (SLOs) examples that apply at three different Cloud levels, between the End-User and the SaaS, the SaaS and the IaaS, or the IaaS and the EaaS:

- $SLA_S$ (End-user – SaaS provider): Service Response Time, Service Availability.
- $SLA_R$ (SaaS provider – IaaS provider): Resource Availability, Green Resource (percentage of used green resource).
- $SLA_E$ (IaaS provider – EaaS Provider): Brown energy Availability, Green energy Availability.

The Listing 1 presents an example of code in CSLA[7], a SLA language to finely express SLA and address SLA violations in the context of Cloud services. CSLA allows defining SLA in any language (e.g., XML, Java); we use XML as a representation format for sake of simplicity. This code describes the guarantee terms and penalties for SLA between a IaaS provider and its customer (SaaS provider).

In this example, we focus only on one SLO about the percentage of green resource (lines 1-5). The SLO states that at least 30% of green resource should be guaranteed, with confidence, fuzziness and percentage fuzziness of 83.33%, 5% and 30%, respectively. These CSLA features (confidence, fuzziness) have been introduced to deal with QoS uncertainty in unpredictable Cloud environment [8]. In concrete terms, it means that the percentage of green resource measured within an observation period may be i) lower then 25% in 16.67% of the observation periods, ii) between 25% and 30% in 24.99% (83.33% of 30%) of the observation periods and iii) greater or equal to 30% in 58.33%. A violation of the *GreenResource* SLO implies a penalty that depends on the green percentage not respected (lines 6-13). For each penalty, a procedure (line 10) indicates the actor in charge of the violation notification (e.g., provider), the

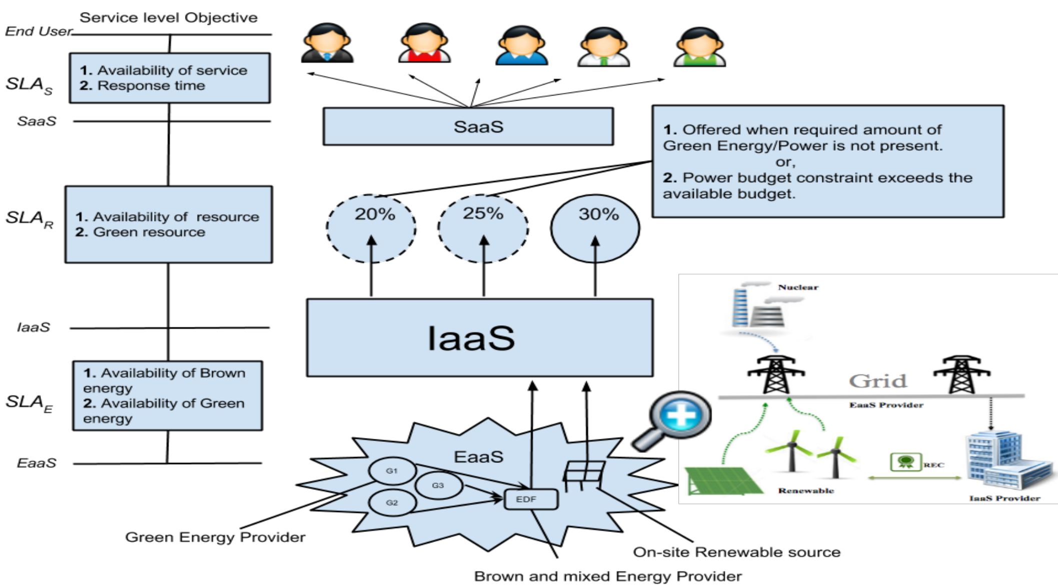


Figure 1. Cross-layers SLA

notification method (e.g., email) and the notification period (e.g., 7 days).

Listing 1. CSLA example.

```
<csla:terms>
<csla:objective id="GreenResourceSLO" actor="provider">
<csla:expression metric="Gr" comparator="gt" threshold="30" unit="
     \%" monitoring="Mon−1" Confidence="83,33" fuzziness−value=
     "5" fuzziness−percentage="30"/>
</csla:objective>
</csla:terms>
<csla:penalties>
<csla:Penalty id="p−Gr" objective="GreenResourceSLO" condition="
     violation" obligation="provider">
<csla:Function ratio="0,5" variable="GreenPercentage" unit="\%">
<csla:Procedure actor="provider" notificationMethod="e−mail"
     notificationPeriod="7 days">
<csla:violationDescription/>
</csla:Procedure>
</csla:Penalty>
</csla:penalties>
```

## III. Green SLA-driven framework

In this section, we present the modeling details of our energy broker considering one IaaS and multiple energy provider available in the REC market. The objective of this broker framework is to determine the best trade-off between IaaS provider profit while using renewable energy and customers satisfaction (i.e, respecting green SLAs).

### A. GEaaS and REC Market

In deregulated electricity market, prices vary significantly during the day depending on the energy generation and demand in the wholesale market [3]. As a consequence, most of the energy distribution company introduce different non-flat tariffs. Day-ahead pricing (DAP) is one of the many pricing methodology which is widely used in deregulated electricity pricing market. Therefore, this pricing strategy also applies for Green energy market. So, in our framework we consider different GEaaS providers, which update DAP with information including probable generation of Green energy, price and availability. Energy availability gives them relaxation of their prediction error for energy generation.

In REC market, Green energy is sold as commodity separately from electricity to multiple consumers. As, consumer (IaaS provider) has to pay for their Green energy consumption directly to GEaaS providers, hence the real SLA has to be contracted between IaaS provider and GEaaS providers and that SLA should be validated by two supporting SLA via Grid providers monitoring system (e.g. ERDF in France) mentioned in the Figure 2(b). As example from our framework, contracted $SLA_{I-G_1}$ between IaaS provider and $G_1$ (Green Energy provider) should be validated by two supporting SLA named $SLA_{I-P}$ and $SLA_{P-G_1}$ by Grid monitoring Infrastructure.

### B. Planning life cycle

The selection of a GEaaS Providers can be abstracted as a succession of operations in a planning phase (see Figure 2(a)). The complete lifecycle includes both IaaS and GEaaS providers information: forecasting power demand of IaaS, day ahead pricing (DAP) data of GEaaS providers, selection of best GEaaS provider, buying dynamically Green energy from GEaaS providers. Moreover, planning framework is divided into two time frames: hourly and daily (e.g., m hours,k days). The first phase ends with step 4, where buying Green energy is dynamically done hourly. Once GEaaS providers update DAP information, a new schedule is initiated, thus concludes the process for n hours. In addition, the second phase resolves the process by step 5 for k days.

### C. Components of Energy Framework

We present in Figure 2(b) the main components of our broker framework.

- *Information repository*: this component stores DAP information published by different GEaaS providers. The information is updated instantaneously if any change has been made at DAP information of GEaaS providers, otherwise the information is updated periodically in a given time frame.
- *Forecaster*: the amount of Green energy required for IaaS can be forecasted for the next few minutes (short time forecasting) or the next few hours (long time forecasting) based on $k$ days energy usage. Once the requirement of Green energy is forecasted the component can calculate the maximum Green power budget from the history or from IaaS provider's power budget information.
- *Optimizer*: both Information repository and Forecaster forward their information to the Optimizer component. Therefore, the Optimizer provides pareto optimal solution for dynamically selecting GEaaS provider based on respective information for each time interval, for example 1 hour.
- *SLA Negotiator*: after selecting desired GEaaS provider/providers for each time frame (1 hour), this component establishes a SLA contract between IaaS and GEaaS provider. In addition, the SLA negotiator also makes a SLA contract with the Grid infrastructure for monitoring the violation of contract in the case where Green energy is not delivered to the Grid.

### D. Planning phase in details

*1) Monitoring and Forecasting:* Predicting power demand in Cloud computing environment is very arduous as ratio of power consumption at different infrastructure (e.g. servers, cooling, lighting etc.) level are very divergent. Therefore, using Power Usage Effectiveness (PUE) helps to get better understanding about power demand of a datacenter. For a datacenter, PUE is defined as the ratio of the datacenter's total power consumption to the data center's power consumption at the computer servers [3]. Hence, if we can measure the power consumption at server level, it becomes easy to calculate the total power consumption of a datacenter for certain time frame. As CPU consumes majority of the power compared to memory in server level, in our investigation we ignore the power consumption by memory in the power model. Furthermore, future demand of power consumption can be generated by using efficient forecasting method. The output of the forecasting phase is $E$, where $E$ represents the requirement of green energy for next 'm' hours.

*2) Optimizing:* The goal of our optimization framework is to the find optimal amount of energy from GEaaS provider or composition of best GEaaS providers while respecting the budget. We address our optimization problem as Constraint programming (CP) [13], since CP accepts any type of relations to formulate constraints consisting of linear inequalities. So, variable $X_i$ represents the amount of Green energy required by IaaS provider for each time interval from $G_i$ (where, $i \in [1, ...., n]$) provider, whereas Domain $D(X_i)$ demonstrates the DAP information published by GEaaS providers.

- Variable: $X = \{X_i \mid i \in [1, ...., n]\}$
- Domain: $D(X_i) = \{1, ....., e_i\}, \forall i \in [1, ...., n]$

Therefore, we introduce our objective function which tries to maximize both the amount and the availability of Green energy to meet the exact Green energy requirement.

$$Maximize \quad : \quad (\sum_{i=1}^{n} X_i . \prod_{i=1}^{n} Av_i) \qquad (1)$$

where, $Av_i$ symbolizes the availability of $X_i$.

$$Subject\ to \quad \sum_{i=1}^{n} C_i \leq B^{max}, \quad B^{max} \in \Re_+ \qquad (2)$$

IaaS provider requires to have upper bound of budget for each interval to buy Green energy which is stated at constraint (2) as $B^{max}$, where $B^{max}$ is computed by ($E$)*($St.Price$) and $C_i$ represents the cost for buying green Energy from provider $G_i$. The term ($E$) and ($St.Price$) represent the required Green energy and average green energy price from historical window respectively.

## IV. Experiments

This section presents the results obtained from an experimental scenario used to evaluate the proposed broker. The objective is to show a real utilization case of the forecaster and the optimizer.

### A. Experimental Testbed

We consider a datacenter which has an average PUE of 1.77. Though some of the state-of-the-art techniques claim to have reduced this value closer to 1.20, still

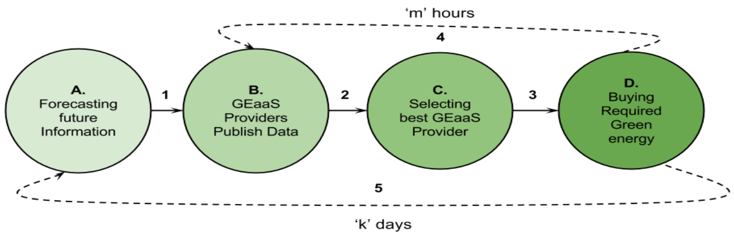


(a) Life Cycle

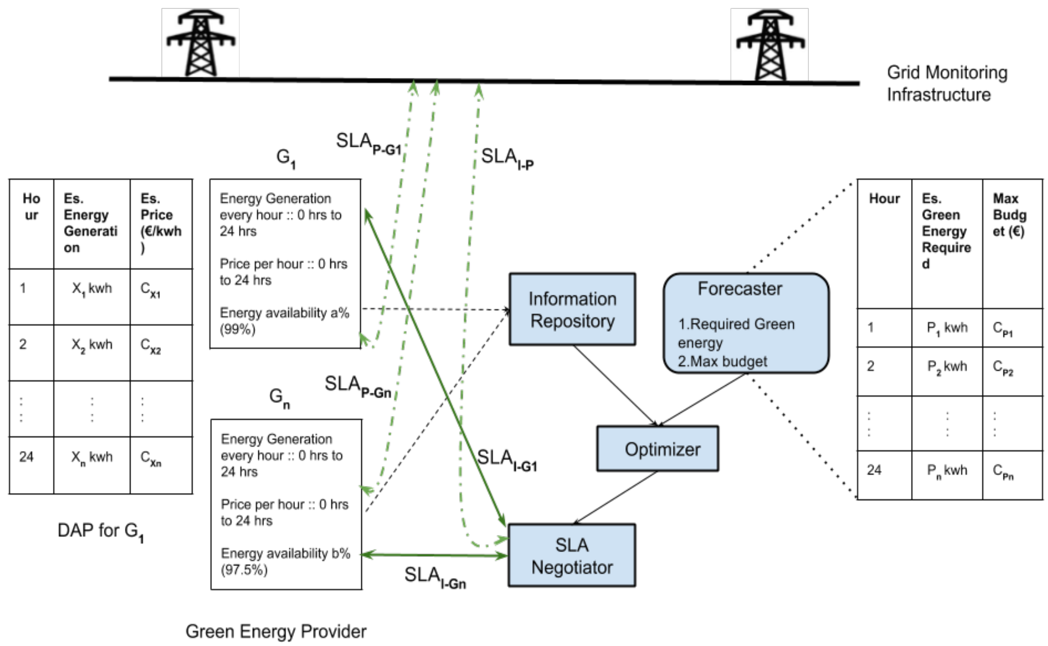


(b) Components

Figure 2. Top level view of the framework

Table I
POWER CONSUMPTION BY THE SELECTED SERVERS AT DIFFERENT LOAD LEVELS IN WATT

| Servers | 0% | 10% | 20% | 30% | 40% | 50% | 60% | 70% | 80% | 90% | 100% |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Dell Inc PowerEdge M620 | 688 | 1151 | 1322 | 1494 | 1671 | 1848 | 2061 | 2289 | 2499 | 2765 | 3239 |
| IBM NeXtScale nx360 M4 | 550 | 873 | 999 | 1123 | 1251 | 1380 | 1525 | 1673 | 1887 | 2116 | 2404 |

most of today's datacenter have higher PUE values than 1.7 [15]. Therefore,for transforming CPU utilization to power consumption, we traced CPU utilization for 7 days of 30 servers from PLANETLAB [11] where CPU Utilization has been traced for 500 different servers from across the world. As, building precise analytical models for modeling power consumption by modern multi-core CPUs makes a complex research problem, instead of using an analytical model of power consumption by a server, we utilize real data on power consumption provided by the results of the SPECpower benchmark [1]. Theoretically, researchers assume the increment of power consumption is linear to the increment of CPU load, whereas practically the power consumption in a server is not linear and increase significantly beyond 80% CPU utilization.

[1]https://www.spec.org/power_ssj2008/

### B. Forecaster Evaluation

We have selected two server configurations with multi-core CPUs published in November 2013: Dell Inc. PowerEdge M620 (Intel Xeon E5-2660,8coresX2.2 GHz,64 GB), and IBM NeXtScale nx360 M4, (2 Intel Xeon E5-2600,10 coresX2.2 GHz,256 GB). The configuration and power consumption characteristics of the selected servers are shown in Table I. In addition, we use OpenForecast [2] to forecast power demand for next 24 hours based on traced last 7 days power consumption. Single variable polynomial regression, Simple exponential smoothing and Double exponential smoothing method are used as forecasting method.

**Results**: Figure 3(a), 3(b), 3(c) shows 7 days traced CPU utilization, transformed power consumption from CPU utilization for 7 days and power consumption prediction for next 24 hours respectively. As our goal is to make the datacenter implicitly 30% green, we scale down the power requirement demand to 30% which is shown in Figure 3(d).

### C. Optimizer Evaluation

We consider 4 GEaaS providers exist in REC market for the purpose of our evaluation but it can be extended to more than 4 providers. As demonstrated in Figure 4(a), every GEaas provider has different level of availability of energy in kwh over time which is published at DAP information. The level of availability differs for various reason including different wind speed over time, unavailability of cut-in wind speed, different solar irradiation over time and the capacity factor of the plants. Furthermore, some providers might use more than one or different sources to produce green energy, which also results different level of energy generation. Using *Riemann sum*, we calculate the energy consumption demand from Figure 3(d), as the billing or cost for consumption is always calculated over energy consumption in kwh rather than power consumption in kw.

Finding market prices of each kwh produced by green sources are extremely difficult as most of the today's wind or solar power infrastructure or plants receive enormous incentives either from government or different policy making organizations. Hence, to model a realistic price for energy of different GEaaS providers and energy purchasing budget for IaaS providers, we investigate information of CAPEX-OPEX, levelized cost, fixed O&M cost, variable O&M cost of different sources of energy (e.g.; Nuclear, Wind, Solar, Hydro etc) [3] and find that the ratio of energy consumption cost between nuclear and green energy is *1:1.68* approximately. Therefore, we consider, the price of green energy sold by GEaaS providers will be around .19 - .25 cents/kwh while the price of Nuclear or mixed energy provided by EDF [4] is .13 cents/kwh.

In our experiment, we compare our optimization framework with two greedy approaches based on availability and cost. The first one tries to find the GEaaS providers such a way that it can satisfy the near optimal green energy demand whereas second ones seeks to select certain GEaaS providers which offer lower cost for selling their energy.

**Results**: Figure 4(b) compares the green energy demand by forecaster, meeting the demand by our optimization framework and the cost aware greedy approach. While cost aware greedy approach fails to meet the energy demand by 14%, our optimization framework performs better by providing 98% of the total demanded green energy within the green energy budget of IaaS provider. Furthermore, availability aware greedy approach incurs 5% more cost than the green energy budget of IaaS provider, while our approach follows the budget strictly and fails to provide only 2% of demanded Green energy showed in Figure 4(c).

### D. Discussion

The accuracy of our forecaster depends on the input window size and the prediction interval. The possibility of occurring error in the run-time phase can not be ignored. To address this issue, we can use CSLA features (confidence, fuzziness) or the VM consolidation in the run-time phase.

## V. Related work

Using renewable energy and offsetting carbon footprint in datacenters have been widely studied in recent years. GreenSLA was first introduced by [6], where they propose hardware and software technique for reducing energy consumption and integration of renewable energy. They also have shown the difference between Green and normal SLA, but validation process and approaches to satisfy SLA was not discussed. In contrast, we define different SLO objectives in all the Cloud computing layers while using renewable energy to provide Green Service.

In [4], Haque et al. proposed explicit integration of Green energy at the data center having different rack of servers and infrastructures for accommodating brown and green energy to provide Green computing services to the client where Green SLA specifies the requirement for percentage of Green energy needed to run a job. As they have integrated on-site renewable energy, providing green computing services is not possible when the generation of green energy is scarce or not present. On the contrary, we exploit REC market to make our datacenter Green by 30% to provide green services to the customers(e.g. SaaS layer, end user) where realistic SLA negotiation and execution model exist in the framework. Apart from that many papers have focused on datacenters that exploit green energy [1], [10], [9], [14]. Of these, [12] studied carbon-aware energy capacity planning for datacenter where they

[2]http://www.stevengould.org/software/openforecast/index.shtml
[3]http://www.eia.gov/forecasts/aeo/pdf/electricity _generation.pdf
[4]http://entreprises.edf.com/entreprises-45638.html

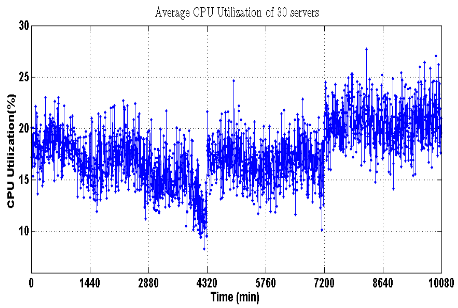


(a) CPU Utilization

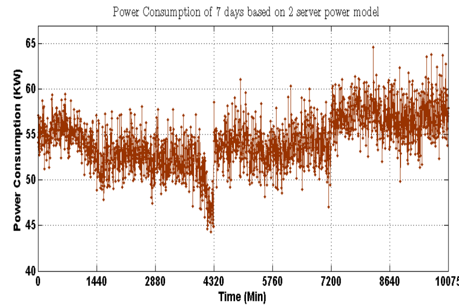


(b) Power Consumption

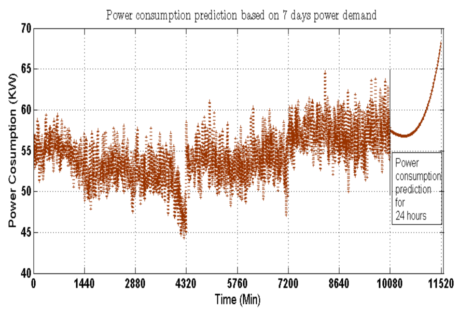


(c) Predicted Power

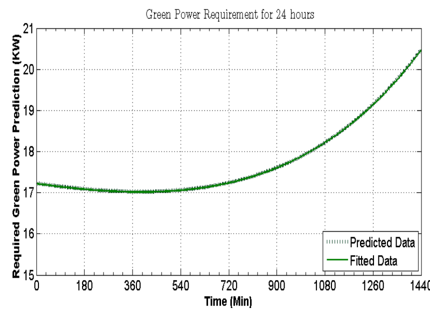


(d) Green Power Prediction

Figure 3. From CPU utilization to Green Power Prediction

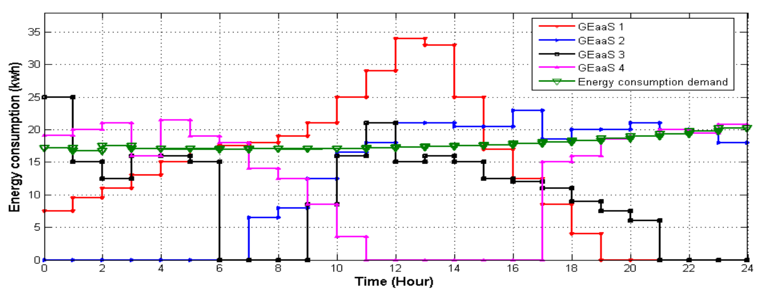


(a) Energy distribution

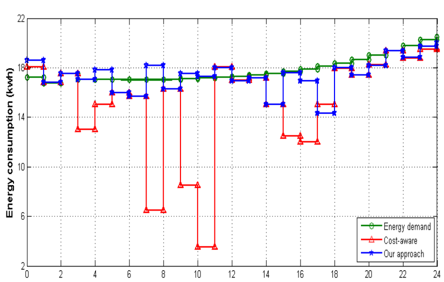


(b) Cost aware

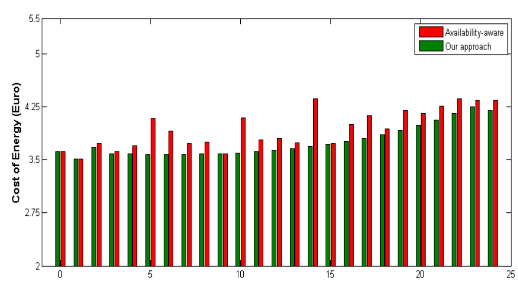


(c) Availability aware

Figure 4. Our approach vs Cost aware vs Availability aware

propose optimization framework to reduce datacenters cost in the presence of different carbon footprint reduction goal, renewable energy characteristics, policies and utility tariffs. While they investigate several cost-effective approach for using different renewable energy models, didn't consider any kind of SLA or Green SLA framework to justify the realistic implementation in Cloud computing environment.

A substantial number of prior works have [2], [5] also addressed Green data center or Green Cloud Computing by reducing total power consumption in server level either by VM consolidation, efficient VM packing or server consolidation with bounded performance loss. But none of these work considered integrating green energy in the datacenter.

## VI. Conclusion and future work

In this paper, we proposed that, IaaS provider can offer Green computing services to SaaS providers or end clients by exploiting REC market. Also, we have described a broker framework by which cross-layer SLA could be easily validated. Functionality of broker components and planning phase were briefly explained throughout the paper. The energy broker was evaluated with a real workload traced by PlanetLab. We also proposed two greedy heuristics policies for achieving the same goal at planning phase. But our optimization framework out performed those greedy policies in terms of availability and cost. While, we only have discussed the planning phase in this paper, our future plan is to provide efficient solution for run-time phase.

## Acknowledgement

This work is supported by the EPOC project within the Labex CominLabs (http://www.epoc.cominlabs.ueb.eu/).

## References

[1] B. Aksanli, J. Venkatesh, L. Zhang, and T. Rosing. Utilizing green energy prediction to schedule mixed batch and service jobs in data centers. In *Proceedings of the 4th Workshop on Power-Aware Computing and Systems*, HotPower '11, pages 5:1–5:5, New York, NY, USA, 2011. ACM.

[2] N. Bobroff, A. Kochut, and K. Beaty. Dynamic placement of virtual machines for managing sla violations. In *Integrated Network Management*, pages 119–128, May 2007.

[3] M. Ghamkhari and H. Mohsenian-Rad. Energy and performance management of green data centers: A profit maximization approach. *Smart Grid, IEEE Transactions on*, 4(2):1017–1025, June 2013.

[4] M. Haque, K. Le, I. Goiri, R. Bianchini, and T. Nguyen. Providing green slas in high performance computing clouds. In *Green Computing Conference (IGCC), 2013 International*, pages 1–11, June 2013.

[5] M. S. Hasan and E.-N. Huh. Heuristic based energy-aware resource allocation by dynamic consolidation of virtual machines in cloud data center. *KSII Transactions on Internet and Information Systems*, 7(8):1825–1842, 2013.

[6] S. T. Klingert, S. and C. Bunse. Greenslas for the energy-efficient management of data centers. In *International Conference on Energy-Efficient Computing and Networking*, May 2011.

[7] Y. Kouki, F. Alvares De Oliveira Jr., S. Dupont, and T. Ledoux. A Language Support for Cloud Elasticity Management. In *IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, CCGrid 2014*, pages 1–8, Chicago, United States, May 2014.

[8] Y. Kouki and T. Ledoux. CSLA: a Language for Improving Cloud SLA Management. In *2nd Int. Conf. on Cloud Computing and Services Science (CLOSER)*, 2012.

[9] A. Krioukov, C. Goebel, S. Alspaugh, Y. Chen, D. Culler, and Y. Katz. Integrating renewable energy using data analytics systems: Challenges and opportunities. In *In Bulletin of the IEEE Computer Society Technical Committee*, 2011.

[10] Z. Liu, Y. Chen, C. Bash, A. Wierman, D. Gmach, Z. Wang, M. Marwah, and C. Hyser. Renewable and cooling aware workload management for sustainable data centers. In *Proceedings of the 12th ACM SIGMETRICS/PERFORMANCE Joint International Conference on Measurement and Modeling of Computer Systems*, SIGMETRICS '12, pages 175–186, New York, NY, USA, 2012. ACM.

[11] K. Park and V. S. Pai. Comon: a mostly-scalable monitoring system for planetlab. In *SIGOPS Operating Systems Review*, volume 40, pages 65–74. ACM, 2006.

[12] C. Ren, D. Wang, B. Urgaonkar, and A. Sivasubramaniam. Carbon-aware energy capacity planning for datacenters. In *Modeling, Analysis Simulation of Computer and Telecommunication Systems (MASCOTS), 2012 IEEE 20th International Symposium on*, pages 391–400, Aug 2012.

[13] F. Rossi, P. v. Beek, and T. Walsh. *Handbook of Constraint Programming (Foundations of Artificial Intelligence)*. Elsevier Science Inc., New York, NY, USA, 2006.

[14] N. Sharma, S. Barker, D. Irwin, and P. Shenoy. Blink: Managing server clusters on intermittent power. *SIGARCH Comput. Archit. News*, 39(1):185–198, Mar. 2011.

[15] J. Yuventi and R. Mehdizadeh. A critical analysis of power usage effectiveness and its use in communicating data center energy consumption. *Energy and Buildings*, 64(0):90 – 94, 2013.