Cloud Sustainability Dashboard, Dynamically Assessing Sustainability of Data Centers and Clouds

Cullen Bash, Tahir Cader, Yuan Chen, Daniel Gmach, Richard Kaufman, Dejan Milojicic, Amip Shah, Puneet Sharma

HP Laboratories
HPL-2011-148

Keyword(s):

Sustainability; Clouds; monitoring; and modeling.

Abstract:

Quantifying and understanding the sustainability impact of large scale systems is becoming more and more critical. We present a Sustainability Dashboard that models and assesses the overall sustainability of Data Centers and Clouds. The dashboard provides a comprehensive view of IT infrastructure and services with respect to economic, ecological, and social aspects based upon an evaluation of multiple resources: servers, storage, networking, power and cooling, IT support, water, carbon, etc. A prototype is deployed and currently running on the Open Cirrus Cloud computing testbed.

External Posting Date: September 6, 2011 [Fulltext]. Approved for External Publication

Internal Posting Date: September 6, 2011 [Fulltext]

Cloud Sustainability Dashboard

Dynamically Assessing Sustainability of Data Centers and Clouds

Cullen Bash, Tahir Cader, Yuan Chen, Daniel Gmach,
Richard Kaufman, Dejan Milojicic, Amip Shah, Puneet Sharma

Hewlett Packard, Palo Alto, CA, USA
[first.lastname@hp.com]

***Abstract*—Quantifying and understanding the sustainability impact of large scale systems is becoming more and more critical. We present a Sustainability Dashboard that models and assesses the overall sustainability of Data Centers and Clouds. The dashboard provides a comprehensive view of IT infrastructure and services with respect to economic, ecological, and social aspects based upon an evaluation of multiple resources: servers, storage, networking, power and cooling, IT support, water, carbon, etc. A prototype is deployed and currently running on the Open Cirrus Cloud computing testbed.**

Keywords: Sustainability, Clouds, monitoring, and modeling.

I. Introduction

As large, consolidated Data Centers and Cloud computing infrastructures [1] are proliferating worldwide as a result of the demand for IT services, the global community is becoming increasingly conscious about the sustainability impact of IT services [2]. In our work, we define sustainability as a weighted sum of economic, ecological, and social metrics. This formulation is consistent with qualitative descriptions of sustainable systems as being those that lie at the nexus of economic, ecological and social benefit [3]. It is difficult yet extremely important to understand Cloud and Data Center sustainability. Cloud computing users may not even know in which region their services run, nor the trends of economic and ecological cost. For example, electricity cost increasing in one region or decreasing in another may impact Cloud provider costs and thereby prices. Additionally, business continuity may be at stake without in-depth knowledge. For example, some users may not want to host Cloud services in areas affected by catastrophes, whereas others may want to host them there out of solidarity with the region. To collect and present this information to customers and providers in a user-friendly manner requires addressing the following hard problems: modeling resources in the Data Center/Cloud, collecting information dynamically, and integrating with tools such as dashboards, closed loops, and alert mechanisms.

Users of Cloud services have different constraints. Some are cost-driven, others are more ecologically conscious, the third care most about the social implications of where the Cloud services are executing. A flexible and configurable solution is required, where Cloud users and providers can easily express and view different aspects of sustainability to better understand implications they care most about. In this paper we showcase a Cloud Sustainability Dashboard that models, measures and reports on the sustainability of Clouds and Data Centers. Besides obtaining insight into sustainability of the IT infrastructure, Cloud Sustainability Dashboard can be integrated with other tools to improve workload management or business continuity. Then both Cloud users and providers can decide where to host services to maximize sustainability.

II. Our Solution

We built a Sustainability Dashboard to evaluate and understand the overall sustainability impact of Data Centers and Clouds. By developing a set of economic, ecological and social models and then factoring in IT equipment information, power source, and regional electricity prices, the dashboard displays real-time metrics and can also send the data to other tools to enable sustainability-aware management.

A. Conceptual Architecture

The conceptual architecture of the Sustainability Dashboard is presented in Figure 1. We use standardized probes from monitoring tools to collect resource usage data for different elements in Data Centers, including server, networking switch, storage, application, power, and cooling. This information is then processed (filtered, converted, aggregated, etc.) to generate sustainability metrics using a set of sustainability models. Some examples of sustainability metrics are power consumption, costs, energy efficiency, carbon emission, water usage and economic development index. The sustainability information can be aggregated at different levels from individual equipment, groups of equipment, application and service level to the entire Data Center. These sustainability metrics are then sent to a dashboard, which dynamically displays real-time sustainability information. The sustainability metrics can then be consumed by different management tools, such as Workload Manager (a tool that monitors the overall load and distributes it based on predefined policies, such as load balancing, best performance, or optimal power consumption) and Business Continuity Evaluator (a tool that assures the services are deployed for optimal business continuity, e.g., on different Data Centers or different racks).

B. Sustainability Models

The evaluation models are constructed as follows:

Sustainability~ f(Economic, Ecological, Social)
Economic ~ f(Servers, Storage, Networking, Facility, Support, Efficiency)
Ecological ~ f(Carbon Emission, Water Use, Resource Consumption)
Social ~ f(Economic Development, Sociopolitical Stability)

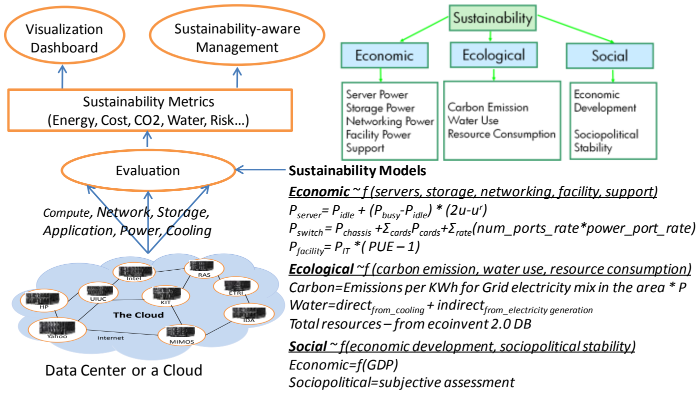

Figure 1. Conceptual Architecture

In the **economic model**, we estimate the costs from computing, storage, networking equipment, facility and IT support as well as energy efficiency. Server, networking, and storage equipment power usage is obtained from monitoring data or estimated using power models based on resource usage [4], [5]. Facility (power delivery + cooling) power is calculated using Power Usage Effectiveness (PUE) [6] or determined from cooling and power delivery models using COP_G (coefficient of performance of the ensemble) [7]. IT support costs consider the number of servers, homogeneity, redundancy, and the need for skilled labor in estimating the frequency and cost of incident service [8].

- *Server Power*. The server power consumption can be obtained from direct measurement (e.g., via HP iLO interfaces for HP iLO-based servers) or using power models. The server power usage is estimated based on the average CPU utilization, and the power consumed at idle and in fully utilized states as follows

  $$P_{server} = P_{idle} + (P_{busy} - P_{idle}) * (2u - u^r)$$

  where u is the average CPU utilization, P_{idle} and P_{busy} are the power consumed by the server at idle and their fully utilized state, respectively, and r is a calibration parameter [4]. This simple model has proved very useful and accurate in modeling power consumption since other components' activities are either static or correlate well with CPU activity. Given certain workload/utilization traces, we use the model to track the aggregate power usage for servers in a Cloud Data Center.

- *Networking Power*. The power consumed by network switches and routers is dependent on the chassis and linecard consumption along with the number of active ports and the link speed at which these ports are set [5]. We profiled the power consumption of a variety of networking switches and routers deployed in our Data Center. Specifically, we varied different switch and router configurations as well as the traffic flowing through these devices while a power meter attached to these devices recorded the power consumption at each switch and traffic setting. Based on these power measurements, we formulated a model to predict the power consumed by any switch or router. We find that a linear model is able to accurately capture the total power consumption of switches/routers currently in use; the switch power is given by

  $$P_{switch} = P_{chassis} + \Sigma_{cards} P_{card} + \Sigma_{rates}(n_{ports\_rate} * power_{port\_rate})$$

  where $P_{chassis}$ is the power consumed by the switch's chassis, P_{card} is the power consumed by a single linecard, the variable *cards* in the summation is the set of linecards in the switch; variable *rates* in the next summation is the set of rate settings for the ports (10 Mbps, 100Mbps, 1 Gbps, etc.), n_{ports_rate} is the number of ports set at that rate, and $power_{port\_rate}$ is the power of an individual port at that rate.

- *Storage Power*. The power consumed by storage systems is estimated based on the number of disk drives. This is a simplification, ignoring the additional power of equipment such as host bus adapters (HBAs), disk shelves, storage controllers, storage area networking (SAN) equipment, and power supply (in)efficiencies thereof. Nevertheless, it is a reasonable baseline estimate. Drives were assumed to be either large form factor "nearline" drives, large form factor enterprise drives, or small form factor enterprise drives.

- *Facility Power*. Facility (power delivery + cooling) power is calculated using Power Usage Effectiveness (PUE) [6], an industry standard metric that measures operational efficiency in the Data Center infrastructure. Specifically:

  $$PUE = (IT\ Power + Facility\ Power) / (IT\ Power)$$

  If the PUE can be measured empirically, then the facility power can be determined as the difference between the total and the IT power (top-down approach). Conversely, if the infrastructure description is known based on a model or experiment, then the facility power can be determined from cooling and power delivery models by simplifying the facility power into cooling and power delivery components. In doing so, we obtain:

  $$PUE \approx 1 + 1/COP_G + \%\ Power\ Delivery\ Losses$$

  where COP_G is the coefficient of performance of the ensemble [7], a measure of the cooling power required to remove heat from unit power input to the IT equipment.

- *IT Support Costs*. IT Support costs consider the number of servers, homogeneity, redundancy, and the need for skilled labor in estimating the frequency and cost of incident service [8]. We also take into account the cost of applications/services incidents and account system manager.

The **ecological model** includes *carbon emissions*, *water use*, and *resource consumption* (e.g., natural gas).

- *Carbon emissions* associated with a Data Center arise in several ways. We currently consider carbon emissions resulting from the amount of electricity consumed by the facility and the IT equipment, which are calculated based on the average grid electricity mix or the average carbon emissions per KWh for the grid electricity generated in the geography where the Data Center is located. Carbon emissions from transportation and heating (e.g., gas usage) in the Data Center are neglected for this first-order model.

- *Water use* is approximated as a function of direct use (e.g., water required for infrastructure cooling) and indirect use (e.g., water required in electricity generation), based on the approach outlined by Sharma et al. [9].

- *Total resource consumption* (e.g., natural gas) is measured in terms of MJ-equivalent (Megajoule) based on the direct resource use during Data Center operation, using the ecoinvent 2.0 database [10].
- *Energy Efficiency*. The Energy efficiency is defined as the amount of useful work done per Watt hour. The definition of useful work typically is application specific. However, for the general case useful work can be defined as number of CPU instructions. The energy efficiency is then defined as the number of billion instructions per Watt hour [BI/Wh]. We note that energy efficiency is highly correlated to the machine's utilization but also reflects power demand and performance of the machine.
- *CO_2 Efficiency*. The CO_2 efficiency is defined similar to the energy efficiency measured in billion instructions per kg of CO_2 emissions [BI/kg].

For all of the above, we only consider effects related to runtime (operation). Future work will include life-cycle (i.e., cradle-to-cradle, from component creation to its retirement) effects as well as additional environmental parameters such as ozone depletion, waste management, toxicity, etc.

In the **social model**, we consider two parameters qualitatively. First, we consider the potential for economic development and use the overall GDP as well as GDP per capita to represent this opportunity. Second, we consider sociopolitical stability. These are provided as examples only; additional metrics—such as the impact on human health, as well as legislative or legal requirements to comply with certain sustainability criteria—will be included in future work.

III. Implementation

A. Modeling Tools Implementation

We have implemented a number of individual tools that calculate sustainability metrics.

Server power monitoring. For HP iLO based servers, we obtain 5-minutes average power consumption for each server via the iLO interface. More generally, the server power model described in Section II.B is used to estimate the total power consumed by computing equipment. The model parameters P_{idle}, P_{busy} and r are identified from specification or through calibration experiments.

Networking power monitoring. Based on the linear switch and router power model in Section II.B, we have built a web-based tool that predicts the power of switches and routers in our network. Our tool polls the specified switches in the network and obtains relevant configuration information from them using standard entity MIBs over SNMP. Based on this information, our tool uses the appropriate switch power model to predict the power consumption of each switch along with the total power consumed across all switches. These predictions are within a 2% error margin of the actual power consumption.

A *cooling and facility power tool* is based on a model which utilizes the thermodynamic coefficient of performance of the ensemble (COP_G) [7], takes into account the cooling power required as heat is dissipated and transported from components to systems and racks, through the Data Center airspace and air-conditioners, and into the facility (chiller, cooling tower, etc.). By combining such a model with power delivery losses across the infrastructure, the total facility power can be estimated as discussed in Section II.B. Other monitoring tools have been discussed that allow for empirical estimation of PUE [11], which can then be scaled by the above IT power to obtain the facility power.

The *ecological modeling tool* utilizes the above total power input as a starting point. By linking the electricity use to a database containing relevant environmental impact factors—such as the carbon emissions per unit kilowatt-hour for a given site—the total CO_2 emissions can be obtained. Similarly, by linking the electricity use to water consumption as well as other resource consumption factors, we obtain a complete view of the ecological impact. Thus, given an input from the power assessment toolkit described above, the ecological tool would query the environmental database for the appropriate impact factors; link the two pieces of data together; process the linked data using the models discussed in Section II.B; and then return the output value.

Social modeling tools. GDP and other country specific development data are obtained from World Bank World Development Indicators (WDI) database [12]. We leverage our experience on the stability of different regions for an initial qualitative assessment of the political stability of a region.

B. Sustainability Dashboard Prototype

Based on the conceptual architecture and the sustainability models, we have designed and implemented a sustainability dashboard that integrates the model tools together. A prototype has been deployed and is running as a web service on one of Open Cirrus Cloud computing testbed sites—Open Cirrus HP site [1]. Open Cirrus Cloud computing testbed federates fourteen heterogeneous distributed Data Centers, each consisting of physical hardware, virtual machines, and global services, such as sign-on, monitoring, storage, and job submission. HP Site is located in HP Labs data center in Palo Alto, CA. The dashboard is currently monitoring and displaying real-time sustainability information for >200 servers and 18 networking switches. The dashboard prototype provides a web service interface for access and it can interoperate across sites.

Salient features of our implementation are high configurability, flexibility and extensibility. For example, additional metrics from other data sources can be added. New metrics can be defined using models and a model itself can be defined using available metrics and models. This enables us to plug in new models, e.g., a new and more accurate server power model or a new calculation of energy efficiency. Further, each metric and model has a scope associated with it

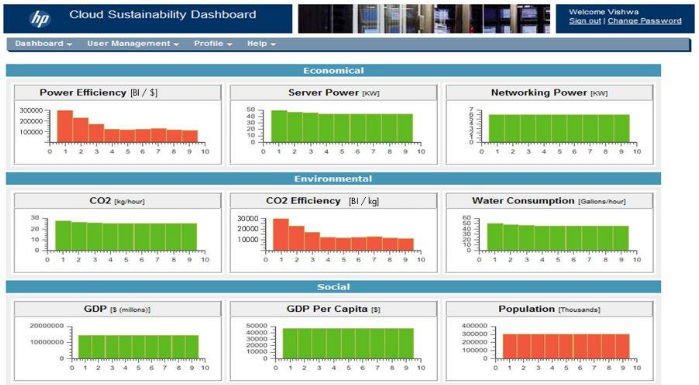

Figure 2. Data Center Level Sustainability Dashboard

such as site, server, or networking switch. Model expressions are evaluated dynamically and the system will automatically choose the correct model parameters for an element (e.g., server) based on its type and scope. We use a dynamic graphical user interface for displaying sustainability data. A user can specify what metrics to display and how to display them. For example, a metric can be displayed at different levels with different update intervals. A category can be assigned to a metric and threshold values can be specified. The dashboard will automatically organize the data according to its category and level, and use different colors to display based on the real value relative to the specified threshold. Thus, our approach allows qualitative estimation of the sustainability of highly dynamic environments, based on quantitative, real-time measurements and allows subjective specification and value judgments to be made by the end user.

Figures 2 and 3 show sample data collected by our prototype at Data Center level (Open Cirrus HP site) and user level (group of servers).

Figure 2 shows the Data Center level view presenting sustainability metrics over 10 hours, including power consumption of servers and networking switches, power cost efficiency for economic aspects, carbon emission and water usage for environmental aspects and population, GDP and GDP per capita and population for social aspects. A user can add/remove metrics from the dashboard and customize the way the metrics are displayed. In this example, the power cost efficiency is a custom metric, which is defined as billion instructions (BI) executed per $. So far, our prototype is deployed on an HP Data Center and shows the sustainability information for one Data Center only.

Figure 3 is another output of our prototype and shows the average power consumption, energy efficiency, and total carbon emission for three different user groups on Open Cirrus HP Site: BookPrep—a print-on-demand service, Sustainability Testbed—our Cloud Sustainability Dashboard service itself and other users. Though the average power consumption per server for all three user groups is similar, the third user (our testbed) has the highest power efficiency due to higher server utilization. The dashboard can also show individual server level metrics, providing a more complete and detailed view.

Figure 3. User Level Sustainability Dashboard

Such information can be used by workload management tools to optimize resource allocation and workload scheduling to reduce cost and carbon footprint, and satisfy sustainability requirements.

IV. Use Case Study

A. *Open Cirrus Cloud Sustainability Dashboard*

In this section, we describe an illustrative sustainability dashboard for Open Cirrus Cloud and how a user could use the dashboard for decision making.

We have applied our models to evaluate the sustainability for nine Open Cirrus sites, which are distributed across North America, Europe, and Asia with more than 2,000 computer servers (10,000 CPU cores) and approximately 4PB storage in total. Server power is estimated based on the assumption of 40% CPU utilization. Networking power consumption is based on an assumption that each site contains two rack switches, each having 48 1-Gbps ports and four 10-Gbps ports. At the next level, for every 10 racks there is a pair of tier-2 switches; and there is a pair of routers that connect all the tier-2 switches to the outside Internet. For each layer, actual equipment power measurements are used for representative switches at 75% utilization. For certain number of sites where the number of drives was not reported accurately, storage power estimates were made based on the total storage. SAN storage was assumed to be RAID 6 with 2 parity drives for every six drives; real overheads may vary considerably. The power dissipation of LFF “nearline” drives was taken to be 8W idle and 12W maximum, LFF enterprise drives to be 10W idle and 15W maximum, and SFF enterprise drives to be 6W idle and 8W maximum. Support costs are based on the assumption that a large number of servers (tens of thousands) with significant homogeneity and redundancy at the application level (e.g., triple redundancy in Hadoop) significantly reduce the cost of incident service. Manual labor can be performed by a less qualified workforce. The cost of repair is reduced due to virtualization and redundancies. Services will be more SLA-aware to clearly identify the boundaries of SLA violations and responsibility. For all of the above, we assume identical configurations at the different sites. In future work, this assumption of homogeneity can be replaced using site-specific data. An illustrative dashboard for Open Cirrus testbed is shown in Table 1.

The value of such a dashboard is best realized by considering several different users, each with their own sustainability preferences. For example, a user concerned primarily with economic costs may prefer to host their services at Site 5, which has the lowest IT, cooling, networking, and support costs. A second user who prefers a balance across all sustainability parameters might select Site 4 or 8, which has slightly higher economic cost, similar ecological footprint, and better social assessment than Site 5. A user most concerned with minimizing their environmental footprint would also prefer Site 5, which has the lowest carbon emissions, water use, and resource consumption. However, if Site 5 currently has insufficient capacity, the dashboard would inform the user that Site 9 has the next lowest ecological costs. The dashboard could also enable service delivery customized to user preferences. For example, a user could be advised when to start a large compute job that will satisfy the user's sustainability requirements. Such automation is not possible in a traditional enterprise computing environment, where service delivery is typically centralized and has fixed sustainability impact.

B. *Evaluating Sustainability of A Data Center with Different Workload Management Policies*

We consider a data center powered by Grid power and Solar power, e.g., photovoltaic (PV) power. Figure 4 shows the time varying renewable power supply from the solar panel and dynamic electricity price from Grid power over 24 hours for a small data center. Intuitively, different workload management will have impact on the data center sustainability (power, cost and carbon footprint). For example, given the lower electricity price at night, the jobs should be scheduled to run at night to minimize the cost. On the other hand, because more renewable power is available around noon, running jobs in the daytime can reduce the environmental impact. Due to complexity and frequent changes of the power supply and electricity price, most data center operators struggle with determining sustainability impact of different workload management policies. With the Sustainability Dashboard, data center operators can get a comprehensive view of their data center sustainability with different management policies and hence have an opportunity to optimize their workload management.

Figure 5 shows the sustainability of the data center in terms of power consumption, electricity cost and CO_2 reported by our sustainability dashboard for two different workload management policies. The first policy uses a simple First-Come-First-Serve (FCFS) scheduling while the second scheduler takes into account the electricity price and the availability of PV supply, and shifts the demand accordingly (e.g., running more jobs when the PV is available). The results show different schedules have a great impact on data center sustainability. Figure 6 summarizes the normalized results of power consumption, electricity cost and CO_2 footprint (normalized to the FCFS policy). As shown in the figure, given the same amount of demand, the second scheduler can significantly reduce both the electricity cost and CO_2 footprint via supply-aware scheduling and demand shaping.

Table 1. An Illustrative Dashboard for Open Cirrus (per quadrillion CPU cycles)

Open Cirrus Site	Economical ($)					Ecological				Social		
	IT	cooling	ntwk	suppt	econ ovrl	CO_2 tonnes-eq	water mill. Gal	resrc GJ-eq	ecol. ovrl	state of dev	risk of instabil	social ovrl
Site 1	$0.72	$0.35	$0.16	$0.43		6.0	2.6	83		high	low	
Site 2	$1.27	$0.59	$0.21	$1.11		6.8	3.3	96		high	v. low	
Site 3	$1.05	$0.47	$0.12	$1.07		5.9	2.3	81		high	low	
Site 4	$0.75	$0.35	$0.12	$0.61		6.1	2.7	85		high	v. low	
Site 5	$0.27	$0.13	$0.05	$0.09		4.3	2.4	59		low	high	
Site 6	$1.82	$0.77	$0.11	$1.17		10.2	4.3	142		high	low	
Site 7	$1.23	$0.54	$0.11	$0.98		15.0	4.4	192		high	low	
Site 8	$0.55	$0.26	$0.10	$0.16		6.9	2.6	95		med.	low	
Site 9	$1.01	$0.44	$0.10	$0.83		5.3	2.5	74		high	v. low	
Brick/ Mortar	$0.58	$0.70	$0.12	$0.83		9.0	2.1	127		high	v. low	

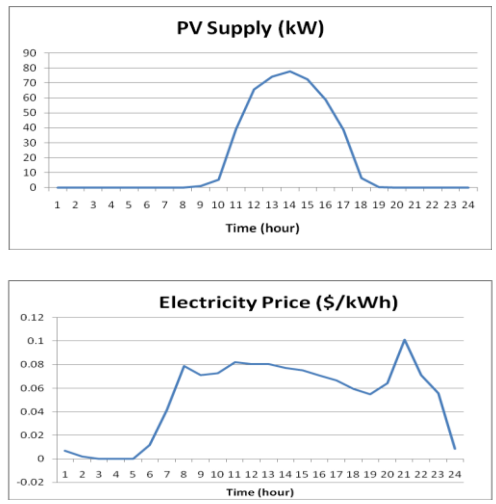

Figure 4. Power Supply Information

V. RELATED WORK

There are several recent efforts assessing Data Center sustainability. Microsoft has launched "Environmental Sustainability Dashboard" for SMBs [13]. IBM has launched a program "Ready for IBM Energy and Environment" [14]. Computer Associates is unveiling the "EcoSoftware for Green IT" [15]. However, these efforts only focus on one or two sustainability parameters and none are dynamically scalable for the use in a Cloud. Earlier solutions tend to rely on manual calculation, and thus have high labor costs and are error prone. Our dashboard uses automatically measurable quantities to provide a more accurate portrayal of Data Center sustainability. Other tools, such as those offered by Synapsense [16], focus on environmental metrics (temperature, pressure, humidity). Compared with these tools, sustainability dashboard provides a comprehensive sustainability view of infrastructure and services for traditional Data Center and Clouds.

VI. SUMMARY AND FUTURE WORK

We have developed a set of sustainability models and implemented a prototype on the Open Cirrus Cloud testbed. The prototype integrates the tools and models to dynamically

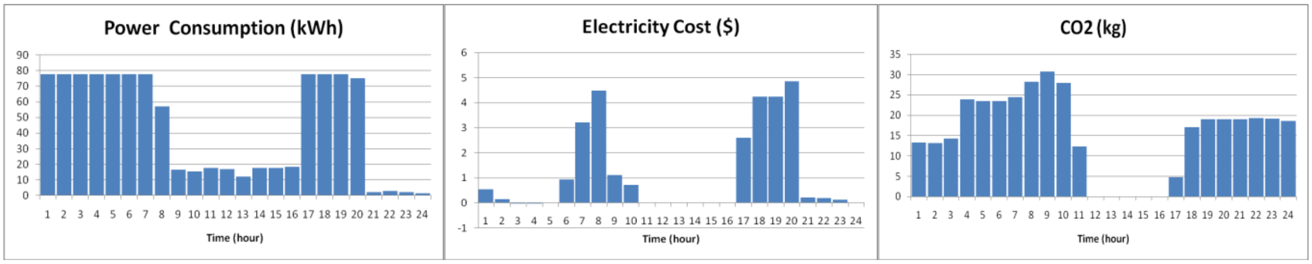

(a) FCFS Scheduler

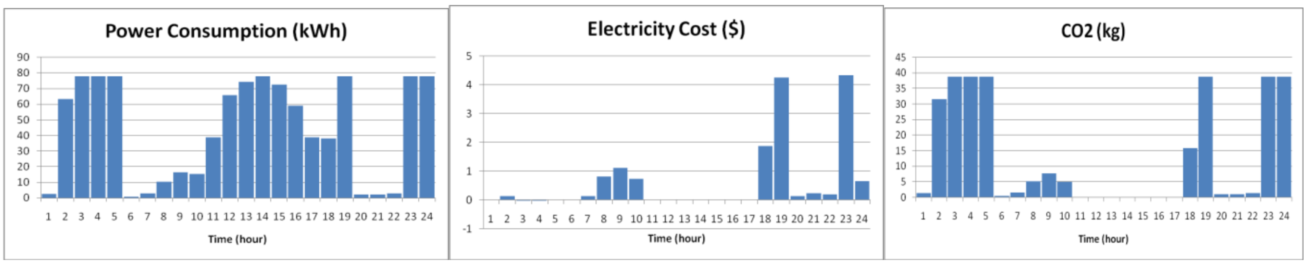

(b) Supply-aware Scheduler

Figure 5. Workload Management Impact on Sustainability

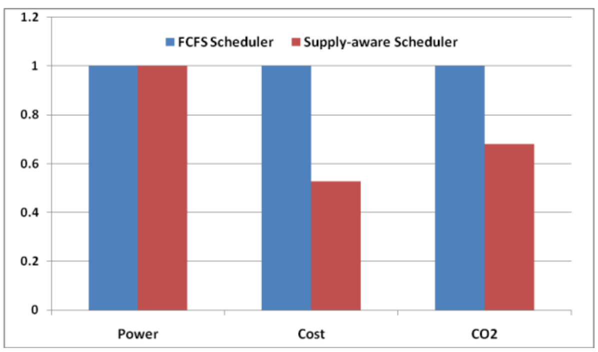

Figure 6. Comparison of Power, Cost and CO_2

generate and display sustainability metrics from low level data. We are verifying our models with real data. We plan to partner with other Open Cirrus sites to have a cross-site deployment and collect sustainability information from their Data Centers. We also plan to extend our model to integrate sustainability metrics into the Open Cirrus Node Reservation system. Once the Cloud Sustainability Dashboard becomes more widely deployed and used, we will be able to verify credibility of the metrics and tune it to accurately represent the sustainability of Cloud Data Centers and Cloud services.

ACKNOWLEDGEMENTS

Priya Mahadevan contributed significantly to the networking part of the paper. Arpita Kundu made significant contributions to Sustainability Dashboard implementation.

REFERENCES

[1] Avetisyan, A., et al., “Open Cirrus A Global Cloud Computing Testbed,” IEEE Computer, vol 43, no 4, pp 42-50, April 2010.

[2] Adams, W. “The Future of Sustainability: Re-thinking Environment and Development in the Twenty-first Century”, Report of the IUCN Renowned Thinkers Meeting, January 2006.

[3] Savitz, A.W., Weber, K., 2006, The Triple Bottom Line, John Wiley & Sons, Inc., New York, NY

[4] Xiaobo Fan, et al., “Power Provisioning for a Warehouse-sized Computer”, ACM, ISCA 2007.

[5] Mahadevan, P. et al, “A Power Benchmarking Framework for Networking devices", IFIP Networking Conf, May 2009.

[6] Belady, C., et al, “Green Grid Data Center Power Efficiency Metrics: PUE AND DCIE,” White Paper 6, The Green Grid.

[7] Patel, C.D., et al., “Energy Flow in the Information Technology Stack”, Proc. of IMECE, 2006.

[8] Silliman, R., “Vendor Survey Analysis: Benchmarking Hardware Support Operations, N. Am., 2009,” ID G00165983.

[9] Sharma, R., et al., “Water Efficiency Management in Data Centers: Metrics and Methodology,” ISSST, May 2009.

[10] The ecoinvent v2 database, PRé Consultants, available http://www.pre.nl/ecoinvent/default.htm

[11] Manos, M., Belady, C., “What’s your PUE Strategy?” July 2008. http://blogs.msdn.com/b/the_power_of_software/archive/2008/07/07/part-3-what-s-your-pue-strategy.aspx.

[12] World Bank World Development Indictors (WDI) database, http://data.worldbank.org/indicator.

[13] http://www.microsoft.com/en-us/dynamics/environment.aspx

[14] http://www-03.ibm.com/press/us/en/pressrelease/26604.wss

[15] http://community.ca.com/blogs/greenit/

[16] http://synapsense.com/go/index.cfm/data-center-solution/?gclid=CLrojOfR76gCFcnt7QodVCmsTQ