GreenIT Service Level Agreements

IN SERVICE LEVEL AGREEMENTS IN GRIDS WORKSHOP COLOCATED WITH IEEE/ACM GRID 2009 CONFERENCE, IN BANFF CANADA, OCTOBER 13, 2009

Gregor von Laszewski
Pervasive Technology Institute
Indiana University
2729 E 10th St.
Bloomington, IN 47408
U.S.A.
laszewski@gmail.com

Lizhe Wang
Service Oriented Cyberinfrastruture Lab
Rochester Institute of Technology
Bldg 74, Lomb Memorial Drive
Rochester, NY 14623-5608
U.S.A.
Lizhe.Wang@gmail.com

Abstract In this paper we are introducing a framework towards the inclusion of Green IT metrics as part of service level agreements for future Grids and Clouds. As part of this effort we need to revisit Green IT metrics and proxies that we consider optimizing against in order to develop GreenIT as a Services (GaaS) that can be reused as part of a Software as a Serice (SaaS) and Infrastructure Infrastructure as a service (IaaS) framework. We report on some of our ongoing efforts and demonstrate how we already achieve impact on the environment with our services.

Keywords: Green Service Level Agreements, Service Level Agreements, Green IT, Green Grids, Green Clouds.

1. Introduction

With the increased attention that green information technology (Green IT) is playing within our society it is timely to not only to conduct service level agreements for traditional computer performance metrics, but also to relate the effort of conducting agreements while incorporating their environmental impact. Much attention has recently been placed on reducing the environmental and operation cost impact of information technology [3],[4],[5]. This includes activities by the US government to especially target data centers, computers, and electronic equipment [1],[24]. As this world wide trend intensifies [6],[16],[8] Green IT will become even more important.

To provide a small motivational overview let us consider a typical desktop computer consumes 200-300W of power. One typical measurement for environmental impact is the resulting CO_2 emission. For the desktop computer this may result in about 220Kg of CO_2 per year. In contrast a typical data center produces 170 million metric tons of CO_2 worldwide currently per year. The expected emission data centers worldwide annually by 2020 will result in 670 million metric tons of CO_2. The average American car emits about seven tons of CO_2 per year. To put this in perspective, the average American family emits about 24 tons. Thus the data centers produce about the same as 28 million people or more than 95 million cars per year.

Another common measure is to simply use the power consumption and relate it to common values. To give an example, today's state-of-the-art supercomputer with 360-Tflops with conventional processors requires 20 MW to operate, which is approximately equal to the sum of 22,000 US households power consumption. In conjunction servers consume 0.5% of the world's total electricity usage and total energy usage is expected to quadruple by 2020. The total estimated energy bill for data centers in 2010 is $11.5 billion. However, 32% of all servers are running at or below 3 % peak and average utilizations, wasting energy spinning and cooling, and doing virtually no work. Thus, it is obvious that any efforts in reducing power and environmental impacts are significant.

1.1 Impact Factors

We distinguish for our purposes a number of impact factors that are essential for our Green IT efforts in regards to service level agreements. These impact factors are targeting hardware, software, environment, and behavior (see Figure 1).

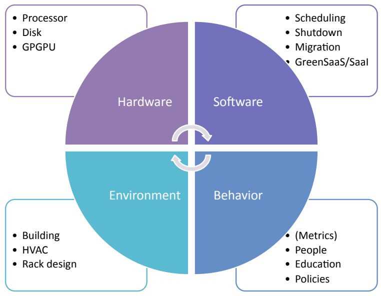

Figure 1. Green IT impact factors

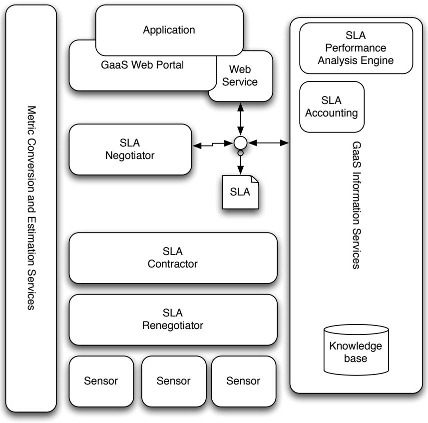

Figure 2. GaaS SLA Framework

Hardware. First, we have to recognize that any hardware used in a data center consumes energy, produces heat and thus has an immediate impact on the environment. While a single computer has only little impact, millions of machines as combined in many data centers have a significant impact.

Furthermore, we have to consider that the newest generation of hardware includes a variety of sensors and the capability to ingest energy efficient considerations in their uses either through automatic features provided by the hardware itself, or through enhanced software services, which we will discuss later. Presently, obtaining temperatures from processors, motherboards, disk drives as well as modern power supplies and fans and even motherboards have become the norm. Using these capabilities becomes thus most natural to be considered as part of the usage of the equipment. Examples of such integrated features include the shutdown of monitors when not in use (together with software controlled by the operating system), the dynamic voltage scheduling of processors, and the automatic adaptation of cooling equipment.

Furthermore, we also have to consider that recently as part of what has been termed disruptive technologies, general-purpose graphics processing units have been exploited to conduct numerical calculations. The availability of this specialized add on components can provide a quite significant performance boost. For example, we demonstrated in [12]a speedup of seventy times in a biostatistical application for flowcytometry in. This is in contrast to a processor that was available to us as part of our standard hardware infrastructure. Thus, it becomes apparent that is better to use a single system with a general-purpose graphics processing unit instead of building a cluster of seventy compute nodes. Not only do we save energy but also substantial capital costs in hardware. However, one has to develop and find applications that fit such special purpose hardware and it is at this time unclear how to factor the software engineering effort into account. Yet even on a cluster we would have to make sure we obtain or develop proper software that utilizes the hardware in optimal fashion.

Software. From the previous section, it is clear that the software plays an increasingly important role. We need to distinguish software that we term to be at the “close to metal layer,” software on the “operating system layer,” software on the “middleware layer”, and software on the “application layer.”

Each of these software layers can contribute towards the energy efficiency of the underlying hardware. While including sensors and metrics and providing proper optimization mechanisms in order to reduce the environmental impact of the hardware at runtime.

This includes for example the development of sophisticated energy efficient scheduling algorithms [23]. Another aspect of such scheduling frameworks may be the migration of applications or calculations onto more energy efficient hardware. A primitive but effective component can also be the automatic shutdown of hardware or software services if they are not in use. Helping to enhance this trend are Software as a Service (SaaS) and Infrastructure as a Service (IaaS) efforts.

While shifting the traditional programming efforts from developing stand alone programs towards the creation of services, it provides us with the unique opportunity to augment these services not only with capability descriptions in regards to computing such as its functionality or performance, but also integrate Green IT metrics that are relevant for establishing service level agreements between green as a services (GaaS).

We will describe our efforts to develop a GaaS in a later section of this paper.

Environment. Another important impact factor is the actual environment in which the compute and IT resources are located. This environment includes the building, machine room, offices, heating, ventilating, and air conditioning (HVAC) and reaches as far down as the computer rack and case design. Much effort has been recently spent in this realm including container based compute centers and redesign of backup power systems as for example introduced by Google.

Furthermore, on has to consider where the power for large-scale datacenters is obtained from. Some companies have increased their efforts to obtain power from alternative energy resources. Using the geographical location of the center can also be used to drive down the energy consumption as recently has been proposed by Google moving one datacenter to Brussels. With virtualization, the opportunity exists to engage in service level agreements that take into account such environmental impact factors and schedule Virtual Machines based on agreements which include not only the cost of the electricity, but also how the electricity is produced and which CO2 value it has.

Behavior. One of the most important factors that sometimes gets ignored the awareness of the environmental impact one has. Often programs are developed correctly for performance and accuracy reasons, but we have recently shown that we can reduce indirectly the amount of energy used by scheduling calculations on a supercomputer in such a way that lowers the overall temperature. Furthermore, it is not that essential on how fast the calculation is performed but that the overall set of scientific experiments is conducted in a high throughput fashion. Hence the overall round trip time of many experiments is more important than the actual performance of an individual experiment. This concept is well known as High Throughput Computing [14] has been employed by for example the Condor project.

Other behaviors such as using single precision or a mixed precision environment for obtaining results that are good enough could lead to a significant reduction in resources used. Most importantly we observed as part of our educational efforts that the availability of tools to visualize and give fast feedback about the environmental impact of a calculation is a significant feature to change behaviors of the scientists. Reporting on the overall environmental impact of a calculation in terms that can be related to daily activities (such as driving a car from New York to San Francisco) help motivating changes. Exposing this information as part of a computational portal and gateway is an essential component of GreenIT services. Together the activities need therefore included not only in educational processes at an institute level, but also in policy or operational considerations to reduce the carbon footprint. Efforts such as being charged for the environmental impact (for example a carbon emission tax) will accelerate such behavioral changes more quickly.

1.2 Service Level Agreements

As we are reaching more mature frameworks and middleware for Grid and Cloud computing it is becoming more important to incorporate service level agreements (SLA) into these frameworks that act automatically. Several specification, frameworks and software systems have recently been developed that target service level agreements [15],[17].

A good elementary introduction of the topic of SLA can be found in [15]. Important to note are the requirements and phases for an encompassing SLA framework which includes: (a) the SLA template specification, (b) publication & discovery in regards to the quality of service, (c) negotiation between the parties, (d) provider resource optimization to fulfill the contract, (e) monitoring of the agreement, (f) re-negotiation in case the original agreement is violated, (g) evaluation of the agreement to detect violations, and (h) accounting to offer appropriate payment of reimbursement of the services offered. Additionally, one has to also consider that in future one may want to introduce additional monitoring capabilities that allow (i) self regulation and policing as suggested in [21].

A good example of the need for this is a system which makes false advertisements (willingly or unwillingly) for a service that it can not deliver. Engaging in such a contract will be useless and before a contract is established such factors as reputation may have to be put into consideration.

Hence, a SLA framework will have to support these requirements and the associated phases with specialized services that are typically integrated in a service oriented architecture framework. This includes but is not limited to the negotiation of the contract and the monitoring of its fulfillment in run-time, as well as the policing of the contracts.

2. GreenIT Metrics

As part of our efforts we need to utilize Green IT metrics, such as Data Center Infrastructure Efficiency (DCiE) [18], [11], Power Usage Effectiveness (PUE) [11], Data Center energy Productivity (DCeP) [13], Space Watts and Performance (SWaP) [2], storage, network, and server utilization. Furthermore, we need to expose information about proxies that provide effective indicators and are often easy to implement. Such proxies include productivity sample workload, and bits per Kilowatt hour [13]. We will incorporate these and other metrics into the GaaS framework. This framework can than be used as an incubator to innovative algorithms and strategies while utilizing cyberinfrastructure in an attempt to minimizing the impact of science on the environment.

One of the most important aspects of a SLA for Green IT is to establish proper metrics that can be used for an agreement. There are many metrics available that would be appropriate for consideration. However, we have to recognize that a metric may not be exactly correlated to the final goal of reducing the overall environmental impact. This is especially true if we consider just runtime metrics and not lifetime metrics that include creation, recycling, and disposing of a resource or a system that generates energy to operate a resource. For the purpose of this discussion we do not consider at this time the later.

Furthermore, we notice that the issue of SLAs is actually a multi-scale problem that reaches from the smallest components over a server to an entire datacenter or even an agglomerate of data centers as part of a Grid or cloud.

The most common metrics that we can consider for SLAs are green house gas emission temperature [10] power consumption. Often humidity [7] is also considered as it has a significant impact on HVACs.

Whatever metric we chose, we must be careful if it indeed offers the required environmental benefit. For example let us consider metrics that measure the efficiency of a datacenter, such as the Data Center Infrastructure Efficiency, DCiE or DCE [18], [11], and the Power Usage Effectiveness (PUE): where *PUE = 1/ DCiE = Total Facility Power /IT Equipment Power*.

As seen above, PUE shows the relation between the energy used by IT equipment and energy used by other facilities such as cooling needed for operating the IT equipment. However it does not directly compare the environmental impact between data centers that must be measured in a different way. We find the PUE inadequate for our purposes.

We prefer however to apply the concept of metrics that DCxP introduces [13]: the amount of a unit x consumed to produce a particular work item. Hence, DCeP refers to as *Useful work/Total energy consumed to produce that work*.

To correlate such activities to Grids and Clouds has just been initiated and for example as part of efforts at the Open Grid Forum (OGF) [9].

In general we believe that multiple metrics and proxies that are easier to measure than the actual environmental impact will dominate for the next years the SLAs in regards to Green IT. However, if improved and integrated between each other, we may be able to extrapolate a holistic approach of choreographed services on multiple scales to assess the entire environmental impact of a given calculation, task, or scientific experiment.

3. GreenIT SLA Specifications

The specifications that we need to formulate include typical unit comparison and regular expressions to allow for maximum flexibility. Thus we will allow specifications such as

> Establish an execution service for 3 hours if the total carbon emission of the service is below x number of tons.

However we discovered while talking to application users that such a metric, although correct and useful for machine oriented services, may not be enticing the application user to actually use such a service as the unit metric ton may be too unfamiliar or can not be related to a real scenario.

Thus our framework will also include inquiry services that allow to query or even measure a subset of the calculations to be performed the impact of the calculation on a machine. Over time we may obtain a comparative factor for a number of different calculations that can be used for picking a resource with the least amount of environmental impact. This is similar to our efforts that have been conducted more than a decade ago [20][19].

In addition we will have a metric to proxy service that can correlate different environmental measurements that are more easily understood. An example query to this service would be

> How many miles can I drive with my car in order to use the same CO_2 value as my supercomputing application that I ran on my super computer.

Clearly such a common formulation of the impact of users not necessarily being experts in environmental units is important to entice the community and to build a stronger relationship to environmental impact of scientific calculations.

4. GreenIT Services

Next, we describe our services for a Green IT enabled service level agreement framework.

Besides providing services that give immediate feedback about environmental conditions as part of a sophisticated sensor network, we also provide services that provide feedback in regards to better utilization and cost reduction of existing infrastructure. These services can be integrated into a framework that over time provides valuable input in how we design the next generation datacenter and encourage behavioral changes how we can balance performance requirements with environmental concerns.

The GreenIT services for SLA are based on a number of integrative activities. This includes (a) analyze and leverage from current efforts in Green IT, (b) project and improve best practices including metrics, (d) provide green monitoring and auditing services, (e) provide services supporting Green IT service level agreements, (f) provide a green portal component, and (g) develop a coordinated infrastructure for Green IT.

Together these coordinated efforts impact how we conduct science with the help of modern cyberinfrastructure such as the TeraGrid and be able to project clues about its environmental impact. As a result we have a positive impact on how modern cyberinfrastructure and datacenters are utilized when clearly confronted with metrics not related to performance or typical QoS methods, but with metrics related to the environment.

Concretely, we design services for each of the metrics we are concerned with. These SLA metric services and than be integrated as part of carefully choreographed web services to utilize the information as part of specialized environmental services to reduce the environmental impact of the infrastructure whale at the same time achieving the tasks to be performed on user requests. Two examples of such services have already been developed. A thermal aware task scheduling service [25]and a dynamic voltage scheduling service [23].

4.1 Thermal aware task scheduling service

We developed thermal aware task scheduling algorithms, by predicating resource temperatures based on online task-temperature profiles [25]. The algorithms can be incorporated into a service providing a better carbon utilization for a set of tasks to be scheduled. Important for this service is that an accurate record of the improvements are kept in order to expose the environmental impact to the requesting consumer service. This should be done not only on a single system basis, but on a Grid wide basis so we can pick an environmentally friendly high throughput scheme to minimize the environmental impact while completing all tasks. We are currently developing some artificial intelligence techniques based temperature prediction methods and data center thermal operation patterns.

4.2 Dynamic voltage frequency scheduling service

Similar to our thermal aware task scheduling service we can leverage an algorithm that we describe in [23]as part of a sophisticated service to to reduce power consumption via the technique of Dynamic Voltage Frequency Scaling (DVFS) of scheduling virtual machines by dynamically scaling the supplied voltages. Some operating system level support for scientific applications are being studied and developed, for example, DVFS enabled MPI applications and multiple module compute intensive applications.

4.3 Integration services

To be successful, we need close integration with other service level agreement component services to provide contract establishment, contract fulfillment, and contract evaluation. The creation of a long-term knowledge base that we can mine through service invocations is an essential part of this integrative activity. Through this service we are able to assess potential service level agreement candidates while at the same time minimizing failures during service execution and service fulfillment. The integrative service will also have the ability to establish a reputation of various services that are offered through it as part of a brokering strategy. Middleware developers are able to register their own services as part of the integration service and their effectiveness will be measured automatically for a multitude of scenarios.

In general our services using this data are under development to serve as a basis for a new generation of more efficient and environmental friendly data centers and supercomputers. The access of the information is managed through control list and group memberships within a federated security framework. Data is exposed through enhanced Web Services following WS-* or the restful service paradigm.

Our architecture to representing our GaaS services and framework is depicted in Figure 2 and demonstrates the integration of the services that are needed to choreograph a successful SLA.

4.4 Portal

A user portal is essential to display the information conveniently and to interactively or programmatically mine it. Users of the GaaS portal will be able to save a *customized and state full view.* We are using state-of-the-art content management portals that are used by companies and millions of users. The Portal components hosted as part of this effort will allow a much more enhanced and dynamic experience while dealing with GreenIT data, in contrast to just a simple portal framework that has become common today. One of the important features will be that the community can contribute custom designed *green* components that can be integrated, even while using frameworks such as Google gadgets. However, we have demonstrated over the last year advanced features that allow us to expose an entire desktop framework through our portal. Hence, we will develop a GeenIT GaaS desktop into a browser rather than just focusing on developing portal tabs and components as shown in [22].

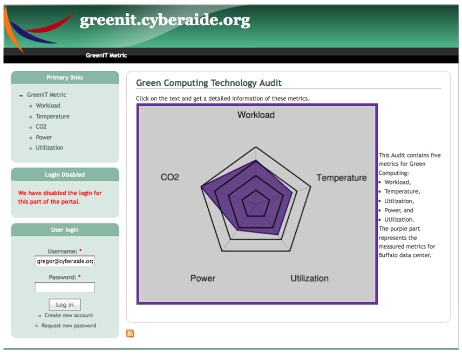

Figure 3. Cyberaide Green Portal

A screenshot of the integrative monitoring abilities of the portal is depicted in in Figure 3 that shows a snapshot of the environmental control data from the University of Buffalo as well as the calculations associated with it within a simulation that predicts the environmental impact for a calculation to be performed.

5. Conlusion

In this paper we documented our efforts on integrating GreenIT as part of a SLA framework that considers environmental impact as part of the agreements. We have conducted a number of significant efforts in developing algorithms that can be integrated as part of services decreasing the environmental impact without the knowledge of the users. However, it is important to also develop contracts and services in which explicitly specify metrics unique to environmental impact factors. This allows us to address SLAs not only from an individual server but also from a data center while solving not only a single calculation but a suite of experiments. A history service is integrated into a knowledge base that allows usto learn from past similar scheduled agreements and provide better services in the future.

Acknowledgments

Work conducted by Gregor von Laszewski is supported (in part) by NSF CMMI 0540076 and NSF SDCI NMI 0721656. We like to thank Tom Furlani and the members of the CCI at University of Buffalo to provide us with the temperature data. We like to thank Andrew Younge and Casey Rathbone for their help in finalizing this paper.

References

[1] Report to Congress on Server and Data Center Energy Efficiency. Available from World Wide Web: http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf.

[2] SWaP (Space, Watts and Performance) Metric. Web Page. Available from World Wide Web: http://www.sun.com/servers/coolthreads/swap/.

[3] Green IT definition. Webpage, May 2008. Available from World Wide Web: http://en.wikipedia.org/wiki/Green_computing.

[4] Green IT definition. Webpage, May 2008. Available from World Wide Web: http://thefutureofthings.com/articles/1003/green-computing.html.

[5] Green IT definition. Webpage, May 2008. Available from World Wide Web: http://www.webopedia.com/TERM/G/Green_IT.html.

[6] California government goes green. Website, 2009. Available from World Wide Web: http://www.arb.ca.gov/oss/green/greencites.htm.

[7] Data center humidity levels source of debate. Website, June 2009. Available from World Wide Web: http://searchdatacenter.techtarget.com/news/article/0,289142,sid80_gci1261265,00.html.

[8] EU energy star. Website, 2009. Available from World Wide Web: http://www.eu-energystar.org/en/en_database.htm.

[9] Metrics and an API for Green Computing on the Cloud. Website, May 2009. Available from World Wide Web: http://www.ogf.org/gf/event_schedule/index.php?id=1704.

[10] Recommended Data Center Temperature & Humidity. Website, June 2009. Available from World Wide Web: http://www.avtech.com/About/Articles/AVT/NA/All/-/DD-NN-AN-TN/Recommended_Computer_Room_Temperature_Humidity.htm.

[11] Christian Belady. The Green Grid Data center Efficiency Metrics: PUE and DCIE. Technical report, The Green Grid, Feb. 2007.

[12] Jeremy Espenshade, Andrew Pangborn, James Cavenaugh, Gregor von Laszewski, and Doug Roberts. Accelerating Partitional Algorithms for Flow Cytometry on GPUs. In *The 7th IEEE International Symposium on Parallel*

and Distributed Processing with Applications (ISPA-09), Chengdu and Jiuzhai Valley, China, August 2009. IEEE, IEEE.

[13] Jon Haas et. al. A Framework for Data Center Energy Productivity. Technical report, The Green Grid, Feb. 2008.

[14] Scott Fields. Hunting for Wasted Computing Power: New Software for Computing Networks Puts Idle PC's to Work. University of Wisconsin Research Sampler, 1993. Available from World Wide Web: `http://www.cs.wisc.edu/condor/doc/WiscIdea.html`.

[15] Gridpedia. Service Level Agreements. Available from World Wide Web: `http://www.gridipedia.eu/sla-article.html`.

[16] J. Kirk. EU to study energy use by data centers. Website, Feb. 2007. Available from World Wide Web: `http://www.networkworld.com/news/2007/022307-eu-to-study-energy-use.html`.

[17] SOI. Service Level Agreements. Available from World Wide Web: `http://sla-at-soi.eu/`.

[18] Gary Verdun. The Green Grid metrics: Data center infrastructure efficiency (DCiE) detailed anaysis. Technical report, The Green Grid, Feb. 2007.

[19] Gregor von Laszewski. An Interactive Parallel Programming Environment Applied in Atmospheric Science. In G.-R. Hoffman and N. Kreitz, editors, *Making Its Mark, Proceedings of the 6th Workshop on the Use of Parallel Processors in Meteorology*, pages 311–325, Reading, UK, 2-6 December 1996. European Centre for Medium Weather Forecast, World Scientific.

[20] Gregor von Laszewski. A Loosely Coupled Metacomputer: Cooperating Job Submissions Across Multiple Supercomputing Sites. *Concurrency, Experience, and Practice*, 11(5):933–948, December 1999. The initial version of this paper was available in 1996.

[21] Gregor von Laszewski, Beulah Alunkal, and Ivana Veljkovic. Toward Reputable Grids. *Scalable Computing: Practice and Experience*, 6(3):95–106, September 2005.

[22] Gregor von Laszewski, Fugang Wang, Andrew Younge, Xi He, Zhenhua Guo, and Marlon Pierce. Cyberaide JavaScript: A JavaScript Commodity Grid Kit. In *GCE08 at SC'08*, Austin, TX, Nov. 16 2008. IEEE. Available from World Wide Web: `http://cyberaide.googlecode.com/svn/trunk/papers/08-javascript/vonLaszewski-08-javascript.pdf`.

[23] Gregor von Laszewski, Lizhe Wang, Andrew J. Younge, and Xi He. Power-Aware Scheduling of Virtual Machines in DVFS-enabled Clusters. In *IEEE Cluster 2009*, New Orleans, 31 Aug. - Sep. 4 2009. IEEE. Available from World Wide Web: `http://code.google.com/p/cyberaide/source/browse/trunk/papers/09-greenit-cluster09/vonLaszewski-cluster09.pdf`.

[24] David Wang. Meeting green computing challenges. In *International Symposium on High Density packaging and Microsystem Integration*, pages 1–4, June 2007.

[25] Lizhe Wang, Gregor von Laszewski, Jai Dayal, Xi He, Andrew J. Younge, and Thomas R. Furlani. Towards Thermal Aware Workload Scheduling in a Data Center. In *Proceedings of the 10th International Symposium on Pervasive Systems, Algorithms and Networks (ISPAN2009)*, Kao-Hsiung, Taiwan, 14-16 December 2009. Available from World Wide Web: `http://cyberaide.googlecode.com/svn/trunk/papers/09-greenit-ispan1/vonLaszewski-ispan1.pdf`.