

Usage Centric Green Performance Indicators

Doron Chen, Ealan Henis, Ronen
I. Kat and Dmitry Sotnikov

IBM Research – Haifa, Israel
{cdoron,ealan,ronenkat,dmitry}
@il.ibm.com

Cinzia Cappiello, Alexandre Mello
Ferreira, Barbara Pernici and Monica Vitali

Politecnico di Milano,
Department of Electronics and Information, Italy
{cappiell,ferreira,pernici,vitali}@elet.polimi.it

Tao Jiang, Jia Liu and
Alexander Kipp

High Performance Computing Center
Stuttgart (HLRS), Department Intelligent
Service Infrastructures, Germany
{jiang,liu,kipp}@hlrs.de

Abstract

Energy efficiency of data centers is gaining importance as energy consumption and carbon footprint awareness are rising. *Green Performance Indicators (GPIs)* provide measurable means to assess the energy efficiency of a resource or system. Most of the metrics commonly used today measure the energy efficiency *potential* of a resource, system or application usage, rather than the energy efficiency of *the actual usage*. In this paper, we argue that the way that the resources and systems are actually used in a given data center configuration is at least as important as the efficiency potential of the raw resources or systems. Hence, for data center energy efficiency, we suggest to both select energy efficient components (as done today), as well as optimize the actual usage of the components and systems in the data center. To achieve the latter, optimization of usage centric GPI metrics should be employed and targeted as a primary green goal. In this paper we identify and present usage centric metrics, which should be monitored and optimized for improving energy efficiency, and hence, reduce the data center carbon footprint.

Keywords *power consumption; energy efficiency; carbon footprint; energy benchmarks; KPIs, GPIs*

1. Introduction

The purpose of defining Green Performance Indicators (GPIs) is to provide measurable means to assess the energy efficiency of a resource or system. If appropriate metrics are defined, then making changes to the system which results in better metrics levels should improve energy efficiency. Hence, defining “good” metrics is a central step towards energy efficiency optimization. But which will serve as good GPI metrics?

Currently, a myriad of GPIs have been defined and are used by various organizations (see more details below). Many of the common metrics measure the energy efficiency of a single resource (e.g., Hard Disk Drive – HDD, or CPU), whereas others

measure resources and entire systems behavior. Some metrics assess the energy efficiency of a “raw” device, and others assess the energy efficiency of their usage. All metrics commonly used have merit to some extent, and are useful tools to compare components and systems, and make informed decisions and choices based on comparative analysis. However, in this paper we claim that usage centric metrics have an added value that justifies promoting them to be primary GPI metrics.

Consider a simple example of capacity energy efficiency of a 300GB HDD measured in GB/Watt. Clearly, a disk with a higher value of this metric provides more efficient raw storage than one with a lower value. Note that a large percentage (60-70%) of the energy consumed by this HDD is expended whether the entire or a small portion of the space is used, due to the energy cost of spinning the disk. In one system configuration the space provided by this HDD may be fully used by applications (say 90% of the available 300 GB raw storage), whereas in another configuration only 10% of the capacity is actually used, out of the same HDD. Are we satisfied to use the raw capacity metric (which has identical value) in both cases? Had we measured the energy cost (GB/Watt) of the used capacity of the HDD rather than measuring the raw capacity, we would have reported a much higher efficiency value for the 90% usage case, and correspondently a reduced carbon footprint.

Based on the above motivation, we argue that defining and using usage centric GPIs has merit as energy efficiency metrics, and should be defined and used at all system levels as a primary means of assessing a system’s energy efficiency. The goal of this paper is to generalize this idea and provide the rationale and initial draft for the metric formulation in a wide range of scenarios.

2. GPIs State of the Industry

One of the most efficient drivers for energy efficiency is the governmental regulations.

The U.S Environmental Protection Agency (EPA) submitted in 2007 a report to the U.S congress on server and data center energy efficiency [4]. In this report the EPA projected that from 2006 to 2011 the energy consumed by data centers will be doubled. The EPA suggested that there is a significant potential for energy savings in data centers. The EPA report identified servers as the primary source of Information Technology (IT) energy consumption

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GreenMetrics 2011 June 7, 2011, San Jose, California, USA.
Copyright 2010 ACM 1-58113-000-0/00/0010...\$10.00.

in data centers. Only the site infrastructure in data centers consumes more energy than the servers. Therefore, the EPA introduced its ENERGY STAR Data Center Energy Efficiency Initiatives: The ENERGY STAR Rating for Data Centers, which provides tools to evaluate the efficiency of the data center, and the ENERGY STAR Data Center Product Specifications, which focus on requirements for data center products. The ENERGY STAR product specifications includes at this time the ENERGY STAR Computer Server Specification program version 1.0, published May 2009 [3]; the ENERGY STAR Data Center Storage Specification, which is under development and is expected to be published in 2011; and the ENERGY STAR Uninterruptible Power Supplies Specification, which is currently under development.

The Green Grid is a global consortium dedicated to advancing energy efficiency in data centers. The Green Grid publishes energy efficiency metrics that are widely used in data centers. The most known is Power Usage Effectiveness (PUE) [17], which measures the data center infrastructure energy overheads for running the IT equipment. The Green Grid set out to chart data center productivity metrics, which define the energy cost of useful work. Unfortunately, defining what is useful work is not simple. Therefore, the Green Grid defines proxies [10] for useful work productivity metric.

The Storage Networking Industry Association (SNIA) is an industry association dedicated to developing standards for storage. SNIA is actively pursuing energy efficient storage through its Green Storage Initiative (GSI) [11]. SNIA is developing the SNIA Emerald™ [14] program which is expected to be available to the public in 2011. SNIA is developing specific benchmarks for evaluating the energy efficiency of the storage system. SNIA Emerald will be based on the Storage Power Efficiency Measurement Specification [16] currently open for public review. The specification has 3 components: (1) an Idle measurement that focuses on what SNIA defines idle-ready state of a storage system. The idle metric is a capacity metric measured in GB/Watt, (2) a workload performance metric based on I/Os per second. This metric is defined as IOPS/Watt, and (3) a bandwidth performance metric based on the throughput of the storage systems, and measured in Megabytes per second per Watt (MBPS/Watt).

The Storage Performance Council (SPC) [15] is a non-profit corporation with the goal to define storage benchmarks and to provide an authority for verifiable performance data. SPC published two benchmark specifications: SPC-1, which focuses on online transaction processing and SPC-2, which focuses on sequential data processing. In 2009 SPC released an energy extension for SPC-1 that provides a performance/energy metric. SPC also provides an insight into the expected energy consumption of the storage by defining heavy, moderate, light and idle workloads. Using the data from the benchmark, the expected energy cost and the amount of hours per day spent in each workload type, it provides an estimate of the annual storage energy cost.

The Standard Performance Evaluation Corporation (SPEC) [12] is a non-profit corporation that defines and creates a standardized set of relevant benchmarks that can be applied to the newest generation of high-performance computers. SPEC provides a set of benchmarks that address topics such as CPU, graphics, Java, Power, Virtualization. SPEC started to work on a benchmark that evaluates both performance and power, and resulted in SPECpower_ssj2008. In the process of creating this first benchmark, SPEC, has created a power and performance bench-

mark methodology [13] that is now used to add power metrics to existing benchmarks, instead of creating energy specific benchmarks. As of writing this paper, SPEC has adapted SPECweb2009 (web server performance) and SPECvirt_sc2010 (virtualization in the data center) to include power metrics. SPEC is currently developing a Server Efficiency Rating Tool (SERT). This tool is developed to address the requirements of the EPA version 2.0 of the computer server specification. EPA specification version 2.0 will include means to evaluate the efficiency of the server while performing actual work (EPA version 1.0 addressed only idle power consumption of servers).

3. Metrics Survey and Classification

The government and industrial consortia described in the previous section base their measurements on a set of Green Performance and Quality of Service indicators. In the literature many indicators have been suggested to measure the performance and greenness of a system. Note, however, that some of these metrics cannot be measured efficiently. The subset of the most relevant indicators has to be identified on the basis of the business context and the characteristics of the system itself.

Indicators can be classified on the basis of what they monitor, estimating the energy consumption of the application itself and the environment in which the application is run. We identified four clusters of metrics for monitoring: (i) facility and data center, (ii) energy impact, (iii) application, and (iv) resources. The first cluster is focused on environment measures, such as air conditioning, temperature and humidity. In this paper we focus on IT metrics therefore, we skip the facilities cluster.

3.1 Energy Impact Metrics

This cluster of metrics is strictly related to the energy efficiency of hardware and software components. Many energy related parameters can be considered, such as power supply, consumed materials and emissions. Hardware components from different vendors have different energy efficiency properties. In order to measure these properties it is possible to use performance benchmarking tools with incorporated energy efficiency measurement. Some representative metrics included in this cluster are:

Capacity metric: measures the energy consumed within the storage facility [5] and it is computed as:

$$\text{Capacity Metric} = \text{Capacity}_{\text{storage}} / \text{Watt}$$

I/O throughput metrics: this indicator expresses the energy efficiency of I/O operations (i.e., data read and write) [5] considering the number of operations performed, computed as:

$$\text{I/O throughput} = \text{Number of I/O operations per second} / \text{Watt}$$

Data Transfer throughput metrics: expresses the energy efficiency of I/O operations considering the (rate of transfer of the) amount of data involved in the I/O transactions [5], measured as:

$$\text{Data Transfer} = \text{MB moved per second} / \text{Watt}$$

Application Performance indicator: measures the energetic efficiency of an application, calculated on the basis of the number of transactions and the energy consumed in a time period [2]. It is measured as:

$$\text{Number of transactions} / \text{KWh or FLOPS} / \text{kWh}$$

3.2 Application Metrics

Application metrics do not directly reflect the greenness of an application or of a data center. These indicators measure the performance of an application, and express the quality of the process and the efforts needed to design and maintain it. These metrics include quality of service indicators that refer to the analysis of the non-functional properties of the application, including [8,1,6]: response time, throughput, availability, reliability and recoverability.

3.3 Resource Metrics

In the attempt to measure the energy consumed by a server or an application the resources they are using within the system must be considered. Intuitively, the efficiency of the system depends on the efficiency of resources allocation among the running applications. The cluster of resource metrics measures the efficiency of resource usage relative to available resources. Consequently, their values are expressed as a percentage. Examples of indicators in this category are [9]:

CPU Usage: this GPI relates to the CPU utilization, and can be measured as

$$\text{CPU Usage} = \text{Used}_{\text{CPU}} / \text{Allocated}_{\text{CPU}}$$

Memory Usage: this GPI refers to the usage of the main memory (RAM) by a server or an application, and it is computed as:

$$\text{Memory Usage} = \text{Used}_{\text{Memory}} / \text{Allocated}_{\text{Memory}}$$

Storage Usage: this GPI refers to the entire storage utilization percentage for data read and write operations on the corresponding storage device, and can be computed as:

$$\text{Storage Usage} = \text{Used}_{\text{DiskSpace}} / \text{Allocated}_{\text{DiskSpace}}$$

Deployed Hardware Utilization Ratio (DH-UR): this indicator measures the amount of active deployed servers as [7]:

$$\text{DH-UR} = \text{Number of active servers} / \text{Number of Deployed Servers}$$

Deployed Hardware Utilization Efficiency (DH-UE): this indicator measures the efficiency of the deployed servers taking into account the amount of servers needed to face a peak load and it is measured as [7]:

$$\text{DH-UE} = \frac{\text{Minimum number of servers per peak load}}{\text{number of deployed servers}}$$

For this cluster of indicators, we observe a complicated situation. On one hand, it is obvious aim that the less a resource is used to achieve a result, the more energy efficient is the application, and hence, a lower value for these indicators is desirable. On the other hand, often it is not possible to reduce the amount of resources used by the system and their energy cost has already been paid (i.e., resources have already been configured and allocated). Hence, improving system energy efficiency means that a higher value for the indicator could be preferable, matching applications needs with allocated system resources.

3.4 GPIs Relations and Granularity

In general, indicators are not independent, and can be related to each other. Each Indicator I_k can be defined in terms of a metric $I_k = f(D_k)$ that defines the formula used to evaluate it, where $D_k \subseteq D$ is the set of raw data which are used in the assessment. For exam-

ple the GPI $I_1 = \text{CPU usage}$ can be associated with $D_1 = \{\text{CPU used, CPU allocated}\}$. In some cases an indicator can be also defined as a function of other indicators. The defined relationship is called Operational Dependency and can be expressed as $I_k = \text{Op}(\text{Ind}_k)$ where $\text{Ind}_k = \{I_1 \dots I_n\}$ is the set of indicators used to compute the indicator I_k . In this case there is a direct relation among indicators that is expressed in the evaluation formula. In summary, indicators can be defined as aggregations of different elementary variables or indicators, which are defined at the same or at a lower level. In other cases indicators can be indirectly related. This relation is called Qualitative Dependency, for which an indicator $I_k = \text{Q}(\text{Ind}_k)$ where $\text{Ind}_k = \{I_1 \dots I_n\}$ is the set of indicators that is correlated with the indicator I_k . Qualitative dependencies can be found using data mining techniques.

The indicators discussed above can also be computed at a different granularity level. Granularity depends on the level of detail used to measure the factors which compose the formula of an indicator. Typically, a system is designed to store the indicators values at the lowest level of granularity, but the values can be aggregated along different dimensions in order to enable users to achieve business goals. It is common that in order to understand the reasons why a data center system has performed in a certain way, it is often necessary to analyze data starting from a summary followed by detailed level analyses (i.e., drilling down operation), and vice versa. Possible dimensions that can be used to aggregate indicators values are time, infrastructure component, middleware component, application, geographical areas, etc. An intuitive example of analysis in which exploring indicators at different granularity levels could be significant is the resource metrics. In this case the indicator measures the amount of a resource used relative to the amount of resource available. Each of these indicators can be measured at the server level (including all the running applications) or with a finer granularity at the application level, measuring the amount of resources actually used by a single monitored application (relative to the amount of resources reserved for it). Aggregation can also be applied to clusters. In this case measurement can be made at each node and aggregated using a weighted sum in order to obtain a single value for the whole cluster.

Similar reasoning can be applied to the other clusters of metrics, and it paves the way to the distinction we make between two perspectives for measuring indicators: the manufacturer perspective, which is focused on the efficiency of the elements of the data center, and the usage perspective, which focuses on the efficient usage of these elements. We re-emphasize that our effort in this paper is to highlight the importance of usage perspective metrics. The *usage perspective* is further discussed in the next section, where some metrics and their implementation are discussed.

4. Usage and Metrics Definition

In the following we provide examples for usage metrics for servers and storage. This examples aim to strengthen our argument regarding the importance of usage centric energy efficiency metrics.

4.1 Server Level Usage Metrics Example

In this section we apply a subset of the metrics defined above to a concrete usage case, namely a cancellous bone simulation application scenario. Within this scenario, we show how some of our Green Metrics can be applied in order to determine the optimal system configuration, thus allowing for the best performance vs. energy consumption ratio. In particular, we assume that this High Performance Computing job is deployed on two different machine

cluster configurations, and we analyze the greenness of the system.

	System A	System B
Application Service Usage	(avg) 95% CPU usage	(avg) 83% CPU usage
I/O Usage Metric	(avg) 8GB/s	(avg) 6 GB/s
Memory Usage Metric	(avg) 93% Memory Usage	(avg) 60% Memory Usage
Application Performance Metric	20 GFLOPs/KWh	18 GLOPs/KWh

Table 1. Estimated Energy Footprint of a Cancellous Bone Simulation in two Configurations.

The simulation job is executed several times on the available resources of the HPC Center. The jobs are deployed on two different Systems (A and B in the following). According to the above mentioned metrics, the values reported in Table 1 have been monitored (where the Application Server Usage Metrics stands for CPU Usage). The corresponding energy consumption graph is depicted in Figure 1.

The *Application Server Usage* Metric allows getting an overview about the CPU utilization, and the *I/O Usage* Metric measures the amount of I/O operations. The former allows determining whether the deployment decision for the job, under the current setup, has been appropriate. For example, if the job just produces a medium CPU load, this might imply a possible reduction of the clock speed of the CPU without affecting the entire job execution time whilst reducing the entire energy consumption.

The latter *I/O Usage Metric* allows determining the average memory operational performance in terms of I/O operations, which is a relevant indicator for HPC jobs. In particular, this allows determining the impact of a compute job on the IO memory performance. For example, if a job with a supposed high CPU load just produces a medium CPU load during execution whilst outbidding the available memory I/O bandwidth, we consider that this implies that the memory I/O channel might be a performance bottleneck. This also implies that deploying this kind of job on a machine with improved memory I/O bandwidth might allow for a smoother execution of the job, and thus for better energy efficiency for the entire job execution.




Figure 1. Resource and Energy Consumption. A typical example of CPU, memory, temperature and power usage of a server node.

The *Memory Usage Metric* allows determining the average memory usage, considering also peak loads. In particular, this metric determines if potentially too many memory modules are active in the system. In that case, for future deployments, unnecessary memory bars can be turned off and switched on again on demand, thus allowing for additional energy saving. The Performance Metric provides a value relevant for comparison among different data centers, by relating performance values to the corresponding energy consumption. According to the measurements reported in Table 1, System A performs better for our test case job than System B, although System A has a higher CPU load during job execution. Moreover, System A provides better memory interaction, and consequently a better and more energy-efficient processing. Finally, System A allows for 20 GFLOPs/KWh whilst System B allows for only 18 GFLOPs/KWh; hence, for future deployments, System A should be selected, in order to allow for a higher energy efficiency.

Summarizing this example, the evaluation of these measured values shows a better energy efficiency of System A for this job type, due to a better Memory I/O ratio and a better usage of the available memory, whilst System B provides memory that is relatively less used during the job/application execution. A potential solution to improve the Performance Values for system B would be to switch off some of the unused memory modules.

Applying these metrics to HPC jobs is a first step towards the consideration of Green Metrics for the entire IT data center. These measured metrics allow deriving high level and more abstract views on the system involved, and in particular, on the energy consumption from a fine-grained viewpoint, while also considering the required QoS properties.

4.2 Storage Level Usage Metrics

While the storage metrics definition is very similar to metrics described by the organization referenced in Section 2, the main difference is how the metrics are computed. Therefore, the metrics provide the actual usage perspective and allow applications and administrators to evaluate and improve the energy efficiency.

The **Capacity Metric (GB/Watt)** represents the energy efficiency of storing the data of the user's applications. The metric is defined as the ratio between used storage space by applications (GB) and power (Watts). Used storage space is defined as the space used by files written and stored on the storage system. Power is measured as the average power of the storage system under typical usage, measured over a representative (long enough) time period.

The **I/O Throughput Metric (IOPS/Watt)** represents the energy cost of the storage system while running the user applications workload. The metric is defined as the ratio between applications I/O rate (IOPS) and power (Watts). Applications I/O rate is measured as the number of I/O operations per second that an application executes. Power represents the average power consumed by the storage system while running typical workload.

The **Data Transfer Throughput Metric (MBPS/Watt)** represents the same metric as the above I/O throughput, but for data transfer (MBPS). The metric is the ratio between applications data transfer rate and power. The two throughput metrics emphasize different aspects of the same efficiency. I/O throughput is commonly used for evaluating random workloads, while data transfer

throughput is commonly used for evaluating sequential workloads.

The above energy efficiency storage usage metrics describe how the user application utilizes the storage system. Hence, the above usage metrics are defined via terms that are directly linked to a user application, and represent how the application is taking advantage of the storage system. While the metrics defined above refer to applications files and I/O operations, these metrics can be easily adapted to accommodate object based storage system and databases. For example, transactions execution can be measured instead of I/O operations, and capacity can be measured by the amount of data in a database table.

5. Concluding Remarks

Green Performance Indicators provide measurable means to assess the energy efficiency of a resource or system. Many of the commonly used IT metrics today measure the potential energy efficiency of a resource, typically from the manufacturer's point of view. Using this approach allows comparing the energy efficiency potential of various alternative resources, and favors the more efficient one over others. Given the selected resources, however, in this paper, we argue that the way that the resources are actually used in a given data center configuration is at least as important as the efficiency potential of the raw resources. Optimization for data center energy efficiency must target both the selection of energy efficient components (as done today), as well as the optimization of the actual usage of these components. For the latter, optimization of usage centric GPI metrics should be employed and targeted as a primary green goal. The Green Grid has also identified the importance of IT system usage for improved energy efficiency in data centers, and outlined a framework for energy efficiency productivity metrics.

The usage metrics can be evaluated continuously, as opposed to potential metrics which can be computed once only and on a reference system. Data center and administrators can review the changes in the metric efficiency over time and also compare, to some extent, with other data center.

In practice, comparing and selecting resources based on their energy efficiency is relatively easy. The manufacturers test their components and advertise their energy efficiency results. Measuring the usage centric metrics for a given data center configuration, and optimizing the configuration to improve the metrics levels is much more complex. It requires expertise and configuration/workload planning that many data center operators are only now beginning to practice. In this paper we identified and presented the need to properly measure (via usage centric metrics) and optimize system usage for energy efficiency and carbon footprint.

Acknowledgments

This work has been partially supported by the GAMES project (<http://www.green-datacenters.eu/>) and has been partly funded by the European Commission's IST activity of the 7th Framework

Program under contract number ICT-248514. This work expresses the opinions of the authors and not necessarily those of the European Commission. The European Commission is not liable for any use that may be made of the information contained in this work.

References

- [1] Anbazhagan Mani and Arun Nagarajan. Understanding Quality of Service for Web Services. <http://www-128.ibm.com/developerworks/library/ws-quality.html>, 2002
- [2] Asanovic, K. et al., The Landscape of Parallel Computing Research: A View from Berkeley. Technical Report No. UCB/EECS-2006-183. <http://www.eecs.berkeley.edu/Pubs/TechRpts/2006/EECS-2006-183.pdf>. 2006
- [3] ENERGY STAR® Computer Server specification, May 2009. http://www.energystar.gov/ia/partners/product_specs/program_reqs/computer_server_prog_req.pdf
- [4] EPA Report to Congress on Server and Data Center Energy Efficiency. 2007
- [5] Greg Schulz. The Green and Virtual Data Centre (Chapter 5). <http://thegreenandvirtualdatacenter.com/greenmetrics.html>
- [6] Gregor von Bochmann, Brigitte Kerherve, Hanan Lutfiyya, MohamedVall Salem, and Haiwei Ye. Introducing QoS to Electronic Commerce Applications. In Proceedings of the Second International Symposium on Topics in Electronic Commerce, pages 138-147, 2001
- [7] John R. Stanley, Kenneth G. Brill, and Jonathan koomey. Four Metrics Define Data Center "Greenness". Uptime Institute. 2007.
- [8] Liangzhao Zeng, Boualem Benatallah, Marlon Dumas, Jayant Kalagnanam, and Quan Sheng. Quality Driven Web Services Composition. In Proceedings of the 12th World wide Web Conference (WWW), 2003
- [9] Mathias Weske. Business Process Management: Concepts, Languages, Architectures. Springer Berlin 2007
- [10] Proxy Proposals for Measuring Data Center Productivity. <http://www.thegreengrid.com/~media/WhitePapers/White Paper 17 - Proxies Proposals for Measuring Data Center Efficiencyv2.ashx?lang=en>.
- [11] SNIA Green Storage Initiative (GSI). <http://www.snia.org/forums/green/>
- [12] SPEC - Standard Performance Evaluation Corporation. <http://www.spec.org>.
- [13] SPEC Power and Performance Benchmark Methodology. http://www.spec.org/power/docs/SPECpower-Power_and_Performance_Methodology.pdf
- [14] Storage Networking Industry Association (SNIA) Emerald. <http://sniaemerald.com/>
- [15] Storage Performance Council (SPC). <http://www.storageperformance.org>
- [16] Storage Power Efficiency Measurement Specification. http://www.snia.org/tech_activities/publicreview/Storage_Power_Efficiency_Measurement_Spec_v0.2.10_DRAFT.pdf
- [17] The Green Grid Data Center Power Efficiency Metrics: PUE and DCiE. 2008.