



A FRAMEWORK FOR DATA CENTER ENERGY PRODUCTIVITY

CONTRIBUTORS:

DAVE ANDERSON, VERARI

TAHIR CADER, SPRAYCOOL

TOMMY DARBY, TI

NICK GRUENDLER, IBM

REMA HARIHARAN, AMD

ANNE HOLLER, VMWARE

CHRIS LINDBERG, INTEL

CHRISTINE LONG, APC

PHIL MORRIS, SUN

ANDY RAWSON, AMD

FREEMAN RAWSON, IBM

VIKRAM SALETORE, INTEL

JIM SIMONELLI, APC

HARKEERET SINGH, BT

ROGER TIPLEY, HP

GARY VERDUN, DELL

JOHN WALLERICH, INTEL

EXECUTIVE SUMMARY

Recent studies by industry and government [EPA - Report to Congress] have highlighted the issue of rapidly escalating data center energy consumption. Faced with the limited availability of additional electrical energy, many CIO's are forced to consider relocating their data centers or adding additional capacity situated in remote locations. For those who have an abundant supply of power, there is still the issue of the ever increasing power bill. A fundamental issue facing the data center manager is, "how do I control my energy costs without impacting the delivery of critical IT services my customers demand?" Required are new tools that allow the continuous monitoring of the work product of the data center as a function of energy consumed. Up to now, standardized tools to measure the productivity of a data center have not been available. This paper presents:

- a technical analysis of the problem of assessing the energy efficiency of a data center and the IT equipment that composes the data center,
- an examination of various power and energy efficiency metrics that have been proposed, and a discussion of their attributes and applicability, and
- an analysis of ways in which those attributes fall short of providing the comprehensive tools necessary to optimize data center energy utilization.

This paper introduces a new family of data center resource optimization metrics designated collectively as Data Center Productivity (DCP) metrics and presents the first derivative metric within this family called Data Center energy Productivity (DCeP). The DCeP metric provides a unique analytical tool that may be used to track the overall work product of a data center per unit of energy expended to produce this work. While DCeP in its current form is only applicable to improvements in a single data center, it is hoped that this work will provide a framework to develop similar metrics for comparing across different data centers.



TABLE OF CONTENTS



INTRODUCTION.....	4
CONTEXT FOR THIS PAPER.....	4
CONTEXT FOR THE DATA CENTER PRODUCTIVITY METRIC.....	5
WHAT IS A METRIC?.....	5
WHAT IS DATA CENTER PRODUCTIVITY?.....	5
ROLE OF POWER-PERFORMANCE BENCHMARKS.....	5
ROLE OF OTHER CURRENTLY DEFINED EFFICIENCY METRICS.....	6
POWER USAGE EFFECTIVENESS (PUE).....	6
COMPUTE POWER EFFICIENCY (CPE).....	7
DESIRED ATTRIBUTES OF A PRODUCTIVITY METRIC.....	7
INTRODUCING A NEW FAMILY OF METRICS.....	8
DATA CENTER PRODUCTIVITY.....	8
DATA CENTER ENERGY PRODUCTIVITY METRIC.....	8
MEASURING ENERGY CONSUMED.....	9
DEFINING USEFUL WORK.....	9
DEFINING A TASK.....	9
DEFINING AN ASSESSMENT WINDOW.....	10
ASSIGNING A VALUE TO TASKS.....	10
DEFINING A TIME-BASED UTILITY FUNCTION.....	10
TRANSACTIONAL OR THROUGHPUT-BASED WORKLOADS.....	11
EXPERIMENTAL RESULTS.....	12
SETUP.....	12
RESULTS.....	13
SUMMARY.....	14
FURTHER RESEARCH.....	14
REFERENCES.....	16

INTRODUCTION

Recent studies by industry and government [Kooimey, 2007; EPA - Report to Congress] have highlighted the issues of escalating energy consumption by data centers. Up to now, managers of data centers have not had a comprehensive set of tools to answer the question of how energy resources supplied to the data center are being used to efficiently generate the economically valuable work product of the data center. What has been lacking is a metric that captures and quantifies the useful work output that a data center provides in relationship to the amount of any specific resource expended to produce this useful work output. This paper begins to fill this gap by supplying a framework for quantifying useful work.



The paper begins by establishing the context of this discussion and describing the purpose and use of metrics. It goes on to explain the role of power-performance benchmarks and then examines a couple of extant power efficiency metrics. The desired attributes of a productivity metric are explored. The concept of the Data Center Productivity (DCP) metric is introduced. Essential to assessing DCP is the quantification of the useful work that a set of IT equipment produces. A framework for this process is proposed. It should be noted that this framework has been designed to allow the customization of the useful work measurement based on the specific characteristics of the workload of an individual data center and is not currently applicable to the formulation of a metric that would allow the comparison of the productivity of one data center to another. Using this framework an example assessment of Data Center energy Productivity (DCeP) is presented using laboratory conditions and a well-known synthetic workload. The paper demonstrates that this process of assessment is simple to understand and simple to implement given the appropriate instrumentation of the equipment within the data center. Experimental results are presented and further research is outlined.

CONTEXT FOR THIS PAPER

The equipment contained in a data center may be divided into two main categories: 1) infrastructure equipment and 2) Information Technology (IT) equipment. Included in the infrastructure category are equipment for conditioning and distributing electrical power (e.g., transformers, switch gear, uninterruptible power supplies, power distribution units, circuit breakers, and distribution wiring), as well as equipment used to remove waste heat from the data center (e.g., Computer Room Air Conditioners [CRACs], Computer Room Air Handlers, Direct Expansion (DX) coolers, chillers and pumps for circulating chilled water, and cooling towers). IT equipment includes computers called servers that run the operating system and application software that produce the primary work product of the data center along with support hardware such as storage devices and networking equipment.

Studies have shown that there is a significant opportunity to conserve energy in the data center by improving the energy efficiency of infrastructure equipment within the data center [Rasmussen, 2006]. This paper, however, will focus on IT Equipment, and more specifically, servers. It will leave the consideration of tools for optimizing the energy use of storage and networking equipment to another paper.

The next section lays the groundwork for the remainder of the paper by defining the term metric and discussing the proper use of metrics.

CONTEXT FOR THE DATA CENTER PRODUCTIVITY METRIC

WHAT IS A METRIC?

Simply put, a metric is a measuring stick. In other words, a metric is a scale for measuring some important characteristic of an object or system and includes a procedure or methodology for making this measurement. Why are metrics important? Again, to state it simply, one can not manage and improve something that is not measured. Implementing a metric allows the manager of a system to know how well the system is performing at some point in time. This makes it possible to adjust one or more tunable parameters of the system and assess the impact on the system when measured again utilizing the same metric. In this way it becomes possible to optimize whatever aspect of the system that the metric quantifies. To obtain a specific desired goal for the system, however, the particular metric one utilizes must be chosen carefully. Using the wrong metric in this process will lead to either erratic or invalid results.



WHAT IS DATA CENTER PRODUCTIVITY?

While efficiency and productivity are closely related, efficiency tends to focus on reducing costs by eliminating unnecessary expenditures of various resources required to produce a work output. Productivity, while not ignoring this aspect, focuses on increasing the amount of useful work produced for a given expenditure of resources. When applied to the data center, productivity is the quantity of useful information processing done relative to the amount of some resource consumed in producing the work.

ROLE OF POWER-PERFORMANCE BENCHMARKS

Power-performance benchmarks (which are one type of metric) have an important role in the life cycle of the data center. Knowing when to use a power-performance benchmark as opposed to some other metric is very important. The following discussion develops the distinct context of use of power-performance benchmarks. Consider the design phase of a data center (either new construction or an upgrade). Traditionally, the data center designer has carefully considered the peak processing capacity of the IT equipment to be installed comparing it to the total processing requirements of the data center. For example, the peak processing capacity of a server will determine the number of these servers that need to be deployed in order to meet current customer Service Level Agreements (SLAs) and to satisfy anticipated future peak compute capacity growth requirements. The designer has also traditionally considered the peak power usage of the equipment to be installed in order to size the power delivery and cooling equipment required for the new data center. As the focus on the energy consumption of data centers has drastically increased, assessing the power efficiency of both infrastructure equipment and IT equipment is now an essential part of data center design best practice. While there are standard metrics for assessing the power efficiency of both electrical power conversion equipment and cooling equipment, up until recently there have not been standard metrics for assessing the power efficiency of IT equipment.

Power efficiency of IT equipment is best determined by assessing the processing capacity of the piece of equipment (in its specific or typical use in the data center) in relationship to its power consumption. Since this is prior to actual selection, the data center designer needs an objective method of comparing the power efficiency of two or more otherwise equivalent pieces of equipment. In this phase, being able to examine and compare publicly available standardized power-performance benchmark scores of specific models of computers or other equipment being considered for selection is invaluable. Benchmark tests are carried out using well-defined test conditions, standardized configurations, specific software, and a standardized

synthetic workload. While this workload is by necessity not the exact workload that the server will run in any actual installation, it serves as a proxy for a specific workload type and enables objective comparisons without actually purchasing the server systems being considered, installing them, and testing them in their actual application.

The Standard Performance Evaluation Corporation (SPEC[®]) organization has just recently begun to fill this void with the public announcement of the SPECpower_ssj[™]2008 benchmark [SPEC, 2008]. The SPECpower_ssj2008 benchmark assesses Java Virtual Machine (JVM) performance of a server at a number of different loading levels while measuring the system-level power consumed. This benchmark is useful for making comparisons of the JVM power-performance of two or more otherwise equivalent server systems. Future power-performance metrics based on other workloads will round out the set of standardized benchmarks needed to select power efficient servers for a range of different applications.



Now consider the operational phase of a data center. In the operational phase of the life cycle of a data center, power-performance benchmarks are of little practical use. Once a piece of IT equipment has been purchased and installed, it serves no purpose to attempt to rerun a synthetic workload on that specific piece of equipment. Assuming that the data center infrastructure has been right-sized and is being operating efficiently (these are issues to be considered in other white papers), the issue at this point is how one measures the operational performance or productivity of the installed IT equipment relative to its energy use based on its real-life day-to-day workload with the goal of improving this quantity. This requires a completely different methodology. A proposal for just such a methodology will be presented below. But first, a couple of existing efficiency metrics will be considered.

ROLE OF OTHER CURRENTLY DEFINED EFFICIENCY METRICS

This section examines two metrics that have been proposed to allow data center managers and operators to determine whether their data center and/or specific equipment within their data center has been configured and is being operated in a power efficient manner. Below we consider Power Usage Effectiveness and Compute Power Efficiency.

POWER USAGE EFFECTIVENESS (PUE)

The Green Grid in its white paper, “The Green Grid Data Center Power Efficiency Metrics: PUE and DCiE” [Green Grid, 2007], has defined the metric Power Usage Effectiveness (PUE) and its reciprocal Data Center infrastructure Efficiency (DCiE) as a tool to assess the amount of power used by a data center to deliver a certain amount of power to the IT Equipment and remove the waste heat generated. As described in the reference, DCiE was formulated to assess the overall power efficiency of the data center infrastructure and not the ability of the IT equipment to convert electrical power into useful information processing work.

The referenced paper notes the appropriate application domain of the DCiE metric, and points to the strategic need for a new class of metrics that quantify the useful work that a data center produces in relation to the power it consumes. It names this metric Data Center Productivity (DCP).

This paper refines the definition of DCP by generalizing it to describe the amount of useful work that a data center produces relative to any specific resource it might consume to produce this work. This change expands DCP into a family of productivity metrics. In the sections that follow this paper defines a framework for quantifying the useful work that a data center produces and presents the first member of the DCP family

to be defined. This metric quantifies the amount of useful work a data center produces relative to the amount of energy it consumes to produce this work and is therefore called Data Center energy Productivity (DCeP).

COMPUTE POWER EFFICIENCY (CPE)

The paper “Metrics and an Infrastructure Model to Evaluate Data Center Efficiency” [Belady and Malone, 2007] proposes a new metric that seeks to quantify the overall efficiency of a data center while taking into account the fact that not all electrical power delivered to the IT equipment is transformed by that equipment into a useful work product. Some of the equipment within a data center consumes power while it sits idle. Other equipment is being used but not at 100% of its capacity. We will look at the equation for CPE:

$$\text{CPE} = \text{IT Equipment Utilization} / \text{PUE}$$

OR equivalently

$$\text{CPE} = (\text{IT Equipment Utilization} * \text{IT Equipment Power}) / \text{Total Facility Power}$$



This metric does have some interesting features. It is unit-less and has a maximum value of 1.0 or 100%. This allows the CPE of two different data centers to be compared since there is no need to convert the units of measurement. In addition, the ultimate goal for this metric is clear – to reach a CPE of 100%.

This metric has a number of difficulties, however. First, it uses “utilization” as a measure of useful work. There is no clear definition of utilization that works for all IT equipment in all applications. It can be assumed for servers that what is intended for utilization is actually CPU utilization, but as described in the cited paper, it remains merely conceptual. Using CPU utilization as a proxy for the normalized amount of useful work a server produces, however, does not account for processing overhead which produces no useful work. It also does not adequately account for I/O intensive workloads where the CPU is essentially idle waiting for I/O requests to complete. Also it should be noted that power management, if enabled, drastically skews the measurement of CPU utilization.

In preparation to discuss the new metric Data Center Productivity, we discuss the desired attributes of a data center productivity metric.

DESIRED ATTRIBUTES OF A PRODUCTIVITY METRIC

As discussed above, power-performance benchmarks are most appropriately employed during the equipment selection phase of the lifecycle of a data center. The focus of this paper is in optimizing productivity in the operational phase of the data center life cycle. Therefore, a metric that quantifies the useful work that a piece of equipment (or set of equipment in aggregate) produces and then relates it to the amount of resources it consumes, provides the correct tool. In the case where the resource to be tracked is energy, the metric should attempt to quantify the useful work that the data center performs per kWh of energy it consumes.

A metric that requires a piece of equipment to be taken offline to run the measurement is not appropriate for an operational data center. The metric measurement process should be minimally disruptive to the normal work flow within the data center. Ideally, it should have no impact on normal production. The ideal metric should work on whatever workload the data center is currently processing and not substitute or inject any sort of synthetic workload. The metric should take into account that not all tasks that a data center performs have equal value to the end user or the business interests of the owner of the data center. It should account

for the fact that the value of a task is a function of time. Based on the needs of the customer (which may be formalized in a Service Level Agreement [SLA] contract), some tasks have more or less value depending on the elapsed time required to run to completion. Other tasks have a constant value up to some absolute time deadline. When the deadline passes, the value drops, perhaps to zero. This metric should also be able to account for these concepts. The metric should recognize the fact that each data center operator will have a different take on the value of each of the various tasks that it runs during a given period of time. Finally, the metric should utilize LEAN principles, in other words, it should only “give credit” for work that is useful to the end customer.



While it is important to have a metric that allows one data center to compare its productivity to another data center and also to some objective standard applicable to all data centers, this is not the first priority. The first priority and the current scope of this paper is to give the data center operator a tool that allows him or her to benchmark the productivity of the data center and facilitate comparisons of this characteristic of the data center before and after actions are taken with the intent to optimize the resource consumption of the data center being managed based on its current workload.

Now that we have laid the groundwork, let’s examine the definition of data center productivity and a new family of metrics used to quantify it.

INTRODUCING A NEW FAMILY OF METRICS

This section refines the definition of the new metric Data Center Productivity (DCP) that has already been publicly disclosed by the Green Grid. Herein the definition of DCP is generalized to quantify the useful work that a data center produces as a function of any of the various resources that might be used to produce the work. Therefore, based on this generalization, it becomes an entire family of new metrics.

DATA CENTER PRODUCTIVITY

As noted above, Data Center Productivity is simply the amount of useful work that a data center produces in relation to the amount of a consumable resource that it employs to produce this work.

Mathematically this can be expressed as:

$$\text{DCP} = \text{Useful Work Produced} / \text{Total Quantity of a Resource Consumed Producing this Work}$$

From this parent metric an entire family of metrics can be derived based on the specific resource that is to be optimized. For example, one might be interested in the amount of work a data center produces per peak power consumed or per square foot of floor space utilized. The metrics of this family will all be designated by a name of the form Data Center [resource] Productivity with the corresponding acronym of the form DCxP. This paper will focus on useful work produced relative to the energy consumed producing this work. This metric is called Data Center energy Productivity (DCeP).

DATA CENTER ENERGY PRODUCTIVITY METRIC

The goal stated above is to define a metric that quantifies the useful work that a data center produces based on the amount of energy it consumes.

Mathematically this can be expressed as:

$$\text{DCeP} = \text{Useful Work Produced} / \text{Total Data Center Energy Consumed Producing this Work}$$

Note that since we are considering energy and not power, the period of time over which energy is measured must be specified to make this metric meaningful. This time period shall be called the assessment window. Energy is measured by the integral of instantaneous power over a specific time interval.

MEASURING ENERGY CONSUMED

Before we examine how to quantify useful work, let's look at measuring the quantity that constitutes the denominator in the above equation - the energy consumed by a data center during the assessment window. This paper assumes that either the electrical power feed to the entire data center is instrumented or that each piece of equipment that makes up the data center including its power conditioning and distribution and cooling infrastructure equipment is separately instrumented, and that it is capable of reporting its current power utilization. Note that total data center energy may also be estimated based on a measured value of the total energy consumption of the IT equipment multiplied by the current data center PUE value given that this value is available. Decreasing PUE or equivalently increasing DCiE has the effect of improving DCeP.



The details of this instrumentation and the standard protocols used to communicate this information are being addressed within the Technical Committee of The Green Grid Association. Assuming this instrumentation exists, a data center management system can sample the power consumption of the entire data center (or poll each piece of equipment and accumulate the sum) at regular intervals during the assessment window and calculate the power via a numerical estimation of an integral function. Also, if it is available, the average power consumption of the entire data center or the aggregate power of all the equipment over the assessment window may be used. In this case, multiplying this quantity by the length of the assessment window provides a reasonable estimate of the actual energy consumed.

DEFINING USEFUL WORK

The DCP metric and all its derivative metrics require the quantification of useful work. Useful work may be defined by the equation:

EQUATION 1

$$\text{Useful Work} = \sum_{i=1}^M V_i * U_i(t, T) * T_i$$

where M is the number of tasks initiated during the assessment window, and V_i is a normalization factor that allows the tasks to be summed numerically, and $T_i = 1$ if task i completes during the assessment window, and = 0 otherwise.

$U_i(t, T)$ is a time-based utility function for each task, where the parameter t is elapsed time from initiation to completion of the task, and T is the absolute time of completion of the task.

Note that Useful Work is defined to be the sum over i of all tasks 1 through M initiated within the assessment window multiplied by a time-based utility function $U_i(t, T)$. The factor V_i assigns a normalized value to each task so that they may be algebraically summed. T_i eliminates all tasks that are either initiated prior to the assessment window or are initiated within the window but do not complete.

The following sections will discuss the key terms introduced in Equation 1 above.

DEFINING A TASK

To execute this measurement, all the tasks initiated within the assessment window must be known. Here it



is useful to discriminate between the concept of a task type and the concept of a task instance. A task type is descriptive of a specific class of processing that the data center provides and involves the invocation of a specific piece of application software installed in the data center. A task instance is a single invocation of this software with a specific set of input parameters or data and a specific resultant output. While a given data center may in the course of one day carry out a very large (perhaps in the millions) number of task instances, it will normally process a much smaller number of task types. Task types are defined prior to the assessment of DCeP based on the installed equipment and software within the data center. To simplify the quantification of Useful Work, tasks are aggregated according to task types. Task instances within a given task type will all have the same relative value and must comply with the same service level agreement. This means that the parameters V_i and $U_i(t,T)$ are determined on a per task type basis instead of per task instance.

The formulation of Useful Work leaves the definition of what is considered a “task” up to the person personalizing the metric for use in a given data center. This makes the metric applicable to any workload. However, if a task is defined too broadly, for example, “maintain data base XYZ” it will not be possible to determine if the task completes within the window. This is solved by redefining such a task at a finer level of granularity. For example, the task “maintain data base XYZ” could be broken down into a number of typical subtasks involving data base XYZ such as “satisfy query against data base XYZ” or “load a new record into data base XYZ,” or “run standard report A against data base XYZ.”

DEFINING THE ASSESSMENT WINDOW

To quantify the energy consumed while executing the measurement, a specific time window must be established. This time window is called the assessment window. The length of this window is arbitrary, but to obtain accurate results in executing the measurement, it should be no shorter than about 20 times the mean run time of the any of the tasks initiated in the assessment window. An assessment window should be sized in accordance with the nature of the workload and the purpose of the measurement. For example, a useful assessment window could be as short as a few milliseconds or as long as a month or more.

Note that this methodology as defined ignores any tasks that may have been initiated prior to the start of the assessment window and those that are initiated within the window but do not complete prior to the end of the window. These tasks will consume energy during the assessment window, but will not contribute to Useful Work. These effects lead to an error in the measurement of Useful Work. This error, however, may be minimized by appropriately sizing the length of the assessment window.

ASSIGNING A VALUE TO TASKS

It is clear that not every task that the IT equipment in a data center performs has the same value. Yet, in order to aggregate the useful work that a server or group of servers produce, the tasks must be normalized. This is the purpose of the factor V . The value of V_i must be assigned prior to the assessment of the DCeP metric for each task type so that the value of the task is normalized to some standard task that the data center performs. In this way, more important or valued tasks receive greater weighting in the calculation of Useful Work and the completion of less important tasks receive a lesser weighting. When appropriate, a straight-forward simplification of this process of determining the V_i weights is to set them all to 1.0. This indicates that all defined task types have approximately the same value to the end user or the owner of the data center.

DEFINING A TIME-BASED UTILITY FUNCTION

Note that the function $U_i(t,T)$ must also be specified for each task type prior to running an assessment of the

metric. This function handles the time dependent nature of the value of each task. t is relative run time while T is the absolute time of completion. A given utility function can ignore one or both parameters. In other words, U_i can be a constant, a function of the relative run time of the task (in which case the function may be denoted by $U_i(t)$), a function of only the absolute completion time ($U_i(T)$) or a function of both the run time and the absolute completion time ($U_i(t,T)$). If the value of completing a given task is time invariant, U_i for that task should be expressed as a constant.

The following will discuss a typical run time based utility function.

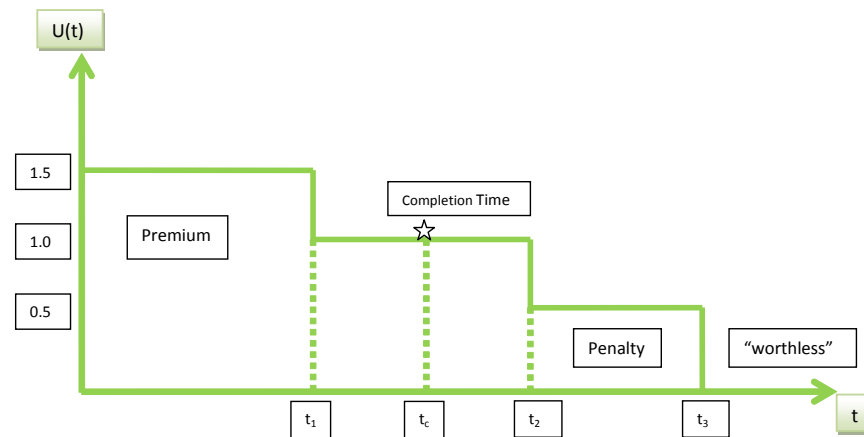


Figure 1. Example Utility Function

Refer to Figure 1 above. This example shows how the value of a task may change as run time increases. Assume that a service level agreement (SLA) applies to this task. In this case, the time based utility function merely represents this SLA in mathematical form. Note that this SLA provides for a premium to be paid to the service provider if the task is completed early (prior to t_1). If it is completed in the interval t_1 to t_2 , the SLA is met. If it does not complete by t_2 the service provider has agreed to a penalty of 50%. At t_3 this penalty goes to 100%. This is, of course, a hypothetical example, but it displays the power of the utility function to capture mathematically the terms of an SLA. If no SLA exists, other means may be used to decide the form of the utility function.

For other tasks, the most important determinant of utility is absolute completion time (e.g., payroll must complete by 3:00 AM Friday morning.) In this case, $U_i(T)$ would be a function of absolute time alone. Note that an easy simplification is to assign a constant value of 1.0 for the utility function.

With the forementioned definition of DCeP we now have a mathematical basis for measuring the total work output of a data center and relating it to the energy consumed to produce this work output.

TRANSACTIONAL OR THROUGHPUT-BASED WORKLOADS

Several of the simplifications that may be made to the equation for useful work (Equation 1) have been discussed above. One specific simplification appropriate for transactional or throughput-based workloads is to set both V_i and U_i to 1.0. With these assumptions, Equation 1 represents the simple act of counting the tasks (transactions or operations) that complete during the assessment window. Note that for this case all tasks are deemed to have equal value. This simplification is often used when measuring the performance of a transaction processing application, for example. Typically transaction-based applications will be

instrumented to measure and report the average number of transactions processed per second. Other applications may report performance in terms of operations per second, queries per second, web-pages served per second, or jobs completed per second. Using this reported performance value useful work may be estimated by the equation:

$$\text{Useful Work} = \text{Average Transaction Rate} * \text{Length of Assessment Window}$$

where any rate-based performance metric may be substituted for Average Transaction Rate



Note that the denominator of the metric remains the energy consumed during the assessment window. Note also that the length of the assessment window could be any length of time that the data center operations staff deems to be suitable for their business.

The following section describes the results an experiment that demonstrates the measurement of the DCeP metric under controlled laboratory conditions using a pre-defined synthetic workload.

EXPERIMENTAL RESULTS

SETUP

To provide a concrete example of how the DCeP metric is evaluated, we ran an experiment on an instrumented system of IT equipment diagramed in Figure 2.

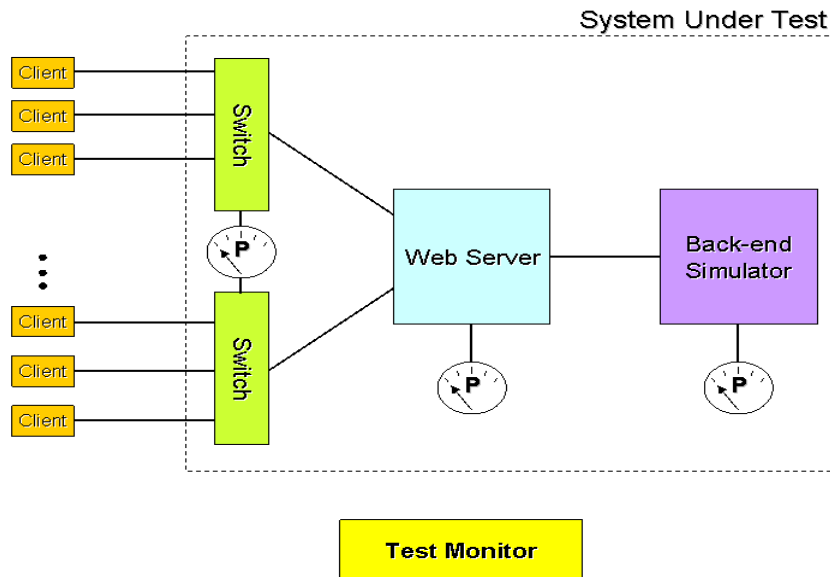


Figure 2. DCeP Demonstration Experimental Setup

The Switches are SMC TigerSwitch 8624T Ethernet switches. These are connected to the Web Server which is a dual socket Barcelona (quad-core AMD Opteron™) based server. The Back-end Simulator runs on a



SunFire V40Z server and is connected to the Web Server via Ethernet. This experimental setup is configured to run the eCommerce workload used in the SPECweb®_2005 standardized benchmark. Three Exttech model 380803 power meters are hooked in series with the AC line cords of the equipment as shown to measure the power consumed by the equipment. The block Test Monitor represents those computer systems outside the system under test that are used both to run the test and gather the instrumented task completion and power data.

We use a synthetic workload not because it is a prerequisite to assessing the DCeP metric, but rather to allow running the experiment used to generate the data in a reproducible manner. We chose the eCommerce workload because it is easily accessible to researchers and is relatively simple.

The eCommerce workload simulates a typical Web-based direct marketing application. A configurable number of simulated Clients generate requests to a Web Server that services them. In order to display products to be viewed by the end users (simulated by the clients) and aid in their customization and purchase, the Web Server must interact with a back-end data base. This data base is simulated by the SunFire V40Z server (labeled Back-end Simulator) in the figure.

The eCommerce workload is generated by a plurality of simulated Clients each generating requests to the Web Server driven by a statistical process. Each client goes through a selection and ordering session based on a defined state transition scheme. During this simulated ordering session each client generates a number of request types and receives simulated responses in terms of composed Web pages. The eCommerce workload defines 13 request types. These are: *index, search, browse, browse_productline, productdetail, customize, customize2, customize3, cart, login, shipping, billing, and confirm.*

These request types satisfy all the needed attributes of task types as defined by the DCeP metric. We decided for the sake of running this demonstration that the relative value of all 13 task types would be weighted the same. Therefore the V_i parameter is set to 1.0. Only conforming request / response transactions (those that complete within a specified minimum response time) are counted by the test monitor software. Therefore we can ignore the time dependency in evaluating the relative value of tasks. This is carried out by setting the utility function U_i,T to the constant 1.0. As a result of both of these decisions, the equation for Useful Work becomes

EQUATION 2

$$\text{Useful Work} = \sum_{i=1}^M T_i$$

Note that this is merely a count of the tasks completed within the assessment window. The test logging software provides this count directly.

RESULTS

We first established the maximum number of clients that could be supported by the system under test without violating the response time requirement. This was determined to be 6000 clients. We then backed off the number of clients to 66% of the maximum. After a warm up phase which allows memory and processor caches to be primed and the mean response time to reach steady state, we triggered the start of the assessment window. We set the assessment window at 30 minutes, since this is the typical run time for one iteration of the SPECweb benchmark. Both the number of completed request / response transactions and the power consumed by the equipment under test were logged during the assessment window. Time stamps allow correlation of all data logged. Power measurements are converted to energy amounts via a

simple numerical approximation of an integral in which the power level measured at each sample point is multiplied by the sample time and these products are summed for the entire assessment window. For the run at 66% of maximum supported clients 675,075 individual request / response transactions were completed during the 30 minute assessment window. During this time the equipment under test consumed a total of 0.4582 kWhrs of energy. We then uplift this number by an estimate of the total “data center” energy consumed by assuming that an additional 0.7 Watts of power is required to condition the power delivered to the IT equipment under test and to cool it for each watt consumed. This is equivalent to a Power Utilization Effectiveness (PUE) of 1.7.



Therefore we have:

$$DCeP_{exp} = 675075 / (0.4582 * 1.7) = 866,658 \text{ Normalized Tasks/ kWhr}$$

This quantifies the useful work produced by the energy consumed by the data center during the assessment window. The metric can be evaluated using a much shorter assessment window without introducing significant measurement error. Shortening the assessment window allows the metric to be more closely approximate an instantaneous assessment of productivity.

SUMMARY

The proposed DCeP metric meets all the requirements discussed in the earlier Section “Desired Attributes of a Productivity Metric” To summarize again, these are:

- Quantifies useful work relative to the amount of resource the data center consumes producing the work
- Does not interfere with normal operation of the data center
- Does not require equipment to be taken offline
- Generally applicable; works on any workload
- Although a synthetic workload was used in the experiment shown, in a real application of this metric, no additional workload needs to be injected into the data center make the measurement
- Corrects for differences in the relative value of the various types of tasks that a data center runs
- Accounts for the time sensitive nature of the tasks a data center must process
- Handles both run time and absolute completion time
- Utilizes LEAN principles, i.e., only “gives credit” for work that is useful to the end customer or owner of the data center

As will be discussed in the following section, further research is required to extend the DCeP metric to allow the comparison of productivity of different data centers.

FURTHER RESEARCH

This work provides a firm theoretical foundation for the quantification of data center useful work which in turn provides a basis to measure the useful work produced relative to the amount of any specific resource that the data center consumes in producing this work. A current weakness of this formulation is that the definition of the metric requires some manual intervention to define tasks types and assign values V_i and utility functions $U_i(t,T)$ for each task type.

The first extension of this work will be to investigate tools and techniques that might enable the automation of the process of defining tasks and assigning the task type values and utility functions.

The scope of this first formulation of DCeP was purposely limited to the context of a given data center. In other words, it is useful for comparing the current energy productivity of a given data center with an earlier benchmark productivity that has been established for that data center. Ultimately The Green Grid would like to define a Data Center Productivity metric that will allow the productivity of one data center to be compared to the productivity of another. This work will be the topic of significant research effort in 2008.



REFERENCES

[Belady and Malone, 2007]

Belady, Christian L. and Malone, Christopher G., “Metrics and an Infrastructure Model to Evaluate Data Center Efficiency”, Proceedings of IPACK2007, July 2007.

[EPA – Report to Congress]

U.S. Environmental Protection Agency, “Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431”, August 2007.

http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf



[Green Grid, 2007]

The Green Grid, “The Green Grid Data Center Power Efficiency Metrics: PUE and DCiE”, December 2007.

http://www.thegreengrid.org/gg_content/TGG_Data_Center_Power_Efficiency_Metrics_PUE_and_DCiE.pdf

[Koomey, 2007]

Koomey, Jonathan G. Ph.D., “Estimating Total Power Consumption by Servers in the U.S. and the World”, February 2007.

<http://enterprise.amd.com/Downloads/svrpwrusecompletefinal.pdf>

[Rasmussen, 2006]

Rasmussen, Neil, “Electrical Efficiency Modeling of Data Centers”, 2006.

http://www.apcmedia.com/salestools/NRAN-66CK3D_R1_EN.pdf

[SPEC, 2008]

For information on the SPECpower_ssj2008 Power-performance benchmark refer to :

http://www.spec.org/power_ssj2008/

[Stanley, 2007]

Stanley, John R., Brill, Kenneth G., and Koomey, Jonathan, “Four Metrics Define Data Center ‘Greenness’”, Uptime Institute, 2007

[http://uptimeinstitute.org/wp_pdf/\(TUI3009F\)FourMetricsDefineDataCenter.pdf](http://uptimeinstitute.org/wp_pdf/(TUI3009F)FourMetricsDefineDataCenter.pdf)