Energy Benchmarks: A Detailed Analysis

Meikel Poess
Oracle Corporation
500 Oracle Pkwy, Redwood Shores
CA-94065, +1-650-633-8012
meikel.poess@oracle.com

Raghunath Othayoth Nambiar
Hewlett Packard Company
11445 Compaq Center Dr. W, Houston
TX-77070, +1-281-536-2728
raghu.nambiar@hp.com

Kushagra Vaid
Microsoft
One Microsoft Way, Redmond
WA 98052, +1-425-722-4038
kvaid@microsoft.com

John M. Stephens, Jr.
Gradient Systems
643 Bair Island Rd, Redwood City
CA-94063, +1-650-556-9380
jms@gradientsystems.com

Karl Huppler
IBM Corporation
3605 Highway 52 N Rochester
MN 55901
huppler@us.ibm.com

Evan Haines
ICF International,1725 Eye St NW
Suite 1000 Washington,
DC 20006 +1-202-862-1158
ehaines@icfi.com

ABSTRACT

In light of an increase in energy cost and energy consciousness industry standard organizations such as Transaction Processing Performance Council (TPC), Standard Performance Evaluation Corporation (SPEC) and Storage Performance Council (SPC) as well as the U.S. Environmental Protection Agency have developed tests to measure energy consumption of computer systems. Although all of these consortia aim at standardizing power consumption measurement using benchmarks, ultimately aiming to reduce overall power consumption, and to aid in making purchase decisions, their methodologies differ slightly. For instance, some organizations developed specialized benchmarks while others added energy metrics to existing benchmarks. In this paper we give a comprehensive overview of the currently available energy benchmarks followed by an in depth analysis of their commonalities and differences.

Categories and Subject Descriptors

B.8.2 [Performance Analysis and Design]: Hardware – Performance and Reliability – Performance Analysis and Design

General Terms

Measurement, Performance, Standardization

Keywords

Energy Measurement, System Performance, Performance Analysis

1. INTRODUCTION

The most prominent industry standard consortia for performance measurements are the Transaction Processing Performance Council (TPC), the Standard Performance Evaluation Corporation (SPEC) and the Storage Performance Council (SPC). All major computer and system vendors are members of these organizations. Each of these consortia addresses different, often unique aspects of computer system performance. Due to the massive increase in energy cost and a boost in energy consciousness, these organizations have added methodologies and metrics to measure energy efficiency in addition to existing performance metrics. Although all are aimed at aiding customers in the purchase decision process they follow slightly different approaches and philosophies in defining and policing the measurement of power consumption of computer systems.

The Transaction Processing Performance Council (TPC), founded in 1988, defines transaction processing and database benchmarks and disseminates objective, verifiable TPC performance data to the industry. While TPC benchmarks involve the measurement and evaluation of computer transactions, the TPC regards a transaction as it is commonly understood in the business world: a commercial exchange of goods, services, or money. The TPC offers currently two benchmarks to measure On Line Transaction Processing systems (OLTP), namely TPC-C and TPC-E [4] and one to measure decision support performance (TPC-H [7]). In 2007 the TPC formed a committee to add energy metrics (TPC-Energy [16]) to all its benchmarks, which will coexist with existing metrics of all its benchmarks including a comprehensive power measurement for all the components including the database server, middle tier, and storage subsystem and connectivity devices. Earlier work allowed the power consumption estimation of TPC-C and TPC-H systems [8][9].

SPEC has been the front-runner to announce the first industry standard benchmark that measures power consumption in relation to performance for server-class computers. It is a non-profit corporation formed to establish, maintain and endorse a standardized set of relevant benchmarks that can be applied to the newest generation of high-performance computers, including processor-intensive benchmarks, benchmarks to measure graphics and workstation performance, high performance computing benchmarks, Java Client/Server benchmarks, mail server benchmarks, network file system benchmarks and SPECpower_ssj2008 [10], a benchmark focused on the relationship of power and performance. SPECpower_ssj2008 reports power consumption for servers at different performance levels, from 100 percent utilized to idle in 10 percent segments, over a set period of time.

The Storage Performance Council (SPC) is a non-profit corporation founded to define, standardize, and promote storage subsystem benchmarks as well as to disseminate objective, verifiable performance data to the computer industry and its customers. Since its founding in 1997, SPC has developed and publicized benchmarks and benchmark results focused on storage subsystems as well as the adapters, controllers, and storage area networks that connect storage devices to the computer system. Simultaneously to the above efforts, the U.S. Environmental Protection Agency (EPA) is working on benchmarks to label computer servers as energy efficient under broad operating conditions. First established in 1970, the charter of the EPA is to consolidate in one agency a variety of federal research, monitoring, standard-setting and enforcement activities to ensure environmental protection. EPA's mission is to protect human health and to safeguard the natural environment—air, water, and land—upon which life depends. The EPA recently released the first version (1.0) of its Energy Star® specification for servers [1] [2]to help consumers identify systems that deliver performance while reducing energy consumption. On average, computer servers that earn the Energy Star label will be 30 percent more energy efficient than standard servers. In parallel with this and other product efficiency efforts, EPA has also launched the development of a building efficiency metric for data center facilities.

The wide variety of application areas and methodologies of these benchmarks become increasingly confusing for data center managers. This paper aims to compare the currently available energy tests, how results should be interpreted, and how they can be used to make energy-conscious purchase decisions.

Sections 2 through 5 give detailed descriptions of the current approaches and philosophies to energy measurements by SPC, TPC, SPEC and EPA. It is followed by a discussion of the commonalities and differences in Section 6.

2. TRANSACTION PROCESSING PERFORMANCE COUNCIL

The TPC has two active benchmarks: TPC-C and TPC-E for online transaction processing and TPC-H for ad-hoc decision support. While these benchmarks serve existing market needs for allowing comparisons between system configurations for transactional application performance and price performance, there is a strong urge from the industry community to even better serve end-users needs by allowing the reporting of energy efficiency for each of the above benchmarks. Having energy efficiency reporting will allow IT organizations to evaluate the merits of a particular system based not only on transactional performance and price performance, but also on how much power was consumed to deliver that performance level. With increased awareness in the industry for power optimized designs and corporate initiatives to reduce carbon emissions, adding an energy efficiency component to TPC benchmarks will stimulate greater innovation in power systems design by ensuring a consistent methodology for reporting efficiency metrics, and ensuring that end-users would have important knowledge about the impact to their bottom-line resulting from energy costs when evaluating future system purchases [3]. With this goal in mind, the TPC has developed a consistent energy usage component for all its existing TPC benchmarks.

The TPC-Energy [16] specification augments the existing TPC benchmarks by allowing for optional publications of energy metrics alongside their performance results. Wherein the reported TPC performance metrics correspond to the amount of work completed per unit of time, the TPC-Energy metric measures the energy consumption corresponding to the amount of work completed. The metric is represented as the ratio of the Energy (typically measured in Watts-seconds) consumed by all components of the benchmarked system – this includes servers, storage, clients, network switches – to the work completed (typically measured as number of Transactions) over the benchmark interval. After moving the time element to the denominator, the TPC-Energy metric is plainly represented as Watts/Performance.

Since the TPC-Energy specification is common to all TPC benchmarks, it needs to accommodate their differences in workload characteristics (time based vs. task based), number and type of system components (1 Tier, 2 Tier or Tier 3) and metrics. TPC-C and TPC-E follow a time based benchmark model. They report performance as the transaction throughput during steady state condition. The performance is measured during the measurement interval, which must begin after the system reaches steady state, be long enough to generate reproducible throughput results that would be representative of the performance that would be achieved during a sustained eight hour period and extend uninterrupted for a minimum of 120 minutes.

TPC-H follows a static task benchmark model, which is divided into three distinct measurement tests (load, single-user and multi-user). All three tests exhibit an oscillating system utilization behavior. Most TPC-H tasks do not utilize the entire system the entire time. That is, not all resources (e.g. IO, CPU, Memory, Network) are fully used during the execution of each task. This is embedded in the nature of the different tasks. For instance, a hash join is CPU bound during the build phase of its hash table and, usually, IO bound during its probe phase. TPC-H's 22 queries are so diverse that it is impossible to setup a system that assures 100 percent utilization 100 percent of the time during the power test. Consequently, the system consumes more power in the IO subsystem during some time of the power test and more CPU power during other times of the power test.

To deal with such scenarios, the TPC-Energy specification requires measuring power P_i[W] of the entire system for each interval I $(i_1, i_2, \ldots, i_n)$ in addition to the performance measurements T_i [tpmC], [tpsC],[QphH], and then independently determining the combined value for power P [W] and performance for all intervals using weights corresponding to the duration of each interval: $\text{Overal Work} = \sum_{i=0}^{n} \mathrm{T_i} * \mathrm{S_i}$; $\text{Overal Energy} = \sum_{i=0}^{n} \mathrm{P_i} * \mathrm{S_i}$.

$$\mathrm{P} = \frac{Overall\ Work}{Overall\ Energy} = \frac{\sum_{i=0}^{n} \mathrm{T_i} * \mathrm{S_i}}{\sum_{i=0}^{n} \mathrm{P_i} * \mathrm{S_i}}$$

Most TPC benchmark configurations tend to be large scale setups that can be challenging from a power measurement standpoint since the different subsystems cannot be all measured by a single power meter. In such cases, TPC-Energy allows configuring the test harness (see Figure 1) with separate power meters to enable the energy measurements of individual components (e.g. server chassis, storage subsystem(s), clients, network switches etc.). The energy data may be gathered from multiple power devices and combined by either aggregating data at the same points in time or by gathering the data in multiple runs and combining the totals. In case the energy data is gathered over multiple runs the TPC Performance Metric of each run must be within 2% of each other, and the duration of each run must be at least 98% of the duration of the run for which performance metrics are reported. This capability aids cases where it is technically not feasible to measure all subsystem components simultaneously during the performance run. It also allows for reporting secondary metrics.

The TPC-Energy specification also requires measurement and reporting of several other aspects of the system environment that can affect the energy computation. A temperature sensor is required at the server input to ensure that the ambient air is not below 20^{o}C since lower temperatures typically result in lower fan power consumption and better heat dissipation. Humidity and altitude (above sea level <1.1 Atm) measurements also need to be logged to ensure that these environmental parameters are within bounds to not affect power calculations. Most importantly the power meter used in the measurements need to satisfy the accuracy, defined in the specification. All power meters must be NIST [6] certified less than one year before the benchmark run and have an average inaccuracy of not more than 2%. The overall Accuracy Factor af is calculated using the vendor provided accuracy factors af_i for all devices as: $af = 1 + \sum_i^n af_i$. These devices are typically the power analyzer, the coupling device (inline, feed-thru, clamp-on), and range selection. The overall accuracy factor is then applied to the Overall Observed Power Meter Reading P_o to calculate the Compensated Power Value P_c as: $P_c = af_i * P_o$. For example, consider an analyzer using clamp-on coupling with a wattage range selection of 5-50W, with a vendor-specified power analyzer accuracy of +/- 0.05% of reading, used in a test with a wattage value of 15KW, and a clamp-on probe with an accuracy of +/- 1.0% and 1500:1 input/output ratio. The overall accuracy is computed as: (0.05%+1.0%)*15,000W=157.50W/15,000W = 1.05% at 15KW. This calculation is performed on each measurement before combining, scaling, or averaging multiple measurements. This technique allows the use of a wide variety of measurement devices and to pessimistically compensate for their varying accuracy characteristics.

2.1 Primary and Secondary Metrics

The primary metric, reported by TPC-Energy, is in the form of "Watts per Performance" for the overall System Under Test (SUT) where the performance units are particular to each TPC Benchmark. For example the primary metric for TPC-E would be Watts/tpsE. The energy consumption is measured for all subsystems active for the duration of the benchmark run – this includes servers, storage, clients, network switches. The TPC-Energy Specification also defines optional secondary metrics. The purpose of these secondary metrics is to allow more detailed comparisons and analysis of the result for system components such as server chassis, storage system, network gear etc. The secondary metrics are represented in similar units as the primary metric, i.e. Watts/Performance, and the summation of all individual secondary metrics equals the primary metric. This is because both the primary and secondary metrics share a common value for the denominator – the performance value. This was done by design when developing the benchmark specification to allow end-users to see the contribution of the subsystems (represented by the secondary metrics) to the overall system results (represented by the primary metric).

In addition to these primary and secondary metrics, the TPC-Energy specification also calls for reporting the Idle power, which is defined as the energy consumption of the SUT within 30 minutes of the completion of the benchmark run. The intent is to represent the amount of energy consumption of a measured system in a state "ready to accept work". This is useful to customers who have systems that have periods of idle but require the system to respond to a request for work at any time.

2.2 Benchmark Harness Overview

The Energy Measuring System (EMS) is a TPC provided software package designed to facilitate energy measurements. It provides:

- Interface to instruments (power analyzers, temperature probes)
- Reliable logging of power and temperature readings
- Standardized output log of data collected
- Report generation
- Interface to external software for logging events and messages, and for displaying real-time data

The EMS consists of a single execution of the EMS Controller (EMSC) and multiple Power-Temperature Daemons (PTD) Managers and PTD executions. Each device (power and temperature) taking measurements requires one PTD and one PTDM, which is controlled via the EMSC. The EMSC also provides connections for real-time displays and an interface to the Sponsor's benchmark driver.

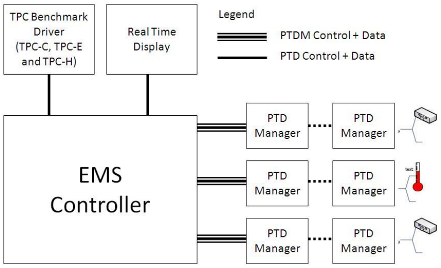

Figure 1: TPC Power Measuring Harness

2.3 Results Reporting

The published Full Disclosure Report (FDR) for TPC-Energy requires an executive summary that displays all the key metrics for performance and power measured for the benchmark run in the following way:

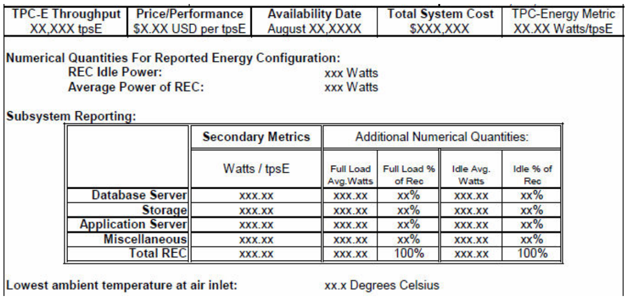

TPC-E Throughput	Price/Performance	Availability Date	Total System Cost	TPC-Energy Metric
XX,XXX tpsE	$X.XX USD per tpsE	August XX,XXXX	$XXX,XXX	XX.XX Watts/tpsE

Numerical Quantities For Reported Energy Configuration:

REC Idle Power: xxx Watts

Average Power of REC: xxx Watts

Subsystem Reporting:

	Secondary Metrics	Additional Numerical Quantities:			
	Watts / tpsE	Full Load Avg.Watts	Full Load % of Rec	Idle Avg. Watts	Idle % of Rec
Database Server	xxx.xx	xxx.xx	xx%	xxx.xx	xx%
Storage	xxx.xx	xxx.xx	xx%	xxx.xx	xx%
Application Server	xxx.xx	xxx.xx	xx%	xxx.xx	xx%
Miscellaneous	xxx.xx	xxx.xx	xx%	xxx.xx	xx%
Total REC	xxx.xx	xxx.xx	100%	xxx.xx	100%

Lowest ambient temperature at air inlet: xx.x Degrees Celsius

Figure 2: Example TPC Power Metric Reporting

3. STANDARD PERFORMANCE EVALUATION CORPORATION

In May of 2006, SPEC declared its intent to develop energy metrics for servers. SPEC took on the substantial task of changing the way that performance analysis is done on computer servers – shifting the focus away from absolute best performance at the highest possible utilization of system resources and toward a view of the relationship between the power used by a system and the performance delivered at graduated workload levels. To achieve this, changes were made to the way the base benchmark was structured. An additional tool-set was created to manage power-related information that is measured from devices that are external to the system under test. A control tool was also created to control the benchmark execution and merge the performance and power measurement information.

The result of these efforts was SPECpower_ssj2008 Version 1.0 [14], which was delivered to the general public in December 2007. The application of the SPECpower_2008 benchmark includes many features that are important to server-side Java business applications. On examination, several similarities between SPECpower_ssj2008 and SPECjbb2005 can be observed. However, there are also some very significant differences that both make results from the two benchmarks non-comparable and allow SPECpower_ssj2008 to achieve its goals of measuring the relationship between power and performance.

Just like SPECjbb2005, SPECPowerssj2008 measures the performance of a Java based middle tier emulating the Client and Database tiers. Consequently the benchmark can be run on a single server without setting up large hardware installations. It measures processor and memory performance. It executes neither disk nor network I/O. On the software side it measures the performance of the Java Virtual Machine (JVM), just-in-time compilation (JIT), garbage collection, user threads and some aspects of the operating system. The main metric, measured in business operations per second, the number of transactions (New Order, Payment, OrderStatus, etc.) divided by the elapsed time.

The benchmark execution begins with a set of calibration cycles, in which the benchmark determines the maximum throughput that is achievable for the configuration being measured. Once the maximum throughput is determined, a set of measurement points are initiated by scheduling batches of transactions in such a way that measurements can be made at 100% of the maximum throughput, 90% of the maximum, 80% and so on until reaching a final “active idle” measurement point, where the system is ready to complete work, but no work requests are scheduled.

Figure 3 shows a portion of a result disclosure for SPECpower_ssj2008 [13]. The eleven measurement points are shown in tabular and graphical form. The metric of the benchmark is the quotient of the sum of the throughput results for all points and the sum of the power measurements for each point. This places importance on the overall performance of the system, the power needed to deliver that performance, the power consumed at idle and the shape of the power characteristics curve between maximum and minimum performance points.

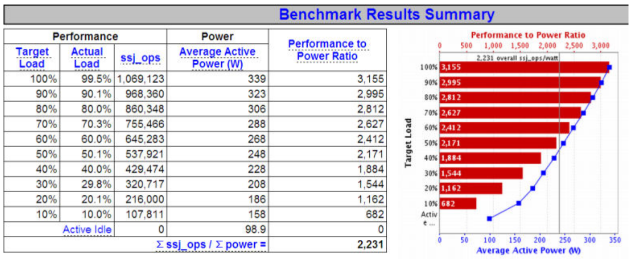

Benchmark Results Summary

Performance			Power	
Target Load	Actual Load	ssj_ops	Average Active Power (W)	Performance to Power Ratio
100%	99.5%	1,069,123	339	3,155
90%	90.1%	968,360	323	2,995
80%	80.0%	860,348	306	2,812
70%	70.3%	755,466	288	2,627
60%	60.0%	645,283	268	2,412
50%	50.1%	537,921	248	2,171
40%	40.0%	429,474	228	1,884
30%	29.8%	320,717	208	1,544
20%	20.1%	216,000	186	1,162
10%	10.0%	107,811	158	682
	Active Idle	0	98.9	0
			Σ ssj_ops / Σ power =	2,231

Figure 3: Sample Result Disclosure for SPECpower_ssj2008

As can be seen from the chart, although the metric is a computation based on all eleven measurement points, sufficient data is provided to examine the relationship of power to performance at any of these points. Thus, if it is known that a server or group of servers typically operate between the 10% and 30% ranges, some indication of their operating characteristics can be gleaned from the SPECpower_ssj2008 data.

The initial release of SPECpower_ssj2008 included support for single servers. In April 2009, support for multi-node measurements of blade servers in associated enclosures was made possible with Version 1.10 of the benchmark. Other enhancements were also provided to improve the ease of measuring, reporting and verifying benchmark results.

In June of 2009 SPEC announced the delivery of SPECweb_2009, its second benchmark with power metrics [14]. SPECweb2009 is used to assess web servers. It uses three different loads to calculate the primary performance metric. The loads are typical for the following application areas: banking, e-commerce and support. The primary performance metric states the factor at which the measured system is more powerful than a previously defined reference system. It is calculated as follows:

Similarly to SPECjbb2005 SPECweb2009 measures the performance of the middle tier emulating the Client and Database tiers. Consequently the benchmark can be run on a single server without setting up large hardware installations. Where SPECpower_ssj2009 is strictly designed as a performance-per-watt benchmark, SPECweb_2009 is designed to be both a performance benchmark and a performance-per-watt benchmark. As with SPECweb2005, there are three workload scenarios that are executed to develop a performance metric. Power is also measured for all three workloads at their rated speed. In addition, a separate power and performance set of steps, similar in concept to those measured for SPECpower_ssj2008, are measured to account for possible changes in the relationship between power and performance at different load levels.

SPECweb_2009 represents a different set of workloads and configuration requirements than those represented by SPECpower_ssj2009. As the name implies, http serving functions using network interfaces are measured. An examination of the existing results [13] shows that the configurations also require a moderate (6-25) number of disks to support the workload.

By examining information from multiple workloads, the data center manager can get a better understanding of how specific configurations are likely to use power in specific conditions.

3.1 SPEC Infrastructure Support for Power-Related Benchmarks

The SPEC organization recognized from the beginning that a single benchmark could not represent the entirety of computing environments. The application used for SPECpower_ssj2008 measures important features and functions: processor, cache, some memory use, Java Virtual Machine operations and more. However, it does not focus on use of storage or network, use of large memory configurations, computation and many other compute functions. To accommodate and encourage development of performance per watt metrics for multiple benchmarks, SPEC developed a power and performance benchmark methodology document [14] in parallel with the development of SPECpower_ssj2008. The controls associated with SPECpower_ssj2008 (Control and Collect System – CCS) were also designed to be adaptable to multiple benchmarks.

Key to measurement of power and performance is the tooling needed to control measurement parameters in power analyzers and temperature sensors and to synchronize measurements from these devices with the measurements of performance data. The SPEC Power and Temperature Daemon is a separate tool that is used by SPECpower_ssj2008 and other SPEC benchmarks. Arrangements for its use by other benchmark consortia are in place.

3.2 Future SPEC Benchmarks

SPEC has a strong reputation for maintaining the currency of their benchmarks. Examples can be found in the revisions from SPEC CPU92 to SPEC CPU95, SPEC CPU2000 and SPEC CPU2006 and also in revisions of many other benchmarks[13]. With the adoption of the SPEC Power and Performance Benchmark Methodology, we can expect that most future revisions of SPEC benchmarks will include enablement functions to allow power measurements to be taken. In particular, SPEC has already announced the inclusion of power metrics in benchmarks developed for measurement of graphical and high-function workstations [14]. Individual benchmark design teams will decide how important power measurements are for their benchmark environment. For some future benchmarks, it is likely that power measurements will be required, as they are for SPECpower_ssj2008 and SPECweb2009. For other benchmarks, it may be that power measurements are options available to the benchmark sponsors. In the final analysis, it may be consumer pressure that helps benchmark sponsors to make the decision to measure power on a wide variety of benchmark environments.

4. STORAGE PERFORMANCE COUNCIL

The Storage Performance Council (SPC) published its first benchmark, SPC-1 [12], in 2001. Since then it has developed a robust methodology for the production, validation, publication and comparison of performance results for the storage industry. With four benchmarks published and over 145 results filed, the SPC has developed a suite of benchmarks that are providing reliable, verifiable, vendor-neutral results that aid in product positioning and comparison, and guide the purchase decisions for business-critical purchases.

When the specific performance needs of storage components became clear, SPC-1C [12], was developed. Building on the work done to define a formal process to produce verifiable, vendor-neutral results, it made focused changes to the SPC-1 specification, that provide a targeted arena for the comparison on smaller, configurations and component-level storage products. Similarly, rather than develop an entirely new benchmark to address the energy consumption of storage systems, the SPC opted to build focused extensions that built on the success and infrastructure developed for its existing workloads.

The result is a set of optional energy extensions. SPC-1C/E, which builds on the benchmark definitions for SPC-1C™, was released in June 2009. SPC-1/E™, which builds on the benchmark definitions for SPC-1™, was released in October 2009, allowing results on larger, more complex systems to include energy measurements.

4.1 Overview of SPC-1 and SPC-1C Metrics

SPC-1 and SPC-1C both rely on an application-level IO load built around an identical mix of 8 different IO patterns modeled after common OLTP applications (e.g., DBMS, email server). In both cases, a benchmark execution consists for multiple phases that illustrate the number of IOs that can be serviced by the Tested Storage Configuration (TSC) within a response time constraint of 30 milliseconds, as well as the sustainability and repeatability of the performance measurement. The differences that exist between the SPC-1 performance tests (used for SPC-1/E measurements) and the SPC-1C performance tests (used for SPC-1C/E measurements) are limited to the nature of permissible configurations, with SPC-1C limited to smaller storage configurations, the duration of test phases, with SPC-1C using shorter test phases to match configuration size and complexity, and the increment used to scale the IO workload applied to the SUT (SPC-1C increases load in increments of 5 IO/sec, SPC-1 in increments of 50 IO/sec).

While a detailed review of the execution rules for the benchmarks is beyond the scope of this paper, both benchmarks assess the IO performance of a commercial storage product, which is generally available (or will be within 90 days of result publication), and validate those results with a rigorous audit and peer review requirements. Neither benchmark imposes any architectural or implementation constraints on the TSC, other than requiring that data be persistent and that the tested product be generally appropriate for the modeled business environments (i.e., OLTP-like applications). Rather than requiring or favoring a particular implementation, it is the goal of SPC benchmarks to provide robust, verifiable, reproducible means to assess the relative strengths of differing design and configuration approaches. While the metrics and data reported by SPC-1 and SPC-1C differ slightly, they produce a common set of metrics and other reported data:

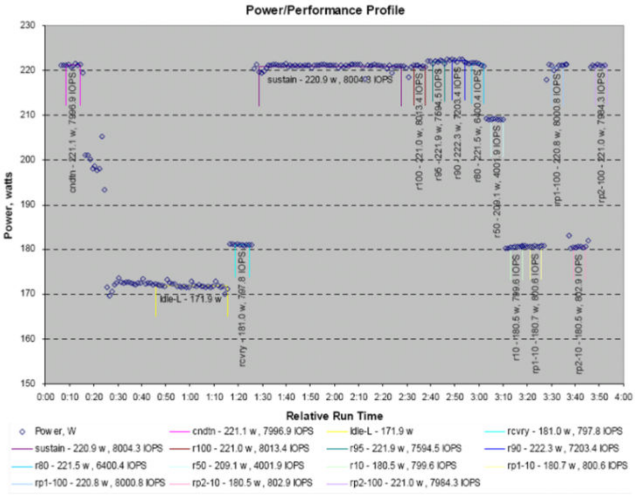

Figure 4: Sample Power/Performance Profile

- SPC-1 (SPC-1C) IOPS: the maximum IO throughput of the TSC subject to response time and operations mix requirements ;
- Total storage capacity and data protection level of TSC;
- Total price for TSC, with line-item detail (SPC-1 includes an explicit Price-performance metric)

4.2 Foundational Workloads with Appropriate Enhancement

As optional extensions, both SPC-1C/E and SPC-1/E serve as supplements to their underlying benchmarks (SPC-1C and SPC-1, respectively). That means that the configuration requirements, measurement criteria, timing constraints, audit protocols, and disclosure requirements defined for the base benchmarks remain in place. The underlying performance tests for SPC-1C/E and SPC-1/E are identical to those used in their parent benchmarks. Further, the performance metrics that were defined for the parent benchmarks are still reported in an energy result, and provide the important linkage between performance and power consumption.

The energy extensions focus on defining appropriate enhancements to these core benchmark components. Each energy extension defines:

- Measurement methodologies for power consumption (e.g., approved power meters, required accuracy and tolerance);
- Additional disclosure requirements specific to measuring energy consumption (e.g., input power characteristics)
- Appropriate metrics and expanded disclosure requirements (see Section 4.3);

4.3 Additional Testing Requirements

The energy extensions add an additional test phase at the start of the benchmark execution to characterize the power consumption of an idle storage system. Minimally, this includes:

- A conditioning phase of 10 minutes, with an IO load equal to 10% of the peak load reported during the benchmark execution;
- A stabilization Phase of up to 10 minutes. No work is applied by the workload generator;
- Phases Idle-0, Idle-1, … Idle-(L-1), optional, progressively "deeper" idle states, each with a duration of at least 10 minutes. No work is applied by the workload generator. While transitions between states may be scheduled, no operator intervention is allowed during the test;

- Phase Idle-L, the deepest idle state, with a duration of at least 30 minutes. No work is applied by the workload generator;
- A recovery Phase, with a duration of 10 minutes. The workload generator applies 10% of the peak IO load reported during the benchmark execution.

4.4 Energy Metrics

Energy measurements are taken during each phase of the benchmark execution, Since end-users are likely to have a broad range of storage duty cycles and differing raw energy costs, SPC-1/E and SPC-1C/E provide a two dimensional matrix that defines three usage levels, based on differing levels of IO load presented during the benchmark run:

- Idle: Power consumption of the SUT during idle test phase;
- Moderate: Power consumption of the SUT while processing 50% of the peak IOPs reported during the performance test;
- Heavy: Power consumption of the SUT while processing 80% of the peak IOPs rate reported during the performance test.

The measured energy consumption for these three load levels are combined to illustrate three common duty cycles for the SUT from a lightly loaded system which is predominantly idle to a heavily utilized system that is never completely idle.

The disclosure required with each result uses these power and performance metrics to provide a number of aggregated metrics:

- Nominal Power and Traffic: a weighted average of the power consumption or IOP rates at the differing load levels, weighted according to the daily usage pattern for the three defined duty cycles, showing the power usage and IO completion estimates for a 24-hour day;
- Nominal IOPS/W: the ratio of Nominal Traffic to Nominal Power, for each of the defined duty cycles;
- Composite Metrics: the average for Nominal Power, Nominal Traffic and Nominal IOPS/W across the three duty cycles;
- Annualized Energy Use (kWH/yr): an annualized figure, estimating the power consumption of the SUT;
- Annualized Energy Cost ($/yr): an estimated annual energy cost, based on a standardized price of $0.12/kWH).

Table 1: Reported Data Summary

	HOURS PER DAY			**NOMINAL METRICS**		
Daily Usage	Idle	Moderate (50% peak IOPS)	Heavy(80% peak IOPS)	Power (W)	Traffic (IOPS)	IOPS/W
Low	16	8	0	161.50	7498	46.43
Medium	6	14	4	162.64	19118.11	117.55
High	0	6	18	164.01	32601.39	198.77
Composite Metrics				162.72	19739.49	121.31
Annual Energy Use (kWH)				1425.41		
Energy Cost, $/kWH				$0.12		
Annual Energy Cost				$171.05		

A composite view of the benchmark result is provided by a unified graph that presents the IO traffic and the energy consumption for each of the defined phases in an SPC-1C/E benchmark execution (Figure 4).

4.5 Market Scope and Next Steps

The primary producers of SPC 1C/E and SPC-1/E results are storage vendors, though there may be additional interest in result production by end users or data center administrators as the need to quantify performance/energy trade-offs becomes more widespread. The testing to date (from three different vendors) has shown that the benchmark can differentiate a broad range of product design trade-offs in balancing performance, capacity and energy consumption [11]. Also, the SPC is working to expand is energy measurement methodologies to ever-larger configurations (current SPC-1/E results are bounded by the need to use a single power meter), and to a broader range of IO workloads (e.g., SPC-2 and its predominantly sequential workloads).

The SPC is also considering the development of end-user tools which would allow a potential customer to employ published audited results in estimating the costs that would result from their particular workload and duty cycle. The goal would be to address the challenge of applying a specific set of benchmark results to a more general user situation. By basing the sort of duty-cycle power calculations outlined above on published, audited results, the tool should provide both verifiability and flexibility, as the need to assess energy/performance trade-offs becomes commonplace.

5. UNITED STATES ENVIRONMENTAL PROTECTION AGENCY (ENERGY STAR)

The server benchmark community is rapidly expanding into areas of server efficiency from the more traditional realm of performance. This progression is natural and necessary as benchmark audiences become ever more aware of the energy impacts of the performance they seek in data center IT equipment. Given estimates that the cost of running data centers is increasing by as much as 20 percent a year, while overall IT spending is increasing by only 6 percent [3], efforts to do more with less will continue to fuel this trend. It is important that results are presented in the proper context as efficiency assessments become more uniformly integrated into performance-based benchmarks. A particular product deemed efficient when operating a database-oriented workload may be more or less efficient when operating a shared-file workload, for example. Because many performance benchmarks are tailored to a relatively specific type of application or usage case, there remains a need in the market for tools to give a *general* sense of efficiency in active modes while encompassing multiple end uses of the server.

The ENERGY STAR program for servers [1] [2] was finalized and established in May 2009. This initial set of requirements was focused on establishing an efficiency foundation, achieved through a mix of component efficiency, idle state power requirements, power management, and reporting criteria 9. Starting in fall 2009, EPA has begun efforts to build on this foundation through development of tools to evaluate general active mode efficiency of servers. EPA plans to work toward a tool that provides a broadly-applicable assessment of efficiency using inputs from multiple test modules addressing the different subsystems of the server. This assessment will be developed to allow for insight into how underlying activity contributes to overall results. It is hoped that this transparency, in turn, will encourage further development of the efficiency benchmark ecosystem by steering more users toward application-specific power-performance benchmarks to supplement the information acquired from EPA's tool. In this way, ENERGY STAR will build upon and compliment the continuing efforts of the server benchmark community to incorporate efficiency metrics in their own benchmarks. Though this section will discuss ENERGY STAR for servers, a similar development effort is underway for data center storage products.

5.1 Context and Approach

Context is critical to the proper understanding of any benchmark results. Figure 5 presents a hierarchy of benchmark approaches that range from general to specific, trading an increase in accuracy as the approaches progress downward with increased resources required to generate the data.

Results from a first-order benchmark would best be understood as broadly applicable to different systems and usage types, but not an appropriate source from which to develop expected efficiency of specific applications; results from an "application proxy" benchmark would provide an excellent understanding of efficiency from a single-use server, but the results would not be applicable once the application load was modified.

EPA intends to provide a first-order evaluation of efficiency. The ENERGY STAR program as a whole is structured with the goals of covering the widest range of products within the scope of each product program. In the context of evaluating server active mode efficiency, this structure is most appropriately channeled through tools that can be applied across the range of products covered by the program. Evaluating servers at this level will serve both users interested in the general efficiency and those with the resources to do more application-based, specific testing. The former group will be able to use ENERGY STAR evaluations as a basis for efficiency-based purchase decisions; the latter will be able to use the general active efficiency information as a first step in identifying the products they wish to test further using their applications or related benchmarks.

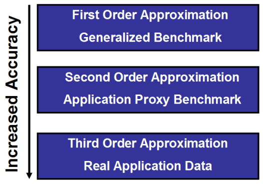

Figure 5: Hierarchy of Benchmark Approaches

5.2 ENERGY STAR Testing Process

The existing ENERGY STAR specification includes a set of testing conditions intended to ensure comparable results across the program. Given that the audience for ENERGY STAR results is not expected to review testing disclosures as would normally be anticipated for a benchmark's audience, stabilizing these conditions will be of primary concern as requirements incorporate active mode measurement. The current ENERGY STAR test procedure includes:

- *Power Analyzer Attributes* that include minimum capabilities and accuracy levels. The attributes list is based on existing and established requirements from other areas of the ENERGY STAR program, updated as relevant to server measurement based on industry feedback from development of the specification. Missing from the existing attributes are data interface requirements, which are viewed as important factors to ensure that future active mode testing can be an automated and uniform process; this was less of an integral consideration under the idle/non-active structure of the current specification.
- *Test Conditions* including specifications for input voltage and frequency, THD, temperature, and humidity. The voltage conditions reference established procedures for the program updated as applicable to the expected operational capabilities of server power supplies. The *2008 ASHRAE Environmental Guidelines for Datacom Equipment* and *ANSI ATIS-0600315-2007* were additionally referenced in establishing conditions.
- *Test Configuration* considerations are briefly addressed through specifications on network activity.

In the next version of this test procedure, conditions will be added such that tuning factors and software settings for the workload or benchmark are specified. Again, the disclosure model present for most established performance benchmarks is not anticipated as relevant for the program, and establishing clear and uniform guidance on tuning and reporting of results will help minimize misinterpretation.

5.3 Initial Focus

EPA will focus on activity integral to the server, consistent with the existing scope of the program, and evaluate systems based on power required to complete the pre-determined workloads in the test. EPA seeks to strengthen the conceptual connection between this productivity (including server performance) and efficiency, placing attention on the fact that performance comes at an energy cost; optimizing this balance is the key.

While EPA is in the early stages of developing tools to meet these goals, it is anticipated that CPU, memory, and integral storage activity will play a primary part in development of appropriate workloads. Secondary considerations could include network I/O activity. The long term intent is to evaluate the capability of the server to translate energy into its three fundamental tasks – operating on data, storing data for further operation, and communicating data post-operation – by incorporating processes that showcase this capability under stable, realistic conditions. Further, access to information across a utilization curve, similar to the reporting process established by SPEC, will be strongly considered for incorporation into program requirements. Lastly, idle power measurement will remain part of the program, either as a reporting requirement or limit, and will ideally be a component of the overall active mode efficiency evaluation.

5.4 Key Considerations in EPA's Approach

To meet EPA's goals of providing a broad view of efficiency for the range of products covered in the program, a number of key considerations are fundamental to creating active mode efficiency tools. The list that follows includes a combination of factors general to the ENERGY STAR process as well as specific considerations related to roll out of a software-based evaluation tool in an efficiency program.

Limited Barriers to Implementation: Encouraging broad testing of systems remains a fundamental component for the ENERGY STAR server program. Moving the market toward more efficient products can only happen if a wide variety of servers are compared against the program's stringent requirements. Architecture and operating system compatibility will be considered to ensure as broad participation as possible, in the near-term and into the future as the program continues to develop.

Broad Relevance: Workloads underlying the efficiency rating will be developed with broad applicability in mind. Specific end uses are too numerous to universally include in a benchmark tool. However, success of the workload underlying an evaluation tool will be judged on the basis of covering the types of operations expected under a broad variety of expected end-use scenarios, applicable to a widest range of products within the scope as possible.

Underlying Transparency: While the focus of a first order tool will be general, EPA seeks to encourage further investigation of product capabilities by providing insight into the underlying activity forming the efficiency rating. An example of this transparency could take the form of presenting power data across the utilization curve, similar to the approach set forth by SPEC in its Power and Performance Methodology [14].

Technology-neutral/architecture-agnostic: While technologies reasonably differ in capability under actual use, EPA seeks to incorporate active mode efficiency tools that do not introduce distortions of these differences, creating advantages/disadvantages within the efficiency assessment where they would not otherwise exist.

Standardization: The program could present an opportunity for the industry to work with EPA to establish baseline assumptions and testing conditions. Moving toward such standardization would create the foundation for valid and trusted comparisons of efficient performance. Creating this standardized environment will rely on a combination of factors:

- automation in active efficiency rating tools, reducing run-to-run variability;
- established testing and reporting conditions in ENERGY STAR specification testing methodologies;
- providing reasonable context to acquired results through reporting structures.

The EPA has worked with the data center industry and research community to set the stage for this work in the 2006 release of an initial *Server Energy Measurement Protocol* [5] and in the 2009 release of the *ENERGY STAR Test Procedure for Determining the Power Use of Computer Servers at Idle and Full Load*, as Appendix A to the ENERGY STAR specification for Computer Servers 9. Additional sources available to meet this consideration include SPEC's existing *Power and Performance Methodology* and *Power and Temperature Daemon*.

5.5 Developing a Road Map to Established Efficiency Ratings

The considerations of the previous section are, in part, tradeoffs with available time and resources to develop an appropriate tool and update the ENERGY STAR specifications. The next version of the ENERGY STAR specification for servers will accordingly set forth a roadmap for the program to move further toward an ideal solution for the market, industry, and data center community to acquire the tools to compare products based on efficiency considerations. The specification will serve to communicate EPA's long term goals for the program, including broad coverage of the server market, requirements and evaluation tools that apply to the complete scope of the program, and ensuring fair and realistic comparison structures.

Through the closing months of 2009 and into 2010, EPA will work with its stakeholder participants in development of the second version of ENERGY STAR server requirements. Work has begun with the SPECpower committee to discuss creation of an active mode rating tool meeting the guidelines presented, as well as a parallel effort to update and "roll over" existing ENERGY STAR companion criteria from the existing program.

6. BENCHMARK COMPARISON

This section compares the different approaches along various dimensions: Hardware components, workload and type of application, metric attributes and accuracy and calibration requirements. The benchmarks covered in this comparison are those that were available for energy comparisons by year-end 2009 (TPC-C, TPC-E and TPC-H from TPC, SPECpower_ssj2009 and SPECweb2009 from SPEC, SPC-1/E and SPC-1C/E from SPC) and anticipated structures of EnergyStar for Servers from EPA.

6.1 Hardware Components

Software engineering refers to multi-tier architectures (a.k.a. n-tier architectures) as client-server environments in which the presentation, the application processing, and the data management are logically separate processes. For instance, applications using middleware software to service data requests between users and databases employ multi-tier architecture. The most widespread use of this format is the three-tier architecture, which we assume when analyzing the different energy benchmarks. In this architecture the user interface runs on a desktop or workstation (Tier 1), functional process logic typically runs on an application server (Tier 2), and a Relational Database Management System (RDBMS), running on a database server, contains the data storage logic (Tier 3).

Table 2: Measured Hardware Components

Measured Hardware Component		**TPC-C**	**TPC-E**	**TPC-H**	**SPEC Power _ssj 2008**	**SPEC web2009**	**SPC-1/E**	**SPC-1C/E**	**EnergyStar for Server[1]**
Tier 1	System enclosure[2]								
	Secondary storage[3]							✔	
	Storage controller							✔	
	NIC							✔	
Tier 2	Server enclosure[2]	✔	✔		✔	✔			✔
	Secondary storage[3]	✔	✔			✔		✔	✔
	Storage controller	✔	✔			✔		✔	
	Storage enclosure	✔	✔			✔		✔	
	NIC	✔	✔			✔		✔	
Tier 3	Server enclosure[2]	✔	✔	✔			✔		
	Secondary storage[3]	✔	✔	✔			✔	✔	
	Storage controller	✔	✔	✔			✔	✔	
	Storage enclosure[4]	✔	✔	✔			✔	✔	
	NIC	✔	✔	✔			✔	✔	
Network switch Tier 1 and 2		✔	✔			✔			
Network switch Tier 2 and 3		✔	✔						

Table 2 shows the hardware components that all eight benchmarks emphasize in their benchmark measurements. A check-mark for a specific component indicates that the benchmark handles the component as critical to the system and measures its energy consumption. A lack of a check-mark for a specific component does not, however, indicate that a particular component could not be part of the benchmark measurement. Rather, it indicates that it is not typically used in this benchmark implementation. The benchmarks discussed in this paper emphasize the impact of component efficiency in different ways, based in part on the multi-tier architecture. TPC benchmarks C and E test system performance on Tier 2 and Tier 3, while TPC-H focuses on Tier 3. SPECpower_ssj2008 and SPECweb2009 both emulate Tier 3 while measuring the power consumption of the Tier 2 systems, focusing on power consumption of the server enclosure. In addition to the server enclosure SPECweb2009 measures the power consumption of the secondary storage, storage controller, storage enclosure and NIC of Tier 2. SPC-1/E and SPC-1C/E measure all hardware necessary to implement and support Application Storage Units (ASUs). These ASU provide the persistent non-volatile storage accessed in the course of executing the benchmark. While this generally does not include the host system, embedded controllers would cause the host, including its processor and enclosure, to be included in the measurement. SPC-1/E results include an entire storage solution (shown in Table 2 as Tier 3), while SPC_1C/E results can focus on storage subsystems at any tier. Although the Energy Star column currently focuses only on areas listed for Tier 2, it should be noted that a separate Energy Star specification exists for single-user computers that would fit in Tier 1

---

[1] Tier 2 program for servers

[2] Includes board, fans, processor and DRAM

[3] Includes computer memory that can retain the stored information even when not powered, e.g., flash memory and hard disks

[4] Includes the storage enclosure (fan, caches, etc.) and secondary storage

Table 2 shows that the hardware range for which these consortia measure energy consumption varies from Tier 2 and Tier 3 (TPC) to only Tier 2 (SPEC) or focus on specific components of every tier (SPC). The reason for this can be found in the focus of the different consortia. TPC is focused on database performance, which is inherently storage, network and processor/memory-heavy. SPEC focuses more on operations internal to servers, therefore emphasizing Tier 2 (processor, memory, etc.), while SPC focuses on storage solutions.

6.2 Workload/Type of Application

Rather than defining synthetic workloads, many of today's industry standard benchmarks model a specific real world-environment (e.g. retailer, stock broker). To accomplish this goal these benchmarks utilize workloads that simulate specific computational examples. These modeled scenarios may in turn be applied to any industry exhibiting similar workload patterns. This real world context frames results to help the reader relate intuitively to the components of the benchmarks, though care must be taken on part of the reader to view the results of the benchmark as a relative evaluation of the server against its peers; actual application loads are too unique and numerous to be modeled in a single, representative workload. EPA's developing efforts for the ENERGY STAR program, to the contrary, are likely to be more synthetic in nature to allow the efficiency metric to encompass activities expected in different application types; results in this case will similarly need to be understood as representative, not absolutely predictive, of actual server efficiency. Table 3 gives an overview of the different Application types modeled by the different benchmarks.

Table 3: Workload/Application Type

Benchmark	Application	Workload description
TPC-C	3-Tier OLTP	OLTP system of an order-entry system
TPC-E	3-Tier OLTP	OLTP workload of a brokerage firm
TPC-H	DSS	Ad-hoc, decision support queries
SPECpower_ssj2009	Server side-JAVA	Measures performance of server side Java applications
SPEC-web2009	Web Server	Measures web server performance with simulated database tier
SPC-1/E	Simulated OLTP Database tier	Performance of a storage subsystem while performing the typical functions of a business critical application.
SPC-1C/E	Simulated OLTP Database tier	Performance of a storage subsystem while performing the typical functions of a business critical application.
Energy Star for Servers	TBD	Tools incorporated into future versions of the program will incorporate tests that mirror activity expected in multiple application scenarios.

TPC benchmarks model OLTP and decision support applications using up to 3-tier architectures. Compared to the SPEC and SPC benchmarks listed, the TPC benchmarks can be considered the closest in measuring the full software and hardware stack of customer systems. The SPEC benchmarks tend to be scoped to specific functions that are critical components of a total compute environment, although. SPEC has announced plans for benchmarks in virtualization and services oriented architecture applications that are likely to broaden that scope. SPC benchmarks concentrate on measuring the storage component of an otherwise simulated OLTP system.

6.3 Energy Metric Characteristics

Computer system benchmarks are generally comprised of definitions for workload (data and transactions/queries), execution rules, and a set of measured quantities Q={q_1,…,q_n}. The quantities can take any forms: elapsed time of transactions, number of executed queries, load, and system price to energy consumption are examples from the benchmarks covered in this paper. These quantities are combined into a set of metrics M={m_1,…,m_m}. Each metric m_i is a function of a subset of the quantities $m_i(q_1,…,q_n)$. This section focuses on the energy metric of the various benchmarks, as well as the underlying measurements and quantities that form the foundation for the primary metrics.

Table 4 displays a description of the metric and the units they measure. Note that although the workload of many benchmarks is measured in T (or Q) their workload can be very different. All benchmarks use the term energy, either in describing their metric or benchmarks. However, the unit they measure is Watt consumed during workload execution. TPC benchmarks report electricity [W] consumed per transaction while SPEC and SPC report work completed per electricity [Watt] consumed. In order to convert the electricity consumed into Energy($J = W * s$), in addition to the average electricity consumed, we need to know the length of the measurement interval. This is only available for TPC and SPC benchmarks. All TPC benchmarks report electricity consumed per transaction while SPEC and SPC report work completed per electricity consumed.

Table 4: Energy Metrics: What they Measure

Benchmark	Metric Description	Unit
TPC-C	Electricity consumed [W] per transaction [T] and per time unit [m]	$\left[\frac{W*m}{T}\right]$
TPC-E	Electricity consumed [W] per transaction [T] and per time unit [s]	$\left[\frac{W*s}{T}\right]$
TPC-H	Electricity consumed [W] per queries [Q] and per time unit [h]	$\left[\frac{W*h}{Q}\right]$
SPECpower_ssj2009	Ssj-ops [T] per electricity consumed [W]	$\left[\frac{T}{W}\right]$
SPECweb 2009	Web transactions [T] per electricity consumed [W]	$\left[\frac{T}{W}\right]$
SPC-1/E	Input/Output operations [T] per time unit [s] and per electricity consumed [W]	$\left[\frac{T}{s*W}\right]$
SPC-1C/E	Input/Output operations [T] per second [s] and per electricity consumed [W]	$\left[\frac{T}{s*W}\right]$
Energy Star for Server	TBD	TBD

Table 5 displays detailed characteristics of the energy metrics used by TPC, SPC, SPEC and under development by EPA. TPC and SPC consider the energy metrics as secondary to the primary function of evaluating performance. Overall ranking using TPC and SPC benchmarks is based on performance only, with energy consumption presented as a supplemental consideration. Vendors are free to add context to the performance results through comparison with energy measurements. SPEC and EPA consider their energy metrics primary for their benchmarks. Rankings in this structure are based on efficiency (or "Energy Performance"). Consequently, reporting the energy metric is optional in TPC and SPC benchmarks and mandatory in SPEC and EPA benchmarks. All benchmarks report energy or power consumption during various components of the workload. However, it is notable that SPEC and SPC report energy consumption on a stepped workload (idle, 10%, 20%,.., 100%) while TPC report energy consumption during full and idle load. One differentiating factor between ENERGY STAR and industry consortia is focus: TPC and SPC are adding efficiency measurements to performance evaluations. SPEC issued a standalone energy benchmark that built upon existing performance-oriented workloads. ENERGY STAR is converging on industry benchmarks from the opposite direction by adding a performance component into an efficiency structure as a way of evaluating efficient computational activity. SPC benchmarks are the only benchmarks of those discussed that unite performance, price and energy consumption by reporting an annual energy cost. TPC and SPC-1C allow for the reporting of energy consumption of individual components in addition to the energy consumed of the overall system.

Table 5: Energy Metric Characteristics

Energy Metric Characteristics The Energy Metric	TPC-C	TPC-E	TPC-H	SPEC Power p_ssj 2008	SPECWeb 2009	SPC-1/E	SPC-1C/E	Energy Star for Server[5]
Is mandatory (m) or optional (o)	o	o	o	m	m	o	o	m
Reports total energy consumed during benchmark operation	✔	✔	✔			✔	✔	
Reports energy consumption at peak performance levels	✔	✔	✔	✔	✔	✔	✔	✔
Reports energy consumption at idle state	✔	✔	✔	✔	✔	✔	✔	✔
Reports energy consumptions at intermediate performance levels				✔	✔	✔	✔	
Unite performance and energy	✔	✔	✔	✔	✔	✔	✔	
Unite performance, price and energy						✔	✔	
Reports yearly energy cost						✔	✔	
Allows reporting of energy consumption of individual components	✔	✔	✔				✔	

6.4 Accuracy and calibration requirements

In order to guarantee correctness of measurement results, each benchmark consortia defines its own list of accuracy and calibration requirements. Within each benchmark consortia the measurement methodology is applied consistently to all benchmarks. Table 6 shows a list of accuracy and calibration requirements.

Table 6: Accuracy and calibration requirements

Requirement	TPC	SPEC	SPC	Energy Star for Server
Power Analyzer Accuracy [%]	2	2	2	2
Power Analyzer certification	NIST[6]	NIST[5]	NIST[5]	NIST[5]
Power Analyzer calibration	yearly	yearly	yearly	yearly
Ambient temperature [°C]	22-23	22-23	logged	18-27
Low end dew point [°C]	none	none	none	5.5
High end dew point [°C]	none	none	none	15
Atmospheric pressure [atm]	<1.1	none	none	none
Result certification	auditor	peer	auditor peer	self/audit

All benchmarks mandate calibration of power meters by a standard; examples include NIST in the United States, with equivalents existing in other countries. The meter must have been calibrated within the past year to ensure accuracy and fair comparison of benchmark results and the power analyzer needs to have a minimum accuracy of 2 percent (except for SPC benchmarks). Environmental conditions during the test are also specified in many cases, reflecting the impact the data center environment can have on performance of equipment. TPC, SPEC and EPA define an acceptable ambient temperature range for the power measurement. TPC and SPEC require the ambient temperature to be between 22 and 23°C while EPA requires the ambient temperature to be between 18 and 27°C. SPC does not require the ambient temperature to be in any specific range. However, it requires the ambient temperature to be reported. None of the industry standard organizations require any low or high end dew point, whereas EPA requires a low or high end dew point of 5.5°C and 10.5°C. Only TPC includes an atmospheric pressure requirement (less than 1.1 percent). Benchmark results need to be certified under each benchmark process, though this process varies. TPC benchmarks are certified by an independent auditor, SPEC benchmarks are certified by SPEC peers and SPC certifies benchmarks itself. EPA has included a combination of peer-verification, third-party auditing, and market assessment as parts of the ENERGY STAR program's certification efforts.

In respect to the above accuracy and calibration requirements, TPC and SPEC follow similar guidelines. SPC seems to have the least strict requirements while EPA imposes the strictest rules. While the metrics and structure discussed in previous sections are unique to each benchmark and are not easily compared, the accuracy, calibration, and test conditions noted above are one aspect of energy benchmark evaluation that could be more standardized between benchmarks. Efforts in this regard could result in less variation, test to test, and a point of trust on the part of the benchmark's audience.

7. CONCLUSION

This paper summarizes the responses by major industry-standard benchmark consortia and the US EPA to the industry demand for standardized energy consumption measurement. It highlights the unique values that each metric provides and associated measurement methodologies and requirements. Energy benchmarks, and their continued refinement, will help accelerate the development of techniques and technologies to conserve energy by providing a structure upon which these features can demonstrate energy-efficient performance, helping customers to make informed, energy conscious decisions. EPA's current and future initiatives are expected to both expand upon and further foster the efforts set in motion by these energy benchmarks. Continued development of tools that show the connection between server performance and the energy required to feed that performance will be critical to ensuring that the data center industry serves the public good – minimizing energy consumption while maintaining the broad set of services expected by a data-hungry public.

8. ACKNOWLEDGMENTS

The authors acknowledge the contributions of the members of the TPC, SPEC, SPC and EPA in developing energy standards.

9. REFERENCES

[1] ENERGY STAR Program Requirements for Computer Servers (2009)
[2] Fanara A.,Haines E., Howard A. "The State of Energy and Performance Benchmarking for Enterprise Servers" TPCTC 2009 52-66
[3] Forrest, W., Kaplan, J., and Kindler, N.: "Data Centers: How to cut carbon emissions and costs." McKinsey on Business Technology 14, 6 (2008)
[4] Hogan, T. "Overview of TPC Benchmark E: The Next Generation of OLTP Benchmarks" TPCTC (LNCS Vol.) 2009 84-98
[5] Koomey, J., Belady, C., Wong, C., Snevely, R., Nordman, B., Hunter, E., Lange, K., Tipley, R., Darnell, G., Accapadi, M., Rumsey, P., Kelley, B., Tschudi, W., Moss, D., Greco, R., Brill, K.: Server Energy Measurement Protocol. Analytics Press: Oakland, CA: (2006)
[6] NIST http://www.nist.gov/index.html
[7] Poess, M. and Floyd, C., "New TPC Benchmarks for Decision Support and Web Commerce". ACM SIGMOD RECORD, Vol 29, No 4 (Dec 2000)
[8] Poess, M. and Nambiar, R.: A Power Consumption Analysis of Decision Support Systems. WOSP/SIPEW:147-152 (January 2010)
[9] Poess, M. and Nambiar, R.: Energy cost, the key challenge of today's data centers: a power consumption analysis of TPC-C results. PVLDB 1(2)
[10] SPEC Benchmark Methodology v1.1.1 http://www.spec.org/ power_ssj2008/docs/SPECpower-Methodology.pdf
[11] SPC results http://www.stgorageperformance.org/results
[12] SPC Specifications:http://www.storageperformance.org/specs
[13] SPEC benchmarks: http://www.spec.org/benchmarks.html
[14] SPEC press release http://www.spec.org/gwpg/publish/power_announce.html
[15] TPC Benchmarks http://www.tpc.org/information/benchmarks.asp
[16] TPC Energy Specification Version http://www.tpc.org/tpc_energy/

---

[5] Tier 2 program for servers

[6] NIST for US or equivalent in other countries [6]