



Power Benchmarking: A New Methodology for Analyzing Performance by Applying Energy Efficiency Metrics

June 2, 2006

Elisabeth Stahl
Industry Solutions and Proof of Concept Centers
IBM Systems and Technology Group
estahl@us.ibm.com

Executive Overview

Historically, power has been a commodity with little interest from the IT community. As the price of energy increases and systems grow beyond the capacity of current facilities to supply their power and cooling needs, energy efficiency for data centers has now taken center stage along with traditional areas of interest such as performance, reliability, and availability. Concerns about server energy consumption are growing among senior IT managers because of rising electric rates, higher oil prices, technology trends that result in higher power densities, and hefty increases in new server deployments over the past few years. Clients and manufacturers are now talking about systems that achieve great performance and “eco efficiency” by reducing resource consumption.

This paper will discuss the application of power metrics to legacy systems performance measurements. It will define power benchmarking, discuss the history of server power consumption, review energy efficiency measurement considerations for systems, highlight new metrics in the industry, demonstrate the application of this methodology to an industry standard benchmark, and conclude with recommendations for the future.

Introduction – Why Power Benchmarking?

"Even without factoring in the need for additional systems, due to need for faster processors and the reality of higher energy prices/watt, Robert Francis Group believes power will be the number one issue for most large company IT executives to address in the next 2-4 years. Ignoring this issue will NOT be an option. Power consideration must be incorporated into data center planning. This includes close cooperation with facilities management, intelligent reduction of unneeded applications, systems, and infrastructure, and re-evaluating system design with an emphasis on power efficiency as a fundamental design criterion". (<http://www.ibm.com/systems/bladecenter/powerconfig/>)

Power Benchmarking is a new methodology in which power metrics are applied to legacy systems benchmarking measurements. It is a process for analyzing performance of a system by applying energy efficiency metrics.

IBM and other vendor strategies are to use power benchmarking to show leadership in both performance and power efficiency. The goal is to demonstrate how well the system being measured makes use of the available power to deliver computational performance. In the near future, IT providers will plan, execute and publish power-performance proof points using industry standard and ISV benchmarks in order to demonstrate advantage.

This paper will discuss the application of power metrics to legacy systems performance measurements. It will review energy efficiency measurement considerations for systems and highlight an example of applying this methodology to an industry standard benchmark. It will conclude with recommendations on future steps for Power Benchmarking.

History of Server Power Consumption – The “Greening” of IT

Historically, power has been a commodity with little interest from the IT community. With increased demands on power grids and municipal power systems, power can no longer be ignored. As the price of energy increases and companies want to be more environmentally conscious, energy efficiency for data centers has now taken center stage. Technology has delivered the dramatic strides in computational performance needed to keep pace with the growth in demand, but the new technology increases power consumption and packs it into a much smaller space than before, making it much more expensive to provide the required power feeds and cooling. A paradigm shift is evident. Even recently, processing capacity was the goal and the energy to run it was the commodity. Now processing capacity is the commodity, and the power to run it is the focus.

The environmental focus or “greening” of IT had been more significant in Europe, in part because energy costs in the EU are approximately twice the costs in the United States. Therefore, Europe was an early adopter of this energy efficiency focus. But recently the United States has joined in the interest in energy efficiency for IT, as evidenced by attention from the EPA and a recent Wall Street Power Efficiency Industry Roundtable. The mission of this roundtable was to establish joint technology provider and customer priorities and deliverables for introduction of more efficient technologies for use in corporate data centers.

In many data centers over the last few years, performance has doubled while electricity consumption has also almost doubled. The increase in system density has led to even more dramatic increases in the power expended per square foot and created hot spots inside installations. Studies have shown that if server power consumption grows 20 percent per year, the four year cost of a server's electricity bill could be larger than the initial price of a typical low-end server, ignoring the capital costs that may be required to provide additional power and cooling. If power consumption grows at 50 percent per year, power costs by the end of the decade could dwarf server prices. In fact, it is estimated that each year companies may pay up to a third of the purchase price of a system on powering, cooling, and housing that system. Saving power is key to lowering total cost of ownership (TCO). Energy costs are a top priority in the minds of IT leaders today.

IBM policy is to ensure the responsible use of energy throughout its businesses, including conserving energy and improving energy efficiency. Energy conservation makes both good environmental sense and good business sense. It reduces direct expense, saves resources and lowers emissions of greenhouse gases. It is IBM's goal to sustain leadership in energy conservation and management and to minimize the impact of its use of energy on the environment. IBM is delivering new technologies that allow our systems to use less power, generate less heat, and use less energy to cool the system. With the strong IBM commitment to climate protection, high “eco efficient” solutions are being provided.

Power Metric Considerations

In order to implement a methodology for applying power metrics to industry standard benchmarks, the metric for the power must be considered carefully.

Power metrics are needed to anticipate and implement power supply and cooling requirements in datacenters, determine operational costs and balance performance with energy usage. Metrics for power consumption of different components can help build more balanced systems. Metrics provide the necessary data for high level management software, operating system scheduling decisions and power profiles for changes to compilers and operating systems. Metrics can highlight hardware and help sell power management solutions.

Server power calculators are the newest tool advertised to estimate power used. Examples include the IBM BladeCenter® Power Calculator (<http://www.ibm.com/systems/bladecenter/powerconfig/>) and Sun's power calculator for its latest server systems (<http://www.sun.com/servers/coolthreads/t2000/calc/>). But simple power calculators, while helpful, are simplifications and approximations based on limited data and calculations. Actual practice is more complex, and the complexity begins with deciding how to calculate power or what to measure.

For calculating Power/Space, Watts/Area could be considered. But is data center area, rack footprint area, or some other area measured? Maybe not just the area but the volume is also important to the data center's costs. The server box volume could be used in a formula for Watts/Volume. Or rack could be a standard that is used in a data center. Watts/Rack could be calculated based on standard and custom rack measurements.

For calculating Performance/Power, Server Application Performance per Watt can be calculated. Power performance can be measured at the data center, system, server, or component level. Components of interest could be CPU, memory, I/O, or many others. Nameplate power, which can be the potential maximum, can be found on the system and could be used. Actual power can be measured or an average of 50-70% of peak estimate can be used. Some systems measure power internally or measuring is performed at the plug. Power consumption can be measured at different system performance utilization and workload levels. An OLTP workload may create different power levels than a Business Intelligence (BI) workload or a high performance computing (HPC) workload. Another consideration is that comparing two systems that do not have the same level of RAS features may not be useful. A system with redundant power supplies for example would consume more power than one without.

All of these considerations must be reviewed and analyzed when deciding which power metrics to use and how to calculate power performance.

Applying Power Metrics to Benchmarks

Types of Benchmarks

In order to understand how Power Benchmarking works, it is imperative to understand the different types of benchmarks. Three types of benchmarks exist: Industry Standard Benchmarks, ISV Benchmarks and Customer Benchmarks.

Industry Standard Benchmarks are created by an industry consortium and results are certified by independent parties according to strict rules. These benchmarks utilize one or more software programs which stress a system to provide comparative measures of a specific type of performance such as:

- Processor and memory (SPECcpu)
- OLTP throughput (TPC-C)
- Java™ performance (SPECjbb)
- Web serving (SPECweb)
- HPC performance (LINPACK (TOP500), Stream)
- Business Intelligence (TPC-H)

ISV Benchmarks utilize a software program where the software is the actual ISV application such as SAP, Oracle Applications Standard Benchmark and Lotus® NotesBench. These benchmarks are also used for sizing, and each ISV benchmark has different rules.

Customer Benchmarks are important where there exists demand to prove a 'capability' to a specific customer. Information is needed outside of available data such as Industry Standard Benchmarks or ISV Benchmarks, and the client sets criteria against their own data and workload.

Power metrics can be applied to any of these types of benchmark results based on the energy efficiency of the system. Many different formulas could be used based on the power metric considerations described in the metrics section above.

Sun's SWaP Metric

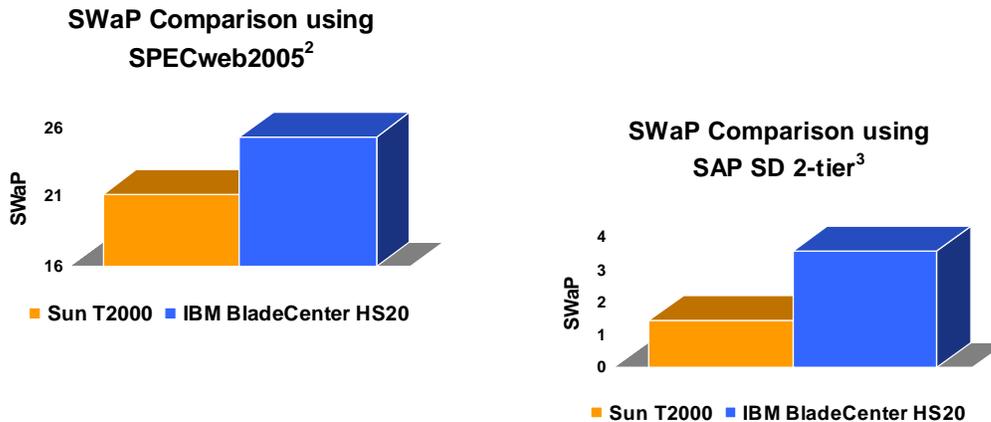
As an example, Sun Microsystems has developed a metric called SWaP which uses space, watts, and performance (<http://www.sun.com/servers/coolthreads/swap/>) as noted in Figure 1. Traditional metrics are good for calculating throughput, but they do not consider these new power and space demands in the equation.

$$\text{SWaP} = \frac{\text{Performance}}{\text{Space} \times \text{Power Consumption}}$$

Figure 1. SWaP Metric

In this equation, performance is derived using industry standard benchmarks. Space is calculated by measuring the height of the server in rack units (RUs). Data center racks are filled with servers, switches,

communication equipment, storage arrays, wireless routers, WAN switches, backup power supplies and other equipment. These devices compete for available space and contribute to the cost of powering and cooling the data center. Finally, power is computed by determining the watts consumed by the system, using data from actual benchmark runs or vendor site planning guides. A power meter can be used that records the watts used by the system during the benchmark run. To avoid inaccurate measurements, it is important to take a sequence of power measurements during the run and average them to get a representative power value. An example of SWaP¹ computations is shown in Figure 2.



1 – SWaP is defined as performance (of an industry standard benchmark) divided by space (measured in U of real estate) times watts (measured in power consumption); reference: <http://www.sun.com/emrkt/itanic/>

2 - Source: www.spec.org; All results as of 01/10/06; Sun results based on: <http://www.sun.com/servers/coolthreads/t1000/benchmarks.jsp#1>; BladeCenter HS20 configuration: 3.8GHz Intel Xeon with 2 cores, 8GB memory, 36GB drive, Fibre Channel, SUSE Linux Enterprise Server 9.2 SP1; 330 watts estimated average wattage; SPECweb2005 measured results = 4177;

Sun Fire T2000 configuration: UltraSPARC T1 processor with 8 cores and 32GB of memory; SPECweb2005 measured results = 14001

3 - Source: www.sap.com/solutions/benchmark/index.epx; All results as of 01/10/06; Sun results based on: <http://www.sun.com/servers/coolthreads/t1000/benchmarks.jsp#1>; BladeCenter HS20 configuration: 3.6GHz Intel Xeon with 2 cores, 8GB memory, 330 watts estimated average wattage; Sun Fire T2000 configuration: UltraSPARC T1 processor with 8 cores and 32GB of memory @ 330 watts; Based on two tier SAP SD Standard Application Benchmark result 2-core IBM BladeCenter HS20 (3.6GHz) (588 benchmark users, 1.99 second average response time, 177,000 dialog steps/hour, certification number 2005015) running DB2 UDB 8.2. on Windows Server 2003 Enterprise Edition, SAP R/3 Enterprise Rel. 4.7 as compared with the 8-core 1.2GHz UltraSPARC T1 Sun Fire T2000 SAP SD Standard Application Benchmark result (950 benchmark users, 1.91 second average response time, TBD dialog steps/hour, certification number 2005047) running MaxDB 7.5, Solaris 10, SAP R/3 Enterprise Rel 5.0.

Figure 2. SWaP Metric Calculations

New SPECpower Benchmark

The Standard Performance Evaluation Corporation (SPEC) is currently working on developing an industry standard metric for power (<http://www.spec.org/specpower/>). The SPEC Power-Performance committee has begun development of the first generation SPEC benchmark for evaluating the energy efficiency for server class computers. The SPECpower benchmark goal is to provide a means to fairly and consistently report system energy use under various usage levels. The goal of the Power-Performance committee is to utilize proven SPEC server benchmarks in order to provide workloads for which energy use can be measured. The drive to create the SPECpower benchmark comes from the recognition that the IT industry, computer manufacturers, and governments are increasingly concerned with energy use of servers. Development of the SPECpower benchmark provides a means to measure energy use in conjunction with performance metrics. The goals of this planned benchmark include:

- System utilization will vary in order to provide multiple usage loads and energy measurement points.
- Energy measurements are made at the AC input to the system under test.
- The benchmark as initially planned addresses a subset of server workloads.

Current SPEC member companies committed to developing the new power-performance measurement standard include AMD, Dell, Hewlett-Packard, Intel, IBM, and Sun Microsystems. SPEC is also getting direct input from organizations such as the EPA's Energy Star Program. Energy Star is a United States government-backed program helping businesses and individuals protect the environment through superior energy efficiency.

Implementing Performance Per Watt

For now, many vendors and clients are simply using performance per watt. As a basic example, select an industry standard benchmark where results exist for the systems to be compared. Divide those performance metrics by the watts used to get performance per watt. This calculation could show that systems with similar performance may be vastly different when power metrics are applied to the equation.

As mentioned in the power metrics section above, there is great complexity when using and applying these metrics. It is imperative that the workload used in the benchmark matches the production type of workload. Also, the benchmark must fit the utilization model; if the benchmark is a varying load benchmark and the system runs in production at high utilization most of the time, then a power efficiency number from this benchmark may not be very useful. In this case, a peak performance/peak power metric would be better than an average performance/average watt metric. Also note that if a vendor publishes wattage numbers at different load levels along with performance at those load levels, it does not mean necessarily that the performance will be acceptable on a varying workload type of machine. If the performance and wattage was measured at each load level individually, it does not take into account the overhead of switching power states. In addition, note that as many performance benchmarks work at maximum utilization and most systems consume at least half the power in idle mode, the real life power efficiency with typical low to medium workloads is expected to be significantly lower.

Conclusion

The greening of the IT industry and ensuring that the IT community becomes more eco efficient has changed the entire landscape of systems performance. A server's overall efficiency calculation now must include power consumption. Vendors and the industry as a whole must respond by providing relevant and applicable data about the systems that are coming to market and standard benchmarks and metrics for power-performance and power efficiency.

Power Benchmarking will be one of the most important issues for IT leaders in the coming years. Balancing performance with power consumption is a top priority now that interest in energy-efficiency is growing in the IT industry.

References

SPEC to develop energy metrics for servers

<http://www.spec.org/specpower/pressrelease.html>

Power could cost more than servers, Google warns

http://news.com.com/Power+could+cost+more+than+servers,+Google+warns/2100-1010_3-5988090.html

Group Creating Tools to Rate Server Power Use

<http://www.eweek.com/article2/0,1895,1949457,00.asp>

EPA Energy Star Program

<http://www.energystar.gov/>

Sun Introduces World's First Metric for Server Efficiency and the Internet of Tomorrow

<http://www.sun.com/smi/Press/sunflash/2005-12/sunflash.20051206.5.xml>

New Eco-Friendly Sun Desktop Consumes Fewer Watts Than a Night Light

<http://www.sun.com/smi/Press/sunflash/2006-04/sunflash.20060412.4.xml>

Sun Microsystems Received 2006 Environment Prize for the UltraSPARC T1 Processor

<http://www.sun.com/smi/Press/sunflash/2006-04/sunflash.20060413.1.xml>

Sun, Industry Leaders, EPA Build Consensus for Important Energy Efficiency Metric for Servers

<http://www.sun.com/smi/Press/sunflash/2006-04/sunflash.20060414.1.xml>

Solaris 10 Operating System and Sun Java Platform, Standard Edition 5.0 Set Multiple World Records on Server-Side Java Benchmark Across Sun Fire Product Line

<http://www.sun.com/smi/Press/sunflash/2006-05/sunflash.20060517.5.xml>

Thanks to Freeman Rawson, IBM Research Power-Aware Systems, Matt Accapadi, AIX Performance, Silvio Weeren, Program Manager Environmentally Conscious Products, and Alan Adamson, Java Performance - IBM OSG SPEC Representative.



© IBM Corporation 2006
IBM Corporation
Systems and Technology Group
Route 100
Somers, New York 10589

Produced in the United States of America
June 2006
All Rights Reserved

This document was developed for products and/or services offered in the United States. IBM may not offer the products, features, or services discussed in this document in other countries.

The information may be subject to change without notice. Consult your local IBM business contact for information on the products, features and services available in your area.

All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only.

IBM, the IBM logo, BladeCenter, Lotus are registered trademarks of International Business Machines Corporation in the United States or other countries or both. A full list of U.S. trademarks owned by IBM may be found at:
<http://www.ibm.com/legal/copytrade.shtml>.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. In the United States and/or other countries.

TPC-C and TPC-H are trademarks of the Transaction Performance Processing Council (TPPC).

SPECint, SPECfp, SPECjbb, SPECweb, SPECjAppServer, SPEC OMP, SPECviewperf, SPECapc, SPECchpc, SPECjvm, SPECmail, SPECimap and SPECsfs are trademarks of the Standard Performance Evaluation Corporation (SPEC).

Other company, product, and service names may be trademarks or service marks of others.

IBM hardware products are manufactured from new parts, or new and used parts. In some cases, the hardware product may not be new and may have been previously installed. Regardless, our warranty terms apply.

This equipment is subject to FCC rules. It will comply with the appropriate FCC rules before final delivery to the buyer.

Information concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of the non-IBM products should be addressed with those suppliers.

All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, to evaluate the performance of a system they are considering buying.

When referring to storage capacity, 1TB equals total GB divided by 1000; accessible capacity may be less.

The IBM home page on the Internet can be found at:
<http://www.ibm.com>.

The IBM System p home page on the Internet can be found at: <http://www.ibm.com/systems/p>.