SUMMARY OF SERVER ENERGY METRICS MEETING (V6)

27 March 2006

Jonathan G. Koomey (LBNL and Stanford), JGKoomey@stanford.edu

Our goal for the meeting was to initiate a process that will yield objective measurements of server energy use attached to all of the widely accepted performance benchmarks. Because arguments about performance benchmarks have continued since the beginning of the computer age, we opted not to create new benchmarks or to rank one metric as better than another one. Instead, we tried to identify the computing performance benchmarks most widely used by server purchasers, and then discussed the key components of a protocol to measure energy use while the performance benchmarks are being run. This process will be a success if server purchasers can soon rank servers by an objective measurement of energy use associated with all performance benchmarks that are appropriate for the intended use of the server.

The meeting began with opening statements by Andrew Fanara of EPA's ENERGY STAR program, describing EPA's interest in fostering industry dialogue on server energy use. Andrew emphasized that the EPA is not currently considering an ENERGY STAR label for servers.

I then summarized the goals of the meeting and how I expected the day to go.

There were three initial talks. Walter Bays and Robert Lane described the processes for developing the SPEC and TPC benchmarks, respectively. David Greenhill then described some of the general lessons learned in the process of creating the SWaP benchmark, and described four categories of applications that are commonly used (and that represent different market/user segments for benchmarks): web hosting, application serving, database serving, and high performance (scientific) computing.

Following the initial talks, we began a brainstorming session about major topics for the afternoon discussion. I proposed five possible topics:

A) Equipment boundary issues (memory, disks, other equipment, big vs. small systems)

B) Physical measurement issues (ambient conditions, energy vs. power, watts vs VA)

C) Real world vs. benchmarks (how well do benchmarks reflect real world performance)

D) End users—What benchmarks do customers use now? How do they use them? What would help them make better decisions on energy use?

E) Compliance/Auditing/Documentation/Replicability/Organization

In the initial discussion, a number of ideas came up. One participant emphasized the importance of establishing consistent time series for equipment, components, application loads, and performance metrics. Others raised the possibility of creating "load curves" and corresponding energy curves, at a minimum showing power drawn by servers at 0% load, 100%, and some points in between (later in the meeting it became clear that SPEC is already trying to create such curves, but it was not known when that information would become publicly available).

The group expressed its preference for focusing on the middle three topics (2, 3, 4) and so we chose those as the topics for the breakout sessions.

Over lunch, Bruce Shaw of AMD described the results of an AMD-funded Ziff-Davis survey of almost 1200 key players in the data center industry. The results were striking: 71% of the respondents identified information technology (IT) power, cooling, or both as the primary issue facing their data center, but 62% of the respondents reported that neither IT power consumption nor cooling concerns had affected server purchases in the preceding twelve months. Bruce summarized the results by saying in essence "Our customers have problems with data center power and cooling that they don't know how to fix. We need to help them."

After lunch we broke into three groups, focusing on physical measurement issues, applicability of benchmarks to real world applications, and end-user concerns. Those groups reported out later in the afternoon.

REPORT OF THE PHYSICAL MEASUREMENT ISSUES GROUP

Rob Snevely of Sun led the reporting out for this group.

Group attendees: Sun, Google, IBM, HP, Microsoft, APC, Others?

First, the group defined the scope of their investigations: They only focused on air-cooled servers drawing AC power.

Next, they determined that measurement conditions (e.g. temperature, relative humidity, air pressure/altitude) should follow ASHRAE Class I standards for data centers. Conditions for frequency, voltage, and total harmonic distortion (THD) of the input power should also be specified.

Then, the group described the output measurements to be presented, envisioning at least two levels of detail: a full suite of measurements of all sorts, and a summary metric in the form of a ratio of some level of service/performance divided by power or energy consumption.

The full suite of outputs of the measurements would include Volts, Amps kVA, watts, Btus, watt-hours over some period, power factor, and THD. Procedures for these measurements should follow the relevant ISO standards. Reported results should include all actual measurements, while separately noting the minimums, maximums, and averages over the testing interval. All of these data should be reported in the full disclosure portion of the "benchmark".

Measurements must be between the power outlet and the server, should assess total box power, and must use all power connections at the same time.

The maximum (rated) power of the installed power supplies in the tested configuration must be reported along with actual power measurements.

REPORT OF THE APPLICABILITY OF BENCHMARKS TO THE REAL WORLD GROUP

Christian Belady of HP led the reporting out for this group.

Group attendees: HP, Intel, AMD, LBNL, Cisco, Dell, Sun, Rumsey Engineers Others?

The group concluded that processor utilization is not a reliable measure of how much work is actually being done. Instead, the group strongly preferred some measure of actual workload for the task at hand (e.g. web pages served).

The group concluded that the measurements should be expressed as a graph of energy consumption or power (as a % of maximum) on the y-axis and a measurement of workload (as a % of full load) on the x-axis. Server customers should then be able to compare curves and assess which servers would be most energy efficient in their applications. The group also discussed the advantages and disadvantages of collapsing such a curve into a few points (e.g., 0% load, 50% load, 100% load, or energy measurements for each 10% load increment).

The group noted that not all performance benchmarks can be scaled as described in the previous paragraph.

The group concluded that more work is needed on which benchmarks customers actually use in practice and how well benchmarks reflect actual uses for the servers.

Finally, it became clear that at least two of the manufacturers at the meeting were already taking the approach advocated by this meeting (measuring power and energy use as a standard performance benchmark was being run) and that SPEC is working on creating the kinds of curves advocated by this group (although it is not clear when the SPEC work will become publicly available).

REPORT OF THE END-USER PANEL

Henry Wong of Intel led the reporting out for this group.

Group attendees: Intel, Sun, AT&T, PG&E, AMD, LBNL, Others?

I. Areas of concerns:

–End user education (particularly related to adoption of power management)

–User Friendliness

–Sub-monitoring

II. Recommendation:

–3-5 data points

–System Workload: Low ≠ 0 ;High≠ 100%

–System Workload ≠ CPU utilization

Presentation matrix for power measurements by market segment and workload

Market / Workload Condition	Datacenter	Web Services	HPC	Financial Services	Other
Low					
Med					
High					

Hypothetical Example Workload, Pwr, Perf. graph

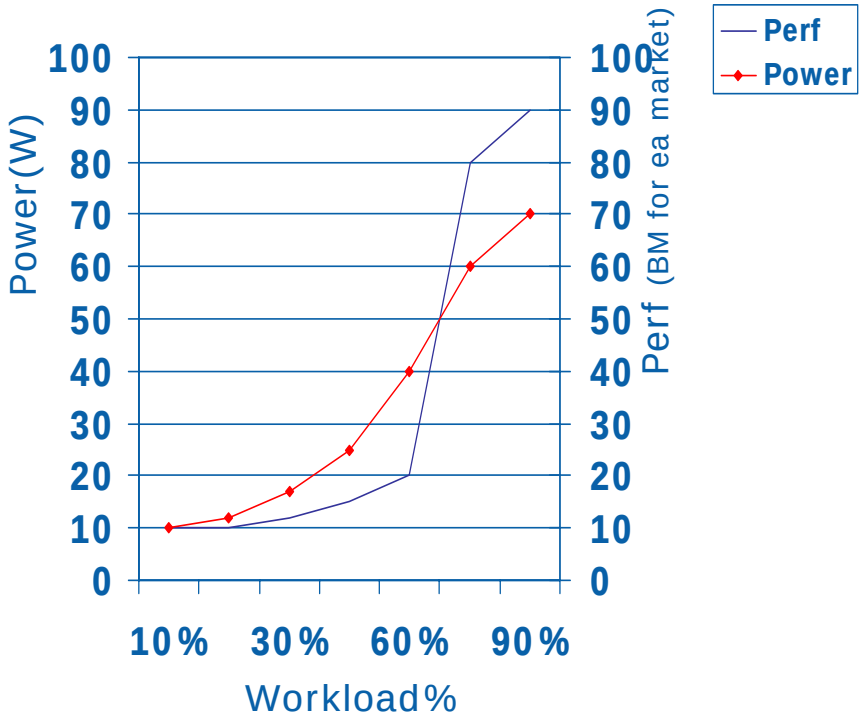

III. Next steps:

–Define Workload conditions and Benchmarks representing each market segment based on feedback from server purchasers, performance metrics organizations, manufacturers, and others.

–Collect sample data to demonstrate feasibility and usefulness of graph (~5 data points per category and condition);

–End user forum to interview, present, and confirm categorization/usefulness of the approach.

–Finalize categories

–Determine derivative metric (composite) to rate the efficiency that will allow a management measure of relative efficiency.

BRAINSTORMING ABOUT WHAT SHOULD BE IN A PROTOCOL

After the groups reported out, all participants brainstormed about what components should be part of a measurement protocol intended to attach energy measurements to existing performance benchmarks.

There was general agreement that we should focus on the simplest benchmarks for small servers (1U, 2U) as a starting point. Consult with users to determine which benchmarks are most widely used for that equipment.

Need to specify boundaries: to keep it simple, focus just on the server box itself (not equipment external to the box). Within the box, we need to have measurements and specifications of the major subsystems (power supplies, disks, memory, motherboard, etc).

Need to consult ISO and ASHRAE standards for measurement practices. Also look at testing protocols for ENERGY STAR computers, some of which may be applicable here. Finally, we need standardized definitions for key terms related to power consumption.

Need to consult with the Storage Performance Council, which has performance measurements for disks and disk arrays.

NEXT STEPS

The next step is for me to draft a version of the protocol for the simplest of 1U and 2U servers, based on the comments people made in the meeting. This will give participants in the meeting something to react to, which will surely result in additional constructive feedback. I expect to have a draft for the meeting participants by early May 2006.

I expect to use the draft protocol and example measurements to solicit feedback from server purchasers to make sure we are serving their needs. They are the ultimate customers for the metrics that will arise from this effort.

VOLUNTEERS TO HELP KOOMEY IN CRAFTING THE DRAFT PROTOCOL

Here's the list of people (in alphabetical order by last name) who volunteered to help me in crafting the draft protocol. I may have missed someone, so if you or someone you know should be on the list, please let me know. Please forgive me if you volunteered and I left you off this list—my memory is (alas) not perfect.

Matt Accapadi, IBM

Christian Belady, HP

Richard Greco, EYP Mission critical facilities

Ed Hunter, Sun Microsystems

Brent Kelley, AMD

Bruce Nordman LBNL

Peter Rumsey, Rumsey Engineers

Bill Tschudi, LBNL

Henry Wong, Intel.

OTHER DISCUSSIONS/ISSUES/UNRESOLVED QUESTIONS

It is still unclear whether a separate protocol will be needed for cluster servers, but my guess is that this will probably be required. I also suspect that the larger computers will need very different protocols from those established for simple 1U and 2U boxes.

The issue of creating consistent and comparable time series of energy and performance of product models over time will continue to come up, and we need to think about this as we move forward. Ideally customers would be able to shop for the most efficient servers at any particular time, and that manufacturers and researchers will be able to use the energy measurements over time to draw meaningful conclusions about the direction of server technology (at least with regards to energy use).

The issue of how reliability is measured has not really been addressed. It is important that any measures taken to reduce energy use not adversely affect reliability of servers. We suspect that many energy savings measures will reduce heat and enhance reliability, but there may be other ways to save energy that don't have that beneficial effect. It will be important in the medium term to address metrics on reliability for servers, but that shouldn't get in the way of moving forward on energy metrics at this point.

APPENDIX: COMPLETE LIST OF COMPANIES ATTENDING

AMD
ANCIS Inc.
Apple
APC
AT&T
Bite Communications
Cisco Systems
Dell
EYP Mission Critical Facilities
Hewlett Packard
Goldman Sachs
Google
IBM
Intel
Lawrence Berkeley National Laboratory
Microsoft
Natural Resources Defense Council
PG&E
Rumsey Engineering
Sun Microsystems
U.S. EPA

OTHERS TO INVITE TO FOLLOW-ON MEETINGS

More end-users

More utilities

Manufacturers of power supplies and disk drives

Software companies

Measurement/instrumentation companies (Agilent, others?)