To: Katharine Osdoba, Osdoba.Katharine@epamail.epa.gov, Energy Star Program
From: ITI Workstation Workgroup for Energy Star '07, Henry ML Wong, Intel Corporation
Date: 9/19/07
RE: Workstation Energy Star '07 Draft 4 Feedback
*** The points and views indicated are those of the industry working group and may not solely represent the views of Intel Corporation.

The industry members consisting of representatives from Apple, AMD, Dell, HP, IBM, Intel, and Sun; discussed the limits and concerns with the workstation proposals. The team included significant work from EPA's consultant in this category, Tom Bolioli from Terra Novum. Tom Bolioli's support was critical due to the insight and impartiality maintained in the discussions and proposal. We'd also like to re-iterate, the learning acquired in the process. The nuances of testing, configurations, and architectural sensitivities has provided key learning's in our progress to a Tier 2 criteria. The details also provided several of the members some of the key considerations in both the type and quantity of information necessary to make a uniform decision.

Unfortunately, we also encountered a number of miscommunications on alternate formula options which delayed our response. The response below references the 8/26 draft that was provided by the EPA.

Outside of the formula and limits, the other feedback areas center around the testing issues uncovered during the testing and data collection. The non-formula feedback is:

- Testing should be allocated a level of variance that reflect the natural error that would occur during the Energy Star compliance testing. The minimum amount is 10% and is believed to be more significant as the power levels and configurations increase. This is consistent with the request on other categories. The 10% test allocation addresses:
    - Testing data aliasing due to the sampling and variability in the waveforms
    - Varying sub-components to represent the category of product with the same configuration.

- A recommended set of testing guideline would aid in providing consistent results. Intel has offered an example proposed.

- System suppliers should be allowed to provide their parameters for Linpack and SpecViewPerf that demonstrates the maximum capability of the system. Such flexibility must follow the guidelines put forth for Linpack and SpecViewPerf (via spec.org). Such parameters must be indicated prior to submitted for Energy Star consideration and allow for external duplication. In addition, the system energy star applicant would provide a reasonable justification for the settings, and the settings would represent the maximum performance conditions a user could observe in use of the system.

Formula and limits.
We recognize the significance of and the desire to address the concerns with a simple percentage in TEC budget. The workgroup has focused its activity on Formula B which constitutes as TEC budget based on 43% max_power and adjusted by a "k" factor.

Although alternatives were being investigated, we would like the EPA to consider :

- Either the adjustments to Formula B listed below or,
- Going back to Formula A, i.e.
    - TEC_budget = Slope * (Max_power[Linpack, Specviewperf] + (#HDD * 5W)) .
    - TEC = (0.7 * PIdle) + (0.2 * PSleep) + (0.1 * PStandby)
    - Slope ~ 35%

Areas of concern for formula B:

- Use of cores will become in appropriate as the industry transitions to larger number of cores (beyond dual) to increase the efficiency of accomplishing the tasks targeted for workstations.
- Some of the data represents historic configurations and may become an inaccurate measure of future platforms demanded by workstation customers. A 25 percentile based on some of these platforms may be inappropriate and would incentivize older or undesirable configurations.
- Given the lack of active_mode contributions as part of the efficiency metric, we need to adjust the metric to compensate for the limits to encourage more energy efficient workload and work-cycle based platforms. We believe this is necessary as a precursor to improved energy benchmarks for Tier 2, in 2009.
- Finally, as a baseline, we do not believe in using the formula for limits on low end. The number of system models listed on the low end unduly weights the percentage of product configurations expected to be desired in the workstation market. The recommendation (from some suppliers) included indicates a limit that drives for a 25th percentile in the high configurations, but, not put that same constraint on the baseline models. These baseline models are unlikely to represent the ongoing workstation demand. The current workstation demand can be inferred by the IDC Workstation market share data.

Feedback and proposal on Formula B:

Restate of formula B:

TEC= (0.6 * PIdle) + (0.3 * PSleep) + (0.1 * PStandby)

Limit: TEC_budget = (slope * Max_power_budget) + (intercept - k)

Slope= 0.43

Max_power_budget= Max_power (Linpack, SpecViewPerf) + (#HDD * 5W)

Intercept = -8

(k restatement not included to avoid confusion)

- Modify formula B by removing the cores from the k-factor to the following:
  - k = ((#cores * cpu_scale) + (#DIMMs * 0.25)) $^{1.6}$

  where cpu_scale = 5

- Min TEC limit of 100. (i.e. Limit is the larger of (Min_TEC_Limit, TEC_budget))
  - Per formula and scatter plot of TEC v. Maxpower, variability can be seen to expand at ≥ 300W max_power
  - At max_power ~ 300W, a TEC budget would be ~ 120, so 100 would be aggressive cut on the TEC limit for these baseline systems.
  - Would inherently encourage current consumers in increased capable platforms (targeted by the current market) to either get systems which meet the curve level or min level.
  - As a compromise to the increase number of workstation models, tighten the slope of scaled systems from 0.43 to 0.41.
  - Side benefit: encourages higher performance configurations to lower both active and inactive power levels, prior to Tier 2 benchmarks.
  - Acknowledges the definition of a minimum level Workstation
  - Acknowledges the fact that the current purchasing criteria for workstations are not based on power or the population represented in the charts. A baseline encourages the overall population to move to systems at this level.

This represents the consensus response. Individual companies may or may not fully endorse or recommend the items above.

Regards,

Henry ML Wong,
Intel Corporation
(408)765-5047

Industry Representatives:
Apple – Dave Cassano
AMD – Sanjiv Lakhanpal
Dell – Steve Donley, Robert White
HP – Scott Deuty, Louis Hobson
IBM – Joe Prisco
Intel Corporation- Henry ML Wong
Sun – Gary Jones, Lowell Sachs, Terry Whatley

Consultant: Terra Novum: Tom Bolioli

Attachment

Testing Guideline Recommendation
Suggested sequence for measurements:
1. Start with system in *SoftOff* state
    - Measure power for 5 min duration
    - Sampling period of 1 sec
    - Average power over 5 min to get Power values
2. Put into *Sleep* state (ACPI S3)
    - Measure power for 5 min duration
    - Sampling period of 1 sec
    - Average power over 5 min to get Power values
3. Turn on system into *idle* state
    - Measure power for 5 min duration
    - Sampling period of 1 sec
    - Average power over 5 min to get Power values

4. [*Power(SpecViewPerf, Linpack)]* measurements
    - Run Linpack and SpecViewPerf at the same time (need to explain how this is run – follow EPA directive)
    - Sampling period of 1 sec
    - Test duration: single run of the combination (Linpak + SpecViewPerf) workloads till both workloads finish.
    - Repeat (Linpak + SpecViewPerf) run ***3*** times with identical system setup and measure system power
    - If all ***3*** measurements fall within a +/-3% error range relative to the nominal measured maximum power, the measurement method is valid.
    - Use the maximum of the measurements to determine the value of [Power(SpecViewPerf, Linpack)] for calculating MaxPower.

----- end of testing guideline ----------