Mar 21, 2008

United States Environmental Protection Agency
Office of Air and Radiation
Energy Star® Program
Washington DC 20460

Dear Energy Star Program,

Intel appreciates the opportunity to comment on the Energy Star Version 5.0, Draft 1 spec proposal dated February 22nd, 2008. We look forward to an ongoing dialogue on the subject. In summary, our comments are aligned to the specific topical elements outlined by the Draft 1 proposal.


General Feedback ..... 1
0.0 Commitments ..... 2
1.0 Definitions ..... 2
3.0 Energy Efficiency and Power Management Criteria ..... 3
- Power Supply Unit (PSU) Efficiency Requirements ..... 3
- Table 2 – Capability Adders ..... 4
- Workstation levels ..... 4
- Desktop Derived Servers ..... 5
- Table 4 – Thin Client Efficiency Requirements ..... 5
- Table 5 – Ethernet Power Management Requirements ..... 5
- Table 5 – WOL Requirements ..... 6
- Table 5 – Network Connectivity ..... 6
- Table 5 – Wake Management ..... 6

4.0 Test Procedures ..... 7
- Table 6 – Test Procedures ..... 7

6.0 Future Specification Revisions ..... 7
7.0 Appendix A: Energy Star Test Procedure ..... 7


# General Feedback

- Intel requests that EPA re-enable the comments feature in the PDF documents for future draft releases. Disabling comments in PDF is unnecessary overhead and slows the document review process.

# 0.0 Commitments

Line 89 – 91: Based on the improvements made in PC boot times, Intel recommends EPA re-evaluate the 5 second logo display requirement which may create an undesirable user experience. The logo should be displayed in a manner such that it does not impact/slow the boot time of the machine.

# 1.0 Definitions

Line 301 – 303: Intel recommends more specificity around the definition of Thin Client to include the following qualifiers:

1) A Thin Client relies on centralized server resources for program execution
2) A Thin Client has no local rotational media devices (i.e.. HDD or ODD) and relies on a wired network connection to a server for all permanent user or application data storage.

Line 315 – 323: The bundling of Tablet PC's with Notebook PC's may create a conflict in subsequent requirements specific to the ability of a Tablet PC to run the EEPA Tool. Intel recommends Tablets be excluded from the v5 specification.

Line 345 – 354: Intel recommends EPA more closely align the low-power mode definitions to ACPI spec which applies to most PC's. The specific changes requested are:

1) Change the "Off" mode to apply to the S5 ACPI state only; Off carries an implication of full system reboot to resume functionality. S4 does not require this level of system reboot.
2) Change the "Sleep" mode to include both ACPI S3 and S4 states which are equally viable operating modes to meet this definition. In this case, a full system reboot is not required to restore normal working order. Intel continues to believe that referencing the low-power mode definitions in the ACPI specification provides clarity to the actual mode of operation being discussed (as opposed to EuP Lot 6). However, Intel also believes that the sleep state definition should not be specific enough as to dictate a specific solution. Intel recommends that the EPA define what operational or user experience (i.e.. maximum resume time to a usable system state) is needed to meet the sleep state definition, then let the OEM pick the appropriate sleep state implementation to meet those requirements. For example, if a given solution can meet the sleep state via an S4 requirement then this should be allowed (and would represent a better-than scenario for actual energy consumption).

Line 356-358: All other operational modes specifically reference the ACPI state that applies. For consistency, Intel recommends that the Idle state definition include the reference to ACPI state S0. In addition, Intel seeks more clarity around what "basic applications" means as referenced in the idle mode definition. Are these constrained to OS resident programs or would it include pre-loaded "trial" user applications as well or other?

Line 360-365: Semantic point – prior to using the acronym EEPA, the term should be defined first. Intel suggests either dropping the EEPA reference from section O, or preceding the Active state definition with both the EEPA (section S) and EEPA Tool (section T) definitions.

Line 395 – 398: EEPA Tool outputs – Why is there is no specific requirement for the EEPA Tool to generate a specific performance score? In this definition the only notion of performance would be asserted by the workload duration which may or may not provide an adequate indication of true system performance. Acknowledging that performance continues to rank in the top-3 purchase criteria for PC's, Intel recommends that an explicit performance score be required as an EEPA Tool output.

Line 406 – 408: the definition of Enterprise Channels being limited to managed server/client environments specifically excludes the vast majority of small/medium businesses (SMB) which presumably value Energy Star but do not contain hierarchically managed network environments. Given the size of the SMB market in the overall client market (>25%), Intel recommends that Enterprise Channels be redefined to something like "sales channels utilized by companies seeking to outfit a professional environment with PC's and ICT technology. These companies are typically in the public (government or education) or private (Large, Medium, or Small businesses) sectors.

# 2.0 Qualifying Products

Line 416 – 417: Intel reiterates it's feedback regarding Tablet PC's which may not run the EEPA Tool to the same degree as Notebook PC's.

# 3.0 Energy Efficiency and Power Management Criteria

## Power Supply Unit (PSU) Efficiency Requirements

Line 429 – 431: Intel generally supports alignment of internal PSU requirements to the Climate Savers Computing Initiative Bronze qualification requirements. It should be noted, however, that this change in PSU requirements may place an undue burden on Channel integrators to procure power supplies that meet these incrementally stringent targets vs the existing 80% efficiency targets.

To aid in the consistency and repeatability of very low platform idle power values, Intel requests the inclusion of the following efficiency targets for IPS.

| % of IPS output load | Target IPS Efficiency |
|---|---|
| 10% | ≥ 67% |
| 5% | ≥ 55% |

The 10% and 5% efficiency targets will provide guidance for long-term platform budgeting

Intel also requests the inclusion of a load vs efficiency requirement set to the 5VSB rail for internal PSU's. This was a rejected request for Energy Star v4.0 but becomes much more necessary as the global energy focus on PC Sx state power increases. To ensure consistency and repeatability of Sx state power measurements, Intel requests the inclusion of the following load vs efficiency scale for Energy Star v5 compliant internal PSU's.

| 5VSB Icc Load | Target 5VSB Efficiency |
|---|---|
| No load condition | TBD |
| 0.100A – 0.250A | ≥ 50% |
| 0.251A – 1.000A | ≥ 60% |
| >1.000A | ≥ 70% |

## Efficiency and Performance Requirements

Line 457 – 462: Intel is concerned with the apparent disparity that exists as the v5, draft 1 language that describes a desire to measure Energy Efficient Performance but never requires performance to be explicitly measured, reported, or factored into the proposed calculation model ($E_{annual}$) despite referencing a Performance Rating on line 483.

Line 467 - 471: The proposed $E_{annual}$ formula has several structural issues which Intel believes need to be addressed.

1) With $P_x$ values measured in Watts but the final $E_{annual}$ unit to be represented in kW·h, the initial annual hours multiplier should be 8.76, not 8760
2) the $T_x$ values make no attempt to account for different usage patterns between corporate and consumer environments, nor is there a provision to account for work days vs non-work days (weekends and vacation days)
3) the provision for $E_{active}$ * $N_{active}$ is structured in such a way that it would be unlikely to ever influence the $E_{annual}$ result by more than 5-10%; To more accurately account for the impact of performance on the energy efficiency of the system, it may make sense to account for $T_{active}$ (time to complete the workload) in some manner. This would show the impact of higher performance systems which can finish the workload quicker and get back into a lower power idle state sooner.
4) $N_{active}$ is likely to be purely a subjective (arbitrary) value with no mechanism to correlate the EEPA workload to actual PC workload densities over an annual period
5) $N_{workload}$ is a typo and should be $N_{active}$ instead
6) This instantiation of $E_{annual}$ makes no provision for explicit performance. Annual Energy Consumption and performance/capability of the system is not really addressed in the equation. As mentioned elsewhere, we believe that performance of the benchmark should be called out separately to help reference the capability of the system that then fits within Annual Energy Consumption targets.

## Table 2 – Capability Adders

Line 473 – 474: Intel believes that the proposed capability adders of memory and network interfaces are insufficient to address the tremendous configurability found in modern PC's. That said, the industry likely needs some time to consider what the full set of capability adders should be, balancing the EPA's timeline, data collection requirements, and impact of said capabilities on the EEPA Tool outputs. This activity may be gated by the delivery of a fully functional EEPA Tool.

Line 492 – EECoMark™ is a trademarked product by BAPCo® and should be referenced as such

## Workstation levels

Lines528-538: Intel agrees with the stakeholders that the EEPA tool would not be applicable for workstations. Key considerations are both the workload and the duty cycle. As highlighted by the industry in Energy Star for Computers v4.0 development the usage model workstations does NOT follow a typical client device, especially in the compute intensive workloads and collaborative computing models previously reviewed. Specifically, an EEPA method combining a benchmark (for active mode) with an annual energy calculation with fixed duty cycles in various modes, would not be consistent with the usage model of these machines.

Intel recommends that the SPEC GWPG benchmark quantify the energy consumption in a typical graphics and compute cycle, while the capacity of the system be used to scale the system. Intel believes that given:

a) the compute intensive nature of these systems and
b) consistency with premium characteristics described for workstations

the workload transition required in the benchmark should replicate the duty cycle of both active and inactive states of the machine. Therefore, an annual energy calculation with specific mode duty cycles does not represent the usage model of these machines. The duty cycle and manageability of the system during inactive states is part of the default power management settings (policy) defined on the platform as shipped. The energy savings from these modes would be reflected in the execution of SPEC's benchmark. Intel also recommends that to capture the variety of configurations, the capacity scaling of the Tier 1 specification be reused. This should result in a specification that compares typical use as a percentage of maximum capability., i.e. TEC = SPEC GWPG energy, and the criteria would be TEC < xx% of max_capability (Peak power at a worst case workload).

Intel recommends a parallel effort to update the Tier 1 specification methods until such time as a scaleable workload and benchmark can be developed. Such a plan will ensure a workstation benchmark would not be a limiter to a v5.0 release. If Tier1 methods are reused due to difficulties with the benchmark, we also recommend adjustments in the testing to allow energy saving activities such as HDD spindown and graphics engine suspends be allowed as part of the idle test suite, so long as these policies are part of the default configuration of the system. We also recommend that a performance energy consumption component be used as part of TEC calculation.

## Desktop Derived Servers

Lines 558-564: Intel agrees that without additional information, the EPA may need to use the previous levels for Desktop Derived Servers (DDS). Intel also agrees that a workshop is needed to both understand the barriers in achieving these levels and investigate the usage model for this class of computer .

Do recognize that by the definition, some of the manufacturers may have simply applied under the desktop categories. As client devices move to a usage model based specification, the DDS workshop should address the usage model difference, including network monitoring and small business operations. Based on anecdotal evidence, night time operations for commerce support, appear to increase the active % of the time. One will also note that there is an emerging usage model of a desktop derived server in the home, as a means of communication and entertainment consolidation. This usage model is not prevalent at this time, but considerations for future investigations may be warranted.

An alternative to the previous methods, as we have noted before, is to use SPECpower™ as an energy efficiency benchmark. The scope of SPECpower™ should provide a sufficient proxy of server operations expected from this class of machine.

## Table 4 – Thin Client Efficiency Requirements

Line 567 – 568: The requirement for Thin Clients to hit ≤ 1.0W seems very challenging given the likely support of WOL and Wake-on-USB HID devices typically supported by Thin Clients. Intel would recommend the base value for Thin Client Off mode be raised to ≤ 1.5W.

## Table 5 – Ethernet Power Management Requirements

Line 589 – 590: The requirement for IEEE 802.3az (energy efficient ethernet) is not viable for a July 1, 2009 start date. The industry EEE specification is scheduled to complete by EOY'2009 making actual

networking product availability 2010. Intel recommends this provision be removed for v5.0 specification but considered for inclusion in some future Energy Star spec revision.

## Table 5 – WOL Requirements

Line 589 – 590: All PC's must now support option to enable/disable WOL from S3. In addition, WOL from S3 is required to be enabled when shipped for all Enterprise systems: including Workstation, DDS, and Thin Client. Consumer PC's get the option to disable WOL in S3. Currently the ability for software or even end-users to determine if the system actually supports the Wake-on-LAN function is non-deterministic. Even a review of the NIC/LOM data sheet may suggest support for the feature while the platform design may not implement the feature (or implement it incorrectly by putting the NIC PHY on the wrong voltage plane to stay powered in Sx states). The only reliable way to determine whether WOL is actually supported is by visually examining the link light of a network switch when the system goes to sleep (link light stays on if supported, shuts off if it was designed incorrectly). It's unclear to Intel how to reliably detect whether WOL has been implemented correctly via the EEPA Tool approach.

## Table 5 – Network Connectivity

Line 589 – 590: This draft 1 spec requires that PC's (including thin clients) maintain full network connectivity while in S3 mode according to some platform-independent industry standard. Intel believes a more complete definition of "full network connectivity" and a reference to a particular industry standard are required before a meaningful, technical response to this item can be generated. What is the system behavioral expectation in this state (i.e….is the system expected to respond to a ping on some minimum time scale, etc..?).

## Table 5 – Wake Management

Line 589 – 590: This section seems to imply that the central management tools (IT console applications) be provided by the system mfg? Is that viable and realistic? Further, as there is no test for this capability in the test procedure – how will EPA or customers know if it has been implemented correctly?

Line 591 – 594: This provision requires that "all directed packet filters be enabled for WOL in S3 according to some industry standard default configuration." In lieu of an industry standard, EPA is requesting industry submit their current packet filter configs to be posted to EPA website and "stimulate industry discussion." The definition of a wake-up packet must be done carefully so as not to create undesirable behavior in actual practice (i.e…define a packet which wakes the system after milliseconds, and then relying on the OSPM settings to take another 30minutes before putting the system back to sleep). In general Intel believes that Wake-On-LAN should be provisioned by enterprises that will be using the features to prevent unintended ill-behavior as discussed.

Line 610 – 622*:* The following language has been removed:
"*systems where any additional management services are, at the customer's request, pre-provisioned by the manufacturer, do not need to test the system with these functions in an active state providing the function is not actually activated until there is specific action by the end user (mfg should test in pre-provisioned state and does not have to consider the power use after full provisioning occurs on site".*
Intel believes industry can provide data suggesting the market prevalence of such features and that said prevalence is meaningful. Therefore, Intel requests the pre-provisioning language NOT be removed from v5 of the specification.

Line 599 – 601: Intel advises that WOL from Off is a very unique implementation intended for use only in very specific enterprise cases. It would be undesirable for Energy Star to require support for WOL from Off and thus add system design complexity, cost, and ultimately power in off states where it is unwarranted. Intel recommends Energy Star explicitly state that WOL from Off is NOT a required support element for compliance.

Line 653 – 658: While voluntary standards may be a good thing in the area of power, Intel recommends EPA reference international or industry-wide standards which benefit from broad industry review. Further, references to standards should be made once said standard is established and beyond the risk of potential for competing proposals to address the same item.

# 4.0 Test Procedures

## Table 6 – Test Procedures

Line 686FF: Is the "TBD" in "Source" column for PC's intended to be a placeholder to reference the ECMA EEP Test and Measurement standard?

# 6.0 Future Specification Revisions

Line 729 – 734: Intel is very concerned with a defacto cadence of 2yrs being established by Energy Star revisions. Intel would recommend EPA and Industry come to some agreement on what the current v4.0 market penetration rate is and what the expected v5.0 penetration rate may be before trying to establish a v6.0? revision timeline.

# 7.0 Appendix A: Energy Star Test Procedure

Line 761 – 793: In evaluating multiple meters and capabilities it has become clear to Intel that the meter accuracy requirements as outlined in this appendix may be too stringent and in some cases insufficient to be applicable to a test procedure that includes the EEPA Tool. Intel will have further input on this topic at a later date.

Line 763 – 794: To aid in enabling Channel integrators to participate in the Energy Star program along with large PC OEM's, Intel recommends a price limit (≤$750 usd) be placed on the power meter required for Energy Star compliance testing.

Line 797 – 798 (Test Conditions Table): Intel believes the 1% tolerance on AC line voltage is too aggressive a standard to hold and will require artificial AC line conditioners or expensive AC sources to ensure the test procedure is followed. This creates artificial divergence from real-world AC environments and thus Intel requests a relaxation on the AC line voltage tolerance to ±5% which is the ANSI standard and is what local utilities like Portland General Electric (PGE) guarantee their line regulation to.

Lines 846-848: This section was added to address the differences between how desktop and notebooks/integrated desktops are measured (display on or off). Intel believes that given the current EEPA Tool direction (which will account for the efficiency of the display while active) that the idle

power of the system should be measured with display blanked on all systems. This will encourage display power savings techniques for all system classes by focusing on the efficiency of the display system while the display is active.

Further, Intel believes that during the active power test, there will be an issue of large versus small displays for notebook computers. We believe that mechanisms need to be put in place such that brightness is set based on the quality of the display versus the actual NIT output brightness (for example, turn brightness down, then turn it up until a grey pattern can be distinguished). Measuring display NITs (illuminous) output of a screen requires an expensive meter and results in manual and time consuming operations.

# Other Comments

With the proposed inclusion of an EEPA Tool which now address the active workload portion of energy consumption (and presumably short-term idle), Intel requests EPA consider a change to the Idle mode test procedure to allow for individual system components to power-manage themselves. This provision would reward industry focus on continued subsystem power management technology deployment to address the system idle time periods where the system is truly unattended prior to the 30min OSPM Sleep timeout. With this change, both short-term idle (included via the active workload and measured by the EEPA Tool) and long-term idle mode power management would be accounted for and rewarded. This change would remove the confusion that existed regarding HDD spindown and other localized device power management from v4.0 of Energy Star.

**We would be happy to further discuss our rationale for these positions at your convenience.**

**Thank you for your consideration of these comments.**

**Respectfully,**

**Jim Kardach, Erik Peter, Henry L. Wong – Intel Corporation**