



March 25, 2009

Dear EPA,

SGI appreciates the opportunity to submit comments on Draft 4 of the ENERGY STAR® Version 1.0 Program Requirements for Computer Servers. SGI also appreciates EPA's efforts in standardizing requirements for Computer Servers for the reduction of energy consumed in Servers.

In general, I thought the specification as it stands is very complete and very well understood. SGI supports the Draft 4 specification; a few minor edits are below. SGI would, however, like to see that Blade servers get included for Tier 2 of the specification as there appears to be not enough data available to get products completed this calendar year to comply as there are not enough guidelines defined as of yet. Also, the power supply efficiencies for >1,000 watt power supplies are based on gold level targets which were targeted for June 2010. Therefore, adding Blade servers for a Tier 2 release in October 1, 2010 better aligns with other initiatives that SGI has been designing to. Ideally, if the >1,000 watt power supplies had an initial efficiency requirement set to silver levels and the Idle Power requirements for Blade Servers could be established in the very near term, SGI would support adding Blade Servers in Tier 1. Then Tier 2 could introduce the Gold levels for efficiency improvements along with other considerations for energy savings.

Minor edits:

- Line 265: Should add technicians/customers as SGI customers have the availability to add/remove CPU's, Dimm's, I/O devices, ect. to their Blade Servers. The wordings add or replace multiple Computer Server boards in the field should state instead: add or replace multiple Blade Servers in the field.
- Lines 340 – 348: SGI believes the definition for Dual-Node is understood and there is a clear separation from Blade Servers definition.
- Line 1059: Should use the same numbers for an example eg., 0.852
- Lines 1081- 1083: SGI truly believes this needs to be specified/investigated. There are many issues using SPECpower in HPC environments, there are also many issues in using Linpack, as it is very configurable/tunable. Linpack can be optimized on both the system hardware and software configurations to highlight different characteristic results, eg. more flops but more power, less watts but less flops, ect.. SGI would like to suggest a max power benchmark like node_perf. SGI is doing research in the benchmarking arena for standard available benchmarks currently. SGI would like to see reporting of a performance per watt measurement for HPC applications, ex. Linpack results with WATTS per FLOP. Linpack is a widely used HPC benchmark that most of our customers are familiar with and we are asked for this result on nearly every proposal. By providing this performance/watt metric it should normalize the true performance against an Energy consumed metric.

Appendix A:

- Lines 1242 -1243: SGI states the requirements are appropriate for Computer Server Idle power testing.

Lines 1280 -1281:

SGI states that the proposed conditions are appropriate for Idle testing of Computer Servers. However, Energy Star should try to align with ASHRAE TC9.9 for Environmental Guidelines for Datacom Equipment closely to provide synergy with their potential changes for computer environmental operating boundaries.

Thank you for the opportunity to review this specification and send Idle power data for Blade Servers.

If you have any questions please feel free to contact me for the above comments and or the Idle data collection for blade servers.

Best Regards,

Bob Mascia
Principal Engineer
Silicon Graphics Inc.