



One AMD Place
P.O. Box 3453
Sunnyvale, California
94088-3453
Tel: 408-749-4000
Fax: 408-774-7216
www.amd.com

September 19, 2008

Mr. Andrew Fanara
ENERGY STAR Product Development Team Leader
United States Environmental Protection Agency
Washington, DC 20460

RE: ENERGY STAR Draft 2 Specification for Computer Servers

Dear Andrew:

AMD would like to recognize the ENERGY STAR team for its efforts to promote data center and server energy efficiency, and to thank EPA for the opportunity to comment on the Draft 2 of the ENERGY STAR Program Requirements for Computer Servers released on August 15, 2008. AMD has the following comments. Please feel free to contact Andy Rawson and/or Donna Sadowy if you would like to discuss our comments, or if you have any questions.

Section 1. Definitions

Section 1.F.

"locked down" is jargon and should be reworded.

Section 1.I.

Suggest the term "Higher Availability Servers" instead of "High Redundancy Servers."

Section 1.J.

Suggest the term "Standard Availability Servers" rather than "Standard Redundancy Servers."

Note for 1.I and 1.J: Need to explicitly include these classes of computer servers in the first paragraph of Section 2.

Section 2. Qualifying Products:

As noted above, need to explicitly include products defined in sections 1.I and 1.J in the program.

Section 3.A. Power Supply Efficiency Requirements

AMD recommends that the EPA adopt the 80 Plus organization or Climate Savers Computing Initiative's PSU efficiency targets rather than specify any standards or test procedures that differ from industry standard initiatives. This recommendation is due to the increased complexity and incremental cost associated with testing to multiple conflicting standards. Finally, the premise that it is technically easier to meet higher efficiency targets on larger (more capacious) power supplies should be validated by more industry data.

Note on Net Power Loss

AMD has technical reservations about introducing the new concept of power loss as an alternative to specifying power conversion efficiency. We believe that the power loss concept is not fundamentally easier to understand or specify. Power loss will be a function of output power just as efficiency is.

Section 3B. Idle Power

AMD supports the establishment of maximum idle power requirements for this program given that appropriate system classification "buckets" can be defined. Idle power is not only sensitive to the number of processors and the amount of DRAM installed, it is also a function of the number and capacity of hard drives currently installed in a server.

On using 16GB of system memory as a dividing line between two classes of server systems: The EPA should consider the fact that 16GB of system memory is already a norm for volume (2P) servers according to our data. In AMD's view, this "sweet spot" will increase to 32GB by 2010. This will tend to limit the number of server systems within the < 16 GB class.

Section 3B. Standard Information Reporting Requirements

AMD supports the concept of a standardized data sheet. At the same time, requirements to provide this data for maximum, minimum and typical configurations may place a significant testing and reporting burden on system OEMs. However, we defer to the server system OEMs on this point.

Section 3C. Power and Temperature Measurement Requirements

On reporting power consumption: Specifications for both the accuracy and precision of the measurement should be added to the specification as well as

the time period for averaging the power (all AC power measurements must be averaged over at least one cycle of the AC input current.)

On reporting processor utilization: The EPA should be aware that the measurement of processor utilization has a number of issues. The precise details of how this value is estimated by different operating systems varies widely and such measurements made by different operating systems are not generally comparable. If the operating system is running in a virtualized environment, the utilization measured by the O/S will generally be incorrect.

The averaging interval for processor utilization must be specified for the value to be meaningful. If there is no O/S running on the server then the value of processor utilization is both unavailable and meaningless. Finally, there is no existing industry standard communication path for the O/S to report processor utilization to a BMC or service processor so that it may be reported via an "out-of-band" manageability interface.

On behalf of AMD, thank you for the opportunity to provide comments. As always, we welcome the opportunity to answer any questions or provide additional information.

Regards,

Andy Rawson
AMD
System Architecture

Donna Sadowy
AMD
Global EHS

cc: Arthur Howard, ICF International