



ENERGY STAR® for Computer Servers Version 2.0 Draft 3 Feedback and Recommendations (October 2012)

The Green Grid Association, a consortium of industry-leading companies, welcomes the opportunity to comment on an early draft of topics under consideration for the ENERGY STAR for Computer Servers specification.

Introduction

A consortium of information technology providers, consumers, and other stakeholders, The Green Grid seeks to improve the energy efficiency of data centers around the globe. The association takes a holistic and comprehensive approach to data center efficiency and understands that developing the ENERGY STAR® for Computer Server Version 2.0 performance/power metric represents a significant challenge, one which requires cooperation among a wide range of industry principals. Participants in The Green Grid include such diverse companies as major server and storage equipment manufacturers, major software providers, and large data center end users/owners.

Summary

The Green Grid appreciates the EPA's investigations and considerations into the topics and recommendations provided by the Green Grid in the draft 2 response (June 2012) and subsequent meetings with industry members.

The Green Grid is providing comments to proposals and questions in the sections highlighted in the ENERGY STAR draft documents, dated 8/25/12. The comments may be similar to some of the individual responses provided to the EPA, but, represent the consensus opinion of the Green Grid participants in the process. The Green Grid is in general agreement with the current draft though there are several issues and recommendations that warrant consideration before issuance of a final draft. We recommend reviewing these considerations via teleconferences before the generation and release of the final draft.

We hope these comments and recommendations will aid in EPA's plans to complete the ENERGY STAR for Servers Version2.0 specification later this year. We welcome the opportunity to work in an industry forum to address the detailed areas for the final draft and release of the specification. We also recommend the development of an ENERGY STAR for Server primer and workshop after the release of the specification. The Green Grid offers its assistance and technical expertise to support the development of this collateral. We recommend completing the collateral and workshop at least 3 months prior to the effective date of the new specification.

Commentary and Feedback by Section Definitions and Scope

HPC System

The revised definition for a high performance computing (HPC) system is sufficient to differentiate these from other server categories. The inter-connect, clustering and highly parallel customized configurations are indicative of HPC systems. Due to customized configurations and specialized applications of these systems, the Green Grid agrees that these systems are not appropriate for consideration in this version of ENERGY STAR for Servers.

Large Server

Though brief, the classification of a large server as a fully integrated resilient, scaled server system does represent this category of product. The extensive I/O subsystem and high speed interconnect are also features indicative of a large server configuration and its energy profile. We also agree that due to the configuration and application, these systems should not be included as part of this version of ENERGY STAR for Servers.

Storage Equipment

The Green Grid appreciates the delineation of storage equipment and definition consistency to the ENERGY STAR for Enterprise Storage specification currently under development. The delineation and consistency will minimize confusion and overlap between these product types.

Product Family Definition by Socket

The Green Grid agrees that available socket predefines the server product family under consideration. The population of sockets is an option encompassed within the product family. The socket based definition aids in describing the supporting platform features beyond just the additional component and reflects the energy profile of the entire system.

Resilient and Scalable Server

Thank you for incorporating the characteristics that distinguishes this class of server, recognizing the distinct energy profile these features create. The resiliency and scalability features along with its energy profile is embedded in the platform architecture and not just observed by the power levels of any single device. Resiliency to data errors, scalability for large datasets, and fault resolution for big data analyses are critical to users and manufacturers of this class of server. The central processing unit and architecture are critical in determining the features and options with the other platform components. These features and architectures are not refreshed as frequently as volume servers. Therefore, the systems reviewed by Energy Star are indeed reflective of systems sold into the market at that time.

Product Family Testing

The Green Grid understands and concurs that the envelope of products could be well defined by the 5point test plan. As stated in our draft2 comments, due to definitions and configurations, qualified products may exist outside of the described points. To avoid additional testing, cost, and product families, we recommend including the terms, "**lower-price or lower performance**" and "**higher price or higher performance**" to the configuration to allow the flexibility of defining the boundaries represented by the product family.

Power Management Requirements

TGG agrees with the recognition and inclusion of in-band power management options. Operating systems and hypervisors are critical parts of the coordination enabling power management in the system.

As noted, much of the documentation for these systems is supported electronically or on-line, in lieu of physical documentation. Physical documentation is generally insufficient to support the installation and configuration of this class of product. Text in lines 395-406 regarding documentation included for blade systems should be clarified to support an electronic only method of distribution.

Active Mode Efficiency Disclosure

TGG concurs that it's prudent to keep section 3.5 as TBD pending the results of the pre-release evaluations of SPEC's SERT™ tool. SERT™ development is proceeding as planned and is expected to be a reporting requirement in the final version of the specification. The Green Grid believes the data collection is very important but does not believe there will be sufficient experience in the worklets to use the results as a comparative indicator of energy efficiency at this time. Premature association of these numerical values to a specific system may cause confusion and inappropriate comparisons. We recommend that data publication by ENERGY STAR remain anonymous (un-tethered) to the actual product family or product until Q2'2014 (assuming current specification release schedules) We suggest that the EPA consider a plan to issue a minor revision to the ENERGY STAR specification identifying the indicator(s) to associate with specific product(s) after this assessment. This method would allow multiple system, operating system, and architectural analyses across all of the SERT™ workloads before identifying which indicators may be representative of a product and/or product family.

SERT disclosures should also list limitations of what the values represent, such as:
“Output values obtained from the SERT utility are intended solely for Energy Star qualification purposes, based on a limited, conservative sample set. Actual results may vary. SERT output values listed here are intended to represent a precise set of configurations, not necessarily reflective of all available configurations.”

Idle Mode and Full Load Efficiency Criteria – one-socket (1S) and two-socket (2S) Servers (Non-blade)

TGG concurs with change to the base idle and allow an 8W adder per hard drive installed to address those systems which do not contain local hard drive storage.

We acknowledge that there remain limitations in reducing the additional idle power introduced by redundant power supplies. As integration and power densities increase, the 20W adder remains challenging despite improvements in PSU efficiency. We see little evidence to suggest the ability to reduce this adder at this time.

Idle Mode and Full Load Efficiency Criteria – Blade Servers

TGG agrees that reporting idle and full load in a half populated chassis and provisions to provide fully populated chassis data in lieu of the partial population is a reasonable compromise. We appreciate the considerations for the expense and time a fully populated chassis testing would take. Please note that even half populated chassis are very expensive and resource intensive to

configure and test. We recommend considerations for reduced or limited verification be required for this class of product and/or all servers due to the complexity and expense.

Other Testing Criteria

We appreciate the recognition and adjustments made to accommodate the emerging market for Auxiliary Processing Accelerators. TGG has no major concerns with designating the additional computational capability adapters as Auxiliary Processing Accelerators (APA). The term is not generally used in the industry. Training documents may wish to include examples to allow more familiarity with the term. Inclusion of specific implementations or product names may be inappropriate in the main ENERGY STAR specifications.

TGG members are, however, concerned with the limit of 46W of additional idle regardless of number of APA's installed. The limit does not reflect the power levels required by these new capabilities either individually or in multiple card configurations. APA's are under investigation for use to address emerging applications such as "Big Data" analyses and analytics. Some configurations are being evaluated for use in lieu of relying on resources at large HPC facilities to conduct these analyses. As described in a previous TGG-EPA meeting, the local high speed interconnect, memory and management functions embedded in these cards are additional requirements beyond the additional processor. In a recent assessment of graphics processors, many of which are also used in APA's, the range of system idle power due to these processors not including additional APA features indicates up to 45W for a single adapter (idle range reference example: <http://www.anandtech.com/show/5699/nvidia-geforce-gtx-680-review/19>).

Given the limited experience with this emerging technology, we recommend compliance testing be conducted without APA's and inserting a requirement of reporting system idle power data with APA(s) installed. This would be consistent with the current approach to collect information to assess the power impact before creating limits on the technology. The data would provide information on the idle power impact of this emerging technology. We do not recommend the fix limit idle power adder of 46W prior to the active use of this technology. We are concerned that the fixed idle limit adder of 46W may stall technological advances and inhibit more energy efficient options in big data analytics. If a maximum idle power limit is absolutely necessary, the limit should reflect the range of solutions available and should be per accelerator to minimize unintended impact to this emerging technology.

Standard Information Reporting Requirements

We recommend that ENERGY STAR conduct several trial runs in populating the new power performance datasheet in addition to aggregating and posting the information. We suggest this may be conducted after the release of the SERT™ tool which is expected to include hardware detection and other reporting tools that would aid in the accuracy and consistency in the reporting. Documentation and data entry expectations can also be reviewed in an ENERGY STAR for Servers testing workshop prior to when version 2 becomes effective.

The request for annual energy consumption estimates is subjective and highly dependent on the system configuration, application, and supported industry the servers are deployed. We request that this be listed as optional and that estimates include use condition assumptions and configuration information.

Using the results from SERT™ to report idle data in addition to performance information appears reasonable. Please note that due to resident workloads, the registered idle power may be higher than those observed in the previous non-application loaded test procedures. The additional power levels are expected due to application maintenance and the multiuser response expectations on servers. The magnitude of the idle power difference is unknown pending investigations on the

final SERT™ tool suite. We expect larger memory and resource configurations may cause a larger incremental residual maintenance energy for these systems. Therefore, despite no numerical change to the idle criteria, testing to idle after running the workloads results in a more challenging specification. However, the methodology is sound and we support ENERGY STAR's direction to base the idle power compliance on these results.

Effective Date

Given that product releases may occur between the release of version 2 and when the specification becomes effective, we request that systems which have submitted qualification information and would be qualified under version 2 be also listed on the product qualification list currently issued for version 1. We recommend that these systems would be on the version 2 qualified product list once version 2 becomes effective, without the need for resubmission. We believe this is reasonable given that version 2 is effectively more stringent than version 1 and this avoids a logistical gap of availability between revisions.

Conclusion

The Green Grid fully supports the development of the ENERGY STAR for Computer Server version 2.0 specification and the target release by the end of 2012. EPA's collaborative development with all industry stakeholders should allow rapid resolution to the remaining concerns. We believe the data collection process, active mode investigation and considerations for the technologies and product configurations currently being produced, will provide valuable insight to products in this industry in addition to advancing energy efficient growth in this sector. The combination and consistency of the ENERGY STAR for Computer Server program and the efficiency initiatives in the EPA and US DOE should help in accelerating the efficiency in operation of the data center. The Green Grid will continue to collect industry-wide inputs to work with the EPA in developing the ENERGY STAR programs on ICT equipment. Please feel free to contact us on any concerns or questions in the development of the specifications or the implementation of the program.
