NRDC

NATURAL RESOURCES DEFENSE COUNCIL

September 2, 2016

Mr. Steven Hanson
United States Environmental Protection Agency
ENERGY STAR Program
1200 Pennsylvania Ave NW
Washington, DC 20460

Subject: NRDC comments regarding Draft 1 Version 3.0 Specification for Computer Servers

Dear Mr. Hanson,

On behalf of the Natural Resources Defense Council (NRDC), we respectfully submit the following comments in regards to the ENERGY STAR Draft 1 Version 3.0 specification for computer servers issued July 27, 2016.

Computer servers (hereafter referred to as servers) are the workhorses of data centers. According to a recent study by Lawrence Berkeley Labs[i], servers draw an average of 100 watts (1-socket servers) and 250 watts (2-socket servers), typically 24/7. When aggregated across the roughly 15 million servers in use in the United States, this amounts to more than 30 billion kilowatt-hours of annual energy use, with an additional 15 to 20 billion kilowatt-hours to cool these servers and keep them from overheating. This adds up to more than 1 percent of total U.S. energy consumption.

NRDC strongly supports the ENERGY STAR program for servers: it is an important tool to enable businesses of all sizes, as well as local, state and federal government agencies, to reduce their energy use through sustainable procurement policies for their data center equipment; it also enables electric utilities to implement efficiency incentive programs that reduce energy use in their service territories, and that help transform the market toward more efficient equipment.

We commend EPA for initiating the revision of the Server specification in reaction to the evolution of server technology since Version 2.0 went into effect in December 2013. We generally support EPA's Draft 1 proposal, and offer the following comments aimed at improving the draft specification. Our comments cover the following points:

1. **Idle and Active States Criteria:** NRDC recommends keeping the Active and Idle State criteria separate, to ensure that ENERGY STAR servers deliver on customer expectations of energy performance in the conditions they are typically operated in;
2. **Proposed levels and adders**: NRDC requests that the dataset used to develop proposed Idle State levels and adders be made publicly available to allow stakeholders to perform their own analyses and inform their comments on proposed levels and adders;
3. **Memory adders**: NRDC encourages EPA to examine alternative approaches to set memory adders in order to ensure that the specification achieves its objective of targeting the top 25% of the market at effective date;
4. **Power supply requirements:** NRDC proposes alternative power supply efficiency requirements that would better identify the products that save the most energy in real-world operation.

Here are our detailed comments:

1. EPA question: Should Active State and Idle State criteria remain separated as is currently proposed in Draft 1, or are there technical merits to combining them?

NRDC recommends keeping the Active and Idle State criteria separate, to ensure that ENERGY STAR servers deliver on customer expectations of energy performance in the conditions they are typically operated in.

While customer operating conditions vary widely, it is common knowledge that many servers spend a lot of time idle or operating at low load (10 percent or less), either because of periods of low activity (such as at night and weekends), or because they are part of spare computing capacity that only gets used a few hours each month or each year. A significant number of servers, between 20 and 30 percent of servers by some estimates, are even permanently unused (the so-called zombie servers). Idle State therefore remains an important metric for the operation of a typical server in both idle and low-load states, and for the energy use that customers can expect.

There have been concerns that high-performance servers typically used as virtualization hosts may be disqualified because of poor efficiency in idle state, which would in turn reduce energy savings from virtualization. We believe this is a false choice: we see no technical reason why virtualization host servers, with the right allowances for the additional memory and for other functional capabilities they require, would not be able to achieve idle state requirements. In fact, servers which cannot achieve good efficiency

in idle state and low load should not be qualified for ENERGY STAR, because when operated in idle and low-load conditions that are typical in data centers, they would be inefficient and use more energy than most of servers in those conditions.

Combining the Active and Idle State metric would allow tradeoffs that may not yield expected energy performance in customer environments, such as focusing on efficiency at high load points that are not representative of typical use. We therefore recommend keeping both metrics separate and requiring servers to meet both requirements, just like 80-PLUS power supplies must achieve efficiency levels at all load points.

2. NRDC requests that the dataset used to develop proposed Idle State levels and adders be made publicly available to allow stakeholders to perform their own analyses and inform their comments on proposed levels and adders

Assuming EPA based its proposed levels and adders on a dataset that contains both existing v2.0-qualified products, and additional non-qualified products, NRDC would like to perform our own analysis of this dataset to inform our comments on these proposed levels and adders. We therefore respectfully request that EPA's v3 dataset be made publicly available, in order to allow NRDC and other stakeholders to perform an analysis of the effects of EPA's proposed Idle State levels.

Setting appropriate levels for both base allowances and functional adders (memory, storage, power supplies, etc.) is critical to getting the specification right and ensuring it remains effective over its intended life. For example, memory allowances proved too high in v2.0, which allowed most servers with large amounts of memory to beat ENERGY STAR levels by 50 percent or more. We would like to ensure we get levels and allowances right in version 3.0.

3. NRDC encourages EPA to examine alternative approaches to set memory adders in order to ensure that the specification achieves its objective of targeting the top 25% of the market at effective date

We question whether setting memory allowances per gigabyte (GB) is appropriate: given the rapid evolution of the market for computer memory, with an exponential increase in the capacity available for a given price, adders that scale linearly with capacity may rapidly render the specification obsolete. EPA's proposal may be appropriate based on the current dataset, which represents technology introduced on the market over the past few years, but given the time required to finalize the specification and then the 9-month period for the specification to become effective, it may be obsolete before it even goes into effect.

We encourage EPA to examine alternative approaches to set memory adders, such as scaling logarithmically with capacity instead of linearly, to account for the exponential growth in memory capacity, and/or to set allowances by memory module (DIMM) instead of per gigabyte.

We look forward to making more detailed proposals once we have access to the 3.0 dataset.

4. **Power supply requirements:** NRDC proposes alternative power supply efficiency requirements that would better identify the products that save the most energy in real-world operation.

We support EPA's intent to strengthen power supply unit (PSU) efficiency requirements to reflect the fact that 63 percent of configurations tested in v2 achieve Platinum level.

However, we recommend that EPA considers alternative requirements that would better identify the products that save the most energy in real-world operation. As explained in our response to question 1 above, the average load of typical servers is in the single digits or low double digits, and it would make sense for PSU requirements to reflect this situation by setting more stringent requirements for 10% and 20% load than for 50% and 100%. 80-PLUS Platinum criteria are more stringent at the 50% and 100% load points than at the 10% and 20% load points, compared to the mean of all the Platinum units in the 80-PLUS database to date (530 units), as shown below:

Load:	**10%**	**20%**	**50%**	**100%**
Platinum req. (single-output)	83.00	90.00	94.00	91.00
Min	78.80	90.19	94.00	91.00
Max	94.41	95.50	96.24	95.79
Mean	88.26	92.71	94.38	92.78
STDEV	2.44	0.96	0.33	0.78
Difference mean/Platinum	5.26	2.71	0.38	1.78

It should be the opposite, to better reflect the average load of servers in real-world operation. We propose alternate requirements, based on the 80-PLUS data, in the table below.

80-PLUS Platinum is also missing a 10%-load requirement for multi-output PSUs. Because of the importance of low-load efficiency in servers, we recommend EPA sets a 10%-load requirement for multi-output PSUs, and we propose a requirement of 86%, in line with other requirements.

Lastly, EPA's proposed power factor requirement of 0.9 at 50% load isn't aligned with 80-PLUS Platinum which requires 0.95 at 50%. We recommend aligning this requirement in the interest of consistency.

We therefore propose the following updated requirements (updates in bold):

Efficiency

Load point	Rated Output power	10%	20%	50%	100%
Ac-Dc Multi-output	All output ratings	**86%**	90%	92%	89%
Ac-Dc Single-output	All output ratings	**88%**	**92%**	94%	91%

Power factor

Load point	Rated Output power	10%	20%	50%	100%
Ac-Dc Multi-output	All output ratings	N/A	0.80	**0.95**	0.95
Ac-Dc Single-output	<= 500 W	N/A	0.80	**0.95**	0.95
	> 500 W and <= 1,000 W	0.65	0.80	**0.95**	0.95
	> 1,000 W	0.80	0.90	**0.95**	0.95

Thank you for the opportunity to participate in this specification development process and for your consideration of our comments.

Sincerely,

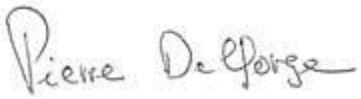

Pierre Delforge
Director, High Tech Sector Energy Efficiency
Natural Resources Defense Council

[i] Shehabi et al., "United States Data Center Energy Usage Report", June 2016