ENERGY STAR® for Large Networking Equipment Draft 1 Version 1.0
Feedback and Recommendations (September 2014)

The Green Grid Association, a consortium of industry-leading companies, welcomes the opportunity to comment on topics under consideration for the ENERGY STAR for Large Network Equipment specification.

Introduction

As 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 Large Network Equipment Version 1.0 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 data center networking equipment manufacturers, 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 of comments from some industry members during the early stages of the project. The Green Grid is providing comments on the Draft 1 Version 1.0 ENERGY STAR® LNE specification and Draft 2 LNE Test Procedure. Some of these comments may be similar to individual responses provided to the EPA by member companies, but represent the consensus opinion of the Green Grid participants in the process.

The Green Grid is dedicated to improving the energy efficiency of data centers and recognizes the importance of large networking equipment to this efficiency. The Green Grid has a body of work covering all aspects of data center efficiency including components and interactions at the broader system level. The Green Grid is concerned that the Energy Star program for Large Network Equipment should have a beneficial effect on the energy efficiency of data centers; however there is a real possibility that some approaches that minimize the energy consumption of single systems may adversely affect the efficiency of the data center as a whole.

The Green Grid offers its assistance and technical expertise to develop and refine future versions of this specification. We hope these comments and recommendations will aid in the management of an effective and timely program. If you or stakeholders have questions, please feel free to contact Gary Verdun or Henry M. Wong.

General Comments on the Program

Overall

For a data center, the productivity depends on the performance of the computing equipment that is performing its primary data processing functions. The data center may be considered efficient if it can perform its primary function effectively while using a reasonable minimum of energy to do so. This energy is used by the computing equipment as well as infrastructure and other equipment (such as networking) that supports the operation of the primary function. In some cases, deficiencies of support equipment may detract from the performance of the primary function and this can be detrimental to the efficiency of the data center.

For example, if the data center network suffers from bottlenecks due to insufficient bandwidth, some of the compute servers will be forced to wait for data and will therefore have to spend more time in high power states waiting to perform their primary functions. In this situation, the networking equipment, that comprises a small part of the data center energy footprint, causes an increase in energy usage for the compute equipment that consumes the majority of energy for the facility. Clearly, minimizing the energy usage of the networking equipment at the expense of efficiency of the overall data center is a poor choice for minimizing total energy footprint. The primary goal for the network design should be for the network to adequately support its primary function, while minimizing the energy consumption of the wider ecosystem.

Similarly for enterprise networks, the total energy consumption of client devices will be much higher than that of networking equipment. Therefore, attempted optimizations of energy consumption for such shared resource equipment should not compromise the overall function at the risk of increasing the overall energy consumption of the operation.

The Green Grid takes a holistic view of energy consumption and therefore urges EPA to take a similar view and develop program requirements for Large Networking Equipment that do not risk false optimizations of specific components at the expense of system-wide efficiency. The Green Grid particularly recommends that energy conserving modes which include reduced functionality or sleep should be avoided for devices that serve many endpoints as they may result in widespread performance problems that can also lead to increased overall energy consumption.

For data center networks and for enterprise networks, the design of the network (and therefore the selection of network equipment) is generally performed by experts who will assess the needs of the application, the equipment interdependencies, and select appropriate equipment to meet those needs. Experts who are evaluating such candidate networking components are able to compare energy efficiency data for those components as a factor in the design decision as long as standard metrics are available. The Green Grid supports the availability of such metrics and has published recommendations for energy efficiency assessment based on the ATIS TEER metrics in the white paper: Energy Efficiency Guide For Networking Devices. The Green Grid does not support setting per unit (subsystem) performance level requirements based on those metrics as they will lead to false optimizations as described above. Comparisons should also always be made between devices that are fit for function and not between dissimilar devices. It would be impractical to create a sufficient number of categories to ensure that all comparisons are valid. Networking systems are available with varying numbers and speeds of downlink ports; varying numbers and speeds of uplink ports; varying complexity of functionality for switching and routing; and varying complexity of functionality for other packet processing. All of these dimensions are independently variable and could be used as classifying characteristics. Furthermore there are newer technologies (such as converged networking and software defined networking) that simultaneously create new categories and gray the distinctions between

categories. A network designer should be encouraged to select multiple candidate networking devices that are appropriate for a purpose and choose between those candidates for optimal energy efficiency.

Scope and range of networking products

The definition and separation of “Small Networking Equipment” (SNE) and “Large Networking Equipment” (LNE) is fundamentally incorrect. The arbitrary selection based on number of physical ports is inconsistent with current and future deployments and usage models. Currently a 10Gb 8 port Ethernet switch would be considered SNE and would fail to be able to get qualified under that specification. Customers that are required to purchase Energy star would be forced to over provision the network by buying 12 port switches and leaving 4 ports unused.
Given networking advancements in speed, port aggregation, software defined networking (SDN) and software defined infrastructure (SDI), number of physical ports would not be able to describe such a separation. We believe that the intent for the separation of LNE and SNE is to describe those systems being used in an industrial or commercial setting (LNE) and those used in a home environment (SNE).

To better separate networking equipment used in commercial environments and home environments, we recommend that total maximum system throughput be used to separate LNE from SNE systems. Based on practical technical limitations, SNE networks are limited to < (less than) 20Gbps, whereas entry level commercial systems are predominantly > (greater than) 24Gbps.
Though there may be some instances to install optical connections to domestic systems, the preponderance of units would support the use of throughput as a separator and an inherent limitation to SNE networks.

Please note that most domestic (one way) gateways have a practical limit at 10Gbps, hence the <20Gbps throughput limitation. For commercial settings, <24Gbps becomes difficult for communications and business operations. Some of the lower entry level switches start at 24-48Gpbs, so 20Gbps seems like a natural limit and designation between SNE and LNE.

This proposed method removes obvious enterprise or data center devices from the SNE specification but there will still be an overlap of commercial grade products and consumer products in that specification that differ primarily in other features such as:

- Power over Ethernet (POE)
- Management Capabilities
- Networking products supporting Campus, Branch, or Teleworker connectivity, with option for VPN
- Deep or statefull packet inspection.
- Role Based Access and Virtual LANs
- Multi-device Redundancy and/or Failover capabilities.
- Products certified for one or more government requirements such as:
    - TAA (Trade Agreement Act)
    - Section 508 of the Rehabilitation Act
    - FIPS 140-2 (Federal Information Processing Standards for Cryptology)
    - National Security Agency Suite B Cryptography

These features drive additional power consumption that will prevent these devices from meeting power limits that are obtainable by consumer products. Without proper accommodation of these kinds of features networking devices supporting campus and office environments will be prevented from obtaining Energy Star qualification.

Sub-categories.

Within the major categories of LNE and SNE, there exist at least 2 distinct functions that may be configuration optimized: core and edge networking systems. Though these may be configurable, there are features that tend to increase a particular system's use as a core or edge system. Within the LNE group, there are additionally classes of devices the energy profile data may create such as 'Core', Top of Rack (ToR), End of Row (EoR), switches integrated in various ways within the server chassis itself (e.g. blade via backplane, physical media (MII), or compute-interconnect (PCIe)). We suggest reviewing such features and defining subcategories based on feature set and capabilities.

Technology considerations

Software Defined Network (SDN) and Network Function Virtualization (NFV) are emerging technologies that provide the flexibility of provisioning network configurations depending on the connectivity demands of the application. The software controlled provisioning optimizes the use of the physical port structures without changing out or dedicating specific hardware. Network virtualization technology also supports the general grouping of product based on total throughput capacity before further sub categorization of the products.

Detailed Comments by Section for Draft 1 Specification

1.a.2.a

Existing Text:
Modular Product: An LNE product in which half or more of the total number of available physical network ports in the product are swappable or interchangeable.
Response:
The combination of the 50% limit and the definition of physical network ports will create scenarios where the determination of whether a product is considered fixed or modular under the specification will become configuration dependent. This will create significant confusion and force manufacturers to qualify a single system twice under this specification to accommodate all configurations. Definition of modular product should be based on capability to accept modules using the specifications definition of a module, not on 50% of ports which can be changed.

1.C.4-5 and note box

Q.
EPA is proposing Core and Edge definitions to separate LNE products which are often not fully loaded or connected from products which face routinely higher loads and are typically fully connected. This distinction is reflected in the accompanying Draft 2 Test Method in sections 6.1 and 6.2. If adopted, it will inform future development of active state energy efficiency criteria. EPA is proposing to allow manufacturers to choose which of these descriptions more accurately reflects their product or product family, and that is the designation that will be referenced for all specification and test method purposes related to that product or product family.

Response:
Core and edge products differ in locality on the network and specifically in proximity to the end of a network branch. Edge devices have end points or non-networking devices connected to their downstream ports. Many networking devices intended for data center type applications can be used as core or edge devices. It is impossible to determine if these devices will be used as core or edge until they get installed by the end user.
Utilization is a side effect of end the user' application of the product and is inappropriate as a part of the definition of core vs edge.
Power management and use of link states will vary between devices configured for core operation as opposed to edge operation. In most cases, networking systems configured as core will likely not enable power management capabilities as the device serves as a critical hub of machine to machine communications. All forms of power scaling with workload have side effects in the form of latencies for traffic appearing after power management has been asserted. Networking devices configured for the edge, will need to logically support edge networked systems and in some cases power down the entire sub-network of networked edge devices. The networking device configured for the edge would not power down unless all devices in its subnet is inactive.

1.D.5

Existing Text:
Line Card: A pluggable module that can be inserted into the backplane of a modular LNE product to provide various forms of connectivity to edge products connected to the network.
Response:
Defining backplane as place of installation is prescribing a physical system architecture. Should instead state as option card or option slot so as not to prescribe:

a) The system has a backplane architecture and
b) The backplane is the location for installation of line cards

1.F.2

Existing Text:
Idle State: The operating state where the product is capable of carrying out work, but is not actively transferring data.
Response:
Considering the type and complexity of networking products a definition such as this one could potentially occur between packets on certain types of products. EPA should probably focus definition on no data packets in process of reception, no data in transit across the device or currently being transmitted on output ports. A clear definition of actively transferring data incorporating activities of the network device inclusive of reception, processing and transmission of packets would also suffice.

1.G.1

Existing Text:
Physical Network Port: An integrated physical connection point primarily intended to accept IP or similar traffic via a cable. Fiber-optic connections are not considered Physical Network Ports for the purposes of this specification.
Response:
TGG would like to point out that the above definition of Physical network ports in combination with the LNE definitional requirement of greater than 11 physical network ports will eliminate a significant portion of fixed core networking products from the specification. Many if not all modular products will be in a situation where some configurations will be in scope and others will be out of scope. This will cause significant complexity and confusion on the part of manufacturers, test labs, Certification Bureaus and customers that are required to purchase Energy Star products.

TGG recommends redefining SNE as networking devices with total maximum throughput of <20Gbps and LNE as networking devices with total maximum throughput of ≥ 20Gbps.

3.2.3.i and 3..2.3.ii

Existing Text:
Partners are required to measure and report PSU power factor under loading conditions of less than 75 watts, though no minimum power factor requirements apply.
Response:
For Power supplies larger than 750W this text forces an additional measurement of power factor for a load condition below the 10% load point.
Alternate Text:
For load points in table 3 that fall below 75W power factor should be measured and reported though no minimum power factor requirements apply.

3.3.1.i

Existing Text:
Port Power Down: An LNE product must have the ability to power down unused physical network ports in an automated fashion, which does not require input from the end-user.
Response:
A definition of power down is needed to clarify whether the power management implementation in a given system qualifies under this requirement. Many implementations exist which reduce PHY power on unused ports. A method needs to be devised to determine which of these qualify as meeting EPA's port power down requirement.

3.2

Tables 2 and 3 state efficiencies and power factors without stating input test voltage. 80 Plus specifications now have different efficiency and power factor values for different input voltages. EPA needs to update these tables per the new 80 plus requirements and identify the input voltage intended for the indicated requirements.

3.4.2

Existing Text:
Active State Efficiency: To certify for ENERGY STAR, an LNE product must meet the following l active criteria [TBD].
Response:
Proposed E* test method and specification are heavily based on ATIS 0600015.03.2013, but this standard was designed specifically to provide network design information and avoid setting limits. The value of the ATIS generated numbers is dependent on the configuration of the network and the device's role in the network. The specification is intended to provide the capability to compare efficiency numbers for functionally very similar products, each of which is capable to replace another in the same application within a customer's network.
Very detailed product classification and feature set compensation would be required to appropriately set limits that account for the functions in networking equipment that drive energy consumption and efficiency scores according to the ATIS test methods.
These classifications do not exist in the ATIS standard. Previous attempts to create such product classifications have not been successful and TGG does not believe such an effort would be successful as part of the LNE spec development.
TGG proposes to follow ATIS standards, ATIS 0600015.03.2013 and ATIS 0600015.2013 and disclose test results for all fixed and modular products. Customers would then have the data available to compare like products and the EPA and industry would have an opportunity to collect and evaluate this data set to determine feasibility and appropriateness of defining limits.

4.1.2.ii

Existing Text:
product characteristics (utilization type, modular vs. fixed, power specifications, etc.);
Response:
Utilization Type is indicated as a data reporting element but is not defined in the specification. All required data reporting elements need to be clearly defined.

7.1.1

Existing Text:
Note: New specifications are generally effective immediately upon completion.
Response:
Industry needs to be provided reasonable and sufficient period of time to schedule and complete testing of systems between the time the spec is published and its effectivity. Immediate effectivity will create a window of time where the EPA process in regards to the LNE specification favors one supplier over for purchases by the federal government..

Specific Questions from Draft 1 Specification

Q.

EPA received several comments to focus on the two states specified above. However, several stakeholders suggest that EPA include a new Low Power State that is an optional state available in some LNE products. EPA is requesting additional information on the functionality of Low Power States and how prevalent they are in LNE products.

Response:
As noted above, systems configured as core will likely not enable autonomous low power states due to connectivity requirements. For edge systems, support should only be needed in the derivative case where all ports are inactive and SLA's will tolerate the delays to resume operation.

Q.
EPA received stakeholder requests to provide further clarity on the delineation between fixed and modular LNE products. Suggestions were provided to create three types of modularity. Due to a lack of data to support the suggested modularity types at this time, EPA is proposing to separate fixed and modular products by whichever port type is more prevalent in the product. EPA welcomes feedback on this proposal.
Response:
See response to section 1.a.2.a above.

Q.
Additionally, EPA welcomes stakeholder feedback on alternative means to delineate between Small (SNE) and Large Network Equipment. The current approach of using port count was created during the development of the SNE V1.0 specification. Should an alternate approach that more effectively delineates be identified, EPA anticipates applying the approach to both product specifications.
Response:
TGG believes that total available bandwidth would be more appropriate to distinguish between enterprise and consumer class networking equipment than the current port count method as described earlier in this document.

Q.
Note: EPA received stakeholder feedback on the Framework Document stating that an unmanaged LNE product is not an applicable characteristic. Rather, the managed product characteristic should be divided by whether the management of the product is occurring internally through the use of a management processor, or externally over the network. EPA welcomes feedback on this proposal.
Response:
While there are certainly different implementations of management capabilities in LNE products TGG does not understand the need to distinguish between products based upon this difference.

Q.
Finally, EPA is proposing Core and Edge definitions to separate LNE products which are often not fully loaded or connected from products which face routinely higher loads and are typically fully connected. This distinction is reflected in the accompanying Draft 2 Test Method in sections 6.1 and 6.2. If adopted, it will inform future development of active state energy efficiency criteria. EPA is proposing to allow manufacturers to choose which of these descriptions more accurately reflects their product or product family, and that is the designation that will be referenced for all specification and test method purposes related to that product or product family.
Response:
TGG responded regarding Core vs Edge definitions previously and noted a large number of products that could serve as core or edge. Allowing the manufacturer the choice of test methods would be a necessity should the spec continue with this level of ambiguity in direction regarding test method.

Q.

EPA has added IT equipment definitions to describe the equipment proposed to be out of scope of Version 1.0. The following definitions are fully harmonized with other ENERGY STAR specifications and should not require revision: computer server, storage product, uninterruptable power supply. EPA welcomes stakeholder feedback on the proposed definitions in Section 1) I) "Other Enterprise and Datacenter Information Technology Equipment".
Response:
SNE definition included here is only a partial copy of the definition in the SNE spec. If EPA does not include entire definitions it should reference the definition source specification for clarity.

Detailed Comments by Section for Draft 2 Test Method

5.1.a.4 (Note)

Existing Text:
Note: DOE received stakeholder feedback stating that including PDU overhead power might be inappropriate, since a system with one PSU should be measured the same way as a system with two. For this reason, the use of PDUs is no longer permitted in the Draft 2 Test Method.
Response:
EPA needs to reconcile this statement in test the test method precluding the use of a PDU and the Spec allowance of one in the data collection section. On relatively low power systems with multiple PSU's the PDU power consumption may be significant relative to the total system power and the number of required power meters is not an issue. On larger systems with many power supplies (8 or more) each rated above 1kW, PDU power consumption would not be significant and the number of required power meters would become an issue. EPA should make the use of a PDU optional instead of requiring or precluding their use.

4 B)

Existing Text:
Ambient Temperature: Ambient temperature shall be 27°C +/- 1°C.
Response:
This is the accuracy limit of many environmental chambers and inconsistent with Servers 2.0 (25 +/-5) and Storage 1.0 ($18 \leqslant Ta \leqslant 28$). A +/- 1°C allowable ambient temperature range will likely force testing to occur in environmental chambers. Servers, storage and large networking products are intended for operation in the same end use environment and should be tested under the same ambient temperature requirements.
TGG recommends harmonizing with one of the above specifications.

Conclusion

The Green Grid fully supports the development of the ENERGY STAR for Large Networking Equipment Version 1.0 specification. EPA's collaborative development with all industry stakeholders should result in a well-constructed specification and should allow rapid resolution of concerns raised.

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.