July 26, 2013

To: Robert J. Meyers
Project Manager, Energy Star for Large Network Equipment
United States Environmental Protection Agency

Re: Hewlett-Packard's Response to the EPA request for feedback on *"LNE Framework Document and Draft 1 Test Method"*

From: Hewlett-Packard Company, HP Networking

Hewlett-Packard has been a strong supporter of responsible stewardship of the environment and will continue to demonstrate its commitment through participation in all applicable green programs, including those sponsored by the EPA. HP also applauds the progress of the EPA's Energy Star program for the improvements in energy savings it has stimulated through fair and reasonable criteria for products across a wide variety of industries and product categories. And HP intends to support and cooperate with EPA's Energy Star program as they extend across new categories and product types, including networking equipment.

In review of the Energy Star LNE program's early drafts, HP would like make both general comments and provide specific feedback on the proposed drafts for framework and test method.

As a general comment, HP is concerned that the categorization of product types may become too generalized and broad ranging for fair and equitable comparison across networking equipment manufacturers' product offerings. In fact even the ATIS test standard, which is being used as a basis for tests metrics for the LNE Energy Star program, has recognized the broad range of networking equipment categorization. In Annex A of the test standard, ATIS has suggested nine classifications for routers and no less than 14 specific categories for switches. Without a clear classification of the differences in product features, many competitive products will be unfairly "lumped" together with less featured equipment, and will measure poorly in any comparison metric solely aimed at "performance per watt".

HP would like to suggest that the definition of a network switch and router requires further clarification. As networking technology has evolved there has become more overlap and confusion in the industry for these definitions in general. The classic definition of a switch only considers the addressing scheme of layer 2 in the OSI model, and today is typically referred to as a layer 2 (or L2) switch. L2 switches tend to have lower cost, consume the least power and have limited capability. A router moves traffic based on the address assigned at layer 3 (L3) or higher of the OSI model. A router normally requires some level of configuration to interoperate within a defined addressing scheme of its network topology. As such, classic routers use more power, have deeper routing tables and require more hardware. Over time, the distinction between classic switches and routers has become blurred, as switches have added features previously only found in routers. These switches are commonly called L3 switches or switching routers. Furthermore, limited L3 functionality has been added to some switches and is referred to as "Light Layer 3 switches" (or LL3). It is HP's recommendation that there be some distinction between L2, LL3 and L3 switches for fair comparison in any Energy Star metric.

Depending on context, another separate class of network devices is also referred to as routers. In this definition, devices that route network traffic to non-Ethernet ports are also called routers. These routers span from the consumer level to core data center equipment. The bridging interface can be a WAN interface, such as DSL, ADSL, Modems, Sonet and other telecommunication protocols, that are not specifically Ethernet. This class of routers requires a separate classification.

Another example of a possible categorization issue is with a class of managed network switches that support a function called “stacking”. More specifically, stacking is the ability for an L3 or LL3 switch to provide link aggregation across multiple switches. A stacked switch essentially can aggregate multiple discrete switches into one autonomous L3 switch. This distinction becomes important for Energy Star because it requires additional hardware for every switch in the stack and likewise consumes more power to support stacking. For example, a 96 port switch that supports stacking may perform similarly to a fixed 96 port switch, when measured purely for performance in a stand-alone configuration, and when placed side by side can look nearly identical in terms of port count and media speed. However, the ‘stacking’ switch will consume more power because of added capability in silicon. This additional hardware for expandability, which is commonly implemented through both proprietary high speed interfaces, and industry standard links allows the bridging of multiple switches together, enabling higher total bandwidth. The resulting array of stacked switches can act as a single autonomous switch with much higher port count. Compared to a non-stacking switch, when measured as a single stand-alone unit, a stacking switch would be not likely achieve the same performance per watt metric that a non-stacking fixed switch could achieve, due to its higher power consumption. Yet the stacking category of switches is highly valued by customers because of it’s inherit capability to scale and expand, a feature not possible with a non-stacking switch. As a total solution, multiple switches stacked together can measure favorably in performance per watt measures when compared against larger single box solutions with comparable port counts. Switches that support stacking, if tested to Energy Star compliance as an individual box, would not perform to its total potential capability. Likewise the results would not be recognized in a performance per watt metric and be unfairly penalized. The recommendation is that stacked switches should be allowed to be tested in the full aggregate configuration.

Modular network equipment brings additional complexity in determining fair criteria for compliance in Energy Star. A modular chassis allows for the insertion of a wide variety of cards or blades that can span a high range of functionality and features, including high port count, aggregate bandwidth, flexible media types and speeds, POE power sourcing, power redundancy, network redundancy, security policies and management. Also the power consumption of each individual card may not always be available or even measurable in some systems. Or if the individual cards did support self reporting of power, any Energy Star test would have to trust the accuracy of the power measurement circuit built into the card by the manufacturer. Furthermore, the power measurement accuracy in cards, because of cost and space to implement, typically will not be acceptable for compliancy testing. Therefore the fairest way to measure a loaded modular chassis is at the AC line cord. Because of the nearly infinite number of permutations of card types, arrangements and slot locations, it would be impractical to qualify every possible configuration. It seems that the only practical solution is to fill the chassis with equivalent cards configured the same and then measure total aggregate performance and power across the entire loaded modular chassis. This would have to be done with each card type for that card to prove compliance. The larger question for modular equipment is: at what hierarchical level should compliancy be applied? Should a card be listed as independently compliant, just the chassis or an entire loaded chassis with cards? With modular network equipment, some specific capabilities could physically reside in either the cards or the chassis hardware, blurring the distinction between the functionality of the chassis and cards. Therefore it seems the only fair approach is to consider an entire loaded chassis for compliancy.

Power over Ethernet (or POE) brings with it additional unique challenges for network equipment qualification criteria for Energy Star. A POE capable switch or router will both consume power from the line cord and is a source of power to third party network devices. The overhead required to power POE devices from a switch is often much higher than the total power required to operate the switch's internal circuitry. For example, a 300 watt switch might require another 1400 watts of power to support POE on all of its ports. Because of the efficiency curves of power supplies, a POE capable switch would use more power than a non-POE switch, even when there is no power being delivered through the ports of the POE switch. Therefore a switch or router capable of delivering POE power should never be categorized or measured against a switch not capable of supporting POE devices. The POE device would be unfairly penalized.

<u>HP's comments on the LNE framework document:</u>

Definition section c.i:

This section states that LNE applies to "wired" ports. Does this imply that non-wired ports, (i.e. fiber-optic ports) are out of scope?

What is a Physical wired port? Does it include ports reserved for management of the device?

Definition section c.ii:

The SNE definition is under review and may not include all equipment with integral wireless capability. The SNE definition calls for equipment with "more" than eleven ports. We think this is a typo, and "less" than eleven ports was intended.

Product types section c.i and c.ii:

The definition of routers and switches is discussed in detail above

Product characteristics:

The characteristic of fixed and modular are orthogonal to managed and unmanaged. In other words a network switch can be fixed and unmanaged.

The criterion for managed network equipment requires refinement. A managed network switch does not necessarily have redundant power supplies. A better definition is that a managed switch can be remotely accessed for configuration, monitoring and telemetry.

Questions for discussion section:

This section asserts that PoE saves energy. In fact if anything the opposite is true. PoE is implemented by customers not to save power, but rather to allow connectivity in places where either conventional AC power is not available, inconvenient or costly to install.

<u>HP has additional comments in regards to the test method draft, listed below:</u>

Section 5.1.A.4:

For UUTs with multiple power supplies, the test method calls for all PSUs to be installed. There is no distinction for whether any of the additional power supplies installed are for supporting redundancy or added capacity. For example, a power supply may be added to a switch to allow for more POE ports, or for redundancy failover in case of power supply or grid failure. If added for power redundancy there may be a loss of efficiency, with each supply only operating at half or less of its capacity. Whereas, if the additional power supply is added for increased capacity, it may be operating closer to its maximum capacity, where it is efficiency could potentially be significantly different. A distinction should be made for any added power supplies, for whether their purpose is for redundancy or added capacity.

Section 5.2.A.b

The section on "dual-group partial mesh topology" requires further refinement. The draft implies that if ports can be functionally partitioned, then traffic is not allowed between ports in the same partition. Many advanced switches have some level of partitioning and can be arbitrarily configured. This requirement suggests that more advanced switches or routers must limit the flow of traffic in the same partition during their test. I don't think this is what the EPA intended.

Section 5.2.B.3

In the proposed draft test method for POE, it was suggested that the ports should be loaded to 90% of their rated power. This is an unnecessary and arbitrary designation. A POE capable switch can provide 100% of its rated power, measured at the switch with 100 meters of cable and the end node loaded to its maximum allowed power consumption. The loss of power in the network cable is built into the POE specification and should be measured as such. To simplify testing, the loss in the cable can be easily simulated with dynamic load boxes, without the necessity of actually having the required length of network cable.

Thank you for your consideration,

Stephen Horvath
Power Engineer Architect
HP Networking
8000 Foothills Blvd MS:5596
Roseville CA, 95747
916-785-5059
steve.horvath@hp.com