General comment regarding the creation of limits

In the section headed "Energy Efficiency Criteria and Test Procedures," it is suggested that during the development of Version 1.0, EPA will explore "the creation of limits based on product type, characteristics, and functionality."

The market for Large Network Equipment is characterized by a large number of distinct products provided by relatively few vendors, compared with other equipment, such as the consumer electronics market.

The differences among LNE products from any given vendor are functional (i.e. there is no market for cosmetic differentiation) and are reflected in differences of energy consumption for these products. Therefore, a detailed categorization will lead to a very large number of categories, each of which will consist of a small number of examples from a few different vendors. Therefore, it will be troublesome to produce a statistically valid set of limits for each of the categories covered. Furthermore, advances in technology have historically produced as much as 40% year-on-year improvement in networking device energy efficiency (for equivalent functions) without the intervention of any energy efficiency promotion. There is a real danger that limits set based on the small number of samples available for any specific category could rapidly become obsolete or could be severely distorted by a mis-categorization.

Alternate approaches may attempt to "linearize" certain differentiating characteristics in order to reduce the number of distinct categories. In certain example cases analyzed by Cisco, such attempts at linearization would inevitably lead to products being passed or failed based on their position on the linear scale (e.g. high port count systems favored over low port count or vice-versa). The existence of such anomalies could inevitably lead to situations where total network energy usage is increased (e.g. by using multiple compliant devices to perform the function of a single device that or by using a compliant larger device where a smaller device would suffice).

In summary:

Cisco recommends that the approach used for Version 1.0 should require a standard representation of power usage compared to performance (often referred to as Power Performance Data sheets) for all Large Networking Equipment (within scope), using standard metrics as a basis. Certain energy efficiency features or characteristics may also be mandated (such as Energy Efficient Ethernet, variable speed fans, 80+ power supplies, etc.). However, the use of pass/fail limits should be avoided.

This approach will benefit the users of Large Network Equipment to a much greater extent than simple pass/fail limits. This also matches the philosophical approach that was intended by the developers of the ATIS TEER standard test metrics. In most cases, the purchase of LNE includes a sophisticated evaluation process that can take into account detailed energy usage information as part of the network design.

Specific comments on Framework Specification

Definitions section

General note regarding IEEE 802 standards

There are a number of instances in this document and in the Test Method where references are made to IEEE 802 standards. Unfortunately the style of reference (although popular) is incorrect.

A published IEEE 802 standard name contains numbers only or numbers and capital letters (e.g. IEEE 802.3 – The Standard for Ethernet; IEEE 802.1D – MAC Bridges). The various functions or devices are specified in clauses of these standards (e.g. IEEE 802.3 Clause 33 – Power over Ethernet; IEEE 802.3 Clause 40 – 1000BASE-T). A specific year of publication may be specified (e.g. IEEE 802.3-2012).

Individual projects within IEEE 802 are designated by adding lower case letters to the name of the standard being amended or revised (e.g. IEEE P802.3az – a project to add Energy Efficient Ethernet to IEEE 802.3). A published amendment takes the same name as the project. However, an amendment is deprecated as soon as the amendment is rolled into a revision of its base standard. For example IEEE 802.3az is no longer a valid standard because the amendment was included in the 2012 revision of IEEE 802.3.

Therefore, there should be no references to IEEE 802 amendments except in the case where the amendment is newer than the latest revision of its base standard. None of the amendments referenced in the Framework Specification and the Test Method are newer than the latest revisions of their standards.

Other Definitions – item ii) – page 3.

Replace “Specified by IEEE 802.3az.” with “Defined in Clause 78 of IEEE 802.3 (originally specified in IEEE 802.3az).”

Other Definitions – item iii) – page 3.

Replace “Currently specified by IEEE 802.3af and IEEE 802.3at.” with “Power over Ethernet is defined in Clause 33 of IEEE 802.3 (originally specified in IEEE 802.3af and IEEE 802.3at).”

Also, add the following after the end of the first sentence:

“The Power over Ethernet specification defines two types of equipment, Type 1 and Type 2. Type 1 Powered Devices may require up to 13.0W; Type 2 Powered Devices may require up to 25.5W.”

Energy Efficiency Criteria and Test Procedures section

Section d) - Specific energy efficient features, item v) – Page 5

Replace “(IEEE 802.3az)” with “(IEEE 802.3 Clause 78)”

Specific comments on Test Method

Semi-modular products

Certain systems accept interchangeable modules that change the port configuration (for example, a module slot may support 8 ports of 1000BASE-T; 16 ports of 1000BASE-T; 8 ports of gigabit SFP; 16 ports of gigabit SFP; 2 ports of 10GBASE-T; or 2 ports of 10 gigabit SFP+). Such systems cannot be described as modular as they are not changing function, but it is unclear how they should be dealt with in this test method. Furthermore it should be noted that these systems are often listed for sale with modules or combinations of modules pre-installed. Should all permutations of module be tested? Or should the manufacturer select an example configuration? It is unclear how such products might be treated in the final program requirements and whether qualification would be attached to the base system or to specific configurations. For the purposes of the test method, Cisco proposes that all combinations of modules should be tested for each base product. Text of the following form should be included in the test method:

"Systems that include the ability to change interface type using modular configuration (e.g. blades, mini-blades, personality modules, etc.) are considered as fixed configuration for the purposes of this test method if the modular configuration changes only the number, medium or speed of interfaces but does not change the networking function of the system. Testing should be performed, and results submitted with all combinations of modules that do not change the fundamental function of the UUT."

It should also be noted that some products include modular power supplies that allow different PSU configurations. The user can select an appropriate power supply configuration for the interface modules installed. In some cases, single or dual power supplies can be used according to the redundancy requirements for the installation. It is important to assess the energy efficiency of a system with the appropriate power supply configuration for the application as a mis-configured system will operate inefficiently. Cisco proposes that the manufacturer should select an appropriate PSU configuration according to the test performed. Text of the following form should be included in the test method:

"Systems that include modular power supply configuration should be tested with a power supply or supplies that are appropriate to the module interface and PoE (if applicable) requirements of the UUT."

3 Definitions Section

Section B) item 3) – Payload

There is often confusion when specifying frame size for Ethernet frames. An Ethernet frame includes overhead for layer-2 addressing and tagging and a typical IP frame contains further overhead for addressing, control and protocol identification. The "frame sizes" quoted refer to IP frame size (i.e. omitting the addresses, the type and the CRC fields). For simplicity, it might be preferable to quote the size of the complete Ethernet frame including the necessary overhead. In this case, the minimum frame size would be 64 bytes (not 46). The maximum frame size would be 1518 bytes for untagged frames; 1522 bytes for tagged frames; and 2000 bytes for envelope frames.

Replace the Payload section with

3) Frame Size: The size of an Ethernet frame, between 64 and 2000 bytes that includes Ethernet framing information and (typically) IP frame payload.

Additionally, add a column to Table 4 that shows Ethernet frame sizes (64, 594, and 1522). The 40 byte payload size (empty TCP/IP packets) requires padding. Note also that 1522 is assumed for the maximum to cater for tagged frames that are more common in large networks.

4 Test Setup Section

Section E) item 4) – Idle-link Period Distribution

It is true that the idle period distribution can have a significant effect on the energy-saving potential of EEE, particularly when the link utilization is relatively high. Modeling performed (by Cisco) to support the development of EEE within IEEE 802.3 used a Poisson distribution of idle period length. However, it is difficult to find test equipment that is capable of controlling the idle periods in such a manner. In most cases, the idle periods will have to be uniformly distributed. It should be recognized that such a distribution will reduce the effectiveness of EEE.

The test method should mandate a uniform idle period distribution for all tests so that comparative results will be valid unless it can be shown that widely available test equipment can support a more sophisticated approach in a normative manner. The additional allowance for EEE should compensate for the reduced efficiency of the test method.

Section F) item 2) – PoE Standards Compliance

Replace “(e.g., 802.3at or 802.3af)” with “(e.g., Type 1 or Type 2)”

5.1 UUT Configuration Section

Section A) item 1) sub-section a) – LNE Requiring Initial Configuration

It is correct to recognize that most LNE requires configuration in order to become operable. The configuration may include features that have a significant effect on energy usage. For valid comparisons, it is important that the configuration used should remain unchanged for all tests conducted. Furthermore, the configuration should become part of any published results. If the final program requirements include publication of test results (e.g. as a power –performance data sheet) then the configuration should be included in that publication. In some cases, an equipment vendor may wish to publish multiple sets of results to demonstrate the effects of specific configurable features.

Add to the clause: A single configuration shall be used for all tests conducted on the UUT without alteration between tests.

Section A) item 4) sub-section a) – UUTs with Multiple PSUs

It should be recognized that some systems with multiple PSUs will support dual a/c inputs while others will support a single a/c input and a separate d/c for backup power. Also, there is no reason why the measurement methods should be different between a system with one PSU and a system with multiple PSUs, therefore including the power overhead of a PDU would be inappropriate.

Replace the last sentence with, “The input power shall be measured separately for each PSU and the total input power summed for the UUT.”

Some systems support modular power supply configuration in order to cater for different configurations of interface modules. The power supply or supplies for such systems should be configured appropriately according requirements of the interfaces selected in order to operate efficiently. In most cases the manufacturer will recommend a power supply based on the interface module configuration.

Add a sentence at the end of the section: “If the UUT supports modular power supply configuration then the PSU or PSUs recommended by the manufacturer according to the interface configuration should be used for all testing.”

Section A) item 5) sub-section a) item ii) – Half port testing

In practical installations, there is a significant difference between usage of uplink ports and edge ports. It may be speculated that some proportion of edge ports (for access network devices) may be unconnected or otherwise unused. However, it is very rare to encounter uplink ports (or ports on non-access networking devices) to be unconnected. If it is deemed necessary to assess the energy efficiency of a device with half of its ports unconnected, then this should be limited to downlink ports. Furthermore, the ATIS TEER test methodology is strongly dependent on the uplink bandwidth (for edge networking devices) therefore, testing with half of the uplink ports unconnected would significantly skew the results.

Replace the second sentence with, “If the UUT has identified uplink and downlink ports, then half of the downlink ports shall be connected, but all of the uplink ports shall be connected; if the UUT has other identifiable groups of ports, then half of each group of ports shall be connected to the Test Equipment.”

Section A) item 5) sub-section c) – Energy Efficient Ethernet (EEE)

Replace “If the UUT has ports which are compliant with the IEEE 201 802.3az standard, which provides EEE functionality,” with:

“If the UUT has ports that support Energy Efficient Ethernet operation (as defined in IEEE 802.3 Clause 78),”

Section A) item 5) sub-section e) – Physical Interface Requirements

It should be noted that it is very rare for systems with pluggable module interfaces to be shipped with specific pluggable modules. Also, it should be noted that for some types of module, the copper interface will represent the highest power whereas for others it will represent the lowest. For example, a 10Gb/s XENPAK module interface may support a first generation 10GBASE-T module with power of 12W whereas a 10Gb/s SFP+ module interface may support an SFP direct attach module, with power less than 1W, as the only copper interface option. If there is an intention to evaluate the energy efficiency of the UUT, then the effect of pluggable modules should be minimized. Similarly, for comparative purposes, the difference between modules used for different systems should be minimized. It may be difficult to make a specification for module configuration for testing that will work for every system tested, therefore it is recommended that the manufacturer should be allowed the choice of the lowest power module (at the highest speed) for each pluggable module interface.

Replace the last sentence of item 1 with, “If no compatible pluggable module is included as a standard component with the UUT, then any module that supports the highest link rate supported by the UUT for that interface shall be used. The module that supports the lowest energy usage may be used”

Section B) items 2) and 3) – Maximum Supported Load

The IEEE 802.3 standard specifies only the maximum supported load per port according to the Type and Class advertised during the classification process. Specific systems may have further restrictions (within the IEEE specification or beyond it) and may use LLDP or other proprietary methods to express those restrictions. Additionally, many power sourcing systems may have a system-wide limit on aggregate power availability that could limit the ability of the system to connect loads to all ports (or all ports of a set) simultaneously. Systems may manage the system-wide power budget using device classes (detected during the classification process); LLDP management; or proprietary methods. It would also be permissible, according to the standard, for a system to ignore power management altogether.

It appears that the DoE is attempting to specify a generalized mechanism for finding and applying the maximum load for a system divided equally amongst all ports. Given the variation amongst systems, it will be difficult to generalize such a method. Furthermore, it should be noted that using homogenous loads attached to all of the ports may not reach the maximum power capability of a system.

Cisco recommends that such complex testing for PoE systems will be unnecessarily onerous for manufacturers and will not produce significantly more benefit than a simple approach using certified PSU requirements (based on 80+ specifications for example). The following methods may be used to find the maximum homogenous loads for various types of systems. It should be noted that other methods may exist. This pre-supposes that maximum homogenous loads are preferred without consideration of alternative approaches.

Systems with unrestricted capability

If the system offers the maximum power for Type 1 or Type 2 devices simultaneously on all ports, then all ports should be populated with the maximum load according to type.

Corollary – some systems offer capabilities beyond those described in the standard (either for Type 1 or for Type 2). The DoE should consider whether to account for such proprietary capabilities.

Systems with fine-grained management

If the UUT supports power management using LLDP then it may be able to benefit from fine-grained power management. All ports of the UUT should be populated with Type 1 or Type 2 loads as appropriate with each load immediately using LLDP to negotiate the smallest power allocation that it can demand. Once all of the ports are populated, then all of the ports can demand increases in power using small increments until the demands are refused. The DoE should consider whether to allow proprietary methods of dynamic power management as an alternative to LLDP.

Corollary – some systems may be limited according to groups of ports. The DoE should consider whether to account for such groupings.

Systems using class-based management

If the UUT only supports power management using PoE classes then two approaches may be used. If it is necessary to test with all ports identically loaded, then the UUT should be populated with identical loads indicating Class 1. If these are successfully allocated power, then the procedure can be repeated using Class 2, Class 3 and Class 4 (for Type 2) until the UUT refuses to allocate power. If it is not necessary to test with all ports identically loaded, then each port can be incremented separately until the maximum net load is reached. Once more, consideration may be given to systems that have limitations based on maximum power for a group of ports.

Systems using empirical management

Some systems manage PoE loads based on aggregate measured power instead of negotiated or classified power requirement levels. For such systems, the UUT should be populated with identical loads indicating Class 0 or Class 4 as appropriate, but the actual power drawn should be reduced to a minimal amount. After all ports are connected, the power drawn can be increased until the UUT denies further increases; disconnects a port; in in some other way indicates that the limit has been reached.

Section B) item 4) – Cabling Requirements

Replace “802.3at or 802.3af” with Type 1 or Type 2.”

5.2 UUT Preparation Section

Section B) – general comment on the PoE load test

In Cisco's experience, it is very rare for PoE equipment to operate for long periods very close to the maximum power capability of the system. Although the use of intelligent power management allows a much lower level of over-provisioning compared to (for example) domestic or enterprise AC power distribution, there are a number of factors that keep the typical usage below 50% of the maximum availability. Firstly, most installations are designed for "future-proofing." The financial (and environmental) cost of replacing equipment as requirements grow can be very large. Therefore, equipment capabilities (including power) are often specified to account for expected growth in demand over several years. Secondly, there is substantial time-of-day, time-of-week and time-of-year variation in demand. The peaks tend to be relatively short compared to the overall cycle and thus the mean power requirement is lower than the maximum and the energy footprint is dominated by efficiency in the typical state. Lastly, advances in energy efficiency of load devices have increased the number of energy saving modes and improved the power scalability of those devices. This has the effect of further reducing the typical power requirement while increasing the ratio between the typical and maximum demand.

It is recommended that Power over Ethernet systems should be evaluated for efficient operation at load points that are substantially below their maximum capabilities. This evaluation is strongly analogous to the approach used for operation network testing in the ATIS TEER methodology. Cisco recommends that PoE systems should be evaluated based on the efficiency at <<25% of the maximum tested capability.

Section B) item 3) – 90% of maximum load (and Note)

The standard defines the maximum power availability in the PSE and the maximum power requirement in the PD in a manner that accounts for the worst case of cable losses. Furthermore, if the test method described in section 5.1 B) uses a practical evaluation of the maximum power (whether defined at the PSE or at the PD) then it is unnecessary to apply a margin to allow for cable losses (that can be assumed unchanged from the evaluation process). However, if the DoE has another reason for keeping this 10% margin calculation, there is no objection to keeping it.

7 References Section

Items D), E) and F) – IEEE 802.3 references

Items D), E), and F) are all deprecated. These should be replaced with:

D) IEEE 802.3-2012 Standard for Ethernet: Clause 33 Data Terminal Equipment (DTE) Power via Media Dependent Interface (MDI); Clause 78 Energy-Efficient Ethernet (EEE).

Cisco responses to LNE Framework, Test Method and Discussion Document