To: Nils F. Nissen and Lutz Stobbe, Fraunhofer IZM
Shailendra Mudgal and Kurt Muehmel, Bio Intelligence Service
Stephan Kolb, DG ENER/C3, European Commission

From: Anita Eide, Director of European Programs

Cc: Bruce Nordman, LBNL
Marie Baton and Eric Gibbs, CLASP

Date: 13 May, 2011

Re: Comments on DG ENER/Lot 26 Preparatory Study on Network Standby Losses draft Task 8 report

These comments, prepared for CLASP by Bruce Nordman of the Lawrence Berkeley National Laboratory, address the content of the draft Task 8 report of the DG ENER/Lot 26 Preparatory Study on Network Standby Losses, but they necessarily derive from and build on the CLASP comments on the Task 1-7 reports. At this time CLASP would like to thank the European Commission (EC) and its consultants for the well-organised stakeholder meeting on February 14 and the process generally. We look forward to exchanging comments and ideas as this topic and related ones evolve and proceed.

The draft Task 8 report includes a number of good observations, principles, and recommendations. In other cases, it takes directions that seem unlikely to save much energy, and ignores some that can. To a great degree, this is due to the overall construction of traditional energy policy. The most promising directions for network issues are not direct extensions of how things are done today in energy policy, but rather take new and unfamiliar approaches. Network technology brings so many new issues to the fore that making small adaptations to traditional policy-making will not be sufficient. Rather, it is necessary to consider a range of policies which depart (substantially in some cases) from current practices. The draft Task 8 report and the earlier reports reference some new directions (e.g. the role of technology standards and user interfaces) but then do not bring them into the policy recommendations.

It is essential to simultaneously consider what can be done in the near term as well as how to craft a longer-term process for more fundamental progress in energy policy in terms of regulation of network standby losses. The inclusion of a Tier 1 and Tier 2 in the recommendations recognizes that as time goes on, there is more opportunity for technology and design to evolve to enable lower power levels and higher performance. Addressing power levels and delay times is certainly a core part of ensuring efficient network-connected low-power modes, but by themselves they are not sufficient.

Horizontal Policy

A premise of the Lot 26 process, which we agree with, is that there are elements of policy regarding network standby that can be applied across many product types, making widely effective policy possible, and making the burden of the regulations manageable for both industry and policy-makers, hence generally supporting the argument for the horizontal approach.

A fundamental question then for policy in this area is *how* to apply the horizontal approach in this instance. For the general standby topic, devices are providing no function at all, or only a very limited set of simple functions. This enabled construction of a singular and cohesive policy approach.

It made sense to consider whether this approach could be transferred directly to the topic of network standby with similar success. However, this should have been a question that was asked at the outset, and not something that was assumed before analysis was undertaken. At this stage, we should be able to answer this question. The consultants made every effort to make a horizontal approach work, and are to be commended for that, but several points are apparent:

While the policy was intended to be universal, *all* of the 21 product groups in Table 8-4 are electronic products. This is in contrast to the wide range of products covered by EC 1275/2008. It is not apparent that the recommendations are suitable for other end uses (see the "Sectoral" section below).

While a single set of requirements was the goal, three different sets were developed, with some process used to allocate devices to these three groups.

So, it seems a horizontal approach of the type used in EC 1275/2008 was not possible in this case. It would be helpful to explicitly recognise this, and then allow the policy structure to adapt to the needs of the topic. A hybrid approach — the term "sectorally horizontal" seems useful — applies horizontal requirements within defined groups of product types.

One of the core differences between network standby modes and the standby and off topics dealt with earlier is that devices in network standby are performing potentially significant and non-trivial functions — just not a primary function of the device. The draft Task 8 report acknowledges this in a number of ways, but does not take the next step of assigning power/energy to these functions. To do so would require identifying, measuring, and rewarding these functions, and note critical technology standards relevant to their efficient operation. Worth quoting is a bullet item from section 8.1.1.2 (emphasis added), "Individual configurations and the resulting performance of a product in conjunction with actual services and service requirements (QoS) however **may demand specific energy conditions for providing certain functionality**." That is, to a great degree it is functions that require power in low-power modes, so functions need to be a core part of the policy structure and power allowances.

Latency Distinctions

With functions apparently discarded as an organising principle for the policy analysis and recommendations, latency ("delay") was chosen. Latency certainly is important in networks. Both network protocols and human beings have limits on how much delay is tolerable in different circumstances.

The primacy of the distinction in network availability by latency seems unconvincing. Other aspects of connectivity — such as type or number as well as protocols involved — seem more important.

In the Task 7 report, high availability (HiNA) was defined as responsiveness within milliseconds; this seemed to be the capability that a product could transparently accomplish within its idle mode. So, this would not seem to be something that would be visible in an energy test procedure. A specific requirement for that HiNA would essentially be just an idle or ready power requirement. It would not be a low power mode because there would be no reduced capability. Responsiveness this quick is not perceivable by people, and is not significant, but is separate from the topic of network standby. This all raises the question of how the values were determined and how suitable they are for different applications.

It would have been helpful to have an explanation of the change to 1 s, and what products fell into the gap between those two values and so no longer can attain a HiNA mode.

Some products can make good use of low network availability modes (LoNA), most notably portable personal computers. However, since ordinary network protocols (those that use the Internet Protocol) expect responsiveness much faster than 10 seconds, this is not a mode that can support ordinary network connectivity. It can support Wake-on-LAN and similar niche applications, but that is a very limited notion of connectivity. There would seem to be few devices that could use such a mode. For most products, such a length of time to availability would be unacceptable to people and so not a mode they would use in ordinary use. Some set-top boxes have such a mode, but it is unlikely these are much used. With the combination of few devices implementing this mode and it not being widely used, it can be expected to have only a limited role in saving energy.

The MeNA mode is where policy can and should be focused, as it leverages the largest amount of potential savings. There are distinctions to be made within network standby, but these are best put to the physical and application layer protocol specifics, not latency. The 1 second new delineation between HiNA and MeNA seems arbitrary.

Omissions

There are several key omissions from the recommendations that reflect an avoidance of mentioning specific technologies (whether they are hardware or software technologies). There is a sense of purity in being generic, but much energy savings potential is wrapped up in particular technologies that are based in industry standards. These omissions are:

Power management is frequently (and correctly) noted as critical, but any sort of definition or mention of specific capabilities is avoided. Without specifics, it would be easy for a manufacturer to point to limited or ineffective power management capability to comply, more on power management below.

“Network” itself is never defined. It is important to clarify the difference between network and data connections (though the policy should cover both).

Network interfaces require power to operate, and the amount of power depends on their number and type. Ignoring these aspects makes it difficult or impossible for devices with many or highly capable interfaces to meet the proposed standards, and too easy for those with few or low-performing interfaces. Section 8.1.1.3 notes an example of a power-hungry interface. In the short term it is necessary to accommodate the real needs of devices that use these as the manufacturer alone lacks the ability to change these. In the long term, standards organizations and relevant manufacturers need to be put on notice that the technology needs to perform better. In addition, the public sector can help attain the improvement.

Other specific network technologies are not mentioned. Oddly, the word “Internet” does not appear.

The importance of the user interface is noted several times, but this is always left in the abstract and never brought to specific items or standards. Section 8.1.1.4 is a good solid one solely on the user interface as background, but it is then left completely out of the recommendations for policy to implement.

Since the importance of standards is not made clear, and because other aspects of energy use lack a similar construct, there is no concept of a “technology strategy” described that can help create technologies which enable or require saving energy (and help avoid creation of those which prohibit or discourage it). The evolution of technology is not fixed; policy can affect and direct it. Network standby is a policy topic well-suited to tackling this.

Power Management

The draft Task 8 report frequently notes the important role that power management has in reducing energy use associated with network standby.

As an example, "The basic improvement strategy is to power down system components or functions which are not actively required by the user (at the moment)" (section 8.1.1.2). This is certainly a good principle, but needs to be made concrete with specifics so that power management is effective. Policy also needs to account for real barriers to doing this, be they operating system platforms or software, technology standards, user interface issues, or cost.

It is important to recognise that there are at least two ways to divide up power management that have significant implications for policy. One is power management that operates *within* a power mode (reducing the modal power level), versus that which involves switching *among* power modes (e.g. on to sleep). A second distinction is power management that can be accomplished by a single company (manufacturer) in product design versus that which requires coordination and cooperation among companies, typically via a technology standard.

"Industry should be encouraged to continue the improvement of energy performance and implement even more ambitious power management schemes" asserts section 8.1.2.3. For power levels, very specific requirements are made and they are mandatory. However, for power management, the requirements are general and not mandatory. A policy is clearly serious when it is specific and elements are required.

Power management is not a binary feature that is present or not, but rather can take on many different forms.

Sectoral Horizontal Policy

It seems doubtful that a completely horizontal approach to network standby, as has worked for "regular" standby, is feasible. There are such substantial differences between products that a limit which is reasonable for all products will be much too high for some. A "hybrid" horizontal approach applies requirements horizontally as far as possible – but no more. This would divide the relevant product types into a number of groups with substantially common characteristics. A possible division could be:

- Appliances (and other devices with limited network functionality)
- Audio/video devices (i.e. consumer electronics), including displays
- IT products (personal computers, printers, etc)
- Network equipment

CLASP thinks it would be quite productive to revisit the data and analysis from the Lot 26 project in the light of this division and a focus only on medium network availability to come to conclusions about workable and effective policy options.

Also, while there is value in having a single power value for each of these groups for simplicity of explanation, there is good reason to modulate the power value based on specifics of the product type and network connectivity, that can be an effective way to better tailor the policy instrument to what products can achieve. That is, functional adders can enable lower overall power values than a single fixed value can. The Broadband Code of Conduct contains such information to draw on for this purpose.

Test Procedures

To implement any sort of power limit, there needs to be test procedure infrastructure to verify that the limits are being met. There is some information in the Code of Conduct, and even more in the draft Energy Star test procedure for Small Network Equipment (and smaller amounts of content in other Energy Star procedures). Creating this material may be outside the scope of this process, but its need should most definitely be highlighted. Development should be harmonised globally as much as possible.

That a section on test procedure content for this topic was included (8.1.2.5) implicitly recognises its importance. In addition, it is clear that it will take some time to create an international standard for this purpose, and there will be considerable learning from experience over the coming few years, so that it will be necessary to use simplified and interim testing for the first regulation of this type. The end of this section references potential difficulties in testing each interface separately; this is an obvious place in which referencing technology standards which require certain types of behaviour can greatly simplify the burden on public policy for testing and verification.

"The study concludes that it is not required to provide/define an exhaustive list of functions, which are "on" or "off" in certain product state." (Section 8.1.1.2). This is certainly reasonable since which functions are in which state in each product mode depends critically on the product type for core functions, and by model for more peripheral functions.

Relation to Vertical Requirements

One issue that we continue to be unclear on from the recommendations is how the horizontal proposed requirements do or do not relate to the vertical requirements in place for some products. It would be helpful to separate the quantification between vertical and non-vertical products. That is, how much energy use and potential savings are there from creating network standby requirements for products which currently have none, versus the amounts obtained by reducing the allowed power for products which already have requirements (or whose savings are substantially or entirely accounted for by the vertical requirements). It is also important to know how the recommended levels compare to those already in place for relevant products.

Role of Technology Standards

Many aspects of network standby are driven by technology standards; these are not stationary but instead evolve over time. The difficulty of assessing the role of BNAT (best not yet available technology) 8.1.1.3 is noted. The biggest uncertainty is the appearance of network technologies that affect energy use, as these could drive energy use either up or down. This is also the area of technology that public policy could play the largest role in. The subsequent bullet point mentions the importance of standards, but does not suggest how public policy can engage the topic.

"Tier 2 … could require inter-market collaboration and standardisation efforts" (section 8.1.2.4) but how those are to be made to happen is never addressed.

Possible Additional Requirements

Power levels are not the only way to address and regulate this energy use in the near term. A suite of other requirements could be part of the requirements without changing the basic policy structure. Examples include:

- Creation of a power (or energy for annual consumption) network adder for appliances (for those that already have a vertical requirement). This could vary by the type of network interface, but should be the same for all types of appliances.
- Require Energy Efficient Ethernet (IEEE 802.3az) for any device with an applicable Ethernet interface.
- Define some basic auto-power down behaviours for audio/video devices. This would be an extension of the requirements for TVs (e.g. to power down after an extended time of no user interactivity).
- Reference IEEE 1621 for user interfaces for power control, at least for electronics, (though some of it applies to other devices as well). This applies to both status indications, as well as control mechanisms.

Whatever is done in the near term, a future regulatory measure needs to be constructed around future characteristics

Task 8 Details

8. The objective is a “clear recommendation” which we think is a good and ambitious goal for this topic.

8.1.1.3 The concept of "resume time to application" is a good one, since it reflects actual usefulness of the product and user experience. However, it could pose challenges for test procedures to the extent that the resuming to application is dependent on the interaction with other devices on the network, so that it may be necessary to specify characteristics of the hardware and/or software of those systems. It also may be dependent on the particular version of the application, and not be strictly comparable across platforms (e.g. the Windows version of an application might be much better than the Mac version, or vice-versa).

8.1.2.7 The report asserts that a general labelling effort on this topic is not a viable approach. This in our view is a sound conclusion.

8.1.1.3 The third bullet on the cost of "performance improvement" is not clear.

8.1.2.1 For CPE equipment, it seems unnecessary to call many of these "Home" devices since they can readily be used in other building types. Energy Star refers to such devices as "Small Network Equipment" to not make or imply such a limitation.

The draft report (section 8.1.1.1) states that it is difficult to anticipate future products. To some degree this is true, but the functions these products have are much better known, and readily addressable via policy.

8.1.1.1 EC 1275/2000 “excludes equipment that is not dependent on energy input from the mains power source”. The next paragraph refers to low-voltage DC powered products, and adding these is certainly appropriate. However, there should be clarity of what is proposed – does it also include battery-only devices? Mains-powered devices when running on battery? For low-voltage DC, there is a need for standard test procedures for these devices, including whether the DC power should be measured directly (and how evaluated), or whether AC power from a suitable source (e.g. USB powered up or PoE Injector) is measured.

8.1.1.1 The lack of coverage of “home automation and ‘white goods’” is reasonable given the small market for these today and in the very near term. Engaging the technology development aspect of these topics is much more important than regulation at this time.

8.1.2.3 "A defined networked standby mode” is referenced. It is difficult to imagine how such a definition would be done except by referencing the specific functions supported (or not required to be supported) in the particular mode.