

Energy Efficient Ethernet: Outstanding Questions

Bruce Nordman (LBNL), BNordman@LBL.gov, 510-486-7089

March 19, 2007

This document is intended to track what is known and what is not (yet) known by the IEEE EEE study group. This is intended to reflect the general consensus of the EEE Study Group but does NOT do so in any official capacity.

What we knew at start of study group

- Energy consumption of IT equipment an increasing concern.
- (added) EEE is good for network industry, IT mfrs, IT users (esp. data centers), economy, and planet
- Goal is to eventually have all new NICs EEE-capable.
- (added) Energy efficiency community (e.g. Energy Star) ready to help
- Typical link utilization very low on average (not a problem — an opportunity)
- Aim is for EEE to be suitable for great majority of links (e.g. >99%) but able to be turned off if necessary.
- Time needed to switch data rates should be low (e.g. 1 ms).

Outstanding questions

Control Policies (CP)

Know

- CP is outside scope of 802 / EEE.
- Existence proof of ≥ 1 good CP needed to move EEE forward.
- Helpful to document good CPs somewhere.

Questions

- Do both NICs on a link need to use same CP?
- Can CP elegance be something manufacturers compete on?
- Should EEE NICs exchange information about policies and perhaps parameters? (e.g. delay times, buffer sizes)
- Where should CP examples be documented? (e.g. Informative annex to EEE standard, Ethernet Alliance, Academic literature, other standards organization, e.g. IETF)

Traces / Use patterns

Know

- Burst pattern time scales of possible interest
 - Sub-second
 - Seconds to 100s of seconds
 - 1000s of seconds (hours) to days

Questions

- Do we need more sample traces (beyond those in presentations from Ken/January and Mike/March) as references for designing EEE?
- Do we need other information about current link use patterns?
- How about usage data for 10 Gb/s links?

Data Rates

Know

- Asymmetric data rates raise questions about signaling, and in most applications would have only slight affect on energy savings, so not a priority for consideration at this time.

- Power difference in PHY between 100 Mb/s and 10 Mb/s is too small to be concerned with for energy savings, however 10 Mb/s may still be of interest for other reasons (e.g., as an implementation of 0 Mb/s)
- Zero Mb/s deserves consideration as a new data rate for EEE PHYs.
- **(revised)** We should not consider PHY data rates other than powers of 10.
- **(NEW)** There is no need to consider a rate other than 1 Gb/s as the low rate for 10 Gb/s
- **(NEW)** There is no need to consider data rates other than powers of 10 Mb/s for non-PHY components.

Questions

- Are there any non-PHY reasons to facilitate both 100 Mb/s and 10 Mb/s for 1G links?
- **(NEW)** Should 10 Mb be the basis for the new 0 MB state?

Transitions

Know

- Candidate signaling mechanisms include: LLDP, OAM, MAC, LLDP (see Frazier January slides).
- Some parameters could usefully be stored to minimize transition time.

Questions

- **(revised)** Where is the “threshold of pain” (for switching time) for RPS? (How do we know if the time needed is “quick enough” for great majority of applications)
- Is switch time fixed or potentially negotiated (potentially influencing control policies)?
- Other than parameter storage, how can we minimize transition time?
- **(NEW)** Is the existing term “settling time” adequate for our needs?

Other Protocols / Layers

Know

- **(NEW)** EEE must not cause perturbations to spanning tree protocol
- **(NEW)** EEE creates an opportunity for energy savings for link aggregation, but we don't want deal with it in the study group
- We should confer with the PoE group on signaling.

Questions

- What interaction could there be with higher-layer protocols?
- Should link rate switching ever be advertised? (e.g. so that links in series can all shift up at same time)?
- Do upper layers ever need to know that link rates are changeable / changing?
- **(revised)** Should EEE be targeted to end host connections?

PHY types

Know

- **(NEW)** EEE should be applicable to any multi-rate PHY, whether or not they are initially specified.

Questions

- **(NEW)** Do we want to consider mechanisms to reduce the data rate of a link (maintaining the same signaling rate) to get other (non-PHY) savings?

- For 10G, is it OK if the EEE 1G PHY is not identical with today's 1G PHY since this state will only occur with 10G EEE PHYs?

Outreach

Know

- We got a flurry of Internet coverage right after January meeting and the high level of press interest appears to be continuing
- Web resources: IEEE/EEE: grouper.ieee.org/groups/802/3/eee_study/
LBNL: efficientnetworks.lbl.gov
USF: www.csee.usf.edu/~christen/energy/main.html

Questions

- Do we want more information about other link technologies that change link rates? (ADSL2+, 1394, ...)
- Any other outreach (companies, countries) of importance to inform or engage?

Miscellaneous

Know

- **(NEW)** There are a number of issues in this document that we want to enable solutions for — and want to note in the PAR submittal — but we don't necessarily want to solve them in this study group.

Questions

- **(NEW)** Might the biggest benefit of EEE be that it will enable bigger energy savings system wide (e.g., within a switch or router by enabling power down of internal parts)?
- **(NEW)** Can we assume that it is not necessary to have hardware experimentation?

(NEW) For EEE Task Force Consideration (so not Study Group)

Questions

- Details and timing for resynchronization at a new rate