

IEEE802.3az task force meeting, IEEE 802 Plenary, Orlando, FL

Tuesday, March, 18, 2008

Attendees

Name	Employer	Affiliation
Mike Bennett	LBNL	LBNL
George Zimmerman	Solarflare	Solarflare
Robert Grow	Intel	Intel
Ahmad Nouri	Broadcom	Broadcom
Hugh Barrass	Cisco	Cisco
Adam Sayer	UNH-IOL	UNH-IOL
Robert Busse	Transition Networks	Transition Networks
Khurram Kazi	SMSC	SMSC
Jim Millar	Force10 Networks	Force10 Networks
Mario Traeber	Infineon	Infineon
Joseph Chou	Realtek	Realtek
Angus (Shang-Ta) Lee	Realtek	Realtek
Dimitry Taich	Teranetics	Teranetics
Aviad Wertheimer	Intel	Intel
Mandeep Chadha	Vitesse	Vitesse
Robert Hays	Intel	Intel
Satoshi Obara	Fujitsu	Fujitsu
Yukihiro Fujimoto	NTT	NTT
Paul Gyugyi	NVIDIA	NVIDIA
Richard Bowers	Micrel	Micrel
Geoff Thompson	Nortel	Nortel
Velu Pillai	Broadcom	Broadcom
Carrie Higbie	Siemon	Siemon
Jeff Lynch	IBM	IBM
Mike Grimwood	Broadcom	Broadcom
Ozdal Barkan	Marvell	Marvell
James Tsai	Intel	Intel
Wiren Perera	Plato Networks	Plato Networks
Jeff Lapak	UNH-IOL	UNH-IOL
Adam Healy	LSI	LSI
Dan Dove	Procurve	Procurve
Kory Sefidvash	Broadcom	Broadcom
Wael Diab	Broadcom	Broadcom
Steve Carlson	HSD	HSD
Brad Booth	AMCC	AMCC
David Koenen	HP	HP
Albert Kuo	Realtek	Realtek
Michelle Gong	Intel	Intel
Charles Cook	Qwest	Qwest
Scott Powel	Broadcom	Broadcom
Howard Frazier	Broadcom	Broadcom
David Law	3Com	3Com
Shimon Mueller	Sun	Sun
Yong Kim	Broadcom	Broadcom

- Secretary appointed: Ahmad Nouri
- Introductions
- The chair reminded people to sign the sign-in sheet
- Minutes from last meeting approved by voice
- The chair displayed the IEEE-SA Standards Board Bylaws on Patents in Standards slides, and read slide #1 aloud.
- The chair called for patents
- Robert Hays: Intel is in the process of submitting LOA.

There was a comment that the chair should include more goals in the agenda, such as adopting a common set of terms for use in our proposals. It was also noted that such a proposal would be presented at this meeting.

Presentation:

Terminology Proposal for LPI EEE (Rob Hays) – hays_01_0308.pdf

Discussion:

Like to see average power to differentiate between active (data transfer) and idle time. Operating power is instantaneous measurement. The problem is energy efficiency should be calculated over time. What matters most is the energy consumption. We are only doing energy efficiency for PHYs and proposing energy efficiency terms does not make sense. There are multiple customers for this terminology. \$/year or how much cost to operate. End users want KWh per year. For 10BASE-T this does not work. Any of the new timing parameters may or may not exist based on the link speed (10/100/1000). This does not apply to 10BASE-T.

Presentation:

Another View of Low Power Idle / Idle Toggle (Geoff Thompson) - thompson_01_0308.pdf

Discussion:

Regarding refresh, 10BASE-T has refresh to indicate link still active. There is a separate requirement for link bit. No need to eliminate link bit. For 10BASE-T, link pulses still come. For EEE, should we shift the link pulse responsibility in the MAC. Normal link pulse is not part of EEE implementation; this is the standard PHY requirement. Putting preamble in the mac was a mistake before. We should not change what we have now. Not sure if we should modify MAC to support EEE. However may be reasonable for MAC (100BASE-TX) to assert extended preamble during low power mode, EEE LPI time. This model may work but there may be refresh symbols without packets for LPI. Except for 10BASE-T, newer PHYs have their internal sync keeping. One thing is missing is the system level energy saving. We should move forward with simplest models. If any change to the basic models,

need detail discussion why a change is needed. Having the preamble time as a negotiable value, it is very important to allow the rest of the system (memory and CPU) to take advantage of the time off period. MAC control level spec would allow more time in for system requirements. This is a MAC solution at high level. If we adopt LPI, and if we want to enable MAC, we need to determine where to buffer the data. In LPI proposal, if we don't have a way to tell above the MAC that data is being deferred, it will require huge buffers.

Presentation:

Low-Power Idle based EEE for 100Base-TX (Joseph Chou) chou_01_0308.pdf

Discussion:

Two people requested to be added to the supporters list. Don't think should use normal idle code for wake up. Should use a different code for LPI. What happens if someone unplugs the cable (related to the state diagram)? It is considered in the state diagram (100msec then drop back to link detect). In slide 7, the proposed transition seems straight-forward and only uses one extra symbol. In 100BASE-TX, curious if need to maintain synch. No reason to maintain sync between octets. Basic 4B5B takes care of alignment. Don't think it's necessary to maintain any refresh for 100BASE-TX. The repeater may need sync. Yes, may only need to keep link integrity (link beat pulse) to maintain link up state. The link pulse for 100BASE-TX may get confused with 10BASE-T link pulse. In slide 9, During the 6ms wake up, the receiver is not in sync. Yes, we need to assume everything re-sync. For IEEE 1588, need sync to maintain between the link partners. There is no requirement for this. This will be discussed in more detail on Wednesday in the 802.1 joint meeting. How long does it take for MLT3 to re-sync? Answer: 23bits to re-sync. Local and remote Frequency offset = 600ppm. Do you mean between local and remote there is 600ppm offset? Answer: the local could be +300ppm and remote could be -300 ppm offset. This seems like this is extremely conservative. No need for legacy compatibility requirement. This will be a new PHY. You have over conservative assumption, and we can tweak this. This is great starting point.

Presentation:

1000Base-T LPI update (Asymmetric sleep operation) (Adam Healey) healey_01_0308.pdf

Discussion:

What is the power saving for (Only Slave sleeps)? Answer: minimal power saving in this mode. Line state undefined? Is this going to be robust enough to identify sleep state and transition to active? What is the wake up time? Slave wants to wake up, it sends wake up request but has to wait before the other side (master) to wake up and send active link. Two people requested to be added to the supporters list.

Presentation:

Low-Power Idle for 1000Base-T (Mario Träber) traeber_0308.pdf

Discussion:

Upper bound for Tq should be capped at 16. Mario Agreed. Slide 18 (Conclusion) disagree on the “Sleep, refresh and Wake state” statement. No reason to change the line code for different states. Need to prioritize robustness. Slide 13: If the system allows 800ps drift, drift of echo is much longer than drift on phase. How long does it take to lock scrambler and echo cancellers? My idea is as before, the PHY is loose phase to wake up. Regarding asymmetric mode: need to be mindful that based on the how fast we need to send refresh rate to avoid frequency drift, may not provide any power saving. There is an issue about Phase drift. We need to consider the phase drift in the new standard. Also, the objective is to never lose bit, but if we do lose sync, need at least 63 bits. Presenter disagreed: It may take longer. De-skew may not be a problem. It takes 64 bit time to do de-scramble, assuming the structure does not change too much during low power mode.

Presentation:

10GBase-T Active/Low-Power Idle Toggling (George Zimmerman) zimmerman_01_0308.pdf

Discussion:

Don't think changing restart sequence would violate objectives. The change is in line with objectives. The prior presentation only occurs in situations which it may violate our objectives. If we were violating our objectives then we would want to use the new slide. What we want to use programmable and what fixed? Not proposing all PHYs should support all capabilities. The Standard does not say all capabilities need to be supported. The market value dictates not to require all capabilities in the standard. We should make it a requirement in the standards with a minimal set and expect all PHYs to support this set. 50 to 1 or 100 to 1 should provide good power savings. 2 is the maximal number. For 10G need a fast wake up, which is 10 to 1.

Presentation:

Negotiation Proposal For LPI EEE (Aviad Wertheimer) wertheimer_01_0308.pdf

Discussion:

Why need two Auto-negotiation mechanisms? Answer: There are MAC control parameters and PHY related Parameters. We can do PHY parameters during AN, then after link up, we can do MAC frame AN. Most Fiber PHYs don't support AN.

Presentation:

Enhancements to the lower power Idle Mode (Dimitry Taich) taich_01_0308.pdf

Discussion:

Some like the idea of K & M proposal. There is a need within 10G to have low power mode and fast wake up mode. By combining two modes into one, we can have one mode with low power and fast wake up time. But don't think will change things. To save a lot of power, it will take longer to wake up. All parameters negotiated during AN time. If MAC wants to take advantage it will need to use the AN time. If day time needed fast wake up time then it can restart AN, but at night may negotiate longer wake up time. It seems to be more flexible to use mac control frames for negotiate the wake up time. This demonstrates one of the capabilities.

Presentation:

Technical Open Issues with LPI (Powell/Diab/Frazier/Thaler) diab_01_0308.pdf

Discussion:

Very interesting if you can provide additional info on traffic patterns. Even using Subset PHY will not solve the traffic patten issue. Need to get more traffic pattern and example to present. The point is in the lower range of traffic, such as AVB, it will not result in going to low power mode. Buffering and burst will work in some application but not all applications. We need to look at all applications and markets to make sure the low power mode covers great range of application. Until we get some samples and examples it is still not clear. Video streams have little burst and then gaps. In reality the IPG varies a lot, in most applications LPI should work fine. Rate shifting not a good idea because, you may have to shift up and shift down at higher rate which results in no power saving. LPI has better power saving even at lower off cycles. Point to make, in real life the IPG being chopped not very good. Let's wait for AVB group and .1 group discussion for traffic pattern behavior. Question on asymmetric noise: this was related to cross talk when waking up from low power mode and the noise from one channel to another. In home AV application: either compressed or uncompressed video traffic. In applications with various streams, the IPG is chopped up and LPI should work fine. Taking the sum total of all links and traffic, the LPI is more useful. The problem with LPI is we have to buffer. This is an issue. Question is if you want higher level of efficiency and throughput or highest level of power saving, or somewhere in between. Streaming protocol for RSVP, avoid bunching, you can disable EEE when threshold exceeds certain level of capability. For switch and data center, a lot of switches are managed switch and know how to configure the ports to disable or enable energy efficient mode. We are running over some confusion on the scaling. All new LPI proposal talks about 10us wake up time. This is 3 orders of magnitude below the traffic requirement. If we use subset PHY and change speed rather than LPI, the lower rate not as efficient as LPI. We have done some simulations. Optimized for desktop links not data centers... We have done simulation in data center applications. The buffer flush and complexity. The idea of orphan packet is not applicable. Assume we continue using same buffering scheme. What is the jitter spec: Did some research; target jitter is 2ms for speakers. Negotiation of "wake-up" time: Agree that there is some complexity added. Traffic

patter example: No need to buffer. Even if some buffering needed, we believe if we just increase the buffer for 10us or more should be ok. Said upfront, the system issues, concern is LPI will affect the .1 we should double check with AV and .1 group. Not opposed to LPI, but we need to make sure it works at system level for all applications. For Video streams, not watched for 24 hours a day. Are we going to utilize when low throughput IPG and we can utilize LPI. Buffer and burst is orthogonal to LPI. Having buffer and burst would have better energy saving. How effective you are saving energy when you have to do buffer and burst, since most applications will turn off the low power feature. Not sure why there is discussion between subset PHY vs LPI. In high speed computer applications, data rate is most important than low power mode. Not sure how we determined if LPI saves more power than subset PHY. Solution is to turn off LPI if not efficient. Big concern is lots of application will turn LPI off and then not utilized at all. Some confusion on transition time to buffer time: 10us is advertised as transition time. It was mentioned the total time is in ms time not 10us. Canned traffic simulation was captured with Orphan packet can not be ignored. It should be send out immediately since have no idea what type packet it is and how important it is to send it. Super computers, vs Home computer, IP telephony, looking at all applications do not have same solution. Is it worth to consider the data center and super computer applications for energy efficient application? We are afraid in some applications EEE will be turned off. We should not exclude certain market center. Objectives say, should include all applications including data centers, indicates data center. We did not say to exclude any application. LPI should fit all applications. Not sure if subset PHY will be any different. LPI seems to be most focused for desk top applications. HPC, moving faster computer center. In different applications LPI or subset PHYs both are useful. Data center, EEE running at 10Gig, LPI may not look good. But what if we run at 1.25 G speed? The objective was to address the growth of power and focus on Data centers.

Presentation:

LPI Synchronization Feasibility Questions (Mike Greenwood)
greenwood_01_0308.pdf

Discussion:

Benefits of refresh period, reoccurring property...Tracking capability is – 7 is the good number. We should specify what number should be. Which system thermal will switch the crystal 50ppm, agree to specify rate.

There was some discussion on preparation for the joint meeting with 802.1

Meeting recessed

Wednesday March 3/19/08

Attendees

Name	Employer	Affiliation
Mike Bennett	LBNL	LBNL
Satoshi Obara	Fujitsu	Fujitsu
Yukihiro Fujimoto	NTT	NTT
Paul Gyugyi	NVIDIA	NVIDIA
Mario Traeber	Infineon	Infineon
Hugh Barrass	Cisco	Cisco
Ahmad Nouri	Broadcom	Broadcom
Joseph Chou	Realtek	Realtek
Velu Pillai	Broadcom	Broadcom
Kory Sefidvash	Broadcom	Broadcom
Khurram Kazi	SMSC	SMSC
Jim Millar	Force10 Networks	Force10 Networks
Richard Bowers	Micrel	Micrel
Mandeep Chadha	Vitesse	Vitesse
Aviad Wertheimer	Intel	Intel
Mike Grimwood	Broadcom	Broadcom
Albert Kuo	Realtek	Realtek
Angus (Shang-Ta) Lee	Realtek	Realtek
James Tsai	Intel	Intel
Scott Powel	Broadcom	Broadcom
Robert Hays	Intel	Intel
Jeff Lynch	IBM	IBM
Shinkyō Kaku	Allied Telesis	Allied Telesis
Adam Healy	LSI	LSI
Robert Busse	Transition Networks	Transition Networks
Geoff Thompson	Nortel	Nortel
George Zimmerman	Solarflare	Solarflare
Robert Grow	Intel	Intel
Ozdal Barkan	Marvell	Marvell
Wael Diab	Broadcom	Broadcom
Dan Dove	Procurve	Procurve

Meeting start: 9 AM

Meeting recessed to resume at the joint 802.1 DCB meeting: 9:50 AM

Joint meeting with 802.1 DCB: See joint meeting minutes

Meeting start time: 1:30 PM

Mike: Were there any show stoppers from .1 meeting? No...

Rob: Did not get as much feedback. Should do straw man or proposal at .1 and get their feedback.

Mike: Agreed, need to show more detail next time.

Wael: Need to show what effects .1 and how it interacts with EEE.

Presentation:

Twisted Pair Subset PHY Presentation (Scott Powell) powell_01_0308.pdf

Discussion:

It is not clear if you can shift up in the middle of packet reception. How is the MAC side can shift up speed? Is the EEE goal is to do anything to the mac? This will directly affect MAC if we need to shift speed during packet reception. How we are going to do the control policy? We need to better define the control policy to see if it effects the MAC or not when changing in the middle of the packet it will effect the MAC reception speed. Agreed, but we have not defined if there is any change to the PHY-MAC interface. This interface may still be running at full speed. This approach will not shut off the MAC. Subset PHY will run the MAC at lower speed. It has serious effect if we need to change MAC speed during TX and RX. 2- Do we have more info in MAC speed change? Any info on how fast speed can be changed? Answer: under 10 usec, independent of the speed. Focus in this group is to power down PHY and not MAC. Maybe we don't need to run all blocks in the PHY all the time. Can we make it vendor independent.

Presentation:

A "Subset PHY" approach for 10GBase-KR Energy Efficient Ethernet (Vivek Telang) telang_01_0308.pdf

Discussion:

Assuming shutting down some analog blocks, during up shift or down shift how do you transition? Answer: We have some timers setup to allow for the transition. Can this be done in 10 usec? Answer: Yes. This will funnel everyone to support lower modes. Not sure if all vendors can do this under 10usec and with same power saving. Slide 7, on receive side does analog need to know when to shift up or down. Answer: No, the receiver will get special code which will decode and determine to change the TX side. When running at lower rate, it is not a normal speed. This will have same issue as 10GBASE-T. 10GBASE-T will have either XAUI or SGMII. This will not have a standard MAC Interface. This will be integrated inside the chip, so no exposed PHY-MAC interface.

Presentation:

Enhancements to the low-power Idle Mode (Joseph Chou) chou_02_0308.pdf

Discussion:

In slide 11, is the frequency drift included? Answer: No. This needs to be included. I don't see Jitter is related to frequency drift. These should both be treated. The Frequency

gets back in lock, so it is not constantly increasing. The drift is only during the power down time. I have not seen any where in your presentation the frequency drift is included in your calculation. Yes, you are correct. The freq drift not included, but it is very negligible. I agree for next time we should include freq drift but it is not important for our initial basic calculation.

If you say this will degrade the performance by 1dB then it will not meet the standard. Need to be careful trading off margins, to achieve this. We loose robustness if we do this. Also, slide 5, maybe able to make asymmetric PHY work but now we have to test 4 PHYs, be mindful additional testing. I prefer Asymmetric mode. Case 3 is the problem (see slides). Could specify Asymmetric but if someone have problem, then don't use it. Can decide as Master if want to go to sleep or not. Everyone can do 1000BASE-T training and then pull-back to lower speed if wanted. Probably distribution of long loops/short loops. This does not need negotiation. You do AN for triple speed. Then go to low power mode when needed. For asymmetric, expectation is power saving in PHY only and not MAC? The MAC doesn't care if it is asymmetric or not. The PHY behavior is transparent. In asymmetric mode, one side is in sleep and one side only receiving. One side sleeps and the other side is still receiving. In congested network, traffic pattern is mostly in one direction, so asymmetric power down works best. It is best to support asymmetric for most flexibility. We need to do cost/benefit analysis for symmetric vs. asymmetric.

Presentation:

European Code of Conduct (Mario Träber) traeber_02_0308.pdf

Discussion:

We need progress in EEE to meet the international power consumption requirements.

Presentation:

Update to zimmerman_01_0308.pdf (George Zimmerman)
zimmerman_02_0308.pdf

Discussion:

Planning to come up with max T_W? I assume a max parameter, but will be determined later. Will be a system wake up time.

Discussion: State Machines (based on question from Khurram Kazi):

Can we come up with all possible schemes to save power? The extreme case is 10Gbase-T. We standardize when it interfaces to another vendor PHY. One interface is the line, but another interface is system behavior. If it effects the system behavior, then we need to standardize otherwise the implementations inside the PHY are transparent to others. Please point out what you want to standardize. If we try to do power saving and it impacts standards then it will not

be compliant. Can we do some modifications to the standard which could help with power saving? If the transmitter amplitude is lowered then it needs to be in the standard in order to interoperate between PHY vendors. We could easily change some of the parameters and modify the standard but it will not interoperate with legacy PHYs. If this group is chartered to do work outside the objectives. Our objectives are to define a mechanism to reduce power during low utilization. The objective does not include reducing power to existing PHY application. If we go outside and define a new mode of operation to reduce power it needs to be reconsidered. In terms of legacy PHYs, but we should consider this for 10GBASE-T. If people want to bring power consumption proposals then we can discuss them, but work related to our PAR and objectives will have higher priority. The scope of the project says: specify PHY enhancements as required for a selected subset of PHY types to improve energy efficiency. Sounds like a good idea to do this.

Mike: The task force will start at 8:30 AM on Thursday.

Meeting recessed: 5:35 PM

Thursday March 20, 2008

Attendees

Name	Employer	Affiliation
Mike Bennett	LBNL	LBNL
Satoshi Obara	Fujitsu	Fujitsu
Mario Traeber	Infineon	Infineon
Robert Hays	Intel	Intel
Paul Gyugyi	NVIDIA	NVIDIA
Dimitry Taich	Teranetics	Teranetics
Richard Bowers	Micrel	Micrel
Shinkyo Kaku	Allied Telesis	Allied Telesis
Khurram Kazi	SMSC	SMSC
Velu Pillai	Broadcom	Broadcom
Scott Powel	Broadcom	Broadcom
Mike Grimwood	Broadcom	Broadcom
Ahmad Nouri	Broadcom	Broadcom
Jeff Lapak	UNH-IOL	UNH-IOL
Angus (Shang-Ta) Lee	Realtek	Realtek
Albert Kuo	Realtek	Realtek
Joseph Chou	Realtek	Realtek
George Zimmerman	Solarflare	Solarflare
James Tsai	Intel	Intel
Aviad Wertheimer	Intel	Intel
Kory Sefidvash	Broadcom	Broadcom
Jeff Lynch	IBM	IBM
Brad Booth	AMCC	AMCC
Adam Healy	LSI	LSI
Hugh Barrass	Cisco	Cisco
Ozdal Barkan	Marvell	Marvell
Mandeep Chadha	Vitesse	Vitesse

Wael Diab
Dan Dove

Broadcom
Procurve

Broadcom
Procurve

Meeting started at 8:30 AM

Floor open for motions

Motion #1

802.3az task force adopt 100BASE-TX Low-Power Idle proposal in chou_01_0308.pdf as part of the baseline solution for Energy Efficient Ethernet.

Moved by: Dan Dove
Second by: Hugh Barrass

Discussion:

Scott: Concerned about the Asymmetric operation. Is this applies to all speed and Asymmetric mode or just 100BASE-TX?

George: It is only for 100BASE-TX.

Rob: No change in control policy and MAC behavior for any of the proposal.

Scott: As long as it does not apply to other speed. Then I am ok.

Wael: Is this the best technique? Important what we do here has a major impact on higher levels. I can not support this motion.

Question has been called. No objection to calling the question

Motion #1

802.3az task force adopt 100BASE-TX Low-Power Idle proposal in chou_01_0308.pdf as part of the baseline solution for Energy Efficient Ethernet.

Moved by: Dan Dove
Second by: Hugh Barrass

Yes: 15
No: 4
Abstain: 6

Tech motion; 75% required to pass

Motion passes

Motion #2

802.3az Task Force adopt 1000BASE-T Low-Power Idle proposal in healey_01_0308.pdf as part of the baseline solution for Energy Efficient Ethernet.

Moved by: Hugh Barrass
Second by: Robert Hays

Discussion:

Hugh: This is a symmetric operation. Because of loop timing Adam needs to do more work.

Wael: Can not support this motion at this time. Adam has done great job. Based on .1 feedback. Not opposed on LPI, but don't know all the impact, latency, buffer and burst issues. Not ready to support this motion.

Scott: One of the main point from Adam proposal, how to design with MAC layer do asymmetric. Is this not asymmetric?

Hugh: yes and no. The physical layer is symmetric. Adam showed getting down the PLA layer, you can force the PHY to be symmetric. Possible to allow asymmetric as subset of this. The MAC is asymmetric. It is not changing the MAC layer. None of the proposals forced the MAC to be symmetric.

Scott: Concerned this is asymmetric proposal. Not sure if it is good to go from Gig to zero speed.

Dan: This is excellent based line to work on. Not perfect but meets the objectives.

Rob: 3 presentations, showing it is ok to go from Gig to 0 speed. We showed below 20% line utilization is more effective to go to 0 speed.

Scott: We have not seen enough info. As discussed in .1 turning down to 0 not a good idea.

Rob: even Wael suggested this is good for low utilization

Wael: This does not meet the objectives. Second line says, this should not affect upper layer. But the proposal will affect upper layer. Even subset PHY does not. How much energy is being saved? Is the goal is for OFF and ON or should cover other applications which throttle at middle throughput?

Dan: Need to come to agreement on the baseline, it helps to move forward with come up with the draft. This is a good starting point.

Hugh: Wael has made excellent points. This is not discussion about data center. They will go to 10Gig. This is most wide deployed PHY. Not for data center application. LPI is most appropriate for gig PHYs at 20% utilization.

Rob: Hard to believe after 6 months task force we don't know what we are proposing. This meets the task force requirements.

George: Going back to the objective, second line, came about how long of a transition time we can handle. Nothing will be broken in upper layers. .1 discussion ... We are trying to make the network to be more energy efficient, and we're bringing in tools and modes to make the network more energy efficient. Best to put some stake in ground and show .1 here are EEE requirements.

Wael: I was a coauthor of the proposal in May 07 which stated the objective was to have EEE implementation to be transparent to upper layers. The objective is clear not to touch upper layers. We have not talked to AVB and other groups. We are touching a lot of PHYs. We need to better understand the impacts before taking the base line.

Mike: Regarding comment on TCPIP, we examined this in July 2007; transitions in tens of milliseconds not detrimental to TCP performance.

No further discussion on the motion

Motion #2

802.3az Task Force adopt 1000BASE-T Low-Power Idle proposal in healey_01_0308.pdf as part of the baseline solution for Energy Efficient Ethernet.

Moved by: Hugh Barrass
Second by: Robert Hays

Yes: 24
No: 7
Abstain: 2

75% required:

Motion passes

Motion #3

802.3az adopt the timeline in eee-timeline_0308.pdf

Rob: Do you mean all details be defined by May? MIB details, etc. Do you mean last feature base line or last proposal?

Mike: Last new proposal. We have been working on two proposals and we need to limit when new proposals can be considered or this project will go on forever. This does not mean we can't continue to develop the proposals we've been working on. I think the MIBs should be defined by the time we go to Working Group ballot.

Hugh: Before going to Sponsor Ballot, need to have MIBs defined. We do clause 30 and the MIB taskforce will take that and implement.

Dan: We need to still bring in proposals 10G and XAUI.

Rob: In May need to have all proposals on the table. By July we need to do final selections. Plan to have some draft by September.

Rob: Concerned EPA may fall back to July 2010. My customer's validation cycle. Can we move in the schedule?

Bob: Gig Ethernet was shipping by first draft ballot.

Rob: Can we move in the balloting process?

Bob: The first sponsor ballot determines if any issues which could delay the process. Takes 2 months process time for each ballot.

Motion # 3

Move that 802.3az adopt timeline presented in eee-proposed-timeline_0308.pdf

Move: Dan Dove

Second: Wael Diab

Yes: 30

No: 0

Abstain: 0

Motion passes

Wael: Adding new feature is OK, but not to bring in totally new proposal for totally new idea for power saving.

Dan: Do you have editors? It would be good to have them by the next meeting so they can begin working on the editor's draft.

Mike: I have two people who have expressed interest in being an editor

Mike: Any other motion to bring forward?

No other motions.

The chair polled the task force regarding meeting attendance

New Attendees: 7

How many plan to continue attending? 7

How many of you plan to attend other 802.3 TF meetings? 0

How many plan to attend May meeting in Munich? 20

How many may attend? 6 (delta, not total)

How many will not attend? 5

How many attend Sept meeting? 13

How many may attend? 16 (total)

How many will not attend? 5

The chair reviewed the future meeting dates and venues with the task force, then asked for a motion to adjourn.

Motion to adjourn the meeting:

Moved by: Hugh Barrass

Second by: Brad Booth

Passed by voice

Meeting adjourned at 10:10 AM.