Nov-13-2007

Name	Employer	Affiliation
Mike Bennett	LBNL	LBNL
Hugh Barrass	Cisco	Cisco
Chris Diminico	MC Communcations	Solarfare.
Robert Grow	Intel	Intel
Kory Sefidvash	Broadcom	Broadcom
Jim Barnette	Vitesse	Vitesse
David Koenen	HP	HP
Gofff Thompson	Nortel	Nortel
Steve Carlson	HSD	HSD
Blaine Kohl	Tehuti Networks	Tehuti Networks
Scott Powel	Broadcom	Broadcom
George Zimmerman	SolarFlare	SolarFlare
Howard Baumer	Broadcom	Broadcom
Dan Dove	Dove Networking	Procurve
Rob Hays	Intel	Intel
Joseph Chou	Real Communications	Realtek
Brad Booth	AMCC	AMCC
Ahmad Nouri	Broadcom	Broadcom
Edward Brown	Broadcom	Broadcom
Adam Healy	LSI	LSI
Ozdal Barkan	Marvell	Marvell
Mark Merrin	Netgear	Netgear
Michelle Gong	Intel	Intel
Wael Diab	Broadcom	Broadcom
Jim Millar	Force10 Networks	Force10 Networks
Jeff Cain	Cisco	Cisco
Sanjay Kasturia	Teranetics	Teranetics
Amit Gattani	Akros Silicon	Akros Silicon
Wiren Perera	Plato Networks	Plato Networks
David Law	3Com	3Com

Bob Grow asked for task force confirmation of Mike Bennett as the chair of 802.3az.
Task force confirmed Mike Bennett as Chairman of 802.3az Task Force at 9:04 AM

Chair appointed Kory Sefidvash as recording secretary for this meeting

Agenda & General information
By: Mike Bennett
See – agenda_2_1107.pdf

The Chair displayed the Patent slides and read patent slides 1 and 5. He verified that everyone in the room had read plenty of time to read the slides that were not read out loud by asking for a show of hands if more time was needed to read them. He called for patents.

Jim Barnette has informed the task force that his company has a patent to submit.

Title: Open Questions for the Task Force
By: Mike Bennett
See: bennett_1_1107.pdf

Mike identified the change of focus for the work of the task force, with the first milestone being to complete a draft. He stated his goal of completing the project in 2009, but that it was more important to produce a good standard than to rush to completion. There was a review of the objectives. Mike presented a set of question he believes the task force needs to answer and ended by presenting a possible timeline for the project. There were comments regarding the timeline seeming aggressive and a suggestion that draft 1.0 should be changed to baseline 1.0. There were comments that we should have a new clause, even if only to refer to other clauses and that doing it this way would make it easier to get through balloting later

Title: EEE Control Protocol Proposal
By: Hugh Barrass
See: barrass_1_1107.pdf

Discussion:

There was a discussion on using LLDP as a communication protocol for EEE. Some people suggested LLDP is not designed for real time use and timing constraints would need to be added. There was concern expressed that timing may not meet the requirements. It was suggested we define the requirements, and then choose a protocol. Some people are confident that LLDP will meet the needs. Others believe it is too early to get into the details of signaling. We were warned that the proposed schedule is too aggressive and if we don't spend enough time to figure things out now, we will end up spending more time in the balloting phase of the project. Another comment was made that we need to consider making this work on fiber and for HSSG.

Break: 11:09 AM
Reconvene: 11:27

Title: Conditions for Backplane PHY EEE Transitions
By: David Koenen
See: koenen_1_1107.pdf

Discussion:

It is good to start thinking about 10GBASE-KR. It is possible the work in this task force may have an impact on the HSSG for backplane and "CX-like" copper PHYs.

Title: Another Piece of EEE: An additional requirement for Energy Efficient Ethernet
By: Geoff Thompson
See: thompson_1_1107.pdf

Geoff presented a requirement for a new interface to EEE that communicates with the upper layers, enables control of the "speed shift", and operates in real time. He warned that the project would fail if the control mechanism is proprietary.

Break: Lunch 11:56
Reconvene 1:30

Title: Energy Efficient Ethernet and 802.1
By: Mike Bennett
See: EEE-and-802_1v7.pdf

Mike reviewed the presentation given to the 802.1 group and the straw poll feedback.

Discussion:

We want to state couple of points clearly to them. The link or latency is not fixed bandwidth. We are going to introduce a new concept a link state than a link speed. If it is a gig link then it is gig link, but it will have a new state that won't have gig speed. We don't need to be political with them and can be blunt with them. Note the last bullet of slide 5 state what Bob and Geoff is recommending. It might not be as blunt. The comment was made that we need to start a liaison interface with 802.1

Title: IEEE 802.3 Clause 30 management, MIB, registers and function
By: David Low
See: law_1_1107.pdf

Discussion:

We will probably need a new MIB. We need to make sure that if we continue to do "do something afterwards" that there is interest and the resources to do the work.

Break: 2:20
Reconvene: 2:40

Title: A Gigabit " Subset PHY" Approach for 10GBASE-T Energy Efficient Ethernet
By: Scott Powell
See: powell_1_1107.pdf

Discussion:

There was a lot of discussion on optimizing power for multi-speed PHYs. There was a concern that creating Subset PHYs will add qualification steps, and the response was that new PHYs always require testing. The response to this was that it is not a new PHY, just that some elements were going to be turned off when not needed.

George noted that in Slide 5, the 8DSQ is not balanced. Scott acknowledged that.

Title: Speed Switching without Communication Interruption
By: Jim Barnette
See: barnette_1_1107.pdf

Discussion:

The interruption is the issue, i.e. down time and how long it will take to change the speed. We could use a faster FLP to have frequency content to provide the training. Why not PAM3 instead of PAM 2? How much better PAM2 with shaping filter than PAM3? The equalizer is the problem and not shaping filter using the PAM3 over the PAM2. Asymmetry is attractive solution

Title: The Path to Working group Ballot
By: Howard Frazier
See: frazier_01_1107.pdf

Discussion:

Last new feature could go between D1.0 and D2.0. Last technical change can go between D2.0 and D3.0. A draft takes approximately one month. Moving from D1.0 to D1.1 is an informal process reviewed by the task force. We should get someone like Howard Frazier or Brad Booth to give a tutorial on the comment tool. Anyone seriously considering being an editor needs to attend the tutorial offered by the Editorial Staff. This tutorial would happen at a plenary meeting so we should do this in either March or July.

Recess: 5:30 PM

Wednesday November 14 2007
Reconvene: 9:30 AM

Name	Employer	Affiliation
Mike Bennett	LBNL	LBNL
Hugh Barrass	Cisco	Cisco
Sylvia Ratnasamy	Intel	Intel
Sergiu Nedevschi	UC Berkeley	UC Berkeley
Robert Grow	Intel	Intel
Ilango Ganga	Intel	Intel
Kory Sefidvash	Broadcom	Broadcom
Jim Barnette	Vitesse	Vitesse
David Koenen	HP	HP
Gofff Thompson	Nortel	Nortel
Steve Carlson	HSD	HSD
George Clasemen	Micrel	Micrel
Scott Powel	Broadcom	Broadcom
George Zimmerman	SolarFlare	SolarFlare
Howard Baumer	Broadcom	Broadcom
Dan Dove	Dove Networking	Procurve
Rob Hays	Intel	Intel
Joseph Chou	Real Communications	Realtek
Brad Booth	AMCC	AMCC
Ahmad Nouri	Broadcom	Broadcom
Edward Brown	Broadcom	Broadcom
Adam Healy	LSI	LSI
Ozdal Barkan	Marvell	Marvell
Mark Merrin	Netgear	Netgear
Michelle Gong	Intel	Intel
Wael Diab	Broadcom	Broadcom
Jim Millar	Force10 Networks	Force10 Networks
Sanjay Kasturia	Teranetics	Teranetics
Amit Gattani	Akros Silicon	Akros Silicon
Wiren Perera	Plato Networks	Plato Networks
David Law	3Com	3Com

Title Reducing network energy consumption via sleeping and rate-adaptation
By-Sylvia Ratnasamy
See-ratnasamy_1_1107.pdf

Discussion:

Typical slide Pidle = .8 Idle state less than 5%. Clock gating to reduce power. Looking at leakage how to control it.

Slide 21 discrete rate shows better performance than decade rate.
1G, 2.5G, 5G 10G will give better performance than 10/100/1000/10G

Question: What are the buffer needs for 10G? Travel time worth of buffer and that does not change. If you are filling your buffer you should not sleep.
Question: which network was used for the data? There were two networks: Abilene and Intel's corporate. Abilene backbone is 15% for slide 13.

Title Active/Idle Toggling with 0BASE-x
By-Robert Hays
See-Hays_1_1107.pdf

Discussion:

There was a comment that a switch needs to know what is coming and how much is coming. If the switch isn't buffering then memory is available. A 24-port switch will at least run off memory for one port and turn on the memory compared with turning off a portion of memory that I need for the buffering.

Regarding add an objective for 0BASE-X, we have an objective which is approved. 0BASE-X could be subset or up coming presentation.

Comment: we need an objective to address period of when there is no utilization. This is sweet sport of saving energy. We need to define a new name for idle over the existing idle of power.

We could have normal idle and low power idle. Moving from low power idle to normal idle is going back to normal idle. This is low power idle signaling the other side to go to sleep. We should provide tools.

Lunch: 11:30 AM
Reconvene: 1.00 PM

Title: Active/Idle Concept for 10GBASE-T
By: George Zimmerman
See" zimmerman_2_1107.pdf

George updated his presentation slightly, replacing Active Idle with Deep Sleep Idle – there were no objections

Discussion:

There was a fair amount of discussion on buffers being needed during the "off time" and that end stations operate differently than switches.

Question: what is T_next_frame? If you look at the wire, it is the time the frame goes over the wire or the LDPC frame time.

Regarding Slide 9, if we want to power off for durations in the order microseconds, then we don't want to power off an entire section otherwise we may not be able to recover fast enough. There is also an issue with turning a lot of power on.

There was a comment regarding Slide 5, that we should look at different solution for each speed. The Subset PHY proposal and this presentation are similar and have synergy

There was discussion that when you send periodic frame it will have 1 bit of information to give indication to link partner to come out of sleep mode. Keeping the XUAI active then it will allow upper layer to tell the PHY to come out of deep sleep – this is similar to wireless with paging channel and waking up and then go to sleep.

There was a question asking if there is any EMI issue. George didn't think there would be

There was a comment that LDPC boundary is aligned with XAUI boundary. Clause 49 is 8 byte. 2048 is for LDPC boundary and should fall in the boundary.

Break 2:30 PM
Reconvene: 3:00 PM

Straw Poll 1 (conducted by Rob Hays)
What term would you prefer the task force adopt to represent the reduced-power idle state discussed in Hays_01_1107

A. Idle
B. Sleep
C. Siesta
D. Low-Power Idle
E. Green State.

A: 0
B: 4
C: 1
D: 14
E: 4

Motion 1
802.3az Task force should adopt the following objective: "Define a low power idle state for use when no data transmission is required for each PHY rate supported by IEEE8023.az

Moved by: Robert Hays

Second by : Sanjay Kasturia
(Technical motion require >= 75%)

Yes: 10
No: 8
Abstain: 5

Motion fails

Most of the discussion on this motion was whether or not a new objective was necessary. Those favoring the motion felt it was needed because there was no objective directing the task force to work on low power idle. Those opposed believed that the objective wasn't necessary and that by adopting the objective the task force would be adopting a proposal. Some also believed that for new objective to be adopted by the task force, the objective would have to meet the 5 Criteria – not all agree that this is necessary.

David Koenen stated that the thought the objective is too strong, but that he supports the concept. He offered "add this to a baseline or draft instead of objective" as a friendly amendment.

Rob Hays did not consider the offer friendly and David withdrew the amendment.

Bob Grow stated, as Chair of 802.3, his disappointment with the level of procedural argument going on. The objective is a way to communicate to the Working Group what the task force is going to work on and that the objective is appropriate.

After much more discussion, Sanjay called the question. There was no objection to calling the question.

Motion 2
Move that the 802.3az task force to adopt the minutes of July and September meetings.

Moved: Howard Frazier
Second: Howard Baumer

The motion passed by voice.

Motion to adjourn.

Moved: Hugh Barrass
Second: Brad Booth

The motion passed by voice.

Meeting Adjourned at 4:06 PM