IEEE 802.3az
Energy Efficient Ethernet

Closing Plenary Report

Orlando, FL
March 18, 2010

Mike Bennett
mjbennett@ieee.org

IEEE 802.3az EEE Contacts

- Task Force Chair
  - Mike Bennett (mjbennett@ieee.org)
- Web Master
  - Bruce Nordman (bnordman@lbl.gov)
- Editor-in-Chief
  - Sanjay Kasturia (skasturia@teranetics.com)

Reflector and Web

- To subscribe to the EEE TF reflector, send your request to:

  ListServ@ieee.org

  with the following in the body of the message (do not include "<>"):

  subscribe stds-802-3-eee <yourfirstname> <yourlastname>
  end

- Send reflector messages to:

  stds-802-3-eee@listserv.ieee.org

- For complete instructions on reflector usage, subscription, and unsubscription:

  http://www.ieee802.org/3/az/reflector.html

- Task Force web page URL:

  http://www.ieee802.org/3/az/

Reflector and Web

- Our latest draft is D2.3

- Task Force *private* web page URL:

http://www.ieee802.org/3/az/private/index.html

Login: 802.3az

Password: xxxxxx

Overview of IEEE 802.3 Standards Process (3/5)
Working Group Ballot Phase

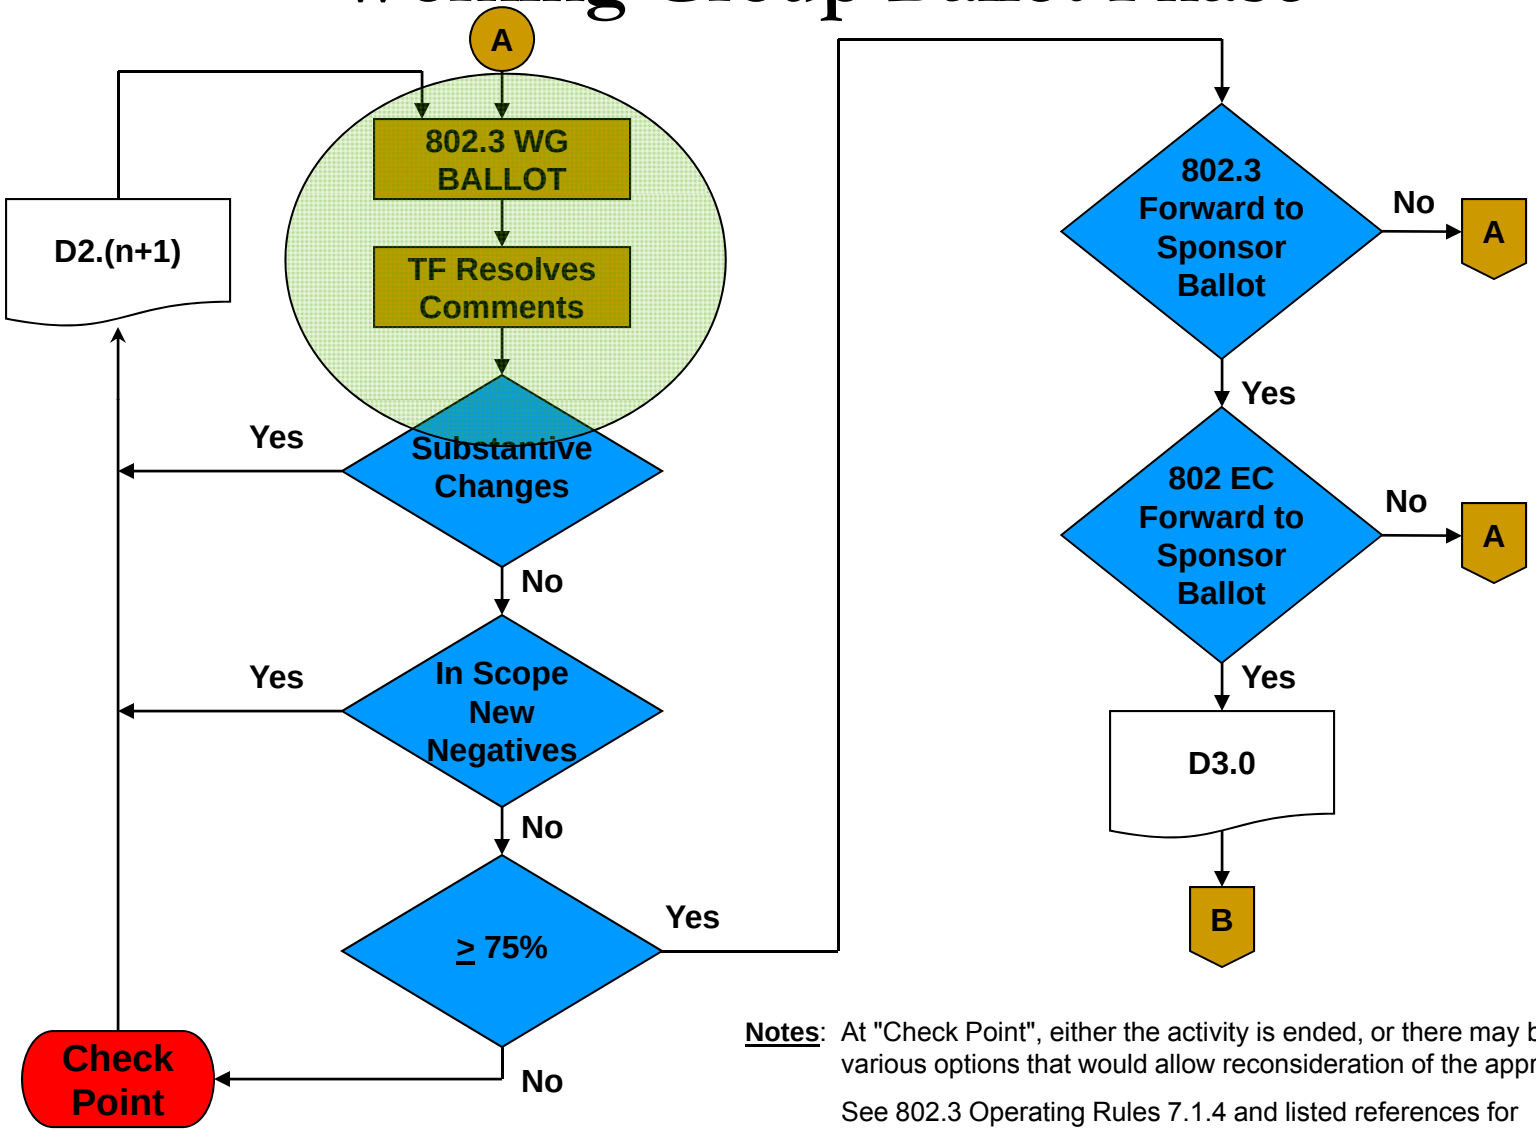

Notes: At "Check Point", either the activity is ended, or there may be various options that would allow reconsideration of the approval.

See 802.3 Operating Rules 7.1.4 and listed references for complete description

Progress this week

- Processed 79 comments
- Fixed a few problems
- Considered Maintenance item 1216
- Directed editors to produce D2.4
- We expect we only need one more recirc
  - Will ask for conditional approval to go to sponsor ballot

802.3az Task Force Report - Motions

Motion #2 M: D. Dove S: V. Pillai

Move to Modify 14.4.1 as follows:

14.4.1 Overview The medium for 10BASE-T is twisted-pair wiring. ~~Since a~~ A significant number of 10BASE-T networks are ~~expected to be installed~~ utilizing in-place unshielded telephone wiring and typical telephony installation practices, the end-to-end path including different types of wiring, cable connectors, and cross connects must be considered.

Y: 10 N: 0 A: 5 (technical, 75%) motion passed

802.3az Task Force Report - Motions

Motion #3 M: D. Dove S: V. Pillai

Move to Modify the second paragraph of 14.4.1 by inserting the underlined text as shown below:

The medium for 10BASE-Te is twisted pair wire. The requirements of the 10BASE-Te simplex link segment (either pure 10BASE-Te or mixed 10BASE-T, 10BASE-Te) are equivalent to the requirements of the Class D channel specified by ISO/IEC 11801:1995. This requirement can also be met by Category 5 cable and components as specified in ANSI/TIA/EIA-568-B:2001.

Y: 10 N: 0 A: 5 (technical, 75%) motion passed

802.3az Task Force Report - Motions

Motion #4 M: J. Barnette S: D. Dove

Move to Modify the second paragraph of 14.1.1.3 by inserting the underlined text as shown below:

The medium for 10BASE-Te is twisted-pair wire. The performance specifications of the 10BASE-Te simplex link segment (either pure 10BASE-Te or mixed 10BASE-T, 10BASE-Te) is a channel meeting or exceeding the requirements of the Class D channel specified by ISO/IEC 11801:1995.

Y: 9 N: 0 A: 5 (technical, 75%) motion passed

802.3az Task Force Report - Motions

Motion #5 M: D. Dove S: V. Pillai

Move to Modify 14.5.2 Crossover function as shown below:

Additionally, the MDI connector for a MAU that implements ~~the~~ <insert>a fixed crossover</insert> function shall be marked with the graphical symbol “X”.

Y: 9 N: 0 A: 5 (technical, 75%) motion passed

Note: This changes text in the original base standard. The Task Force agreed to implement this change rather than sending the request to maintenance and then implement the change since we have Clause 14 open.

802.3az Task Force Report - Motions

Motion #7 M. Hugh Barrass S: Velu Pillai

Accept comment resolutions from D2.3 as recorded in the comment database. Direct the IEEE P802.3az editorial team to generate draft 2.4, based on Draft 2.3, the resolution of comments against Draft 2.3 and motions 2 through 5 of this meeting. Request the Working Group chair to conduct an IEEE 802.3 WG recirculation ballot for P802.3az/D2.4

Y: 13 N: 0 A: 0 (technical, 75%) motion passed

802.3az Task Force Report - Motions

Motion #8 M: Velu Pillai S: Gavin Parnaby

Request the Working Group Chair to submit conditional approval to begin sponsor ballot on P802.3az to the EC.

Y: 13 N: 0 A: 0 (technical, 75%) motion passed

Information required to support request for conditional approval to forward a draft standard

IEEE 802 LMSC Operations Manual Clause 14

P802.3az, Draft 2.0, Monday, 31st August, 2009 11:59 PM AOE

Comments: 483	Initial Working Group Ballot Draft 2.0			Req
	#	%	Status	%
Abstain	25	17.86%	PASS	< 30
Disapprove with comment	21			-
Disapprove without comment	0			-
Approve	94	81.74%	PASS	≥ 75
Ballots returned	140	67.31%	PASS	≥ 50
Voters	208			

P802.3az, Draft 2.1, Wednesday, 4 November 2009, 11:59 PM AOE

Comments 262	1st Working Group Ballot Recirc Draft 2.1			Req
	#	%	Status	%
Abstain	24	16.67%	PASS	< 30
Disapprove with comment	21			-
Disapprove without comment	0			-
Approve	99	82.50%	PASS	≥ 75
Ballots returned	144	69.23%	PASS	≥ 50
Voters	208			

P802.3az, Draft 2.2, Wednesday, 13th January 2010, 11:59 PM AOE

Comments 112	2nd Working Group Ballot Recirc Draft 2.2			Req
	#	%	Status	%
Abstain	22	15.17%	PASS	< 30
Disapprove with comment	17			-
Disapprove without comment	0			-
Approve	106	86.18%	PASS	≥ 75
Ballots returned	145	69.71%	PASS	≥ 50
Voters	208			

- P802.3az, Draft 2.3, March 5, 2010, 11:59 PM AOE

Comments received: 79	Count	%		Status
		Actual	Require	
Abstain	21	14.38%	< 30	PASS
Disapprove with comment	5		-	-
Disapprove without comment	0		-	-
Approve	120	97.60%	≥ 75	PASS
Ballots returned	146	70.19%	≥ 50	PASS
Voters	208		-	-

Note: 2 ballots flipped to approve since the ballot close

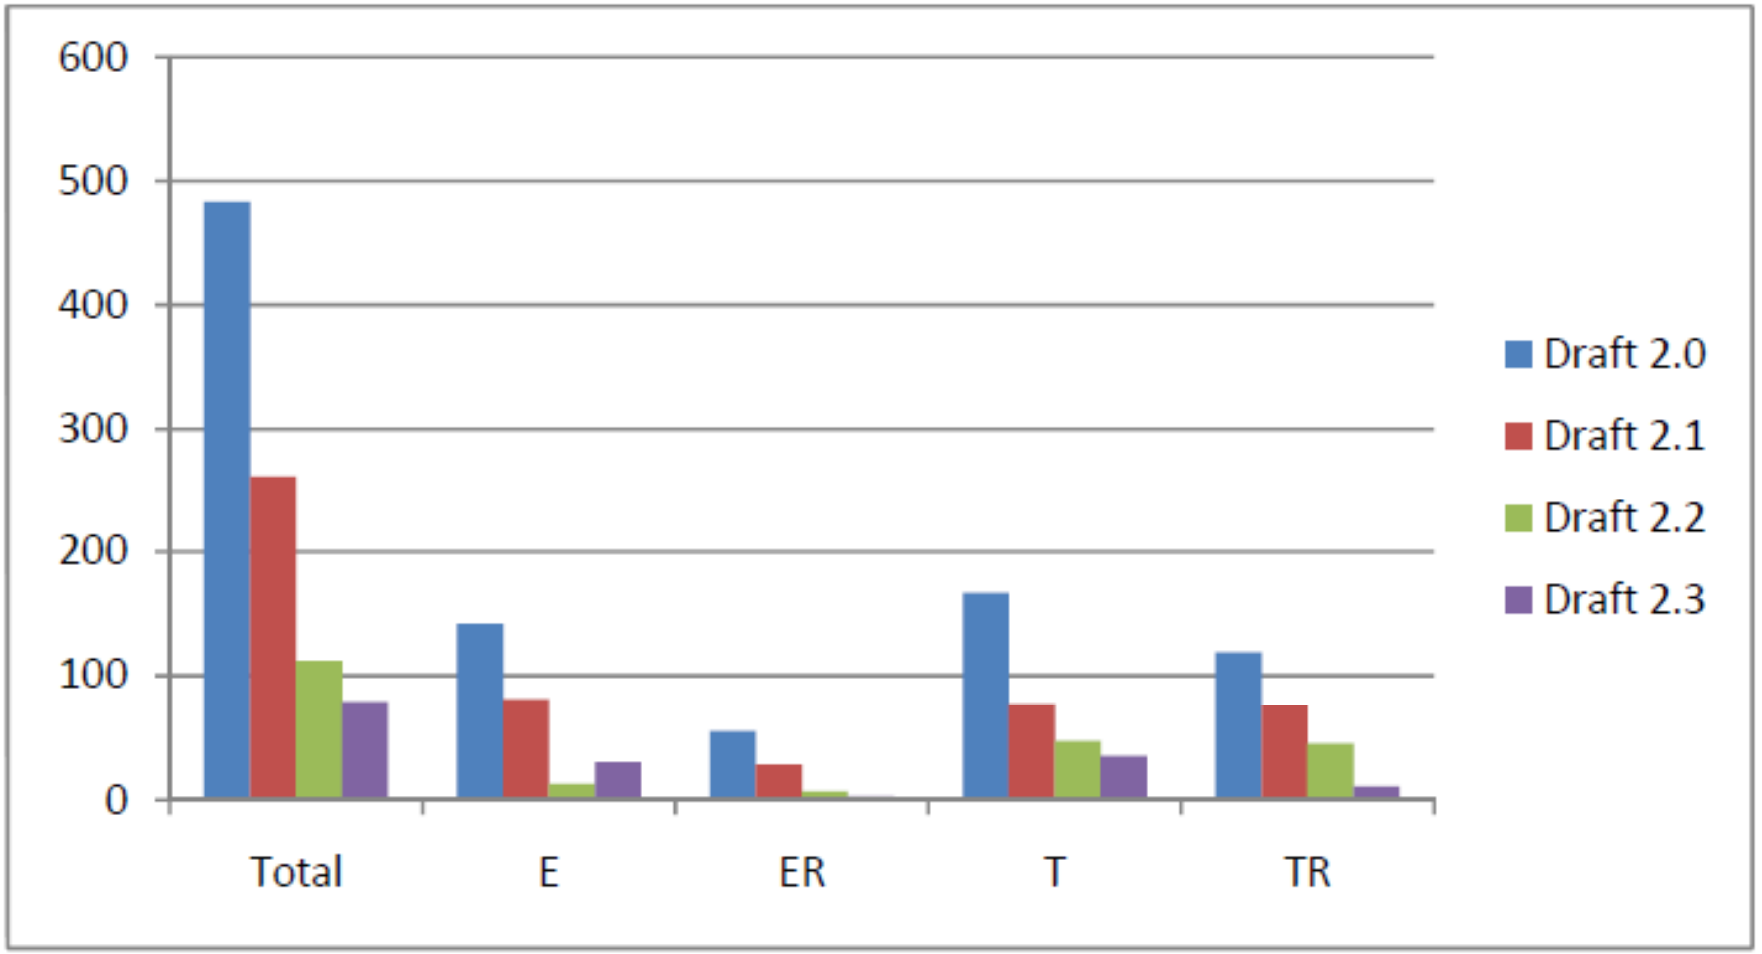
600
500
400
300
200
100
0
Total
E
ER
T
TR
Draft 2.0
Draft 2.1
Draft 2.2
Draft 2.3

Unsatisfied Negative Comments

Cl 49 SC 49.2.9 P 150 L 28 # 127 Dawe, Piers Independent
Comment Type TR
The Lock state diagram, which I don't think is optional, uses the variable "rx_block_lock" where the current standard has "block_lock". Yet 49.2.13.2.2 says "The following variables are used only for the EEE capability... rx_block_lock". Problem - and there may be similar problems e.g. in Clause 36. So I'm piling on to D2.0 comment 190 and 174, we need to preserve the non-EEE material in an undamaged state, by use of annexes like 4A, duplicate state diagrams or other means. Otherwise, users will go back to 802.3-2008 for non-EEE product, and any future maintenance to affected areas will be ignored.
SuggestedRemedy
Preserve the non-EEE material in an undamaged state, by use of annexes like 4A, duplicate state diagrams or other means.
REJECT.
This was discussed at length during the resolution of comments against draft 2.0 and the
task force decided against the suggested remedy.

Unsatisfied Negative Comments

*Cl **00 SC 0 P 1 L 1 # 10174 Frazier, Howard Broadcom***

*Comment Type **TR***

This is a general comment regarding the structure of the draft amendment.

As an amendment to IEEE Std 802.3, the material in this draft will eventually be folded into the base standard. When this happens, the definitions for the 100BASE-X and 1000BASEXPhysical Coding Sublayers will be substantially changed, and the changes will be difficult to discern. The definitions for the MII and GMII will also be substantially changed.

The 100BASE-X and 1000BASE-X PCSs are used for many other port types besides 100BASE-TX and 1000BASE-KX. Among these are 100BASE-FX, 100BASE-LX10, 100BASE-BX10, 1000BASE-SX, 1000BASE-LX, 1000BASE-CX, 1000BASE-LX10, 1000BASE-BX10, 1000BASE-PX10, 1000BASE-PX20, 10G/1GBASE-PRX-D/U1, 10G/1GBASE-PRX-D/U2, and 10G/1GBASE-PRX-D/U3.

These port types are not included in the set of objectives for P802.3az, and the specifications for the PCS and MII for these port types must not be changed or effected in any way by P802.3az. Each of these port types must have a current IEEE Std 802.3 PCS and MII to reference.

SuggestedRemedy

There are many ways to solve this problem. I prefer the following approach:

1. Preserve the definitions for the MII, GMII, 100BASE-X PCS, and 1000BASE-X PCS without change.

2. Define the changes required to support EEE in a set of normative annexes, i.e. Annex 24A for Clause 24, and Annex 25A for Clause 25, etc. Example text for Annex 24A and Annex 25A have been provided by me to the task force chair.

3. Refer to these normative annexes from the body of Clause 78.

ACCEPT IN PRINCIPLE.

See response to Comment #410

Unsatisfied Negative Comments

*Cl **00 SC 0 P 1 L 1 # 10174 Frazier, Howard Broadcom***

*Comment Type **TR continued Response to comment # 410***

The way that EEE operation has been added to the base clauses for PHYs other than 10BASE-T produces a risk that existing non-EEE PHYs and Reconcilliation sublayers will be made non-compliant. The requirements have also been added in a way that will make EEE PHYs incompatible with currently compliant non-EEE devices. My comments on 22.2.2.4 and 22.2.2.7 are examples of where that has happened.

The addition of EEE to IEEE 802.3 should not make existing IEEE 802.3 compliant devices non-compliant. EEE devices should be able to work with non-EEE devices at the xMII and MDI interfaces. It should be optional to support and any new requirements and behaviors should only apply to devices that support EEE/LPI operation. Any behaviors at the xMII or MDI that are outside what is specified for non-EEE devices should only apply when EEE operation is enabled so that EEE devices interoperate properly with non-EEE devices.

SuggestedRemedy

The safest way to do this would be to create separate clauses for behavior when EEE is enabled similar to the creation of annex 4A for full-duplex, though that would greatly increase the size of the document. The alternative is to carefully use the same type of formula any time you change a requirement for EEE. That is, the old requirement needs to be proceeded by something like "When EEE operation is not enabled," and the new requirement by "When EEE operation is enabled,". I have used enabled rather than supported because a device that supports EEE should not exhibit a new behavior when attached to a device that doesn't support EEE. For a PHY, this applies both to the xMII interface when attached to a Reconcilliation layer that doesn't support EEE and to the MDI when the link partner PHY doesn't support EEE or isn't able to enable it because the link partner's Reconcilliation sublayer doesn't support it.

ACCEPT IN PRINCIPLE.

Carefully draw a distinction between requirements/variables/timers that are required for EEE operation. New variables/timers may be kept as a separate list instead of being integrated alphabetically into existing lists. The text should be clear that when EEE is not in use (due to something in the chain -e.g. link partner capability etc) the behavior of the PHY should be identical to that of a non EEE PHY. The text should also be clear that non-EEE capable PHYs need not implement the EEE related counters/timers etc.

Unsatisfied Negative Comments

*Cl **22 SC 22.7a.2.3 P 32 L 15 # 10165 Frazier, Howard Broadcom***

*Comment Type **TR***

A state diagram in the MII clause. Wow. Why can't the PHY assert/deassert the CRS signal to indicate when the transmit path is in LPI?

SuggestedRemedy

Take out the state diagram. The 100BASE-TX PHY with LPI should be responsible for asserting and deasserting CRS, and then implement the Suggested Remedy in my general comment concerning the structure of the draft amendment.

REJECT.

In favor of accepting the proposed reject:

Yes: 15

No: 0

Abstain: 7

The state machine in the Reconcilliation Sublayer was the cornerstone of the baseline (law_01_1108) that was adopted by the Task Force.

It was considered advantageous to have the control of the PLS_CARRIER.indication in the RS for a number of reasons:

1. It keeps the PHY receive and transmit paths separate (the PHY considers CRS to be part of the receive path).

2. It allows the PHY to go to sleep without having to maintain state & control the wake process.

3. It keeps the "data holdback" function close to the MAC and egress buffers, where it would be implemented in most designs.

4. It frees the PHY from having to participate in the wake time negotiation process (that is controled using LLDP frames).

5. It works for PHYs that operate at speeds greater than 1Gbps, so the same mechanism can be used for all speeds.

The state diagram would be present (or deleted according to the comment) whether the proposed changes to the document are accepted or not.

Unsatisfied Negative Comments

Cl 22 SC 22.2.2.6a P 28 L 46 # 10167 Frazier, Howard Broadcom

Comment Type ***TR***

What do the little triangles in Figure 22-6a represent? The figure presents what appears to be a timing diagram that shows the relationship between various logical signals. How does an abstract service primitive fit into a logical timing diagram, and what does a triangle indicate?

SuggestedRemedy

Remove the abstract service primitive from the timing diagram, and then implement the Suggested Remedy in my general comment concerning the structure of the draft amendment.

REJECT.

The diagram is based on the proposal "law_01_1108" that was adopted as the baseline for this section. The representation of PLS_CARRIER.indication adds clarity to the diagram without any ambiguity. This diagram would be present regardless of the document structure chosen.

Unsatisfied Negative Comments

Cl 35 SC 35.2.1 P 65 L 33 # 10201 Grow, Robert Intel

Comment Type ***TR***

I can't figure out what the last sentence is trying to specify. It also seems that the edits treat service primitives as logic signals. Service primitives are not logic signals, they are events and therefore can't remain in any state. Though the value sent in a primitive may have state, the primitive is only generated when the value changes state. So, it may not be best to use the term set in earlier sentences either.

SuggestedRemedy

If I understand the intent right, the following would be more accurate, though I don't believe there is a way to put timing requirements in the service primitives, (only in the layers that cause generation of the primitive) so the following isn't correct either (this needs thought and work):

An LPI_IDLE.request primitive with value ASSERT shall not be generated unless the attached link is operational (i.e. link_status = OK, according to the underlying PCS/PMA). The PHY shall not cause an LP_IDLE.request primitive with value ASSERT to be generated for at least one second following a link_status change to OK.

A similar problem exists in 46.1.7.

ACCEPT IN PRINCIPLE.

Accept the suggested remedy for this clause. Make a similar change for 46.1.7.

Also add a reference to 78.1.2.1.2.

Unsatisfied Negative Comments

Cl 78 SC 78.1.2.1.4 P 228 L 26 # 10202 Grow, Robert Intel

Comment Type ***TR***

Is signaling of LPI between an RS and its link partner, or between the RS and the lower parts of the PHY? If the PHY has no option to signal the request, then the language is appropriate, but it seems inconsistent with MII text describing the xMII signals. The effect of the primitive is to generate signals on the MII and that isn't specified here, but should be.

SuggestedRemedy

Assure MII clause are consistent in what layer is signaling to what peer layer, and that any additional requirements on conveying the LPI request in lower sublayers is properly represented. Add generic text that covers the three MII types -- how the assert or deassert is signaled, can probably be generic using the MII definition of assert low power idle.

ACCEPT IN PRINCIPLE.

The PHY has no option to signal the request so the language is appropriate however editor will look into adding clarifying text as in the suggested remedy.

Editor to check if that this is clear in the xMII clauses.

Unsatisfied Negative Comments

Cl ***78 SC 78.1.2.1 P 228 L 47 # 10203 Grow, Robert Intel***

Comment Type ***TR***

When generated is too generic.

SuggestedRemedy

The primitive is generated because of a change from something (xMII normal Idle to assert low power idle) and vise versa.

ACCEPT IN PRINCIPLE.

Adopt suggested remedy with editorial license to clear typos / grammatical errors.

Future Schedule

Estimated recirculation ballot open date	March 23, 2010
Estimated recirculation ballot close date	April 6, 2010
Proposed interim meeting date	April 8, 2010

If needed:

Estimated recirculation ballot open date	April 12, 2010
Estimated recirculation ballot close date	April 27, 2010
Proposed interim meeting date	May 24, 2010

We only expect to have one more recirculation

Motion

Request that the Working Group Chair request conditional approval to begin sponsor ballot on P802.3az to the EC.

M: M. Bennett S: V. Pillai

Y: 54 N: 1 A: 2 (technical, 75%) motion passed

802.3az timeline

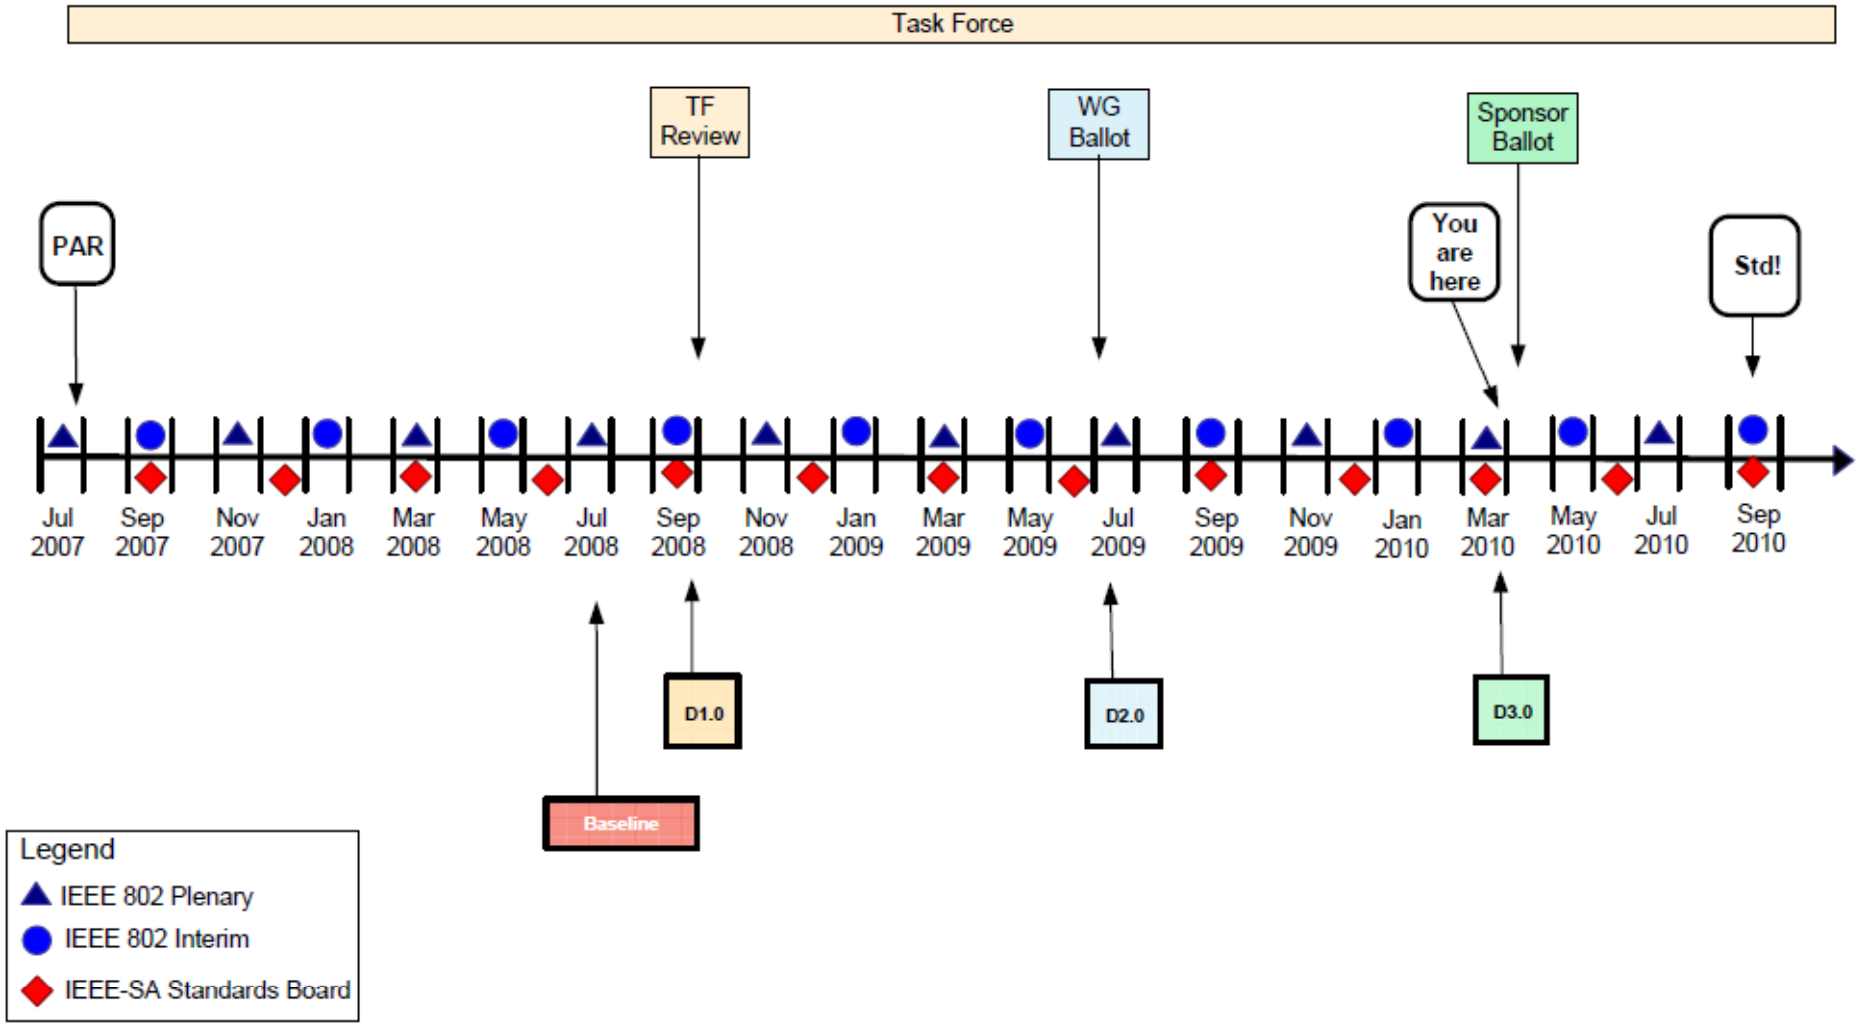

Objectives

Define a mechanism to reduce power consumption during periods of low link utilization for the following PHYs

– 100BASE-TX (Full Duplex)
– 1000BASE-T (Full Duplex)
– 10GBASE-T
– 1000BASE-KX
– 10GBASE-KR
– 10GBASE-KX4

• Define a protocol to coordinate transitions to or from a lower level of power consumption

• The link status should not change as a result of the transition

• No frames in transit shall be dropped or corrupted during the transition to and from the lower level of power consumption

• The transition time to and from the lower level of power consumption should be transparent to upper layer protocols and applications

Objectives

- **Define a 10 megabit PHY with a reduced transmit amplitude requirement such that it shall be fully interoperable with legacy 10BASE-T PHYs over 100 m of Class D (Category 5) or better cabling to enable reduced power implementations.**

- **Any new twisted-pair and/or backplane PHY for EEE shall include legacy compatible auto negotiation**

Thank You!