Edits to 802.3bq D2.2 to add 25GBASE-T

George Zimmerman

CME Consulting

Chief Editor

The road to 25GBASE-T

- Our PAR now has 25G in it! , we need to update the draft
- This presentation is a companion to the draft (and compare draft) posted to the private area:

  http://www.ieee802.org/3/bq/private/8023bq_D2p2_25G.pdf

  http://www.ieee802.org/3/bq/private/8023bq_D2p2_25G_CMP.pdf
- It summarizes the changes, themes & TBDs

Overview of changes to add 25G

- References to 25G in intro & front matter
- Add BASE-T to 802.3by 25G architecture
- Add 25GBASE-T to AutoNeg, EEE, Cl 30 & Cl 45 Management
- Add parallel references to 25G / 40G throughout Cl 113
- Add parallel references to 25GMII & account for differences with XLGMII in Cl 113
- Add frequency scaling for 25G/40G where appropriate

25G Architecture (802.3by)

- From marris_25gbt_01_0315.pdf:

802.3by architecture

- Nomenclature:
  - 25G-MII
  - 25G-AUI
- Roman numerals “XXV” considered too ugly so not used
- 25G-MII is a speeded up version of XGMII rather than a slowed down version of XLGMII. Therefore SOP occurs on 4-byte boundaries rather than 8-byte and local and remote fault encoding is slightly different from XLGMII.
- 25G-AUI is a single lane version of the C2C and C2M electrical interfaces defined in 802.3bm Annexes 83D and 83E

IEEE 802.3 25GBASE-T Study Group Berlin Plenary Meeting March 2015

25G Architecture in d2p2_25G

- 25GBASE-T defined to 25GMII
  - 32 bit 25GMII (2 xfers per 64/65b block) vs. 64 bit XLGMII (1 xfer per 64/65b) should be accounted for
  - Some text pending resolution of d2p2 “Architecture” comments
    - Resolution of these comments would need to result in parallel 25G text

Use of MultiGBASE-T

- Example: 45.2.3.14.2: (p46 of d2p2):

Change the text of 45.2.3.14.2 to include 40GBASE-T as follows:

45.2.3.14.2 Latched high BER (3.33.14)

When read as a one, bit 3.33.14 indicates that the 10/40/100GBASE-R, ~~or the~~ 10GBASE-T, or 40GBASE-T PCS has detected a high BER. When read as a zero, bit 3.33.14 indicates that the 10/40/100GBASE-R, ~~or the~~ 10GBASE-T, or 40GBASE-T PCS has not detected a high BER.

- Becomes (p50, of d2p2_25G):

Change the text of 45.2.3.14.2 to include 25GBASE-T and 40GBASE-T as follows:

45.2.3.14.2 Latched high BER (3.33.14)

When read as a one, bit 3.33.14 indicates that the 10/40/100GBASE-R or ~~the 10GBASE-T~~ a member of the MultiGBASE-T set PCS has detected a high BER. When read as a zero, bit 3.33.14 indicates that the 10/40/100GBASE-R or ~~the 10GBASE-T~~ a member of the MultiGBASE-T set PCS has not detected a high BER.

- Recommend adopt this and instruct editor to copy these edits to 802.3bz

Clause 45 strategy

- 25GBASE-T bit and register allocations roughly follow:
  - http://www.ieee802.org/3/bz/public/may15/Kim_3bz_01_0515.pdf
  - http://www.ieee802.org/3/by/public/adhoc/architecture/anslow_021815_25GE_adhoc.pdf
- Checked with and modified not to conflict with 802.3by d2p0, 802.3bq d2p2
  - 25Gb/s Ethernet used more bits, will need to continue tracking with Arthur Marris

Comment strategy

- Comments on d2p2 are not included in the 25G draft
- Except where noted, comment dependencies are not significant – just implement the same change for 25G is there is parallel text
- Recommend accepting the new 25G “MultiG” clause 45 text into 802.3bz as well (would need motion)

Dependency on Comment Resolution

- The draft, for comparison purposes, has all the existing defects commented on in 802.3bq d2p2
- Changes pending resolution of 802.3bq D2p2 comments on PTS (training reset)
  - Not enough bits if we retain PTS for 40G (& 25G)
  - 25GBASE-T autoneg bits are not defined for either LD or LP. (P51 L50), but PTS for 25G is still in Cl 113.
  - Allocation of bits for 25GBASE-T LP fast retrain ability, LD & LP 25G PTS (see pages 54 & 55 of 25G draft)
    - If we remove 40G PTS, then we do not have to allocate a new AN register to support 25G as well (P54 L1, P55 L3)
    - If we don't remove 40G PTS, then we need a new register, to support 25G. – recommend we reorganize bits along lines of separating 2.5G/5G from 10G/25G/40GBASE-T

Frequency Scaling & Link Segment

- Parameter “S” defined, as in 802.3bz
  - S= 1.0 for 40GBASE-T
  - S = 25/40 for 25GBASE-T
- Transmit PSDs, droop test, test mode frequencies, and SFDR specification all scaled by S
- Link segment parameters, only the upper frequency is scaled to 2000*S MHz
  - No parameter scaling
  - Per July ‘15 straw poll, f< 1250 MHz for 25GBASE-T)

Known oversights

- 25GBASE-T not added to introduction (p12 L19)
- 25GBASE-T not added to definitions (p24 L21)
- Consistency – 25G/40GBASE-T vs. 20/40GBASE-T (various)

Next steps

- Consider text, review until end of meeting
- Add any other known oversights to previous slide list
- Instruct editor to accept http://www.ieee802.org/3/bq/private/8023bq_D2p2_25G.pdf to resolve comment 192 as the basis for d2p3, with:
  - Implementation of comment resolution from WG ballot on IEEE 802.3bq d2.2
  - Editorial license to correct oversights on slide 11