

# EPoC FDD Downstream Spectrum & Channel Bonding

Ed Boyd, Broadcom

Randy Sharpe, Alcatel-Lucent

John Ulm, Motorola

Robert Howald, Motorola

Eugene Dai, Cox (Supporter)

John Dickinson, Bright House Networks (Supporter)

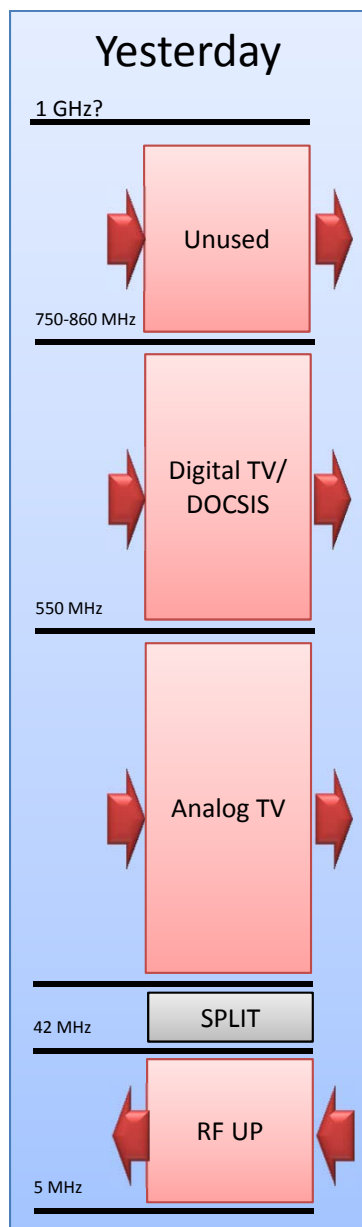
Bill Powell, Alcatel-Lucent (Supporter)

Charaf Hanna, STMicro (Supporter)

Rujian Lin, Luster Teraband Photonics, Shanghai (Supporter)

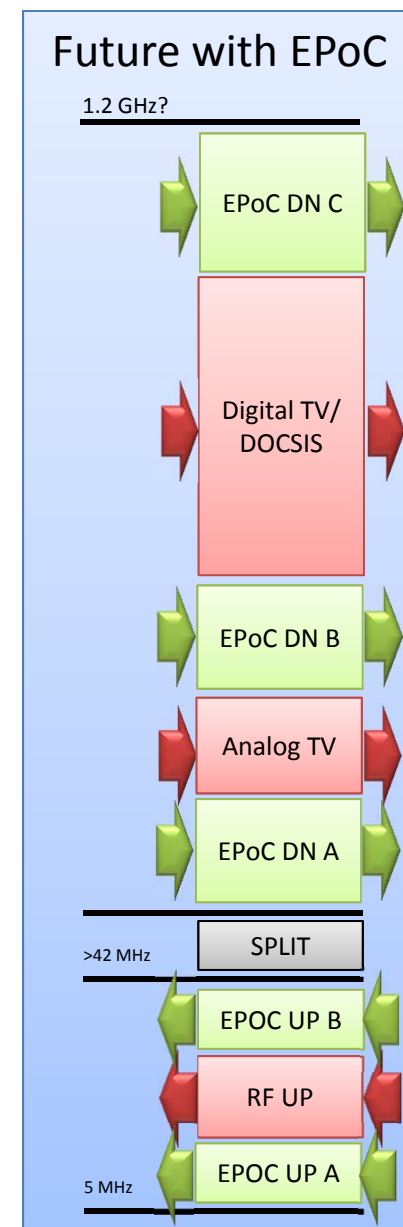
Geoff Thompson (Editor)

# EPoC FDD Requirements



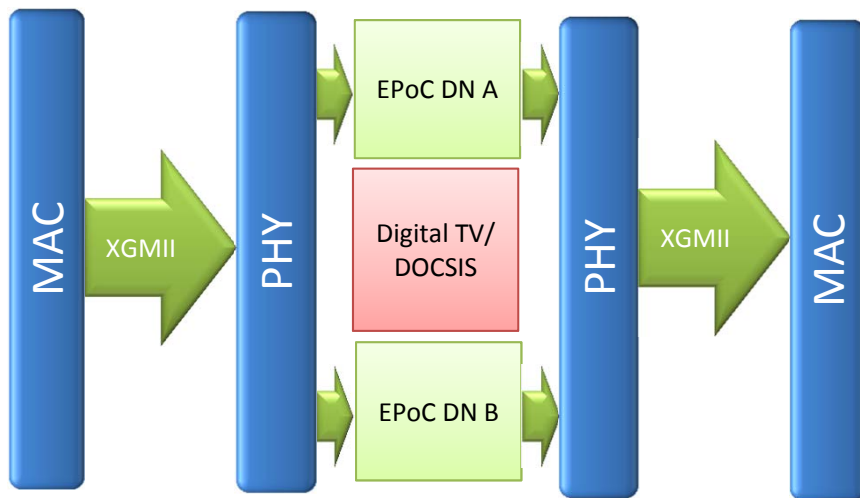
- EPoC should be flexible enough to exist anywhere in today's spectrum and maybe slightly above (up to 1.2GHz?)
- EPoC spectrum should support multiple blocks in different areas of the spectrum that act combined as a single channel.
- EPoC spectrum blocks should accommodate multiple channel bandwidths to fully occupy spectrum allocated for EPoC.

*How do we bond multiple channels?*

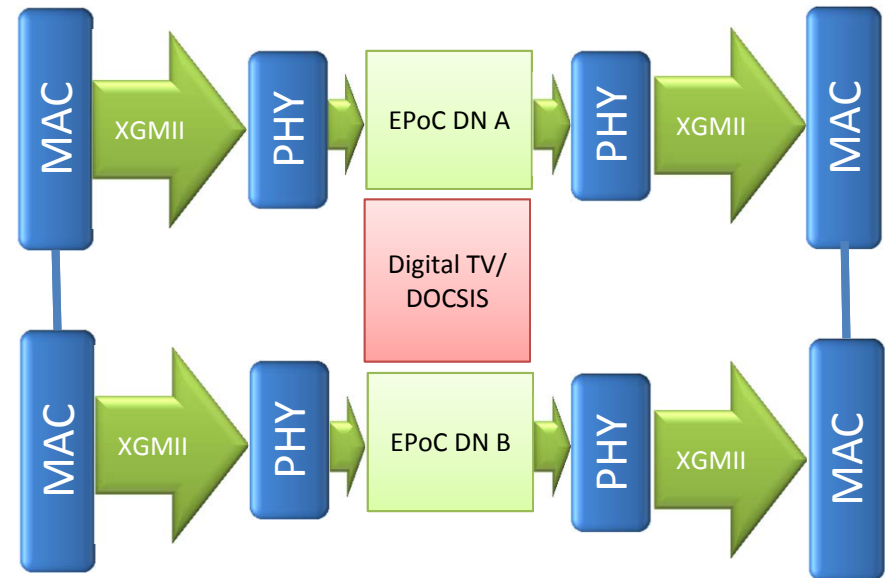


# Bonding Channel Options

PHY Bonding: Single Logical PHY Channel



MAC Bonding: Multiple PHY Channels



- Two options exist for combining multiple blocks of spectrum into a single logical pipe.
  - PHY Bonding: PHY can bond across a split in spectrum and present a single logical PHY pipe to the MAC. One XGMII.
  - MAC Bonding: MAC or higher layer functions can bond multiple MAC channels (more than one PHY and XGMII) on a packet by packet basis. [e.g. 802.1ax link aggregation] (This is out of scope for our task force but we need understand the limitations)

# Bonding Spectrum Blocks (PHY or MAC)?

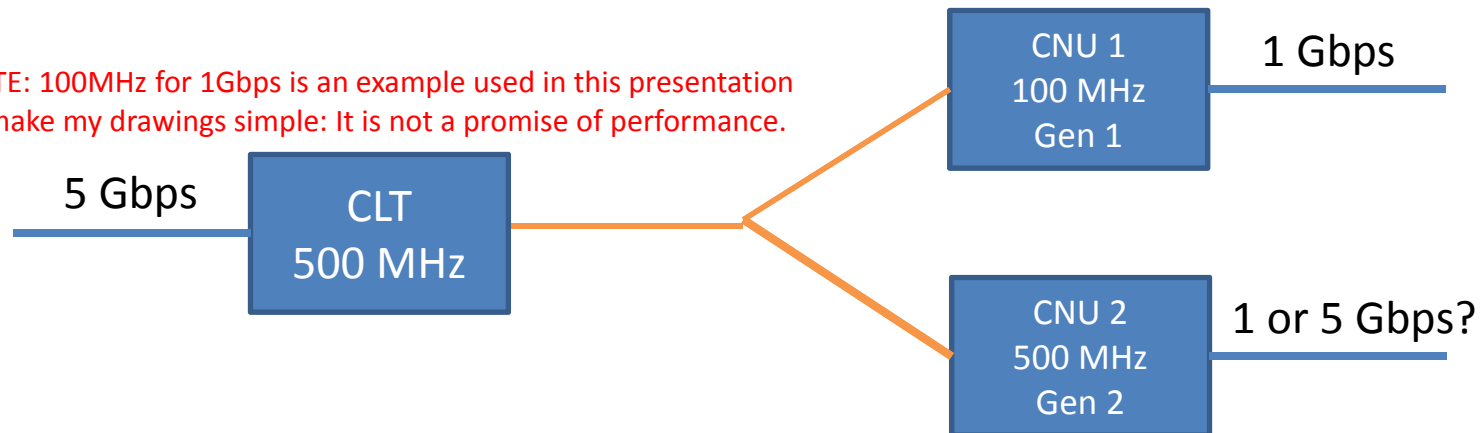
- MAC Bonding
  - Link aggregation (LA) assigns packets to links based on Ethernet DA/SA, IP DA/SA, or a combination of other header fields. (about 70% efficient – See Geoff T)
  - Frame re-ordering on flow is a performance issue in IP and not allowed in other protocols.
  - Link Aggregation requires links of the same speed.
  - GATEs need to have the same delay down different channels or need to stay in a single channel. (Single RTT for a CNU) What about other time control frames (i.e. 802.1AS)?
- PHY Bonding
  - Distribution of data bits into individual spectrum blocks is done irrespective of the Ethernet DA/SA or other frame header.
  - Lower FEC overhead since fewer block terminations.
  - Better statistical multiplexing and predictable performance
- PHY Bonding is much simpler
  - No packet size dependency, no need to look at DA/SA or other header information
  - Wider single pipe has shorter delay than multiple small pipes

*PHY bonding is a better solution*

*(What about devices with different capabilities?)*

# Multiple Rate EPOC Devices


NOTE: 100MHz for 1Gbps is an example used in this presentation to make my drawings simple: It is not a promise of performance.



- If CNUs support different spectrum capacities, how do we handle them in EPoC?
- The drawing above shows 2 generations of CNUs. CNU 1 supports 100MHz (1 Gbps) and CNU 2 supports 500 MHz (5 Gbps)
- PHY layer bonding will not work if full downstream is not received by all CNUs since portions of packets for Gen 1 CNU are not guaranteed to be within the spectrum supported by Gen 2 CNU.

*What are the PHY layer and MAC layer bonding options?*

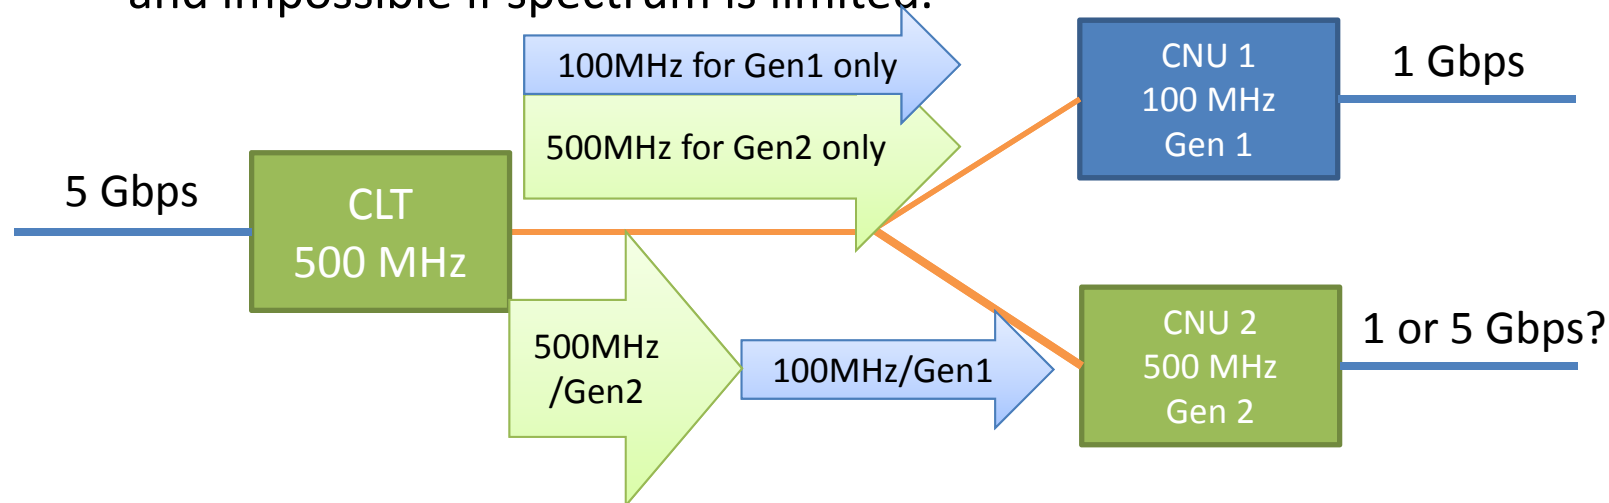
# PHY: Lowest Common Denominator



- In 10/100/1000 Ethernet, a 10/100/1000 Ethernet device goes down to speed of the lowest speed device on the link.
- This solution works well on point-to-point switched network.
- In the EPoC example above, all CNUs must be 500MHz (5Gbps) on the CLT or the network is limited to 100MHz (1 Gbps)
- It is not practical to expect a full network upgrade so this solution is not good. (i.e. replace all Gen1 CNUs at one time)

# PHY: EPON-like Approaches

- In EPON Downstream, 1Gbps and 10Gbps co-exist by carrying both signals on separate wavelengths on the fiber. (e.g. the EPoC CLT would dedicate spectrum for each generation device) – Dedicated Spectrum is not efficient and impossible if spectrum is limited.



- In EPON Upstream, 1Gbps and 10Gbps co-exist by using different timeslots on the fiber. (e.g. EPOC CLT could do 500MHz for CNU2 packets/bursts, 100MHz for CNU1/broadcast/Multicast bursts) – Inefficient use of the coaxial cable spectrum.

*PHY Bonding options for mixed speeds aren't good: MAC Bonding must be considered*

# PHY+nxMAC: 802.3ah DSL Approach

- Designed to aggregate DSL links.
- Provides QoS and adds resiliency.
- PME Aggregation Function (PAF) aggregates 32 PMEs to single logical Ethernet Link (MAC port).
- Cross connect function based on load.
- Transmits Packets are chopped into 64-512 Byte Blocks and a two byte header is added before switching to PME.
- Header contains sequence #, start pkt bit, end pkt bit
- Receiver Blocks are switched to aggregators and re-sequenced back into packets.
- MAC port rate controlled by CRS on MII.

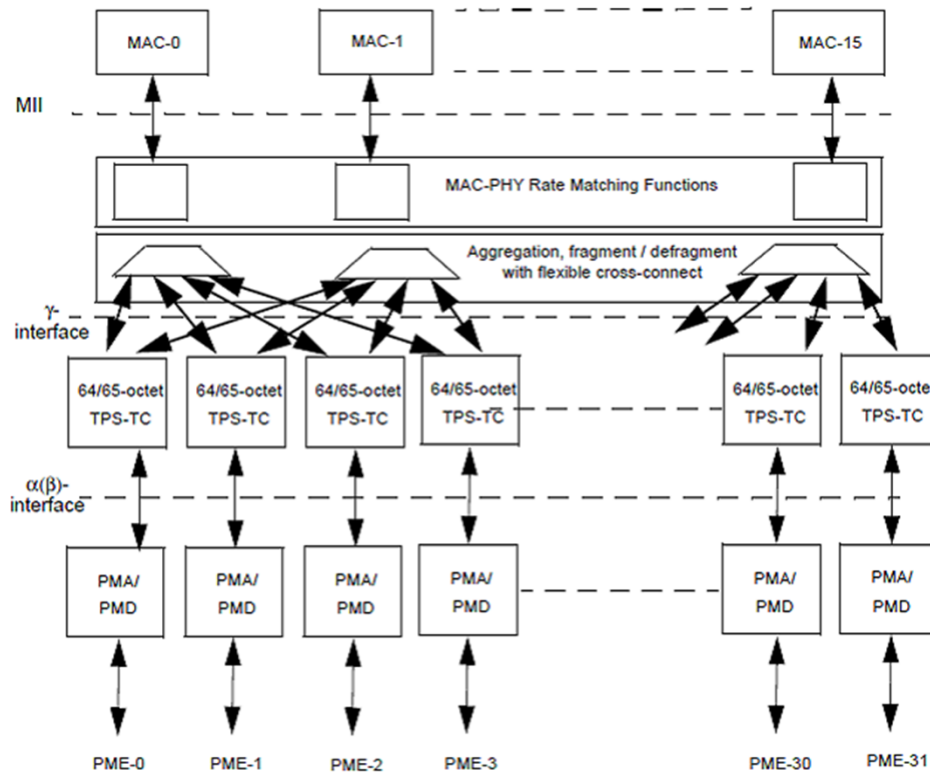
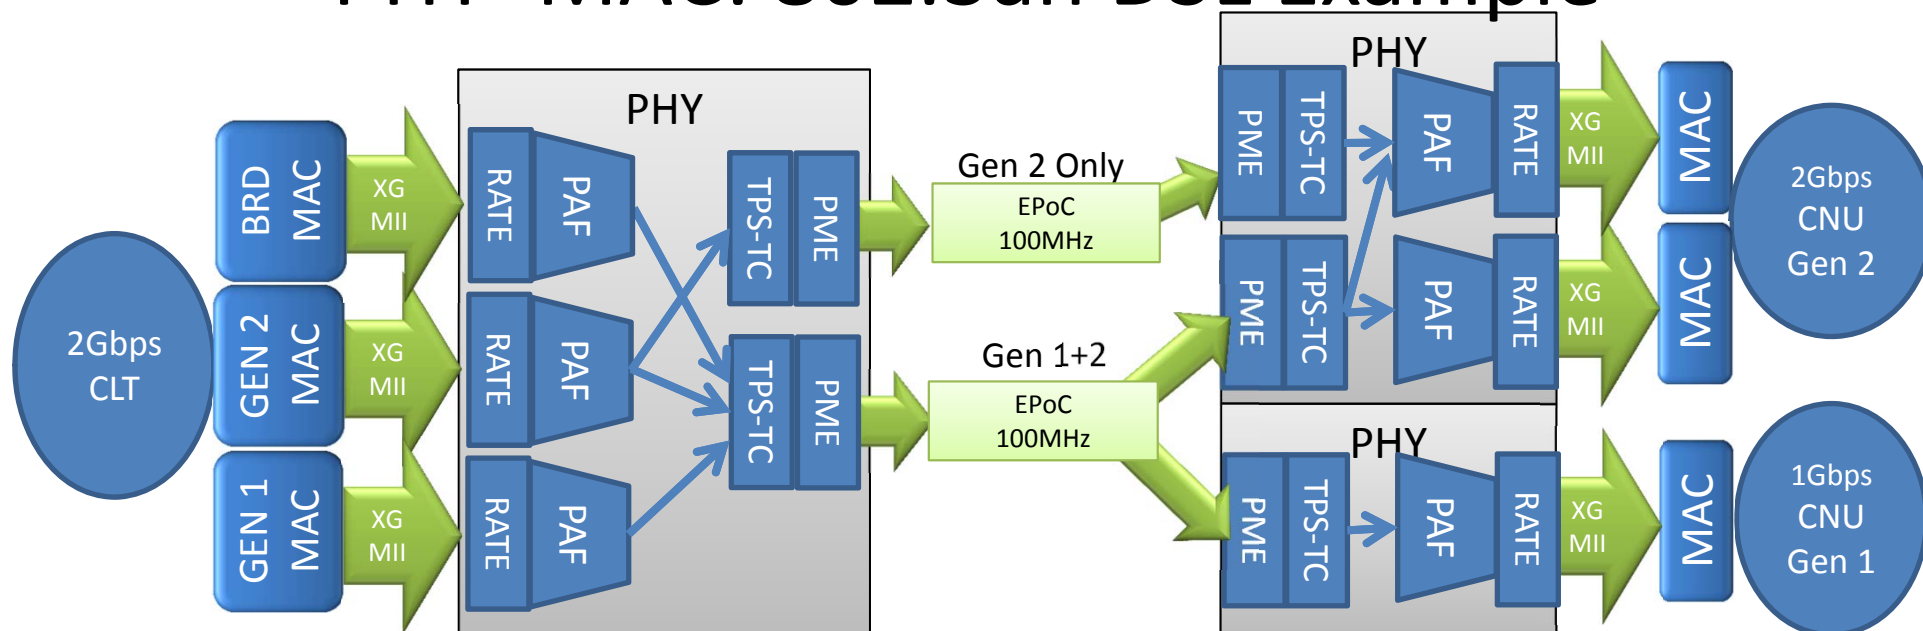


Figure 61-5—4 PME for each 2 MII connectivity

*802.3ah DSL Contains a Cell Switching Function in the PHY*



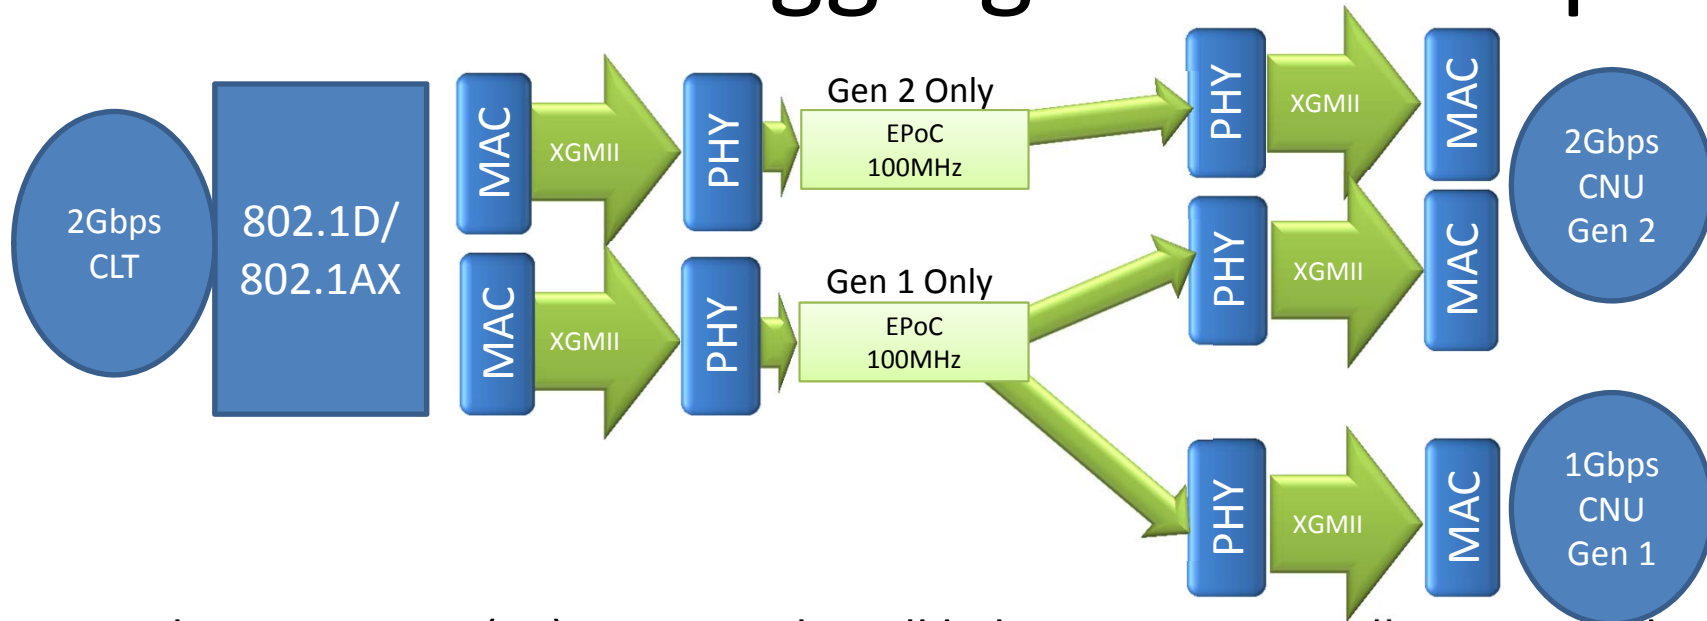
# PHY+MAC: 802.3ah DSL Example



- PME for spectrum blocks.
- PAF (MAC) for each generation of device.
- Load balancing of blocks into PMEs.
- Variable rate of data to XGMII based on load balancing.
- EPoC constant delay requirement would be challenge for this solution.
- EPoC is based on MII (not GMII) so no CRS exists for back pressure.
- Buffering and Delay impact needs to be understood.

*Complex Solution, does this make sense for EPoC?*

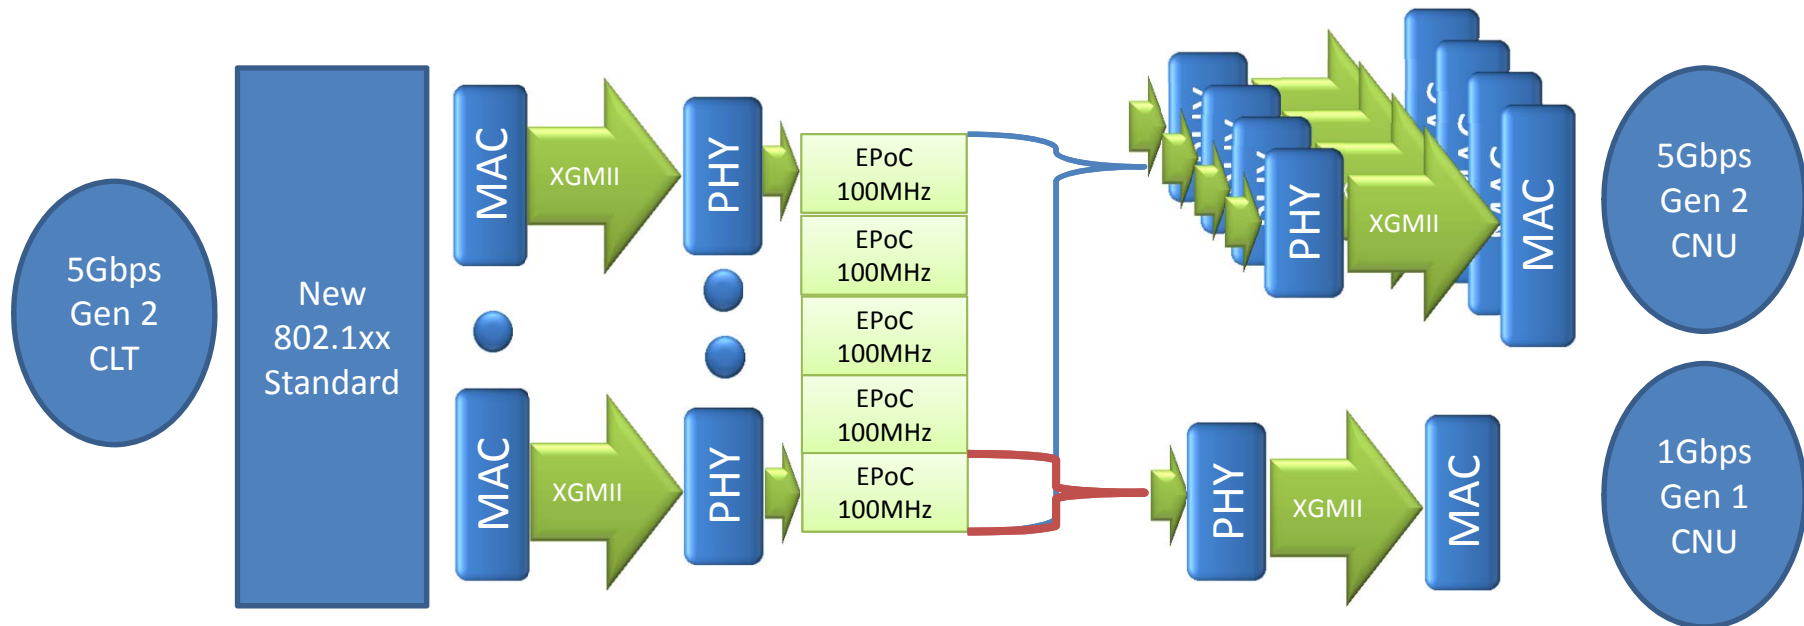
# MAC: Link Aggregation Example



- Link Aggregation (LA) assumes that all links terminate on all ports and the same speed for all links. EPoC support would require modification of LA.
- Broadcast always goes down common link; multicast uses L2/L3 snooping(?)
- LLIDs to Gen 1 CNU will go down common link.
- Packet to LLIDs on Gen 2 CNU can be divided into the 2 links using LA.
- LA based on L2 DA/SA (no L3 due to encryption) will not be an even distribution of traffic. Link Congestion will occur.

*Packet to Channel Assignment should be dynamic based on Channel Load*

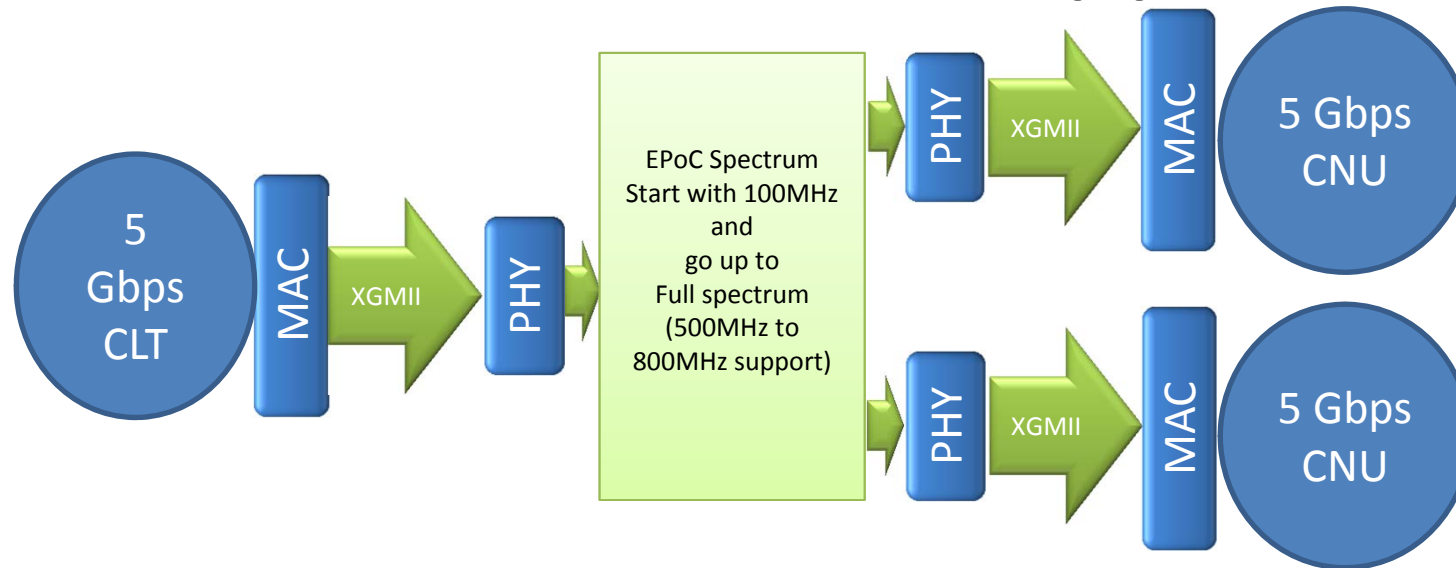
# MAC: Dynamic Channel Assignment Example



- Unicast, Multicast, and Broadcast packets switched based on destination to MACs connected to individual PHY channels
- Packets are switched into least crowded MACs to a destination
- Packets will arrive out of order based on different rate channels and variable packet sizes.
- 1G CNU's can be load balanced across multiple 100MHz channels.
- DOCSIS like solution with packet order marking and re-ordering buffers needed.

*New 802.1 Standards? Do we really want to go in this direction?*

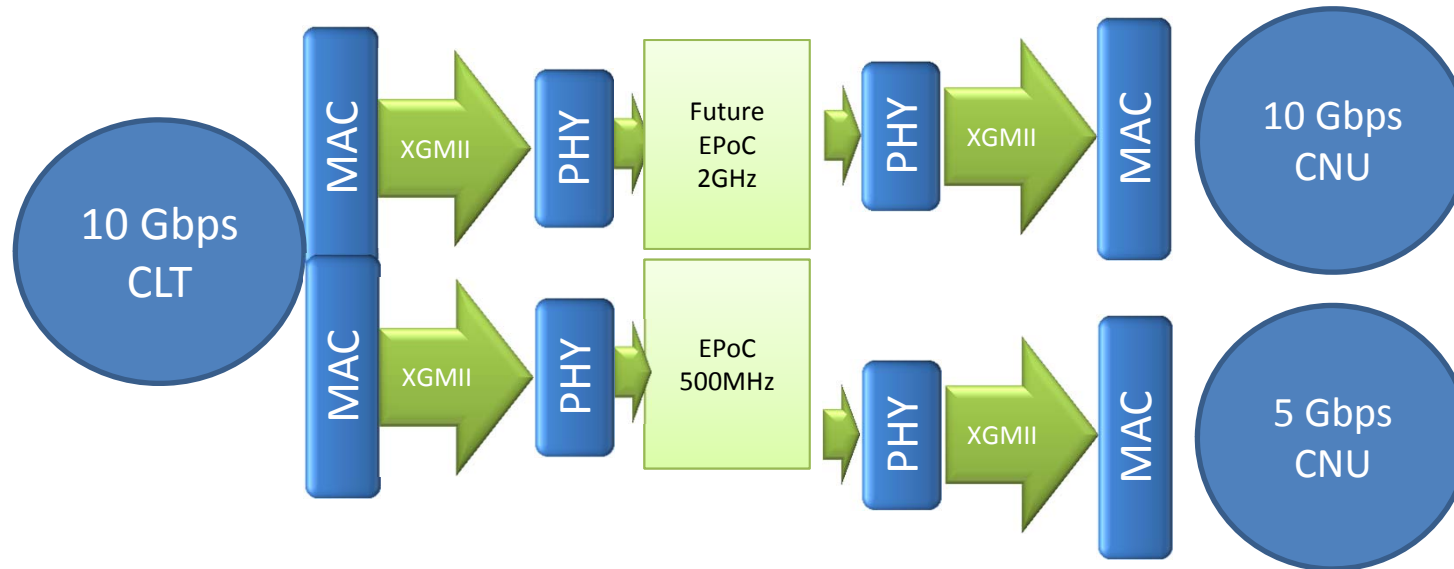
# Full Channel Support



- Specify a full channel of 500MHz to 800MHz. The single wide channel is simple to specify and design.
- EPoC shouldn't be optimized to compete with devices below 1 Gbps.
- It is economically and technically feasible to build devices to cover 100's of MHz of spectrum today at price targets similar to 1 Gbps devices.
- MAC Channel bonding adds additional delay, cost, and complexity and requires a new standard outside the scope of 802.3bn.
- A high performance, flexible, and configurable full band solution gives EPoC a distinct identity.

*Isn't this what the operators want to see?*

# Full Channel Future?



- If we specify a wide channel, what happens in the future if more frequency becomes available (i.e. 1GHz to 3GHz)?
- New PHY standard could use 1 to 3GHz to add 10Gbps downstream.
- 1G/10G EPON method of double downstream channel with CNU only tuning into a single channel would work well in this scenario.
- With a large channel, the penalty from duplicating broadcast and stranded spectrum is balanced with the advantage of CNU tuning in a single channel.

*New 802.1 channel bonding standard could be avoided*

# PHY Channel Aggregation Background:

- Assumption: PHY built from basic FFT building block
  - Single FFT block will probably span 120MHz to 200MHz
    - Translates to capacity range of ~1Gbps to ~2Gbps
  - Sub-carriers may be 'silenced' to create picket fence around existing QAMs
    - Net capacity per FFT block may be limited to spectrum available
      - E.g. 24MHz available may mean only ~200Mbps for particular FFT block

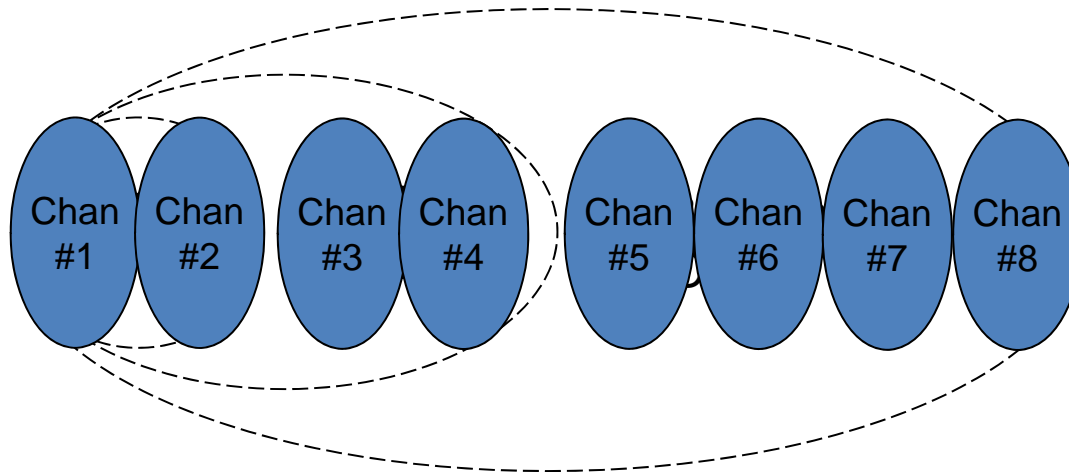
## Other Considerations:

- Frequency agility across 1GHz DS requires multiple OFDM FFT blocks
  - E.g. see slide 2 for example needing three FFT blocks
  - Lots of Frequency agility moves us towards more, smaller FFT blocks
- 2.5G, 5G or 10Gbps total capacity requires multiple OFDM FFT blocks

*Q: Is each OFDM FFT Block a separate PHY channel?*

# One PHY Channel per FFT block:

PHY Channel Aggregation Strategy #1: Each FFT Block is a PHY Channel



## Example Total Capacities:

1<sup>st</sup> Gen CNU => ~2.5Gbps, FFT/Chan 1+2

2<sup>nd</sup> Gen CNU => ~5Gbps, FFT/Chan 1+2+3+4

3<sup>rd</sup> Gen CNU => ~10Gbps, FFT/Chan 1 to 8


# One PHY Channel per FFT block:

- Option 1: PHY Channel per each OFDM FFT block:
  - Capacity for each FFT block may vary significantly depending on spectrum
  - Max Packet skew buffers proportional to difference from min to max
  - What is minimum capacity of OFDM FFT block?
    - Assuming 2KB max packet:
      - 6MHz min means 32:1 difference to 192MHz max FFT block; or 64KB buffer
      - 24MHz min means 8:1 difference; or 16KB buffer
  - Some issues with this approach:
    - Increases latency and packet skew for packet order re-assembly;
    - Doesn't scale well as we continue towards 10Gbps
      - 10G modem with 8 FFT blocks needs 7X buffer requirements (1 small, 7 big)
    - Larger buffers increases modem costs



## One PHY Channel across multiple FFT blocks:

PHY Channel Aggregation Strategy #2: Single PHY Channel seen by MAC



### Example Total Capacities:

1<sup>st</sup> Gen CNU => ~2.5Gbps, Chan 1

2<sup>nd</sup> Gen CNU => ~5Gbps, Chan 2

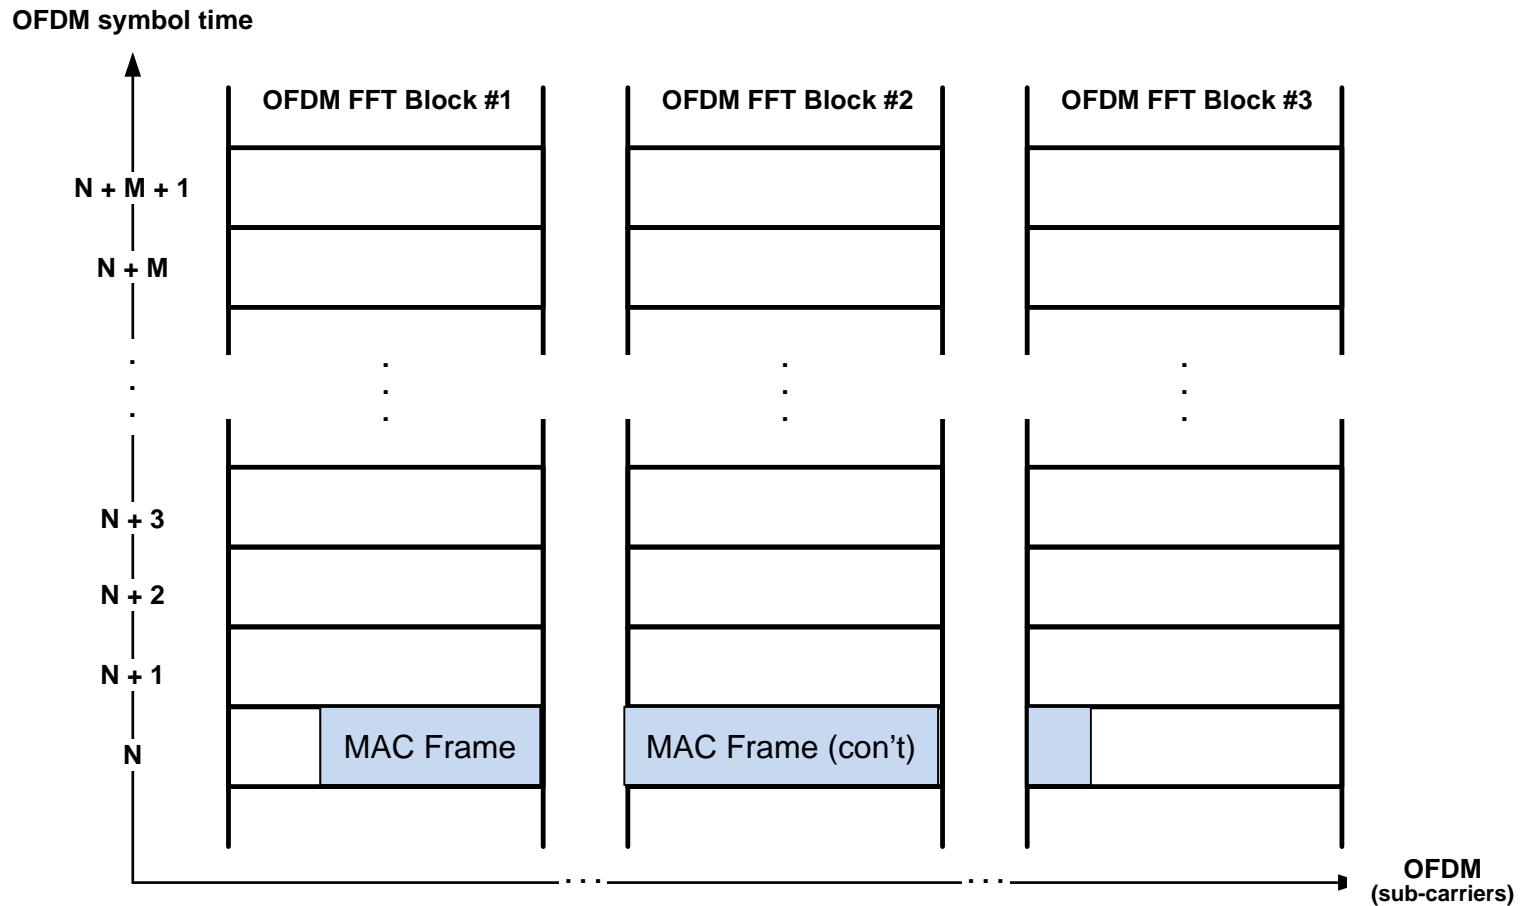
3<sup>rd</sup> Gen CNU => ~10Gbps, Chan 3

## One PHY Channel across multiple FFT blocks:

- Option 2: Single PHY channel across multiple FFT Blocks
  - MAC Frame spread across multiple OFDM FFT Blocks if needed
    - appears as single large pipe to MAC:
      - reduces packet skew + buffers; simplifies QoS, bonding
    - For 1<sup>st</sup> gen CNU, no packet re-ordering buffer needed
    - appears as single pipe with fixed FIFO packet ordering
      - no MAC bonding needed!
  - Questions for future gen CNU: larger overlapping PHY channels?
    - Do we need to aggregate across very large PHY channels?
    - Or can we TDM over shared spectrum?

**Recommendation: MAC Frames can span  
multiple OFDM FFT 'Blocks'**

## One PHY Channel across multiple FFT blocks:



- Challenge: synchronize symbol time across multiple OFDM FFT blocks
  - If FFT subcarrier spacing sufficiently small (e.g. 10KHz), can MAC Frames be contained within single symbol time of a single FFT block??

## One PHY Channel across multiple FFT blocks: FEC

- Challenge: If MAC Frames cross FFT blocks, does FEC also??
  - Requires further investigation
  - Considerations:
    - Keep FEC blocks and MAC frames decoupled if possible
    - Does having FEC cross FFT blocks compromise robustness?
      - need input from channel model
    - Should FEC blocks align with symbol times; or multiple of symbol time?
      - support FEC blocks of different lengths

**Further Study: Common FEC across multiple FFT Blocks**

# Channel Bonding Conclusions

- PHY Bonding is simpler and better performance than any of the MAC Layer (packet) Bonding options.
- PHY Bonding with different generations of devices is not practical.
- MAC Bonding would likely require a new 802.1 standard and is out of scope for 802.3bn.
- A new 802.1 standard would make significant differences above the PHY Layer.
  - CLTs will have different switching than OLTs. CNU's are different than ONUs.
  - Loss of shared OLT & ONU volume will increase the cost for CLT & CNU. (It isn't just a PHY!)
- Covering the entire cable spectrum (Full Channel) is the simpler solution. Bandwidth grows to all CNU's when spectrum is made available.
- Additional discussion needed on minimum spectrum, range of spectrum to cover, and notching spectrum.

*EPoC should support a wide spectrum with PHY layer bonding*

*Thank You!*