EPoC
and
Multiple PHY Generations

802.3bn
Orlando, FL
March, 2013

Edwin Mallette
Bright House Networks

Supporters

- Ed Boyd, Broadcom Corporation
- Matt Schmidt, Cablelabs
- Eugene Dai, Cox Communications
- Marek Hajduczenia, ZTE

NOTE

- This presentation focuses on future products that would come out of the EPoC effort.
- This presentation DOES NOT provide guidance that would help us converge on PHY baselines.
- This presentation DOES intend to provide some food for thought regarding support of multiple product generations.
- This presentation focuses on the downstream/FDD multi-generation case only.

Introduction

- So … what is a technology generation ?
- For this presentation, a technology generation is the time delta between the time we start buying a product feature set and the time we start buying the next product feature set.
  - For illustrative purposes, I use our actual timeframes for DOCSIS3.0 modems.
  - When we started buying 4/4 modems, when we started buying 8/4 modems, and the future look at buying 16/4 (or 24/4 or 32/4) modems.

Generation Purchase Cycle

- At BHN, we still purchase D2.0 modems in some embedded applications including DSG and MTAs.
- We also purchased D3.0 4/4 modems for a short period of time and then shifted to 8/4 modems.
- In the coming year we're looking at shifting again to 16/4 modems.
- There's a possibility that due to market conditions and other intangibles we might end up purchasing some combination of 16/4, 24/4 or 32/4 with a worst case scenario of 3 NEW generations.
  - The timeframe for this worst case might look like 16/4 in Q3, 2013, 24/4 in Q1, 2014, and 32/4 in Q3/2014.

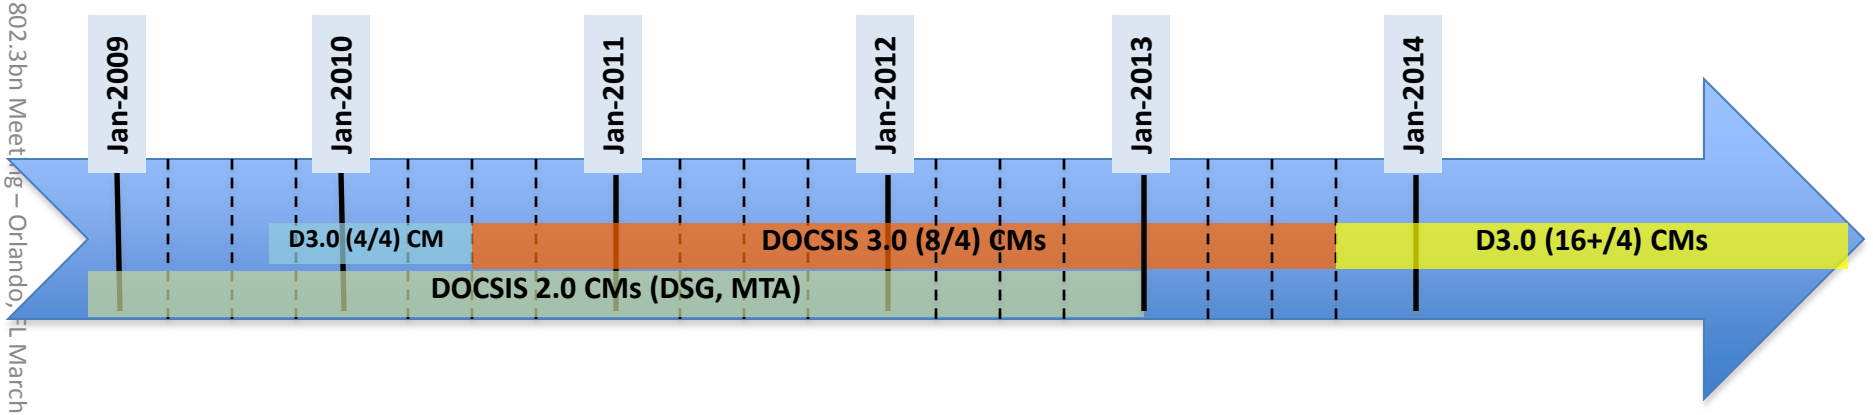

Product Lifetimes

- So product generations are short, but what is a product lifetime ?
- A product lifetime is how long product exists within our production network.
- A product lifetime of a residential CPE is typically around 10 years.
  - Ultimately if we can keep making the product work in the field, we will!
  - We still have settops manufactured in 1999 in the field.
  - We still have DOCSIS 1.1 (specification date of 1999) modems deployed.
  - It doesn't matter that our depreciation schedule for CMs is 3 years.
- So if we have a product generation of 2 years and a product lifetime of 10 years we end up supporting ~5 product generations simultaneously.

Generation Support Timeline

- In general this timeline only takes into account broad categories of DOCSIS capabilities.
- There are a number of finer details that we don't explore (e.g. RF receiver front end, etc.)
  - Some of these finer details do have support implications WRT, e.g. how far apart the downstream channels can be – within 60MHz, within 100MHz.
- These generations are all supported ON THE SAME plant and within the same set of downstream channels simultaneously.

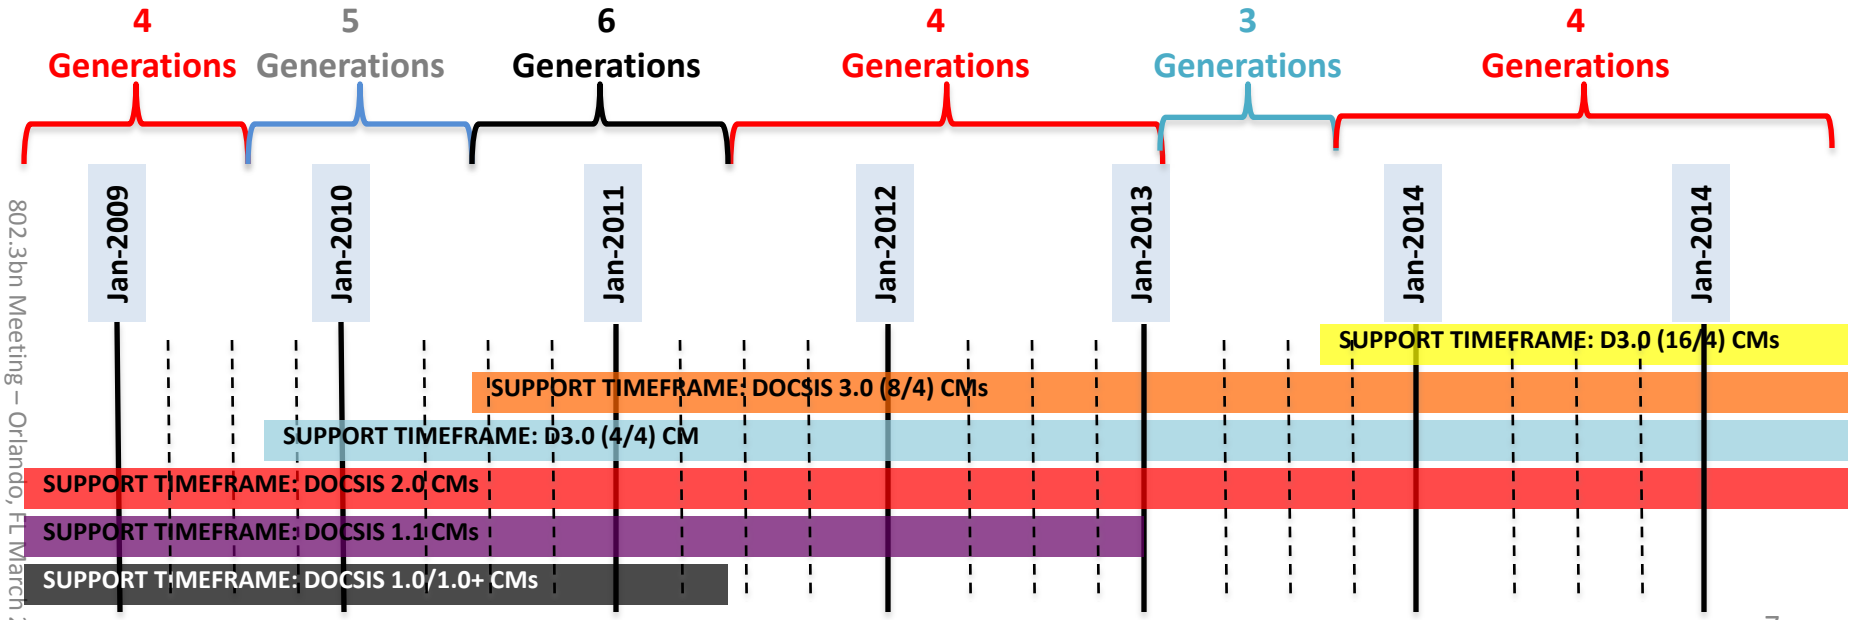

EPoC and CNU Generations

Enough about Docsis … How's this apply to epoC ?

Forecasted BW Top Tier Growth

- Some cable operators are deploying 100Mb/s service on CMs today (Early-2013.)
- If we assume the CAGR growth applies to our tier we have ~50% CAGR.
- For sake of argument, let's assume EPoC products arrive bright and early in 2015.
- Starting from 100Mb/s in 2013, we look 5 years into the future to 2020, we see a service speed of 1.7Gb/s with an associated aggregate speed to a service group of 3x that ~5.1Gb/s.

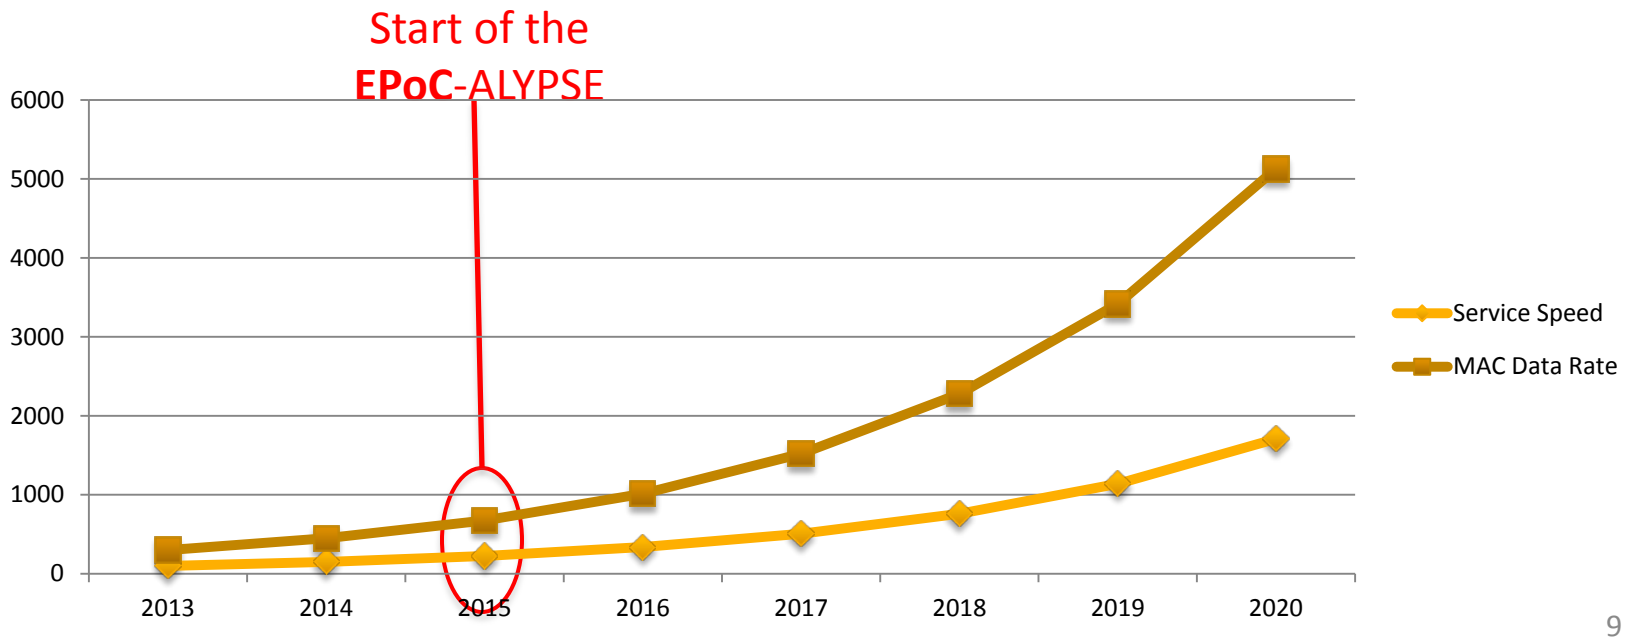

EPoC Shorter Generation

- Suppose on we start with 1 channel (GEN.1) initially, then grow to 2 channels in GEN.2, 4 channels in GEN.3…
- Before we've gone four years we would end up supporting some CNUs with one downstream channel, some with 2 downstream channels and some with 4 downstream channel…
- Sounds familiar?

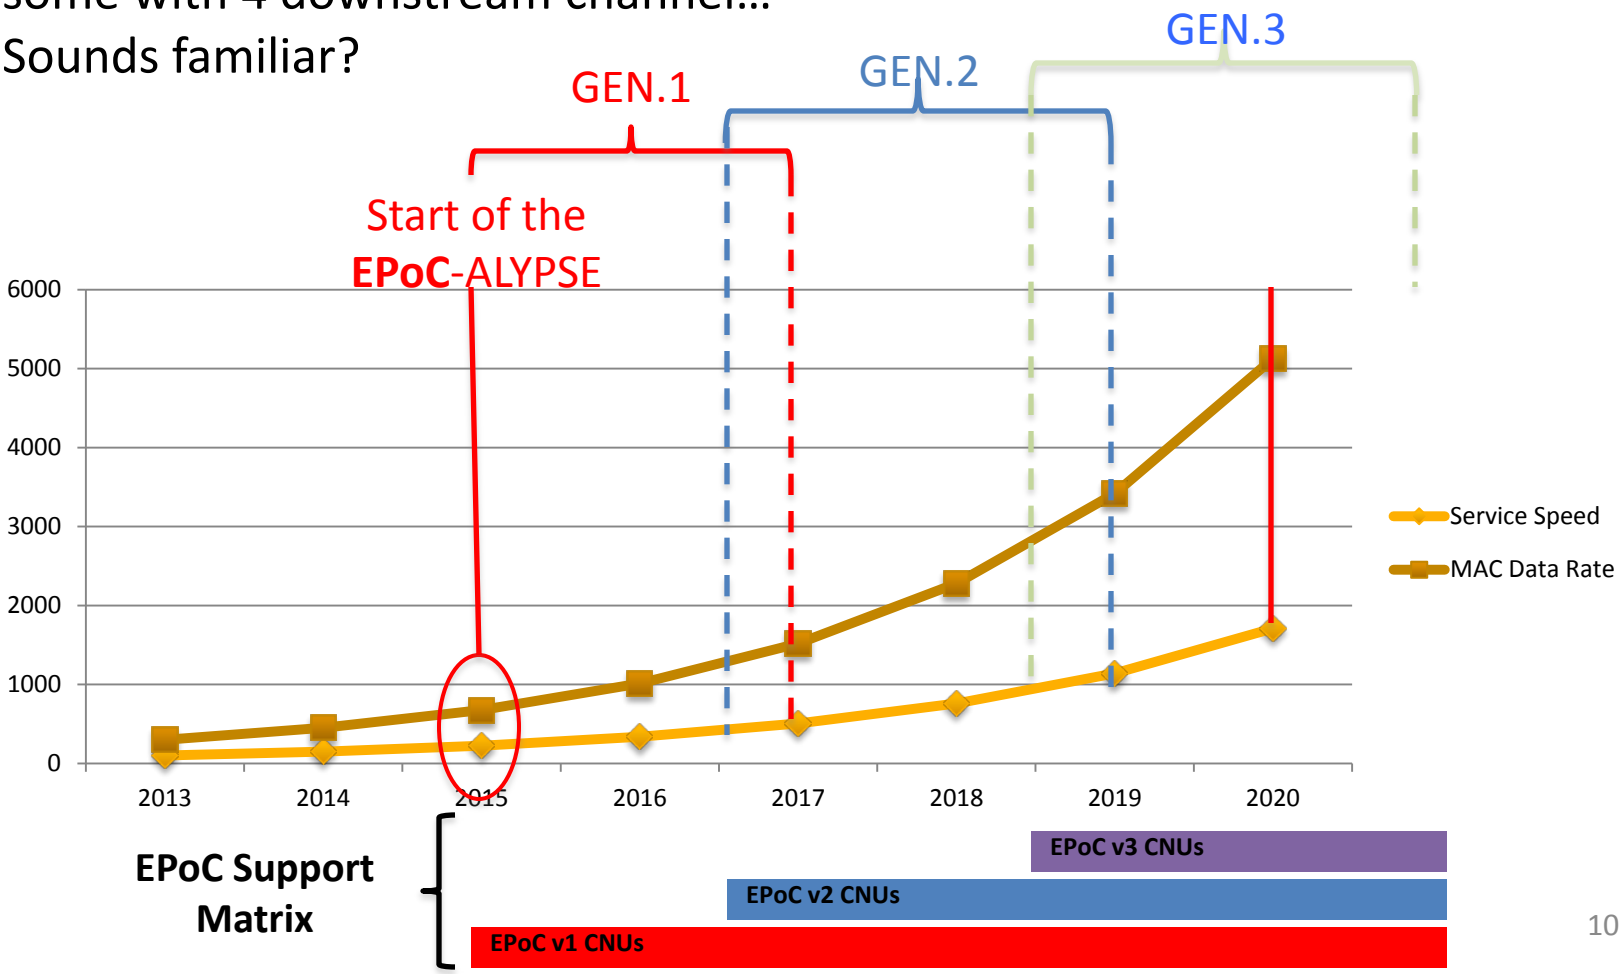

EPoC Long Generation

- On the other hand, suppose we want a generation to be viable to support future tiers up to 5 years from when product acquisition begins.
- This means a CNU would need to support 3 × 192MHz blocks in initial product.
- 2 blocks would get us to 2019 only if we had a very clean plant – 12 bits / Hz.

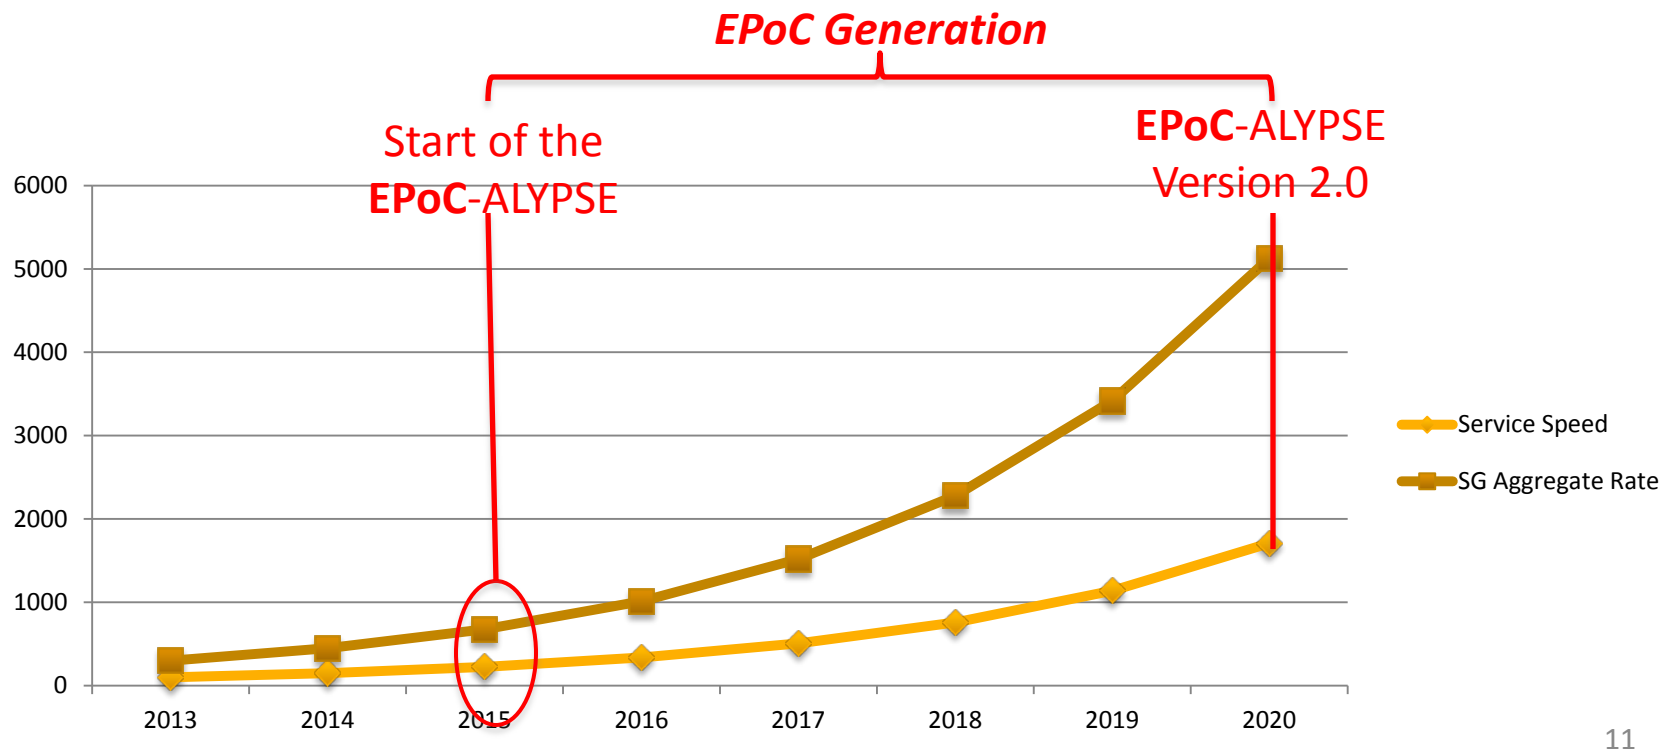

EPoC and CNU Generations

WHAT WOULD MULTIPLE GENERATIONS LOOK LIKE ?

Downstream Multi-Generation PHY GEN.1

Generation 1

All CNUs support single channel.
CLT / FCU EPoC PHY provides single downstream channel

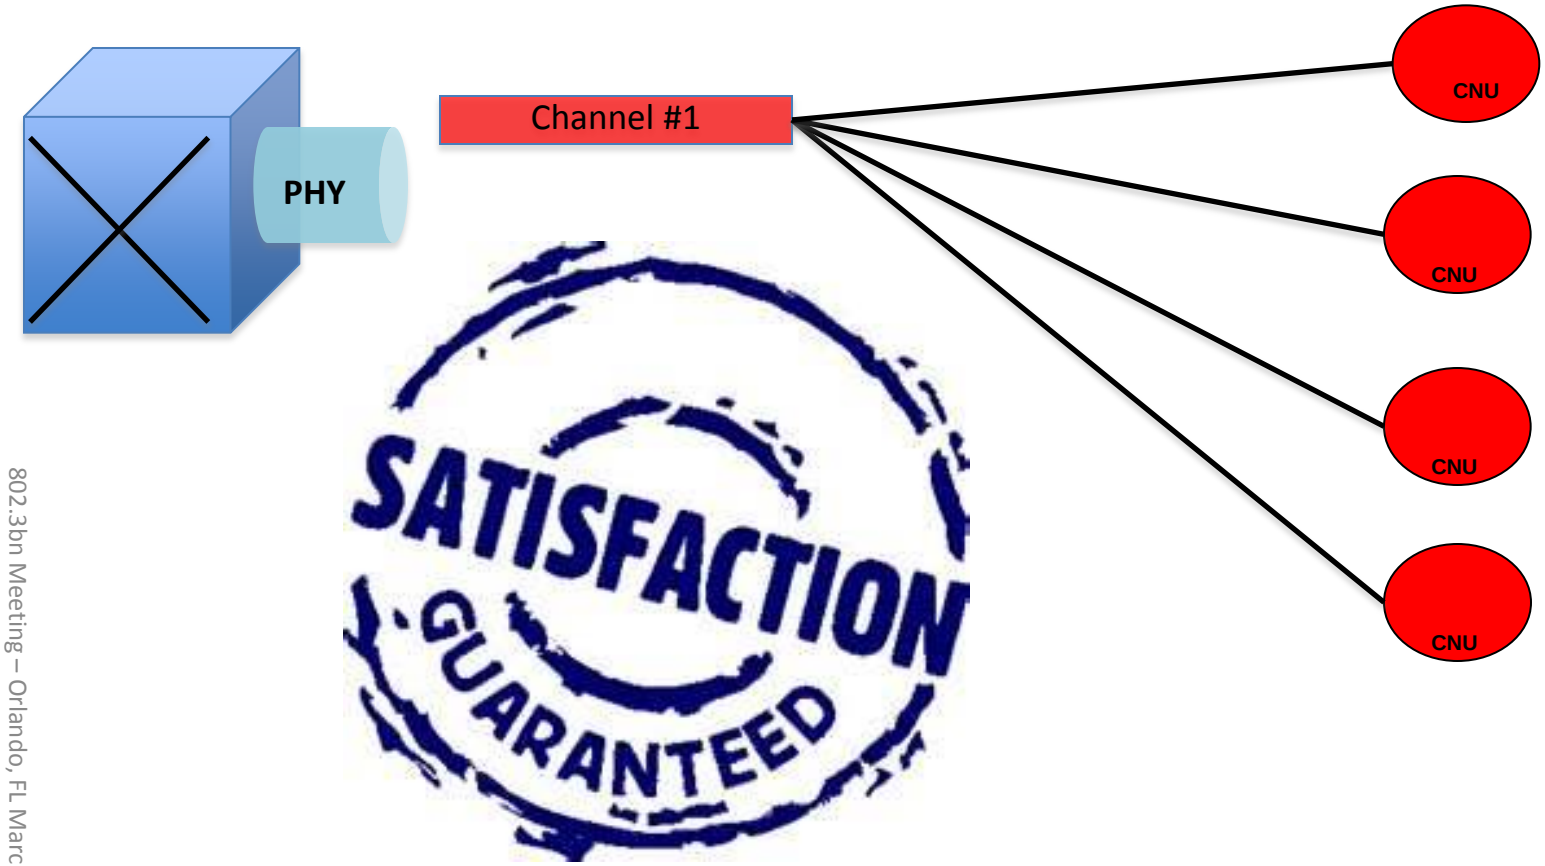

Downstream Multi-Generation PHY GEN.2

Generation 2

Some CNUs support 2 channels, some 1 channel. CLT / FCU EPoC PHY provides 2 downstream channels

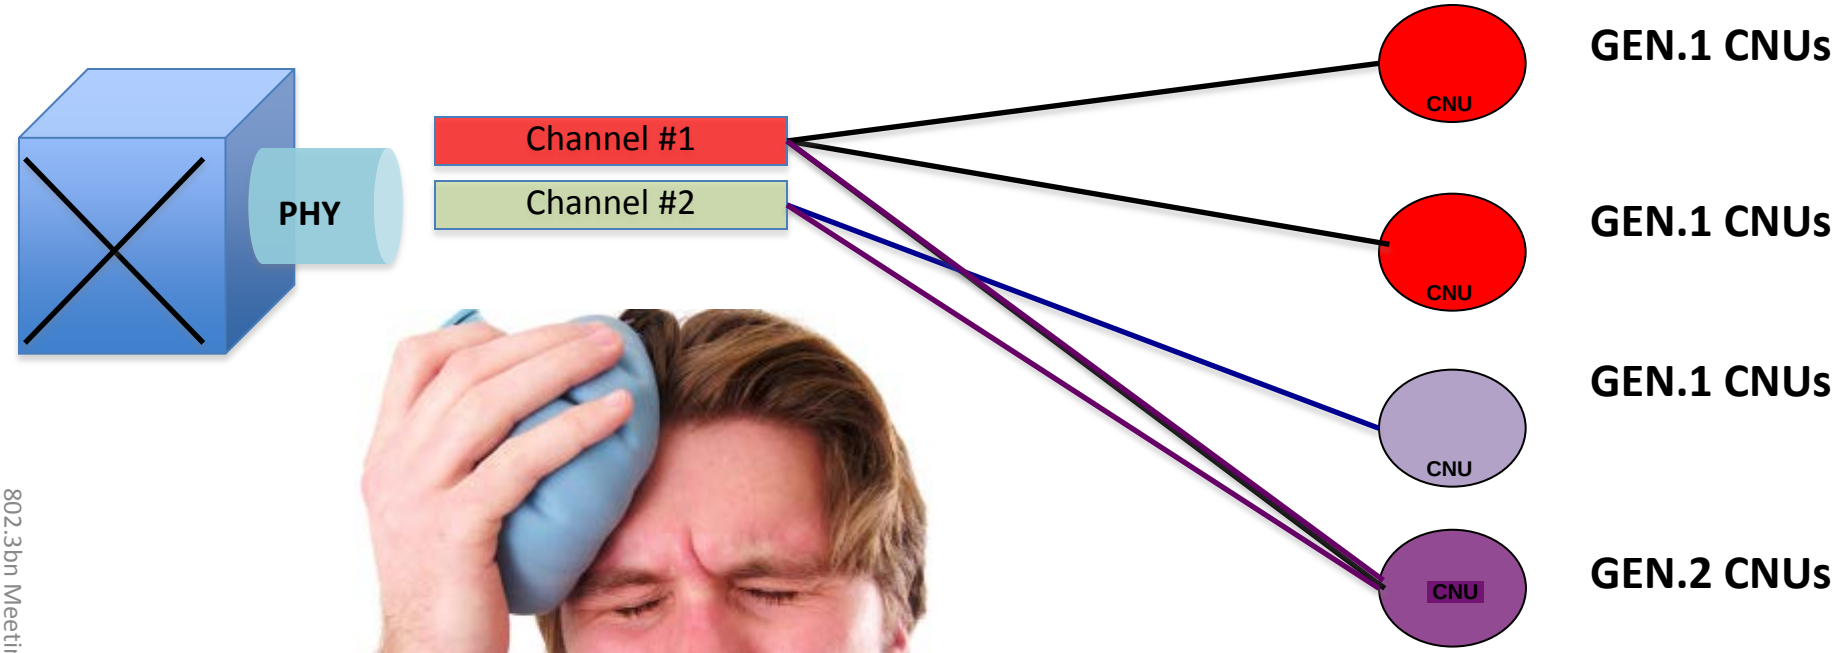

Downstream Multi-Generation PHY

GEN.3

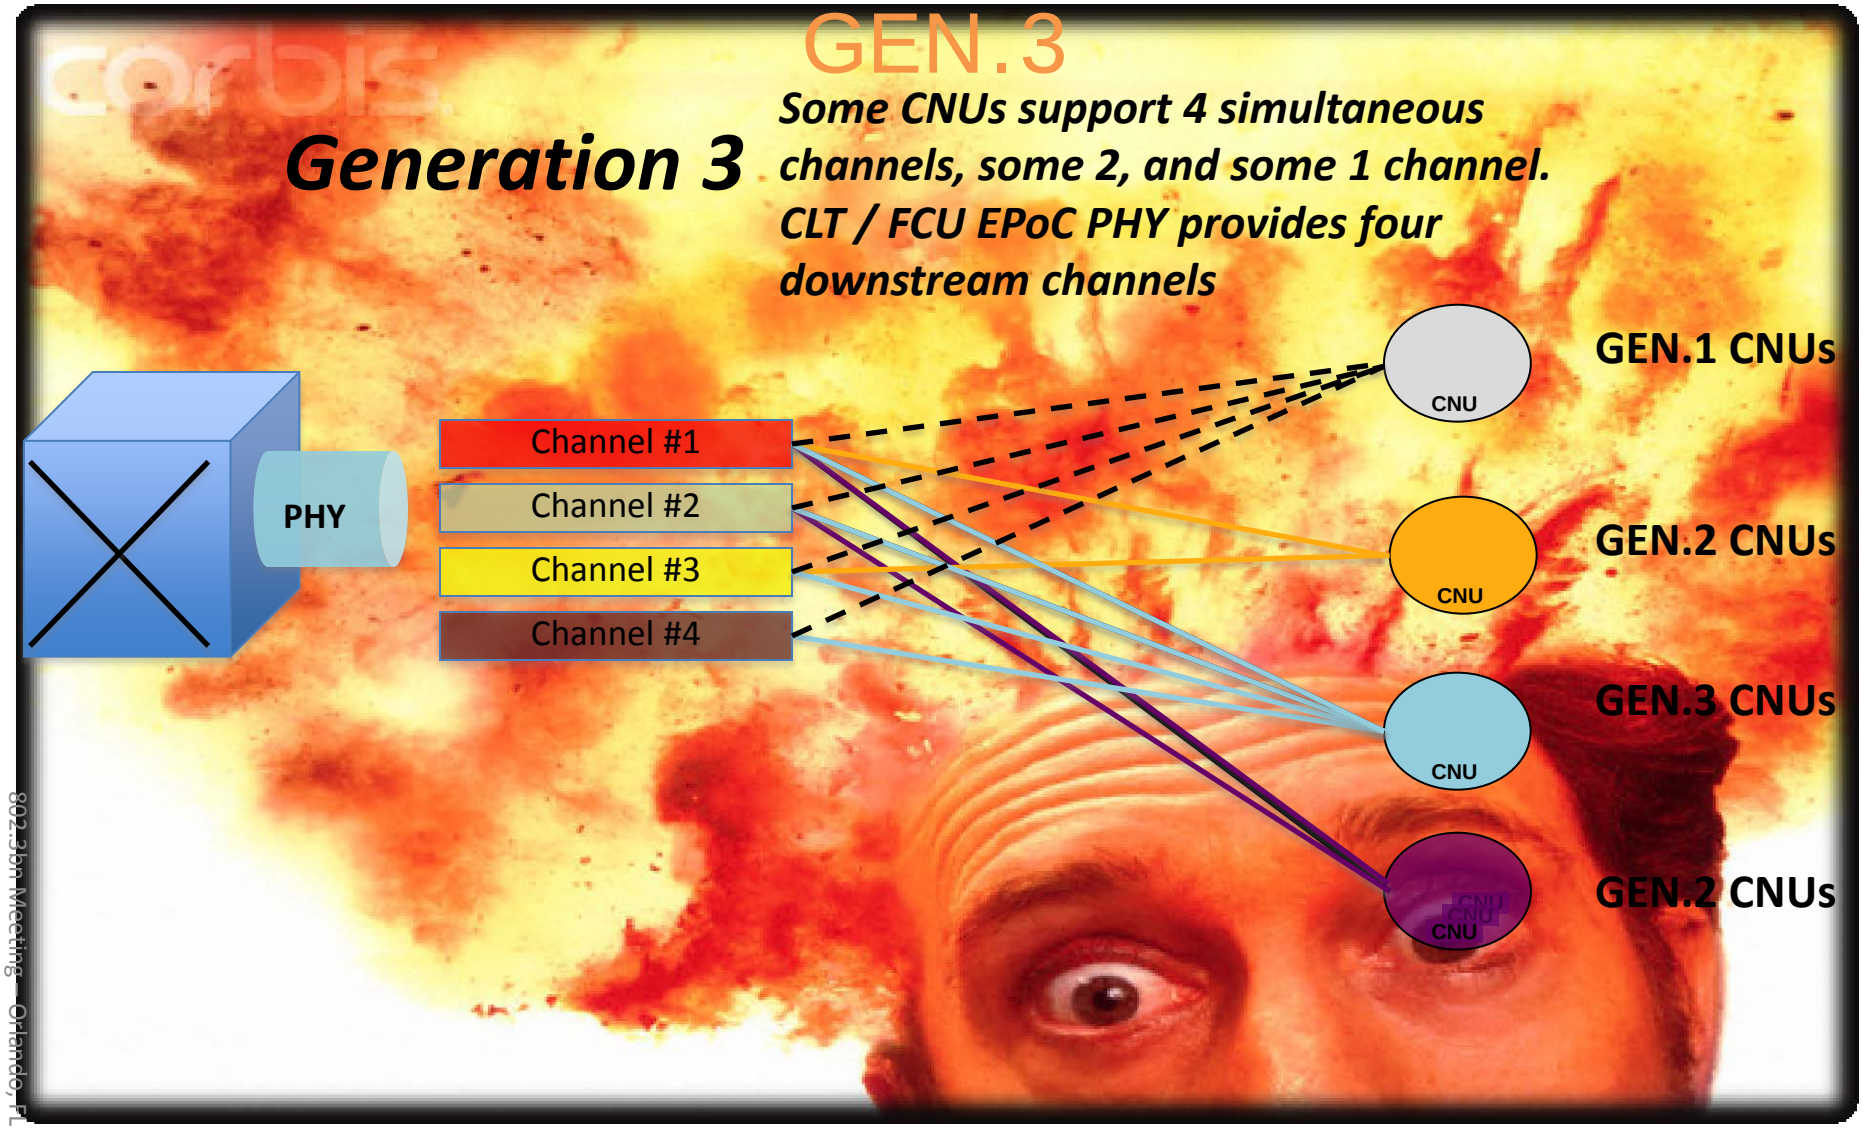

Alternatives

- So there are alternatives including separate channel resources for the different generations
  - Gen.1 in 200MHz – 400Mhz
  - Gen.2 in 450MHz – 900MHz
  - Gen.3 in 950 MHz – 1800MHz.
  - GEN.1, GEN.2 and GEN.3 share a common upstream resource 5-200MHz??
- Is this even reasonable enough to be considered an “alternative ?”

Multi Generational PHY Challenges 1

- Load balancing is complex and ends up creating many operational issues.
- Multiple generations of downstream channel support will result in a more complicated load balancing scenario. Below are a few challenges:
  - Load balancing single-downstream-channel CNUs.
  - Load balancing dual-downstream-channel CNUs.
  - Load balancing actual packets (based on a flow-based hash ? Per-packet load balancing with timestamp ?)

Multi Generational PHY Challenges 2

- Most flow load balancing algorithms in use on production devices (e.g. routers) are *NOT* load aware.
  - Packet based load balancing algorithms (such as might be required to address the load awareness) require significant receive buffering for packet re-ordering.
- Operational issues exist with per packet load balancing restricting certain types of traffic (e.g. Voice bearer) from being load balanced.

Multi Generational PHY Qualification Challenge

- Every PHY generation has to go through an exhaustive set of tests and certifications.
  - Even with a single generation, testing is onerous.
- Legacy generations have to be tested to ensure that they can be commanded to join different channels.
  - This is often not tested in the first release (single channel) because it wasn't necessary and needs to be tested long after the release.
- More complex qualification often results in stretched time to market (for the product and the service).

Service-Based Generations

- No matter what happens with the PHY generations, we will have service-based product generations.
- Service-based generations are easier to manage and make sense in the way we do business.
- Ideally, we would have a single product generation for the PHY with various service product generations including:
  - CNU with embedded router operating at 1G
  - CNU with embedded router operating at NxG
  - CNU with embedded router and WiFi AP.
  - CNU with embedded router and VoIP Agent.
  - CNU with embedded router, WiFi AP and VoIP Agent.
  - CNU with embedded router, WiFi AP, VoIP Agent, and IP Settop Box / Gateway.
  - CNU with IP Settop Box Only
  - CNU with VoIP Agent Only
- We would then be able to focus on testing and certification for what really matters – the service and not multiple versions of the PHY along with multiple versions of service capabilities.

Closing Remarks

- Operators have a desire to minimize the day-one product acquisition cost to the greatest extent possible.
  - However, this approach can have significant long term OPEX impacts.
  - This is a careful balance (short-term versus long-term cost), but it's one we should consider carefully.
- The early decision to mandate multiple generations could have ramifications as described.
  - The set of channel load balancing problems is one example.
- We should continue to investigate the balance of the relative cost versus complexity of enabling a single product generation (for 4 channels) without mandating complex load balancing mechanisms.