Thoughts on 400GbE PMDs

Gary Nicholl, Mark Nowell – Cisco Systems
IEEE 400 Gb/s Ethernet Task Force
Norfolk, May 2014

Overview

- Goal of the presentation
- Review of some data points from 40Gb/s and 100Gb/s Ethernet deployments to date
- Lessons learnt from 40Gb/s and 100Gb/s Ethernet
- Perspective on 400Gb/s PMD objectives
- Summary and Recommendations

Where is this presentation going ?

- We believe that the mainstream PMD for this 400GbE project will be a 2km, SMF duplex solution, and based on a 4x100G lambda architecture.
- We believe that such a solution will meet the initial core networking application for 400GbE, the longer term 400GbE data center interconnect needs, accelerate the availability of a cost optimized 100GbE data center interconnect (single lambda 100GbE), and provide a platform to support dense 100GbE solutions in the future via a 4x100G breakout architecture.

40GbE and 100GbE Adoption Rates

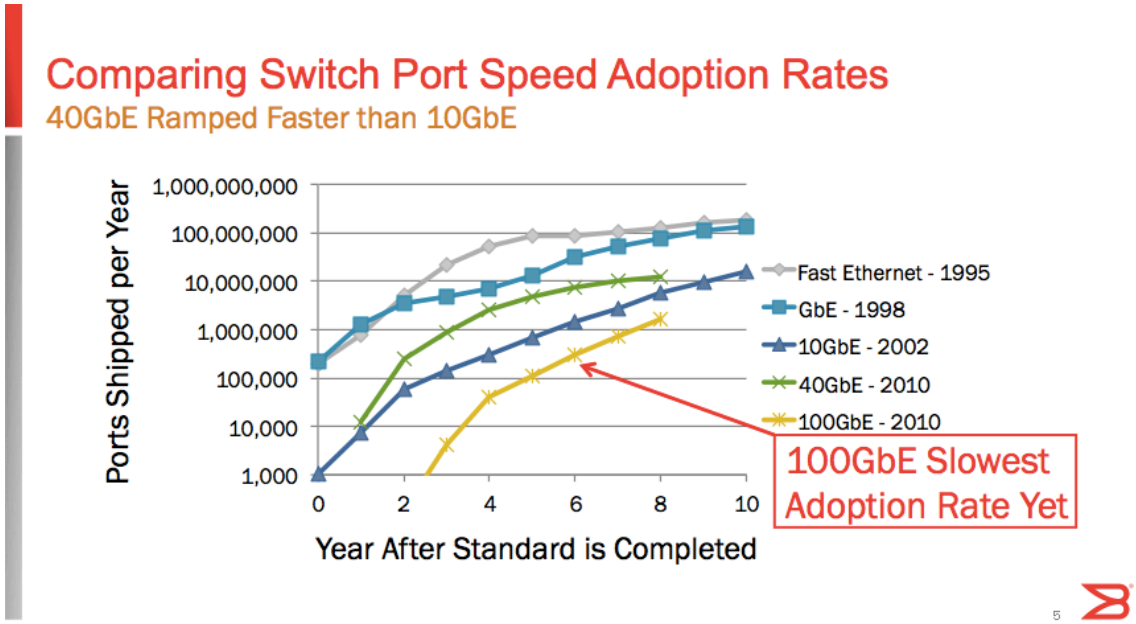

Source: Scott Kipp, Brocade

- Agree with the numbers, but would draw a different conclusion:
- 40G was not a new data rate in the industry. A lot of the early “40GbE” shipments were in fact for 4x10G breakout applications (probably explains why 40GbE ramp appears faster than 10GbE).
- 100G was a new data rate in the industry. As such a fairer comparison of the adoption rate would have been against OC-192 POS or OC768-POS interfaces.

40GbE Deployment Observations

SMF Reach: 2km versus 10km ?

- Q4/2013 shipments[1]: 88% 2km (LR4-Lite)
  1% 2km (FR)
  11% 10km (LR4)
- Clearly 2km (and low cost) is what the industry wants. Is there enough volume of 10km to warrant a unique PMD ?

Fiber Type: MMF versus SMF ?

- Q4/2013 shipments[1]: 65% MMF (SR4)
  35% SMF (LR4, FR, LR4-lite)
- Strong demand for MMF, indicating a more DC centric application space ? MMF numbers may be skewed (on the high side) by the 4x10G breakout application.
- No data yet on MMF BiDi, to give some insight into preference for MMF ribbon versus MMF duplex

(1) Source: LightCounting and all form factors

100 GbE Deployment Observations

SMF Reach: 2km versus 10km ?

- Majority of shipments to date are 10km LR4[2]
- However there is significant industry interest around 2km LR4-Lite and various 2km CWDM based solutions.
- Expect to see the same trend as at 40GbE, when a viable, cost effective 2km solution becomes available ?

(2) Note we don't have any data on 100G "LR10" 2km deployements to compare against.

100 GbE Deployment Observations

Fiber Type: MMF versus SMF ?

- Q4/2013 shipments[1]: 28%[3] MMF (SR10)
  72% SMF (LR4)
- Cisco shipments to date[2]: 13% MMF (SR10)
  87% SMF (LR4)
- Strong preference for SMF is likely a reflection of the initial application space for 100GbE being core networking (i.e. core routers and core routers to transport).
- Discrepancy between LightCounting and Cisco numbers, is likely a reflection that a lot of SR10 modules are only used for internal development and lab use, and are therefore not shipped by OEMs for revenue.
- SR10 may never see significant field deployment, as it is already being replaced by next gen SR4 solutions (802.3bm)

(1) Source: LightCounting and for all form factors.
(2) Total Cisco 100G shipments to date for revenue, on all form factors.
(3) flat growth over past few QTRs.

100 GbE Deployment Observations

Adoption rate: compared with OC768 SR

- By as early as Q1/2012[1] 100G LR4 was shipping in volumes equal to OC-768 SR (which has been shipping since mid 2000s)
- By Q4/2013 100G LR4 was shipping in volumes 8x that of OC-768 SR
- By Q4/2013 100GE was already at at cost parity with OC768 SR, but actually not dropping that fast (which may be the real issue).
- Based on these data points 100GbE has been very successful, and one could argue it may have the fastest adoption rate of any new data rate in the industry to date !

(1) Source: LightCounting

What are the takeaways from all this ?

- 40GbE and 100GbE were very different animals …
  - 40G was a new Ethernet rate, but not a new industry rate
  - 100G was a new Ethernet rate AND a new industry rate
- 400GbE is more likely to resemble 100GbE, in that it is also both a new Ethernet rate and a new industry rate.
- Historically the first (driver) application for any new higher speed data rate in the industry is in core networking (core routers and core router to transport), rather than in datacenters and/or traditional Ethernet switching applications.
  - Such core networking applications have a strong preference for a duplex SMF solution (2km, 4dB).
  - However the same 2km, duplex SMF solution may also address the needs of the datacenter in the longer term …

What are the takeaways from all this ?

- 2km, and not 10km, is the distance sweet spot for SMF PMDs
  - Consistent with historic and recent 40Gbe/100Gbe data points. The industry has clearly shown that it will not support a cost adder for pushing the PMD reach beyond 2km.
- We should be honest with ourselves about what PMDs are really required and when. It doesn't make sense to develop solutions too far ahead of the market need, and where we would likely choose different solutions when the real need materializes. Examples from 100G include the following PMDs; 100G CR10, 100G SR10, 100G ER4.

400GbE PMD Objectives

- 802.3bs has 4 PMD reach objectives;
  - 100m MMF
  - 500m SMF
  - 2km SMF
  - 10km SMF
- How are we going to meet all of these PMD objectives ?
  - 2 (minimum), 3 or 4 (worst case) PMDs ?
- Which is the one we really need to focus on and get right ?

Thoughts on 100m MMF Objective

- Initial application for this PMD is likely to be for internal development and lab use (similar to 100G SR10)
- Don't expect to see significant field deployment during the initial rollout of 400GbE (again similar to 100G SR10).
- The internal lab application could possible be met using 4x 100G SR4 modules (running over 4 separate fibers)
  - A new PMD definition is therefore not technically necessary
- Despite these comments we are not strongly opposed to standardizing a solution for this objective:
  - Solution is assumed to be 400G SR16, and heavily leverage 100G SR4
  - As such it should be a fairly easy and painless to standardize

Thoughts on 2km SMF Objective

- Solution has to be duplex
- 4x100G lambda architecture makes the most sense:
  - leverage a CWDM grid (listen to Matt ☺)
  - accelerates a cost-optimized 100G
  - building block available for future 4x100G breakout
- Advan Mod has advanced !
  - Hot topic at OFC 2014 (papers, presentations, panels, live demo's)
  - PAM4 and DMT remain the most promising candidates
  - Both feasible
    - experimentally demonstrated at OFC
    - decision probably driven by schedule/power trade off
- Provides a longer market life, supporting multiple generations of electrical I/O
  - 16x25G
  - 8x50G
  - 4x100G

  → 4x100G optical

Thoughts on 500m SMF Objective

- Met by the 2km duplex solution (see previous slide)

Thoughts on 10km SMF Objective

- Historic and recent 40G/100G data clearly show that 2km is the sweet spot for SMF reach.
- Is there really enough volume to justify a standalone 10km PMD, in addition to an optimized 2km PMD, and the future need for a 40km PMD ?
- Should we consider merging this objective into a future 40km PMD objective (so that all reaches in the 2-40km range would be met with a single PMD) ?

Summary

- Early market applications for 400 Gb/s Ethernet will be similar to those seen in early market 100 Gb/s Ethernet.
- PMD types seen in these early market 400 Gb/s Ethernet applications are expected to be similar to those seen in early market 100 Gb/s applications.
- Let's learn from history, be honest with ourselves and focus our efforts on the PMDs that really make sense.
- As such we strongly believe that the mainstream PMD for this 400Gb/s Ethernet project will be a 2km, SMF Duplex solution, based around a 4x100G lambda architecture.