Setting Optimal MMF Reach Objectives, Indications from Total Cost Analysis

Paul Kolesar
CommScope
February 28, 2012
Next Gen 100G Optics Ad-Hoc Telecon

Purpose and Approach

- Use total cost analysis (cable + PMDs) to get objective picture of trade-offs that affect reach objectives
- Apply new Kalculator that adds 24-fiber cabling cost to allow 100G-SR10 analysis
  - Posting eligibility now in IEEE legal review
- Look at recent history for guidance

Assumptions and Underpinnings

- At 100G, complete coverage of data center channels necessitates a SM solution
- For data centers the reach of a SM objective is not a critical parameter as long as it permits near 100% channel coverage
  - At least 350m is sufficient to do this job
  - Aiming for 3x this distance will likely not change cost
  - Connection loss budget is more important than reach greater than 350m
- SM solutions will remain more costly than MM, but the cost gap will narrow over time
  - Choices we make affect the rate at which that gap closes
- MM solutions are essential for cost effective data centers and broad market potential
  - Getting MM optimized is a major focus of our studies
  - There are widely varying view points on what is optimal
- Must take a total solution view to find true optimization
  - MM and SM solutions work in concert
- The cost of SM channels is an impediment to data center deployments today
  - The future cost of SM solutions has a direct effect on the percentage of channels that the MM solution must serve
- For Ethernet it's all about cost
  - So cost studies will take a front row seat here

Total Solution Cost for Sw-Sw Channels (2:1 mix)

Kalculator inputs				
PMD description	PMD cost factor	reach (m) OM3	OM4	OS2
SR4 100m	1.3	70	100	n.a.
SR4 150m	1.5	100	150	n.a.
SR10	1.0	100	150	n.a.
SR10 (rgb)	1.3	100	150	n.a.
LR4	2 - 10	n.a.	n.a.	10,000

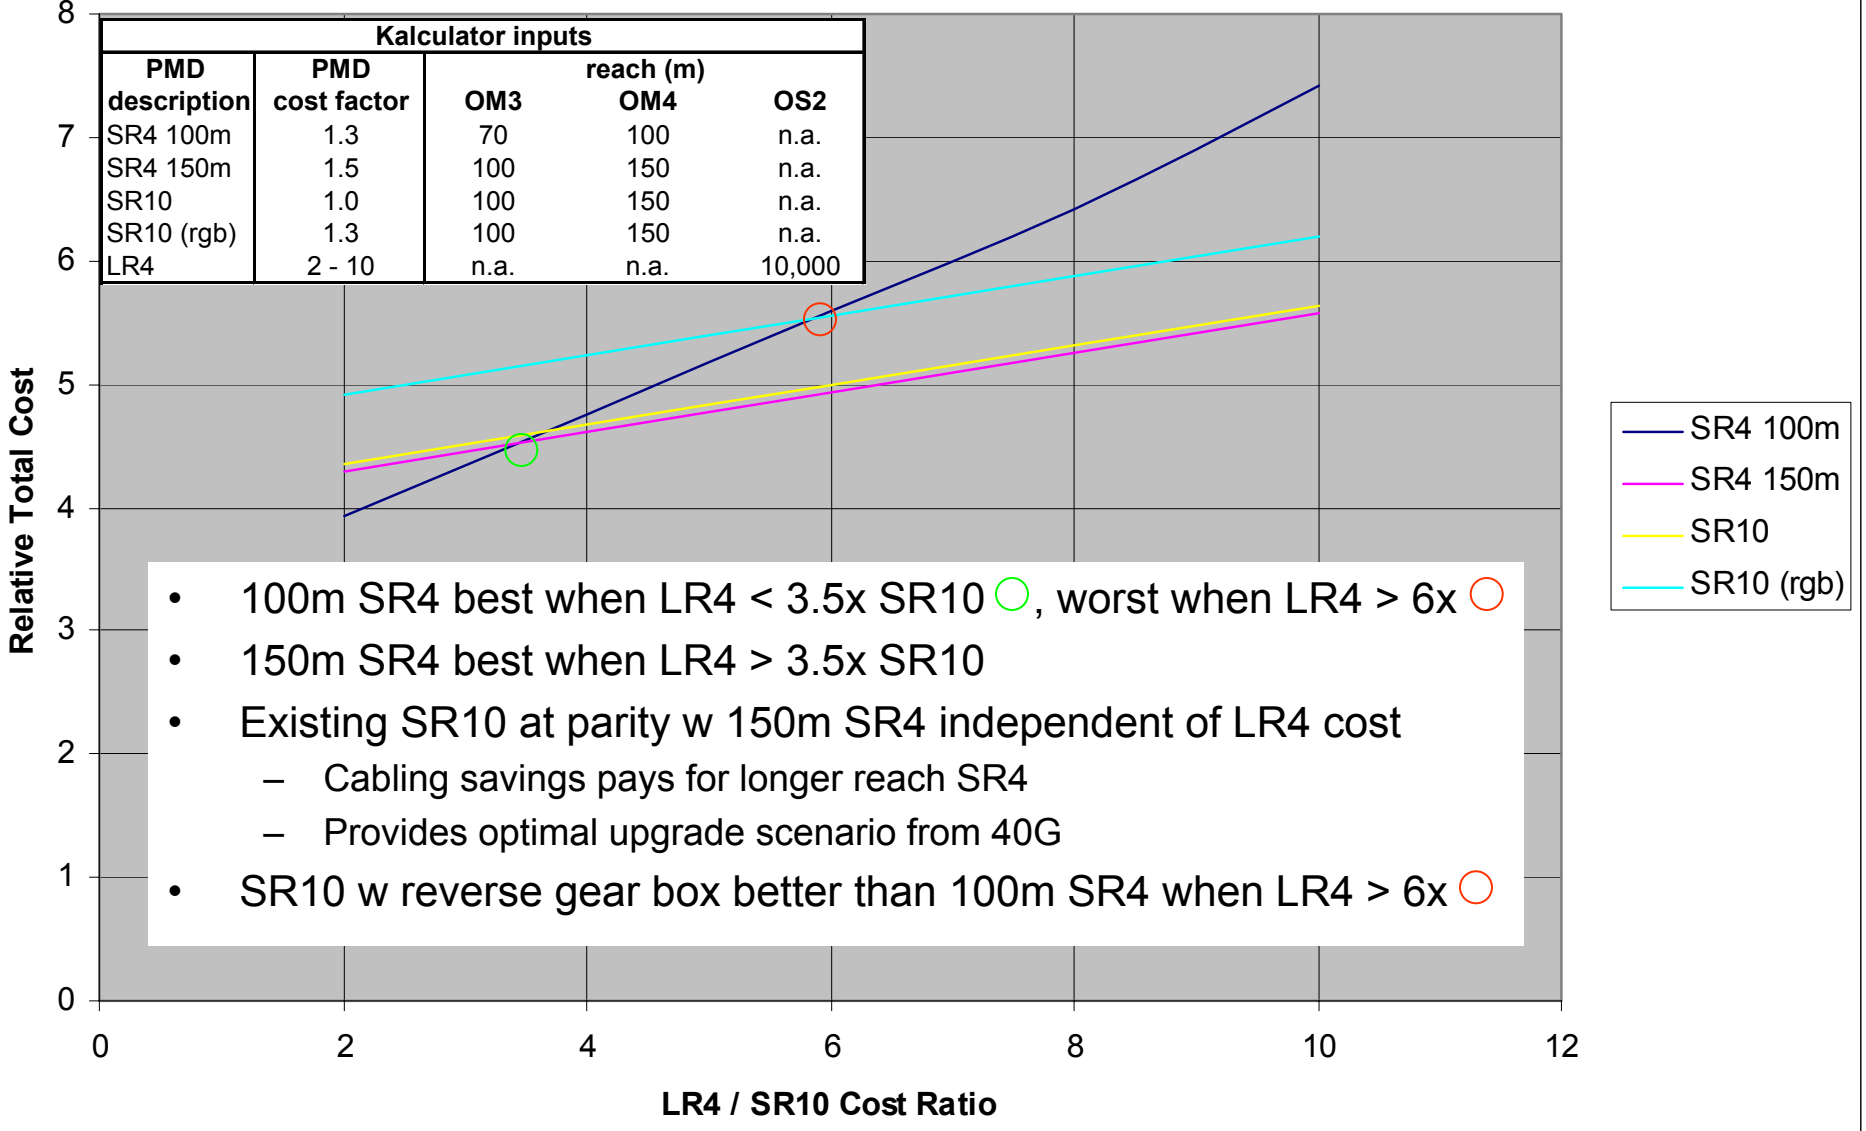

- 100m SR4 best when LR4 < 3.5x SR10 ○, worst when LR4 > 6x ○
- 150m SR4 best when LR4 > 3.5x SR10
- Existing SR10 at parity w 150m SR4 independent of LR4 cost
  - Cabling savings pays for longer reach SR4
  - Provides optimal upgrade scenario from 40G
- SR10 w reverse gear box better than 100m SR4 when LR4 > 6x ○

Total Solution Cost for Sw-Sw Channels (single-link)

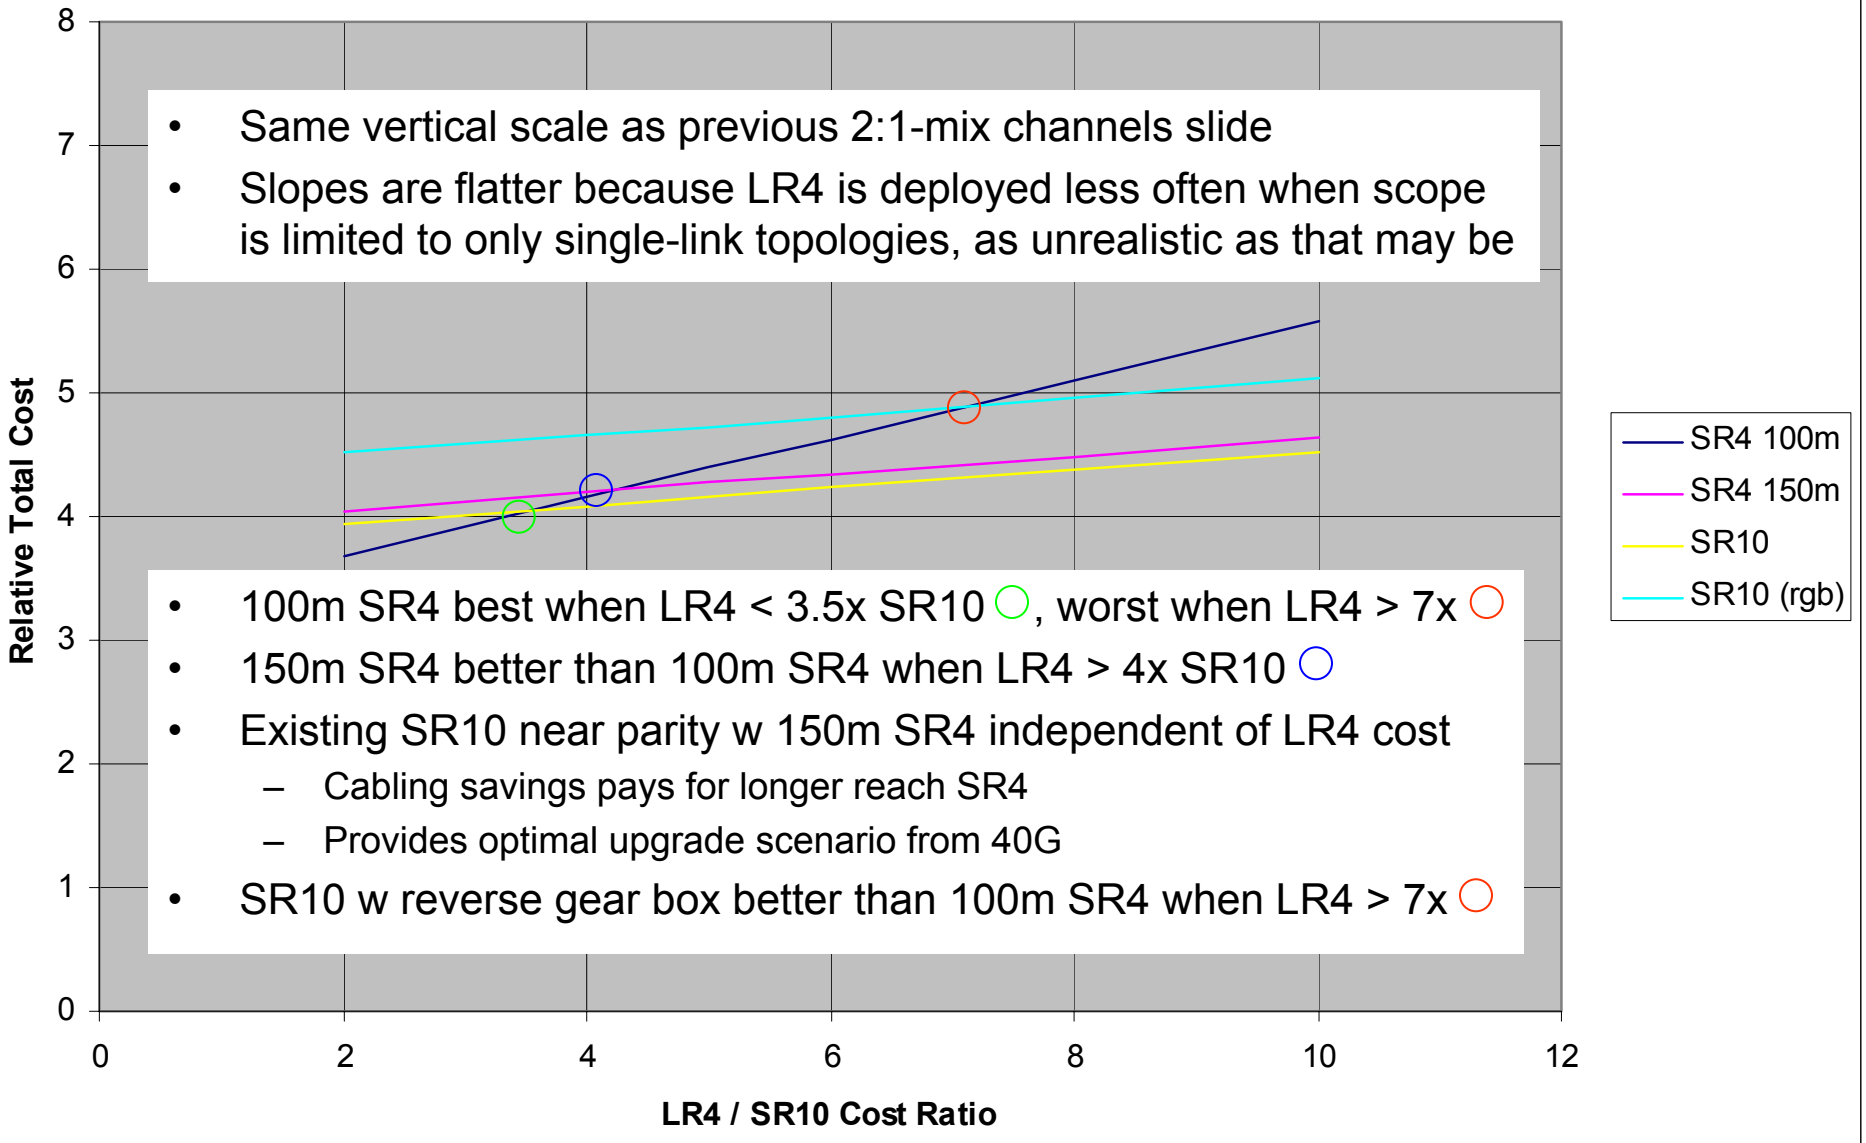

- Same vertical scale as previous 2:1-mix channels slide
- Slopes are flatter because LR4 is deployed less often when scope is limited to only single-link topologies, as unrealistic as that may be

- 100m SR4 best when LR4 < 3.5x SR10 ○, worst when LR4 > 7x ○
- 150m SR4 better than 100m SR4 when LR4 > 4x SR10 ○
- Existing SR10 near parity w 150m SR4 independent of LR4 cost
  - Cabling savings pays for longer reach SR4
  - Provides optimal upgrade scenario from 40G
- SR10 w reverse gear box better than 100m SR4 when LR4 > 7x ○

Total Solution Cost for Sw-Sw Channels (2:1 mix)

Consider the co-evolution of 100GE technology and market need

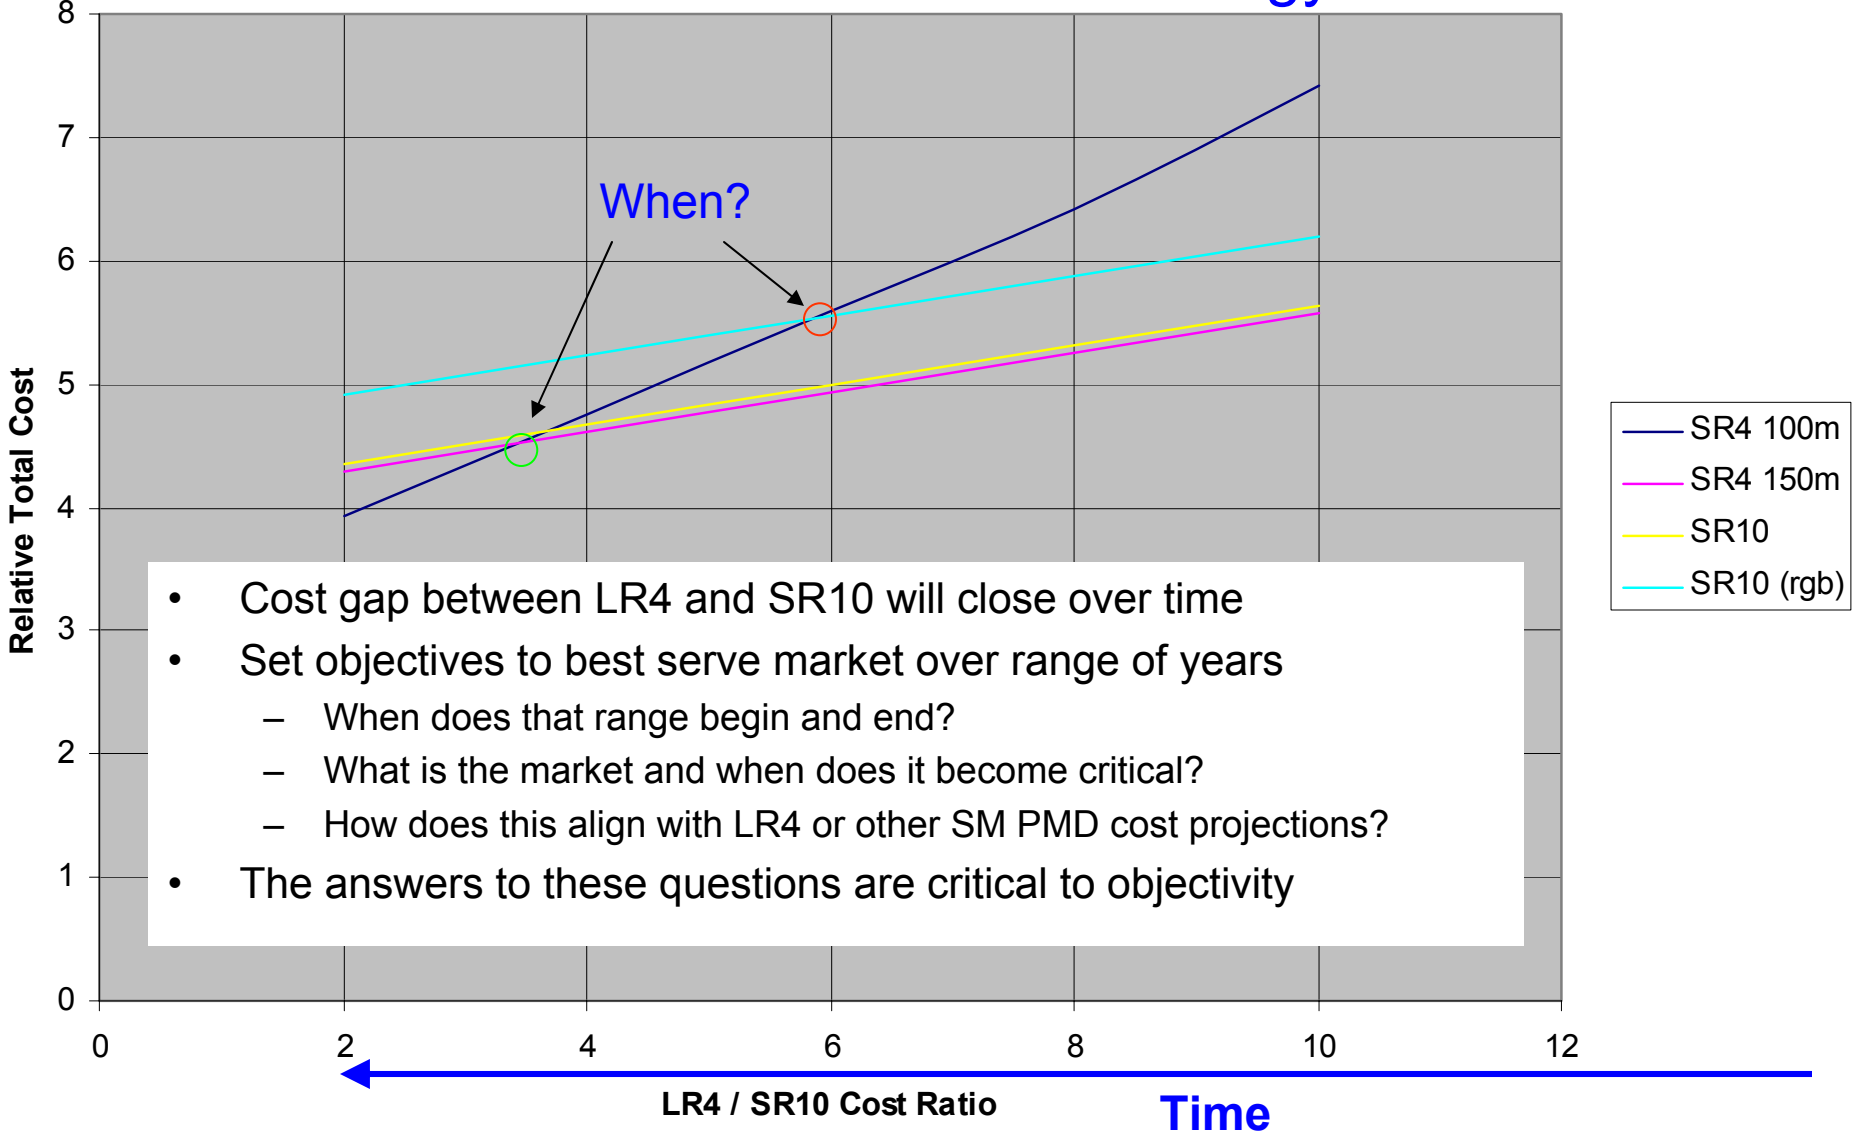

- Cost gap between LR4 and SR10 will close over time
- Set objectives to best serve market over range of years
  - When does that range begin and end?
  - What is the market and when does it become critical?
  - How does this align with LR4 or other SM PMD cost projections?
- The answers to these questions are critical to objectivity

Some Answers / Perspectives

- What are the 100GE data center applications and timeframes?
  - Switch-to-switch applications that aggregate lower-rate server traffic dominate thru at least 2021
  - Server-to-switch applications dominate after that (see backup slide)
- Should Ethernet focus on HPC needs?
  - HPC drove demand for MM optics and thereby lowered cost of 40G-SR4 and 100G-SR10
  - The same benefit can be expected for Next Gen 100G if suitable for HPC
- When does the relevant market window begin and end?
  - Begins when this work is approved: 2014
  - Continues for at least 3 years, likely more, depending on shifts in technology or market
- When does the market broaden?
  - The market broadens continuously from today forward driven by growth of 10G and 40G servers, then potentially explodes if used on 100G servers in the next decade (see backup slide)
- How does this align with LR4 or other SM PMD cost projections?
  - Current costs vs. required costs indicate that it will be many years before LR4 costs decline far enough to allow a 100m MMF reach objective to be cost-optimal
- Do we need another SM PMD or do we wait for LR4?
  - This is a critical open issue and the reason why it would be helpful to establish any SM objective before setting MM reach objectives

Need to do two jobs: 1. reduce total cost for switch-to-switch channels
2. optimize for HPC to drive volume cost reduction

Two-Tier MM Solutions

- What does past experience with 10GE tell us?
  - 10GBASE-S specifies 300m on OM3, will soon specify 400m on OM4
  - Non-standard 100m solutions are said to be successful
  - Two-tier performance choice is working here, but interoperability is not assured
- What about existing 40GE and 100GE?
  - 40GBASE-SR4 and 100GBASE-SR10 specify 100m on OM3, 150m on OM4
  - We failed to specify extended reach, although we had several ways to do so
    - If we had, then there'd be a lower-cost alternative to LR4 filling almost all data center needs
    - If we had, then there would likely be a broader market today
- What are we doing now?
  - Moving up to 25G electrical rates
  - Considering shortening the reach and further constraining the applicability of MM
    - This path will not enable the market without a much lower-cost SM solution
  - Expecting either: 1) LR4 to quickly drop in cost, or 2) a lower-cost alternative
    - The path we take impacts the optimal choice for MM. Will we get a step-function cost decline or not?
- Two-tier MM solution provides the way forward
  - One solution provides lowest-cost and lowest-power for short (< 50m) channels and HPC
    - AOCs can fulfill this need via the electrical interface spec; no optical specs needed; no separate PMD
    - However, AOCs are not interoperable given port lock-outs. Another PMD would resolve that problem.
  - The other solution provides a lower-cost for the remaining bulk of switch-to-switch links
    - The optimal reach depends on the cost of the SM solution

Conclusion

- SM solution cost over time is critical for setting optimal MM reach objectives
  - Must have a clear projection of this trend for SM solutions under study
  - Knowing this will allow determination of the total cost picture over time for various MM choices
  - Once costs of alternatives are known, then consider other influences such as difficulty of development, time to market, power dissipation, density
    - these may drive different choices other than lowest total cost
- Do not make the mistake of “ready, fire, aim”
  - Know where your target is before pulling the trigger
- At this point it seems like two MM reach objectives make the most sense, each one defining the minimum capability of different PHYs
  - 50m on OM3
  - 150m on OM4

Backup Material

Server Market Projection

from http://www.ieee802.org/3/hssg/public/apr07/hays_01_0407.pdf

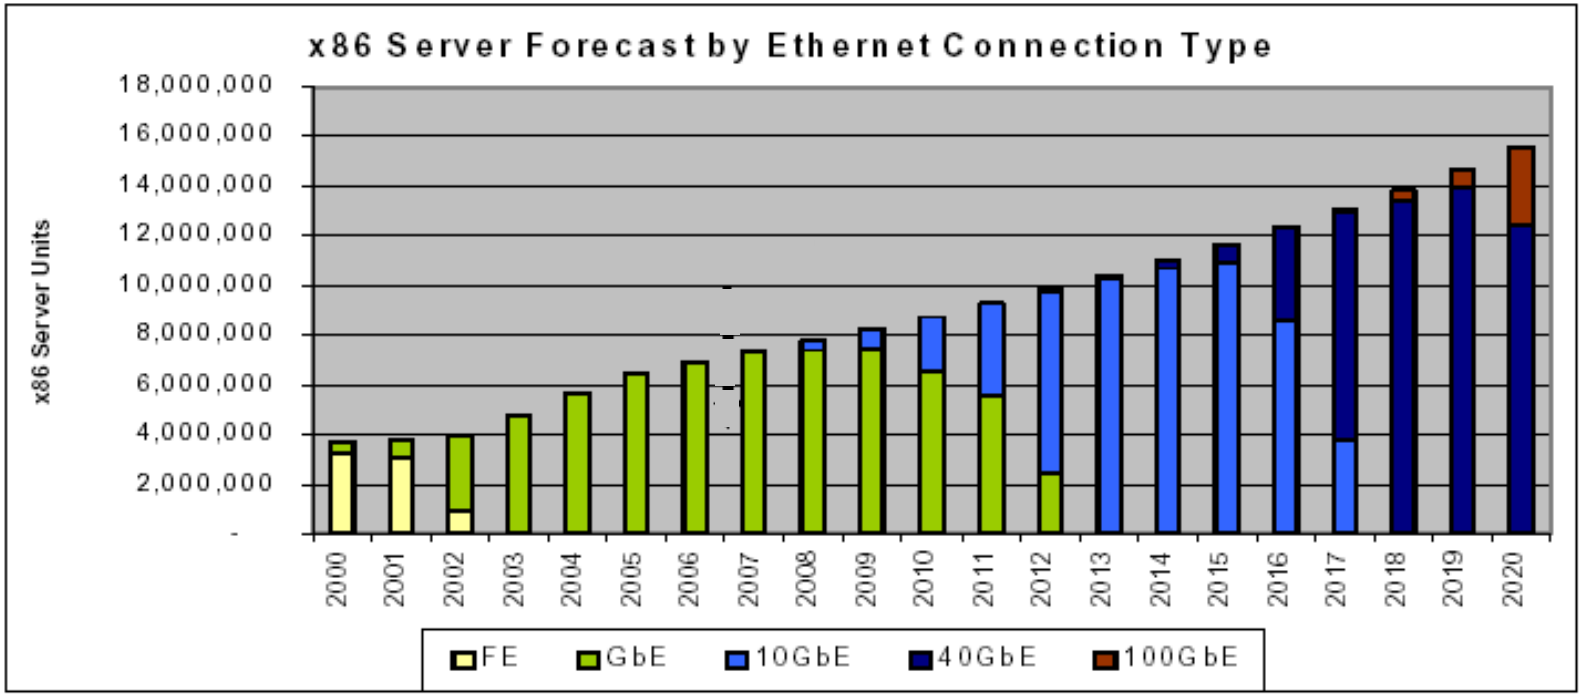

updated:
July 2007

Example of Historical Correlation Between Speeds and Ports Growth

10G server deployment will drive 40GE and 100GE just as 1G servers drove 10GE

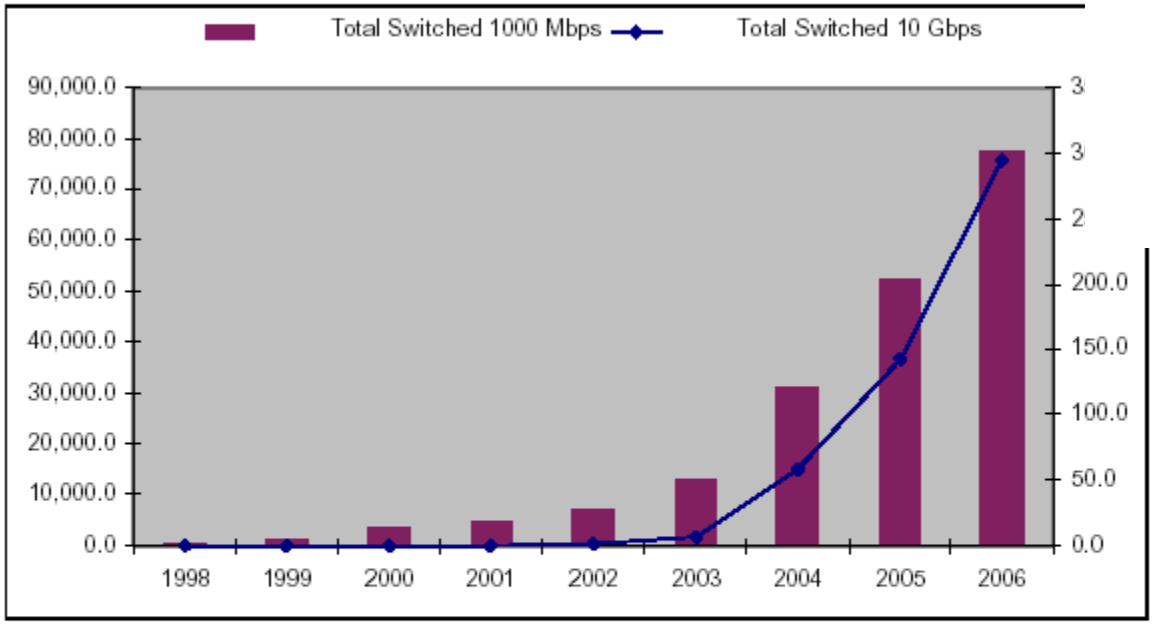

Source: Cisco (barbieri_01_0107.pdf)

- **Between 2003 and 2006, GbE growth and 10GE growth were correlated. Symbiotic relationship.**
- **2007: 10GE growth being constrained by lack of higher speed interface (Sources: Sprint, Yahoo, EDS, Amazon, AMS-IX, Cox, NTT, Equinox)**