

Setting Optimal Reach Objectives, Indications from Total Cost Analysis Over Time

Paul Kolesar
CommScope
March, 2012

IEEE 802.3 Next Gen 100G Optics Study Group


Purpose and Approach

- Use total cost analysis (cable + PMDs) to get objective picture of trade-offs that affect reach objectives
 - For the near term and long term
- Apply new Calculator that adds 24-fiber cabling cost to allow 100G-SR10 analysis
 - Posting eligibility passed 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 and the degree of closure
- 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
 - LR4 is examined first, then "PSM4"
 - Total cost for both is projected well into the future

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



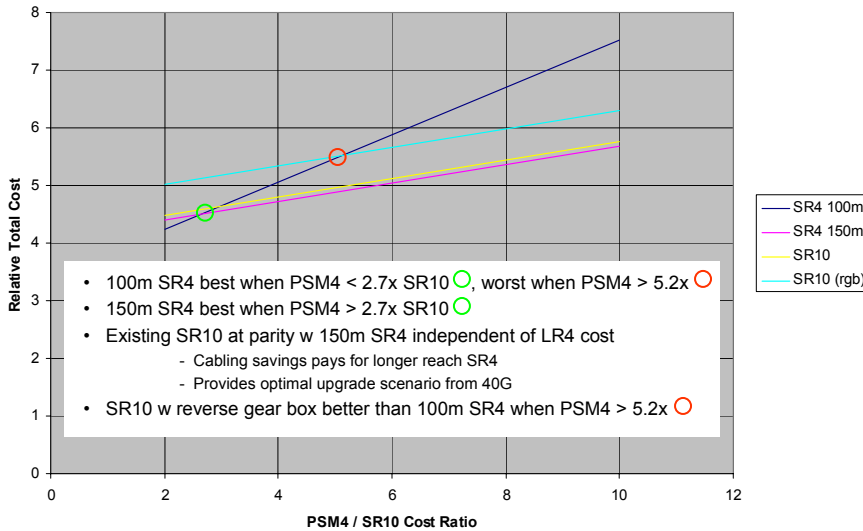
Total Solution Cost for Sw-Sw Channels (single-link)



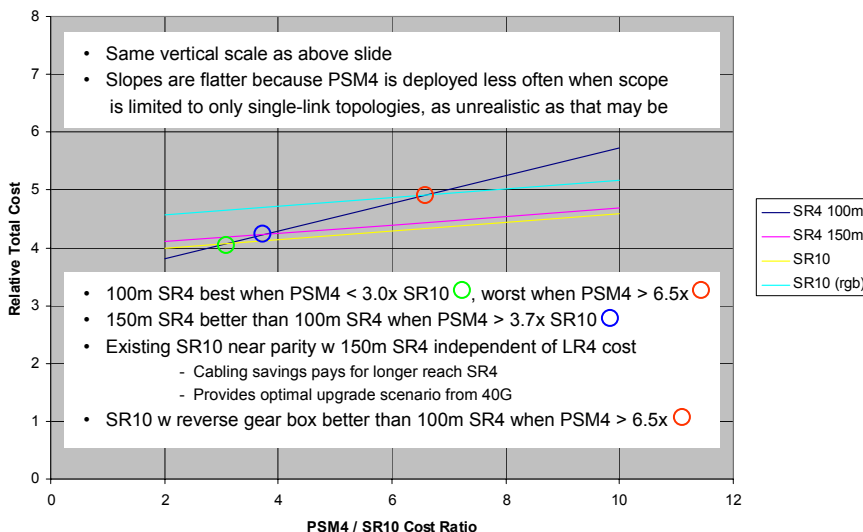
Same analysis for Parallel SM “PSM4”

- Same trends repeat
 - But the PMD cost ratio cross-over points shift lower due to cost of parallel SM cabling
 - In other words the PSM4 PMD cost must be somewhat lower to offset the added cabling cost of parallel fiber

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



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




Calculator inputs - PSM4 study				
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.
PSM4	2 - 10	n.a.	n.a.	1,000

LR4 case:

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

Considering co-evolution of 100GE technology and market need



Best-case for 100m SR4 only results in ~parity w 150m SR4 and requires fully mature optics costs

When?
Likely many years from now.

Lowest-risk lowest-cost frontier

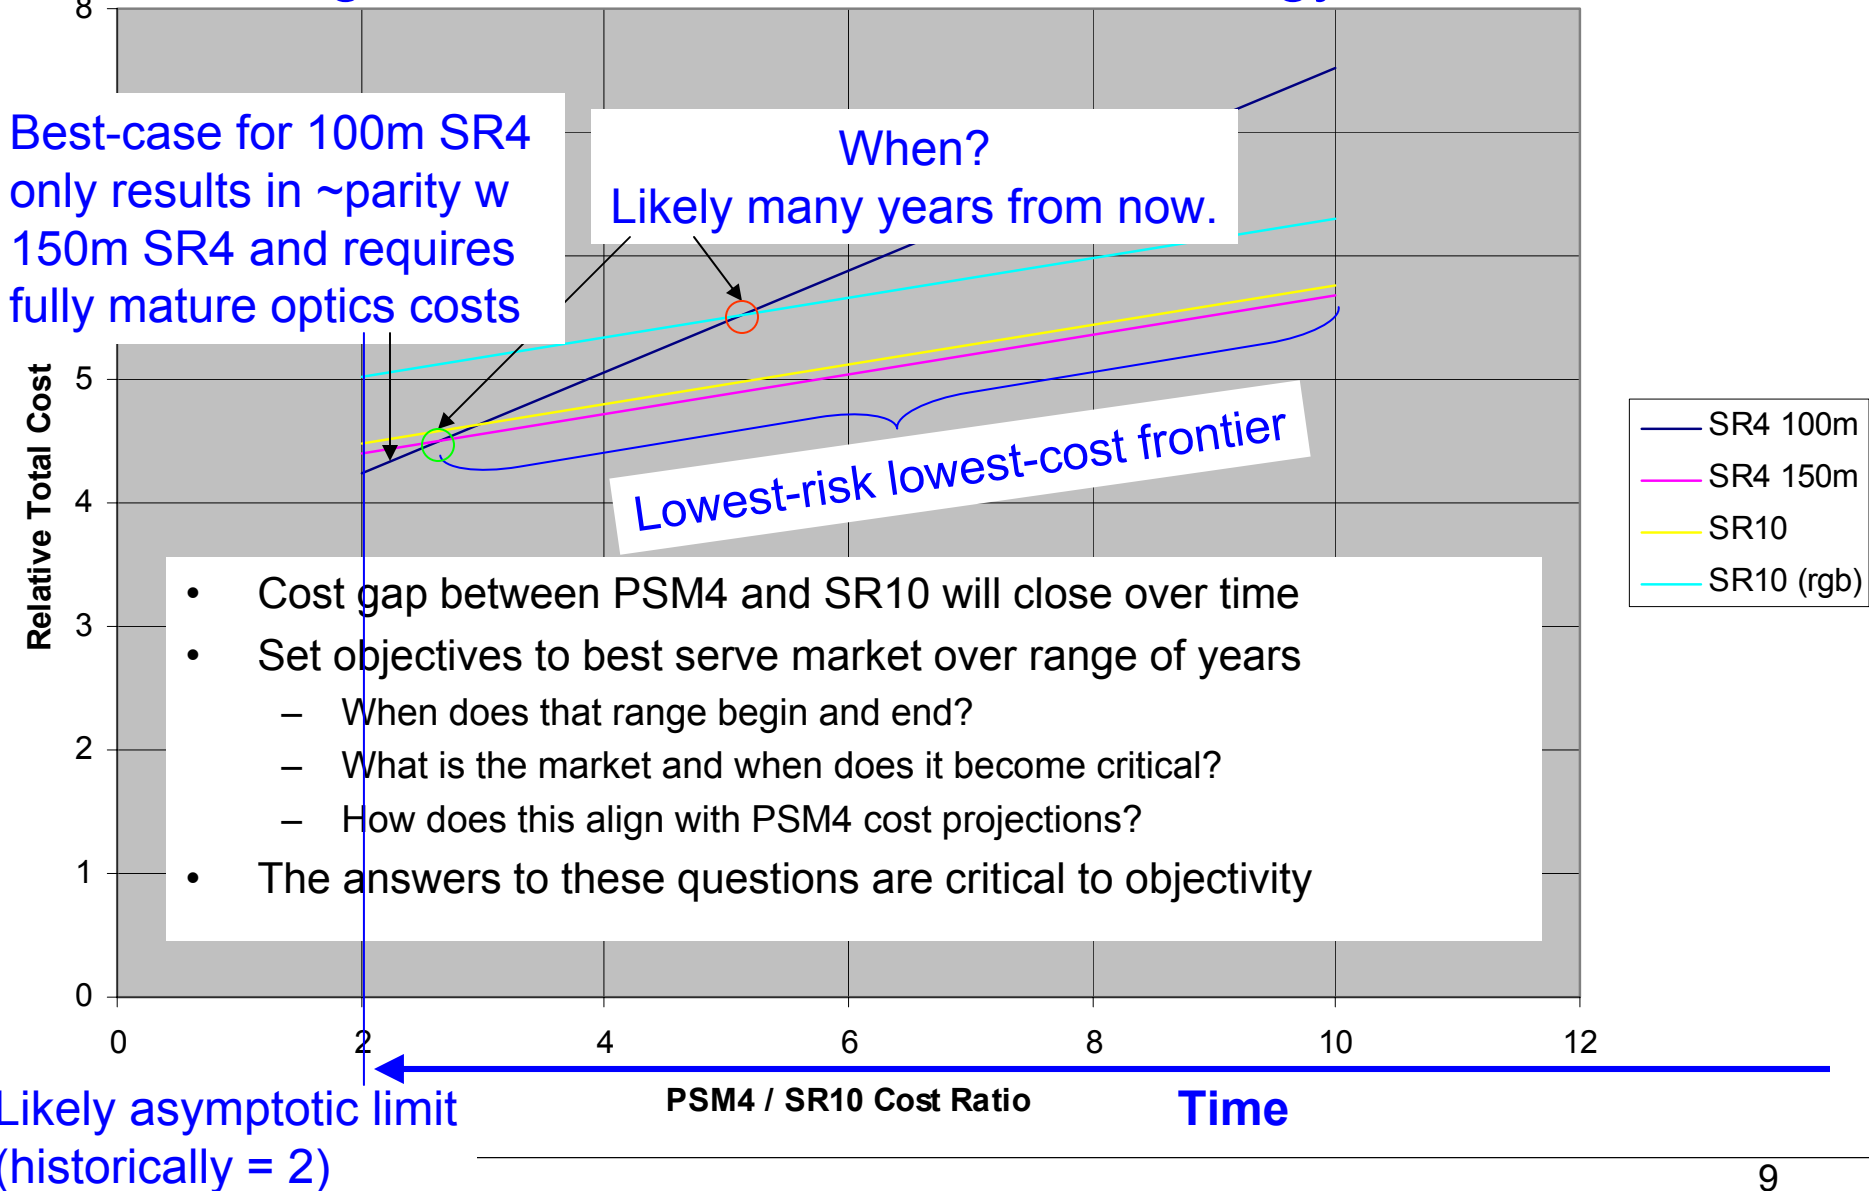
- Cost gap between LR4 and SR10 will close over time
- Set objectives to best serve market over range of years
 - When does that range begin and end?
 - What is the market and when does it become critical?
 - How does this align with LR4 cost projections?
- The answers to these questions are critical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

Likely asymptotic limit (historically = 2, but WDM raises that significantly)


PSM4 case:

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

Considering co-evolution of 100GE technology and market need

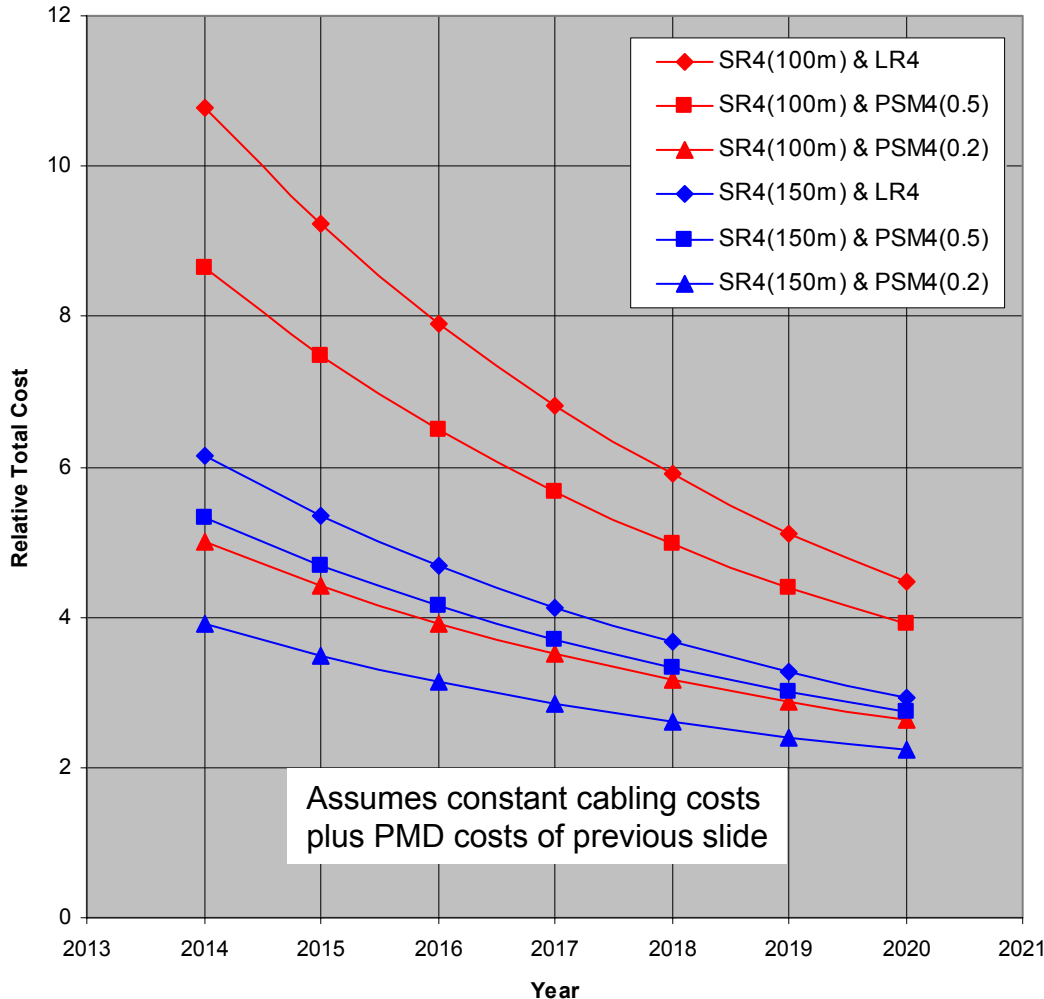


PMD Cost Projections Over Time



Total Cost Projection


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



- Total solution cost is a strong function of the SM PMD cost
- All modeled scenarios show 150m SR4 delivers lower total cost than 100m SR4
 - No SM solution comes close to the cost required to make 100m SR4 a lower total-cost choice than 150m SR4
- 150m SR4 & PSM4 together offer lowest total solution cost
 - LR4 closes the gap over time due to cabling cost difference
 - But even when compared against “conservative” PSM4 the cross-over is more than a decade away

Total Cost Savings Projection

Total Solution Cost Savings vs 100m SR4 Over Time
for Sw-Sw Channels (2:1 mix)



- All modeled scenarios show 150m SR4 delivers substantial total cost savings compared to 100m SR4
 - No SM solution comes close to the cost required to make 100m SR4 a lower total-cost choice than 150m SR4
- Cost savings are proportional to SM PMD cost, typically:
 - 40% w LR4
 - 35% w “conservative” PSM4
 - 18% w “aggressive” PSM4

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, if ever, before SM PMD 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?
 - Cost projection study just presented indicates no likely SM PMD will change the total-cost advantage of 150m SR4 over 100m SR4, but PSM4 appears cost optimal for SM fiber

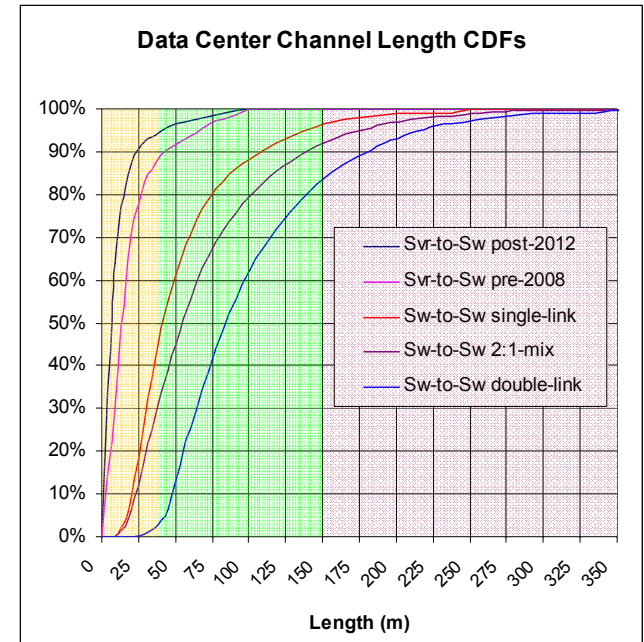
Need to do two jobs: 1. reduce total cost for switch-to-switch channels
2. optimize for HPC to drive volume cost reduction

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
 - But none of the studied choices change the total cost advantage of 150m SR4 over 100mm SR4
- Two-tier MM solution provides the way forward
 - One solution provides lowest-cost and lowest-power for short (< 40m) 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 channels

Conclusion

- SM solution cost over time can influence optimal MM reach objectives
 - However, no SM solution studied becomes low enough in cost to make 100m SR4 a better choice than 150m SR4
 - Although not shown, a two-tier MM PMD set that also includes a no-frills alternative for < 40m will further lower the total cost picture
- Two MM reach objectives make the most sense, each one defining the minimum capability of a different PHY optimized to address one of the two jobs at hand:
 - 30m to 40m on OM3 (interoperable lowest-cost solution optimized for short channels and HPC)
 - 150m on OM4 (optimized for the bulk of switch-to-switch channels)
- A SM PMD objective targeting a cost-optimized solution for data centers to lower overall cost
 - At least 350m on single-mode fiber

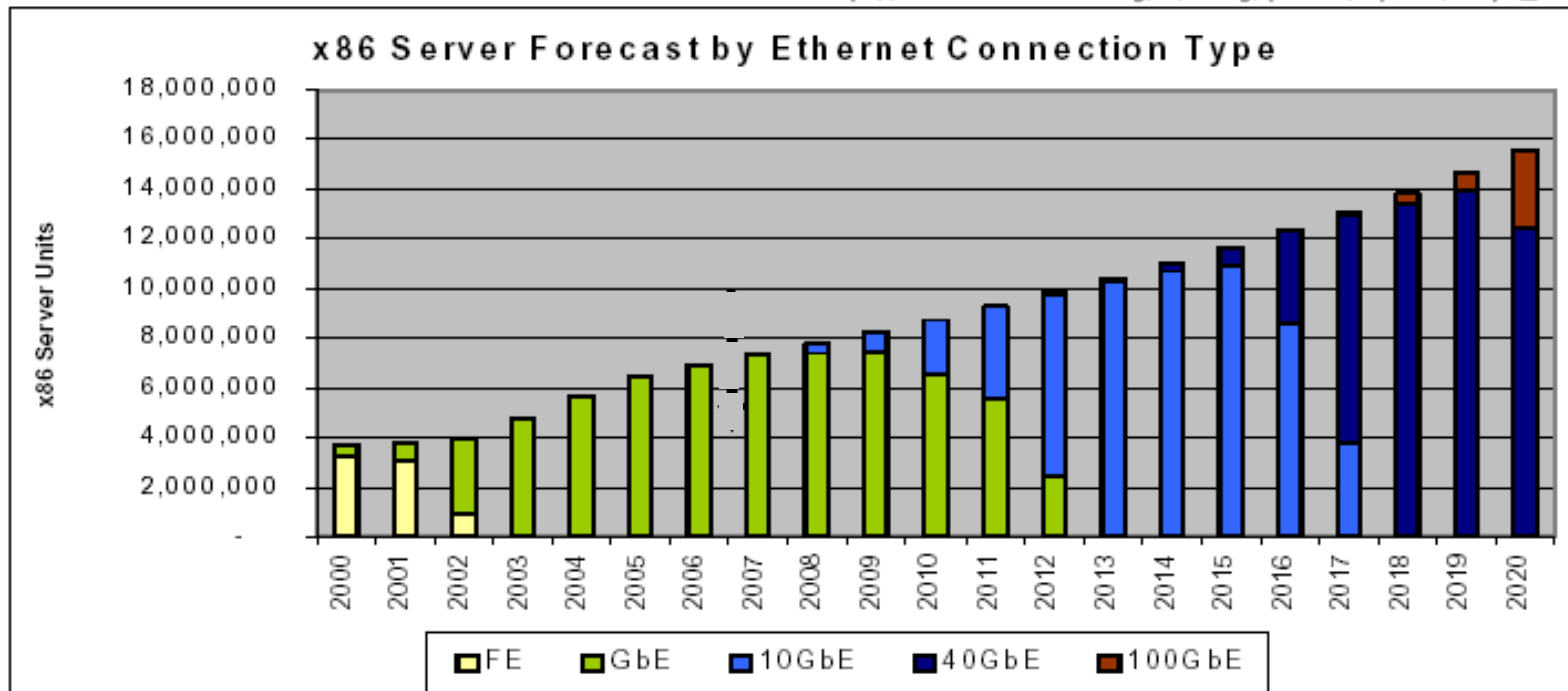


Backup Material

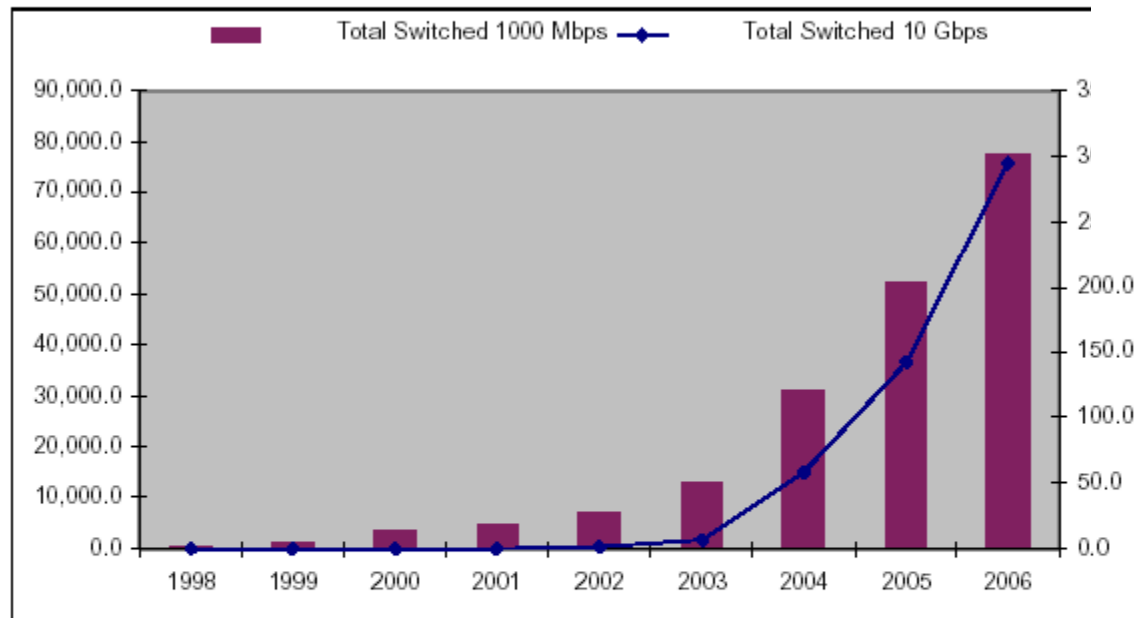
Server Market Projection

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)**

1GE, 2.5G, 10GE PMD Historical Relative Market Cost



- All curves on this graph and all other graphs normalized to 2008 LR XFP cost.
- The cost of Gigabit Ethernet has declined significantly during 10GE adoption (about 50% in 4 years.)
- A similar decline can be assumed for the cost of 10GE during 40GE adoption, which will benefit 40GE-CWDM cost decrease.
- Ratio of 10G PMD cost to 2.5G PMD (OC-48 IR-1 SFP) cost gives a historical basis for estimating ratio of 40G PMD cost to 10G PMD (LR XFP) cost.

Finisar