100 Gb per Lambda: The Building Block to Ethernet's Future

LightCounting
Market Research

Ethernet Alliance /OIDA Workshop

Dale Murray • June 13, 2014

Opinions expressed during this presentation are the views of the presenter, and should not be considered the views or positions of OIDA or the Ethernet Alliance.

Ethernet Optical Transceiver Unit Shipments by Data Rate

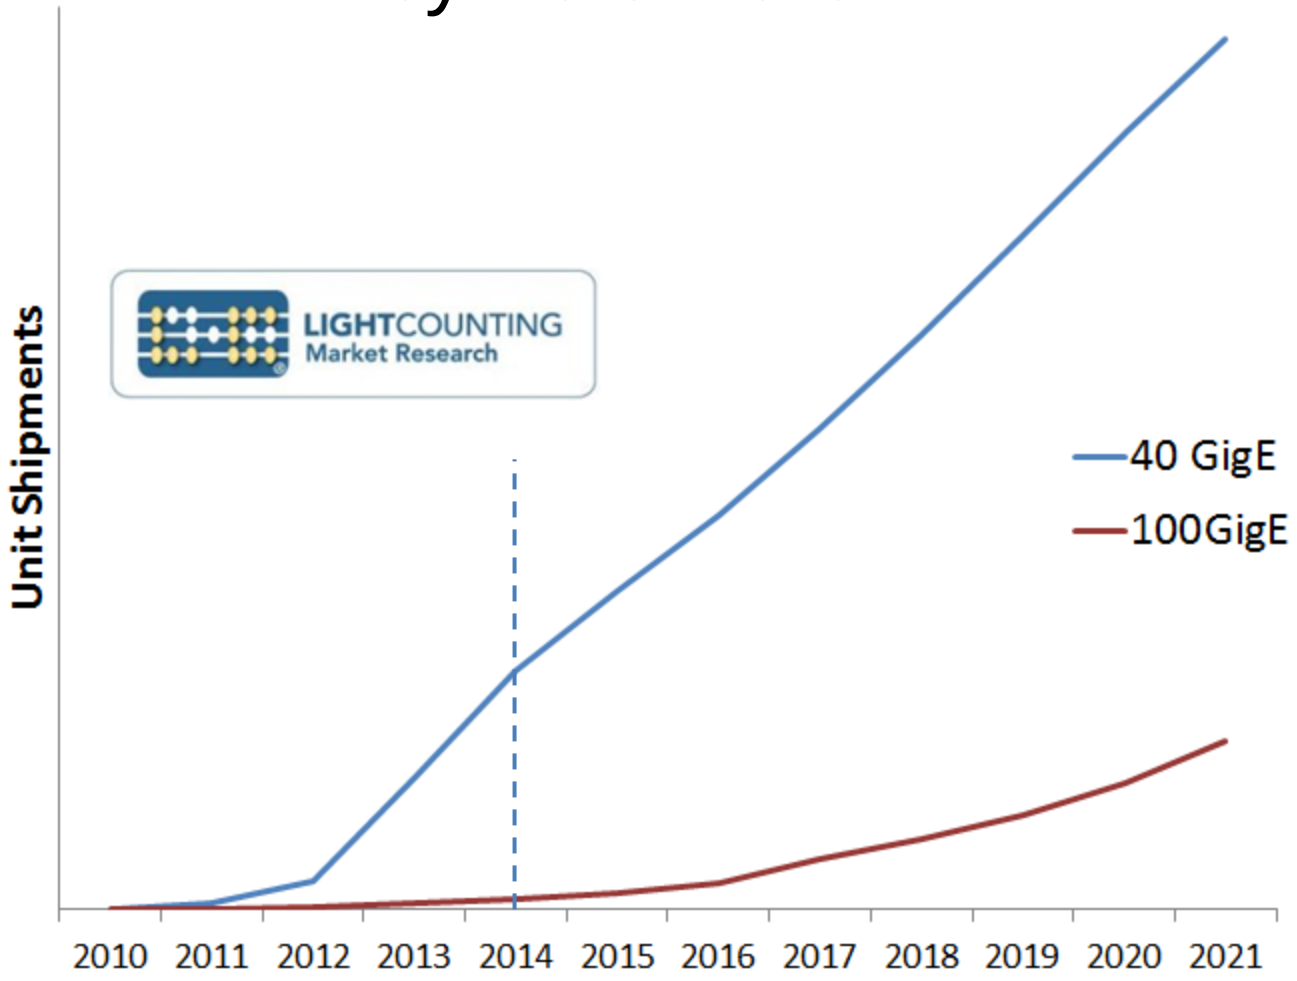

Numerical data for the vertical scale is available in LightCounting's Market Forecast Report

Flatter Leaf & Spine architectures mean more high-speed links

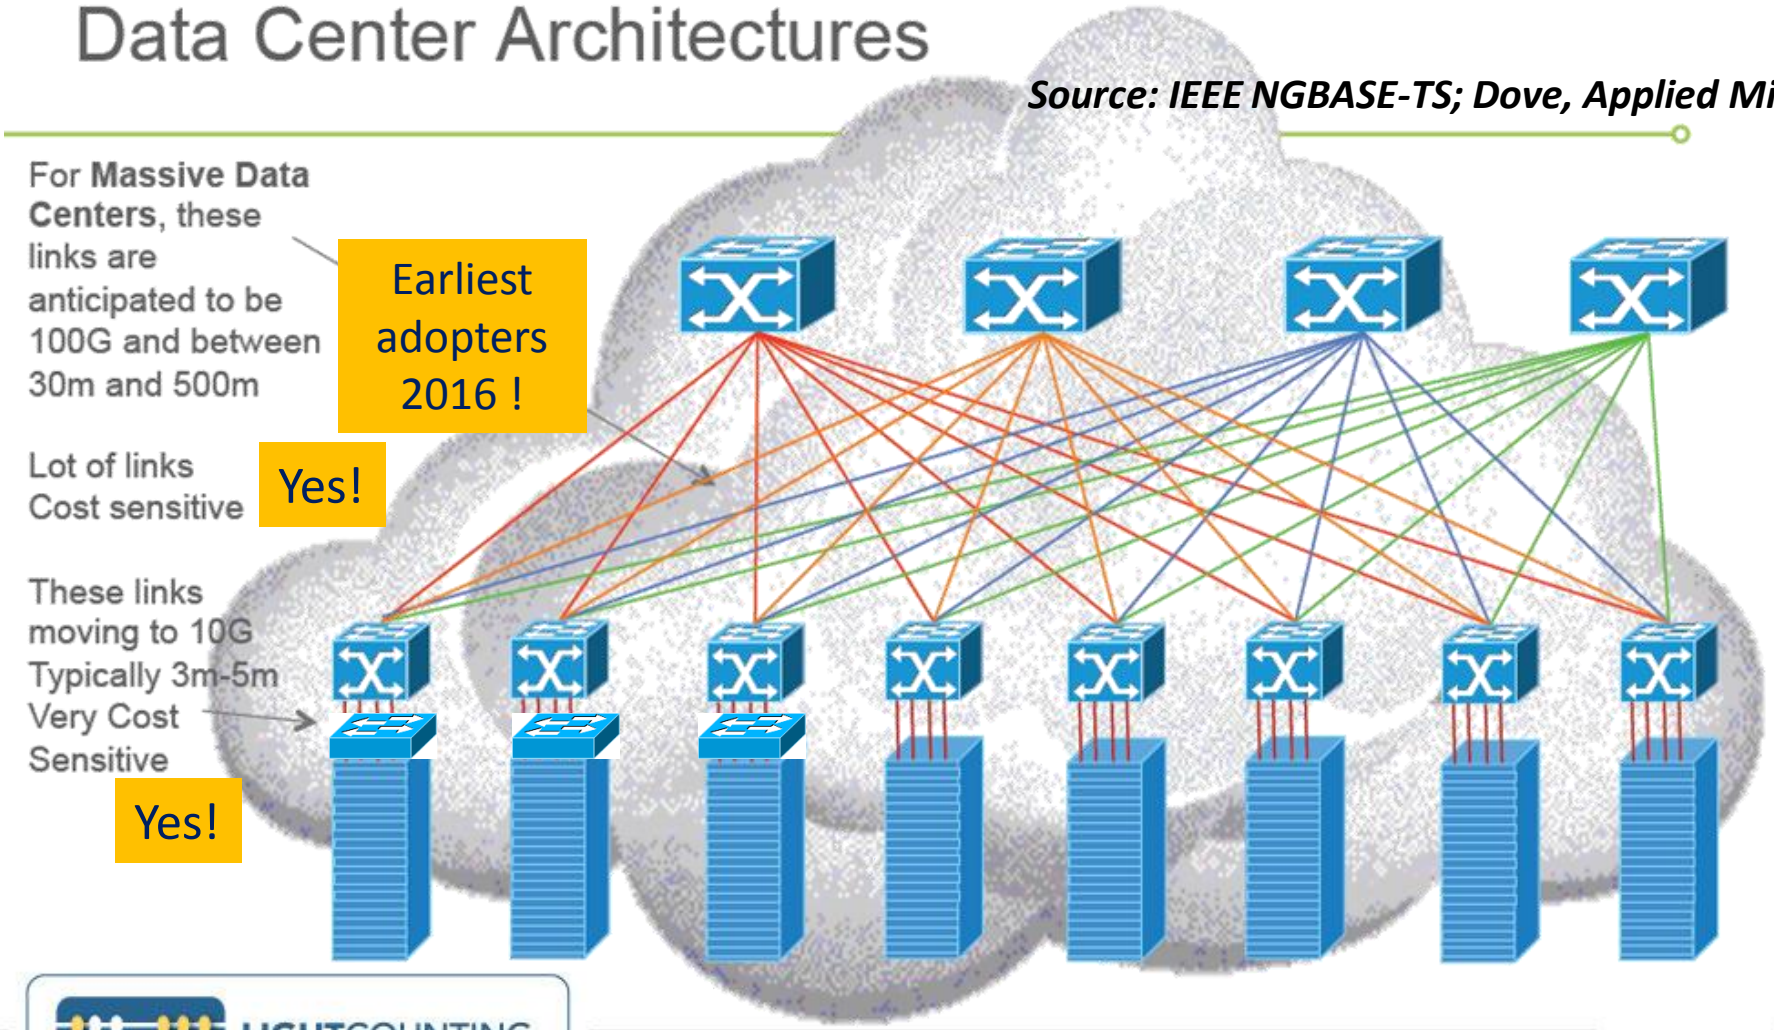

Cost per Optical Gigabit (Long Reach) Ethernet Modules Only

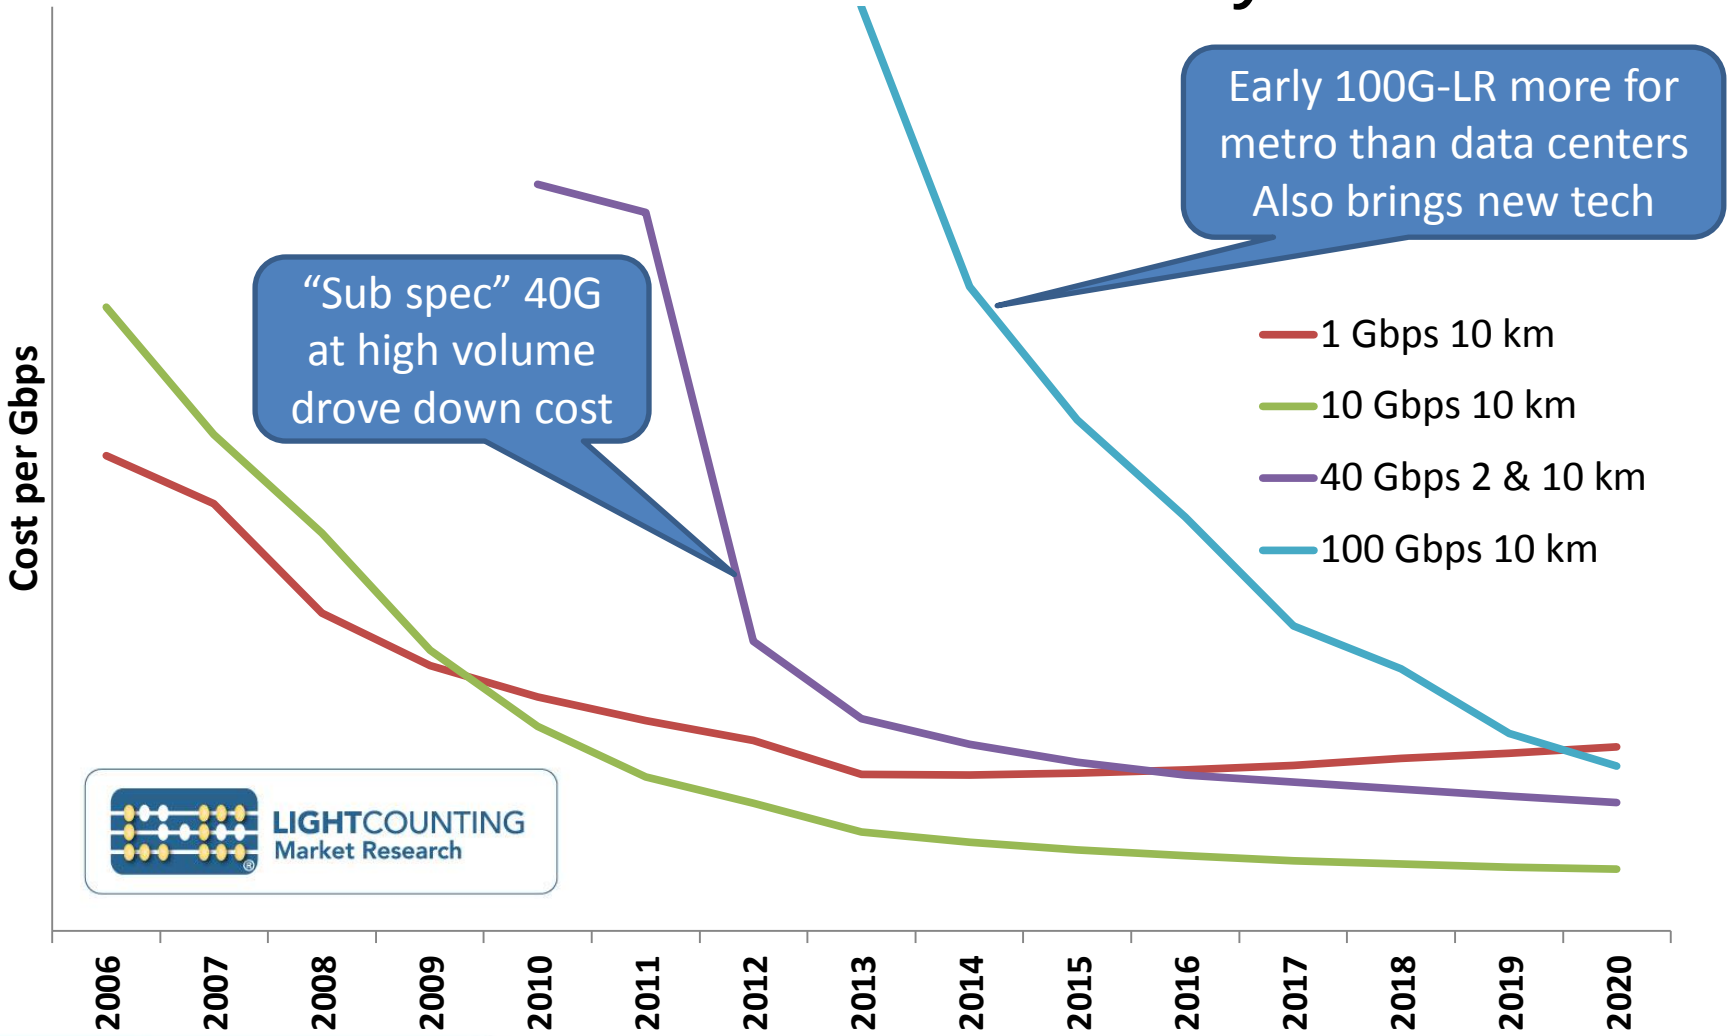

Lessons from 10G-LR

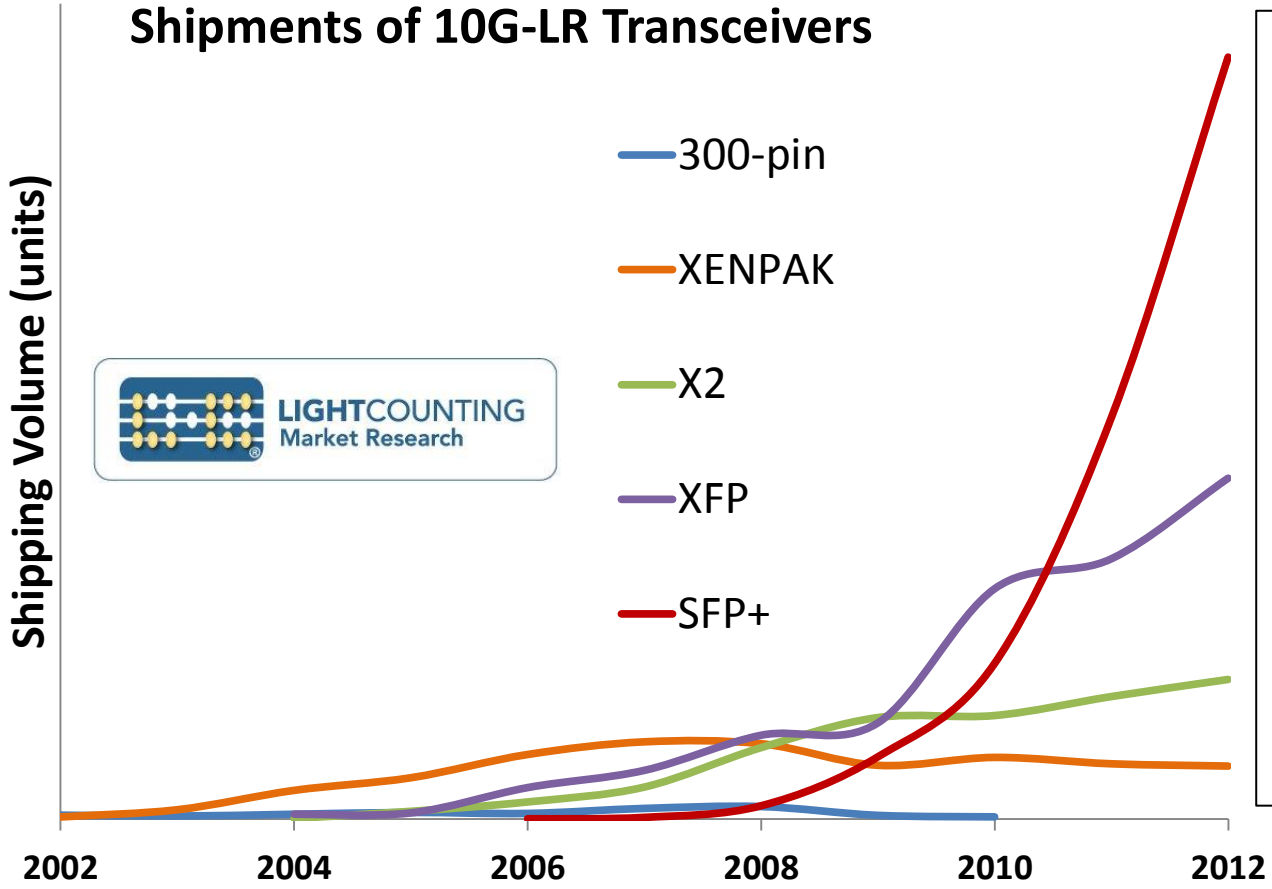

It took:

1. Serial in / Serial Out
2. smaller size,
3. lower power consumption,
4. a settled form factor, and
5. lower cost

for volume to ramp

Industry needs 100G ‘datacenter reach’ but fragmenting the market is not helpful

- 100G-PSM4 – Parallel Single Mode MSA
  - www.psm4.org
- 100G-CWDM4 MSA
  - www.cwdm4-msa.org
- 100G-CLR4 (CWDM with different receiver power)
  - www.intel.com/content/dam/www/public/us/en/documents/presentation/clr4-press-deck.pdf
- 100G 1550nm WDM silicon photonics MSA
  - www.openopticsmsa.org
- 100G silicon photonics Active Optical Cables

100 GbE Fragmentation Effects

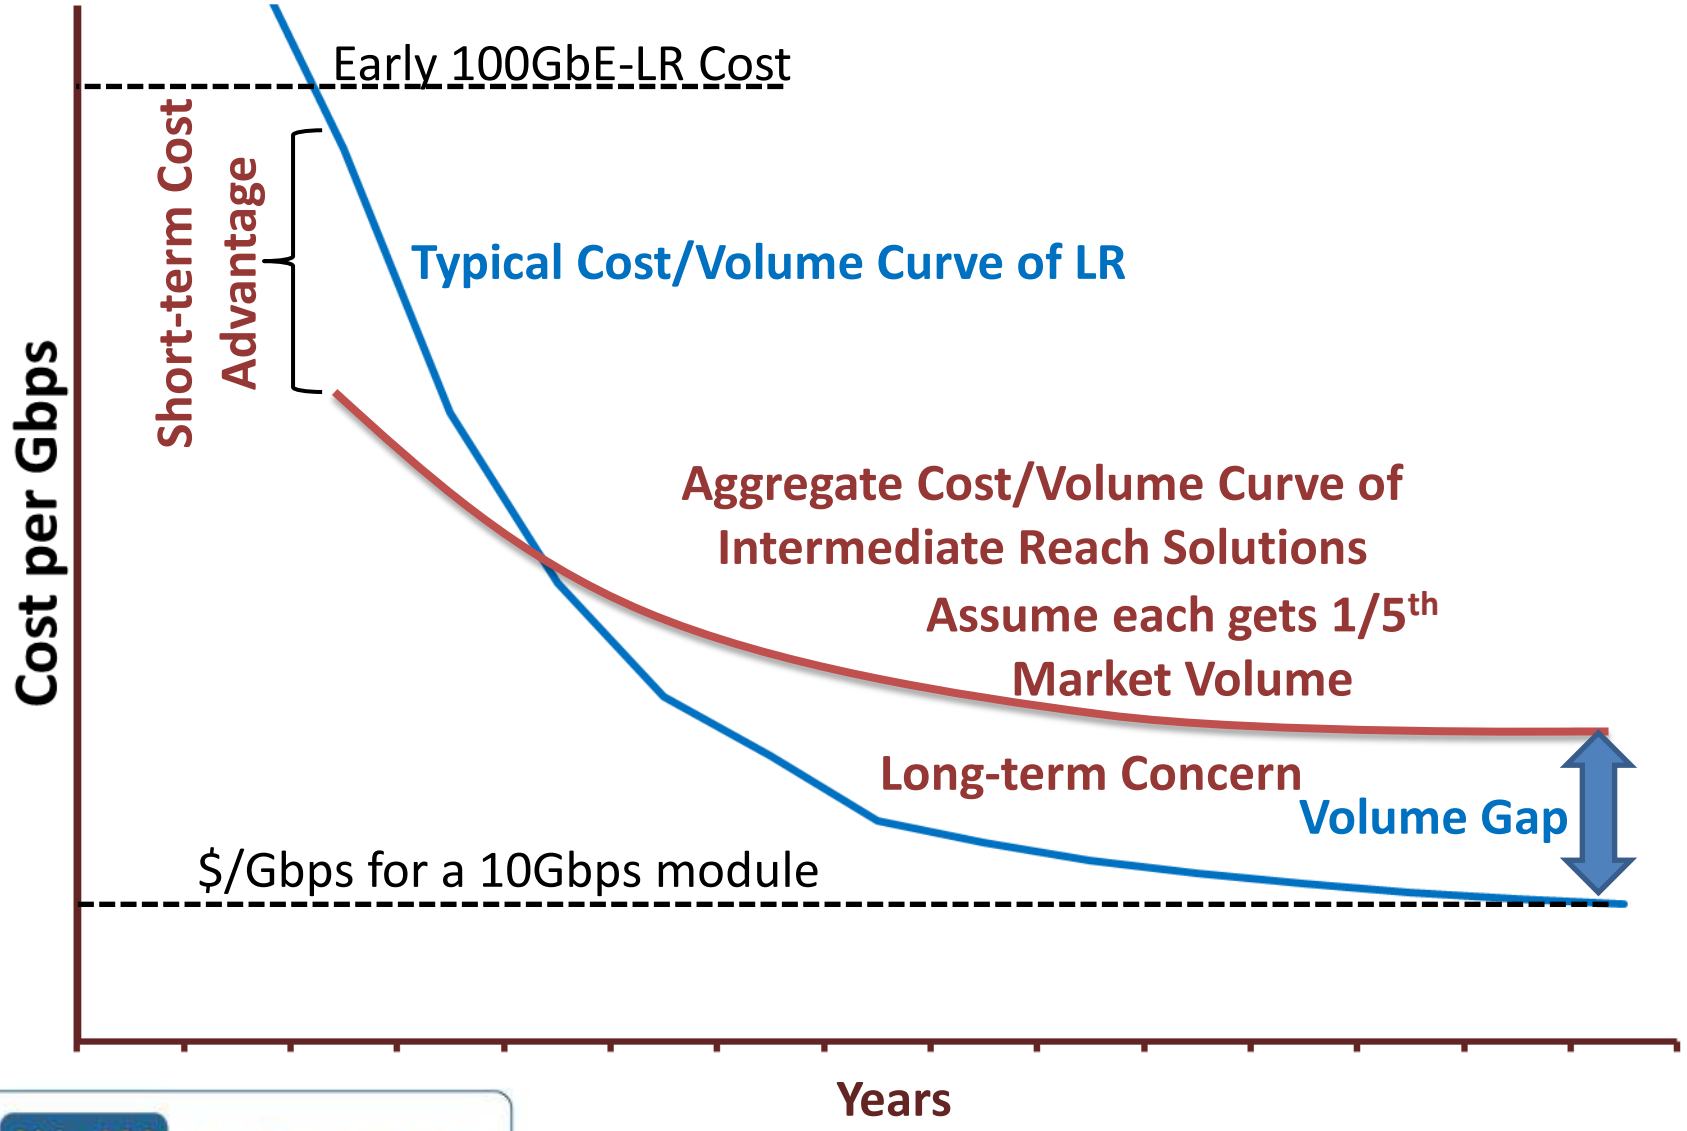

Getting to 400GbE and Beyond
Practical Optical Implementations

4 Lambdas on 4 Fibers

25 Gbps per lambda → **400 GbE**

50 Gbps per lambda → **800 GbE**

100 Gbps per lambda → **1.6 TbE**

End of the Road

Module Interface Width Affects Adoption

400G Density Challenges

- As interface speeds increase, bandwidth density must increase proportionally to maintain the value proposition.
  - Number of interface wires to/from the optics module affects physical dimensions and faceplate real estate.
- 100G Ethernet was initially specified using existing 10Gb/s electrical link technology which hindered adoption, and is now migrating to 25Gb/s electrical link technology.
  - Ten 10Gb/s links required per direction => 40 wires for a full duplex lane.
  - Faceplate density for early 100G systems was worse than 10G systems.
  - Four 25Gb/s links reduces wire count => 16 wires for a full duplex lane.
- Same migration path is expected for 400Gb/s Ethernet:
  - Early specifications will assume sixteen 25Gb/s links => 64 wires.
  - Migration path to eight 50Gb/s links will be necessary to facilitate widespread adoption => 32 wires. *In each direction*
- Reducing I/O counts also reduces power requirements. Power dissipation remains a significant system design issue.

Note:

Lesson: ≤16-bit Electrical Interface Necessary

Source: David Stauffer, Kandou Bus, SA

OIF OPTICAL INTERNETWORKING FORUM

100Gbps per Lambda will be the End Game for Datacenters for some time

Ethernet Rate			
10G	**40G**	**100G**	**400G**
4 x 3G			
1 x 10G	4 x 10G	10 x 10G	
	2 x 20G	4 x 25G	16 x 25G
	1 x 40G	2 x 50G ?	8 x 50G ?
		1 x 100G	4 x 100G

Does not scale

Will Scale Further

Need to Support Multiple FF Generations

400GbE Example

Influence of Wavelength Count

16 x 25G-λ

- One form factor generation supported before reverse mux is required
- CDP does not provide front-panel bandwidth density improvement over 4 x QSFP28
- Suitable for applications not requiring interop over form factor generations such as mega data-center applications

8 x 50G-λ

- Two form factor generations supported before reverse mux is required
- Second generation doubles front-panel bandwidth density
- Not desirable for router applications since front panel bandwidth density is only doubled while preserving interop over form factor generations

4 x 100G-λ

- Three form factor generations supported with progressive cost reduction
- Quadruples front panel bandwidth density while preserving interop over form factor generations

13 IEEE P802.3bs 400GbE Task Force Norfolk, VA, May 2014

JUNIPER NETWORKS

Source: Jeff Maki, IEEE 802.3bs (maki_3bs_01a_0514.pdf)

100Gbps per Lambda could yield an attractive SFP+ solution

The SFP and QSFP Dance

- 10GbE takes off with SFP+, but goes to QSFP+ for density

1X10G
10GbE SFP+

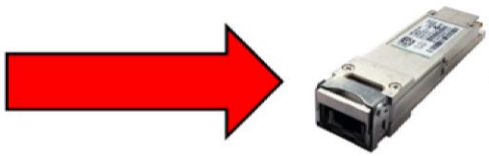

4X10G
4X10GbE QSFP+

- 40GbE takes off with QSFP+, but could shift to SFP+

4X10G
40GbE QSFP+

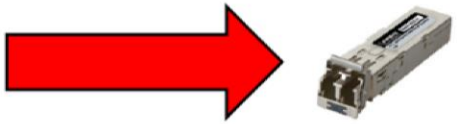

1X40G
40GbE SFP+

- 100GbE takes off with QSFP+, but could shift to SFP+

4X25G
100GbE QSFP28

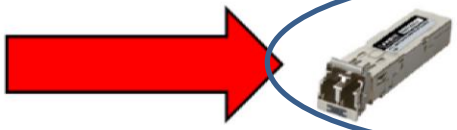

1X100G
100GbE SFP+

5/15/2014 10

Source: Scott Kipp, 40G Serial CFI Planning Meeting: May 13, 2014

We Need Building Blocks for Ethernet's Optical Future

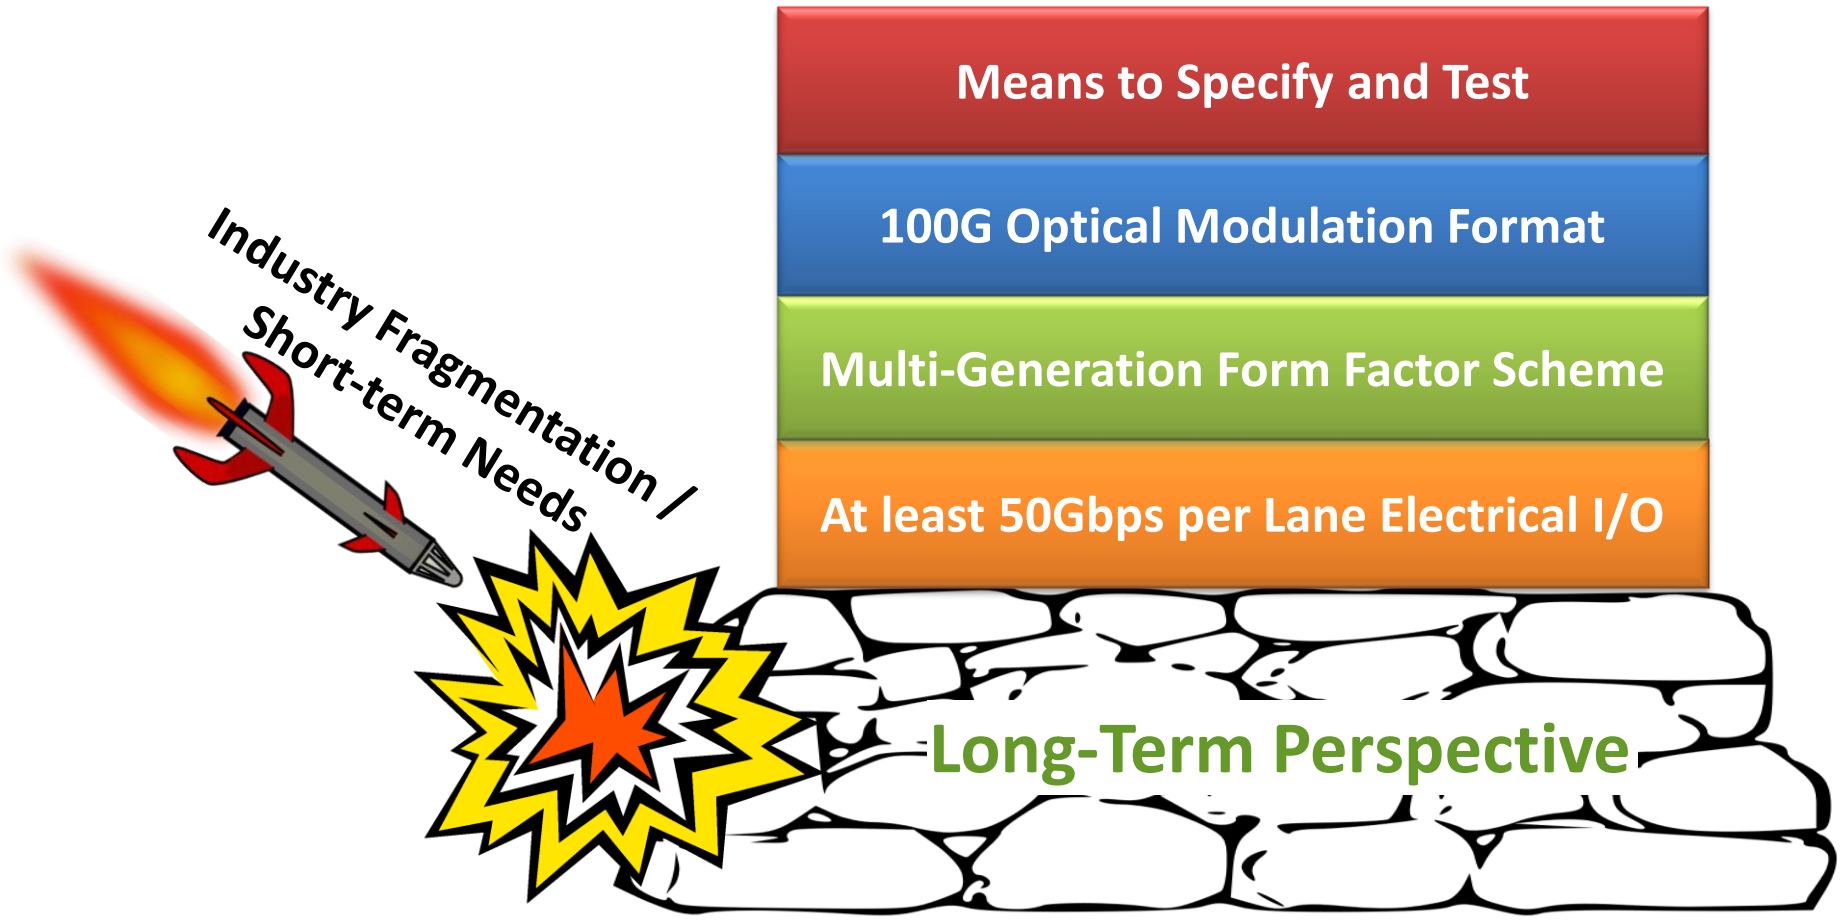

State of the Industry

Average Net Income Across the Supply Chain (in Millions of Dollars)

Companies	2008	2009	2010	2011	2012
Content Providers	$609	$882	$1,083	$1,271	$1,426
Service Providers	$4,628	$7,298	$8,868	$6,312	$5,862
System Vendors	$802	$797	$1,267	$905	$687
IC Vendors IC	$417	$316	$1,059	$1,080	$773
OC and Module Vendors	$19	($8)	($1)	($13)	($20)

Note: Numbers shown in red indicate a loss.

Average Net Margin Across the Supply Chain

Companies	2008	2009	2010	2011	2012
Content Providers	12%	16%	16%	14%	12%
Service Providers	6%	10%	12%	8%	8%
System Vendors	5%	6%	9%	6%	5%
IC Vendors IC	10%	9%	22%	20%	15%
OC and Module Vendors	2%	(4%)	(2%)	(6%)	(5%)

Note: Numbers shown in red indicate a loss.

Source: Published information of public companies

Closing Perspectives

- It takes volume to get competitive $/Gbps. Market fragmentation is not helpful.
- The lessons of 10G, 40G and 100G suggest we should endeavor to take a long-term perspective on PMD choices.
- Anything less than 100GbE on a single lambda limits the practical implementation of higher Ethernet speeds
  - Caveat: A direct jump from 25G to 100G per lambda may not be technically practical at this time.
- 100G per lambda is the only solution that scales and offers long technical life over three form factor generations
- 100G per lambda is one of several key building blocks for the long-term future of Ethernet
- The optical transceiver business can be financially challenging. Roadmap uncertainty results in added investment.