LIGHTCOUNTING
Market Research on High-Speed Interconnects
Datacom, Telecom, CATV, FTTX, Consumer markets

IEEE Study Group 400G Applications Ad Hoc

Ethernet Deployment from an Optical Module Perspective

Dale Murray, LightCounting
dale@lightcounting.com

Flatter Leaf & Spine architectures mean more high-speed links

Data Center Architectures

Source: IEEE NGBASE-TS; Dove, Applied Micro

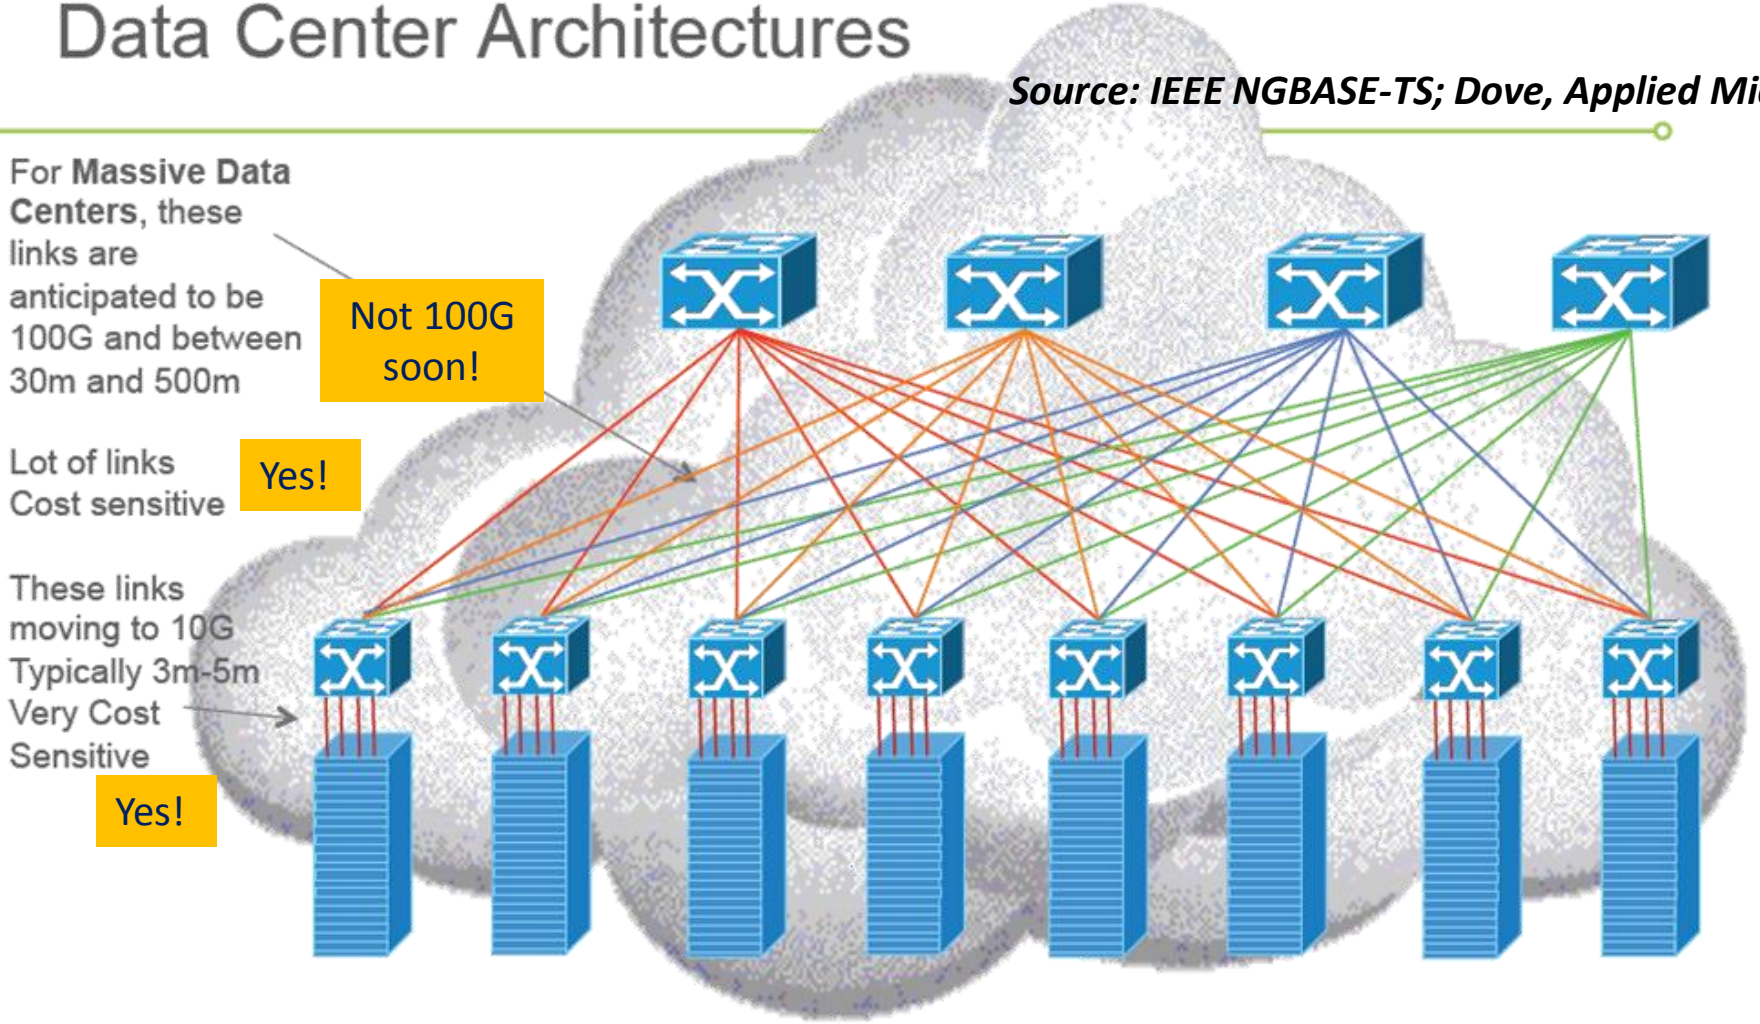

Ethernet Optical Transceiver Market by Data Rate

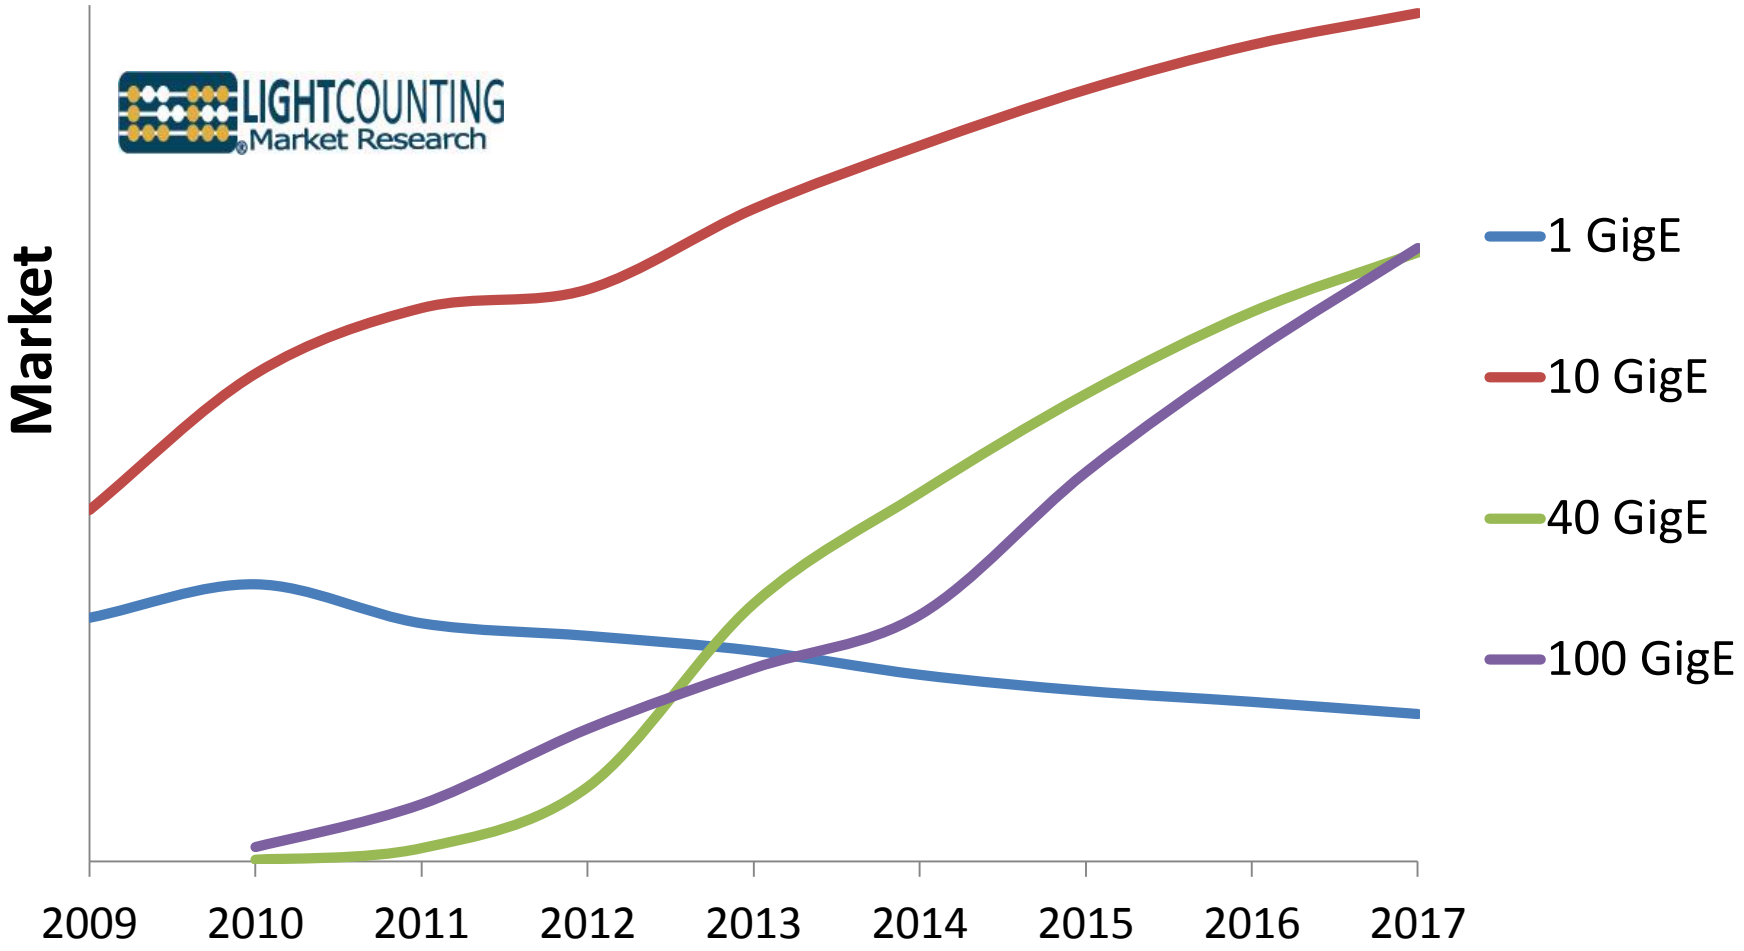

Ethernet Optical Transceiver Unit Shipments by Data Rate

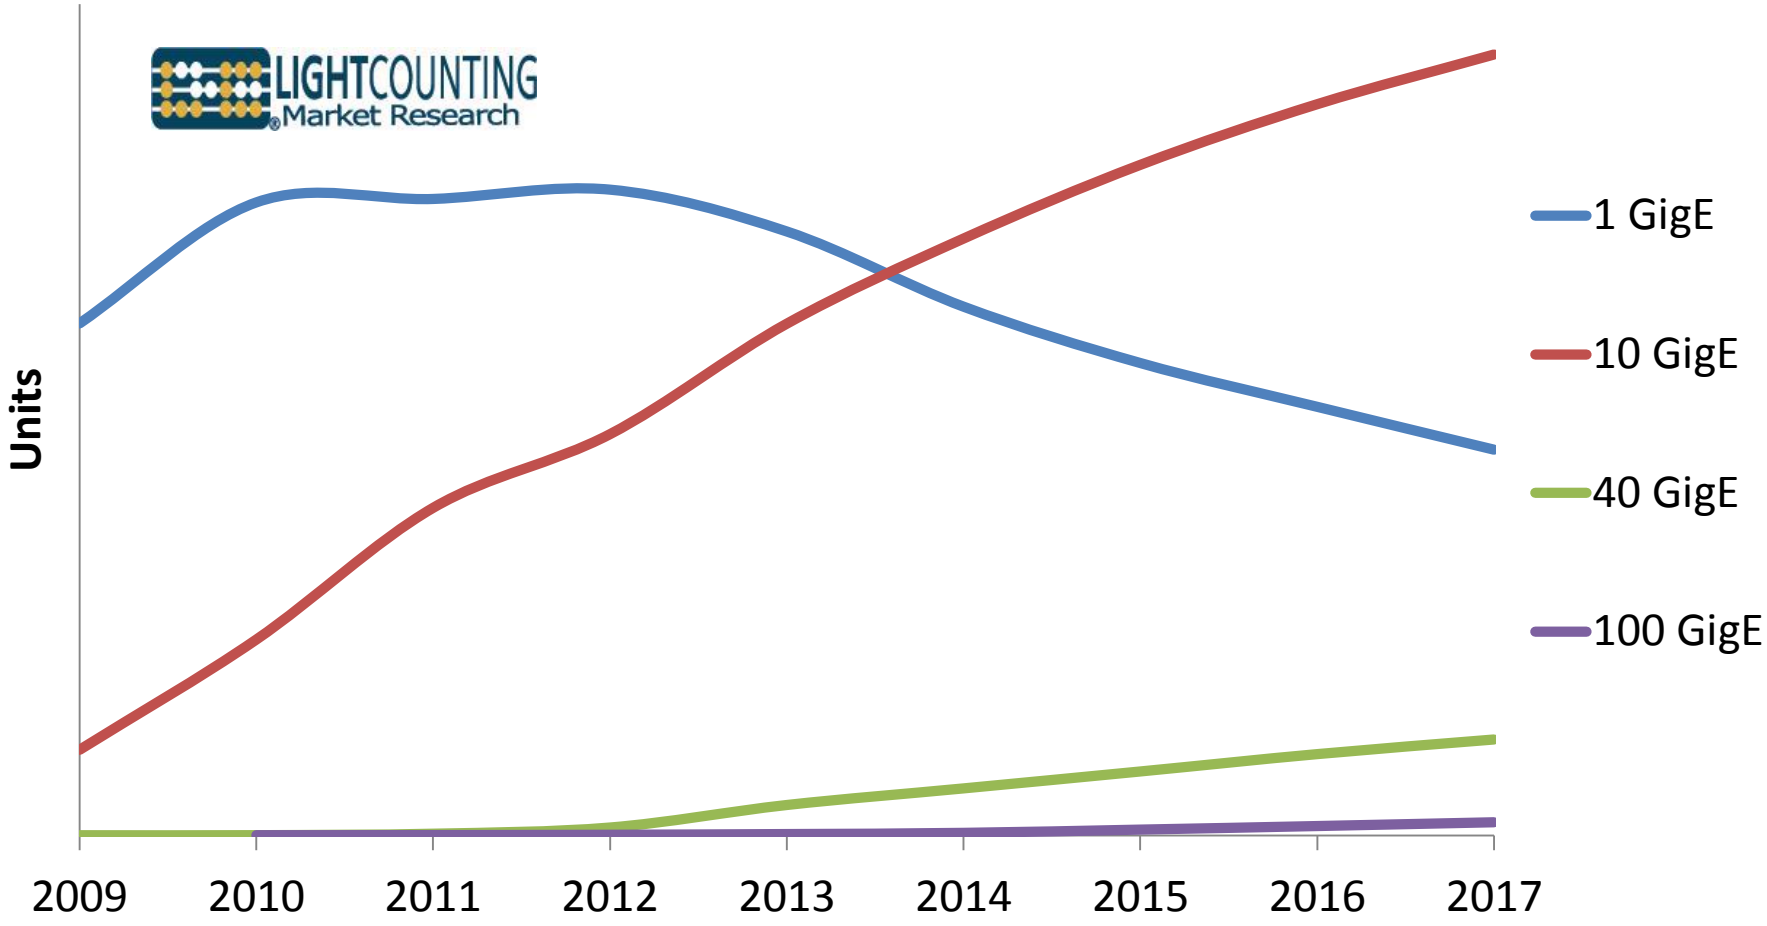

Ethernet Optical Transceiver Unit Shipments by Reach

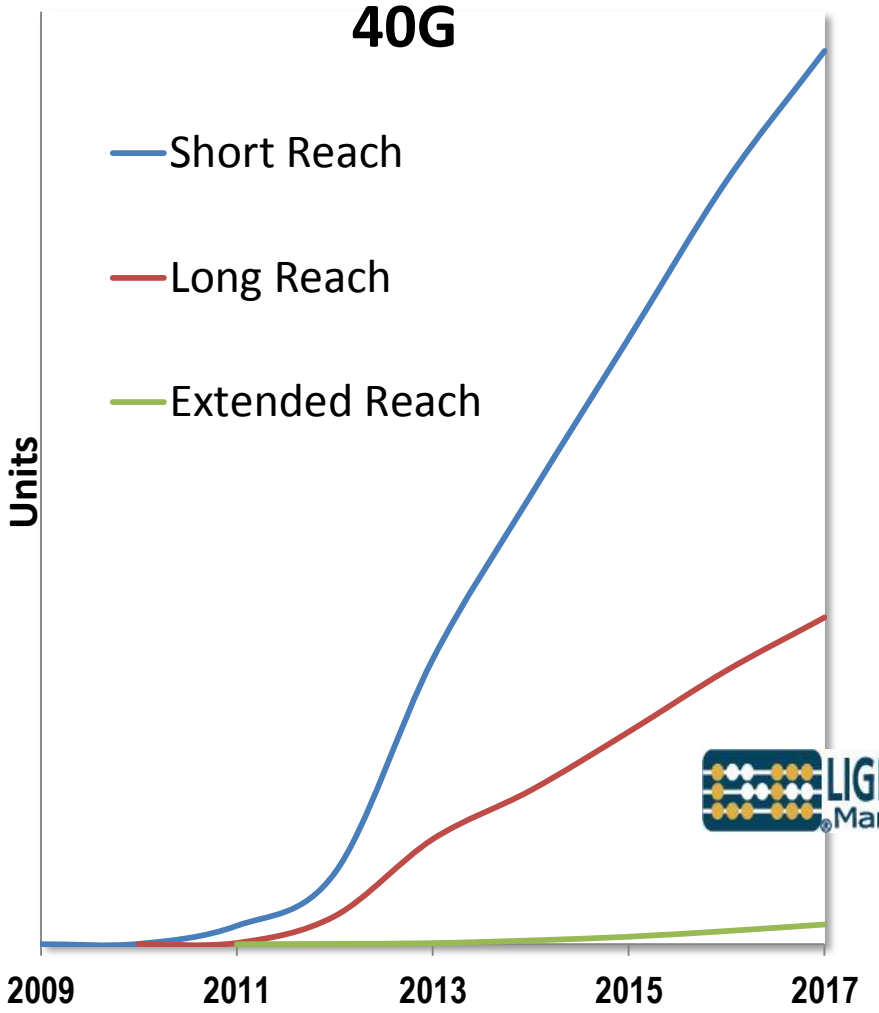

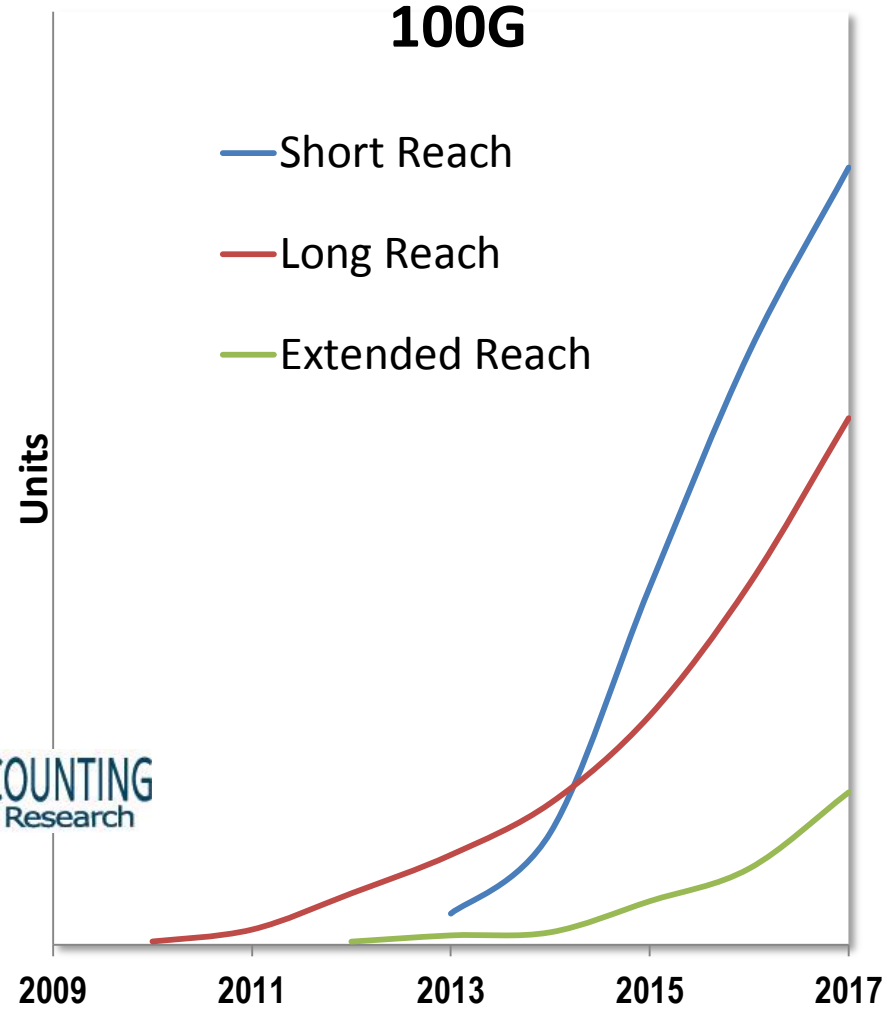

Breaking up 40G to 4x10G

Spine switch

Leaf switch

This appears to be a clear majority for connecting QSFP+ ports.

Difficult to count native 40GigE ports

40G Leaf/Spine

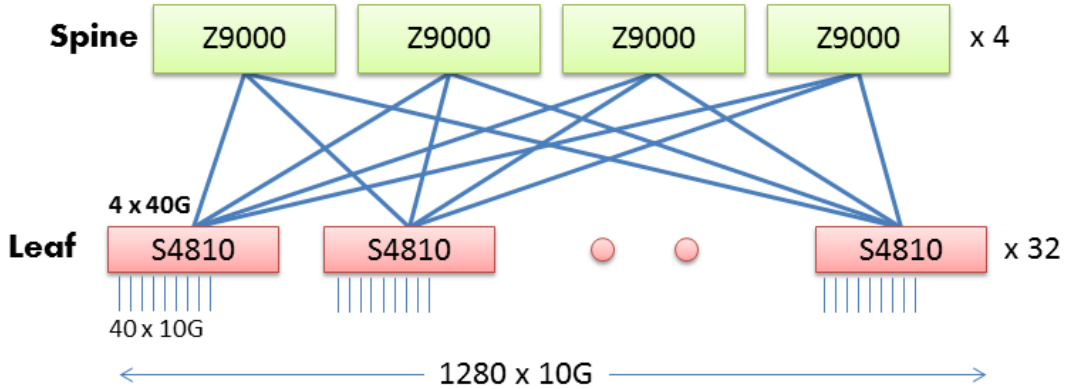

10G Leaf/Spine

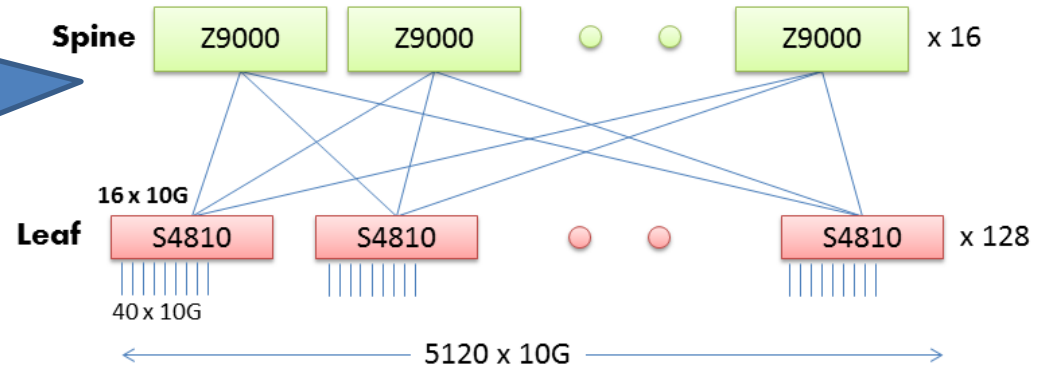

Source: bradhedlund.com

Cost per Optical Gigabit (Short Reach) Ethernet Modules Only

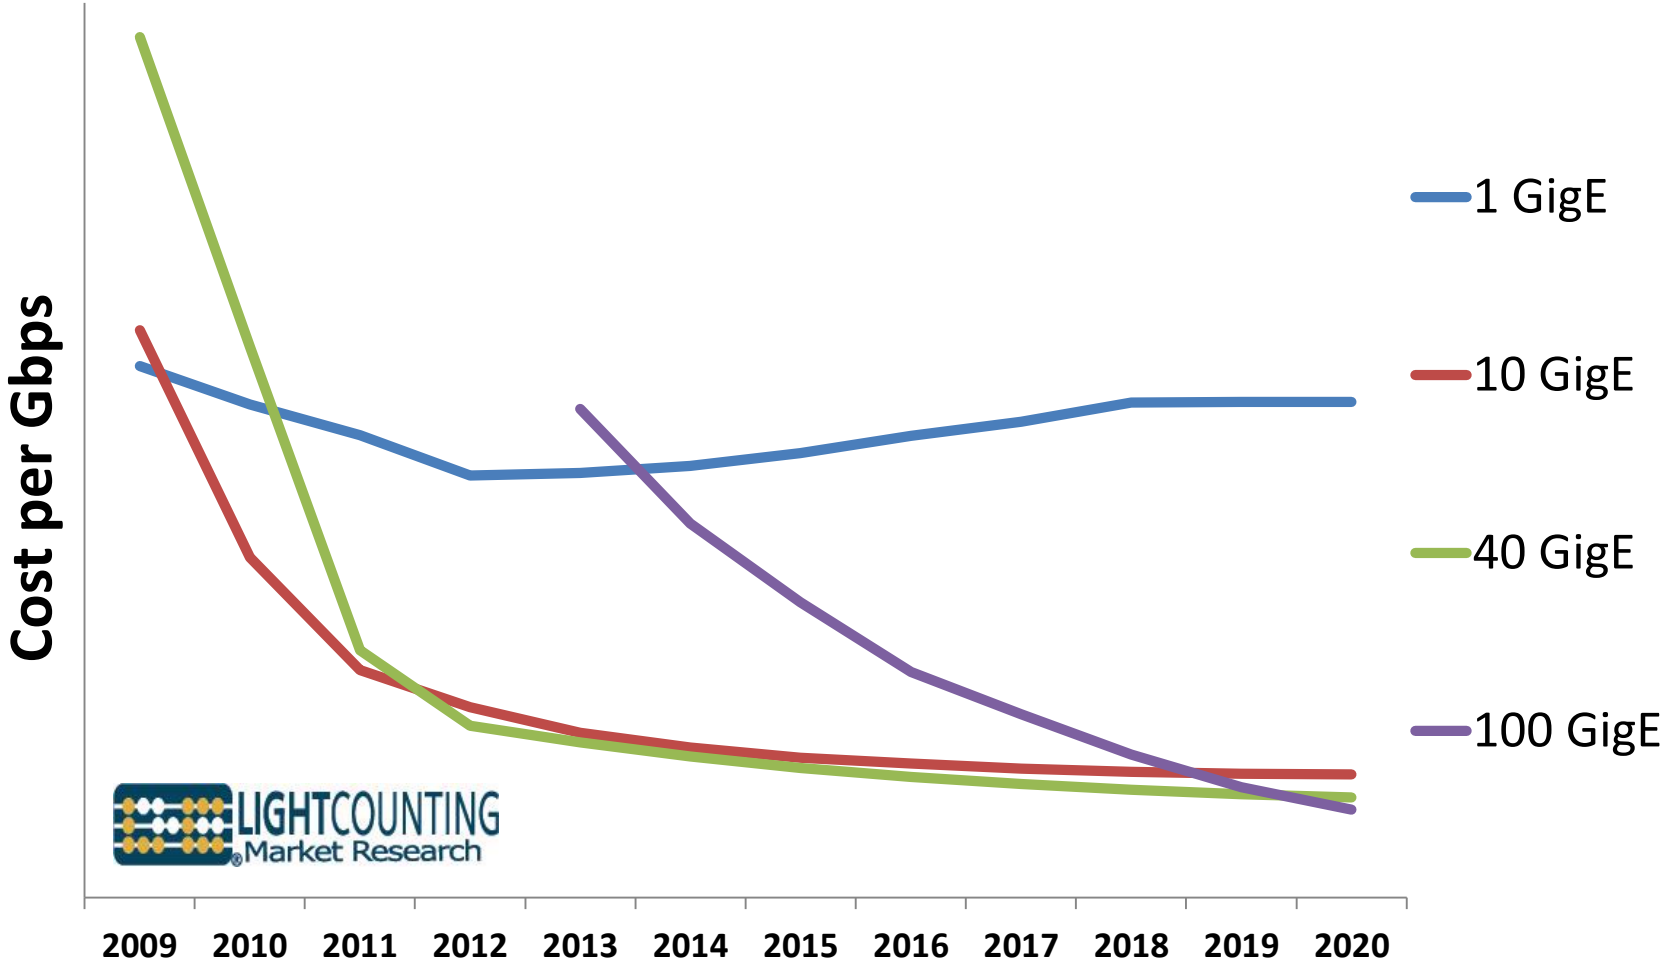

Cost per Optical Gigabit (Long Reach) Ethernet Modules Only

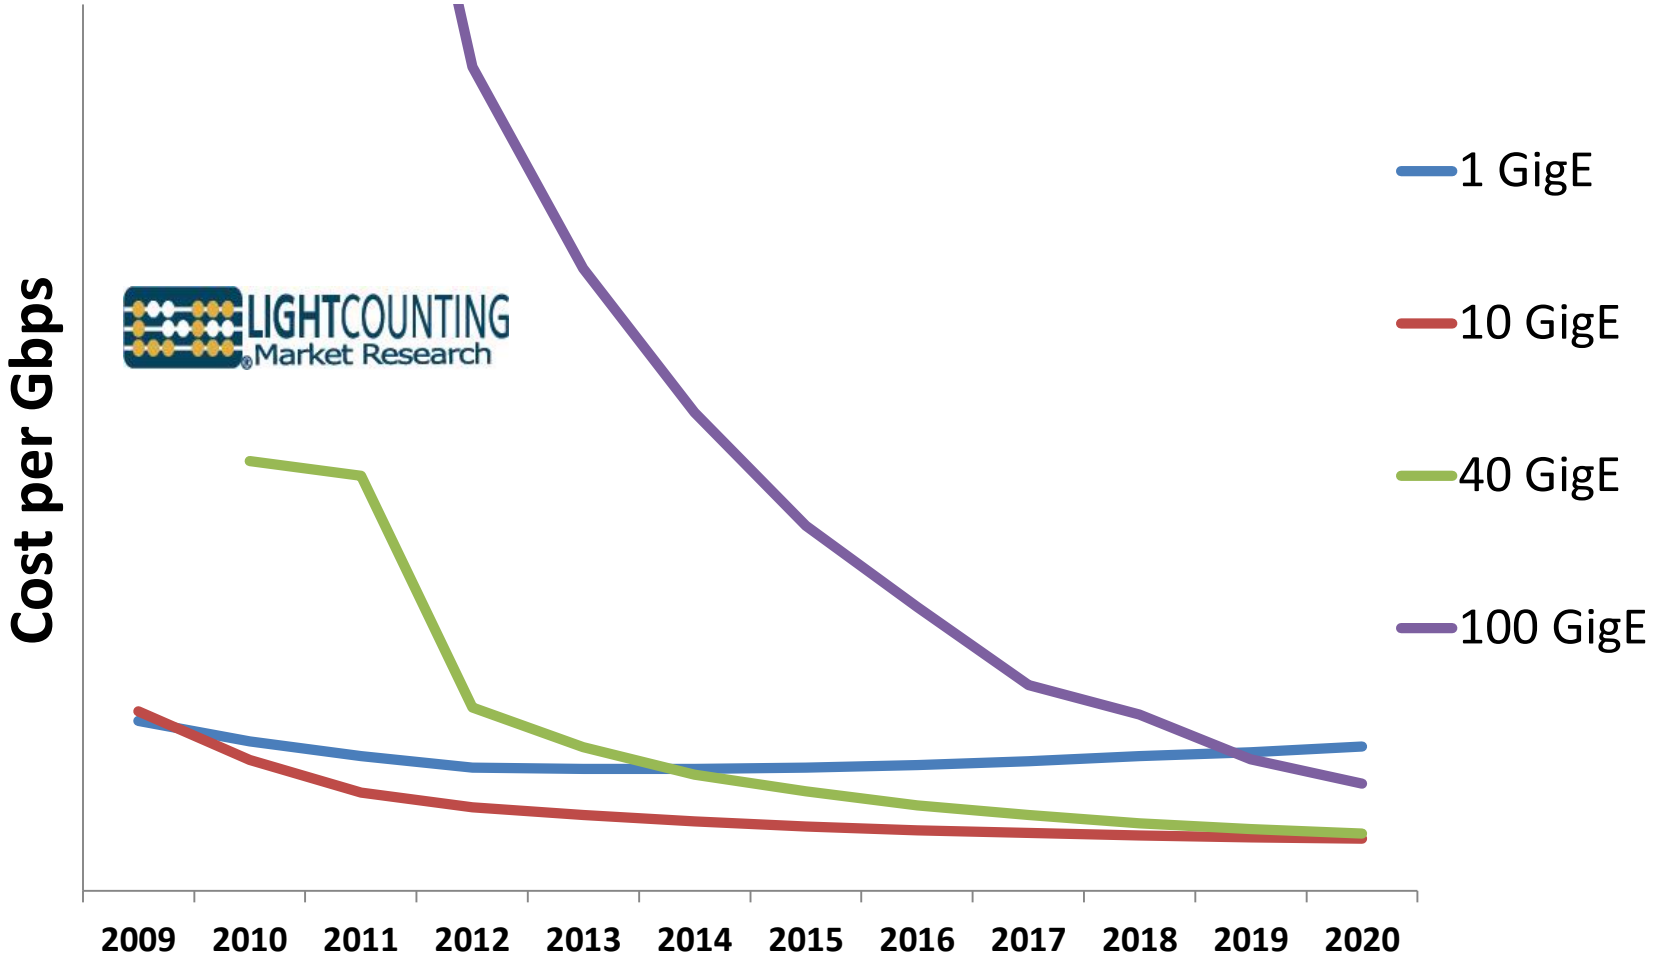

Lessons from 10G

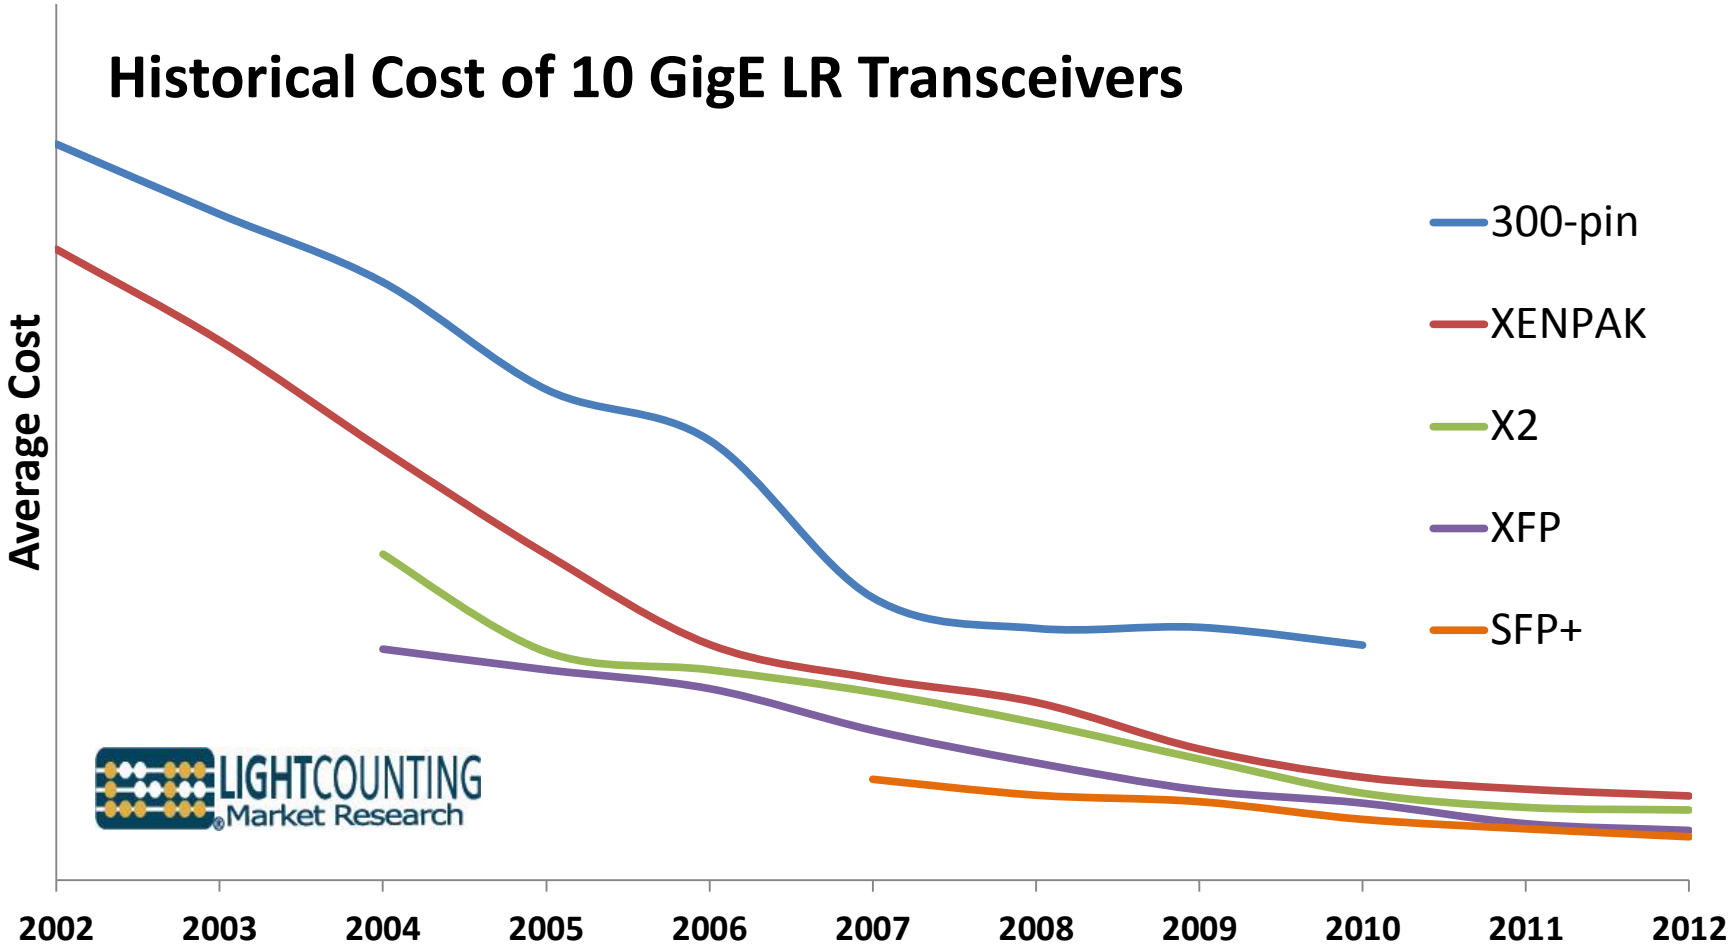

Lessons from 10G

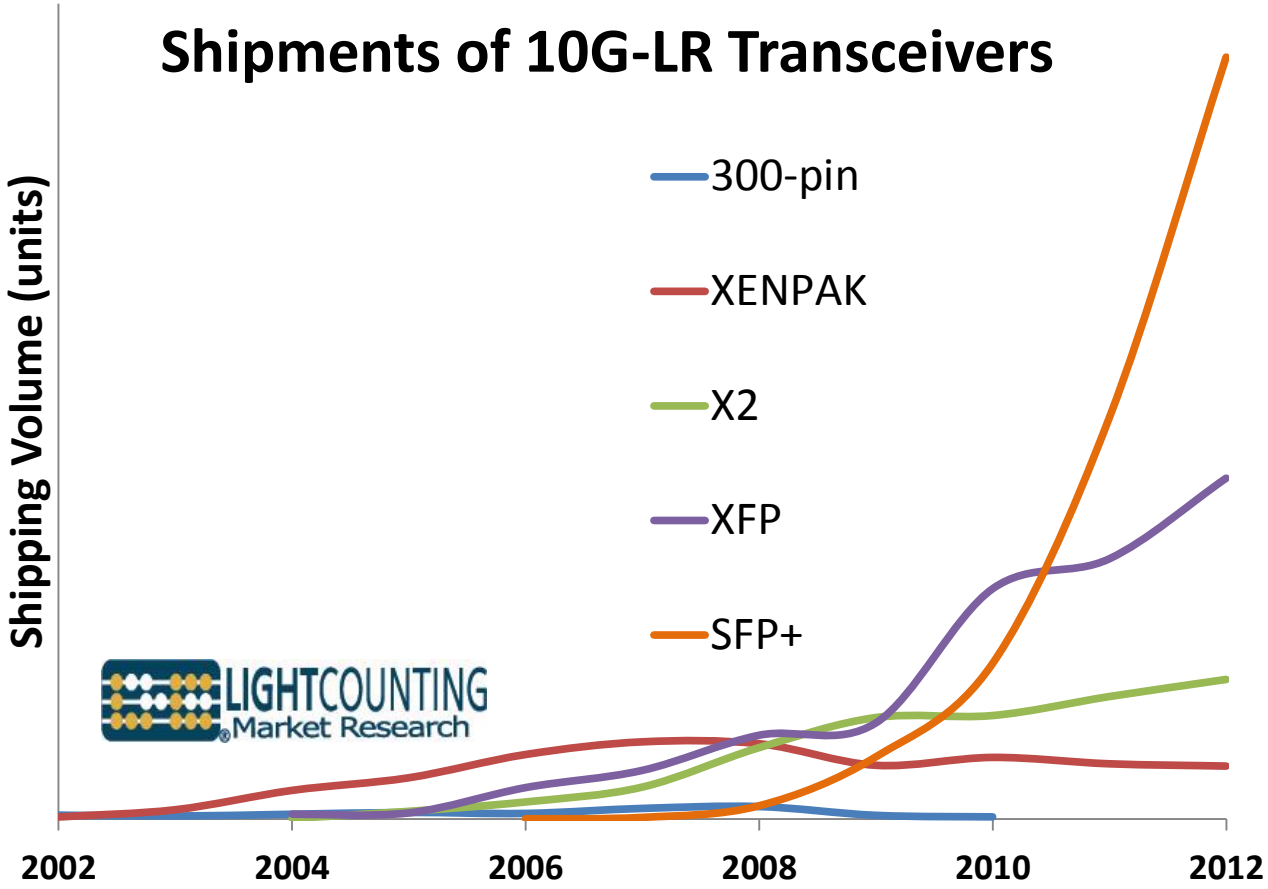

It took:

1. smaller size,
2. lower power consumption,
3. a settled form factor, and
4. lower cost

for volume to ramp

100G-LR will go through form factors similar to 10G-LR history

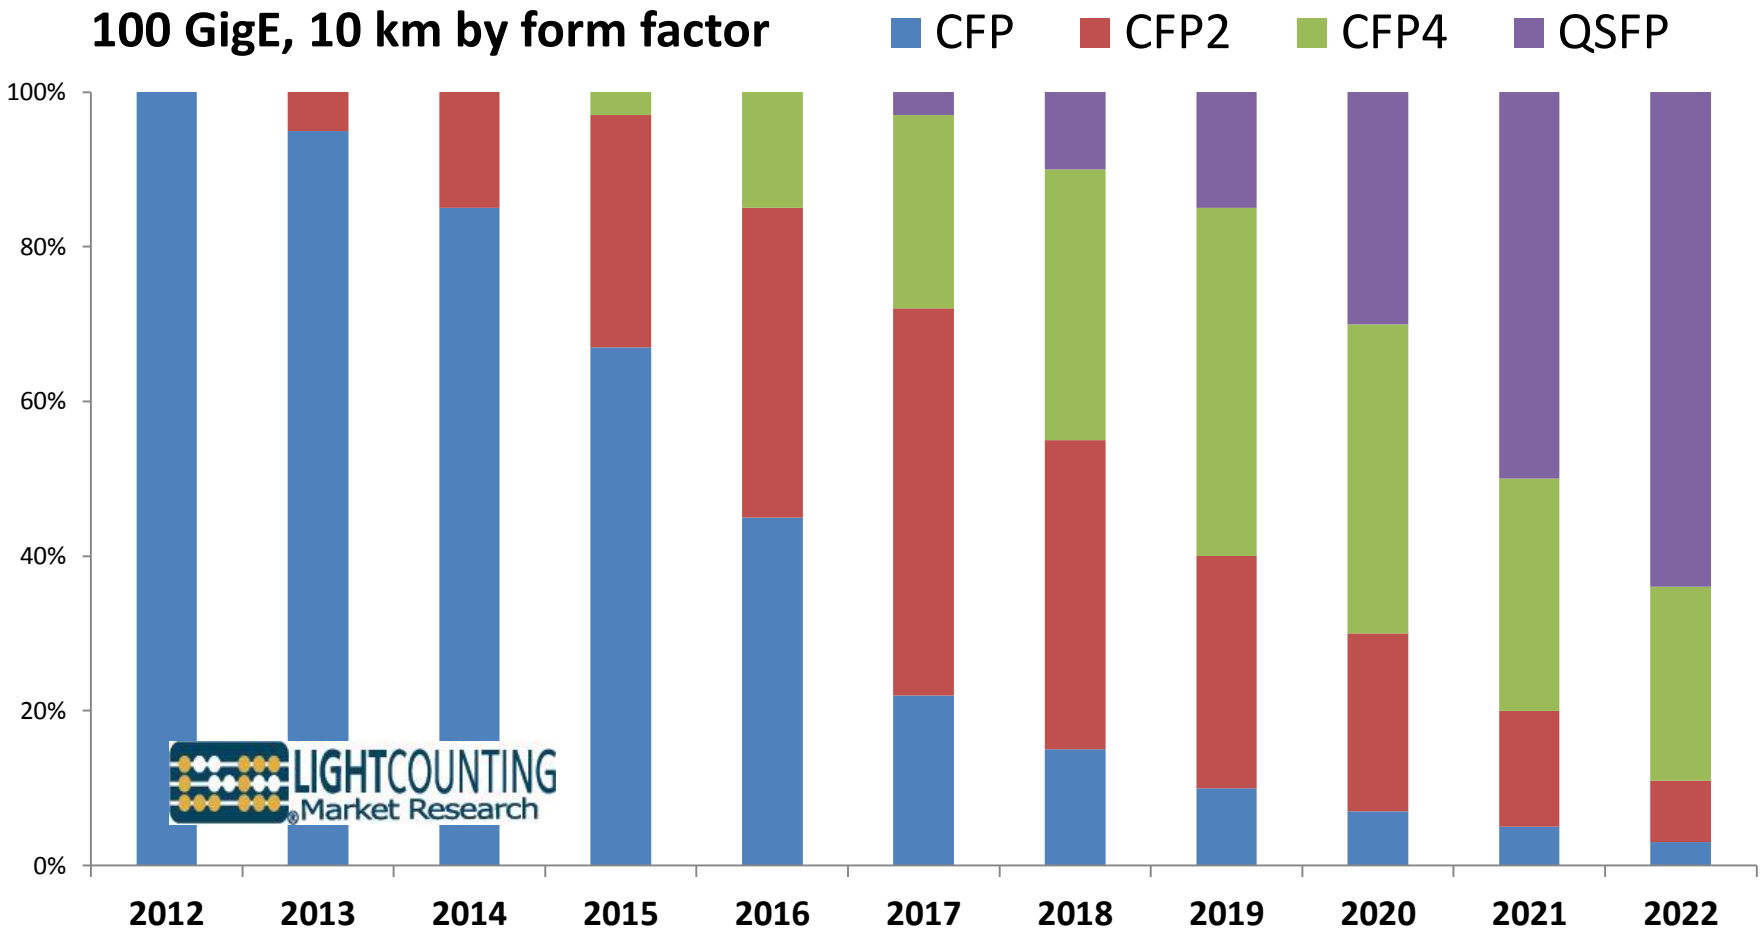

Note that CFP2 may support multi-channel 40G and possibly multi-channel 100G

40 GigE has settled on one form factor

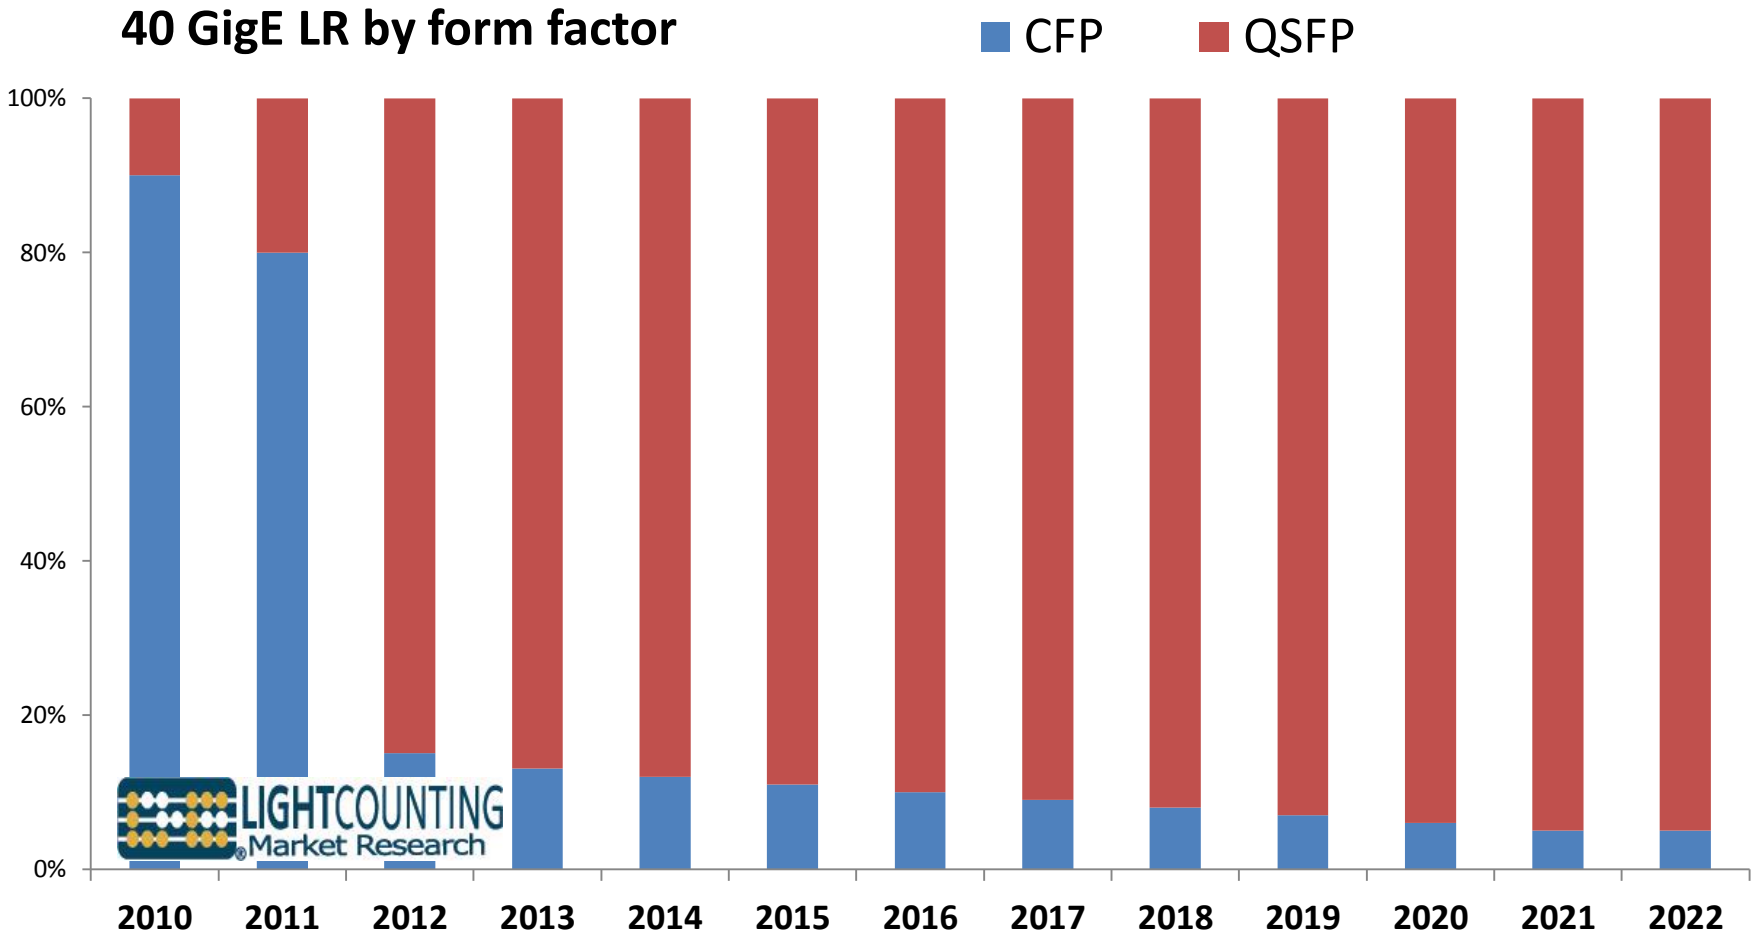

Note that multi-channel 40G may get deployed in the CFP2 form factor.

Lessons from 10G: the LX Story

10 years ago with legacy OM1 fiber:

- Reach of 1G-SR modules was 500m
- Reach of 10G-SR modules was just 100m

Customers demanded a 500m solution

- Suppliers invested heavily and introduced 10G-LX modules (CWDM)
- Volume peaked at <100K units then fell. (LightCounting predicted this ☺)
- Installed base of fiber became obsolete more quickly than many expected

Higher Volumes are key for lower cost

Slide from Julie Eng, Finisar

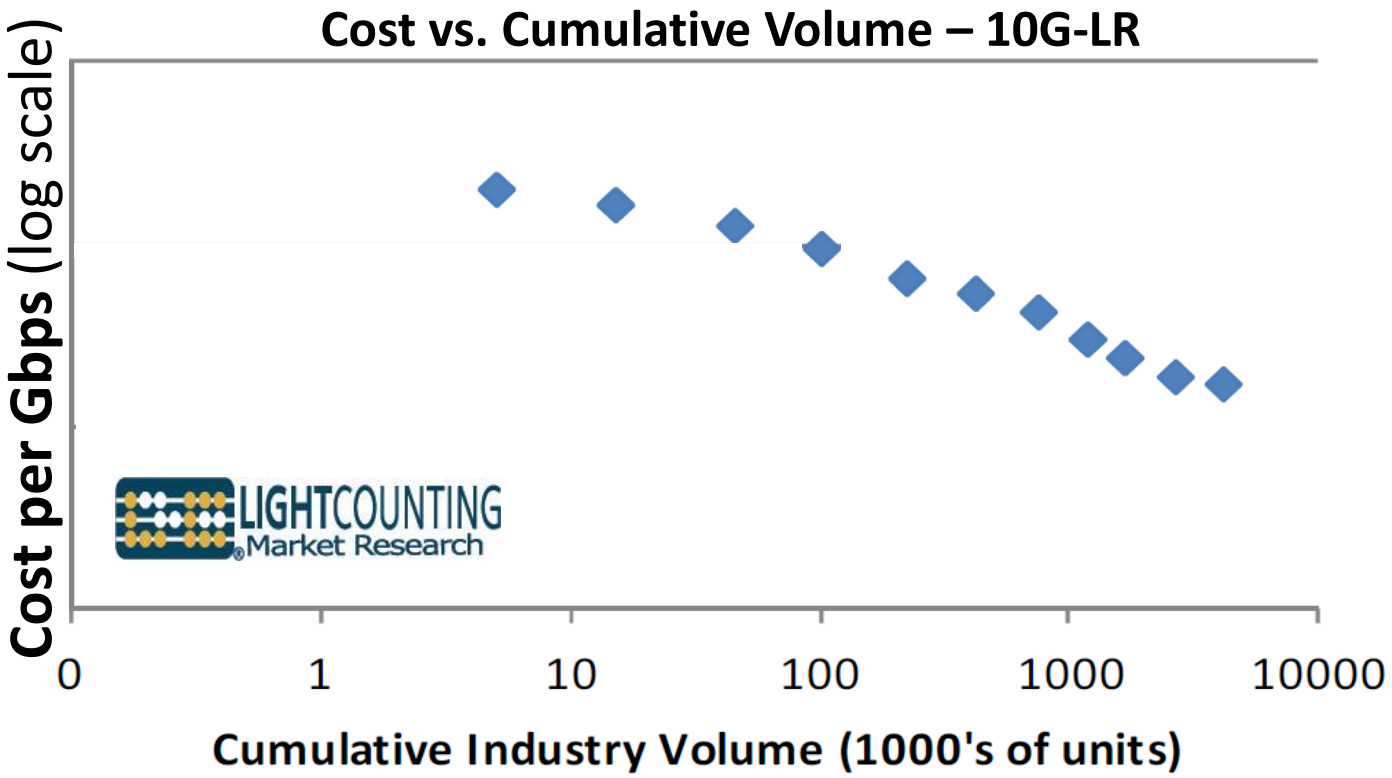

- First generation of any new technology is expensive
- Industry cumulative volume of >1M required to get to desired cost points
- For 10G, this evolution took 6 form factors and10 years

100GigE has to follow the same curve

Slide from Julie Eng, Finisar

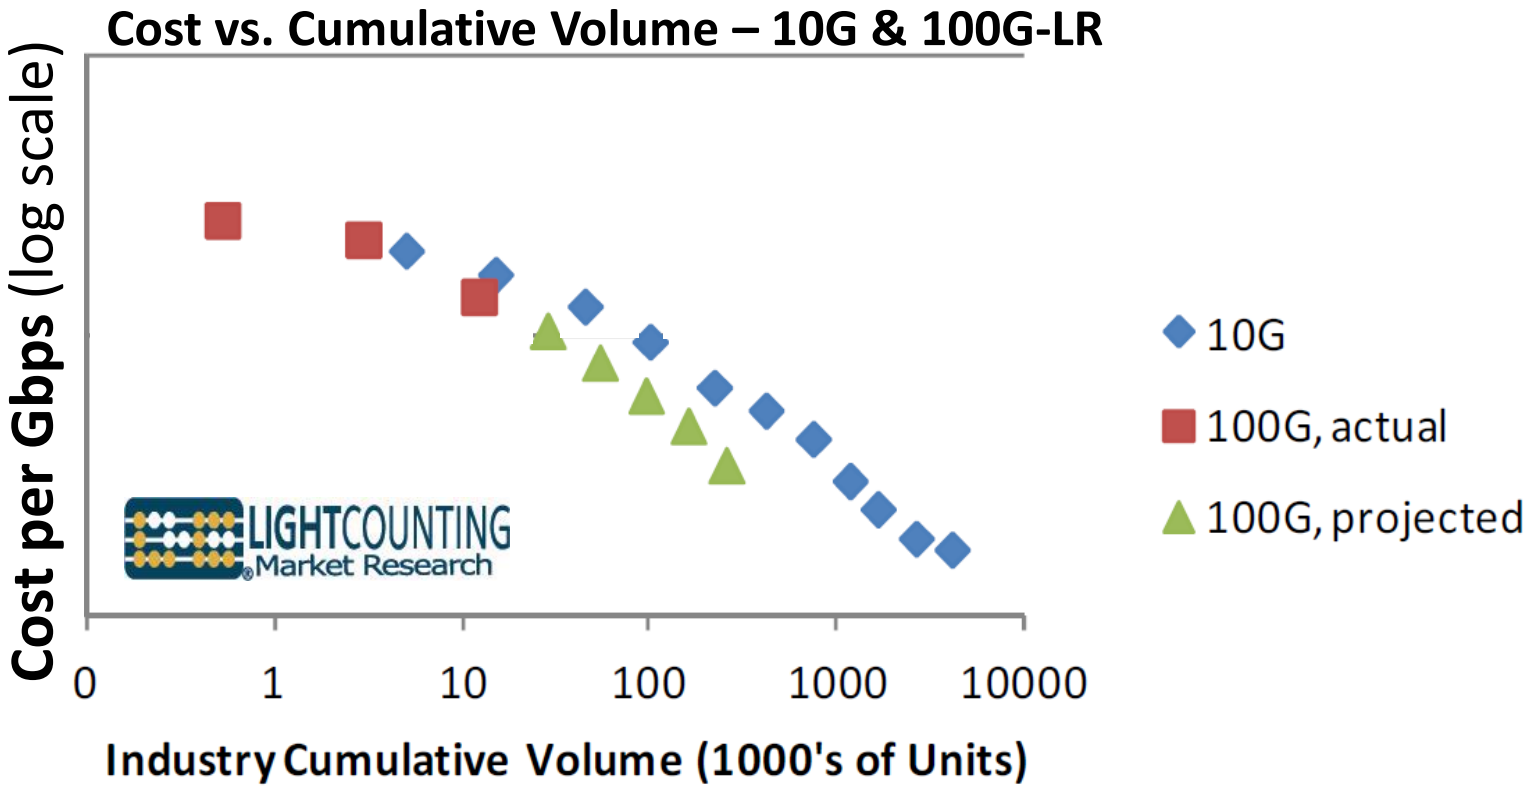

- 100G falls on the same curve as 10G for the volume shipped
- Cost projected to erode more quickly than 10G
- Best way to reduce cost of 100G components: Bring on the Volume!

Technology adoption curves

Telecom:
100G shipments ramping fast now

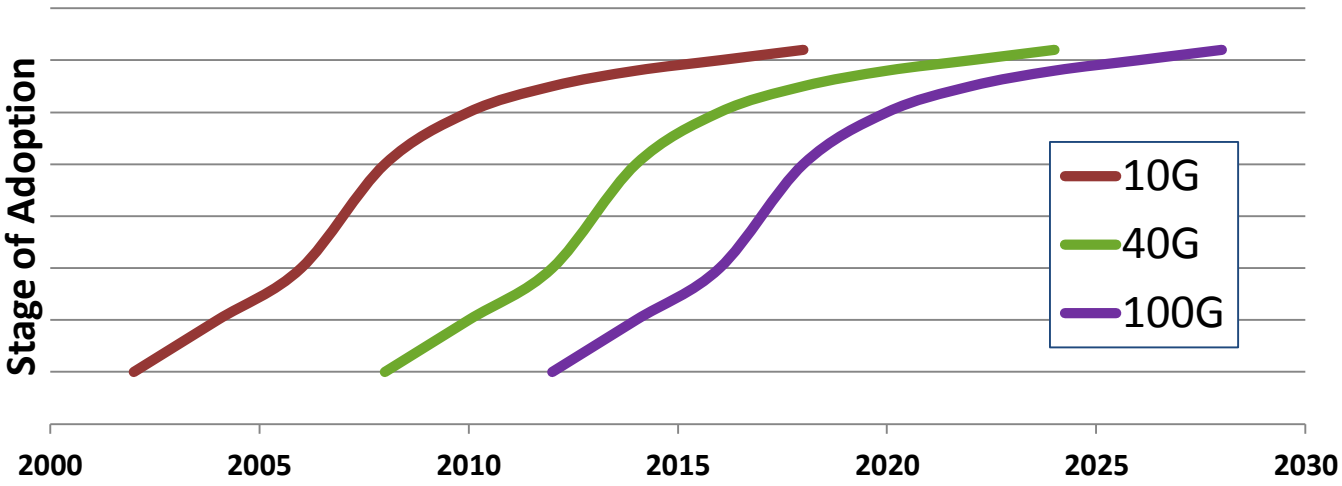

Datacom:
10GigE still ramping fast and 40 GigE is just starting now. 100GigE is not ramping yet.

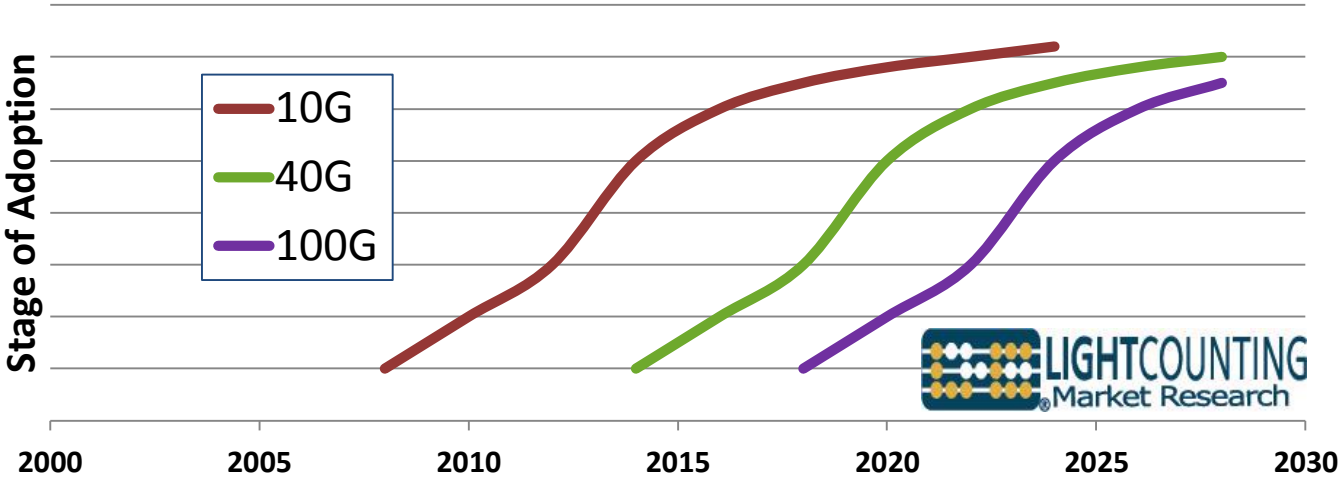

Historical and forecast data; Normalized to volume, Smoothed and Extrapolated

Beyond Standards: Extended Short Reach

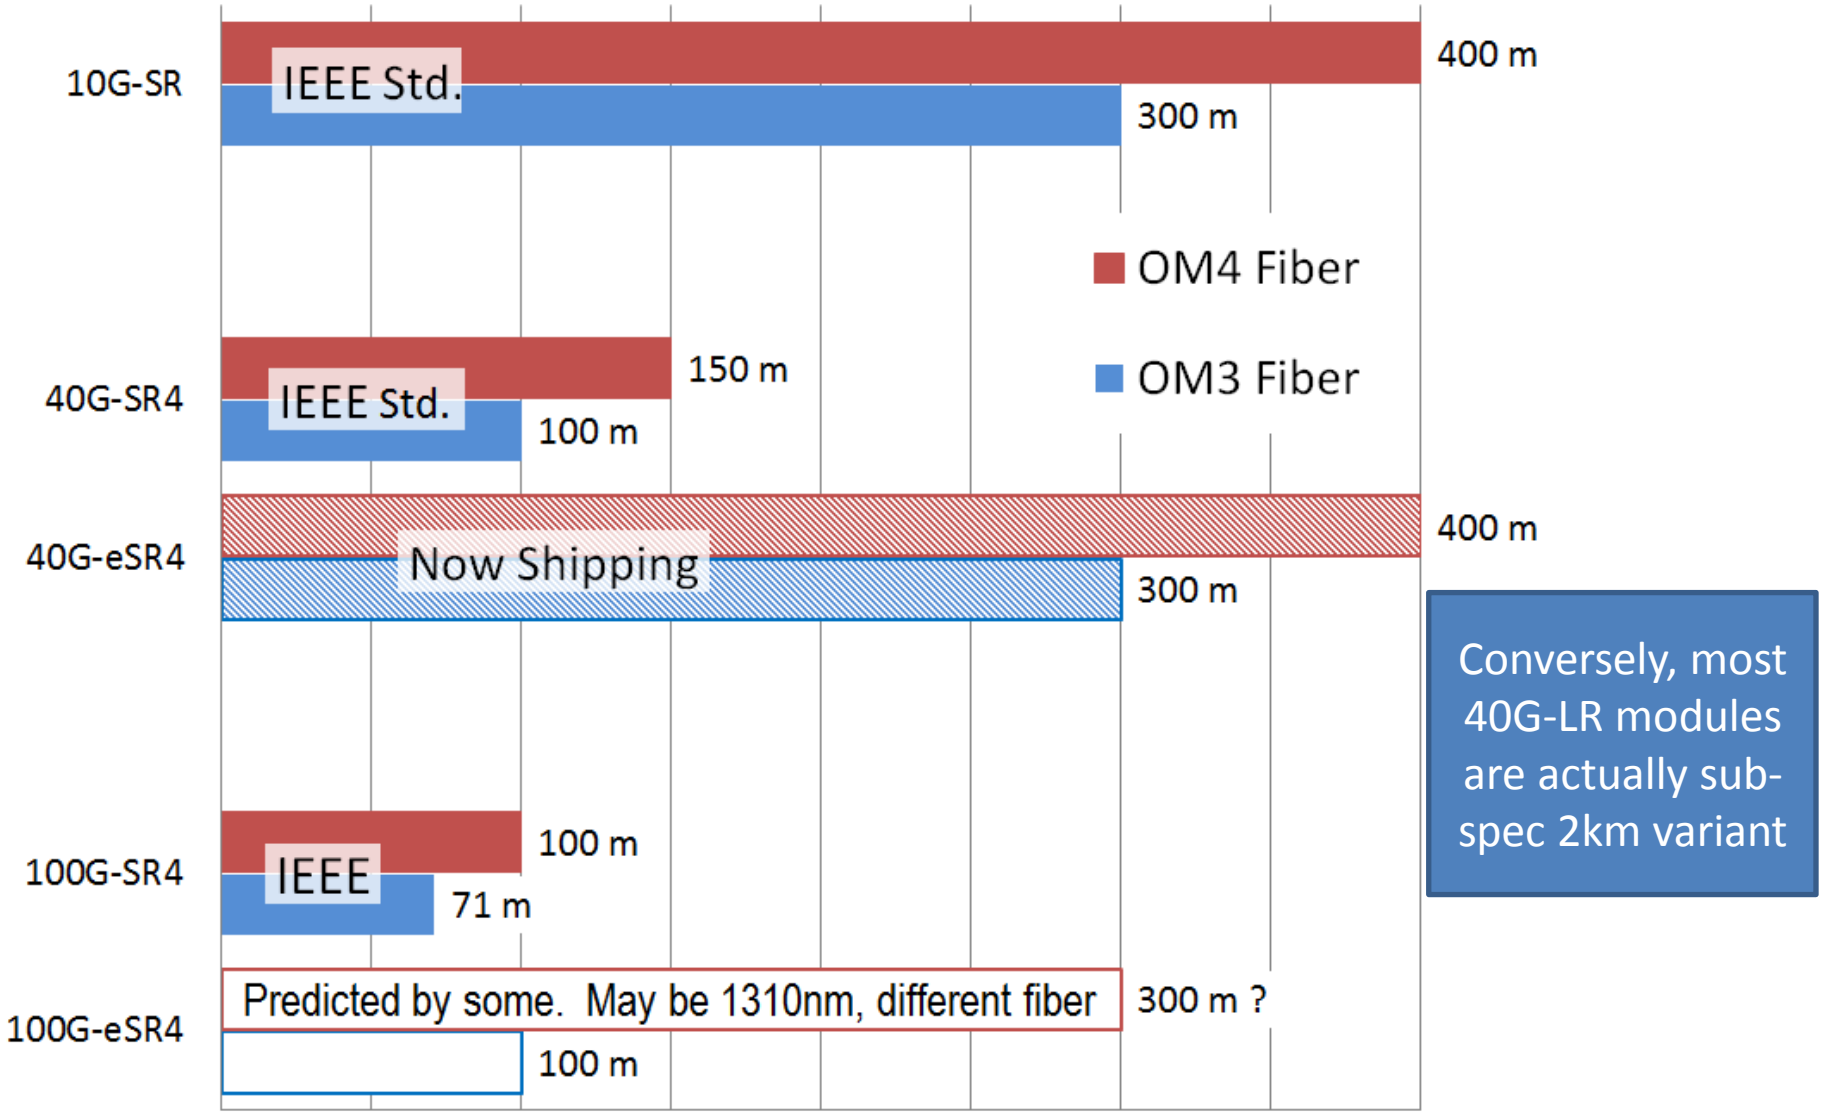

Conversely, most 40G-LR modules are actually sub-spec 2km variant

40G Ethernet Lessons

- 40GigE was a convenient solution
  - Was a practical jump from 10GigE.
    Lane speed stayed at 10Gbps.
  - Very successful in supporting high-speed uplinks and 4x10G leaf & spine installations
- 40GigE is also a transitional solution
  - Cannot multiply current PMDs for next higher speeds
- Are we creating more transitional solutions with 4x25G solutions [potentially 16x25G or 8x50 solutions?]
  - Cannot break out for leaf & spine as with 40G-SR
  - 40G serial or 100G serial might change this picture

400G Call for Interest Slide

Data Center Architectures

100G Core-Aggregation Links? Yes, Soon

> 100GbE need

Fully 100G Fabric? Not before 2020

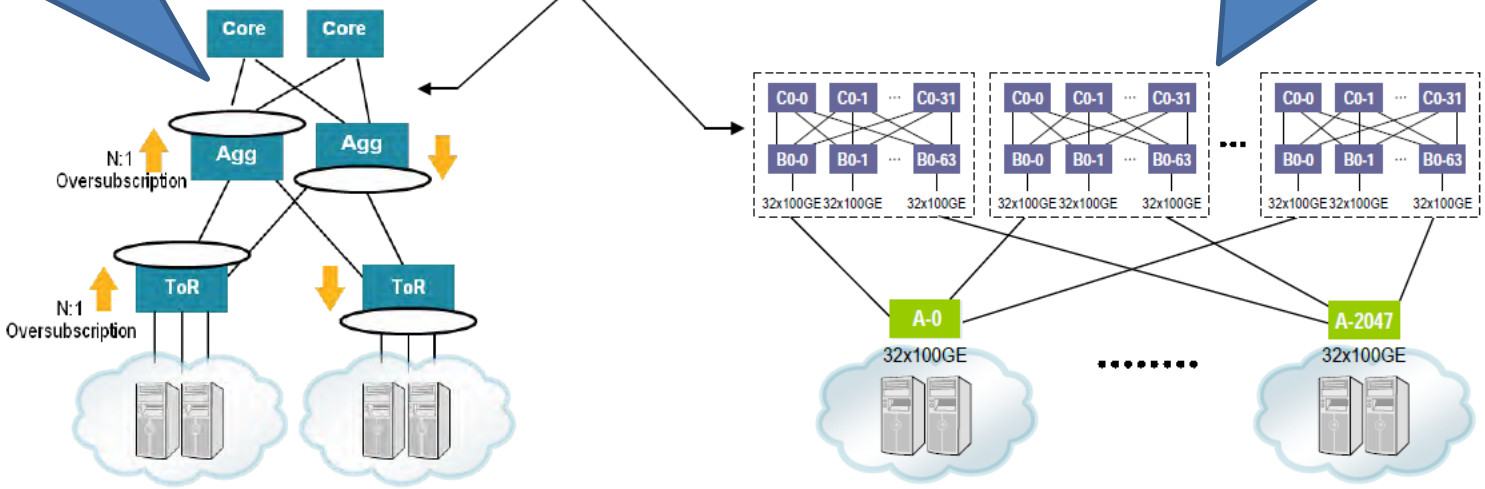

Hierarchical Fat Tree architecture

Non-blocking architecture

Flatter Architectures Driving 4x10G Consumption; Will delay 100GigE Consumption

Optical Line (modulation) Rates (by fiber or wavelength)

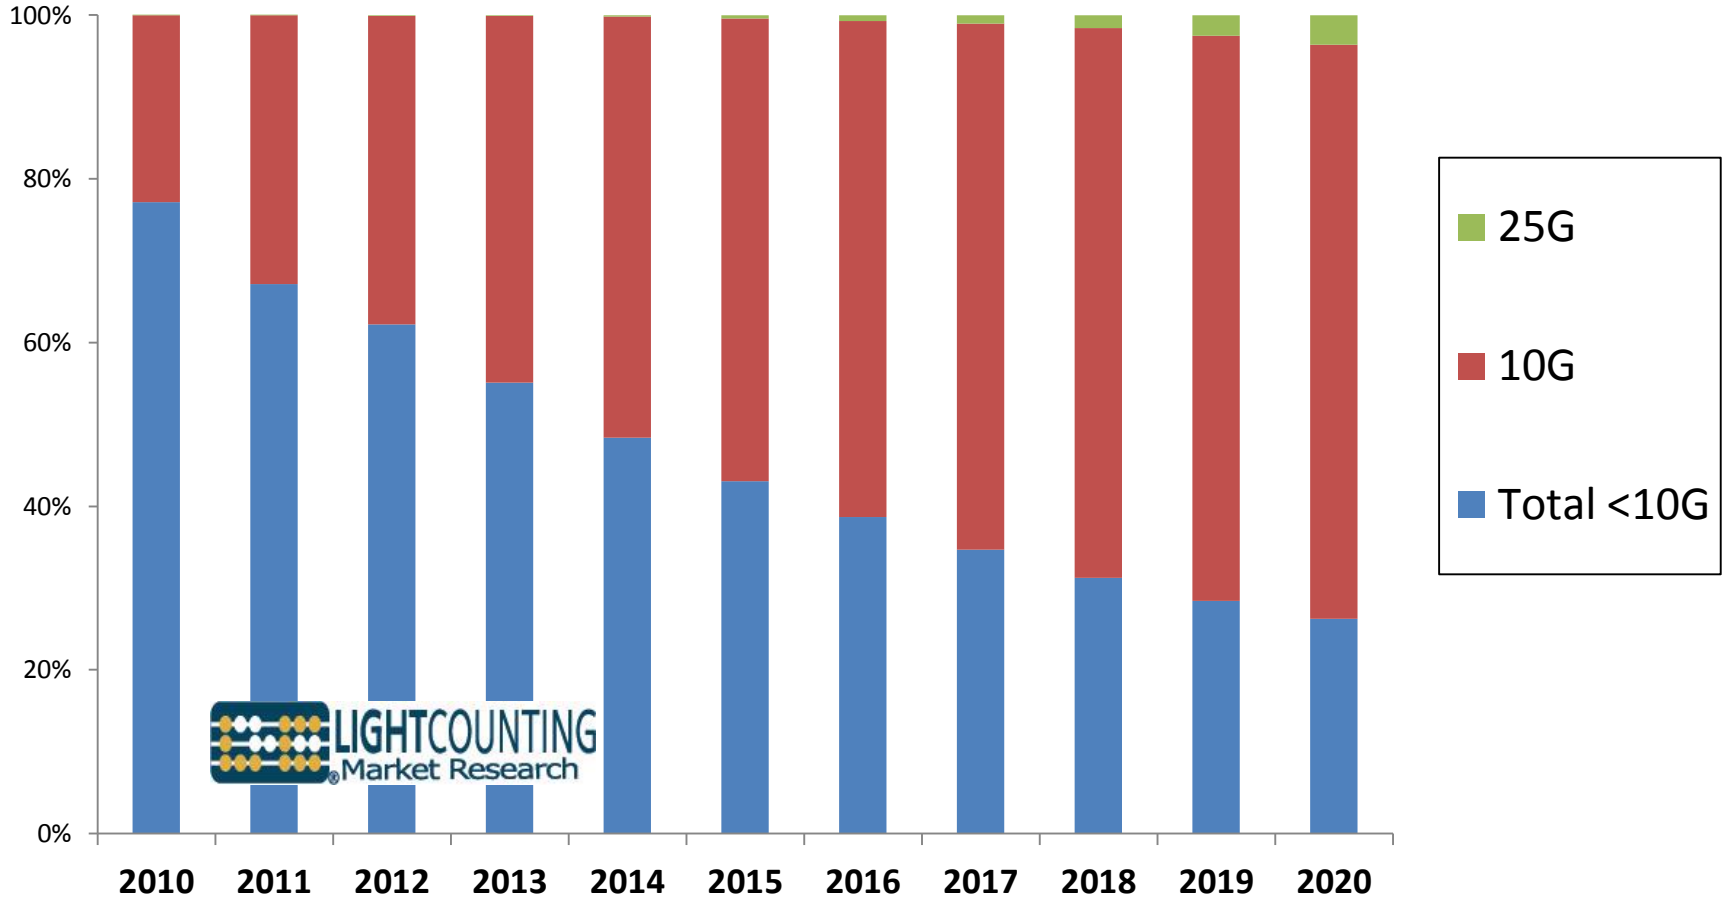

State of the Industry

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

- 10G is the bread and butter of the industry and will remain so beyond 2017
- 40G module sales are soaring, but majority use is for density rather than native 40GigE
- 100G is connecting data centers, but unit volume won't be significant inside data centers until after 2017
  - Not straightforward to transition fabrics from 10G to 40G or 100G without reducing cluster size.
- It takes volume to get competitive $/Gbps
  - Dividing each market by PMDs and form factors extends time-to-volume
- The lessons of 10G and 40G suggest we should endeavor to anticipate consequences of PMD choices, including distance limits. Be cautious of quick-to-market solutions.
  - e.g. Could we have anticipated 300m nx10G solutions and avoided non-standard parts?
  - Will 2x16 fiber/connector infrastructure find acceptance?
  - Time-to-volume economics would suggest minimizing number of PMDs along with a long-term perspective.
- The optical transceiver business is financially challenging.