Direct Attach Copper Cable Broad Market Potential and Economic Feasibility

DAC is not Dead

(Updates)

Tom Palkert (Molex),

3-6-18

Supporters

Outline

- Relative market size of Direct Attach Copper Cables (DAC)
- Challenges with current network use-case
- Network deployment options for 100G serial
- Cable length distribution
- Options for 100G serial switch design
- Relative cost analysis

Relative market size of DAC

Typically 10x larger than equivalent optical market

Interconnection Volume

- Four sections per colo & multiple colos (≥ 4) per data center
- Volumes below are per section (except DCR to Metro)

A End	Z End	Volume	Reach (max)	Medium	Cost Sensitivity	Market Space
Server ǂ	TOR	10k – 100k	3 m	Copper	Extreme	LAN
TOR	LEAF	1k – 10k	20 m	Fiber (AOC)	High	
LEAF	SPINE	1k – 10k	400 m	SMF	High	
SPINE	DCR	100 – 1000	1,000 m	SMF	Medium	Campus
DCR	Metro	100 – 300	10 - 80 km	SMF	Low	WAN

ǂ Server-TOR links may be served by breakout cables

Source: Brad Booth, Microsoft http://www.ieee802.org/3/400GSG/public/13_11/booth_400_01a_1113.pdf

Current deployment use case has 3m DAC for TOR switch and 30-300m AOC for Aggregation switch

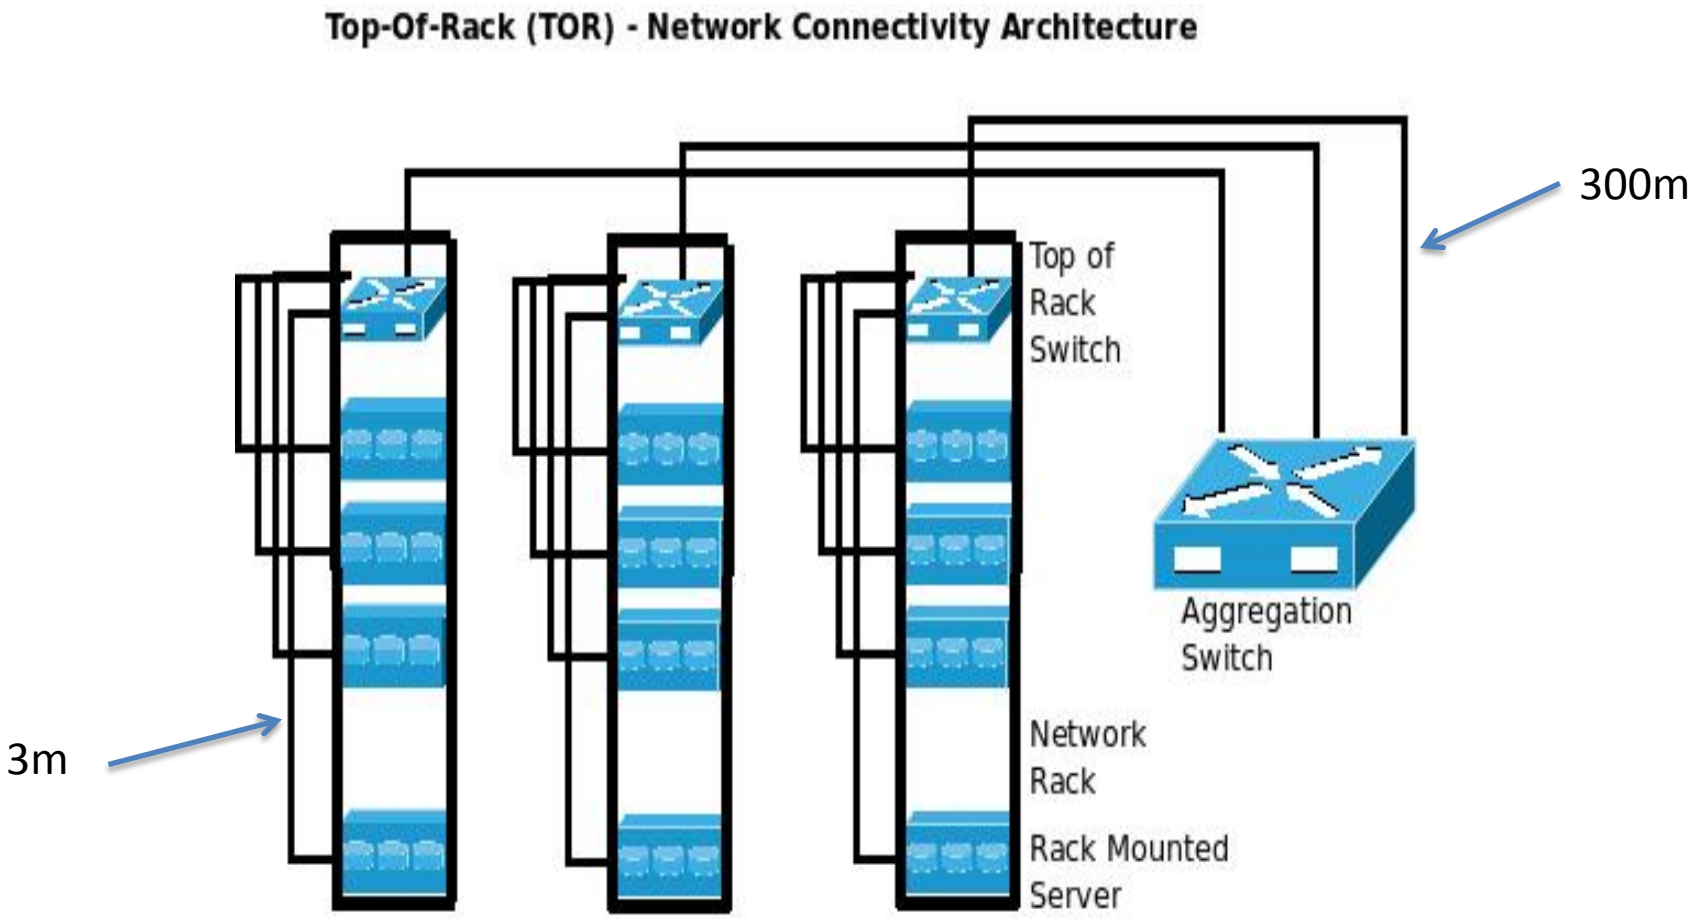

TOR switches can have 32-48 ports of SFP (Single lane module) or 32-36 ports of QSFP (Quad ports)

Improvement option 1: Move Switch to center of rack

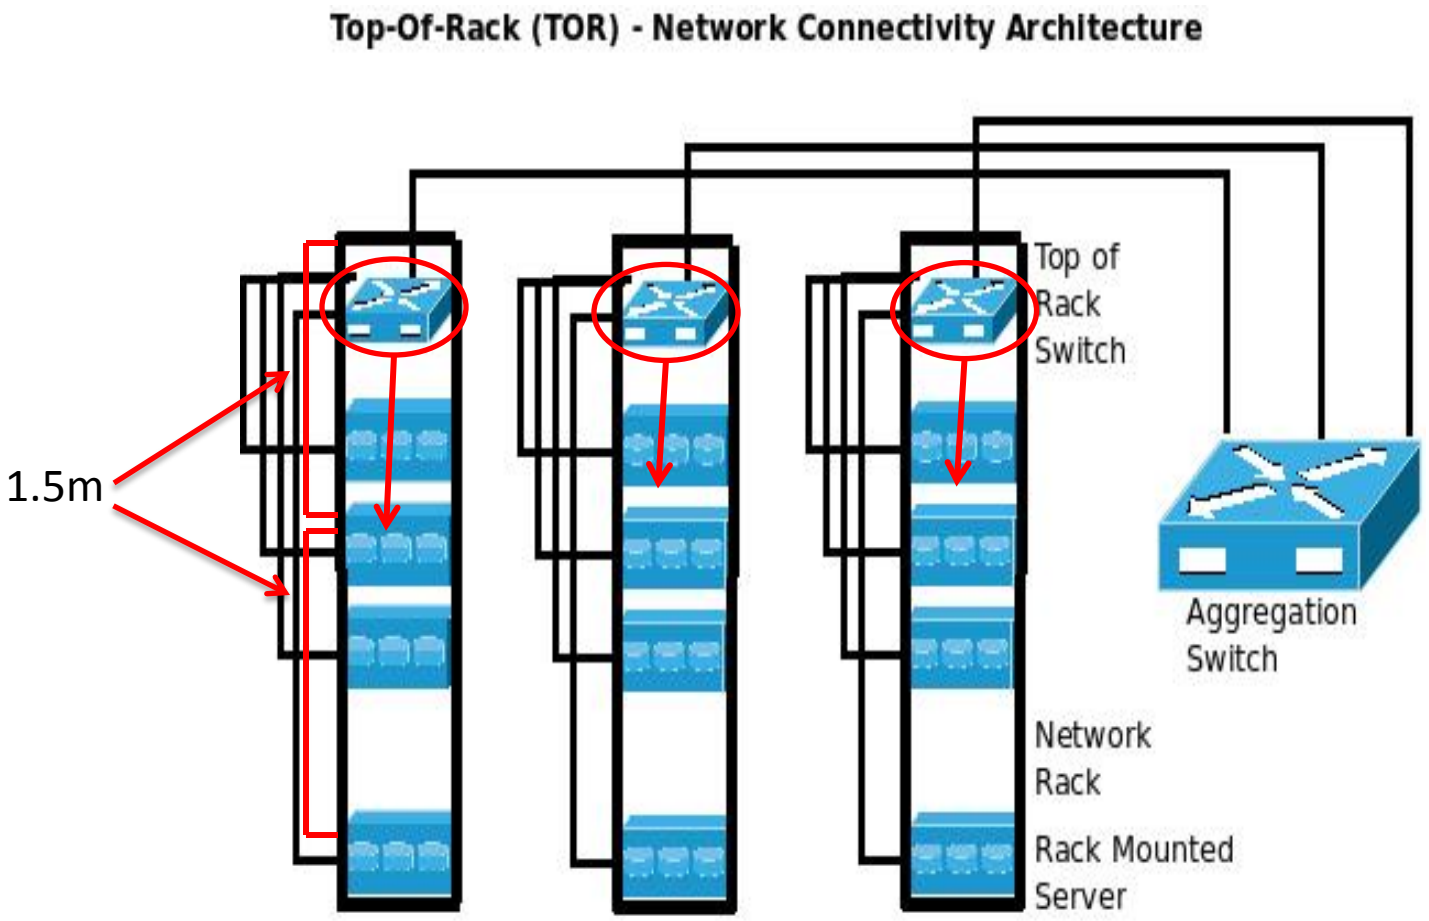

See Joel Georgen Presentation for details of cable length requirements

Cable Length Distribution

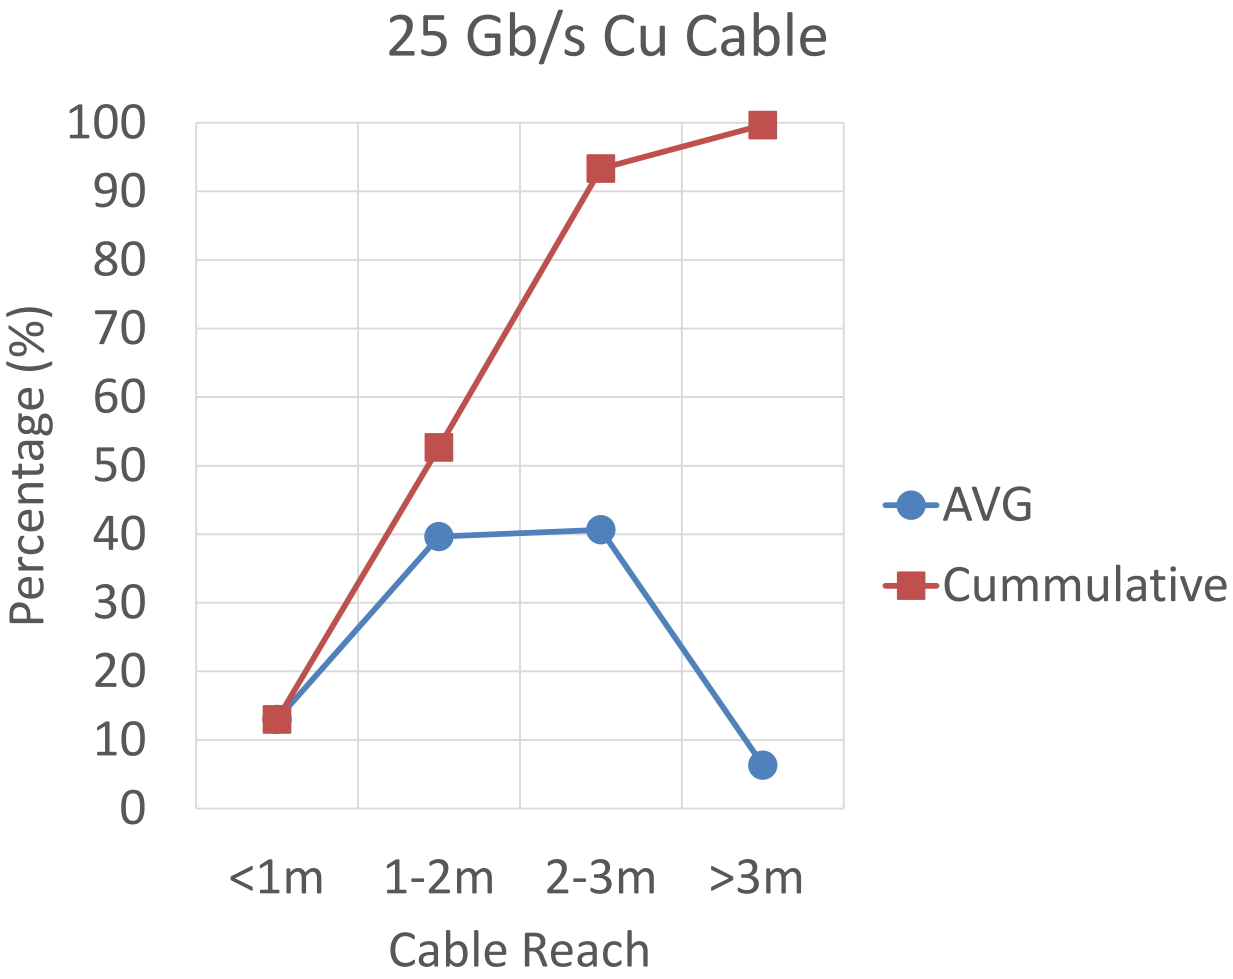

- Cable taken from 3 suppliers.
- Distributions of reaches averaged to anonymize data

Improvement option 2: Low Loss PCB

Low loss switch channels can be achieved with ‘advanced’ PCB materials

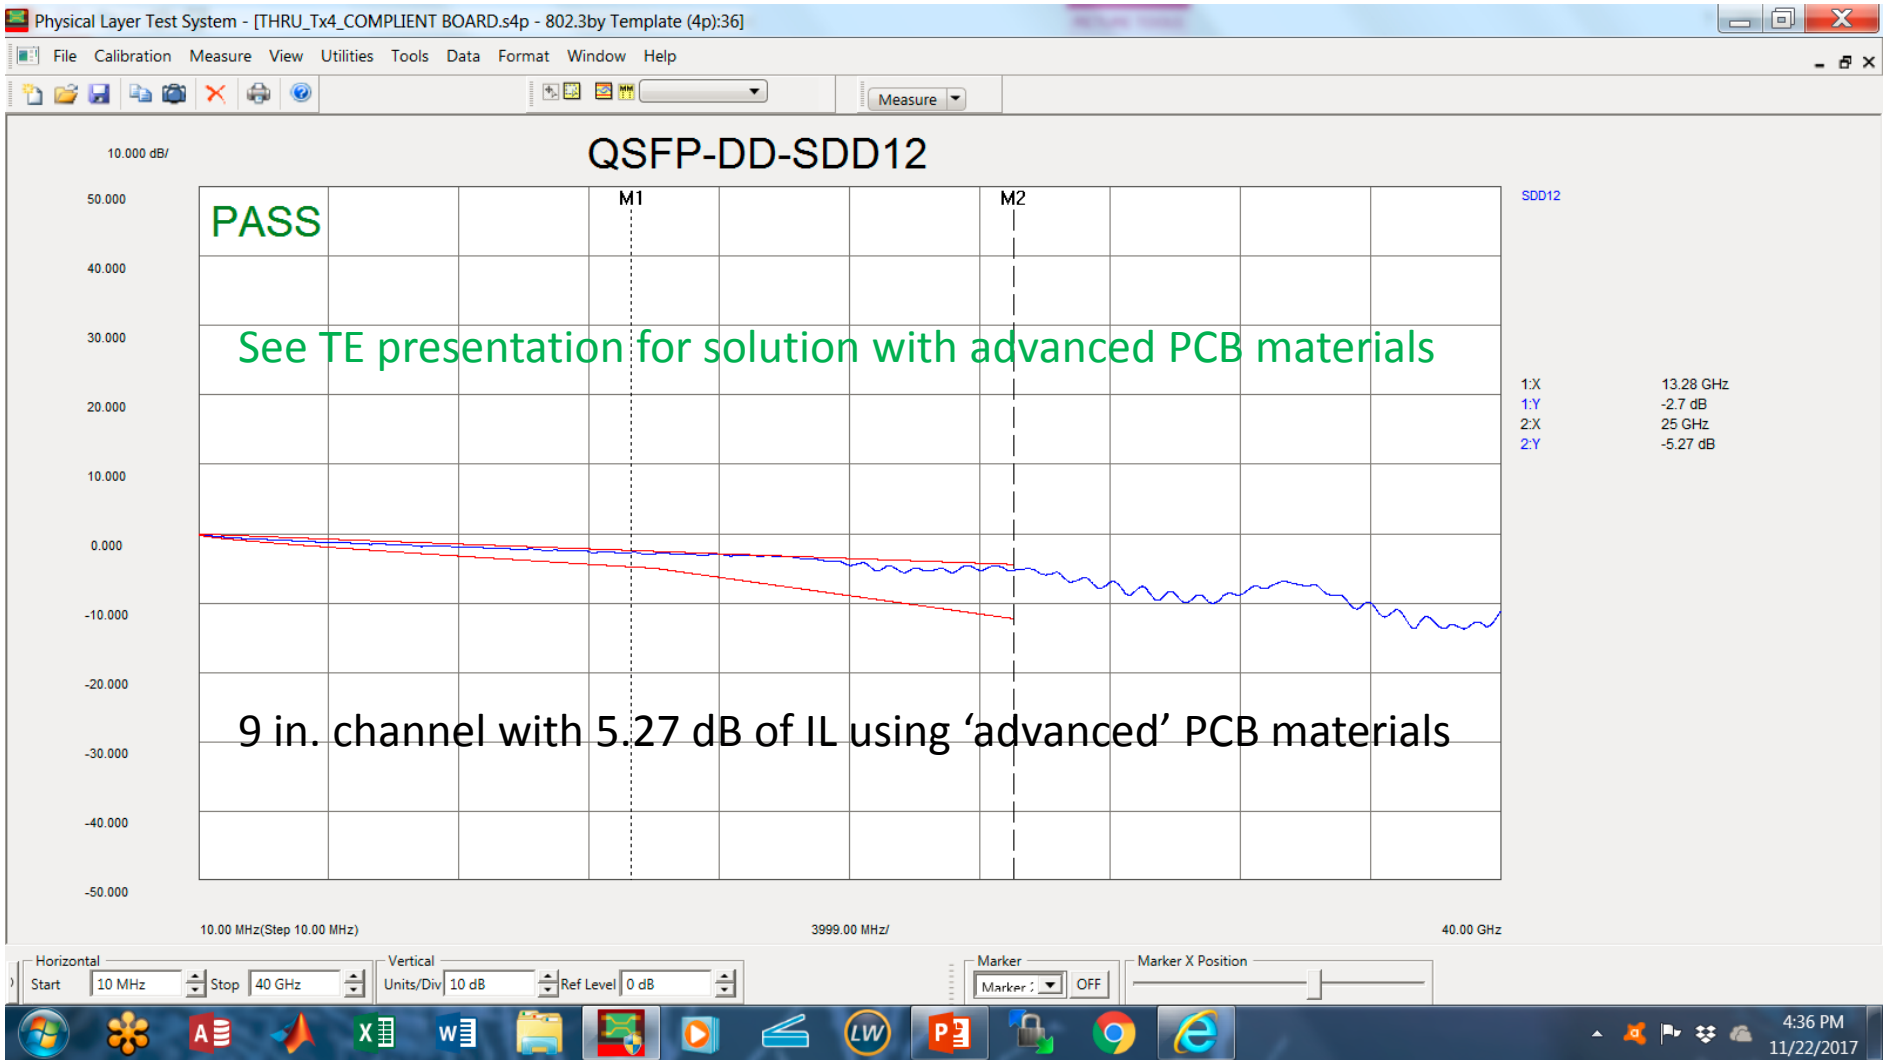

See: diminico_100GEL_01_0118

Internal cables replace PCB traces

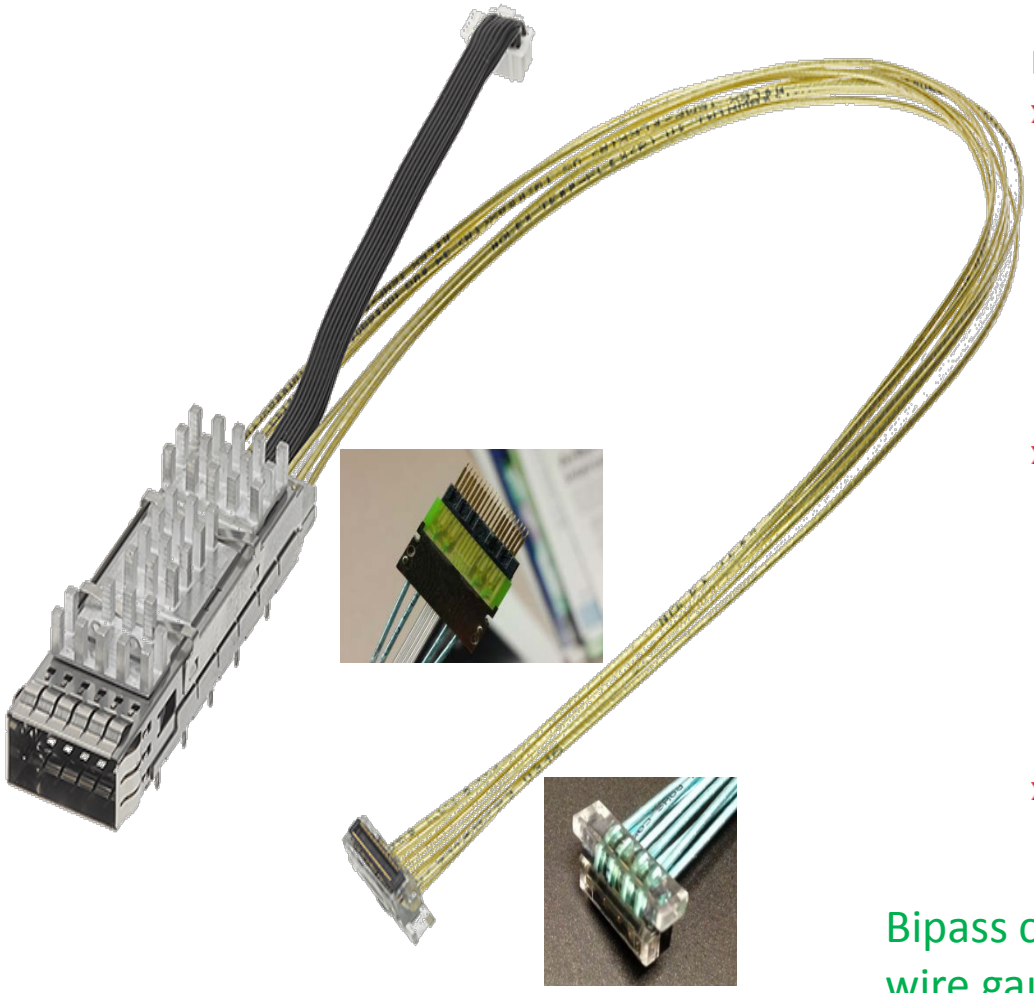

Key Benefits:

- **Lower System-Level Costs:**
  - **Eliminate the need for costly PCB materials (Nelco, Megtron, Tachyon, etc)**
  - **Eliminate the use of additional DFE or retimer chips to drive long traces**
  - **Improved thermal performance with 1x1 cages (air-cooled)**
- **Architectural Flexibility:**
  - **Freedom to locate ASIC anywhere (eg. further from backplane)**
  - **Enable lower power ASIC**
  - **Extended reach from ASIC to I/O**
  - **Enable longer external copper I/O cables**
  - **Cool the ports and the ASIC better**
- **New ways to handle power integrity to large ASICs**

Bipass cables can be made with multiple wire gauges to optimize size vs loss
(See Molex presentation)

Improvement option 4: Define asymmetric switch to server connections

Define end to end budgets that take advantage of short NIC traces

Asymmetric link would require link negotiation similar to 25G-CR-S copper

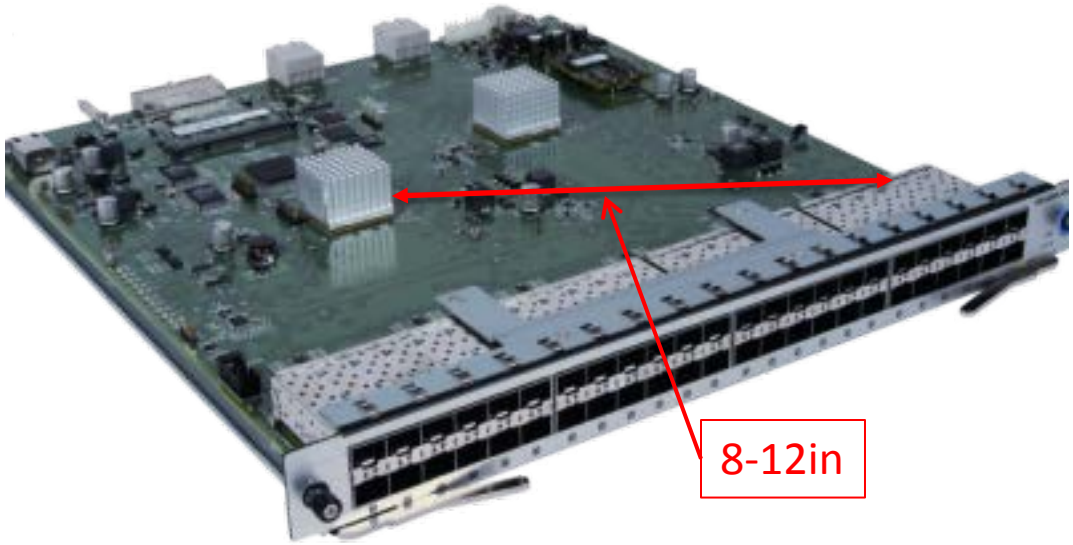

Switch ports require long host to module traces

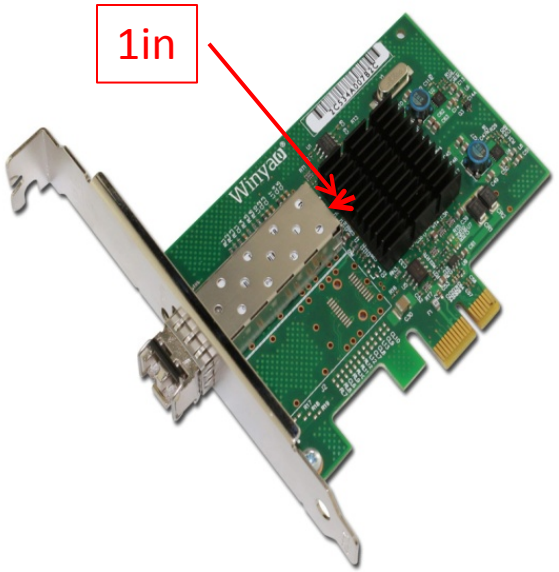

Server ports have very short host to module traces

Improvement option 6: Add retimers

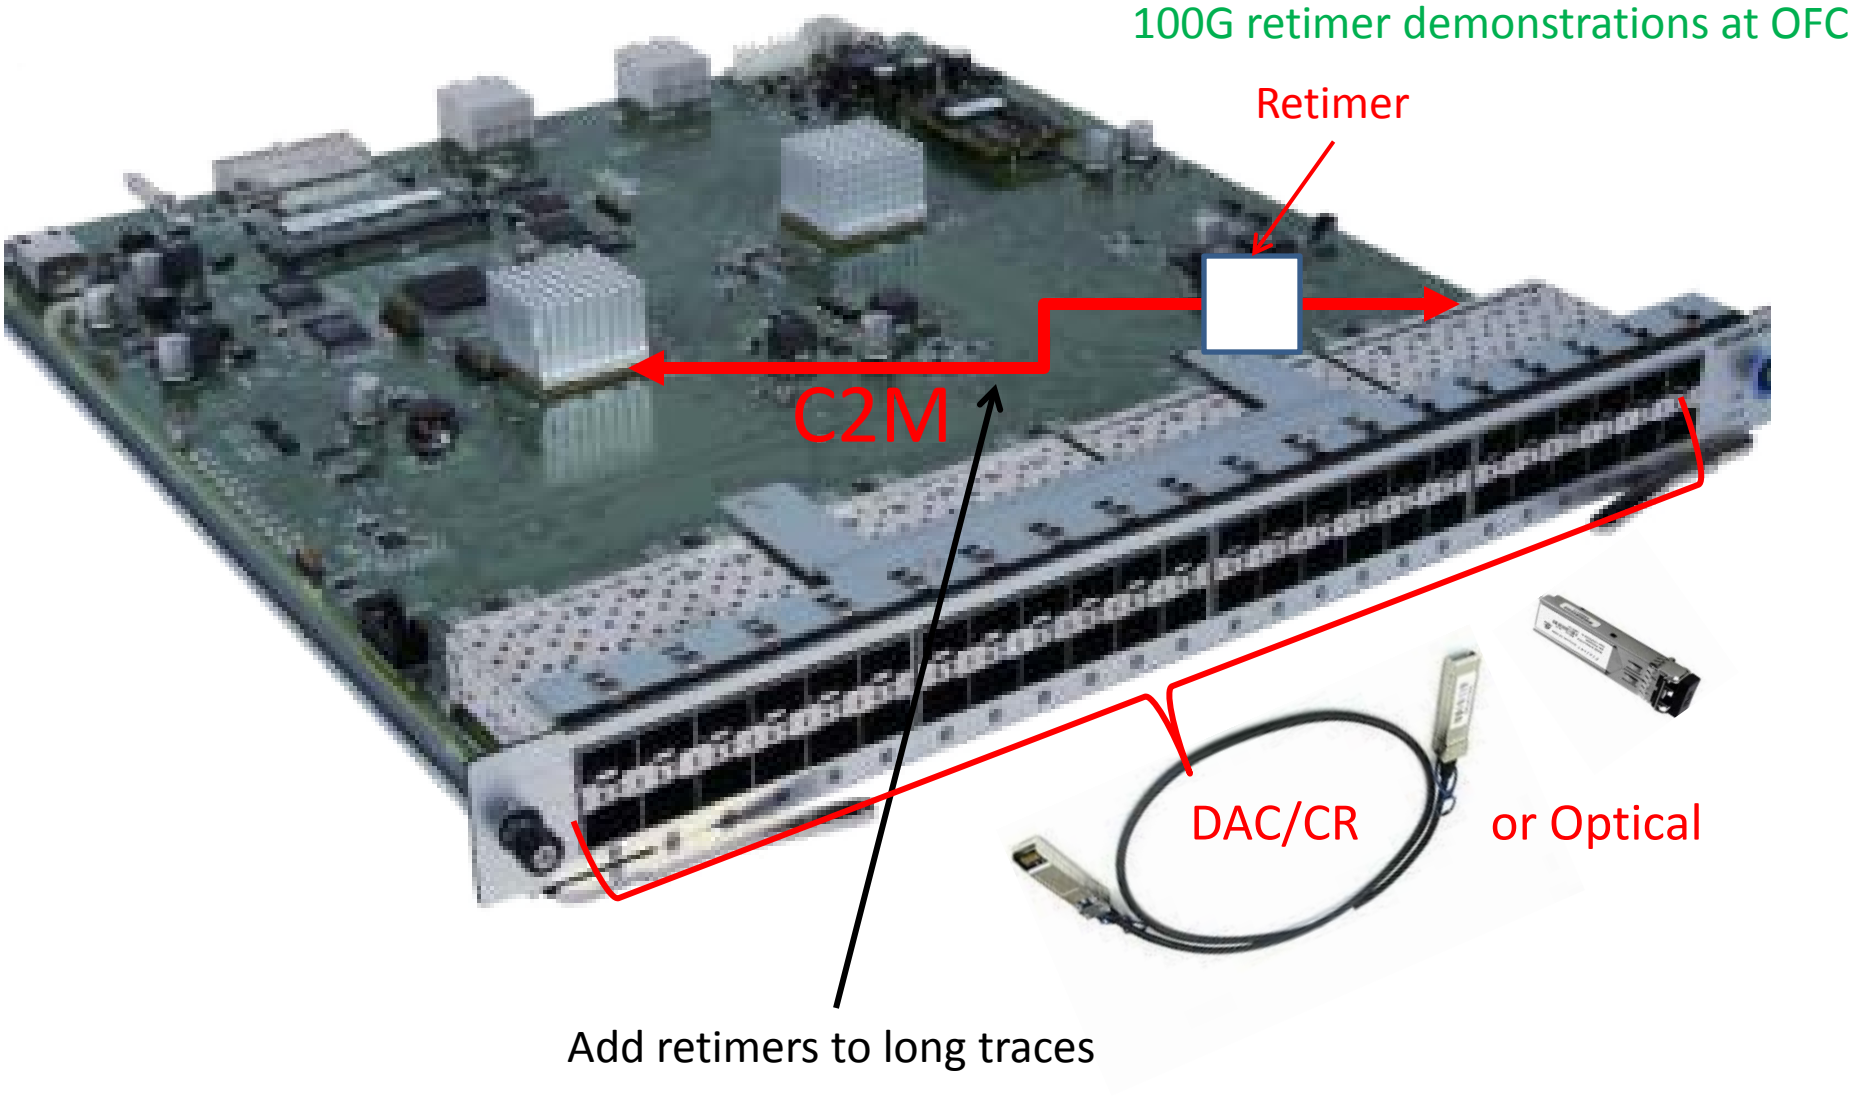

Economic Feasibility

- Reviewed relative costs of 2 network designs
  - Structured cabling
  - Direct attach AOC
- Used both high volume pricing and retail pricing

Economic Feasibility baseline

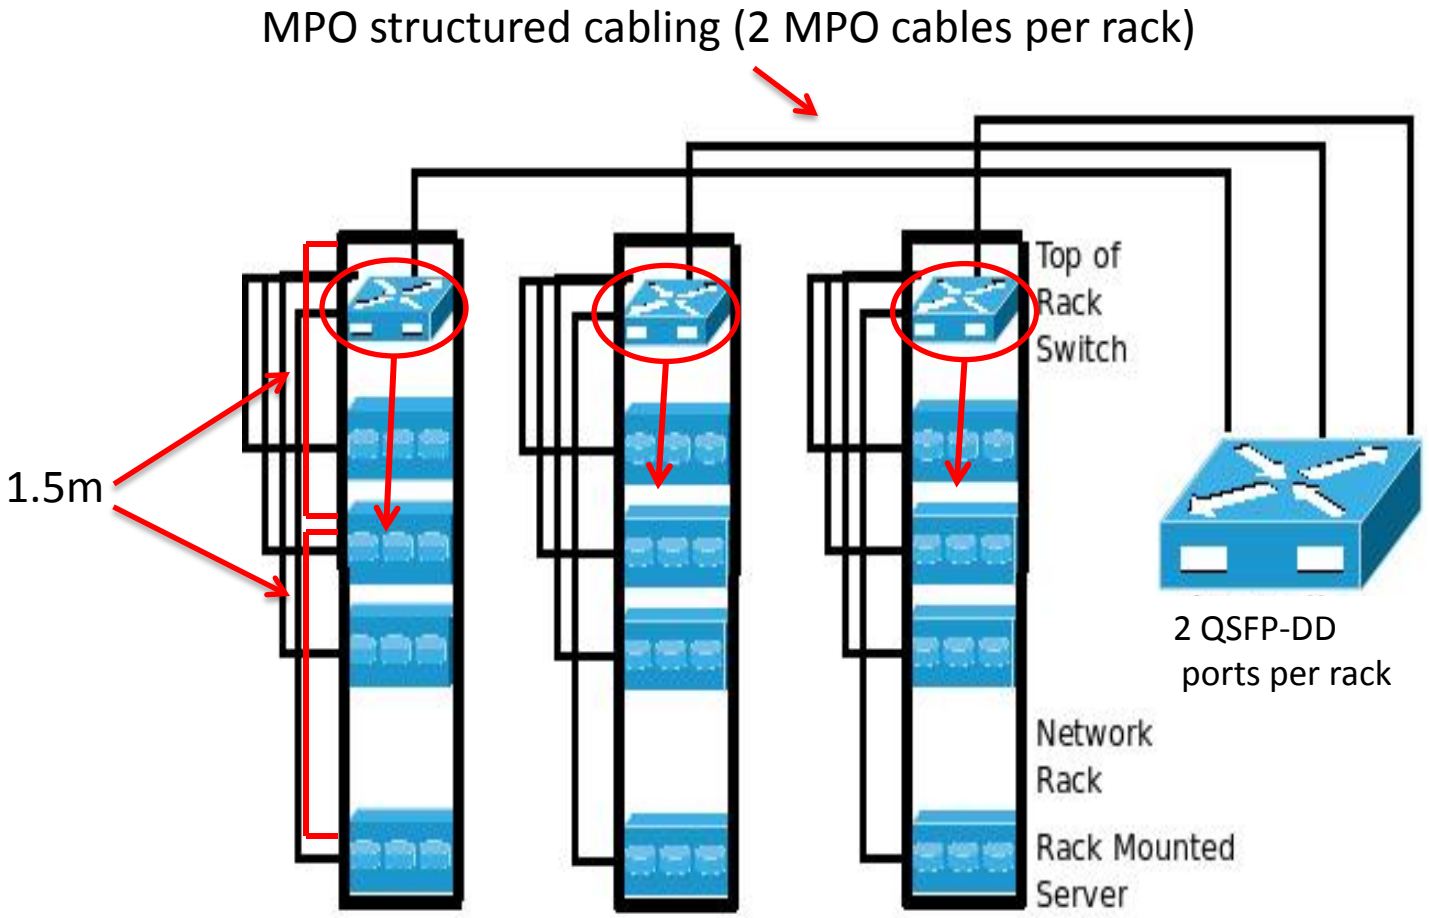

Relative cost analysis Option 1: Replace TOR switch with server to EOR switch optical modules using structured cabling

Relative cost:
7-11x

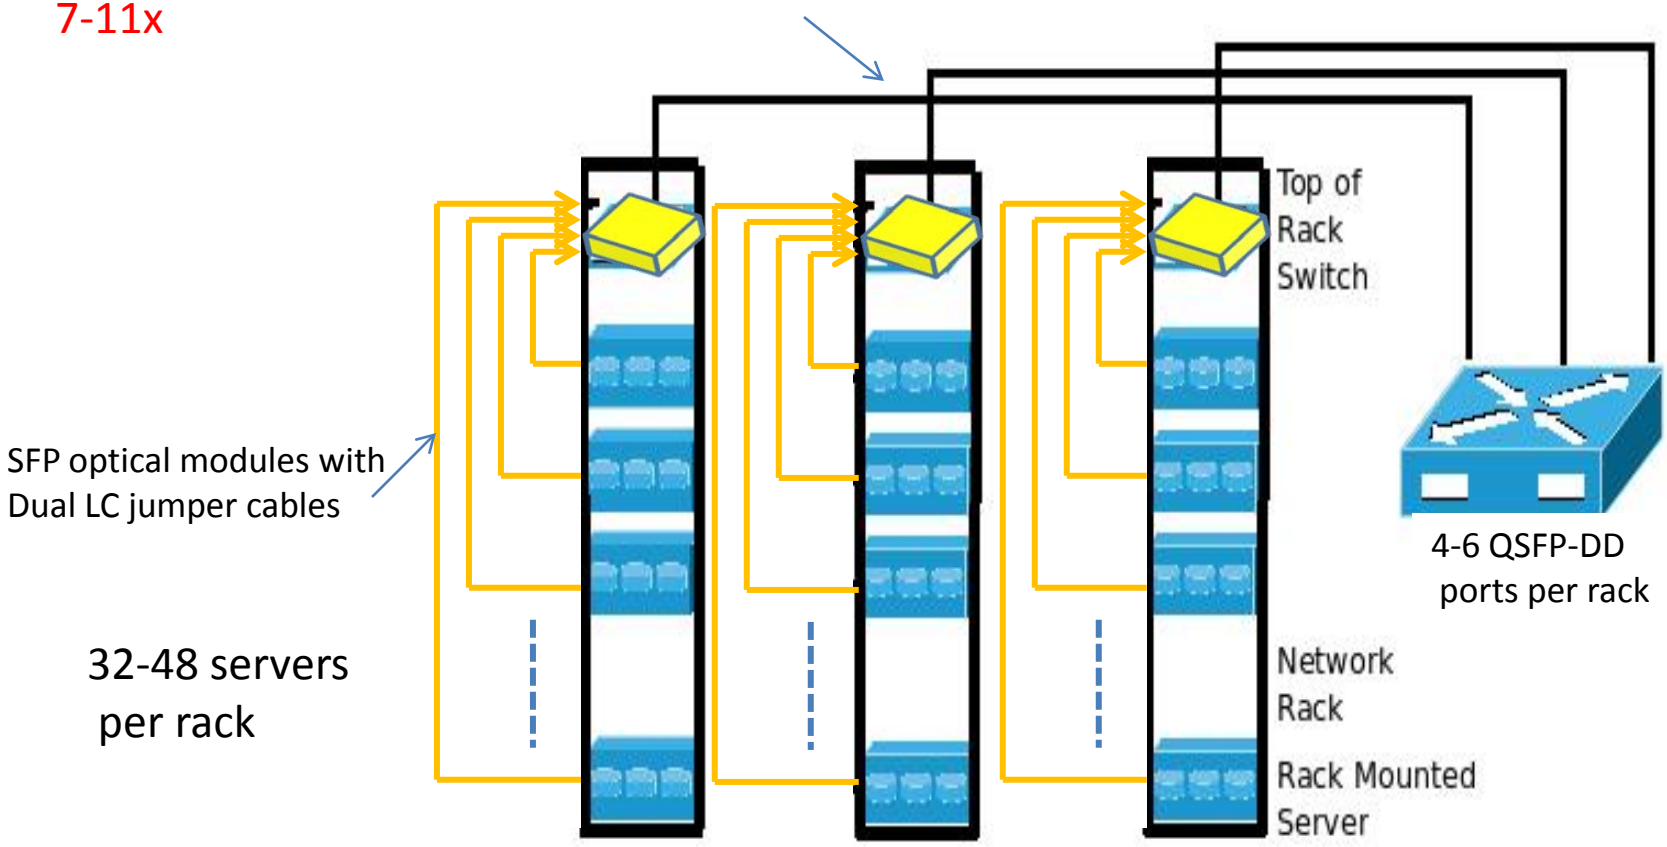

Relative cost analysis Option 2: Replace TOR switch with server to EOR switch optical modules using Active Optical Cables

Relative cost:
5-8x

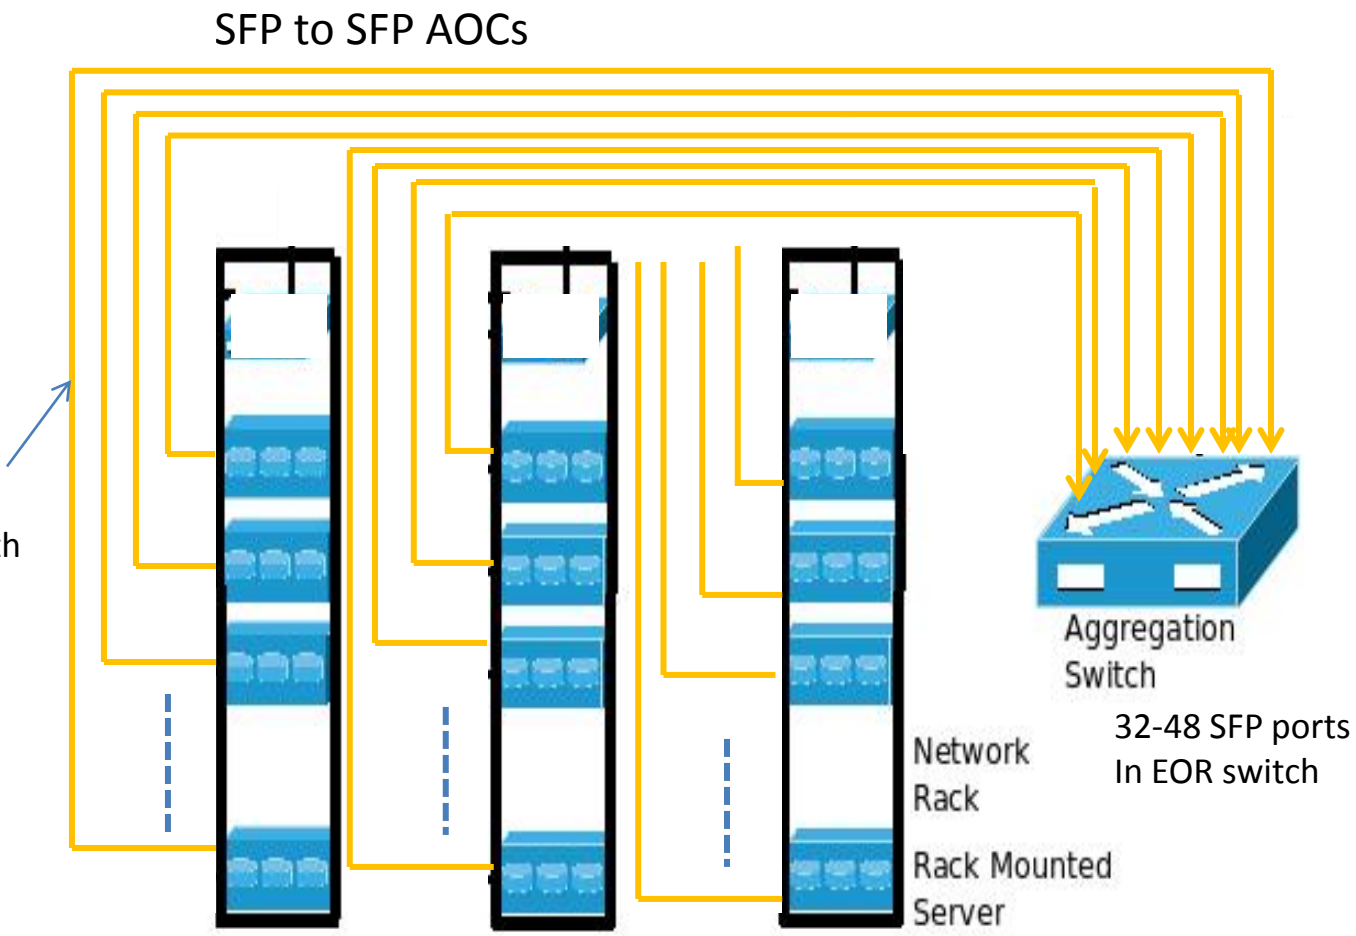

Summary:

- DAC is not dead
  - Copper variants are prevalent and economically critical in the market today because of the cost/performance tradeoff they provide
    - Most cost effective intra rack connection
    - 10x larger market size vs optics
  - Multiple options exist to enable 100G serial DAC to provide the cost benefits demanded by end users
  - 2m DAC cable is the required length