Reach Objective Considerations

Brad Booth, Dell

November 2012

San Antonio, TX

Supporters

- Dan Dove, AppliedMicro
- Robert Winter, Dell
- Alan Flatman, LAN Technologies
- John D'Ambrosia, Dell
- Bruce Tolley, Solarflare
- Dave Chalupsky, Intel
- Scott Kipp, Brocade
- Kent Lusted, Intel
- Mike Bennett, LBL
- Valerie Maguire, Siemon
- Pravin Patel, IBM

Outline

- End User Feedback
- Alternative Topologies
- Relative Cost Analysis
- Power Considerations
- Creating a Reach Objective
- Recommendations
- Questions

End User Cabling Feedback

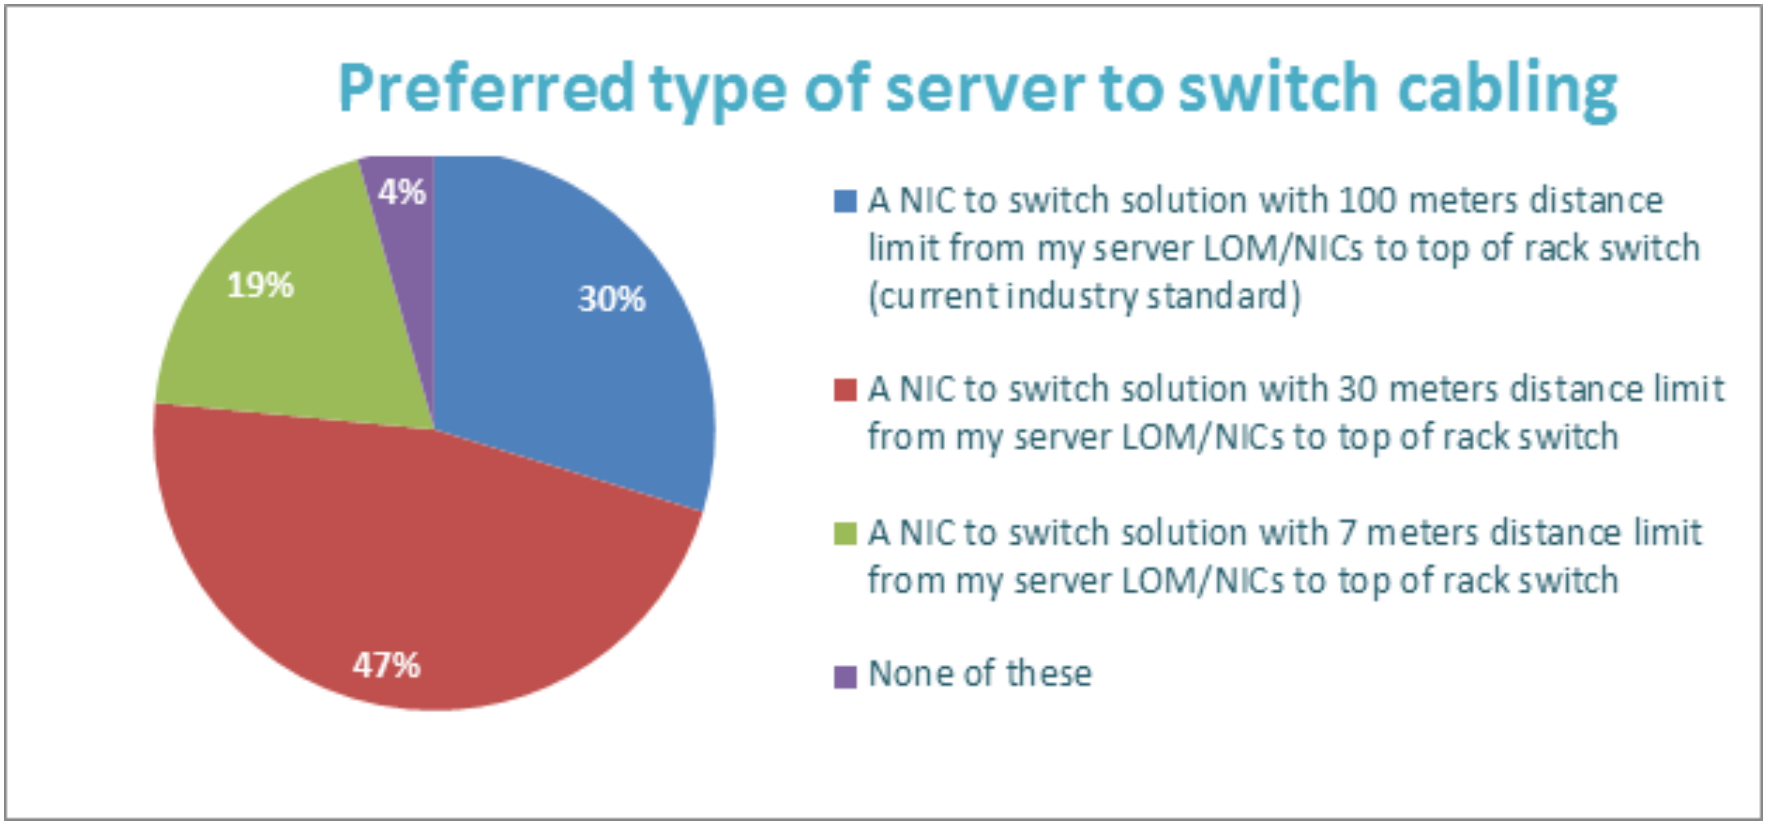

† Source: Dell Networking 2011

‡ "Top of rack switch" is the first switch within the stated reach

10GBASE-T End User Feedback

- Primarily a ToR to server interconnect today
- Reach requirement is 10-15m for a ToR† topology
- EoR is a valid use case for copper
- Most EoR today is over MMF
- Willing to look at alternative topologies if power or cost reduced‡

† ToR is a general term used for the first switch, not necessarily located in the same rack as the server.
‡ Versus traditional Enterprise data center centralized switching.

EoR

- From Nordin presentation (nordin_01a_0912.pdf):

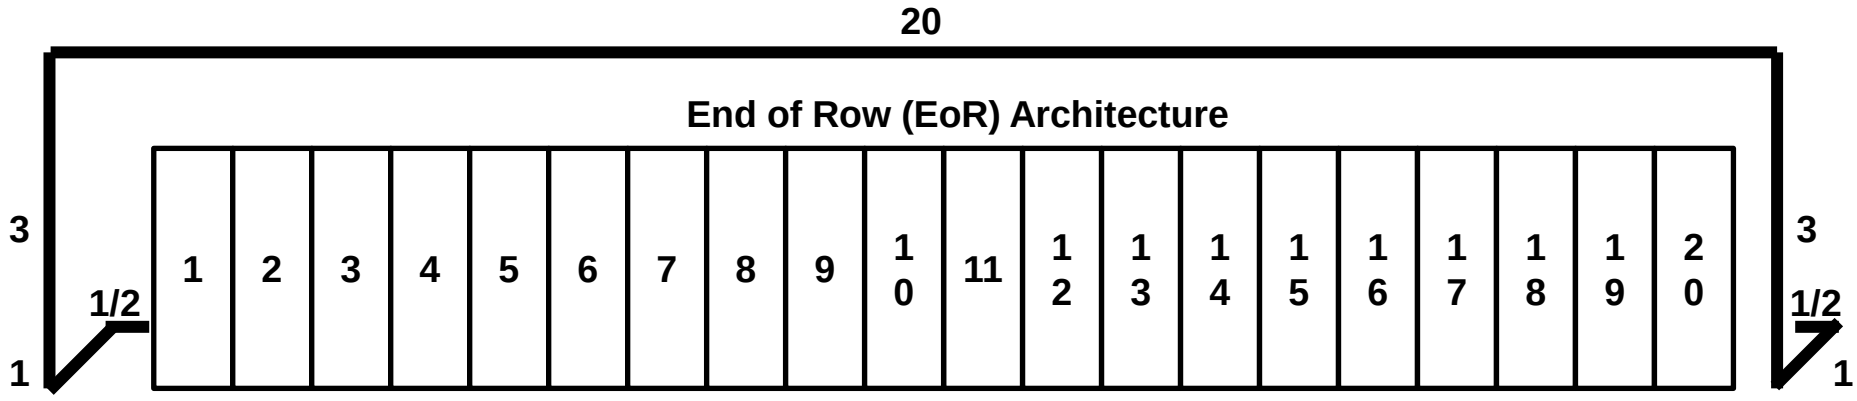

- What if we offered an alternative topology?

Middle-of-Row (MoR)

- Move the switches to be centralized in the row:

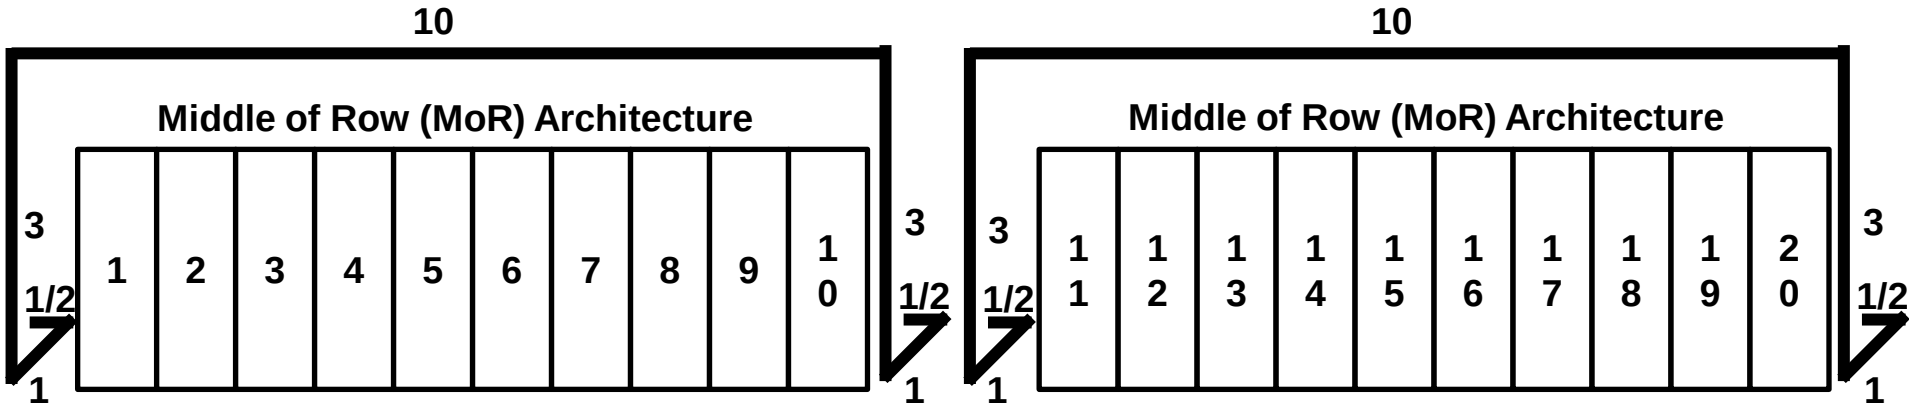

- Reduces the reach to < 20m
  - 12m reduction ≈ 50% reduction in PHY power*

* Per bliss_01a_0912.pdf

Relative Cost Analysis

- Relative cost information was gathered from multiple sources†
  - BSRIA
  - Cable and component vendors
- DC = direct connect
- 2C = two connector w/ two 2 m patch cords
- 4C = four connector w/ two 2 m patch cords
- Only based upon average cost of materials

† Thanks to Alan Flatman for providing guidance on the data collected.

Cable Types†

UTP

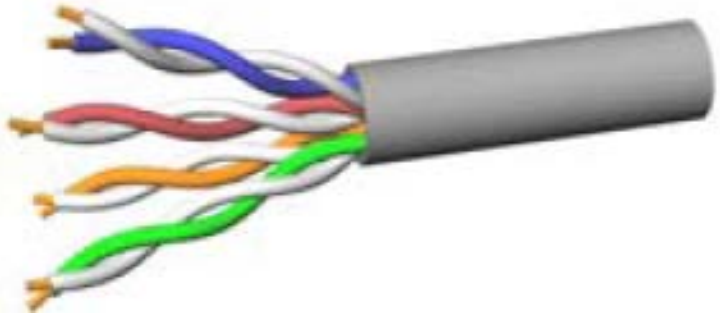

FTP

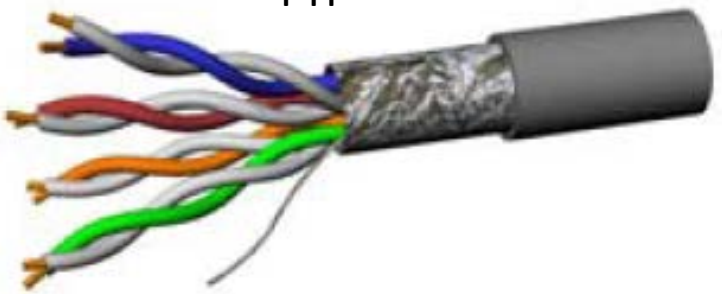

STP or PiMF

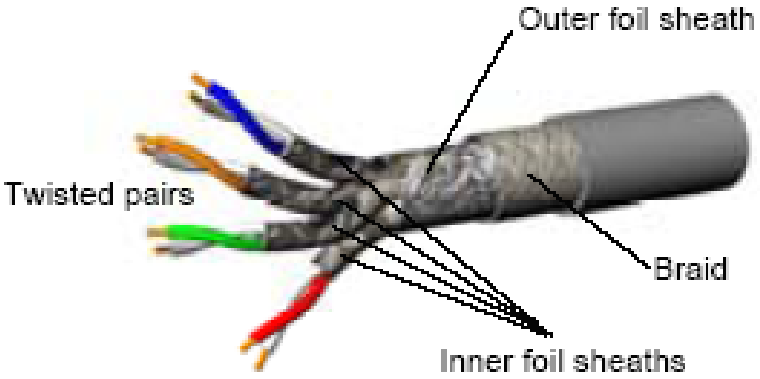

† Source: BSRIA

Plenum Relative Costs

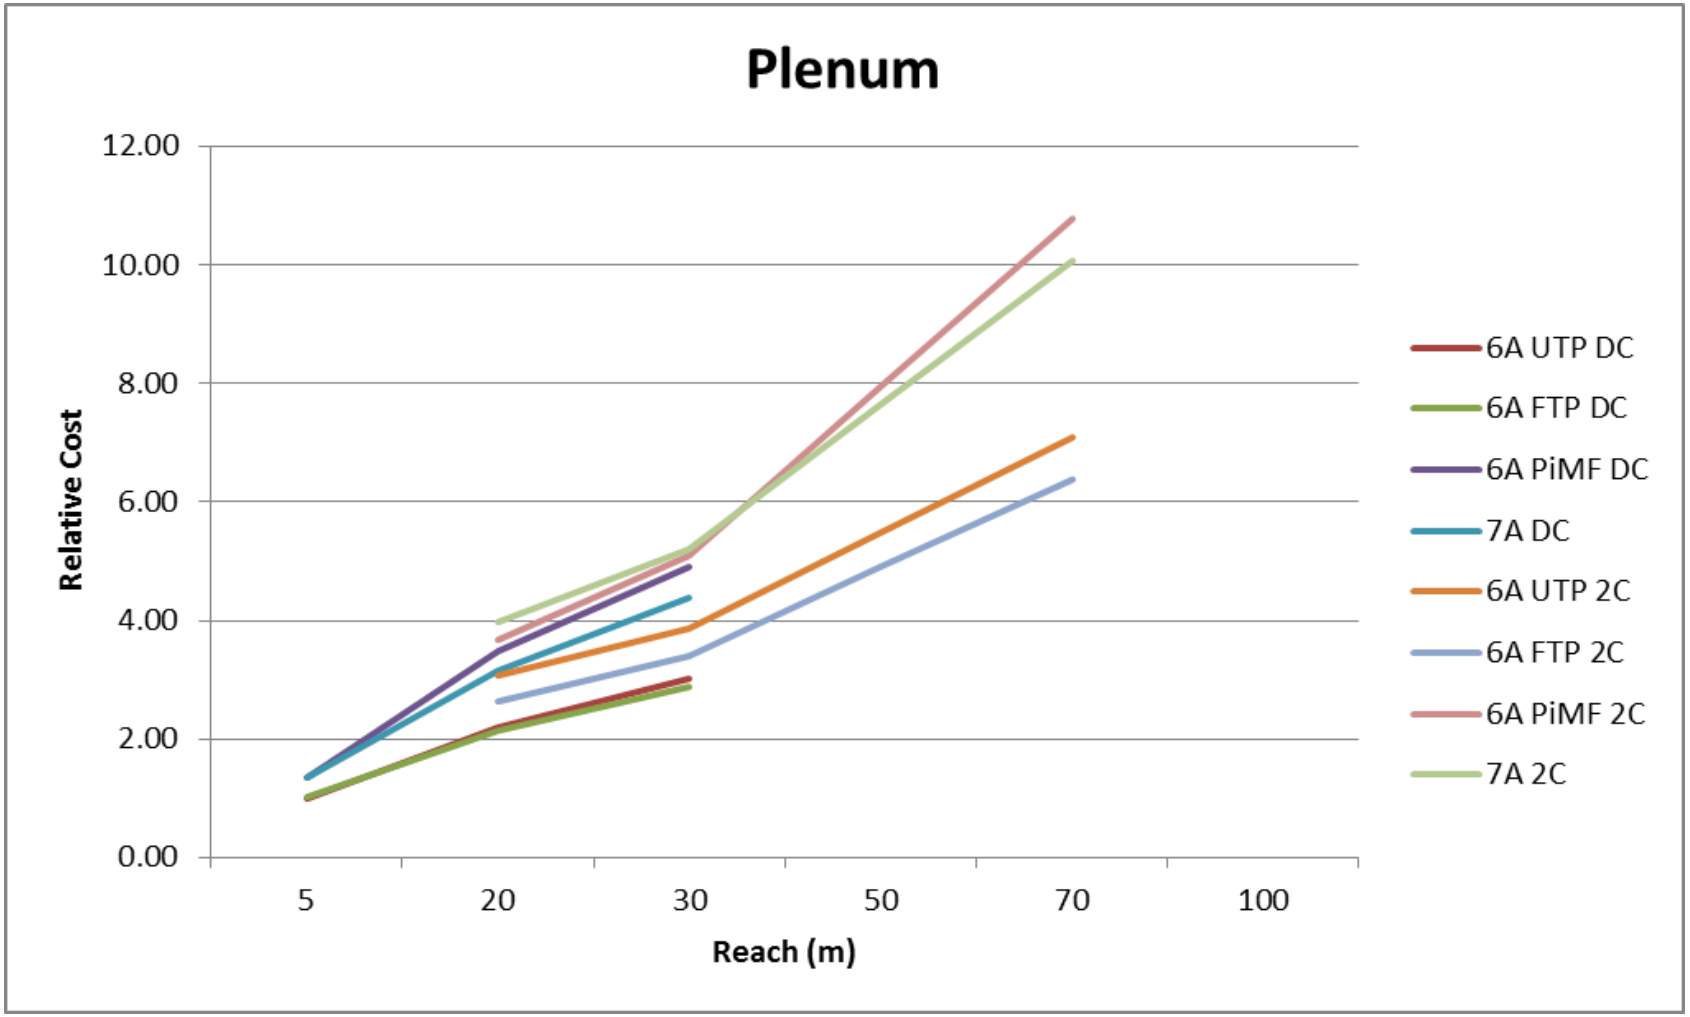

* Data is based upon relative cost analysis per slide 8 with Cat 6A UTP DC as the base cost.

Selected Plenum Data

- Used only 6A UTP and 7A PiMF up to 30 m

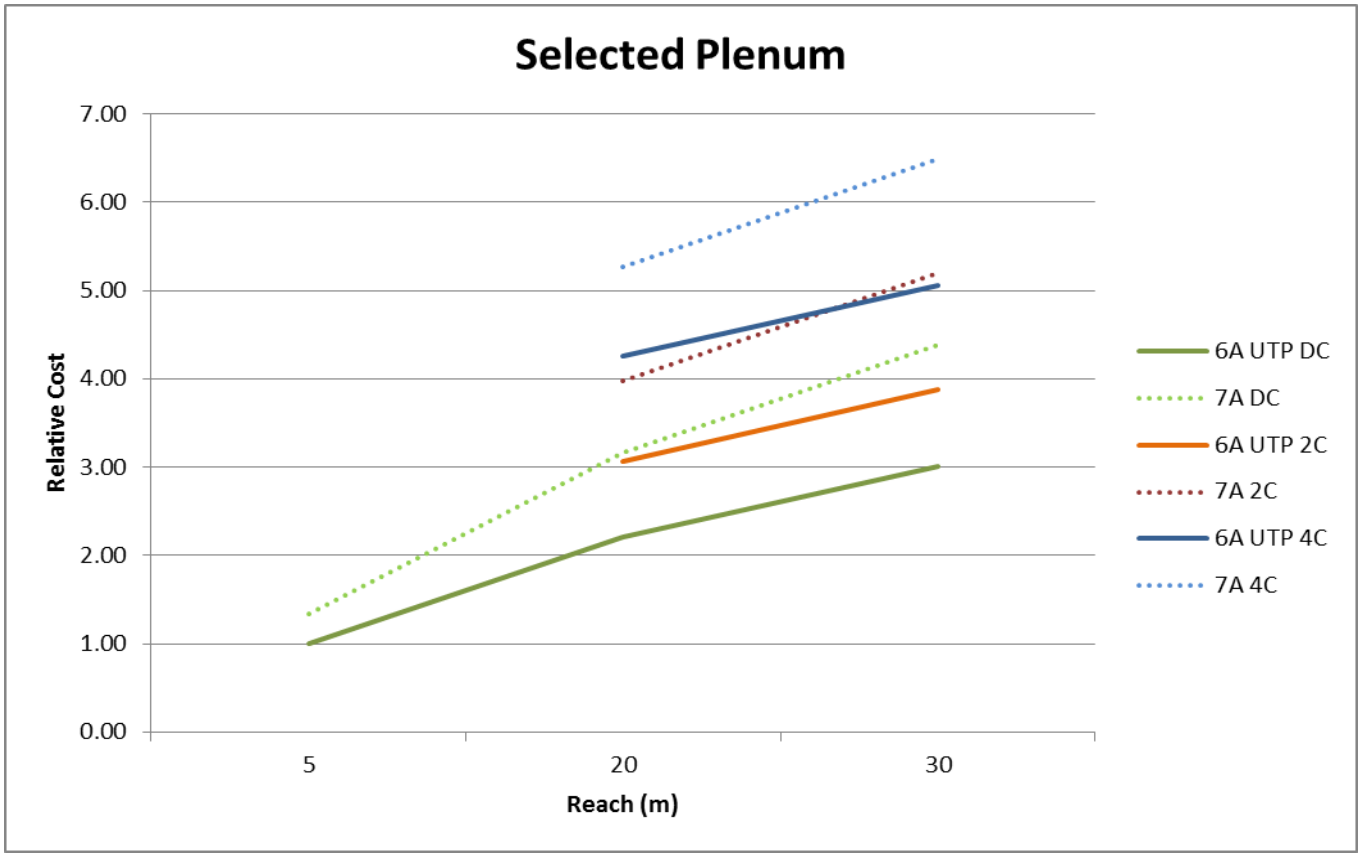

* Data is based upon relative cost analysis per slide 8 with Cat 6A UTP DC as the base cost.

LSZH Relative Costs

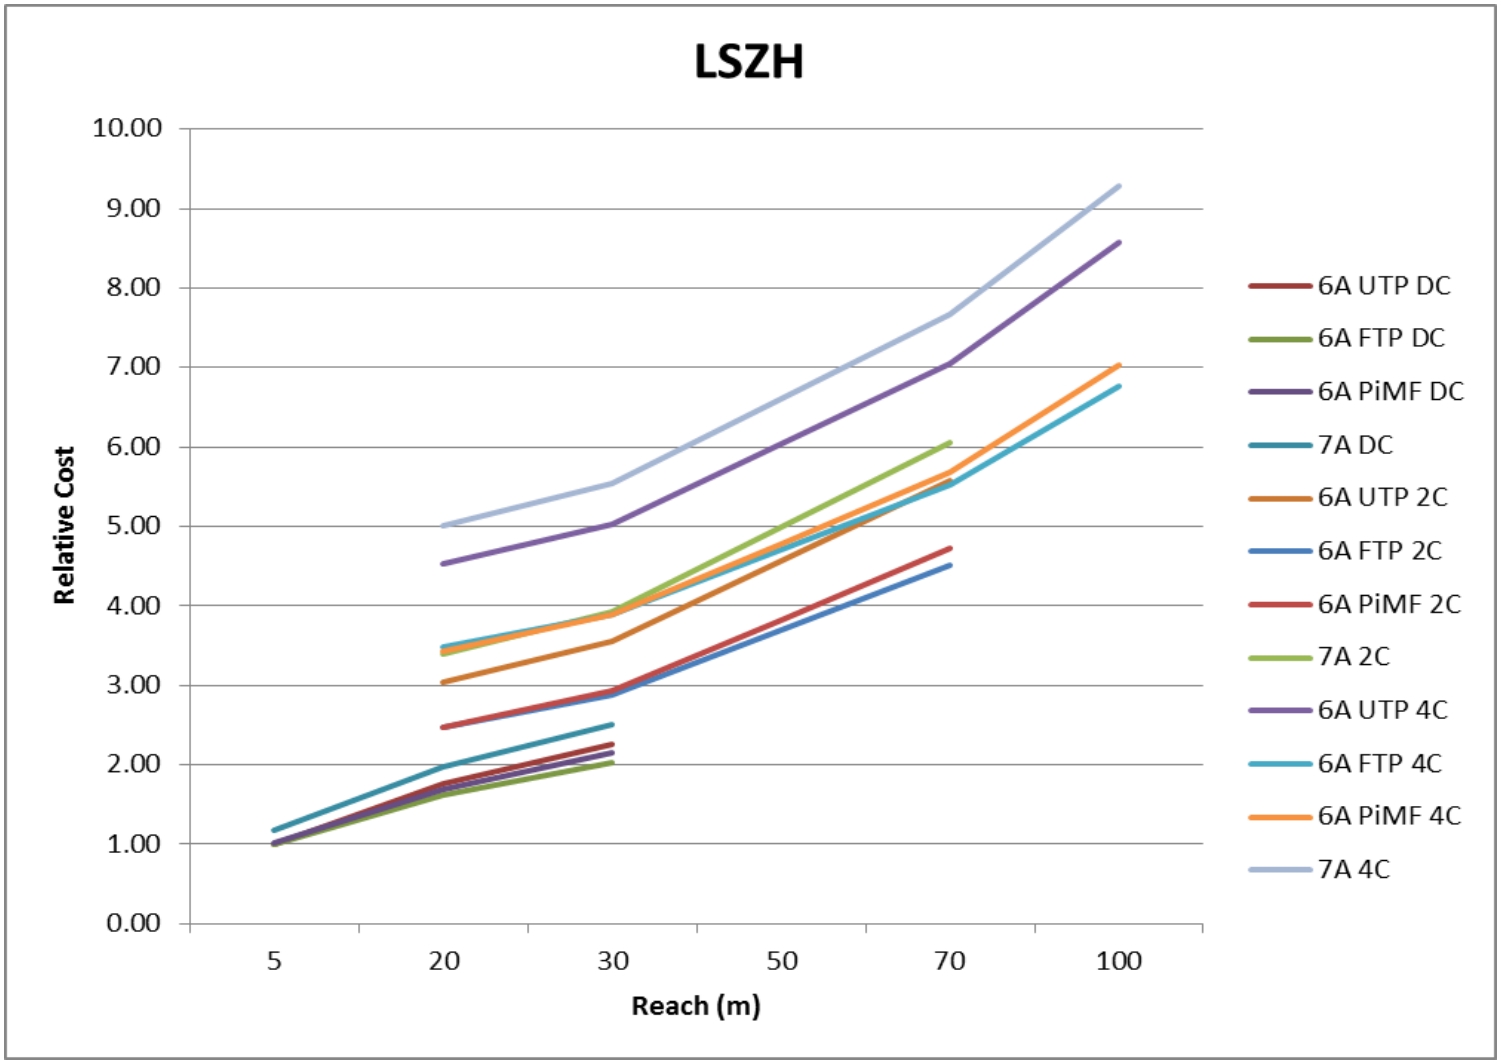

* Data is based upon relative cost analysis per slide 8 with Cat 6A UTP DC as the base cost.

LSZH = Low Smoke Zero Halogen

Selected LSZH Data

- Used only 6A UTP and 7A PiMF up to 30 m

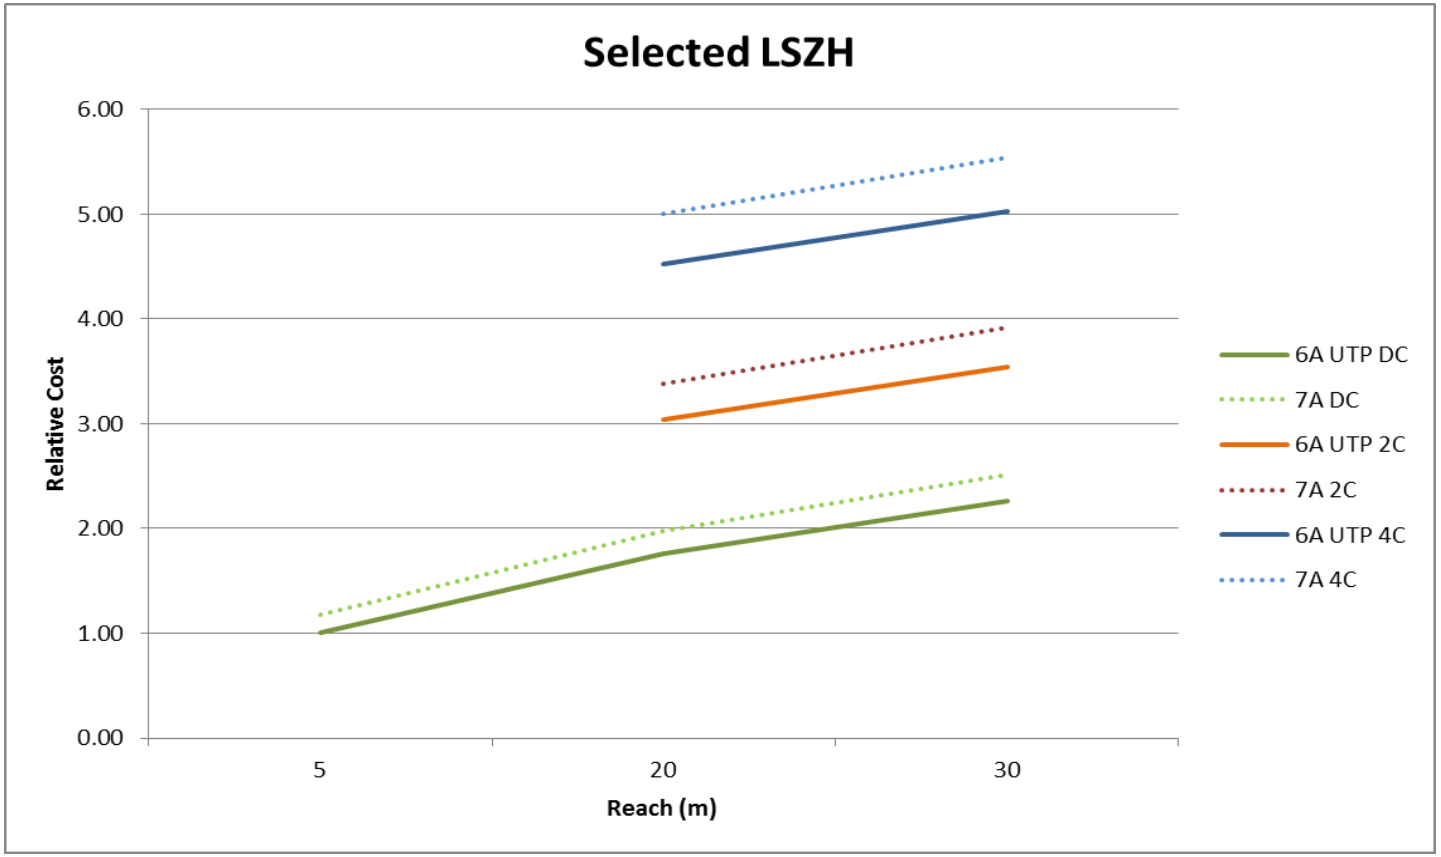

* Data is based upon relative cost analysis per slide 8 with Cat 6A UTP DC as the base cost.

PHY Power Considerations

- Thoughts from bliss_01a_0912.pdf
  - Assumed Class F_A cabling
  - 46m is same power **per bit** as 10GBASE-T @ 100m
  - 22m is same power **per port**
- 10GBASE-T power requirements are considered unacceptable
  - Currently, 3-4 X competing 10GbE technologies
    - Decreases w/ time, but there is a power floor
  - Can only address power through reach objective

Data Center Power Example

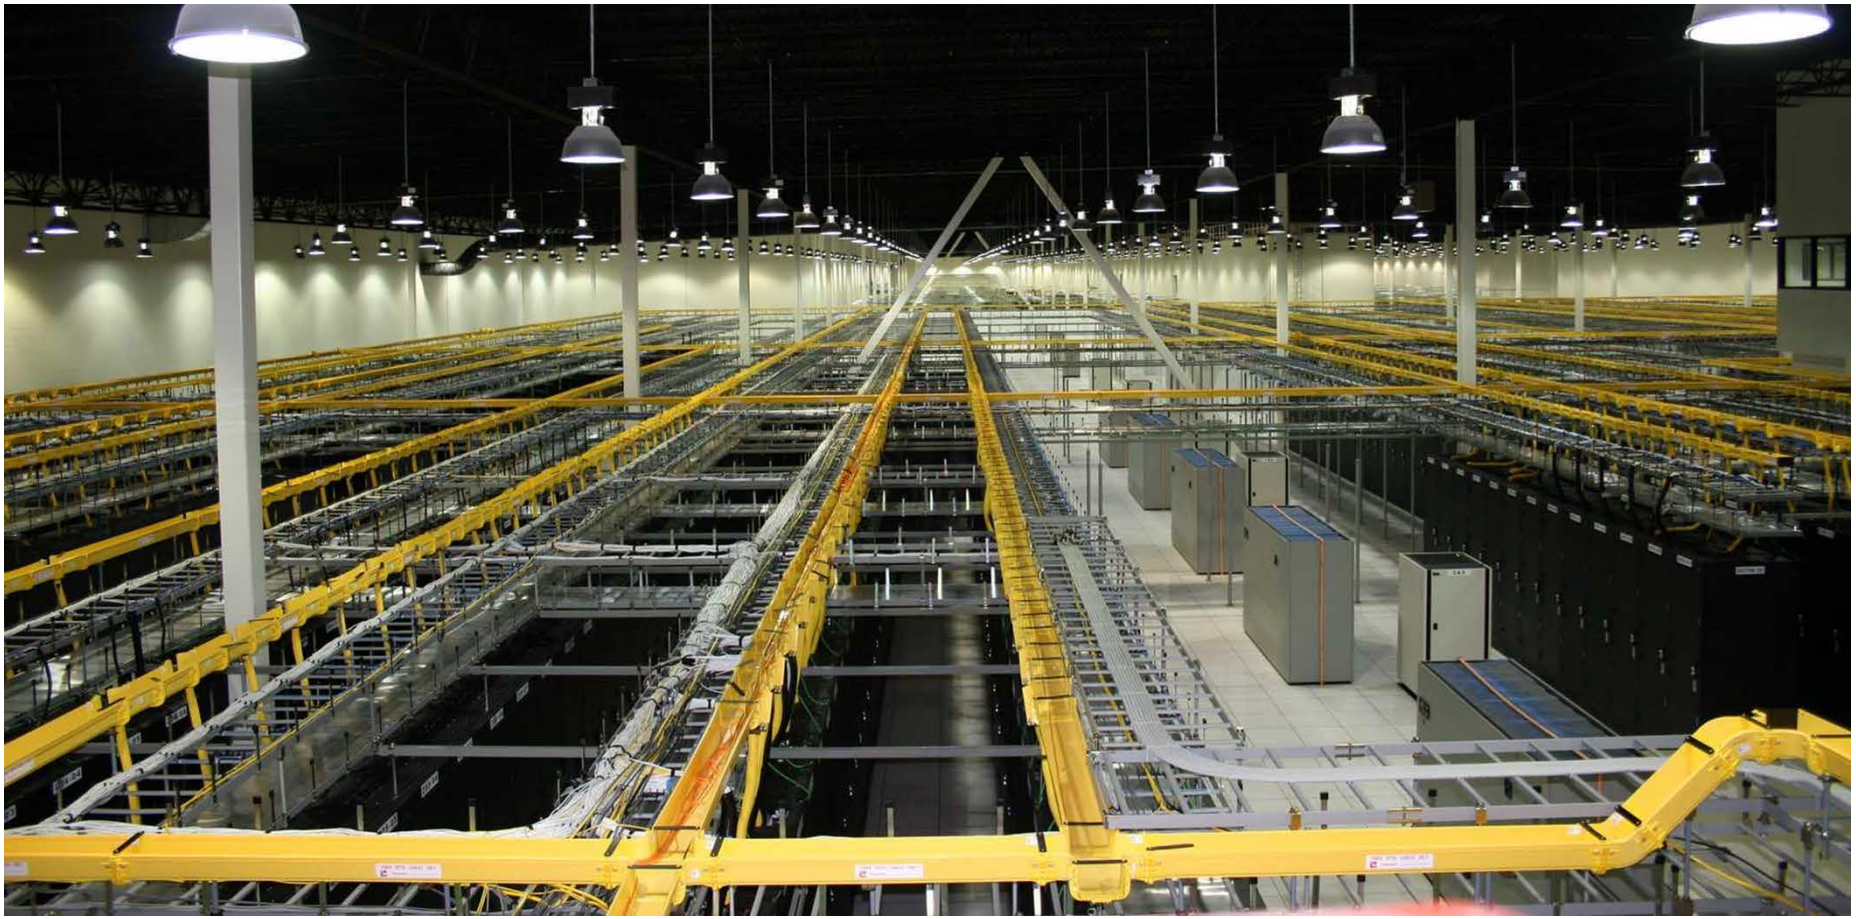

* QTS's data center located in Suwanee, Georgia, USA (www.qualitytech.com)

Data Center Power Calculation

- Assumptions
  - A large data center with 100,000 servers
  - Two PHYs per server-to-switch link
  - Electricity rate is 10.44 cents/kWh†
- Power difference example
  - 1W in PHY power equates to 200 kW
  - Equals $20.88 per hour‡
  - Equals $182,909 per annum‡

† Based on July 2012 average retail price for commercial from US Energy Information Administration (www.eia.gov)
‡ These costs are an example of a possible calculation, are based upon the assumptions made above, and are not actual costs.

More Thoughts on Power

- End users are willing to consider new topologies based on power and cost
  - Occurred with 10 Gigabit Ethernet
  - For example: SR-Lite, SFP+ DAC
- Marketplace factors to consider
  - Data center power is coming under scrutiny†
  - EPA has developed metrics for large network equipment (www.energystar.gov)
  - Green Grid looking into performance/W specs

† Glanz, James, "Power, Pollution and the Internet," New York Times, September 22, 2012.

Creating a Reach Objective

- Volume, volume, volume ☺
  - Selecting reach based upon total available market is a mistake (e.g. 10GBASE-T)
  - Broad market potential occurs when deployment is simplified (e.g. SFP+ DAC)
- Select reach with focus on PHY power
  - Remember: alternative topologies are acceptable
  - Relative cost of cabling does not vary greatly, but noise environment (and power) does
  - Power is a critical consideration in today's market

Recommendations

- Assume only the use of FTP or PiMF cabling
- Pick maximum reach based on Cat 7A
  - Less work to enhance specification vs 6A
- Other reach capabilities can be considered during Task Force
- Forget any reach above 30 m
  - The power is not worth it
- Justify any reach above 20 m

Thank You

Questions?