MACOM™

Partners from RF to Light

RX Reference Receiver Power requirements

Tom Palkert (Macom)

Nov 2019

Palkert_3ck_03a_1119

Contributors:
Tom Palkert (Macom)
Andreas Bechtolsheim (Arista)
John Huang (CIG)

Summary:

- Reduced power is vital to 400G data centers
- Module power is important
- Power savings = $

Data Center operator requirements

- What is the most important issue that needs to be addressed to enable 400G in the data center
  - Power, Power, Power

Brian Welch data

- Welch_3ck_01_0918: Power considerations for 400GAUI-4
- Summary slide:
  - Assuming 2x400G solutions are desired then **very little power is available** for the 400GAUI-4 electrical I/O
  - Taking all reasonable steps to **minimize module side power consumption** in a 400GAUI-4 electrical I/O standard would increase the odds of 2x400G modules being possible

Module power for 8x100-DR-8

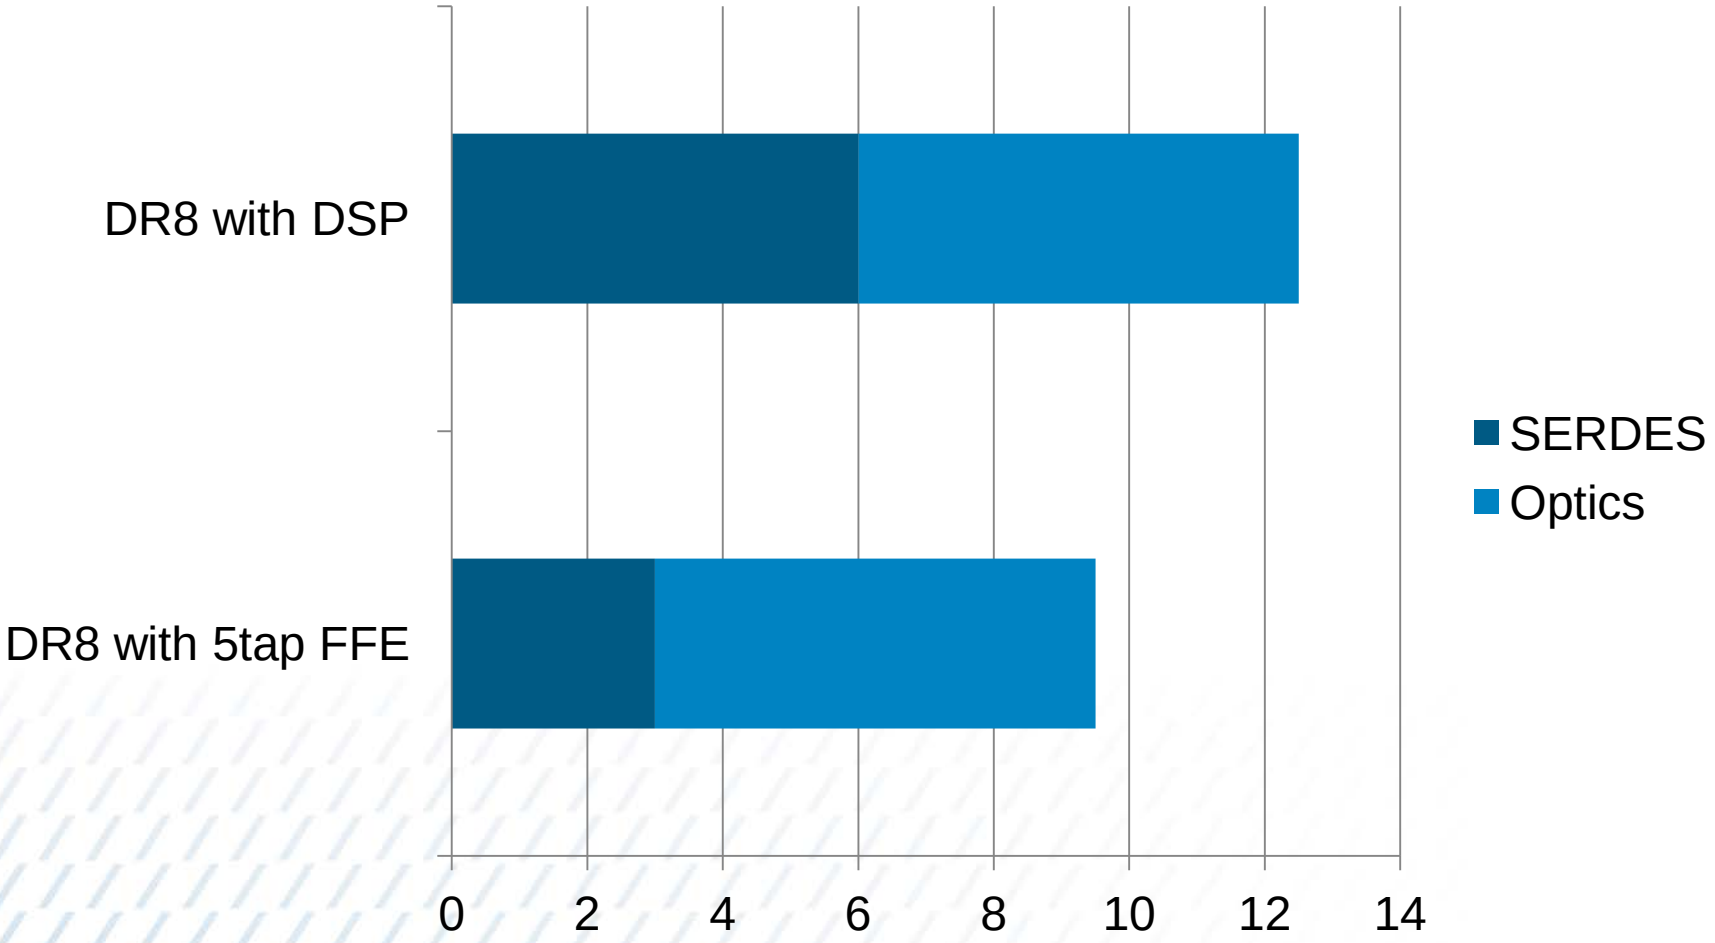

Note: DSP power assumes 5nm cmos

Annual Power savings

> 8760 hours/year

Data Center size (sq ft)	Average annual cost/kw	# of lanes/data center	8 lane DR8 module power savings (W)	Total annual power savings
500-5,000		6,000	3000	.26 Mw
5,000-10,000		12,500	6250	.55 Mw
10,000-25,000		25,000	12500	1.1 Mw
25,000-50,000		50,000	25000	2.2 Mw
> 50,000		100,000	50000	4.4 Mw
Hyperscale		250,000	125000	11 Mw

Lets learn from history

- 10G optical modules moved to high volume when a low cost/low power/non-retimed optical module was available in the market.
- The market will not support a transition from:
  - 10G – Non-Retimed
  - 25G – CTLE only
  - 50G – CTLE only
  - 100G - DSP

Doom and Gloom scenario

- We define a high power reference rx to support legacy channels
- Large data centers define a low power/cost proprietary solution to save power/$.

Conclusions

- Good engineering saves $
- Lets use 100G channels with a low power reference rx.