

Consideration for Advancing Technology in Computer System Packaging

Dale Becker, Ph.D.
IBM Corporation, Poughkeepsie, NY




Motivation

- Modern Computing is driven by
 - Cloud
 - Analytics
 - Mobile
 - Social
 - Security
 - With many data inputs
 - Demanding sophisticated analytics
 - Sent back to distributed users
 - Securely
-
- ➔ More Data Bandwidth
 - ➔ Less Data Latency
 - ➔ Higher integration of computing, networking and storage




Big Data: Why we must move to a new era of computing.

Volume




Terabytes to exabytes of existing data to process

Velocity



Streaming data, milliseconds to seconds to respond

Variety



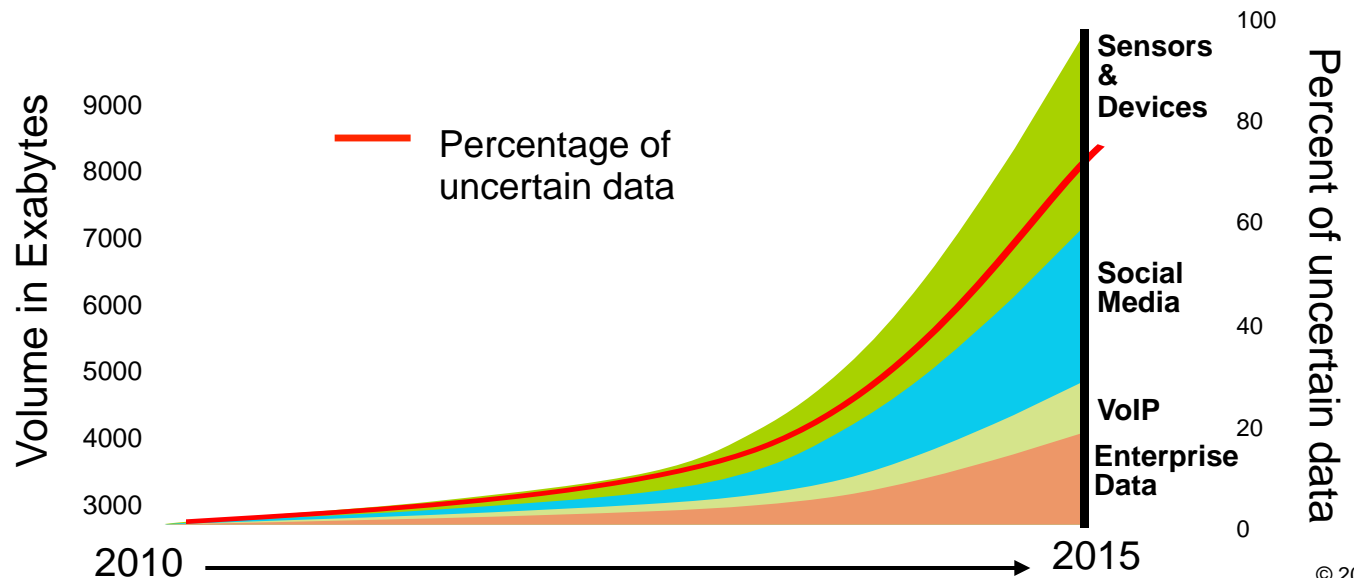
Structured, unstructured, & multimedia text

Veracity



Uncertainty from inconsistency, ambiguities, etc.


This is just the beginning.





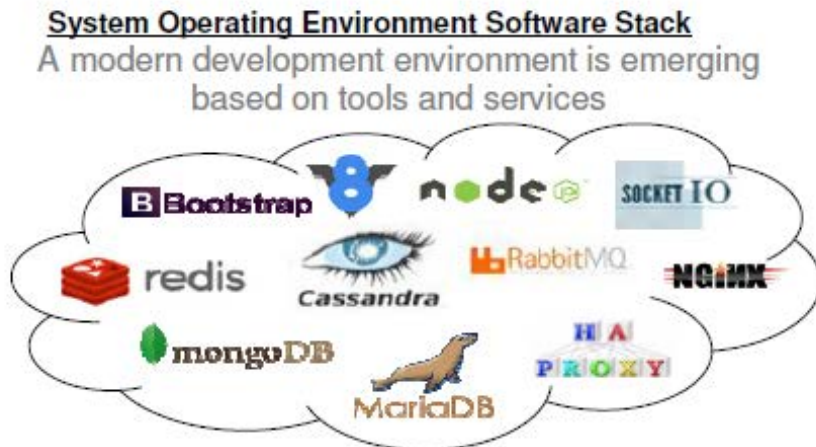
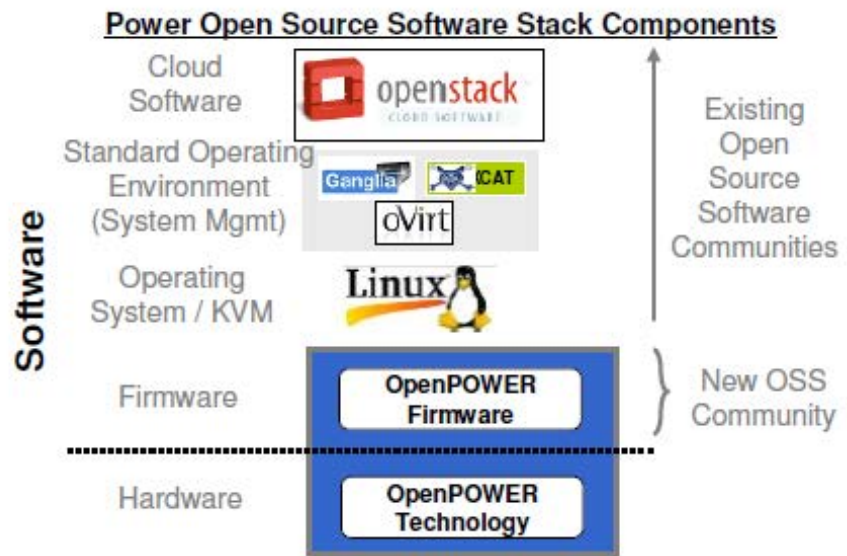
What is Technology?

zEnterprise EC12:

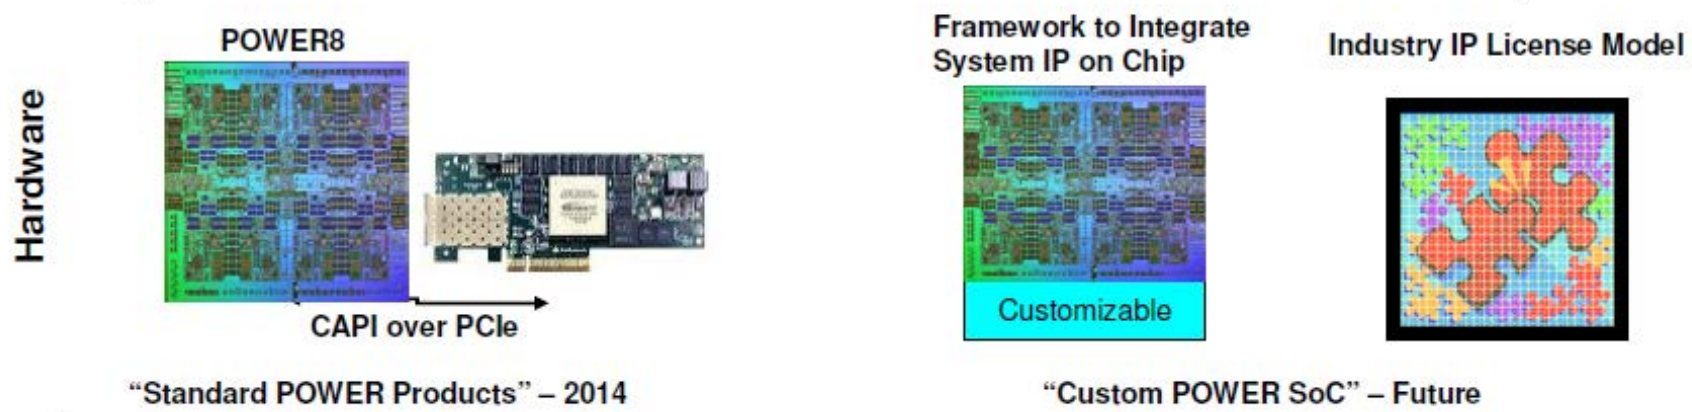




Proposed Ecosystem Enablement



Multiple Options to Design with POWER Technology Within OpenPOWER



50 Years of Mainframe – 1964 IBM S/360

Figure 1 AND/OR INVERT logic module. (a) Completed AOI module, without overcoating. (b) Logic circuit.





Figure 13 Stacked module.



(b)

Power dissipation in mW

	On	Off	
Resistors	28	19	(All R: 5%)
Transistor	7	0	
D_1, D_2, D_3	0	2	
D_4	1	1	
D_5	$\frac{1}{1}$	$\frac{1}{1}$	
Total	37	23	



POWER Die

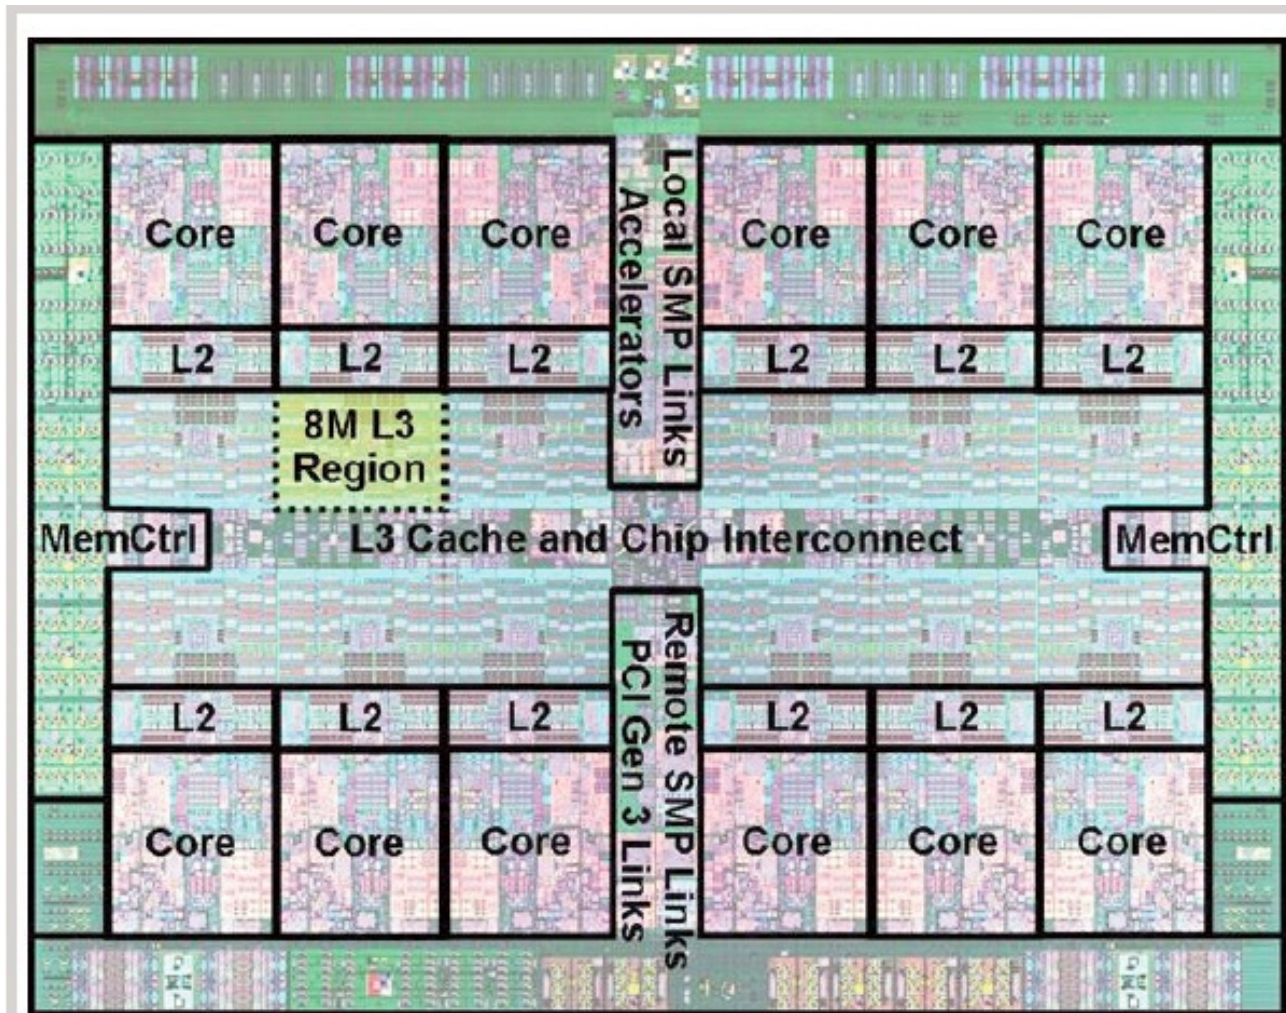


Figure 5.1.7: Annotated POWER8 die photo.



[POWER8: A 12-core server-class processor in 22nm SOI with 7.6Tb/s off-chip bandwidth](#)
 Fluhr, E.J. ; et. al. [Solid-State Circuits Conference Digest of Technical Papers \(ISSCC\), 2014 IEEE International](#)
 Publication Year: 2014 , Page(s): 96 - 97

Ongoing Challenges for system packaging

- Cost – development and product cost
- Physical form factor – Incremental changes
- Signal bandwidth density – Increasing quickly
- Voltage regulation (Power In) – Integrate closer to load
- Cooling (Power out) – Constant power density



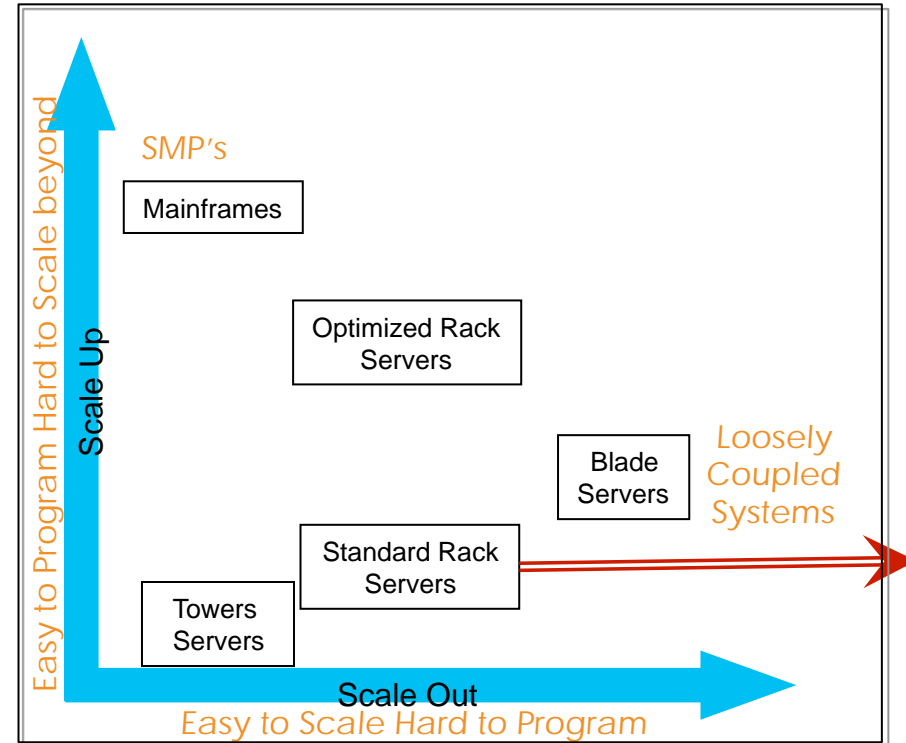
Processor Packaging



Traditional vs. “Cloudified” Hardware

- “Scale-Up”
 - Symmetrical Multiprocessing Systems
 - Large shared memory machines
 - Expensive to scale beyond a certain size
 - 4 / 8 / 16 / 32 sockets
 - 4U/10U/Rack Sized Systems


- “Scale-Out”
 - Loosely coupled systems
 - “Infinite” Scale
 - Mostly 1 & 2 sockets
 - 1U / 2U Form Factor (0.33/0.5/1 wide)



- Significant changes in programming & application paradigms
 - Hadoop/HDFS / NoSQL DB's....
- Open source software community driven
 - Linux / OpenStack ...

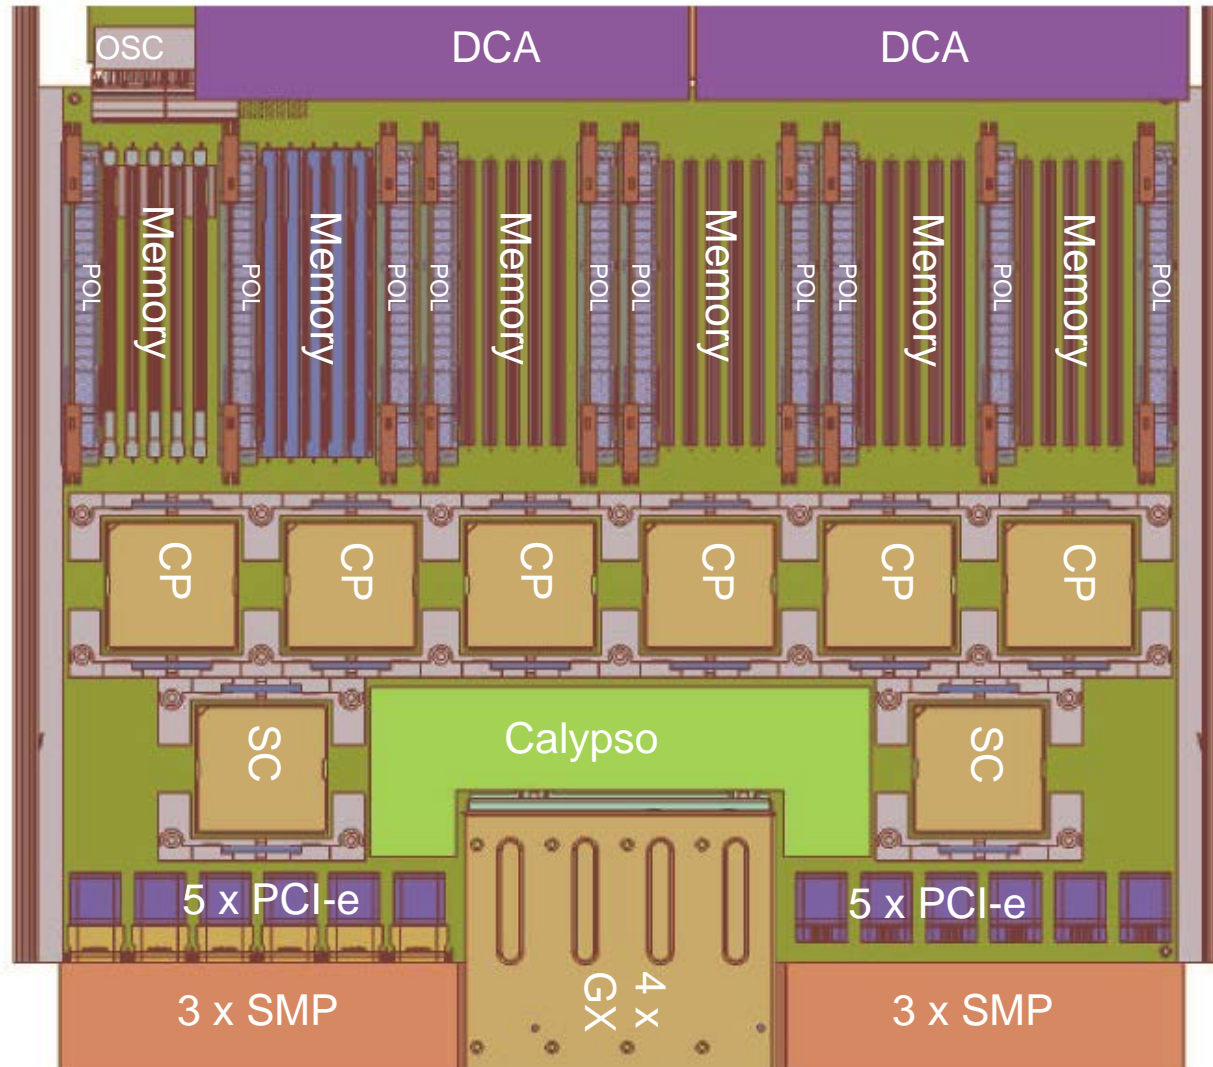


Scale-out





Scale-up



Data rates are increasing

■ Today

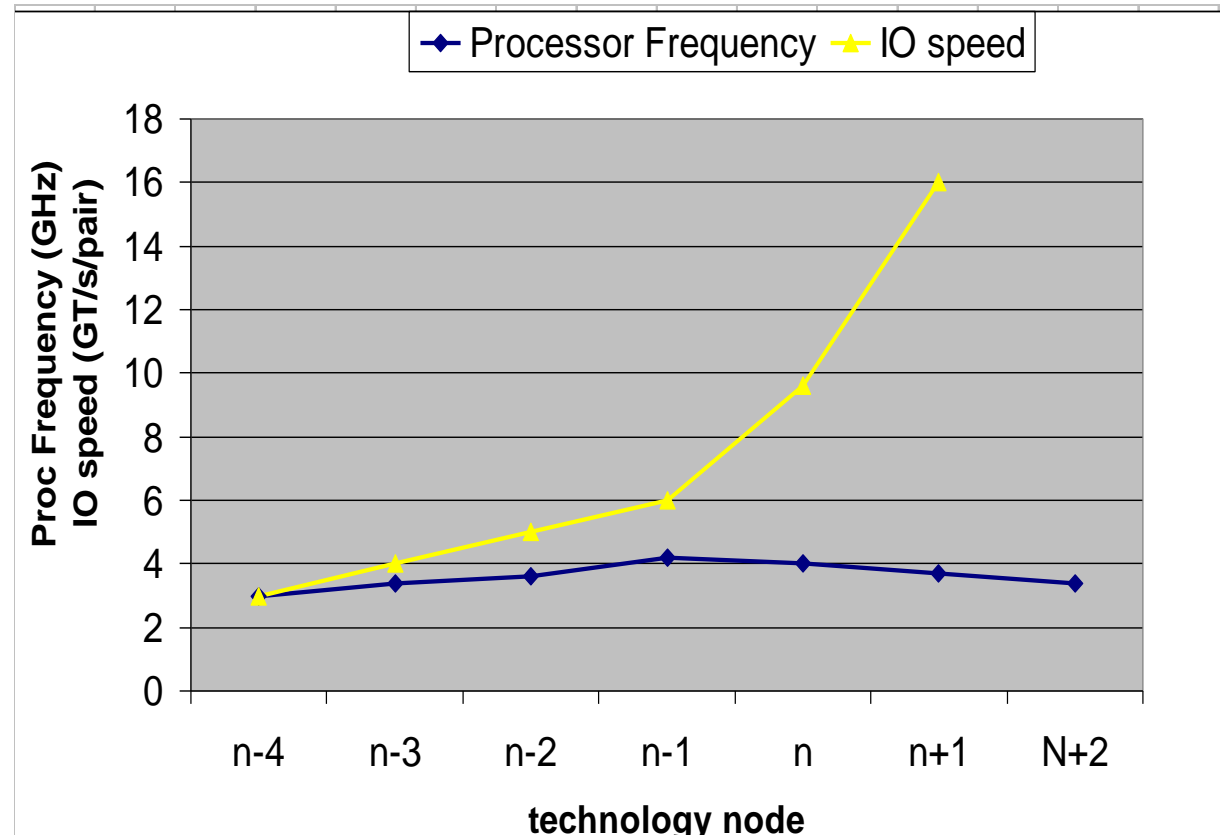
- Proprietary – 12.4 GT/s
- PCIe Gen3 – 8 GT/s
- DDR3 – 2133 MT/s
- SAS – 6 GT/s
- Optics – 10 GT/s

■ Soon

- PCIe Gen4 – 16 GT/s
- Proprietary 28 GT/s
- DDR – 3200 MT/s
- SAS – 12 GT/s
- Optics – 25 GT/s

■ R&D

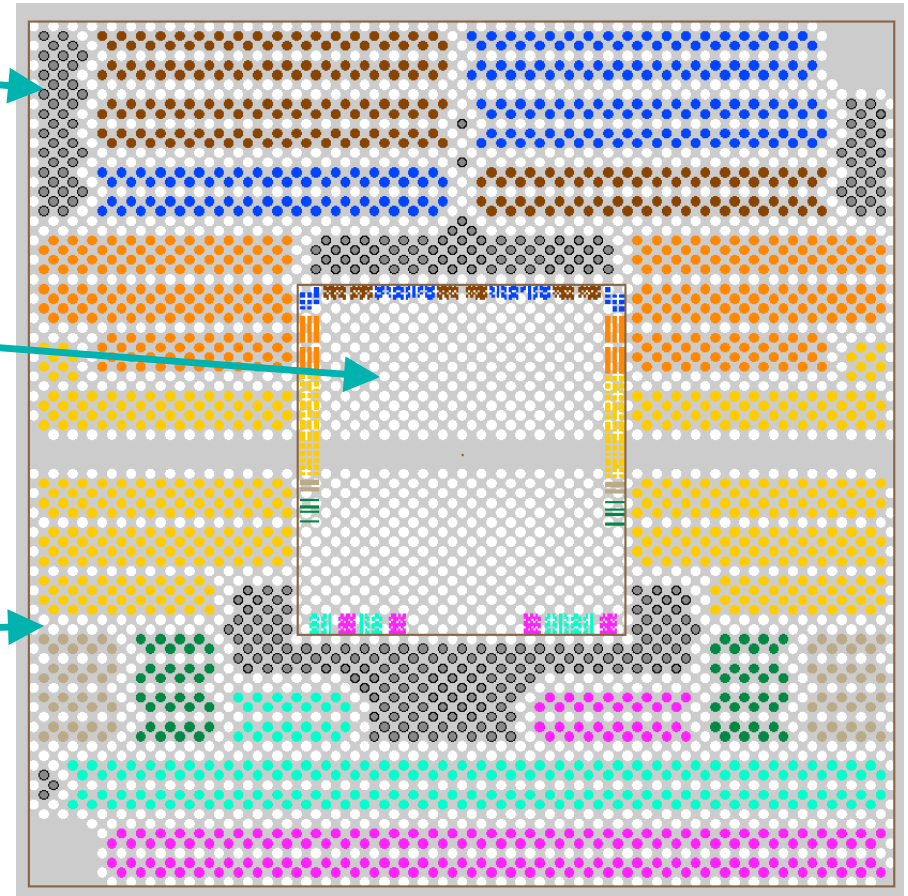
- ~50 GT/s
- PAM4 vs NRZ
- Tighter optics integration






Socket Pin Assignment

- **Signal Pins**
 - More pins = more bandwidth
- **Power Pins**
 - Lower voltage levels require more pins
- **Reference Pins**
 - Higher frequencies require better isolation between signal pins



Reaching a bandwidth breakpoint at the socket level








Pin Density Increases Incrementally

- Over 20 years
 - 1.5 mm min pitch interstitial
 - 50 mil (1.27 mm) square
 - 1 mm square
 - 1 mm min pitch hexagonal

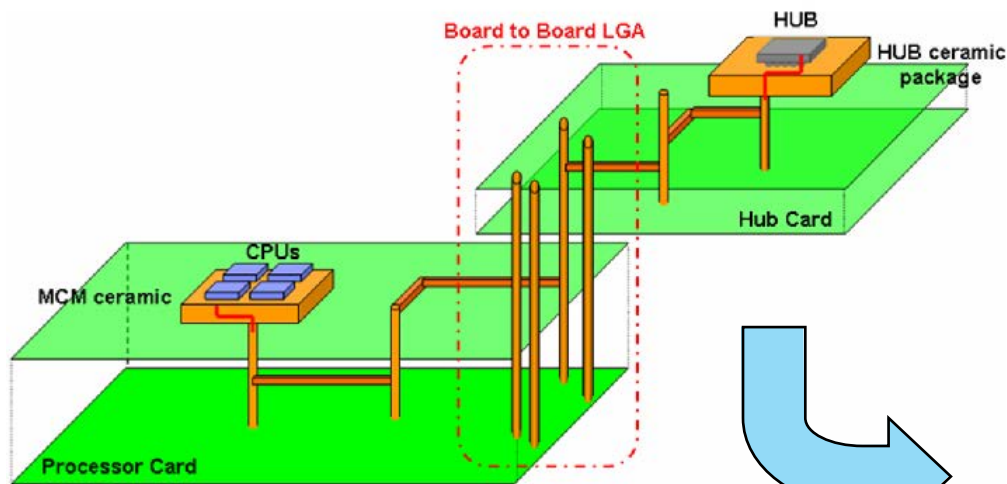
- Sockets are pin limited

- Crosstalk needs to be managed

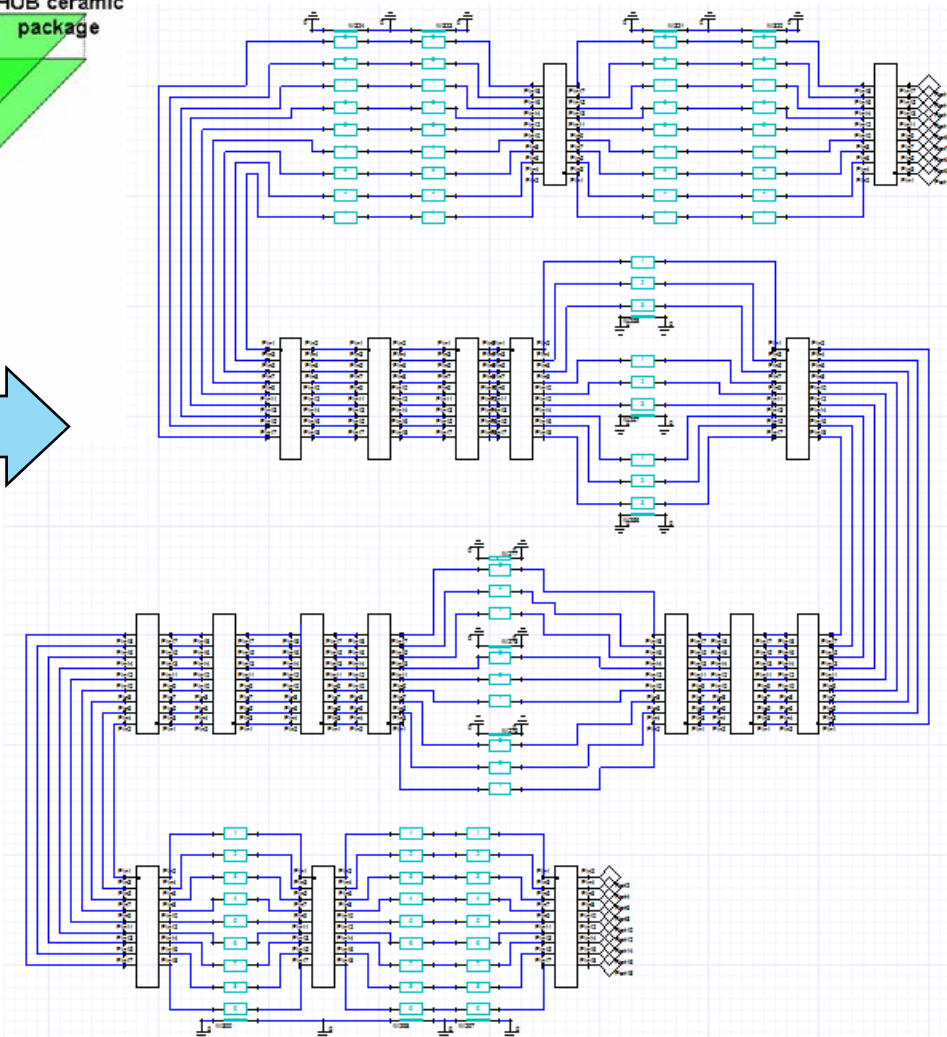




Form factor limitations




- *The system under analysis is composed by two PCBs, two MCMs and three connectors*
- *To represent it adequately 52 models are needed:*
 1. *W-elements to model the TL portions*
 2. *S-parameters (Touchstone) for the 3D parts (Vias and connectors).*
 3. *Mpilog Precompensation Driver macromodel*
 4. *Frequency step for touchstone: 50 MHz*
- Total channel length ~ 70cm



PCB Technology Choices

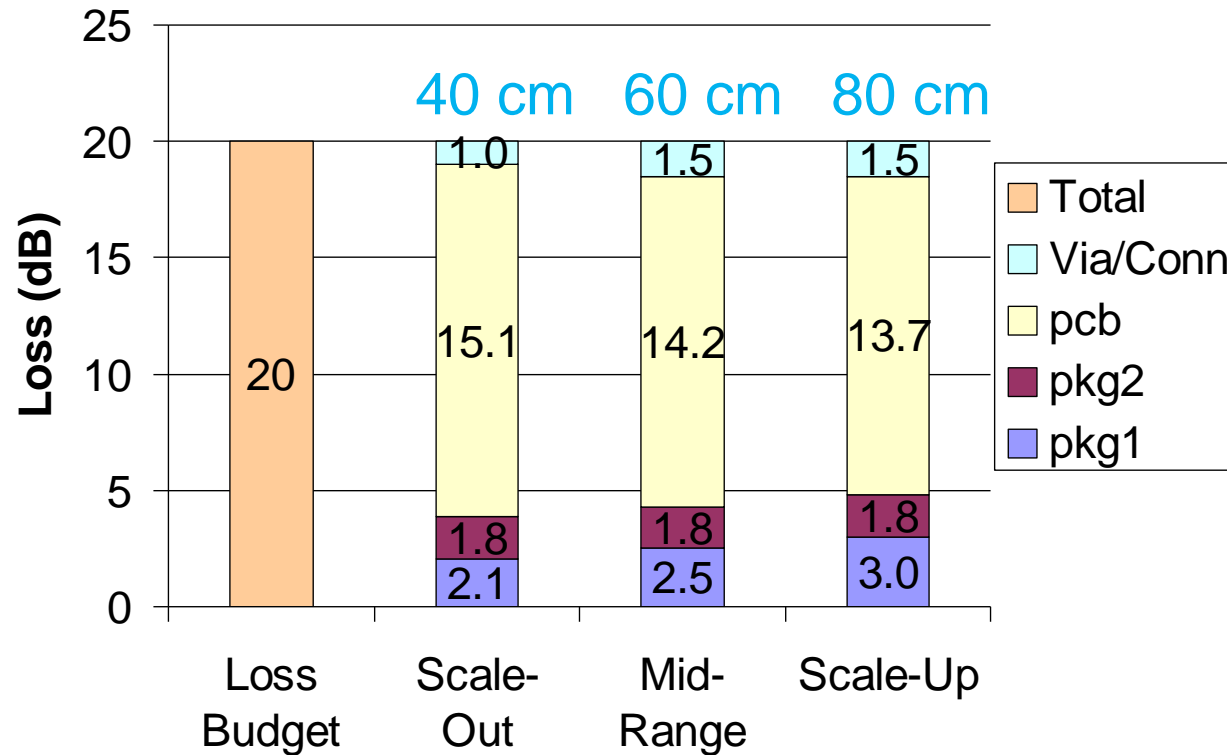
Same speed, different technology



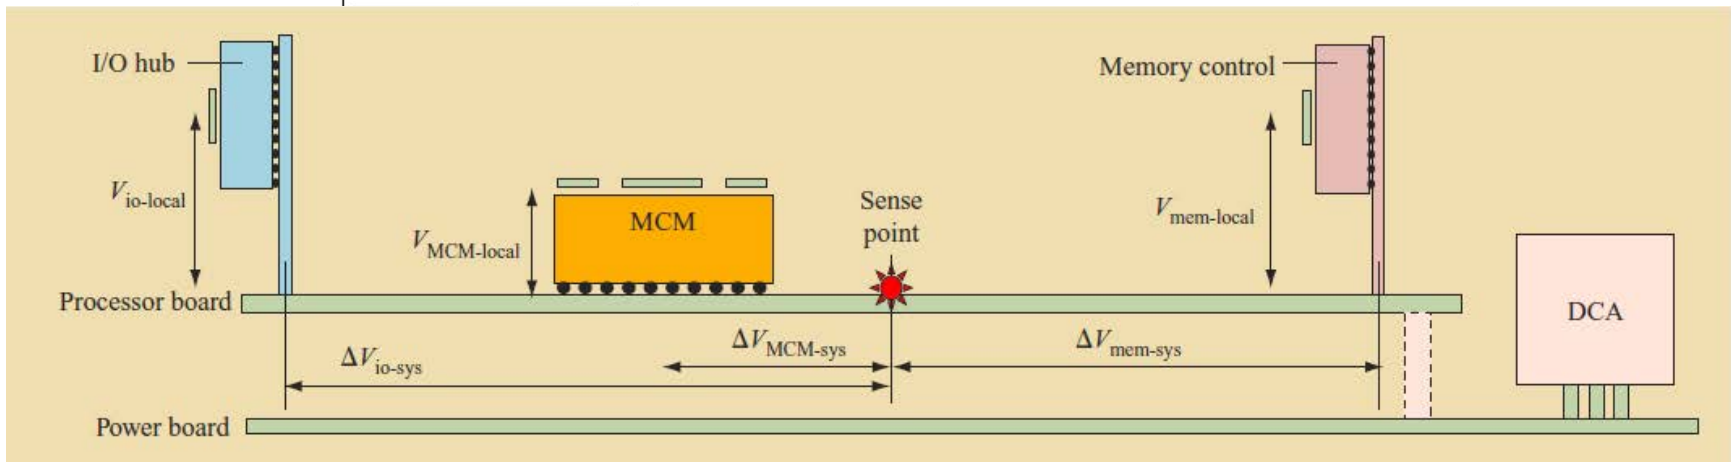

All systems are 8 GT/s meet 20dB total channel loss

Impact of PCB loss – Size of system, length of PCB trace

Loss Allocation Comparison




z196/EC12 Compute Cage




Decoupling improvement

Improved power distribution often counters increased density

zEC12



z13



Deep trench technology – DT capacitors




Figure 2 – Cross section of DT Capacitor (with Word Line and Passing Word Line on top)


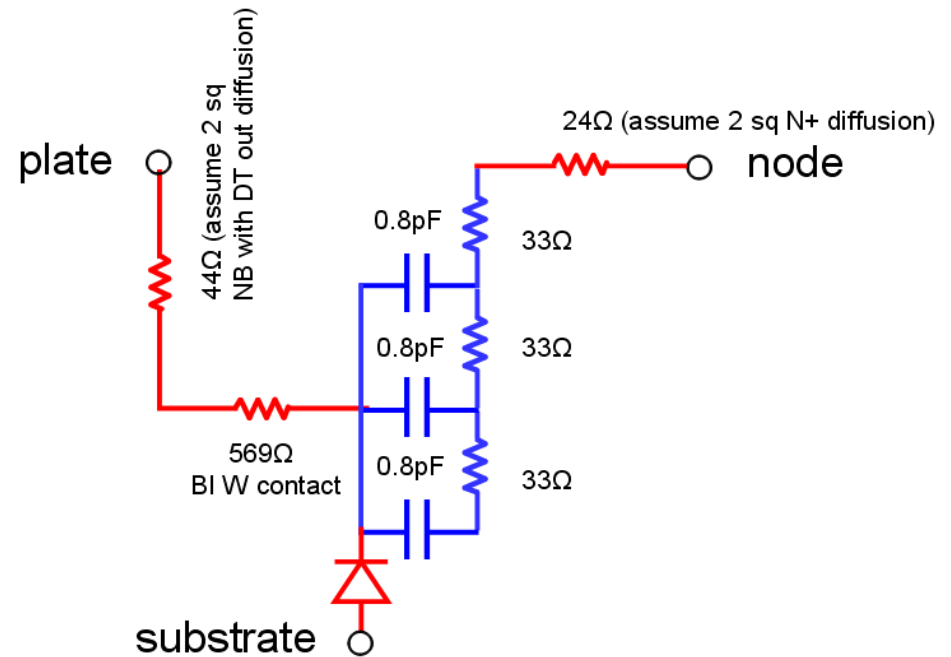



Figure3: Schematic cross section of DZ capacitor





Circuit model for a book having 30 base cells



Power issues : On-chip Caps – Deep Trench

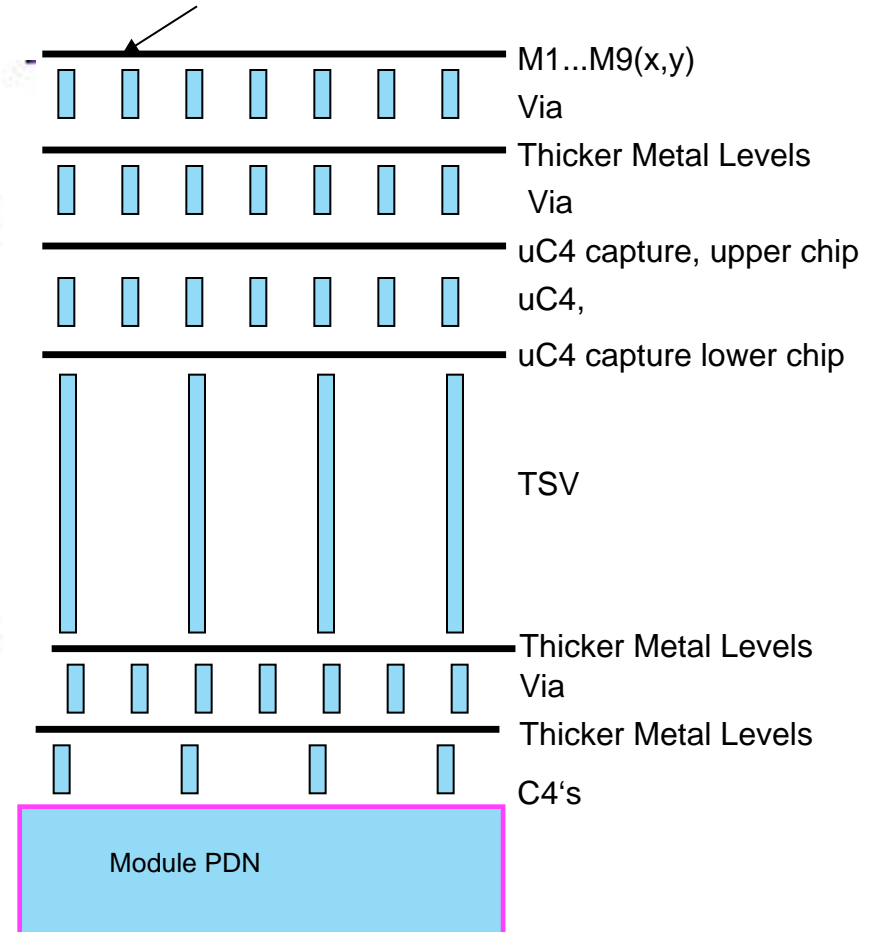

- Deep Trench capacitors
 - Enables significant increase (~20uF) in on-chip decoupling.
 - Mid/high-freq noise significantly reduced
 - Facilitates on-chip voltage regulation





3D chip stacking modeling

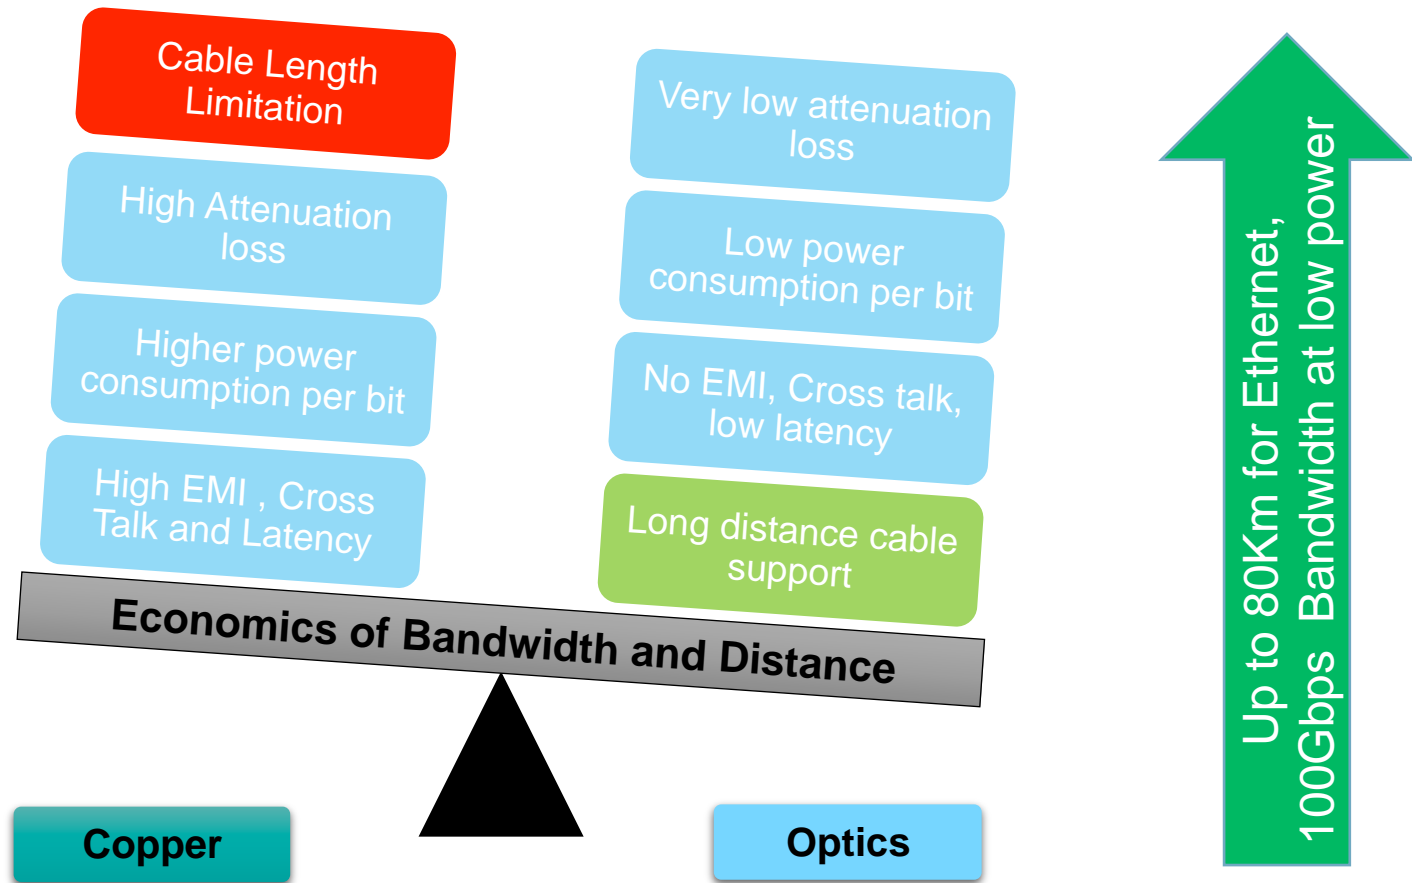
Current sources connect to these vias at small pitch to represent die power map



Knickerbocker, et. al.
2.5D and 3D Technology Challenges and Test Vehicle
Demonstrations
ECTC 2012





The Case for Optics



- Increasing benefits with optical, but products generally cost more than copper
- Optics less expensive when integrated with silicon - Silicon Photonics
 - Photonics integrated into silicon base
 - Reduces cost and provides higher bandwidth



Future Integration



Increasing Bandwidth Density in Future Optical Interconnects

B. G. Lee, C. Baks, F. E. Doany, C. Jahnes, R. John, D. M. Kuchta, P. Pepeljugoski,
A. V. Rylakov, C. L. Schow, S. Assefa, W. M. J. Green, Y. A. Vlasov, J. A. Kash
IBM T. J. Watson Research Center, 1101 Kitchawan Road, Yorktown Heights, NY 10598
bglee@us.ibm.com

CIPS 2014, February, 25 – 27, 2014, Nuremberg/Germany

Laminate with Thermal - Power Insert for Efficient Front-Side Heat Removal and Power Delivery

Dominic Gschwend*, Timo Tick*, Stefano Oggioni†, Stephan Paredes*, Keiji Matsumoto‡, Manish K. Tiwari¶, Dimos Poulikakos¶ and Thomas Brunschwiler*§

Ongoing Challenges for system packaging

- Cost
 - Technology Reuse
- Physical form factor
 - Form factors change slowly
 - Scale-in functional integration
- Signal bandwidth density
 - Frequency per lane
- Voltage regulation (Power In)
 - Regulation moves closer to load
 - More effective decoupling capacitors
- Cooling (Power out)

Conclusions

- IT boundaries are becoming less clearly defined
 - Processor, Storage, Networking Integrated
- Systems
 - Cloud, Analytics, Mobile, Social, Security
- System Hardware
 - Drives system capacity with cores and computing capacity
 - Drives interconnect bandwidth at processor, node, system and network level
- All electrical performance elements must be balanced
 - Bandwidth
 - Power Distribution
 - Thermal
- Technology enables the innovation that systems provide, we must choose wisely.