

# **Cloud&Data Center 2020 Vision**

EPSRC Communications&Networking  
Network meeting on Photonics, Oxford, 28.9.2011  
Jon Crowcroft, Computer Lab, Cambridge  
[Jon.crowcroft@cl.cam.ac.uk](mailto:Jon.crowcroft@cl.cam.ac.uk)

## A Data Center - 1

- Somewhere in Oregon



## A Data Center - 2

- Rows of Racks



## A Data Center - 3

- Coulie scale
- Engineering
- Thanks to
- Dave Clark



## A Data Center - 4

- Trust/Reliability (Woody Guthrie worked on Coulie Dam)



## Interconnect Requirements

- Interconnect Service to external capacity (web/Cloud)
- Interconnect servers to each other (map/reduce)
- Unpredictable load/hotspot ->
  - VL2
  - Fat Tree
  - De Bruijn/BCube
  - Many others
- See: <http://dl.acm.org/citation.cfm?id=1921189>  
A Cost comparison of Data Center Network Architectures, Popa, Ratnasamy, Iannacone, Krishnamurthy, Stoica

## Performance Requirements

- Spint into Internal to Data Center and External
- Availability (MTTR < 200ms – e.g. for skype or youtube)
  - Capacity (Internal v. external)
  - Latency (Internal v. external)
  - Power
- Other? Yes
  - Flexibility

## Server External Capacity Today

- From [www.eurecom.fr/~ennajjar/PUB/YouTube.pdf](http://www.eurecom.fr/~ennajjar/PUB/YouTube.pdf)

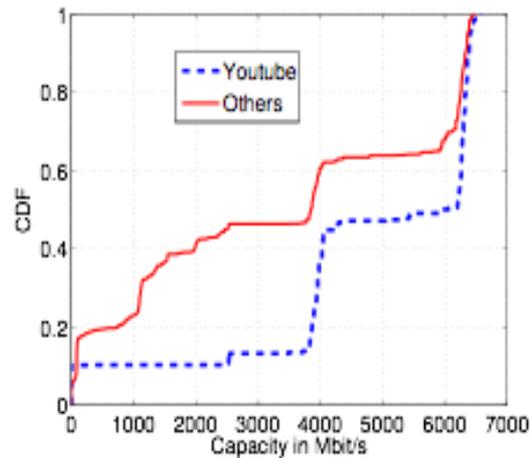


Figure 4: CDF of Servers Capacities: YouTube vs. Others

## Data Center Switching Capacity Internal

- From <http://www.aristanetworks.com/en/products/7500series>
  - 192 \* 10GigE
  - 2.3GB packet buffer per linecard module
  - 4.5usec port to port latency
  - 10 Terabit lossless switch fabric
  - 648Gb/s bandwidth per linecard Slot (1.25 Terabit full-duplex)
  - Less than 10 watts per wirespeed 10GbE port

## Interconnect Capability

- Storage Area Net uniform access
- Fast Migration of Virtual Machines
- Hotspot mitigation
- Zero downtime reconfig/scale up

## Where are the photonics?

- Obviously a lot of the external capacity is met by telecom photonics already
- Internally, only for transmission
- Maybe one more level in (usually 2 layer) could be optically switched
- Future? Who knows
- Integration on processors – reduce energy cost?
- Desperate measures – wireless HDMI – 7Gbps 60GHz
- Alternative – decentralise data centers more

Questions?

[jon.crowcroft@cl.cam.ac.uk](mailto:jon.crowcroft@cl.cam.ac.uk)