



Computer Infrastructure for Big Data Analysis



Computer Architecture



Storage



Hard drives

Bandwidth around 200 MiB/sec

Sizes up to 14 TB

Shinkles are a game changer for performance



SSD

As fast as the bus

SATA 500 MiB/sec

SAS 2GiB

NVME 4GiB

Sizes up to 120TB



JBODS

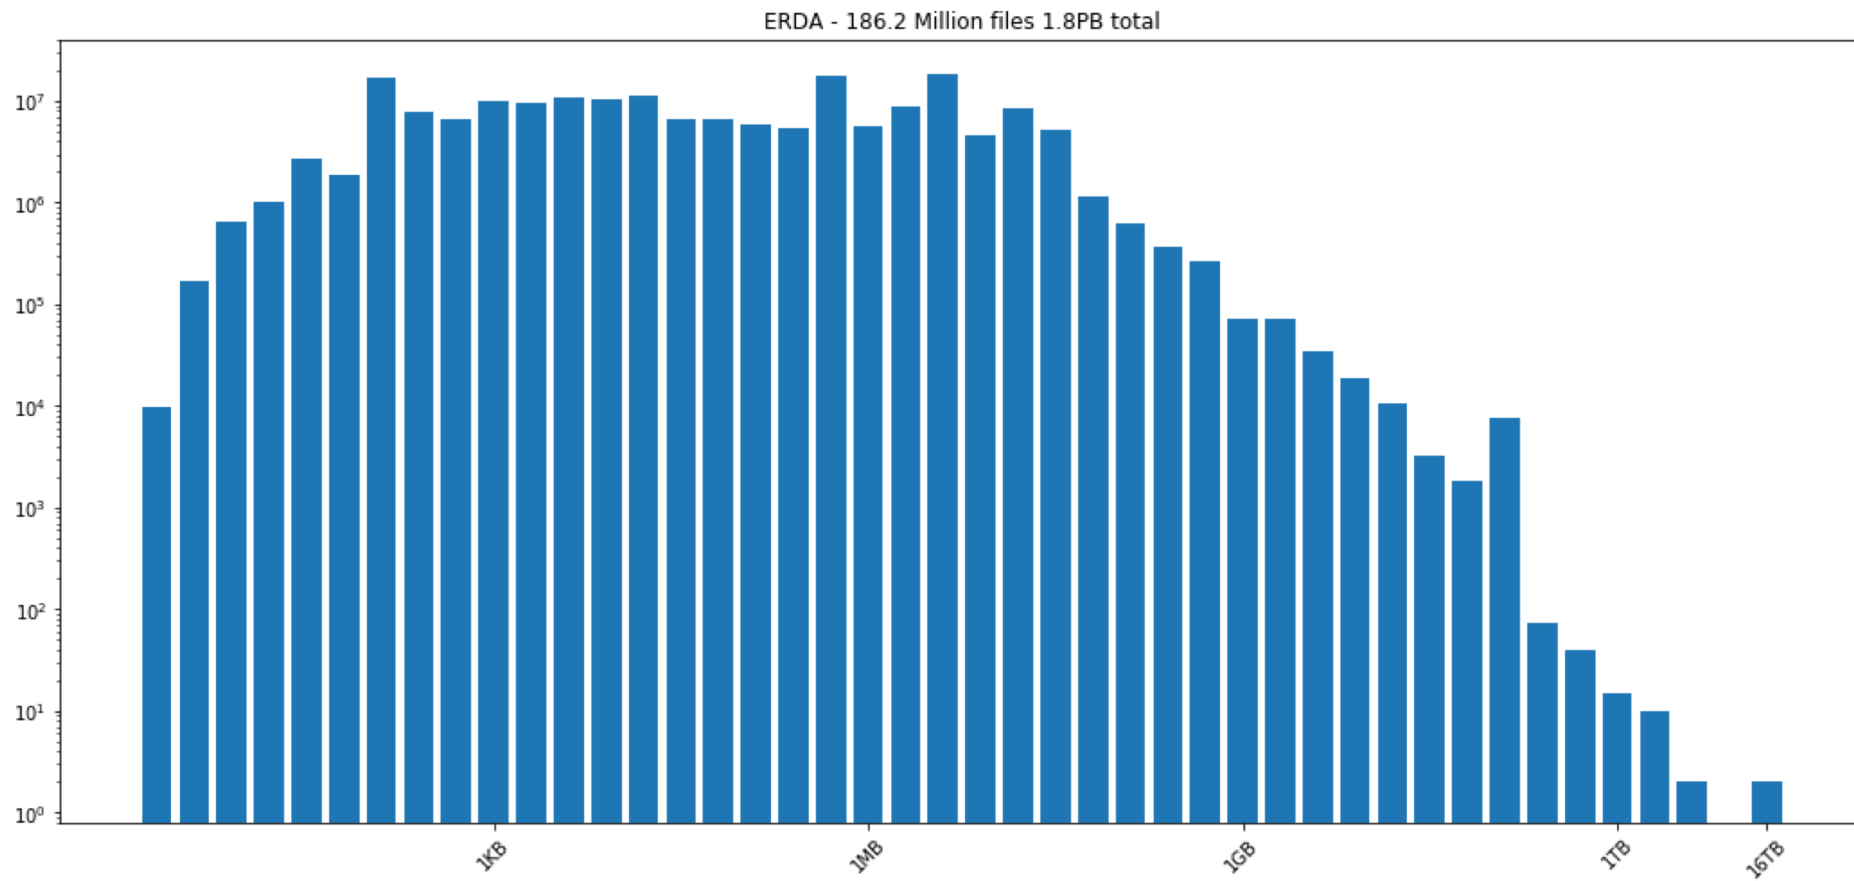
If you have enough data JBODS are fast and cheap

106 drives per chassis

1.7 PB



UCPH/SCIENCE – ERDA (August 2018)



Quiz 2A

What storage parameters matters for machinelearning?

Any affinity wrt the sort of algorithm you use?



Networks



Interconnect approaches

WAN

- 'infinite distance'

LAN

- Few kilometers

SAN

- Few meters

Backplane

- Not scalable



Infinite😊

Nokia and Vodafone will bring 4G to the Moon


The network will allow for HD video streams of the lunar surface.



Rob LeFebvre, @roblef
02.27.18 in [Internet](#)

8
Comments

1618
Shares



Sponsored Li



Physical Cluster Interconnects

Gigabit EtherNet

10 Gigabit EtherNet

- 25
- 40
- 100

Infiniband

Omnipath



10 Gb Ether Spring

Switch (8 port)

- 5.248 kr

NIC

- 2652 kr

Cable 30 kr

Per node: app 3500 kr



40 Gb Ether


Switch (4 port + 48 10Gb)

- 43.464 kr

NIC

- 2.976 kr

Per node: app 3500 kr



100 Gb Ether

Switch (32 x 100Gb)

- 55.000 kr


2m cable

- 1000 kr

NIC

- 4000 kr

Per node: app 7000 kr



Infiniband

New standard

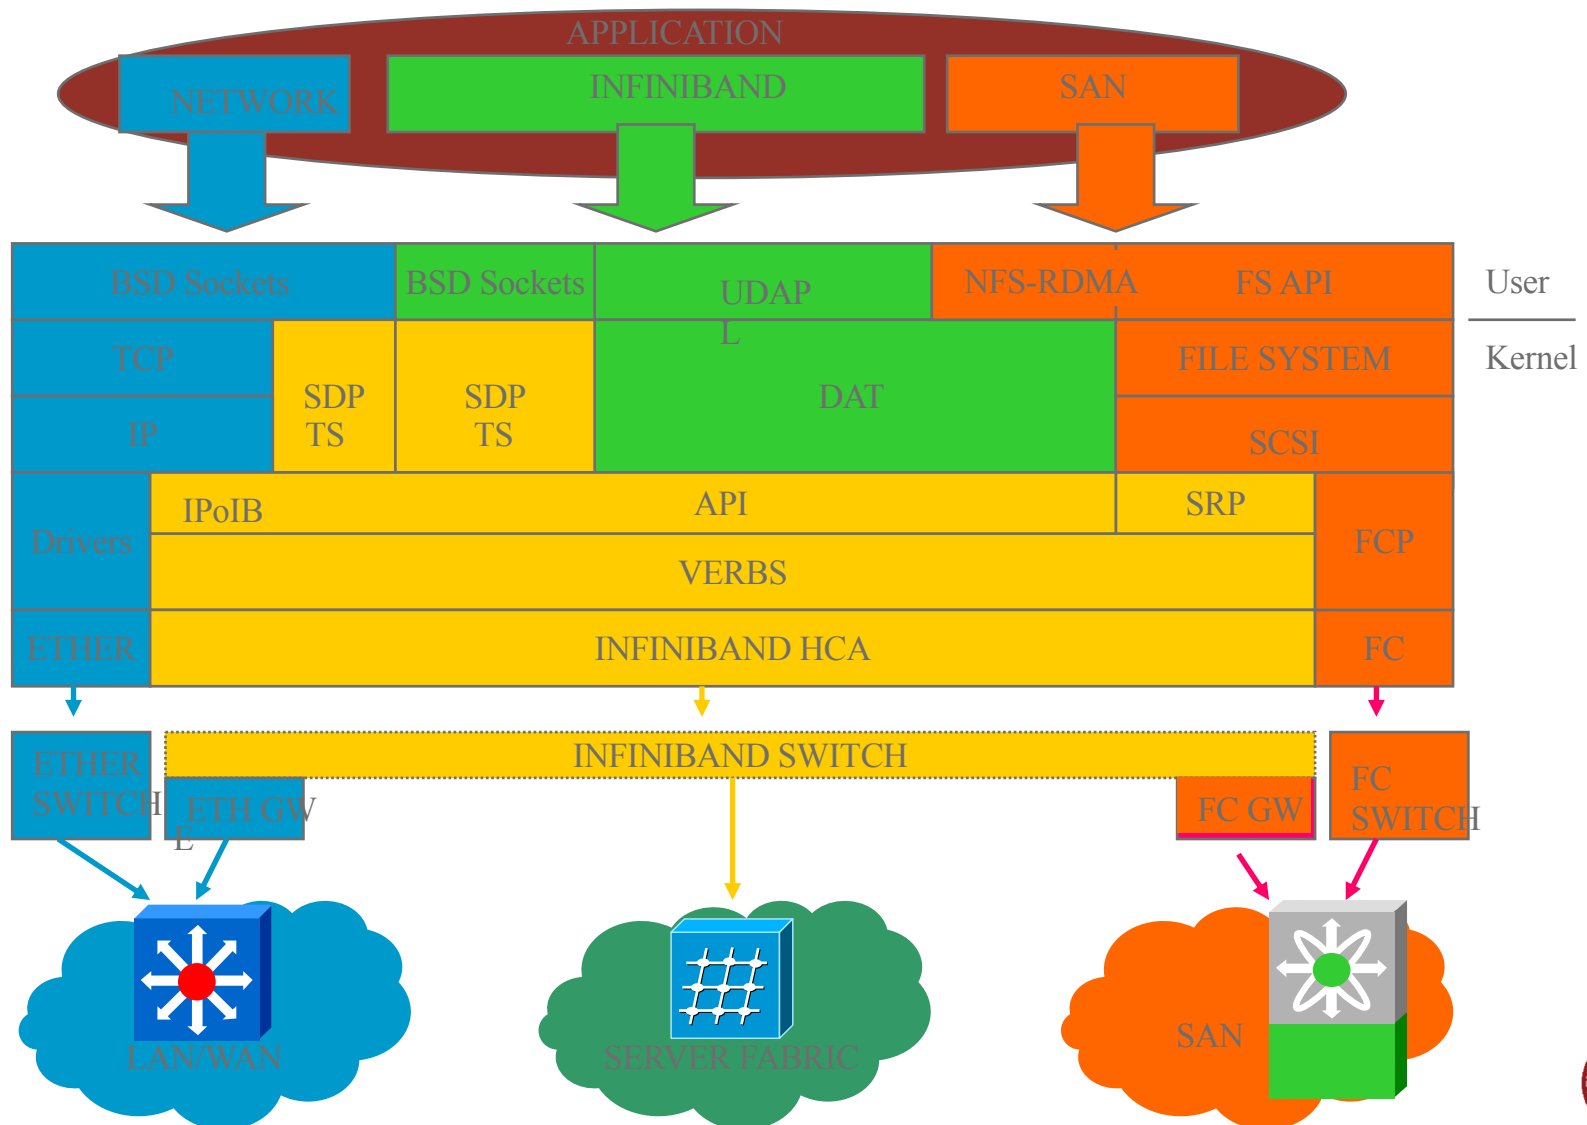
An extension of PCI-X

SDR/DDR/QDR/FDR/EDR

- 1x = 2 / 4 / 8 / 13.64 / 24.24 Gbps
- 4x = 8 / 16 / 32 / 54.54 / 96.97 Gbps
- 12x = 24 / 48 / 96 / 163.64 / 290.91Gbps



The InfiniBand Driver Architecture



InfiniBand Price / Performance

	InfiniBand PCI-Express	10GigE	GigE	Myrinet D	Myrinet E
Data Bandwidth (Large Messages)	950MB/s	900MB/s	100MB/s	245MB/s	495MB/s
MPI Latency (Small Messages)	5us	50us	50us	6.5us	5.7us
HCA Cost (Street Price)	\$550	\$2K-\$5K	Free	\$535	\$880
Switch Port	\$250	\$2K-\$6K	\$100- \$300	\$400	\$400
Cable Cost (3m Street Price)	\$100	\$100	\$25	\$175	\$175

• Myrinet pricing data from Myricom Web Site (Dec 2004)

** InfiniBand pricing data based on Topspin avg. sales price (Dec 2004)

*** Myrinet, GigE, and IB performance data from June 2004 OSU study



• Note: MPI Processor to Processor latency – switch latency is less

InfiniBand Cabling

CX4 Copper (15m)

Flexible 30-Gauge
Copper (3m)

Fiber Optics up to
150m



Infiniband

Interface adapter (100 Gbs)

- 4000 kr

Switch (36 ports)

- 63.000 (135.000 kr last year)

Cable

- 500 kr (was 1000 kr last year)

App 7.000 kr per node

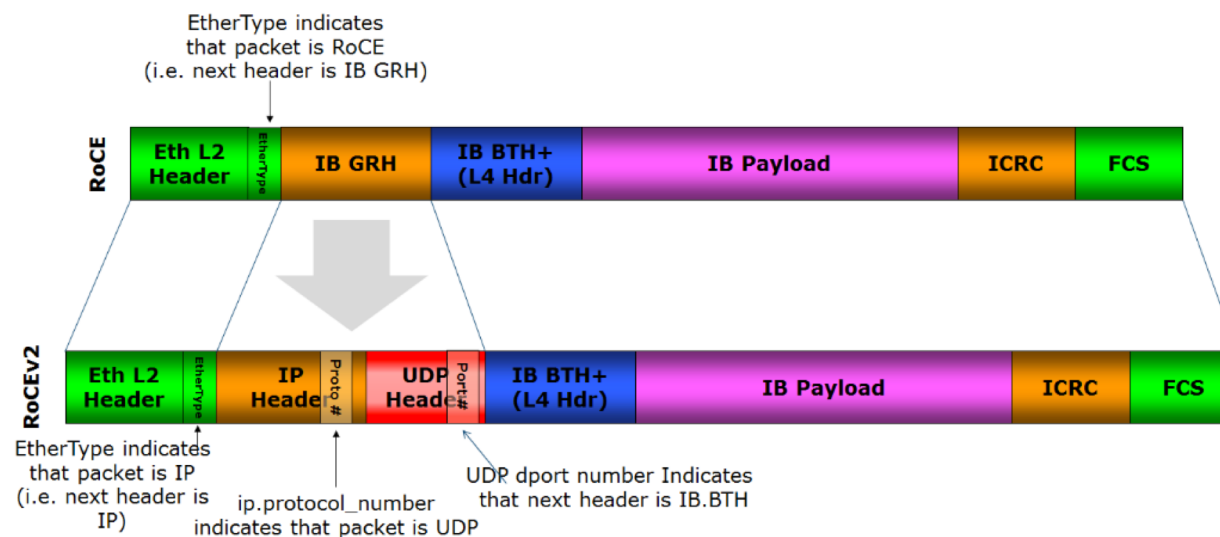


RoCE

Ethernet is not bound to the IP protocol

- We used to make all kinds of protocols on Ethernet

RoCE is a standard for low latency communication over ethernet



25 Gb Ether Spring 2020

Switch (72 x 25Gb)

- 35.000 kr


2m cable

- 300 kr

NIC

- 900 kr

Per node: app 1700 kr



Quiz 2B

What would you like from an interconnect for machine learning?

Any affinity for algorithms?

