

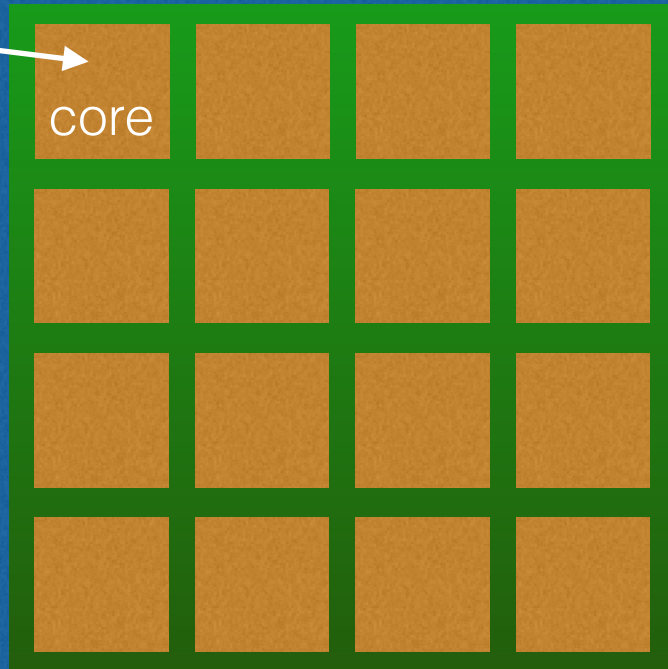
(software agnostic)
Computational
Considerations

The Issues


- CPU
- GPU
- Emerging - FPGA, Phi, Nervana
- Storage
- Networking

CPU

2 Threads



Processor/Chip



Processor/Chip

Computer

CPU

- Threads vs. Cores
- Cores are physical processor units
- Threads take advantage of redundancy to do more than one thing at a time in one physical processor
- 24 cores -> 48 threads -> performance ~ 28-30 cores
- (note that Amazon normally sells threads, not cores)

CPU

- Advantages
 - Generic use, no special software setup
 - Long term support, 8 year old hardware still useful
 - Many cores/machine
- Disadvantages
 - Slower than (well coded) GPU software (for many tasks)
 - GPU speedup typ. 10-100x, 1 GPU vs a single core

Intel Chips 11/2017

CPU	Cores/ chip	Clock Speed (Ghz)	Cost/chip (list)
Intel® Xeon® Platinum 8180	28	2.5	\$10009.00
Intel® Xeon® Platinum 8168	24	2.7	\$5890.00
Intel® Xeon® Gold 6148	20	2.4	\$3072.00
Intel® Xeon® Gold 6154	18	3.0	\$3543.00
Intel® Xeon® Gold 6132	14	2.6	\$2094.00
Intel® Xeon® Gold 6126	12	3.2	\$1776.00
Intel® Xeon® Gold 6126	12	2.6	\$1776.00
Intel® Xeon® Gold 5115	10	2.4	\$939.00

Workstation cost includes 2x1080Ti GPU, 256 GB RAM, 60 TB RAID6

Intel Chips 11/2017

CPU	Cores/ chip	Clock Speed (Ghz)	Cost/chip (list)	Cost/core
Intel® Xeon® Platinum 8180	28	2.5	\$10009.00	\$357.46
Intel® Xeon® Platinum 8168	24	2.7	\$5890.00	\$245.42
Intel® Xeon® Gold 6148	20	2.4	\$3072.00	\$153.60
Intel® Xeon® Gold 6154	18	3.0	\$3543.00	\$196.83
Intel® Xeon® Gold 6132	14	2.6	\$2094.00	\$149.57
Intel® Xeon® Gold 6126	12	3.2	\$1776.00	\$148.00
Intel® Xeon® Gold 6126	12	2.6	\$1776.00	\$148.00
Intel® Xeon® Gold 5115	10	2.4	\$939.00	\$93.90

Workstation cost includes 2x1080Ti GPU, 256 GB RAM, 60 TB RAID6

Intel Chips 11/2017

CPU	Cores/ chip	Clock Speed (Ghz)	Cost/chip (list)	Cost/core	System Cost, 2 chip
Intel® Xeon® Platinum 8180	28	2.5	\$10009.00	\$357.46	\$35450.00
Intel® Xeon® Platinum 8168	24	2.7	\$5890.00	\$245.42	\$25946.00
Intel® Xeon® Gold 6148	20	2.4	\$3072.00	\$153.60	\$19442.00
Intel® Xeon® Gold 6154	18	3.0	\$3543.00	\$196.83	\$20528.00
Intel® Xeon® Gold 6132	14	2.6	\$2094.00	\$149.57	\$17224.00
Intel® Xeon® Gold 6126	12	3.2	\$1776.00	\$148.00	\$19936.00
Intel® Xeon® Gold 6126	12	2.6	\$1776.00	\$148.00	\$16452.00
Intel® Xeon® Gold 5115	10	2.4	\$939.00	\$93.90	\$15168.00

Workstation cost includes 2x1080Ti GPU, 256 GB RAM, 60 TB RAID6

Intel Chips 11/2017

CPU	Cores/ chip	Clock Speed (Ghz)	Cost/chip (list)	Cost/core	System Cost, 2 chip	Total Core Cost
Intel® Xeon® Platinum 8180	28	2.5	\$10009.00	\$357.46	\$35450.00	\$633.04
Intel® Xeon® Platinum 8168	24	2.7	\$5890.00	\$245.42	\$25946.00	\$540.54
Intel® Xeon® Gold 6148	20	2.4	\$3072.00	\$153.60	\$19442.00	\$486.05
Intel® Xeon® Gold 6154	18	3.0	\$3543.00	\$196.83	\$20528.00	\$570.22
Intel® Xeon® Gold 6132	14	2.6	\$2094.00	\$149.57	\$17224.00	\$615.14
Intel® Xeon® Gold 6126	12	3.2	\$1776.00	\$148.00	\$19936.00	\$830.67
Intel® Xeon® Gold 6126	12	2.6	\$1776.00	\$148.00	\$16452.00	\$685.50
Intel® Xeon® Gold 5115	10	2.4	\$939.00	\$93.90	\$15168.00	\$758.40

Workstation cost includes 2x1080Ti GPU, 256 GB RAM, 60 TB RAID6

Intel Chips 11/2017

CPU	Cores/ chip	Clock Speed (Ghz)	Cost/chip (list)	Cost/core	System Cost, 2 chip	Total Core Cost	Cost/ (core*Ghz)
Intel® Xeon® Platinum 8180	28	2.5	\$10009.00	\$357.46	\$35450.00	\$633.04	\$253.21
Intel® Xeon® Platinum 8168	24	2.7	\$5890.00	\$245.42	\$25946.00	\$540.54	\$200.20
Intel® Xeon® Gold 6148	20	2.4	\$3072.00	\$153.60	\$19442.00	\$486.05	\$202.52
Intel® Xeon® Gold 6154	18	3.0	\$3543.00	\$196.83	\$20528.00	\$570.22	\$190.07
Intel® Xeon® Gold 6132	14	2.6	\$2094.00	\$149.57	\$17224.00	\$615.14	\$236.59
Intel® Xeon® Gold 6126	12	3.2	\$1776.00	\$148.00	\$19936.00	\$830.67	\$259.58
Intel® Xeon® Gold 6126	12	2.6	\$1776.00	\$148.00	\$16452.00	\$685.50	\$263.65
Intel® Xeon® Gold 5115	10	2.4	\$939.00	\$93.90	\$15168.00	\$758.40	\$316.00

Workstation cost includes 2x1080Ti GPU, 256 GB RAM, 60 TB RAID6

GPU

- Most current scientific code uses CUDA which is NVIDIA only (GeForce, Quadro, Tesla)
- Gaming (GeForce) vs. Professional (Tesla) GPU
 - \$\$ Price \$\$
 - Less GPU RAM
 - Gaming chip double precision very slow
 - No error correction (RAM)
 - No NVLink, 1 DMA channel (slower transfer)

(V)olta

(P)ascal

	Tesla V100	Tesla P100	GTX 1080 Ti
Lithography	12nm FinFET	16nm FinFET	16nm FinFET
Transistor count	21bn	15bn	12bn
Die size	815mm ²	610mm ²	471mm ²
CUDA cores	5,120	3,584	3,584
Memory	16GB HBM2	16GB HBM2	11GB GDDR5X
Memory bandwidth	900GB/s	732GB/s	484GB/s
Double Precision (TFLOPS)	8	5	0.355
Single Precision	16 125 (mxmult)	11	11
Cost	???	\$5000	\$750

Emerging Technologies

- Phi
 - Intel many-core answer to GPU
- FPGA
 - Field Programmable Gate Array
 - Design your own chips! (Used in K2/3-summit)
- Nervana
 - Intel's emerging deep-learning focused chips


Storage Issues

- Quantity of Data (10 TB - 10 PB)
- Data Bandwidth
- Reliability/Redundancy
- Cost
- Tomography vs High Resolution Movies
- Central archives/databases

Quantity of Data

- 8x x 8k super-resolution counting movie, 30 frames
- 2 gigapixels/per movie
- typically only a few counts per pixel per frame
- 4 bits -> 1 GB movie (plus 256 MB periodically)
- 32 bits -> 8 GB movie
- Compression (slow, but saves a lot of space)
- Compressed counting mode TIFF movie, 0.1-0.2 GB !
- Krios+K2 -> 0.2 - 8 TB/day

Cost Over Time



On average storage cost falls 2x every 14 months !
Most enterprise drives have 5 year warranty

How Much Speed do You Need ?

- Xeon E5-2697v2
 - ~500 GFlops
 - ~200 GOps
 - 100 GB/sec memory bandwidth



- @100 MB/sec:
 - 20000 Fops/float

If a job processes 10 GB of data and takes 1 hour to run, should you worry about I/O speed ?

How about a job where processing 10 GB of data takes only 10 seconds ?


Interface Data Bandwidth

	Speed (GB/s)	Time to Transfer 4 TB
USB2	0.04	28 hours
Gigabit Network	0.1 (0.125)	11 hours
USB3	~0.3	3-4 hours
SATA	0.3, 0.6, 1.2	1-3 hours
10Gb Network	1.0 (1.25)	1 hour
Thunderbolt 2	2.0	30 min
Infiniband	1.0-4.0	15 min - 1 hour
PCIe 3.0	~1.0/channel ~16.0 max	~4 min

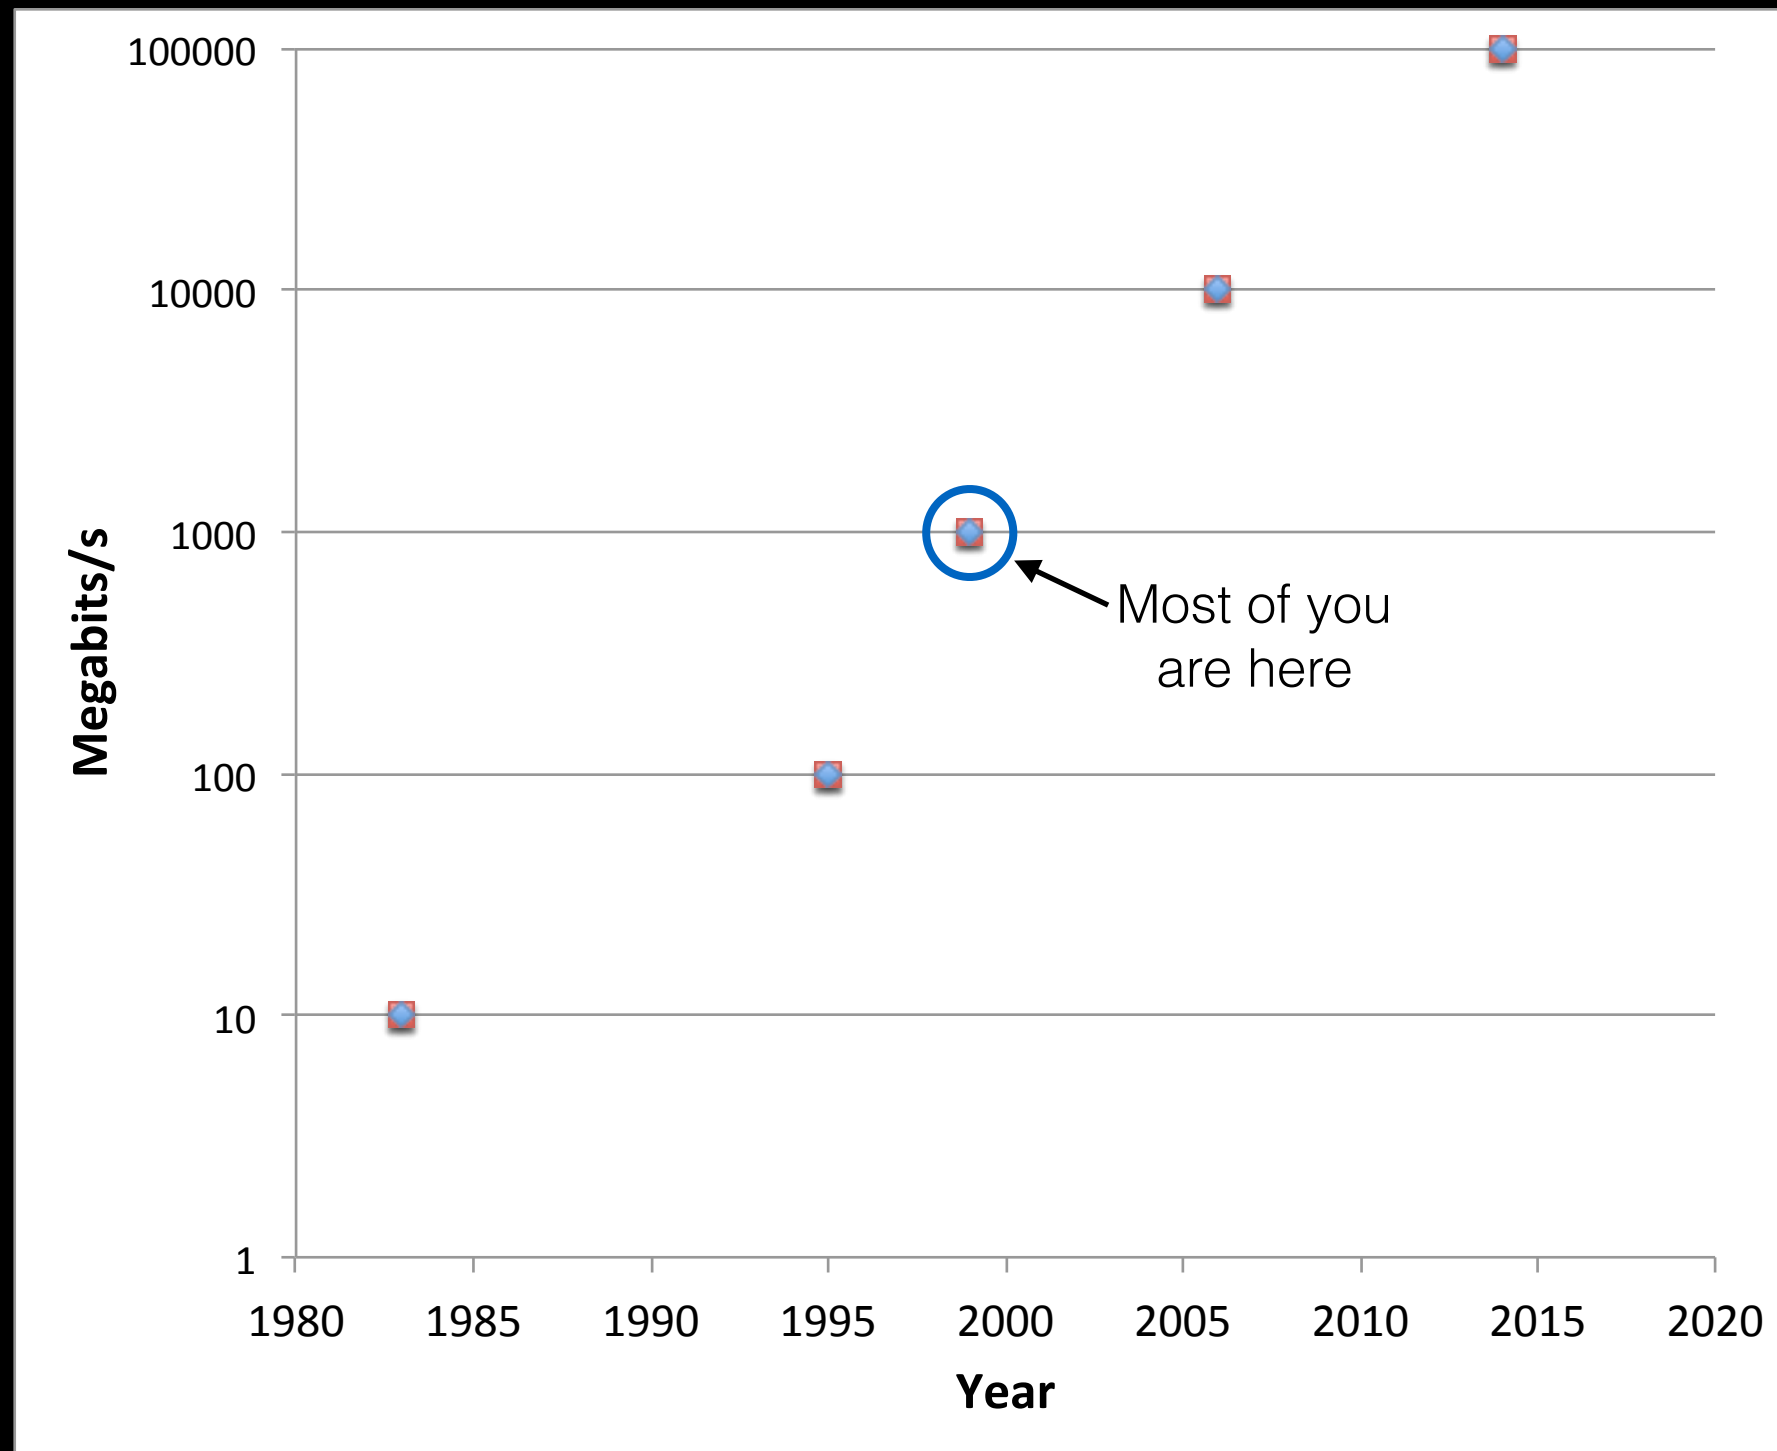
Drive Data Bandwidth

	Speed (GB/sec)	Max Size
2.5" Spinning Platter	0.06	4 TB
3.5" Spinning Platter	0.1-0.2	10TB
2.5" SSD	0.3-0.6	2 TB (16 TB)
RAID (striping)	1.0-1.5 Typ 3.0+ Possible	~80 TB/Array
PCIe/m.2 SSD (\$\$\$)	~2-3 Typ 4-6 Possible	~1 TB

Annual Failure Rates



Network Bandwidth



On average network bandwidth doubles every 27 months
Capacity doubles every 14 months!

Practical Tips

- Gigabit networking -> 125 MB/s absolute max
 - This is slower than a typical hard drive!
- rsync is usually limited by ssh (often 40-50 MB/s)
- NFS has very poor performance unless tuned
- Do not use remote filesystems unless you have no choice (or have good networking, like clusters)
- Do not use external hard drives for processing! You may not regret it today, but you will eventually.
- If you are setting up a big storage system make sure it has adequate (10 Gb) networking

"Standard" Workstation

Workstation with 8-bay chassis + PCIe RAID controller

\$18k = \$8k CPU + \$3k RAM + \$2k GPUs + \$3k Storage + \$2k System

- 40 cores @ 2.4 Ghz
- 2x 1080Ti GPU + Quadro 420 for graphics
- 60 TB RAID6 storage, ~1.5 GB/s

Options:

- 8 core, 2.6 Ghz -> -\$6k
- 96 GB RAM -> -\$2k
- 24 TB Storage -> -\$1.5k



100 TB - an Example

12 Bay Synology - \$1400

12x 10TB Drives - \$4200

RAID6 -> ~100TB

Cost w 5 year warranty:

~\$5600 —> \$1.52/TB-month

~0.1 GB/sec (network limited)

Advantages: Reliable, Easy, Quiet, Cheap

Disadvantages: A bit slow



640 TB - an Example

1x4U computer (\$6500) + 36x 10TB drives (\$12,600)+
1x4U JBOD Chassis (\$3600) + 44x 10TB drives (\$15,400)
Configured as 8x RAID6 volumes → 640 TB usable
~1.5 GB/sec I/O to the attached computer

Cost w 3/5 year warranty ~\$38k → \$0.001/GB-month
x5 → 3.2 PB/rack (usable)



Advantages: Inexpensive, Fast, Includes Computing
Disadvantages: Management, Housing/Noise

Cloud Storage ?

Amazon (S3):

- Standard Storage: \$0.023/GB-month (23x local raid cost)
- Infrequent Access: \$0.0125/GB-month
- Glacier Storage (backup): \$0.004/B-month

+

Download cost:

- \$50/TB (unless you have a waiver)

Advantages: Safe & Reliable, Access to EC2

Disadvantages: Slow Access, Expensive, Legal Issues