What Computer Should I Buy (and maintain) with $ XXX

Context

- CIBR center at BCM, ~$100k/year + PI contrib
- 50-60 users, most infrequent
- ~2500 cores in 5 clusters (1 wGPU), ~1800 TB of storage
- ~70-80% load level

- 20k Ribosome particles -> 4.7 Å,
  - ~1000 CPU-hr in EMAN2, ~10,000 CPU-hr Relion
  - Heterogeneity, large data sets, 10-100x more
  - NCMI (~35 people) uses ~5M CPU-hr/year (mostly EMAN)

Considerations

- CPU Choice (speed, cores/node, GPU/ Phi)
  - $250-500/core in cluster, $300 typical, 12-24 cores/node
- Amount of RAM
  - 64 GB - $600, 128 GB - $1200, 256 GB $4600, per node
- Interconnect (network)
  - 1 Gb, 10 Gb, Infiniband (QDR - 8Gb, FDR - 14Gb)
- Storage System
  - Central RAID(s), Distributed (Lustre), Backup?
- Amount of Storage

Get a Good Workstation

- E5-2690v3 x2 -> 24 cores, 2.6 Ghz ($4000) or
- E5-2640v3 x2 -> 16 cores, 2.6 Ghz ($1800)
- 128 GB RAM -> $1200
- 2 processor motherboard (FCLGA2011) -> $400
- Case with 8-hot swap bays -> $450
- LSI MegaRAID SAS 9271-8i -> $700
- 8x 6tb SATA (speed) -> $2400 (36 TB usable R6)
- NVIDIA GTX980 -> $600 (or cheaper)

$9,750 total (could easily be scaled down)

~200,000 CPU-hr/year

‘Typical’ Compute Node

- 2U Compute Chassis: $36,000 (list)
- 2U -> 4 nodes -> 8 CPUs -> 96 cores
- E5-2690v3 (12 cores, 2.6 Ghz) x2/node
- 128 GB/node
- FDR infiniband (14 Gb)
- 4 TB local scratch (or Lustre) drive
- 2 kW Power supply (~1 kW typ)

‘Typical’ Head Node

- $38,000 (list)
- 2x E5-2690v3 (24 cores 2.6 Ghz)
- 256 GB RAM
- 36x6 TB -> 9 dr RAID6x4 -> 168 TB usable
- Switches, etc ~$30-40k

1 Rack cluster

- 44U standard - ~$750,000 (list)
- ~15M CPU-hr/yr -> ~$0.014/CPU-hr (70% usg, 5yr)
- 4U Head/storage node
- 4U Switches, etc
- 18 x 2U Compute Nodes -> 1728 cores
- ~20 kW actual draw, ~40 kW in planning
- ~$30,000 - 40,000/year in power/cooling

Other Options

- Intel vs. AMD?
  - AMD more cores/$, but cores (much) worse
- NVidia Tesla? Intel Phi?
- Infiniband switches limited to 44 nodes, poor scalability
- The Cloud -> $0.08 -$0.12/CPU-hr

378 TB - an Example

1x4U computer with 36x 6TB drives +
1x4U 45x 6TB drives JBOD* Chassis
Configured as 9x RAID6 volumes —> 378 TB
~1.5 GB/sec I/O to the attached computer

Cost w 3 year warranty ~$36k —> $0.0026/GB-month
x5 —> 1.9 PB/rack (usable)

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

* - JBOD = Just a Bunch of Disks

Cloud Storage ?

Amazon (S3):

- Standard Storage:
  - 1 PB - $0.055/GB-month
  - 1 TB - $0.085/GB-month
- Reduced Redundancy:
  - 1 PB - $0.044/GB-month
  - 1 TB - $0.068/GB-month
- Glacier Storage (backup):
  - $0.01/GB-month

+

Download cost:

- $0.05 - $0.12 /GB

Advantages: Safe & Reliable, Access to EC2

Disadvantages: Slow Access, Expensive, Legal Issues