FY 2007 cluster at Jefferson Lab

Chip Watson

May 14, 2007

FY 2007 Cluster

Goals:

Formal goal: procure 2.9 Teraflop/s sustained <asqtad,dwf>,
& deploy by June 30

Science goal: get the most capacity that $1.4M can buy,
& deploy as fast as possible

Procurement Process

Best Value RFP Process

- Explicitly described a cluster of ~800 processors with infiniband fabric (to guide vendors toward a good solution)
- Allowed for ANYTHING
  (specifically wanted to allow room for a BG/L proposal to compete)
- Included anisotropic clover as one of three benchmarks
- Chose local volumes corresponding to anticipated real jobs
  (not artificial “best performance” numbers)

Single Node Performance as key criteria

- Extrapolate to multi-node based upon past experience, and simple spreadsheet model (for typical clusters, achieve ~80% scaling single node to multi-node on real jobs)

Proposals

Single node performance, showing breadth of potential solutions:

action: local vol:	asqtad 12^4	clover 12x6x6x32	dwf 28x8x8x32	bandwidth per core	<as,cl,dwf> $/MF
1 dual core Xeon 2.66	3530	2487	4800	1500	$0.51
1 quad core Xeon 2.33	2540	3520	4400	715 - 900	$0.78
2 dual core Xeon 2.66	4630	4800	7491	1325	$0.56
2 quad core Xeon 233	4200	12000	5600	400	$0.51
2 dual core AMD 2.6	4900	4040	6560	1750	$0.49
amd w/ 1 GB dimms (not 512)	5140	4400	6950	1975	

Proposals

Single node performance, showing breadth of potential solutions:

action: local vol:	asqtad 12^4	clover 12x6x6x32	dwf 28x8x8x32	bandwidth per core	<as,cl,dwf> $/MF
1 dual core Xeon 2.66	3530	2487	4800	1500	$0.51
1 quad co Xeon 2.3					$0.78
2 dual co Xeon 2.66					$0.56
2 quad core Xeon 233	4200	12000	5600	400	$0.51
2 dual core AMD 2.6	4900	4040	6560	1750	$0.49
amd w/ 1 GB dimms (not 512)	5140	4400	6950	1975	

Single node price/performance needed to be $0.40 to meet project goals. All proposals fell short. (good news 3 slides later)

Proposals

Single node performance, showing breadth of potential solutions:

action: local vol:	asqtad 12^4	clover 12x6x6x32	dwf 28x8x8x32	bandwidth per core	<as,cl,dwf> $/MF
1 dual core Xeon 2.66	3530	2487	4800	1500	$0.51
1 quad c Xeon 2.3				0	$0.78
2 dual c Xeon 2.6					$0.56
2 quad core Xeon 233	4200	12000	5600	400	$0.51
2 dual core AMD 2.6	4900	4040	6560	1750	$0.49
amd w/ 1 GB dimms (not 512)	5140	4400	6950	1975	

We discovered an unexpected improvement in performance when using 1 GB dimms instead of 512 MB dimms (both single rank). Note that the problems did not overfill memory.

Selection Criteria

- Largest weight went to $/MFlops
- Minimum requirements included vendor past performance, and a proposal that was responsive to the guidance.
- For proposals within a few percent in $/MFlops, lower weighted factors became important:
  - physics flexibility
  - installation, operation, serviceability

- The winning vendor scored well on all criteria...

Winning Proposal

Vendor: Koi ("whitebox", same supplier as Kaon at FNAL)

Nodes:

- dual cpu, dual core AMD 2218, 2.6 GHz
- DDR (20g) infiniband, 18:6 leaf switch oversubscription
- ASUS KFN4 motherboard; IPMI 2.0

Interesting Upgrade Option:

Upgrade to quad core Opterons at steeply discounted price

- doubles number of cores
- doubles SSE issue rate / cycle (to match Intel)
- 2 MB L3 shared cache, + ½MB L2/core (effectively doubles cache)
- same power envelope (2.1 GHz vs 2.6 GHz)

Projecting Quad Core Performance

Reasoning:

- dual - dual core Xeons get the most flop/s per MB/sec of memory bandwidth (streams triad) – i.e. Xeons have enough peak flop/s to consume bandwidth
- Raw flops of quads will be 3x faster than Opteron duals
  (2x cores, 2x issue rate, ¾ clock speed)
- With this increase, Opteron will also be memory bound, like the Xeons
- Opterons have 50% - 60% more memory bandwidth than Xeons
- Therefore, quad cores should have ~ 50% performance boost
  20% cost, and delay of ~3 months on just 20% of funds (clear win)
- should achieve goal of <$0.40/MF single node,
  > 3 TFlop/s aggregate sustained performance

In addition:

- Hyper transport bus & Opteron cache protocols are better at multi-threading, yielding additional architectural advantages as code is converted to threads

Modifying the proposal

- Preliminary decision to hold back 20% of the funds to be able to do the quad core upgrades (only path to meeting milestone)

  (no real impact, since continuing resolution kept those funds from Jlab anyway)

  - increase memory / node to 4 GBytes
    (higher density yields better performance for some reason)
  - decrease node count to 399
    (22 switches w/ 18 nodes = 396; 2 interactive, 1 for R&D)
- In May, planned to evaluate quad core chips in one of our nodes, verify that we get price/performance boost
- Order quads (OR order additional 20% nodes if quads are poor)
- Quad chips might take 3 months to receive (high initial demand); install as late as September

Preliminary Results

Status:

- Testing on 1.8 GHz Barcelona yielded LQCD inverter results within 10% of memory bandwidth based projections (60% speedup)
- If the 2.1 GHz performs slightly better (memory bandwidth scales somewhat with clock speed), then we should meet or exceed the 2.9 TFlop/s deployment goal.

We will have obtained the best of both worlds:

- deploy early: 1.8 TFlops as early in the year as funding allowed (more integrated running in FY 2007)
- deploy late: will still achieve ~2.9 by the end of the year based upon latest technology

7n Timeline

April:

machine installation (mid month for racks, later for short cables)

May:

friendly user mode on ~100 dual-duals (using long cables)

separate PBS server & queue, for 64 bit O/S (other systems are 32 bit)

2 of the nodes used as 64 bit interactive & build/test with 8 GB / node

June:

ramp up to production on 396 dual-duals

September:

rolling outages to upgrade to quads

Anticipated Result: 2.9 TFlop/s deployed

optimistically, will be in FY 2007

(may come a bit late, based upon quad core availability)

7n Cluster

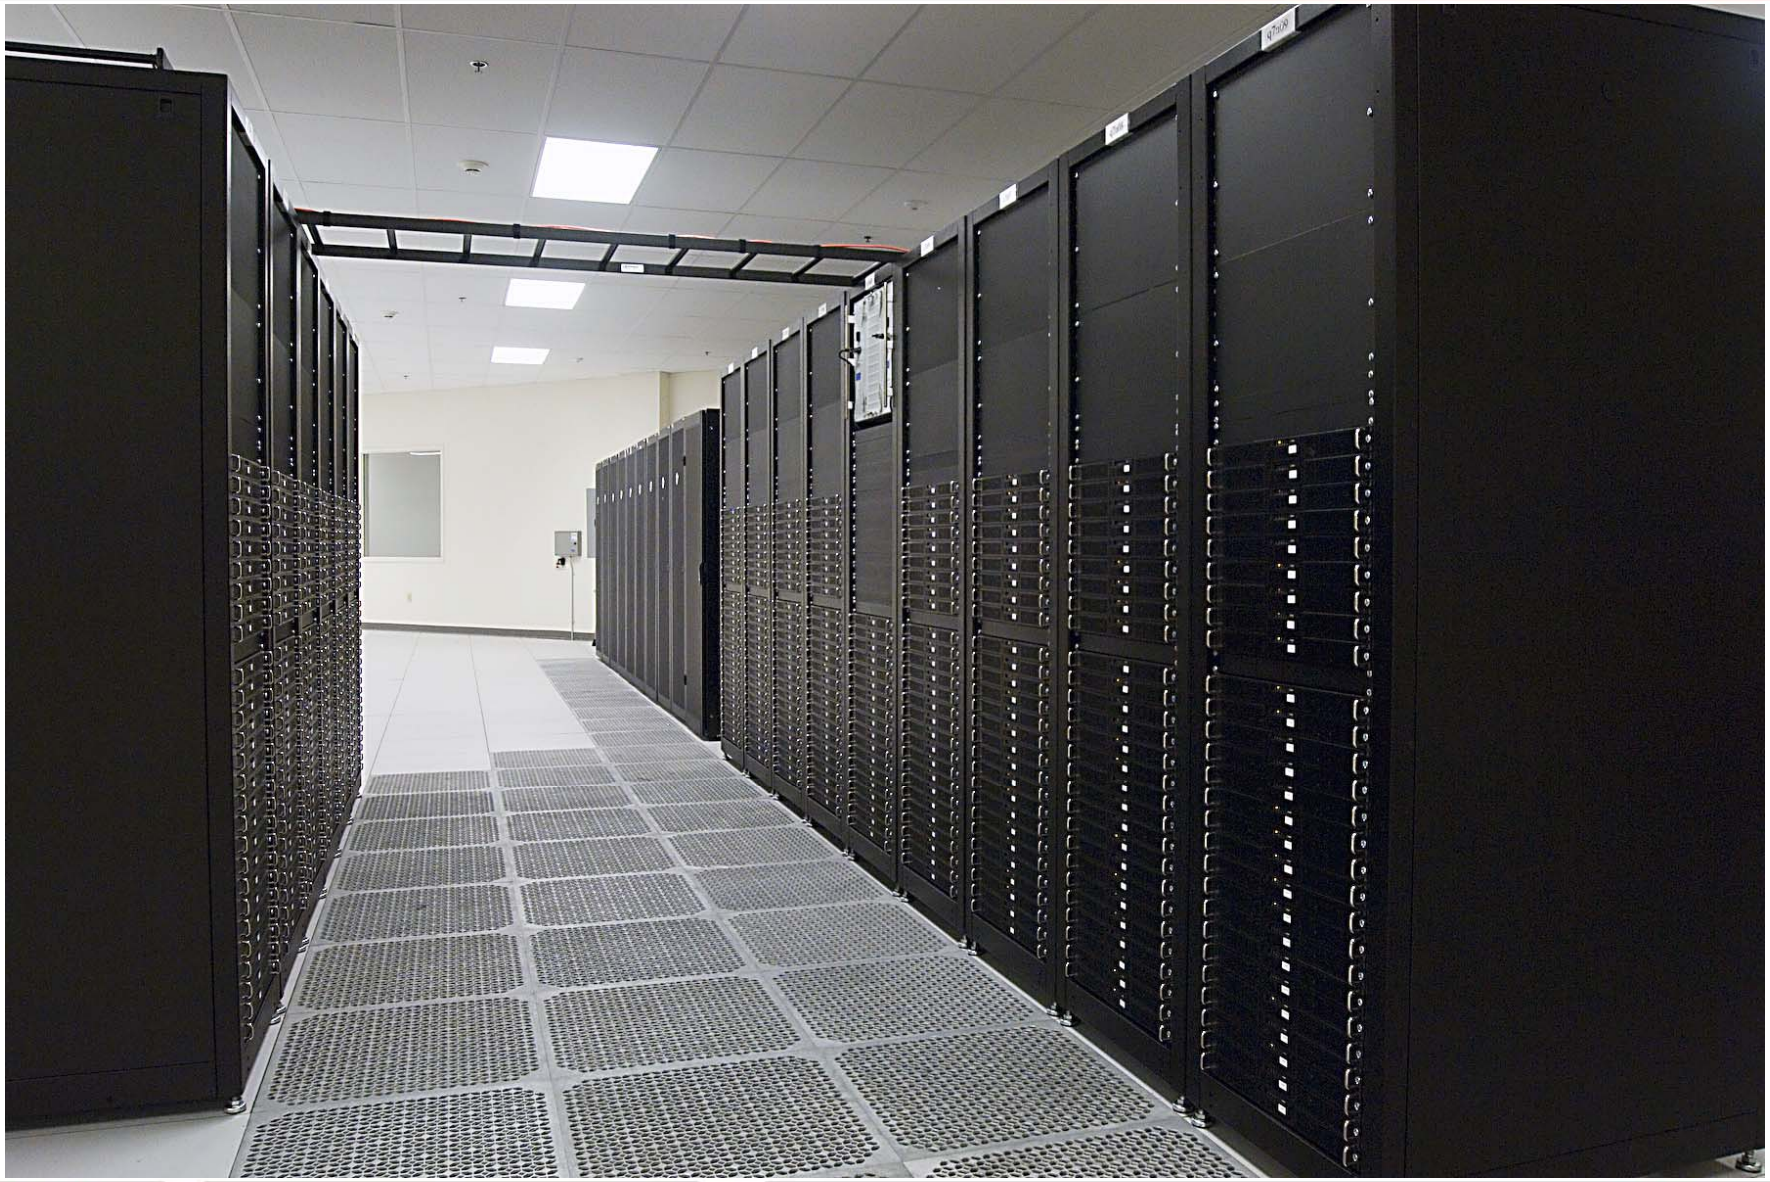

File Server Upgrade

Previously had 5 servers from 1 to 4 years old

total of 15 TBytes, but reliability decreasing.

JLab compute capacity increasing from ~1 to ~4 TFlop/s

thus should roughly quadruple disk capacity to ~60 TBytes total, keeping only 4 TB of the latest server for production use.

Other goals:

- high single stream bandwidth
- ability to handle spikes in bandwidth to/from silo

New Configuration

3 Sun x4500 Servers

- JLab caught Sun's 25th anniversary sale, saving 50% on the x4500, and getting $1,000/terabyte raw (same as cheap raid boxes)
- single stream to Infiniband 120 MB/s/server (2 streams 180 MB/s)
- bandwidth (aggregate) to silo > 200 MB/s/server
- largest partition (single project) up to 18 Terabytes

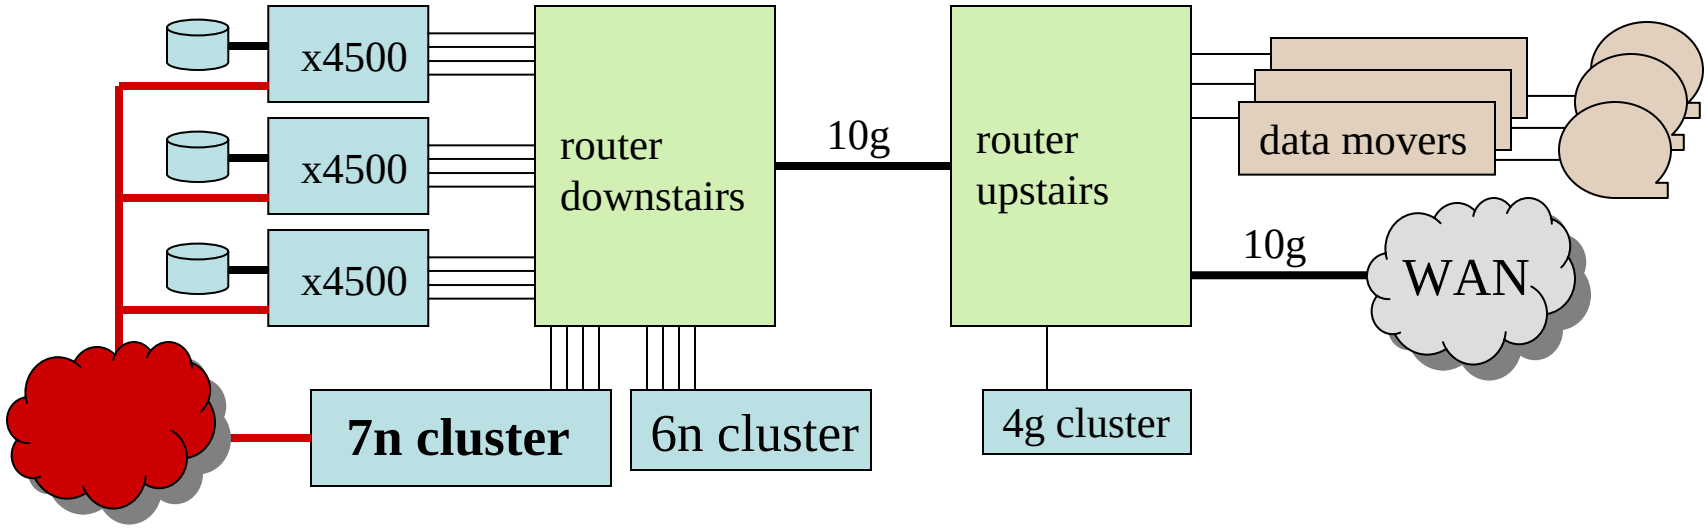

Infrastructure Upgrade Summary

Project funded upgrades:

- **Disk Cache**
  - 36 TB (2*18)
  - (another 18 TB base funded; new total of 58 TBytes)

Base funded upgrades:

- **Wide Area Networking**
  - Upgraded this past year to **10g**
- **Local Area Networking**
  - Bandwidth: file server to silo going to **10g**
- **Planned power upgrades**
  - Over next 2 years add 1 Megawatt UPS
  - Over next 5 years add equivalent cooling capacity

Other JLab Clusters

3g 2003 gigE mesh
2.66 GHz P4, 256 MB / node
½ decommissioned,
now just 128 nodes
no allocation this coming year

4g 2004 gigE mesh
2.8 GHz P4, 512 MB/node
384 nodes, 3 sets of 128
start to decommission this year

6n 2006 infiniband
3.0 GHz Pentium-D
1 GB/node
280+ nodes

(Q & A)