LQCD-ext Technical Performance of FY11 Deployments

Don Holmgren

Fermilab

djholm@fnal.gov

SC LQCD-ext Annual Progress Review
Brookhaven National Laboratory

May 16-17, 2012

Outline

- Overview of SC LQCD-ext acquisitions
- Computational Requirements
- FY11 Budget Delays
- FY11 conventional cluster deployment and performance
- FY11 GPU-accelerated cluster deployment and performance

Overview of SC LQCD-ext Acquisitions

- Plan on approximately five acquisitions
  - Usually one per year in FY10-FY14, but some years will have both conventional and GPU-accelerated cluster purchases
  - FY10 and FY11 conventional cluster buy was “across” the fiscal year boundary so that we could employ a single contract
- Guiding principle: procure the systems that will be the most effective for the planned science, given the portfolio of operating SC LQCD-ext and other machines at that time
  - FY10/FY11 – we have deployed a commodity cluster, purchased across FY10/FY11, and in FY11 we purchased a GPU-accelerated cluster
  - FY12 – Combination of commodity cluster and GPU-accelerated cluster (see next talk)
  - FY13/FY14 – perhaps BG/Q and/or a combination of commodity and accelerated clusters

Overview of SC LQCD-ext Acquisitions

Computational capacity goals by year for SC LQCD-ext:

	FY2010	FY2011	FY2012	FY2013	FY2014
Computing hardware budget (excluding storage)	$1.60M	$1.69M	$1.875M	$2.46M	$2.26M
Capacity of new cluster deployments, TFlop/s Planned/Revised/Achieved	11 / 12.5	12 / 9 / 9	24 / 10-15	44 / 15-22	57 / 22-33
Million "Fermi" GPU-Hrs/Yr Planned/Revised/Achieved	0	0 / 1.02 / 1.22	0 / 2.9-4.3	0 / 4.6-6.9	0 / 7.5-11.2

- Baseline computing hardware budgets are shown
- FY2011 original plan for 12 Tflop/s was changed to 9 Tflop/s plus a GPU-accelerated cluster with 128 nVidia "Fermi" GPUs released to production in FY2012 (152 achieved)
- FY2012-FY2014 revised goals reflect 40%-60% ranges in budget allocated to conventional and accelerated clusters
- FY2012-FY2014 GPU-Hrs/Yr figures are based on FY11-model GPUs (NVIDIA "Fermi")
  - New GPU models will deliver more than 1 "Fermi" hour per wall-clock hour

Computational Requirements

- Either memory bandwidth, floating point performance, or network performance (bandwidth at message sizes used) will be the limit on performance on a given parallel machine

- On single commodity nodes memory bandwidth in the constraint that limits performance

  - GPUs deliver more memory bandwidth per dollar than conventional CPU's, but can only be used for some of our calculations

- On current parallel computer clusters, the constraint is either memory bandwidth or network performance, depending upon how many nodes are used on a given job

  - Network performance limits strong scaling:
    Surface area to volume ratio increases as more nodes are used, causing relatively more communications and smaller messages

  - GPUs require higher network bandwidth than CPUs

Computational Requirements

- We design and/or buy systems that as part of our hardware portfolio will most effectively carry out the current and anticipated scientific programs
- This means:
  - Systems matched to the type and size of LQCD calculations that will be performed
  - Systems with the best price/performance for LQCD applications
  - Machines with the best memory bandwidth
  - High performance interconnects
  - Networks balanced to single node capacities and anticipated job sizes

Typical LQCD Cluster Layout

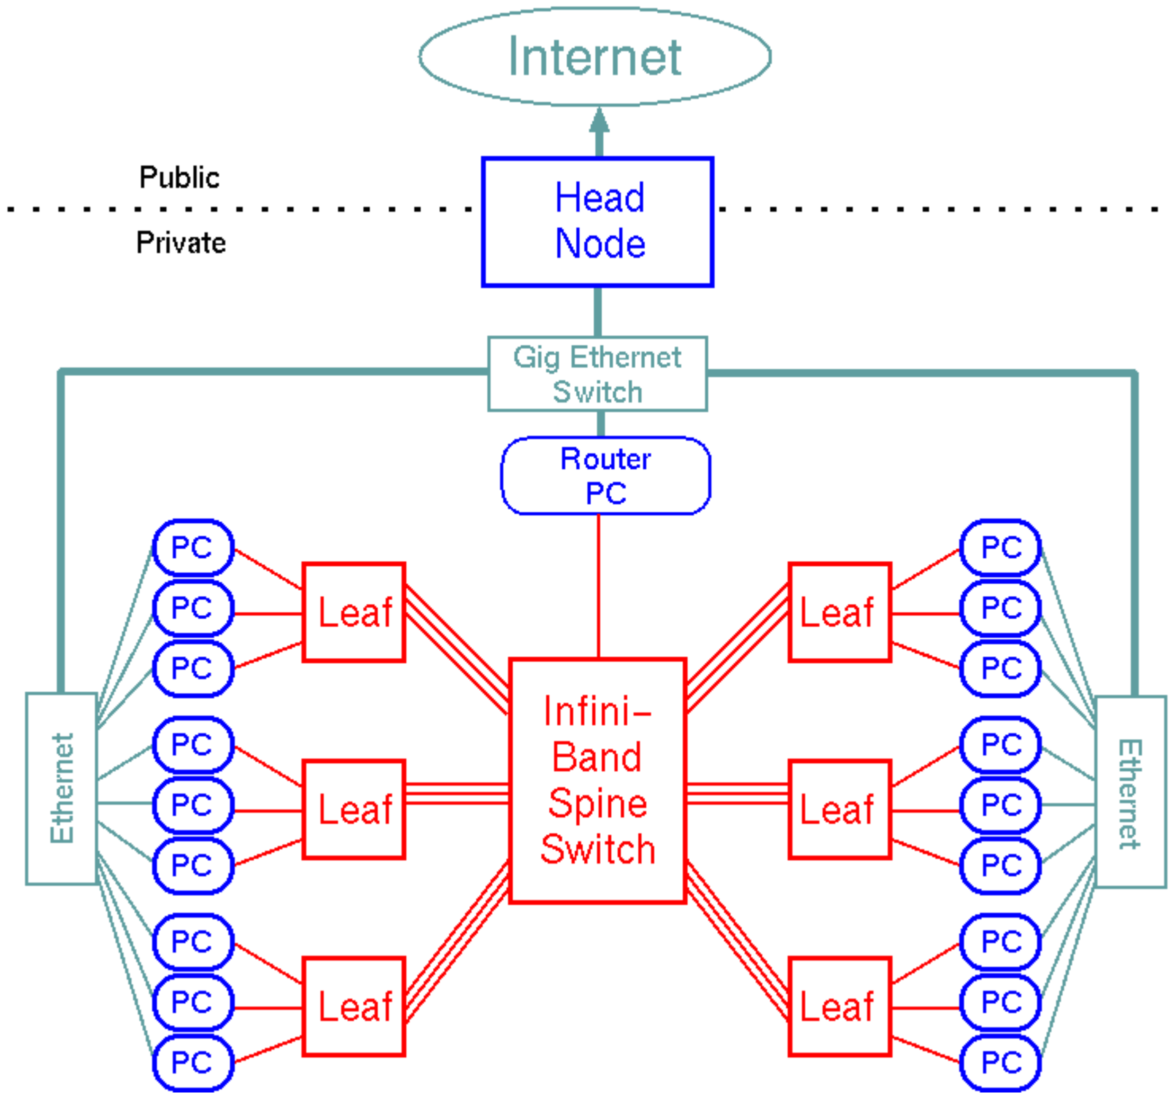

Typical LQCD Cluster Layout

Each PC may contain zero or multiple (2-4) GPUs

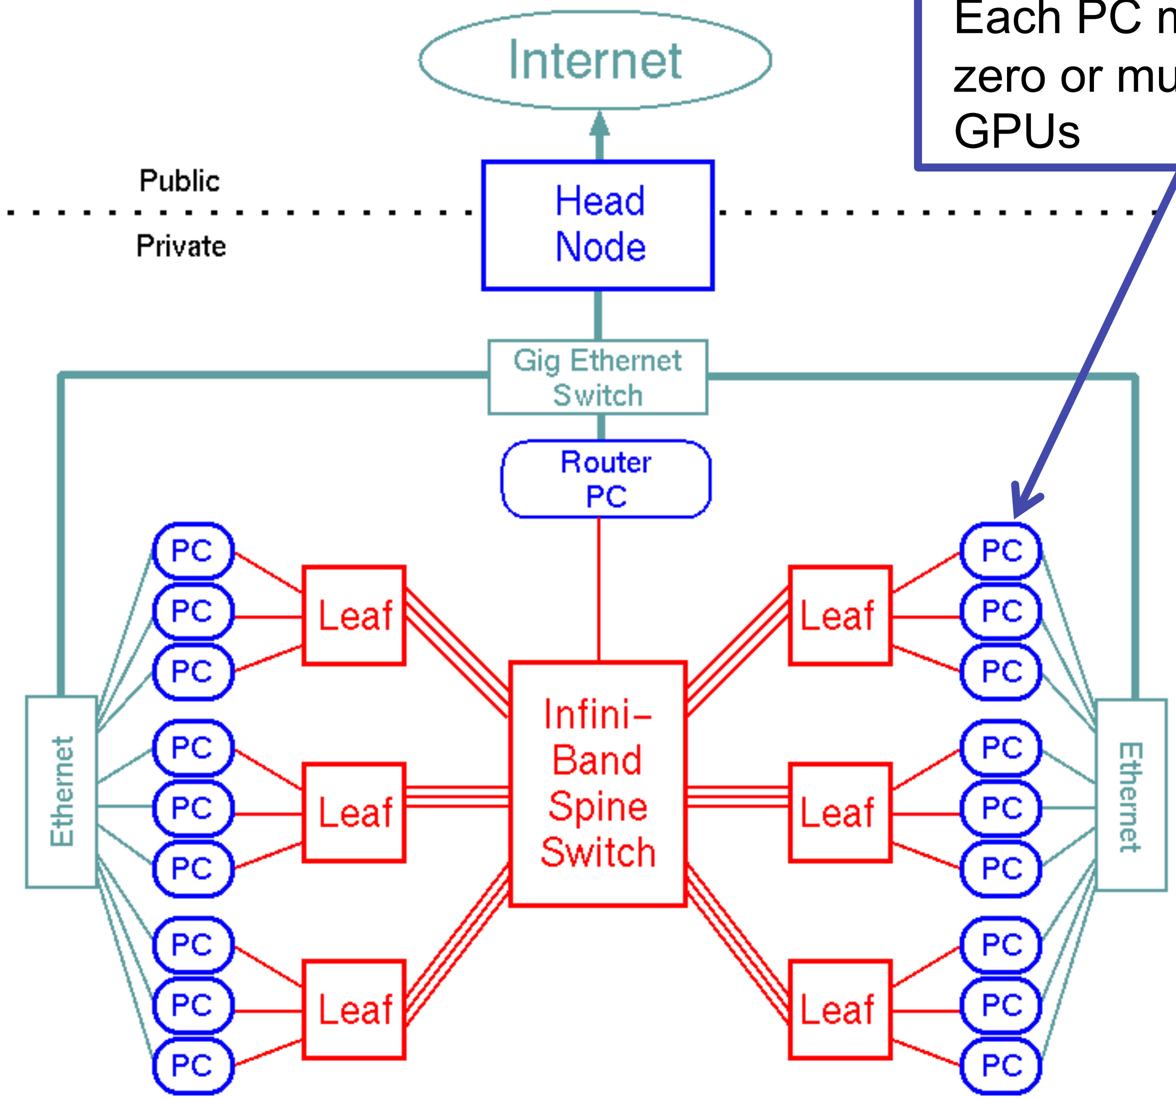

Typical LQCD Cluster Layout

Lustre parallel filesystem accessed through Infiniband

Each PC may contain zero or multiple (2-4) GPUs

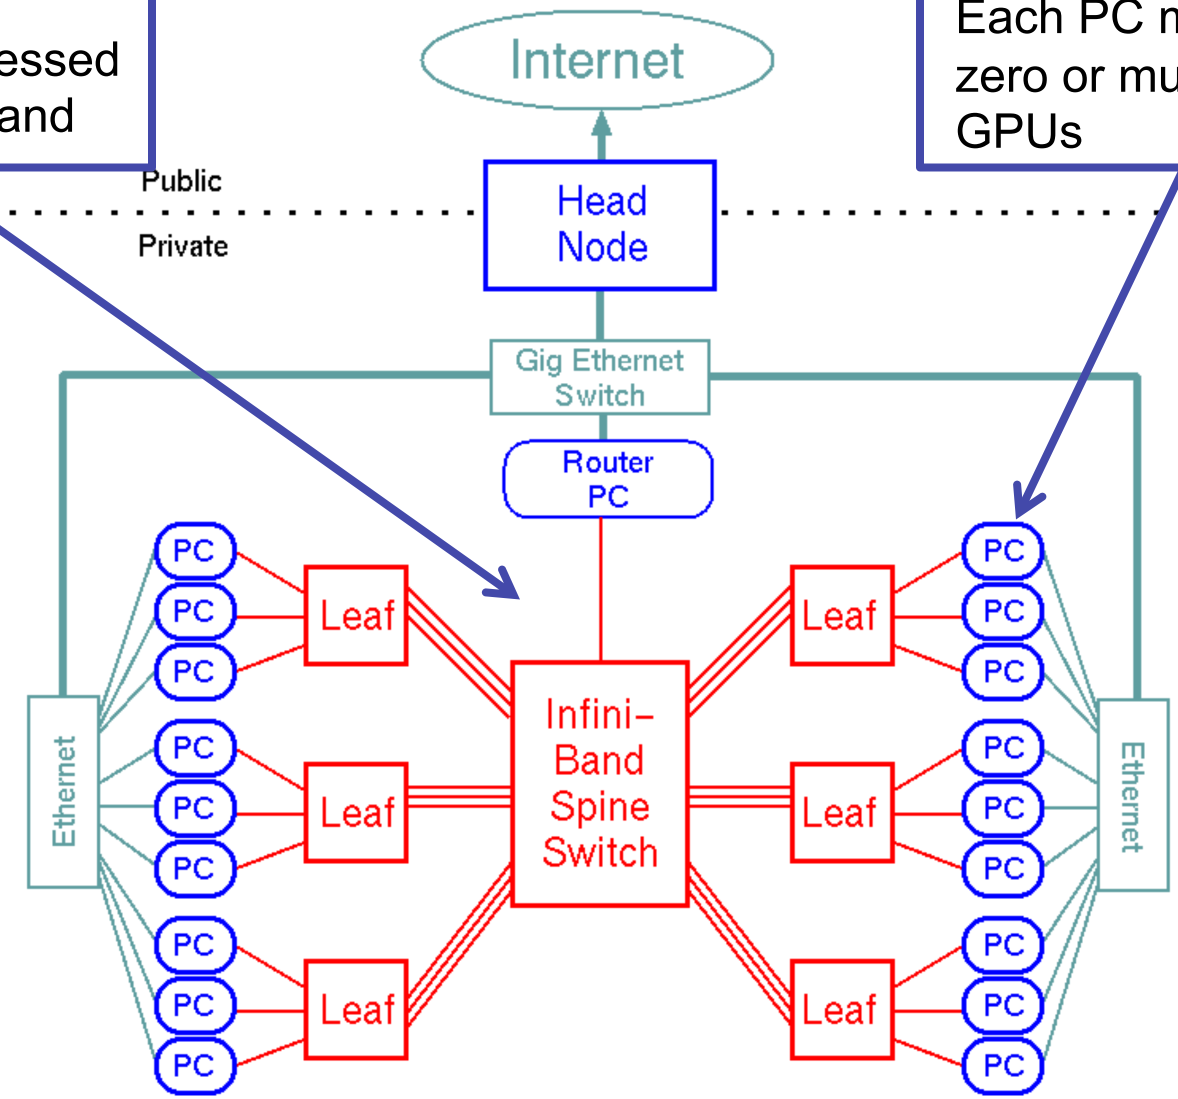

Rating LQCD-ext Computing Facilities

- Definition of the sustained capacity of LQCD-ext computing hardware:
  - The performance of the improved staggered (“asqtad”) and domain wall fermion (“DWF”) conjugate gradient inverters are measured using parallel jobs spanning a significant number of processors (128 cores on clusters)
  - The average of the asqtad and DWF values (per core) multiplied by the number of available cores gives the defined sustained Tflop/s capacity
- Although the inverter is only part of the computing load, and other actions besides DWF and asqtad are used, on clusters and leadership machines the asqtad-DWF average has been predictive of overall computing throughput
- The asqtad-DWF average is not known to be predictive for GPUs
  - Neither DWF or all of asqtad/HISQ are in production
  - For some job types, execution times are not dominated by the inverter
  - Currently requests and allocations for GPU resources are in “GPU-hours”
  - We track “cost-equivalent” GPU capacities by comparing performance and costs of actual jobs run on both conventional and GPU hardware

FY11 Budget Delays

- Continuing budget resolutions affected FY11 purchases
  - 8 total C.R.'s, with the final thru end FY11 signed into law 4/15/2011
  - As of May 2011 LQCD-ext review, final FY11 budget guidance had not been received by Fermilab
- Fermilab throttled the rate of all spending in response
  - By February 2011, enough funds for the project to requisition half of the planned conventional cluster expansion had accrued
  - Lab approval for first half of conventional expansion was delayed until March 2, 2011, the second half until May 20, 2011 and the GPU cluster until June 20, 2011
  - Lab-wide, all Fermilab purchases were affected, and the resulting deep queue of requisitions further delayed the GPU cluster purchasing process
  - Consequences were that half of the conventional cluster ("Ds") expansion, and the GPU-accelerated cluster, did not meet the project baseline schedule milestone, and the Scientific Program Committee adjusted available computational capacities in their 2011-2012 award process

FY11 Conventional Cluster Deployment and Performance

Ds Details

- Award was to best value bid, based on price, LQCD application performance, power efficiency, space efficiency, vendor qualifications and past performance
- Hardware details:
  - Quad-socket eight-core AMD 2.0 GHz “Magny-Cours” processors
  - 64 Gbytes memory per node
  - QDR Infiniband with 2:1 oversubscription
  - 421 worker nodes, plus head nodes
  - $2.55M including G&A ($2.45M for worker nodes and Infiniband)
- Performance
  - Asqtad:DWF 51 Gflop/node (128-process MPI runs)
  - 21.50 Tflop/s → $0.114/Mflop

Cost and Performance Basis

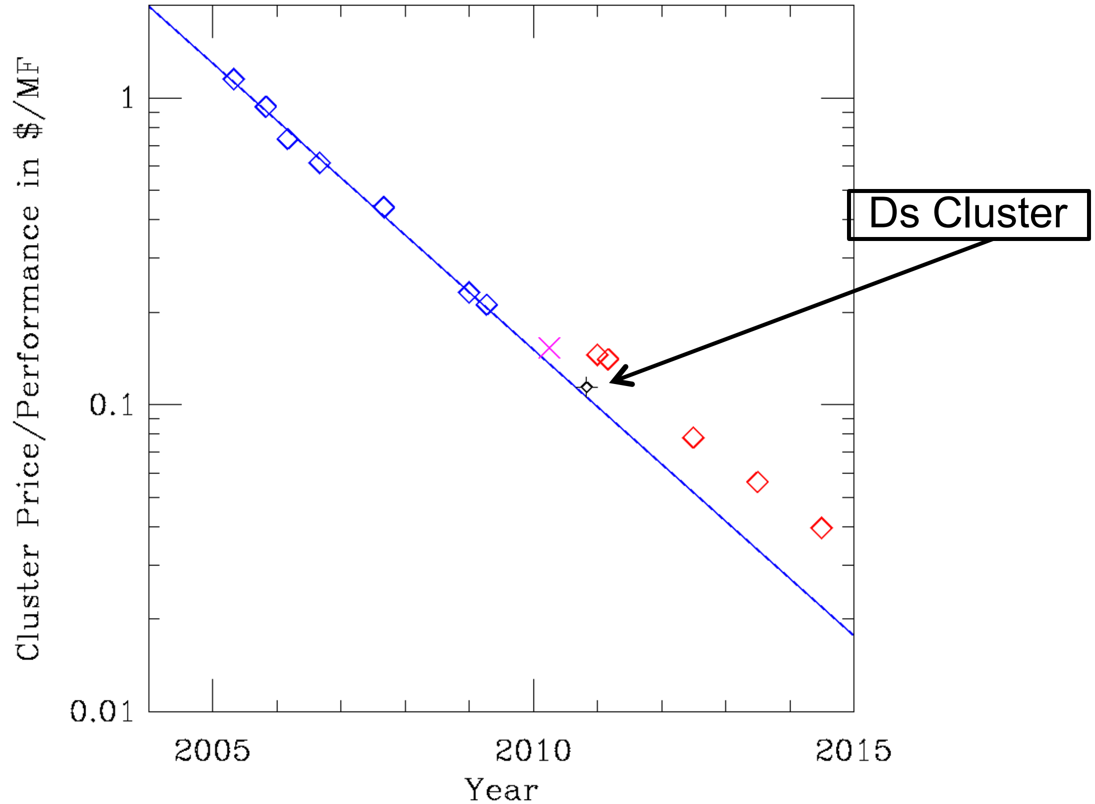

Cluster	Price per Node	Performance/Node, MF	Price/Performance
Pion #1	$1910	1660	$1.15/MF
Pion #2	$1554	1660	$0.94/MF
6n	$1785	2430	$0.74/MF
Kaon	$2617	4260	$0.61/MF
7n	$3320	7550	$0.44/MF
J/Psi #1	$2274	9810	$0.23/MF
J/Psi #2	$2082	9810	$0.21/MF
10q	$3461	22667	$0.15/MF
Ds	$5810	50810	$0.114/MF

Cost and Performance Basis

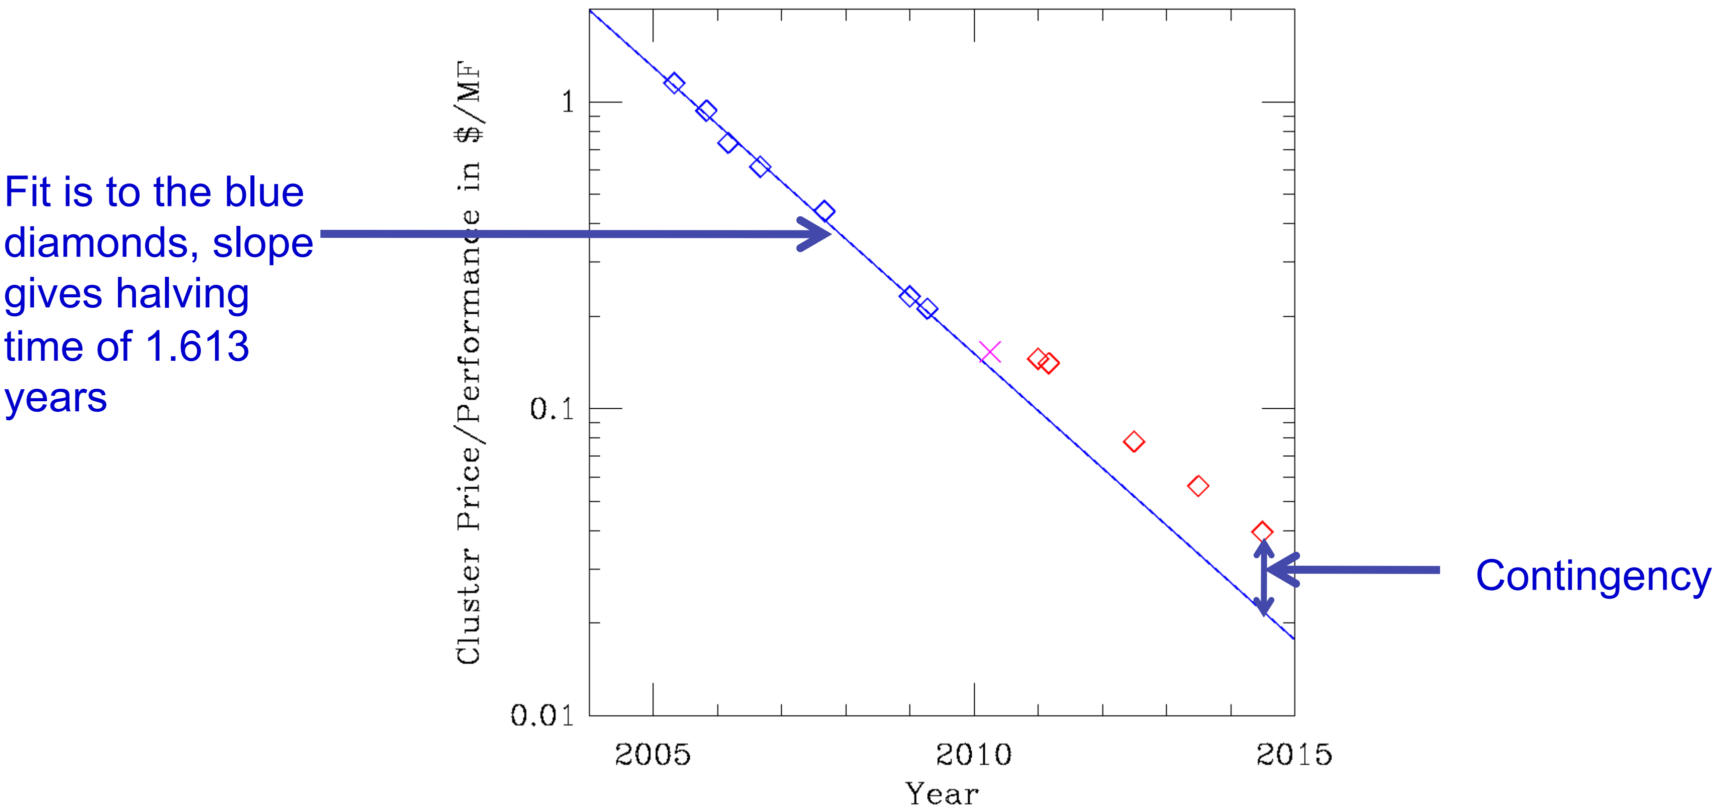

Year	Deploy Date	Price/Perf. Goal	Price/Perf. Trend	Goal (TF)	Contingency (TF)	Contingency (TF %)
2010	2011.0	$0.15/MF	$0.098/MF	11	4.4	40%
2011	2011.2	$0.14/MF	$0.098/MF	12	4.4	36%
2012	2012.5	$0.078/MF	$0.052/MF	24	11.9	50%
2013	2013.5	$0.056/MF	$0.034/MF	44	26.8	61%
2014	2014.5	$0.040/MF	$0.022/MF	57	42.6	75%

The FY11 Ds Procurement
(Planned as of May 2011 Review/Achieved)

2011

- Jan 19 – Advice from the Executive Committee on the split between Ds expansion and a separate Ds cluster
- Mar 7 – Purchase order to vendor for first half of the expansion
- *Mar 31 – Baseline release to production*
- Apr 15 – FY11 Full-year C.R. signed into law by the President
- May 4 – Delivery of nodes
- May 16 – Delivery of remaining equipment for first half
- May 31 (May 26) – Purchase order to vendor for second half of the expansion
- Jun 1 (Jun 6) – Release to production of first half (4.5 TFlops)
- Aug 15 (Jun 30) – Delivery of second half
- Sep 1 (Aug 1) – Release to production of second half (4.5 TFlops)

FY11 GPU-Accelerated Cluster Deployment and Performance

FY11 GPU Cluster ("Dsg") Design

- In consultation with SciDAC software committee, project identified these requirements:
  - GPU count high enough to do large-scale configuration generation: jobs of at least 64 GPUs, and as high as 128; this would require software for cutting along multiple dimensions
  - Sufficient PCIe and Infiniband bandwidth to support these jobs
  - NVIDIA Tesla instead of cheaper graphics cards, because:
    - ECC memory capability, necessary for non-inverter code
    - Warranty issues with non-Tesla hardware that have errors in numerical calculations but pass graphics tests
    - Direct GPU to GPU, and GPU to IB, communications only to be supported on Tesla
    - Larger memory per GPU (3 or 6 GB, vs 1.5 GB for graphics cards)

GPU Cluster RFP

- Vendors were allowed to bid a wide variety of configurations:
  - Hosts could be Intel or AMD based, with 2 GB memory per core, housing a total of 128 GPUs
  - NVIDIA GPUs could be any of Tesla x2050, x2070, x2075, or x2090
  - Hosts servers could house 2, 3, or 4 GPUs, with one Infiniband QDR channel for every 2 GPUs (rounded up, so 3 GPUs required dual QDR)
  - Required options included doubling host memory and a higher GPU count
- Proposals were received from 9 vendors (2 OEM, 7 "whitebox"), with a total of 27 configurations
- Contract award used a best value process that included pricing, synthetic and LQCD benchmark performance, power and space efficiency, preferred components, vendor qualifications, and vendor past performance

FY2011 GPU Schedule

(Planned as of May 2011 Review, and Achieved)

2011 (except as noted)

- May 31 (July 24) – RFP released to vendors
  - Delayed by CR and Fermilab purchasing backlogs
- June 30 (Aug 24) – Bids received
- July 15 (Sept 23) – Purchased order released (commit FY11 funds)
  - Delayed by vendor negotiations and Fermilab purchasing backlogs
- Aug 1 (Oct 13) – Sample unit received
- Sept 15 (Jan 9, 2012) – Delivery of all items
  - Delayed by Thailand flooding, holiday schedules at vendor and Fermilab
- Oct 19 (Feb 5, 2012) – Acceptance test complete
- Oct 31 (Mar 1, 2012)– Release to production
  - Preceded by three weeks of unbilled production (“friendly user”)

GPU Cluster: Cost and Performance

- GPU-accelerated cluster contract was awarded to Hewlett-Packard
  - Cluster is based on their SL390s-G7 blades, with 2 NVIDIA Tesla M2050 GPUs per server host (48 GB host memory), high density (4 GPUs per 2U)
  - Full bisection bandwidth QDR Infiniband (Mellanox)
  - These hosts gave the best host-to-GPU and GPU-to-Infiniband performance of all of the configurations bid by the various vendors (no PCIe switches, and a faster PCIe implementation compared with the AMD-based solutions)
  - Total GPU count: 152 (goal: 128)
  - Cost including G&A: $615K (budget: $640K)
- Capacity: 1.216M GPU-hrs per year (plan: 1.024M GPU-hrs per year)
- 2011 Delivered GPU-hrs: est. 0.423M (allocated: 0.430M)
- Largest allocation (MILC) is achieving a 2.1 speedup on f_π decay constant calculation using *asqtad* gauge configurations, comparing a GPU node to a Ds node
  - For this calculation, Dsg has a cluster-equivalent capacity of 8.5 TF
  - Of the production software ported to date, *asqtad* has lowest acceleration
  - FY11 baseline deploy goal: 12.0 TF   Achieved: 9.0 TF + 8.5 (equivalent) TF

Questions?