

Report on the Clusters at Fermilab

Don Holmgren
USQCD All-Hands Meeting
JLab
May 4-5, 2012

Outline

- Hardware
- Storage
- Statistics
- Possible Summer Outages
- Future Facilities

New GPU-Accelerated Cluster (Dsg)

- Hardware design:
 - Hosts use dual socket Intel 2.53 GHz “Westmere” processors, with 8 cores/host, 48 GiB host memory
 - 152 Tesla M2050 GPUs, two per host machine, with both GPUs attached to the same processor socket (Infiniband interface is on the second socket)
 - QDR Infiniband, full bandwidth
 - Suitable for jobs requiring large parallel GPU counts and good strong scaling
 - GPUs have ECC enabled to allow safe non-inverter calculations
 - ECC can be disabled per job at job start time to increase performance and available GPU memory (from 2.69 GiB to 3.0 GiB per GPU)
- Released to production March 1, 2012 (Planned Oct 31, 2011)
 - Very late because of impacts of continuing resolutions, vendor delays

Hardware – Current Clusters

<u>Name</u>	<u>CPU</u>	<u>Nodes</u>	<u>Cores</u>	<u>Network</u>	<u>DWF</u>	<u>Asqtad</u>	<u>Online</u>
Kaon	Dual 2.0 GHz Opteron 240 (Dual Core)	< 300	< 1200	Infiniband Double Data Rate	4696 MFlops per Node	3832 MFlops per Node	Oct 2006 2.56 TFlops
JPsi	Dual 2.1 GHz Opteron 2352 (Quad Core)	856	6848	Infiniband Double Data Rate	10061 MFlops per Node	9563 MFlops per Node	Jan 2009 / Apr 2009 8.40 TFlops
Ds (2010) (2011)	Quad 2.0 GHz Opteron 6128 (8 Core)	421	13472	Infiniband Quad Data Rate	51.2 GFlops per Node	50.5 GFlops per Node	Dec 2010 11 TFlops Aug 2011 21.5 TF
Dsg (2012)	NVIDIA M2050 GPUs + Intel 2.53 GHz E5630 (quad core)	76	152 GPUs 608 Intel	Infiniband Quad Data Rate	29.0 GFlops per Node (cpu)	17.2 GFlops per Node (cpu)	Mar 2012

Storage

- Global disk storage:
 - 543 TB Lustre filesystem at `/lqcdproj`
 - ~ 6 TB total “project” space at `/project` (backed up nightly)
 - ~ 6 GB per user at `/home` on each cluster (backed up nightly)
- Robotic tape storage is available via `dccp` commands against the dCache filesystem at `/pnfs/lqcd`
- Worker nodes have local storage at `/scratch`
 - Multi-node jobs can specify combining `/scratch` from one or more nodes into `/pvfs`
 - `/pvfs` is visible to all nodes of the job and is deleted at job end

Storage – Lustre Statistics

- 543 TiB capacity, 475 TiB used, 114 disk pools
(2011: 387 TiB used in 110 pools)
- 101M files (59M last year)
42.7M files in /project
- File sizes: 344.8 GiB maximum (a log file!)
4.96 MiB average
8.54 MiB avg if /project excluded (/project avg: 84.15 KiB)
- Directories: 466K (321K excluding /project)
192232 files in largest directory



Storage – Planned Changes

1. Move /project from Lustre filesystem – July 2012
 - Currently /project is stored in /lqcdproj
 - /project has many (40M) small files. We've found that this strains Lustre, particularly for the /project and metadata backups
 - We will very likely move /project to NFS-exported RAID disk during a shutdown at the very beginning of the 2012 project year
 - Will continue nightly backups
2. Deploy additional Lustre storage – now and during 2012 project year
 - Expect to add a net of about ~ 60 TB by mid-June (total to 600 TB)
 - Once funds are available in FY12, will add at least two and perhaps three increments of ~ 100 TB

Storage – Data Integrity

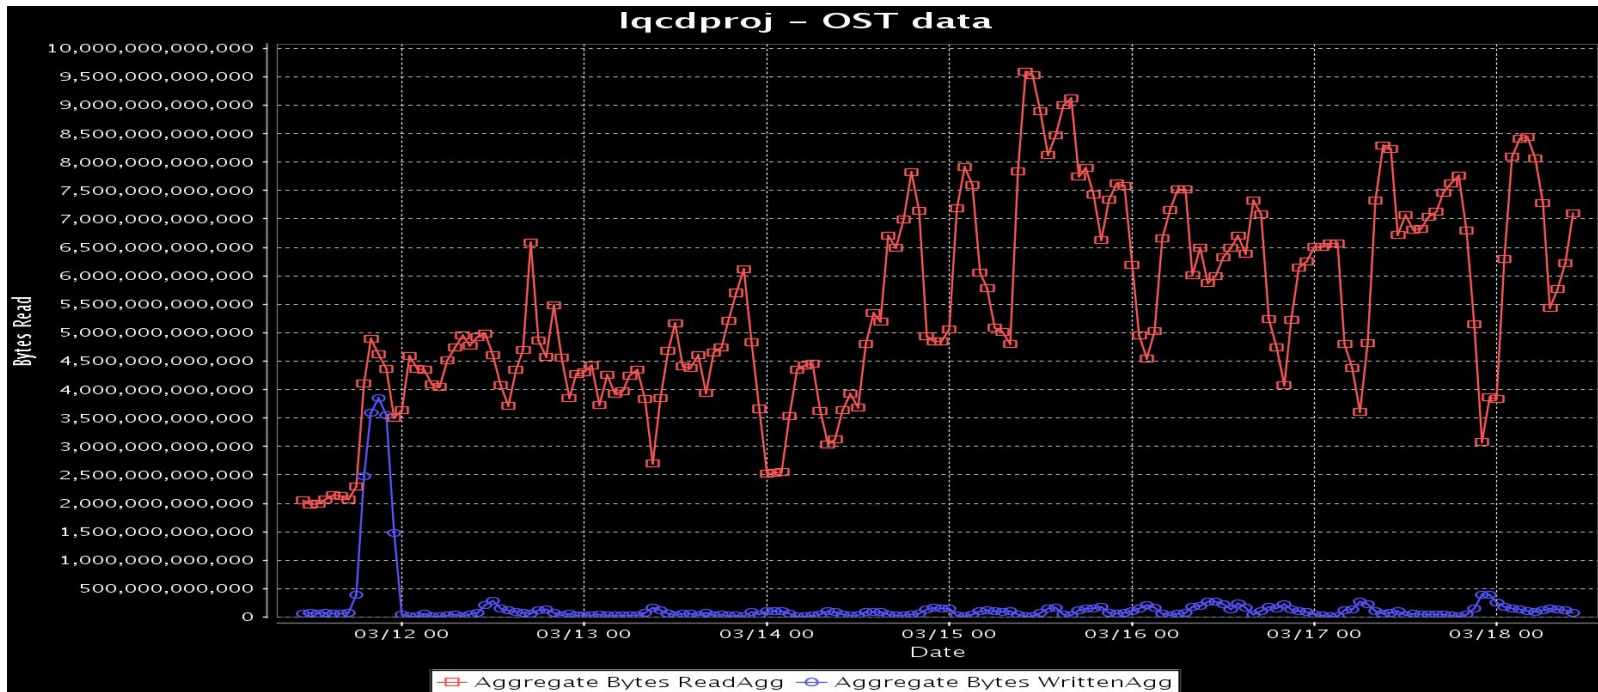
- Some friendly reminders:
 - Data integrity is your responsibility
 - With the exception of home areas and /project, backups are not performed
 - Make copies on different storage hardware of any of your data that are critical
 - Data can be copied to tape using **dccp** commands. Please contact us for details. We can also show you how to make multiple copies that are guaranteed to be on different tapes. We have never lost LQCD data on Fermilab tape (1.09 PiB and growing).
 - At 114 disk pools and growing, the odds of a partial failure will eventually catch up with us – **please don't be the unlucky project that loses data when we lose a pool.**

Storage - Utilization



- Utilization of /lqcdproj will always increase to fill all space. This is a good thing (disk is expensive – we don't mind you using it).
- But:
 - Lustre misbehaves when the pools get above 95% fill. Please be responsive to our requests to clear space. If users prefer, we can set up a transient partition similar to JLab in which older files are automatically deleted to clear space.
 - For our planning purposes, it is critical that in your proposals that storage requests are reasonably (factor of 2) accurate. We have had instances of both large overruns and under-utilization. We can adjust budgets annually, but we need reliable data.
 - For the first time this year, we've seen instances of I/O patterns by some job types that causes total filesystem bandwidth to saturate. This can affect other jobs, and it definitely affects the rates of some critical maintenance activities. We are working to understand this in more detail, and may again this year throttle the number of jobs that can be run at a time by projects that have these I/O patterns.

Storage - Performance




- Hourly aggregate **read** and **write** rates
- Peak hourly rate of **9.5 TiB/hr** corresponds to **2.64 GiB/sec** sustained rate
- We observe that highest read rates when jobs using eigenvector projection methods are running


Statistics

- Since May 1, 2011, including JPsi, Ds, Dsg
 - 442,864 jobs (714,809 including Kaon)
 - 188.0M JPsi-core-hours
 - We did not charge for Kaon (an additional 7.3M JPsi-core-hours)
 - 170 GPU-KHrs (March + April)
- USQCD users submitting jobs:
 - FY10: 56
 - FY11: 64
 - FY12 to date: 51


Fermilab LQCD Cluster Utilization FY10-FY12




Fermilab 2012 Job Statistics




Fermitab 2012 Job Statistics



Fermilab 2012 Job Statistics



Fermilab 2012 Job Statistics



Progress Against Allocations

- Total Fermilab allocation: **170.06M JPsi core-hrs**
433 GPU-KHrs
- Delivered to date: **160.9M** (94.6%, at 83.8% of the year)
170.3 GPU-KHrs (39.3%, at 50%)
 - Does not include disk and tape utilization (**0.48M + 5.8M**)
 - Does not include **6.28M** delivered without charge on Kaon
 - Does not include **64 GPU-KHrs** delivered in February on Dsg (friendly user period)
 - Japan projects: **3.84M** (ended Dec 31)
 - Class A (13 total): 4 finished, 5 at or above pace
 - Class B (9 total): 4 finished, 1 at or above pace
 - Class C: 5 for GPUs, 6 for conventional

Possible Summer Outages



Possible Summer Outages



- Air conditioning condensers are in a valley between the GCC building and the old beamline.
- Last summer, on very hot and humid days, the condensers could not reject sufficient heat. This caused about a week of downtime, and all of JPsi and Ds were powered off.
- The beamline berm is being removed as we speak. A Fermilab engineering study predicts a remaining 30% chance of shutdowns this summer (the best fix would be to elevate the condensers or relocate them to the roof).
- If we have a thermal problem:
 - Stage 1: All JPsi and Ds nodes will have their CPU frequencies reduced (50% drop in frequency → 22% drop in power → 33% drop in LQCD throughput) and some nodes will be powered off (30% load shed)
 - Stage 2: 50% of nodes will be powered off

Future Facilities

- In FY13, the LQCD hardware project will deploy some combination of:
 - BlueGene/Q rack or partial rack at BNL
 - Conventional cluster, probably at FNAL
 - Accelerated cluster, probably at FNAL
- Assuming continued drop in pricing per flop, in FY14 the project will deploy some combination of
 - Conventional cluster, probably at FNAL
 - Accelerated cluster, probably at FNAL
- The project welcomes input on hardware architecture decisions

User Support


Fermilab points of contact:

- Best choice: lqcd-admin@fnal.gov
- Don Holmgren, djholm@fnal.gov
- Amitoj Singh, amitoj@fnal.gov
- Nirmal Seenu, nirmal@fnal.gov
- Jim Simone, simone@fnal.gov
- Ken Schumacher, kschu@fnal.gov
- Rick van Conant, vanconant@fnal.gov
- Paul Mackenzie, mackenzie@fnal.gov
- Please use lqcd-admin@fnal.gov for requests and problems


Questions?

Previous Histograms

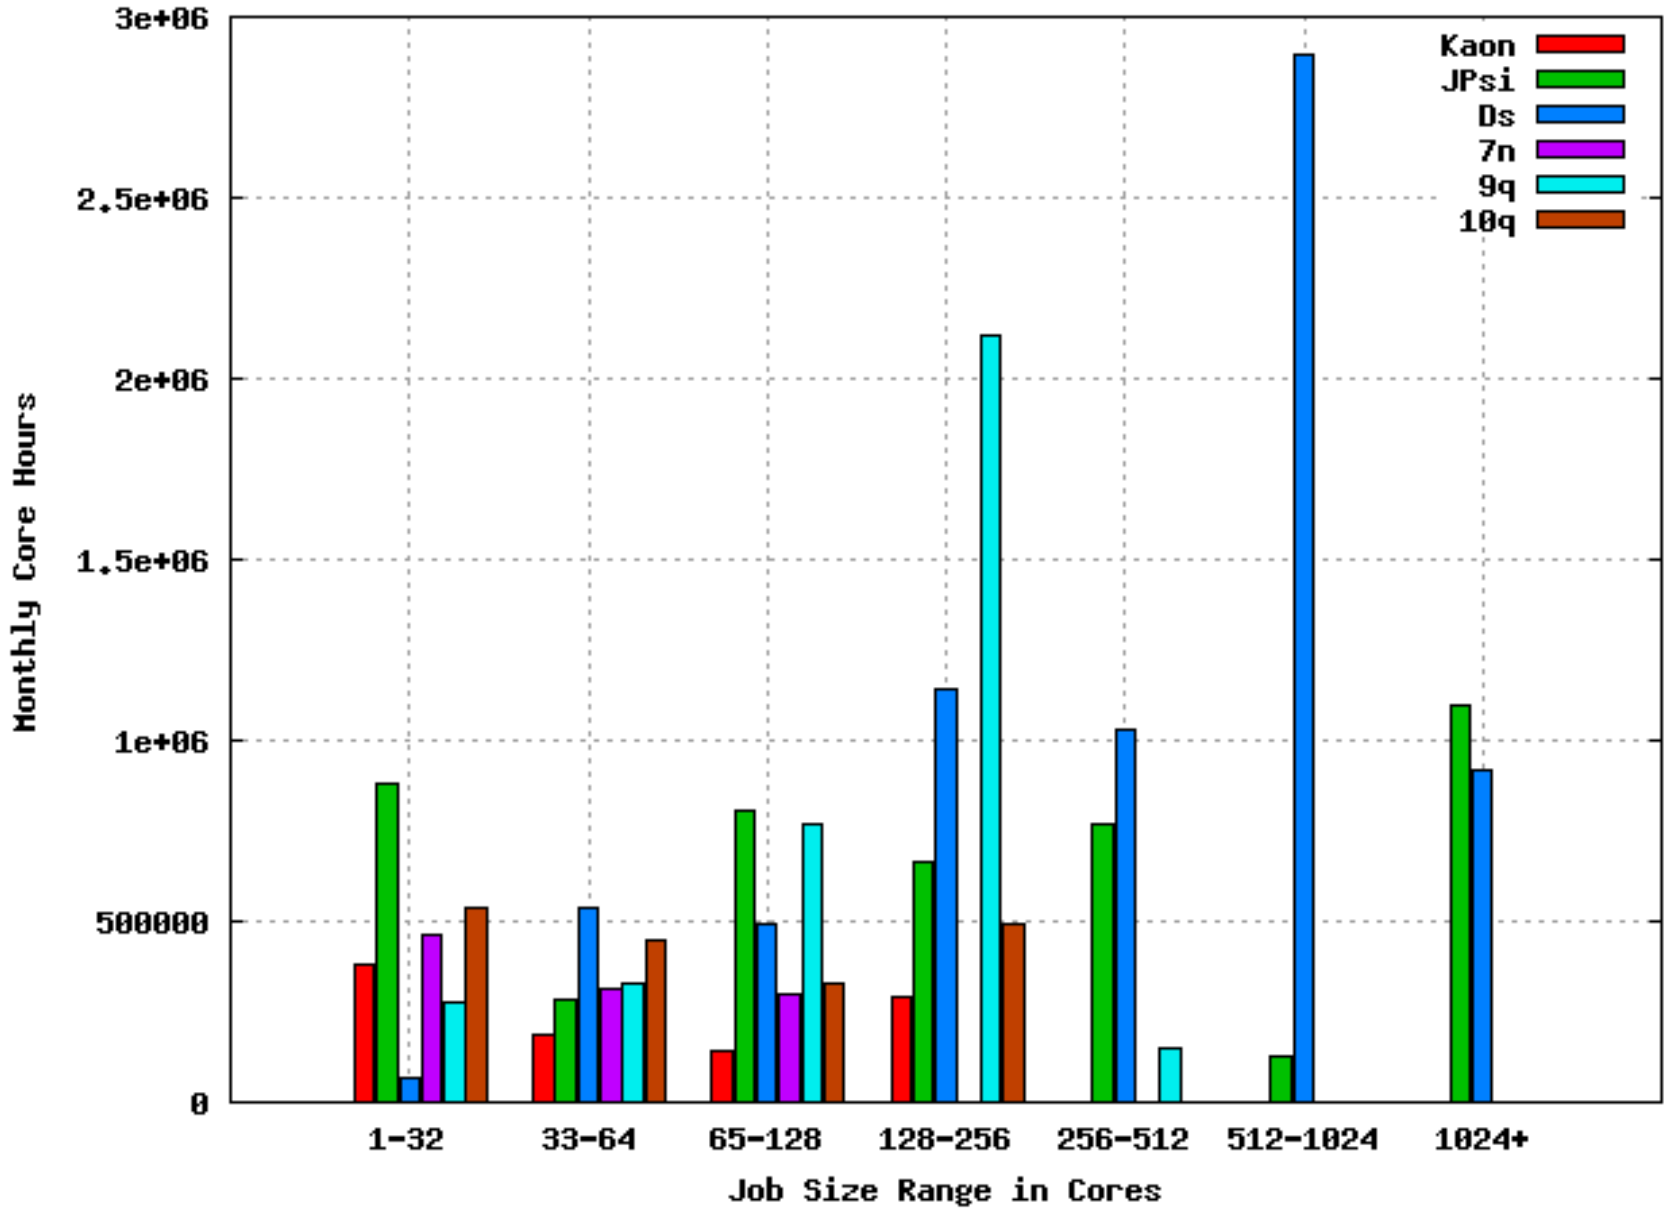
USQCD 2011 Job Statistics, July 2010 - March 2011



USQCD 2011 Job Statistics, July 2010 - March 2011



USQCD 2011 Job Statistics, July 2010 - March 2011



USQCD 2011 Job Statistics, July 2010 - March 2011

