

LQCD Computing at BNL

2016 USQCD All-Hands Meeting
BNL
April 29, 2016

Robert Mawhinney
Columbia University

BGQ Computers at BNL

USQCD half-rack
(512 nodes)

2 racks of DD1
RBRC

1 rack of DD2
BNL



USQCD 512 Node BGQ at BNL

- USQCD SPC allocated time for 3 projects in 2014-2015. Usage as of June 30, 2015.

P.I.	Allocated	Used	% Used
Kelly	42.03	47.65	114%
Kuti	15.42	15.62	101%
Mackenzie	13.35	15.24	114%

- USQCD SPC allocated time for 4 projects in 2015-2016. Usage as of April 24, 2016.

P.I.	Allocated	Used	% Used	Max Usage	Max % Usage
Feng	27.14	31.96	117%		
Kuti	14.10	7.60	54%	14.80	105%
Mackenzie/ Sugar	29.56	21.5	73%	31.0	105%

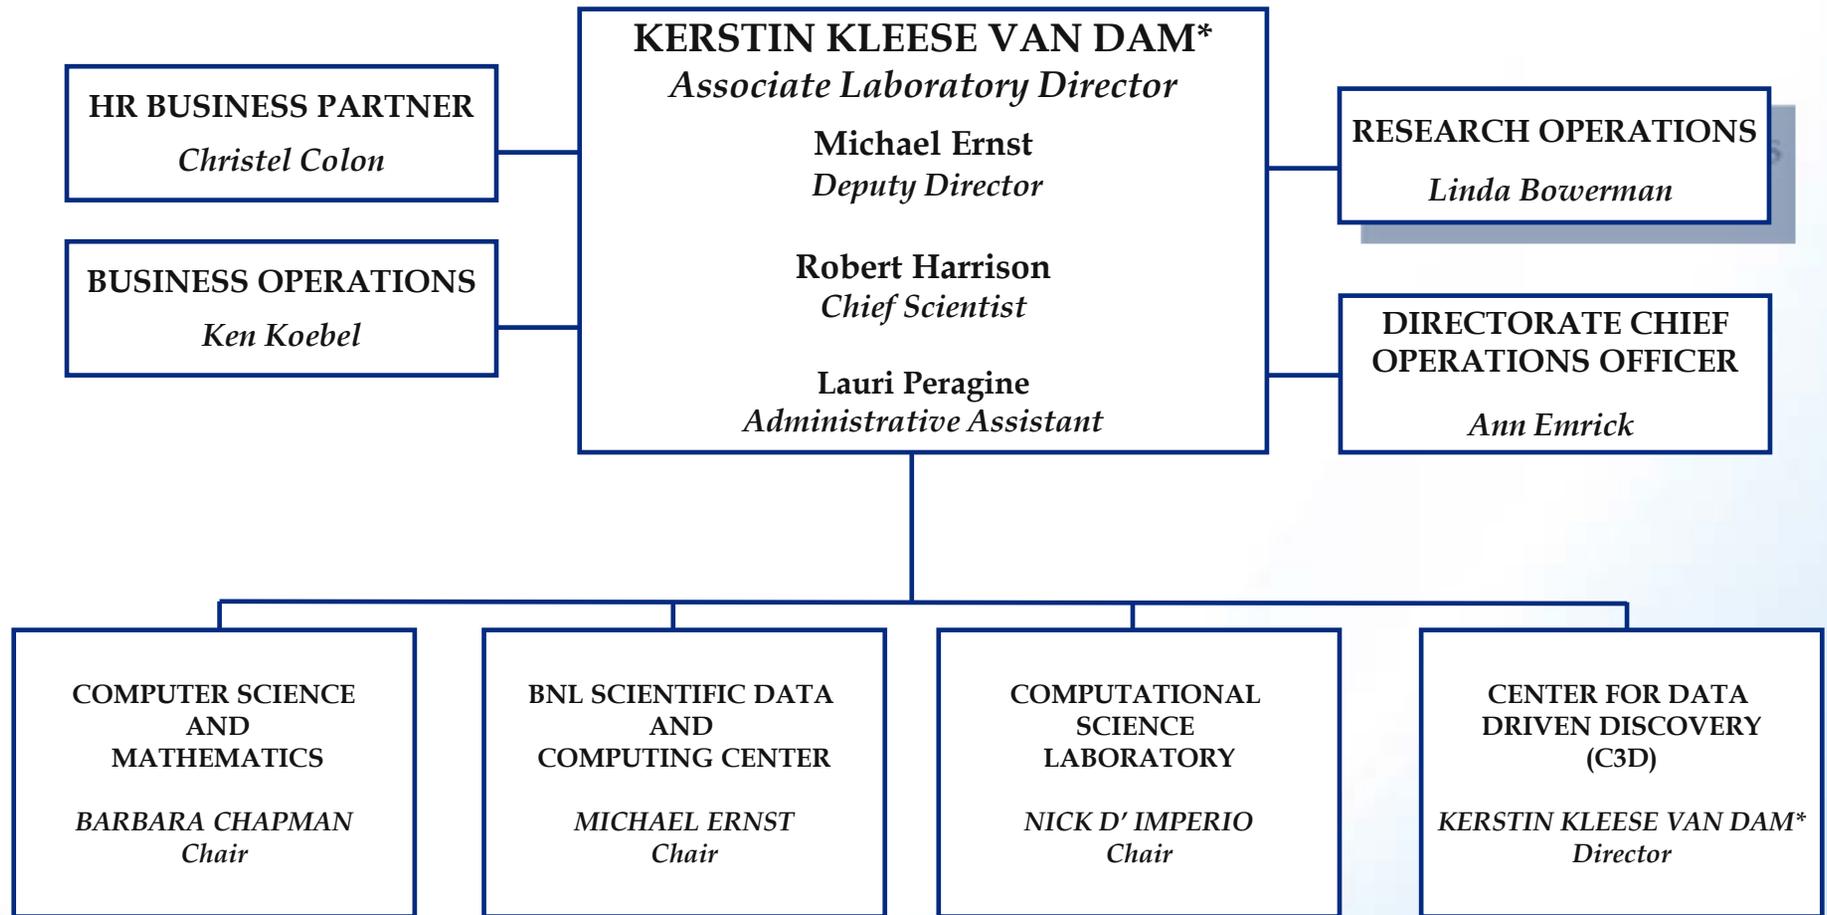
On schedule to deliver slightly above full allocations to all users.

USQCD BGQ Utilization at BNL 2015-2016

2013-2014 allocation month	Utilization	Comments
July	92.6%	1 hour chilled water outage
August	97.9%	
September	98.6%	
October	98.9%	
November	91.3%	
December	96.9%	12 hour hardware problem downtime, Dec 28
January	99.0%	12 hour downtime, January 29-30
February	98.0%	
March	94.4%	
April		
May		
June		

- Utilization reported here is the fraction of the time jobs were running divided by the maximum hours available in the month, with no derating
- Almost all usage has been a single user running on 512 nodes full time.

Computational Science Initiative Directorate



*Reports to Robert Tribble, Deputy Director Science & Technology
ESH Coordinator - Bob Colichio, Pat Carr

Approved: _____
Kerstin Kleese van Dam Date

Computational Science Initiative (CSI)

- CSI aims to harness computing expertise across BNL to foster cross-disciplinary collaborations to deliver operational data processing/analysis capabilities and address future scientific challenges.
- CSI to leverage RHIC-ATLAS Computing Facility (RACF) expertise in forming the new Scientific Data and Computing Center (SDCC) → see www.bnl.gov/compsci.
- New Institutional Cluster (IC) is the first significant hardware component in support of HPC activities at BNL.
 - Expected to be ready for users in July 2016
 - 108 servers with dual-socket Broadwell cpu's
 - 2 Nvidia K80 gpu's per server
 - 128 GB of ECC RAM per server
 - EDR Infiniband inter-connect
 - 1 PB GPFS-based storage
 - SL 7.x (or CentOS equivalent) operating system
 - Slurm workload manager
- Upgrades to data center power and cooling infrastructure to accommodate the IC.
 - Expected to be completed in May 2016

From Tony Wong
LQCD Site Manager as of 3/1/16

CSI support for USQCD

- As of March 1st 2016, SDCC takes over operations for BG-Q and DD2.
 - Support for legacy BG-Q and DD2 until September 2017.
- Enhanced support for USQCD under 3-site (BNL, FNAL, JLAB) deployment model.
 - Fraction of IC to be made available to meet project's computing and storage goals
 - Access to RACF mass storage system (intermediate disk caching and tape)
 - SDCC will support the evaluation, purchase and operation of successor(s) to BG-Q and DD-2 in FY17 and FY18.

Site manager: Tony Wong

Site architects: Shigeki Misawa, Bob Mawhinney

Conclusions and Outlook

- BGQ half-rack and DD2 are running very reliably and no changes are planned.
 - * Joe DePace and Robert Riccobono handle hardware issues
 - * System administration by Zhihua Dong
- BNL Computational Science Initiative strongly supportive of USQCD hardware at BNL
 - * Providing access to BNL insitutional cluster as of mid to late summer (Broadwell + K80's)
 - * New USQCD hardware to be procured and managed coherently with other BNL computing resources, such as RHIC and ATLAS computing
 - * USQCD will have access to high throughput (up to 24 GBytes/sec) shared file system.
 - * Opportunity for USQCD to run jobs on larger partitions for shorter periods of time.
- Exciting opportunities for USQCD and BNL.
- Many thanks to Frank Quarant and Nick D'Imperio for their leadership of and work on the USQCD project at BNL for the last many years.