

BNL Site Report

Chulwoo Jung (Presenter, co-site architect)

Zhihua Dong (site manager)

Costin Caramarcu (co-site architect)

Tony Wong

May 1, 2020

The logo for Brookhaven National Laboratory, featuring a stylized particle detector or accelerator structure above the text "BROOKHAVEN NATIONAL LABORATORY".
BROOKHAVEN
NATIONAL LABORATORY

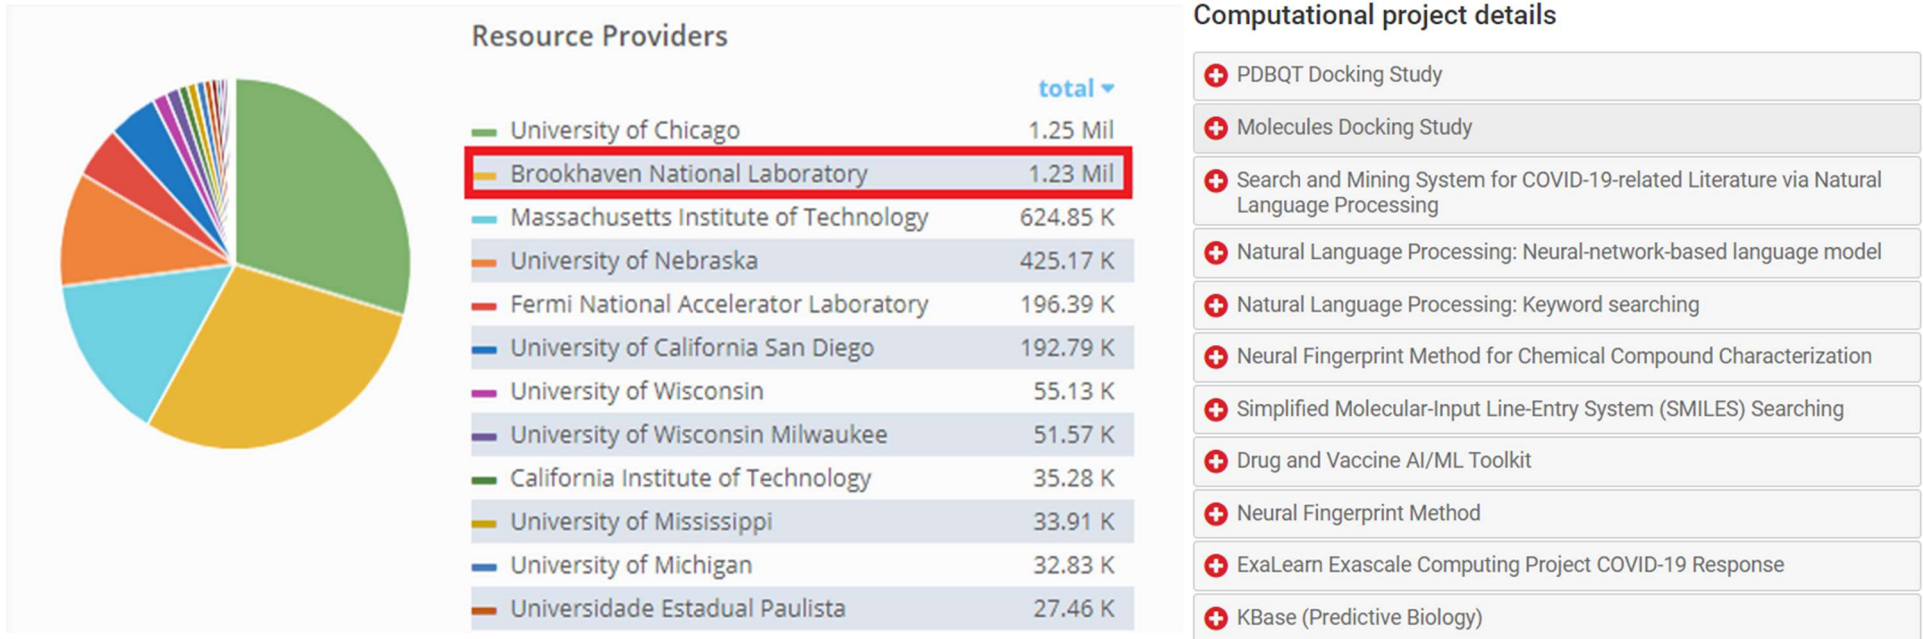
The logo for the U.S. Department of Energy, featuring the official seal of the U.S. Department of Energy to the left of the text "U.S. DEPARTMENT OF ENERGY".
U.S. DEPARTMENT OF
ENERGY

Scientific Data and Computing Center (SDCC)



- Support for various programs:
 - RHIC, LHC ATLAS, BER ARM, LSST, DUNE, LQCD,RIKEN, BES, Center for Functional Nanomaterials(CFN), National Synchrotron Light Source(NSLS) II, National Nuclear Data Center, Simons Foundation,...
- ~2000 users from >20 projects
- Staff
 - 36 full-time regular members
 - 3 current openings

Covid-19 Researches at BNL



<https://gracc.opensciencegrid.org/dashboard/db/covid-19-research?orgId=1>

- Most of current BNL SDCC computing resource usage for covid-19 from Open Science Grid (Most resources provided until recently)
- Other BNL participation in covid-19 projects can be found in <https://www.bnl.gov/science/COVID-working-group.php#projects>

SDCC support for HEP experiments

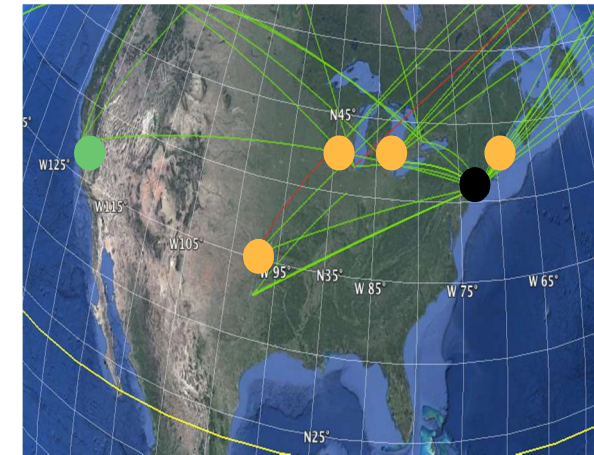
•The RHIC Tier 0

- Store and process data from RHIC experiments
- Provide analysis means for 1'200 users
- Long term data preservation
- Simulation resources for future programs (sPHENIX & EIC)



•The US ATLAS Tier 1

- ~25% of ATLAS Tier 1 computing capacity worldwide
- Store RAW data from LHC and from simulation
- Distribute data to the 4 US Tier 2 sites + analysis site (SLAC)
- **Analysis center for US physicists**



- From 41 institutes (incl. 4 Nat. Labs)
- 600 physicists, 190 PhDs

•A Belle II data center outside Japan


- Initial operations began on Oct. 2017
- Data taking began in Fall 2018



24/7 availability

SDCC Resources Summary

- 90+k CPU cores — 4 PFlops
 - 4 HPC Institutional Clusters (GPU, KNL, Skylake,ML)
- 688 GPUs
 - K80, P100, V100
- ~80 PB of disk storage
 - Central and distributed storage systems
- 165+ PB of tape storage
 - Largest HPSS tape library in the US, 3rd worldwide
- 2x100 Gbps connection to ESnet
 - 100Gbps redundancy
 - Onsite ESNet support

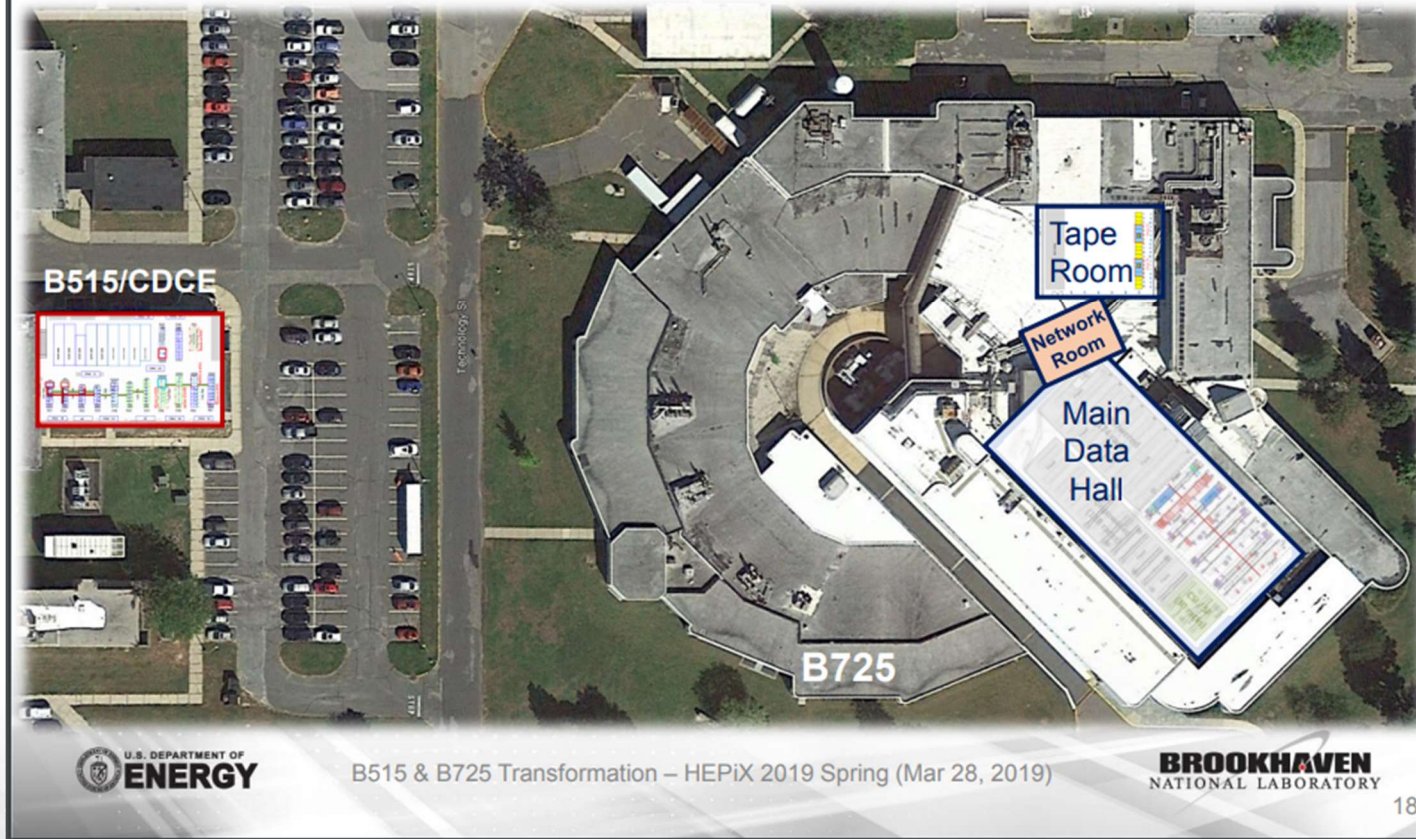


BNL Data Center (CFR)

- Existing data center is full
- Construction of new data center
 - Began last May,
 - Demolition phase close to finish ...halted.
 - Was expect to be online On Feb 2021, but could be delayed...
- Approximately 3x times more floor space and electrical power with room to expand if needed
- Higher PUE (power utilization efficiency) – mandated by DOE
- Any new LQCD-accessible systems post-2021 would be housed in the new data center

New BNL Data Center (CFR)

Proposed solution: Constructing the new datacenter in B725 in FY19-21, migrating all spinning disk storage and compute to it in FY21-23; leaving the B515 datacenter reduced to just one area (CDCE) as a tape storage room



SDCC support for HPC

- Institutional Clusters
 1. CPU-GPU cluster (aka “Annie”) with 216 compute nodes (36 physical core Xeon Broadwell and 2 GPUs each: K80 & P100) inter-connected with dual-rail Infiniband EDR
 2. KNL cluster (aka “Frances”) with 144 nodes (64 cores each) interconnected with dual-rail Intel OPA
 3. Skylake cluster: 64 nodes (36 physical cores each) with single-rail Infiniband EDR
 4. Machine Learning Cluster: 5 EDR connected nodes, each have 8 Nvlink connected V100 GPU
 5. Start In production since January 2017
 - >300 registered users

MoUs (describing level of resources and services) organized with each user community



Monitoring

- Several tools available
 - Graphical interface here (**authentication required**)
 - <https://monitoring.sdcc.bnl.gov/pub/grafana/>
 - Accounting information
 - <https://monitoring.sdcc.bnl.gov/pub/allocation>
 - **LQCD only**
 - <https://monitoring.sdcc.bnl.gov/pub/allocation/lqcd.html>
 - After loading module lqcd, Command line “lquota” for same information

Accounting

<https://monitoring.sdcc.bnl.gov/pub/allocation/lqcd.html>

BNL SDCC LQCD Projects Usage Summary

Institutional Cluster

(Sky Core Hours)

updated: 2020-03-31 00:03:25

Cluster	Account	Start Date	End Date	Allocation	Usage	Usage(%)			
Annie-IC	lqcd-19-20	2019-07-01	2020-06-30	6,728,385	4,783,845	71.10%			
Project	Original SPC Allocation	Adjustment	Adjusted SPC Allocation	Usage	Progress(%)	Remain	30Day Usage	30Day BurnRate	
1	stagnug-2-19-20	1,034,714	425,017	1,459,731	1,609,524	110.26%	0	473,015	32.40%
2	semibdff-19-20	691,747	273,454	965,201	857,264	88.82%	107,937	120,253	12.46%
3	nucstructclover-19-20	842,885	170,303	1,013,188	740,004	73.04%	273,185	125,144	12.35%
4	axialgpu-19-20	697,560	296,936	994,496	676,647	68.04%	317,849	1,299	0.13%
5	sextet-19-20	546,422	156,075	702,497	604,166	86.00%	98,330	49,126	6.99%
6	class-c-19-20	10,765	0	10,765	615	5.72%	10,149	0	0.00%
7	qgpd-19-20	1,499,742	0	1,499,742	295,625	19.71%	1,204,117	295,625	19.71%
8	UnAllocated:	-1,510,507	1,593,273	82,766	0	0.00%	0	0	0.00%

Skylake Cluster

(Sky Core Hours)

updated: 2020-03-31 00:03:25

Cluster	Account	Start Date	End Date	Allocation	Usage	Usage(%)			
Skylake	lqcd-sky-19-20	2019-07-01	2020-06-30	17,278,272	11,118,173	64.35%			
Project	Original SPC Allocation	Adjustment	Adjusted SPC Allocation	Usage	Progress(%)	Remain	30Day Usage	30Day BurnRate	
1	semibdff-sky-19-20	4,500,000	2,414,093	6,914,093	5,208,518	75.33%	1,705,575	407,806	5.90%
2	stagnug-2-sky-19-20	5,000,000	(2,731,934)	2,268,066	2,417,375	106.58%	0	731,201	32.24%
3	etap-sky-19-20	1,500,000	300,000	1,800,000	625,947	34.77%	1,174,053	53,883	2.99%
4	qgpd-sky-19-20	5,000,000	17,841	5,017,841	2,862,086	57.04%	2,155,755	367,215	7.32%
5	class-c-sky-19-20	8,333	0	8,333	4,247	50.97%	4,086	0	0.00%
6	UnAllocated:	1,491,667	(221,728)	1,269,939	0	0.00%	0	0	0.00%

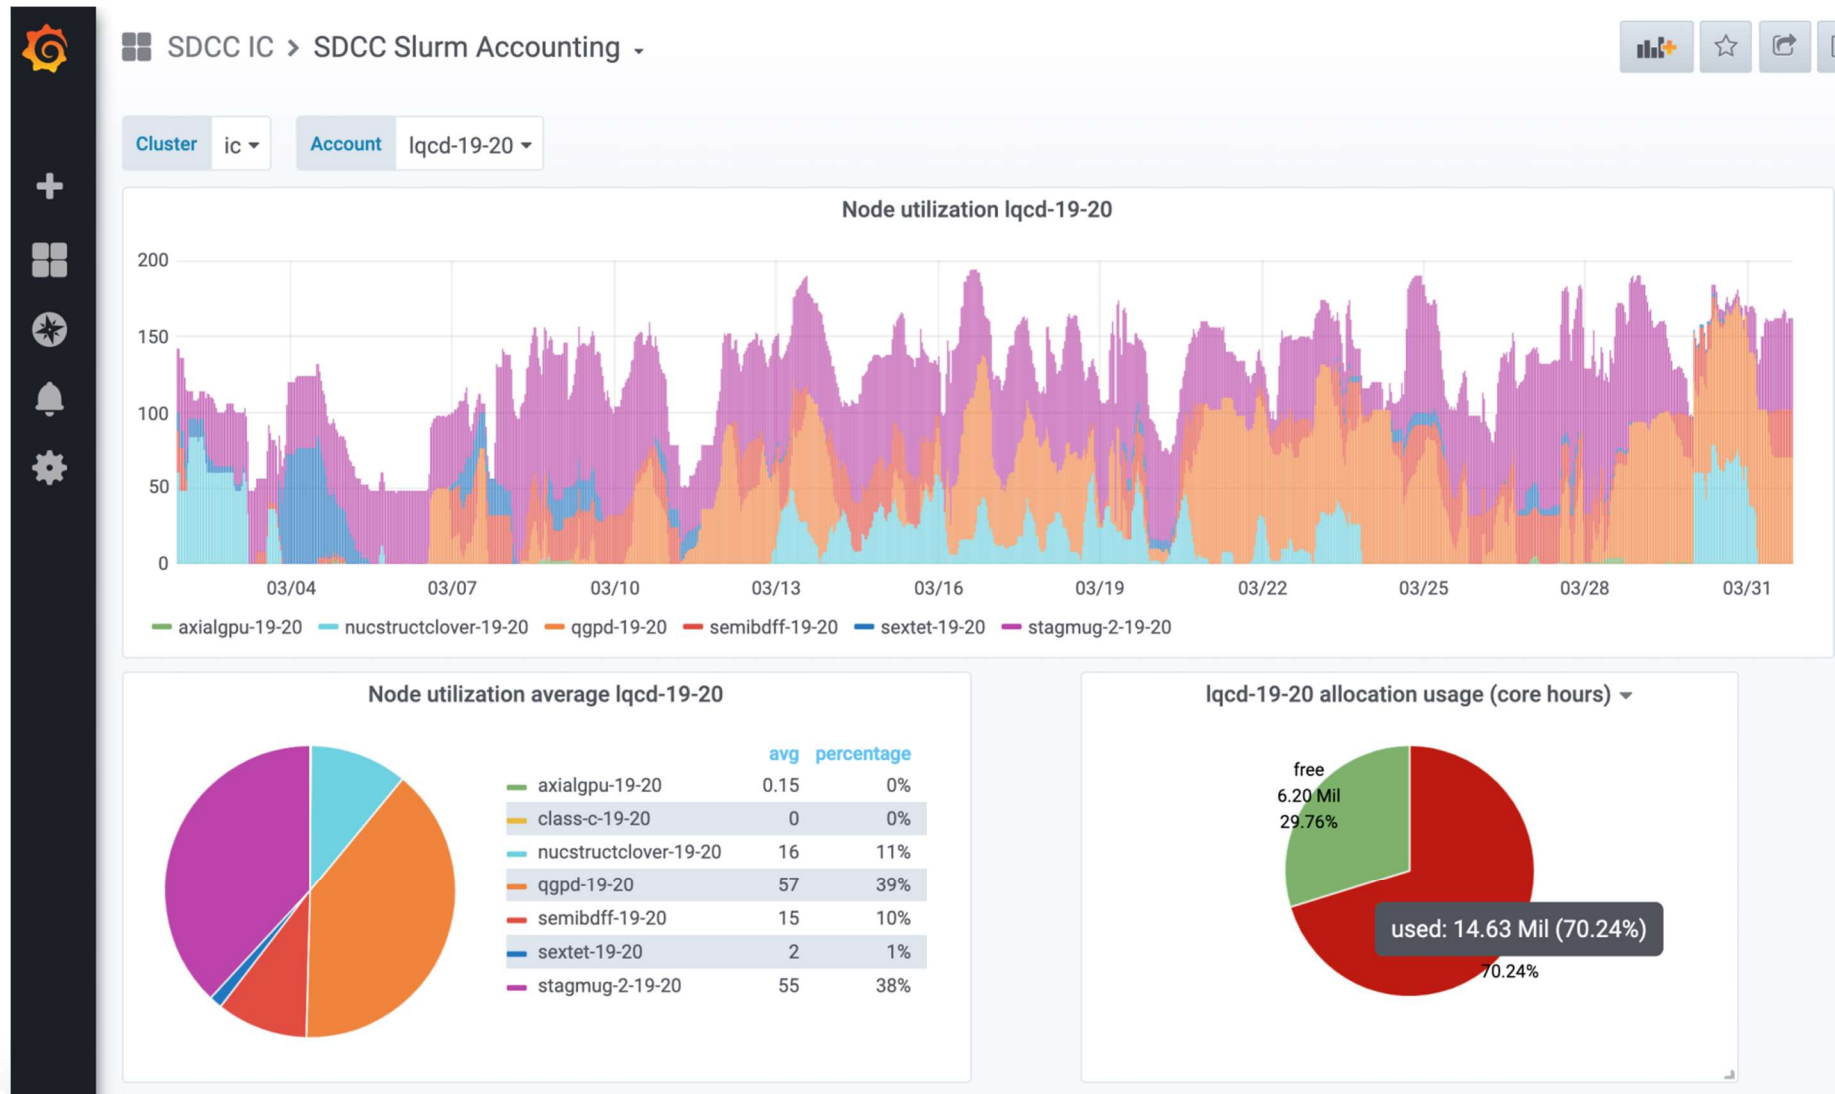
KNL Cluster

(Sky Core Hours)

updated: 2020-03-31 00:03:09

Cluster	Account	Start Date	End Date	Allocation	Usage	Usage(%)			
Frances-KNL	lqcd-knl-19-20	2019-07-01	2020-06-30	11,133,023	9,796,792	88.00%			
Project	Original SPC Allocation	Adjustment	Adjusted SPC Allocation	Usage	Progress(%)	Remain	30Day Usage	30Day BurnRate	
1	qcdqeda-knl-19-20	6,756,000	164,205	6,920,205	7,142,674	103.21%	0	478	0.01%
2	posnpr-knl-19-20	1,689,000	328,417	2,017,417	1,376,230	68.22%	641,186	137,502	6.82%
3	ndbeta-knl-19-20	1,970,500	(492,622)	1,477,878	1,251,523	84.68%	226,355	978,515	66.21%
4	class-c-knl-19-20	28,150	0	28,150	26,365	93.66%	1,785	0	0.00%
5	UnAllocated:	-928,950	1,618,323	689,373	0	0.00%	0	0	0.00%

Accounting <https://monitoring.sdcc.bnl.gov/pub/grafana/>



LQCD Access to SDCC Resources

- Current resources allocated
 - 578k node-hour allocation on CPU-GPU cluster
 - 309k node-hour allocation on KNL cluster
 - 480k node-hour allocation on Skylake cluster
 - 600 TB of GPFS storage
 - 600 TB of Tape Storage service (134TB used so far)
- Usage policy
 - SDCC does not decrement underused allocations as a function of time, but unused allocations are increasingly “at risk” as we approach end of year when resource contention can become an issue.
 - Opportunistic lower priority usage after allocation used up, when cluster have available resource.

HPSS Tape Storage Data Growth View

110 Days Activities

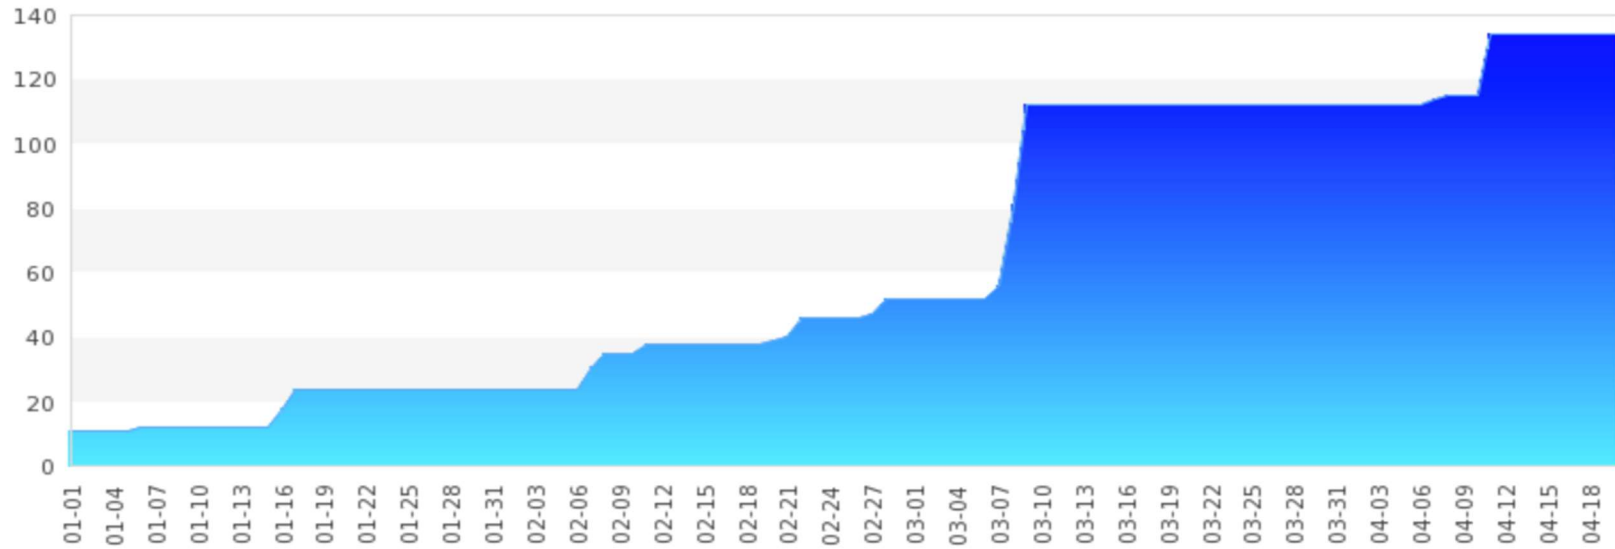


LQCD

HPSS Data Growth - LQCD

Date: [01-01 - 04-20] | Window Range: [11 - 134], Delta: 123 TB

■ Tape Usage in TB



+ - Reset

Start Date: 1 January 2020 End Date: 20 April 2020 GO

USQCD data growth at BNL(134TB as of 4/20)

User Support

- Facility website is www.sdcc.bnl.gov .
 - New accounts
 - Instructions on website
 - Usually ~24 hours to process after verification
 - User support requests
 - **SDCC policy is to respond within 3 business days. Majority is resolved within this period**
 - **In the July 1, 2018 to June 30, 2019, 200 tickets were submitted to ticket queues (majority from LQCD users) and 91% were resolved within 3 business days**
- Bi-weekly meetings between facility staff and program/experimental Liaisons
 - Agenda on <https://indico.bnl.gov/category/169/>
 - Remote access via BlueJeans—Minutes of meeting posted for those who cannot join in person or remotely

Recent Developments

- Globus endpoint @ the SDCC
 - Fast point-to-point data transfer mechanism
 - Available on Institutional Clusters (Endpoint 'SDCC')
 - <https://www.racf.bnl.gov/experiments/sdcc/institutional-cluster/storage>
- **Tape archival services available for LQCD**
 - All hardware installed and tested. Initially 600 TB of tape storage with room to grow
 - Interface mechanism and documentation (including a Data Management Plan) available in early May 2019.
 - Information on USQCD webpage (<https://www.usqcd.org/bnl/tape-archive.html>)
 - **Deletion or renaming creates dead space. Will count against USQCD.**
- **BNLBox service** operational since December 2019 – file sharing and archiving (like CERNBox) between SDCC users
 - www.racf.bnl.gov/docs/services/cloud-storage/using-cloud-storage

Federated User Management

- SDCC moving towards accepting selected federated identity provider (IDP) for user management
 - First step towards Single Sign-On (SSO) with Multi-Factor Authentication (MFA)
 - InCommon and OneID– used at many universities and labs
 - SDCC establishing IDP under InCommon (independent from BNL Active Directory IDP)
 - Some issues still unresolved
 - Trust levels
 - Resources available to BNL users vs. non-BNL users
 - Some applications already enabled (Jupyter, Indico and Invenio)
 - Potentially beneficial to LQCD users at BNL—SDCC account can be used to authenticate to external organizations
 - Evolving cyber-security policy to accommodate federated access to BNL resources

Questions?