BNL RACF Site Report

Ofer Rind

Spring HEPiX, Budapest, Hungary

April 24, 2017

RACF Overview

- Located at Brookhaven National Laboratory on Long Island, NY
- Provides full service computing mainly for the two RHIC experiments — STAR, PHENIX — and for ATLAS (US Tier1), along with some smaller groups: LSST, Daya Bay, DUNE, EIC, etc.
    - RHIC Run 17 (STAR, 255 GeV polarized p-p) in progress

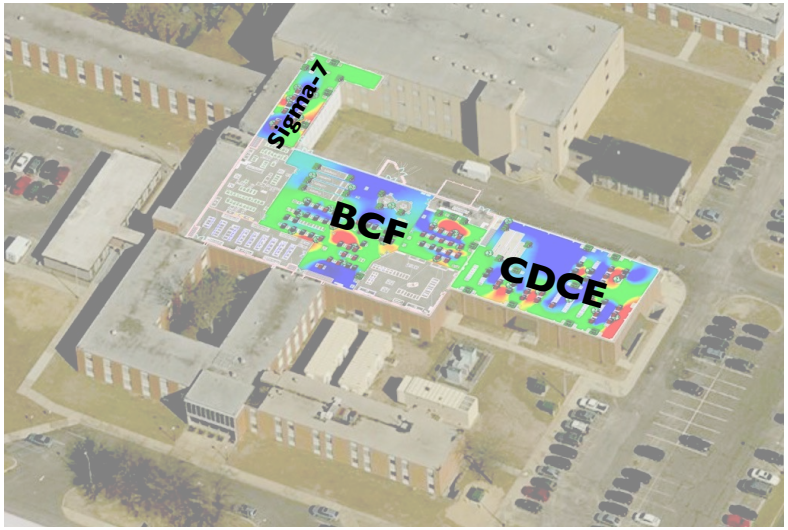

- RACF is the main component of the Scientific Data & Computing Center (SDCC) within the lab's Computational Science Initiative (CSI)
    - Other components include two new HPC clusters supporting research in LQCD, NSLS2, Biology, CFN, etc.
- Looking to hire add'l manpower — if interested please talk to Eric Lançon or Tony Wong

New Data Center Update

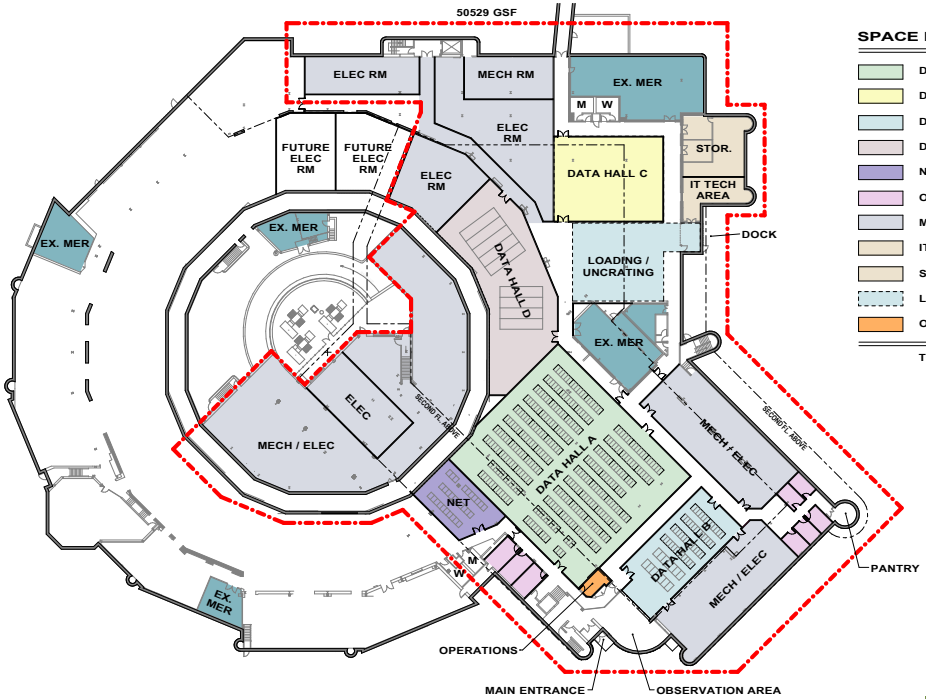

SPACE LEGEND	PROGRAM NET AREA (NSF)	DELINEATED AREA (NSF)
DATA HALL A	10800	9271
DATA HALL B	3450	2988
DATA HALL C	1950	2238
DATA HALL D	3600	4159
NETWORK EQUIPMENT	1000	1156
OFFICES	1680	1156
MECH/ELEC	20200	25007
IT TECH AREA	500	643
STORAGE	1200	1200
LOADING / UNCRATING	560	2579
OPERATIONS	150	132
TOTAL NET AREAS	48090	50529

Phase I : 60% Space Plan, 3.6 MW

- Includes CSI office space already in use
- SDCC operational target is FY22

- *CD-1 milestone approval received on April 17th*

Institutional Cluster (IC)

- The first HPC cluster available to the entire BNL community
  - Operational since January 4, 2017
- 108 compute nodes
  - Dual Xeon Broadwell (E5-2695v4) CPU's with 36 physical cores
  - Two NVidia K80 GPU's
  - 1.8 TB SAS drive + 180 GB SSD for temporary local storage
  - 256 GB RAM
- Non-blocking Infiniband EDR interconnect
- 1 PB of GPFS storage with up to 24 GB/s connectivity

Institutional Cluster (IC)

- Have fixed initial problems encountered with hardware failures, GPU performance issues, and support
- Currently ~120 registered users
- Cluster utilization approaching 95%
- Uptime nearly 100% over past three months
- Expansion under active discussion
  - Extent to be determined by expected demand
  - P100 instead of K80 GPU's

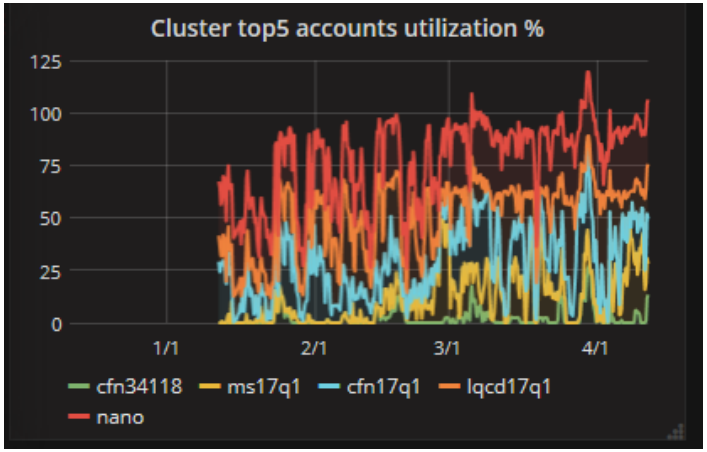

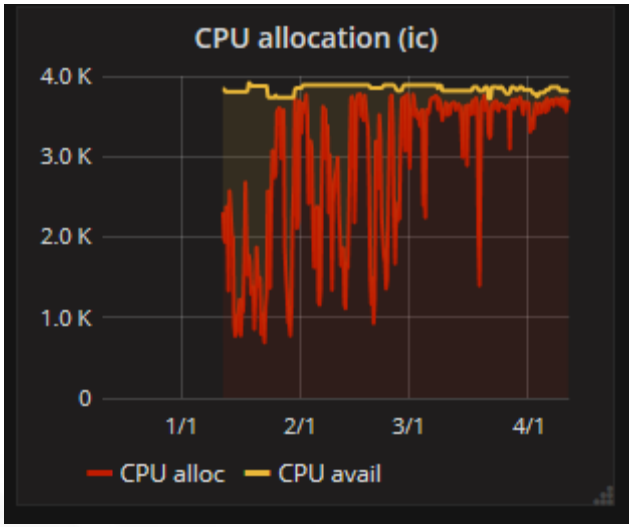

KNL Cluster

- 144 Nodes
- Unlike IC, achieving stability and performance with KNL has required significant dedicated effort
  - See William Strecker-Kellogg's talk on Wednesday for details
- Experimenting with different workloads (ATLAS, LQCD, accelerator physics, etc)
  - LQCD cluster to be purchased this Fall - KNL or GPU-based configurations being considered

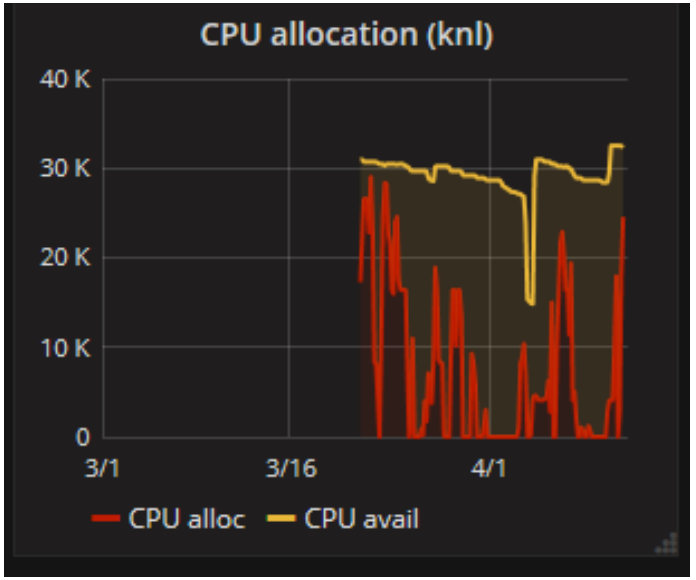

Network Reconfiguration

- "Science Core" network - 2x30 Tbps switching capacity
  - Initial connection to NSLS-II @ 4x40GbE
- Investigating migration to 25/50/100 GbE switching infrastructure (from 10/40/100 GbE)
- US Atlas Tier-1 facility successfully migrated outside of the BNL campus
  - Motivated by need to support IPV6 and evolving DOE cybersecurity requirements
  - Careful preparation ensured smooth migration that was transparent to ATLAS users
  - Remainder of the facility to migrate in CY2017

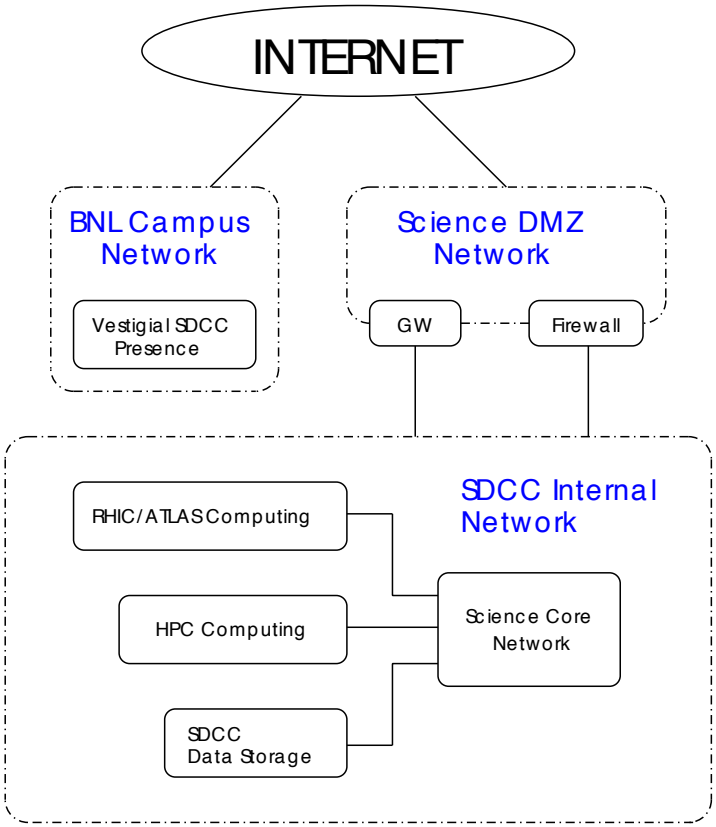

Tier-1 dCache & FTS Upgrade

- dCache currently managing 14.5 PB of unique data
- Upgraded in March from version 2.10 to 3.0.11
  - RHEL 6.x to 7.3; PostgreSQL 9.3 to 9.6
  - Support RFC VOMS proxy; improved SRM/NFS performance
- Some upgrade issues
  - High memory usage on admin node
  - XRootD front-end memory leak required version upgrade 4.4 to 4.6.0
- ***IPv6 dual stack now fully operational***

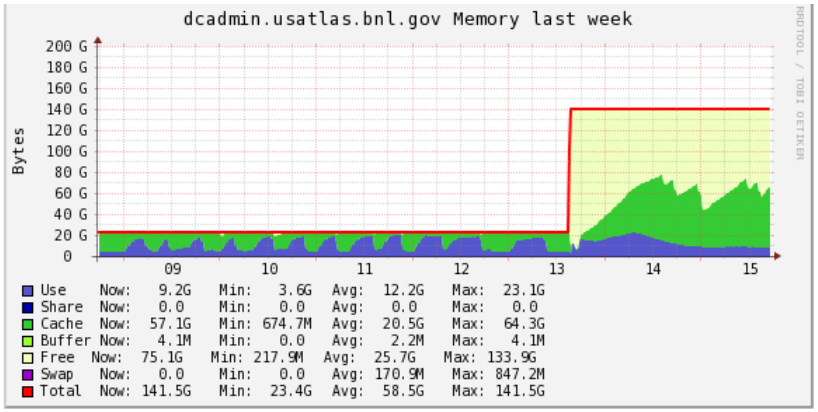

- Testing new features of GridFTP door to support Globus Online (thanks to Paul Millar)
- Will be testing integration with Ceph as well as new replication manager

HPSS

- Upgraded to 7.4.3 in November
- ~90 PB on 65K tapes
- STAR now uses LTO-7 for both Raw and DST; ATLAS migrating to LTO-7 in mid-May
- Converting Phenix Raw from LTO-4 (800 GB/ cartridge) to LTO-7 (6TB/cartridge)
  - Due to the large capacity of LTO-7, we are making 2 copies of these Raw data (mitigate risk of loss due to heavy access)

Tier-1 Staging Performance

- During ATLAS reprocessing campaign, BNL dCache staging limit (concurrent transfers) increased from 3k to 30k, reaching 1GB/s staging rate from HPSS
- dCache v3 allows non-blocking staging requests to be sent to HPSS – expect higher staging rate with much less system load in the future

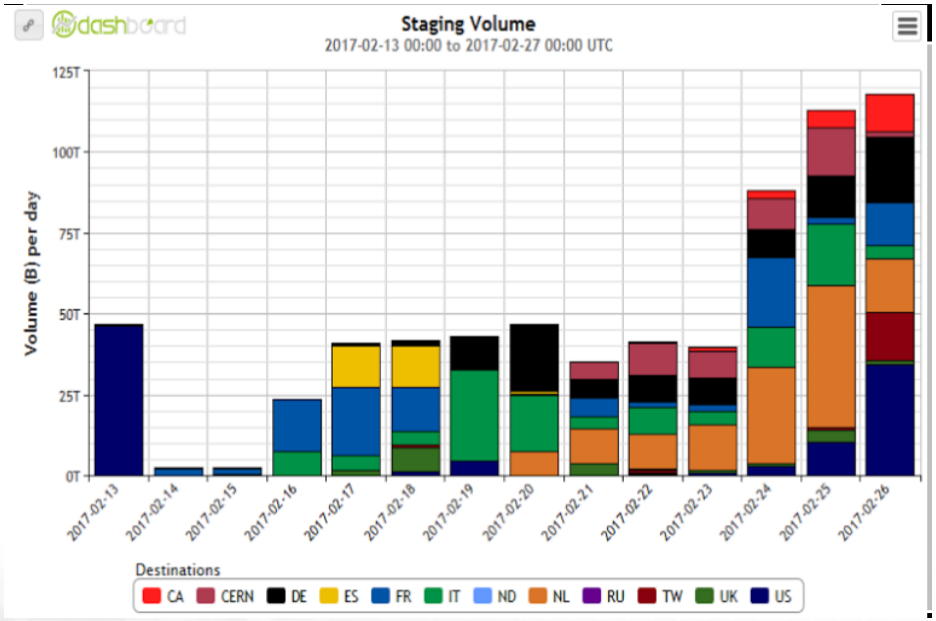

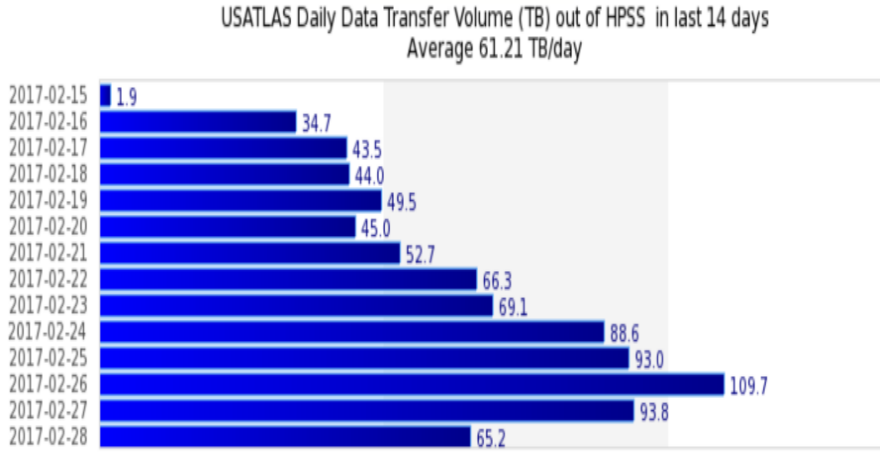

ATLAS AFS Phaseout

- *ATLAS Tier-1 jobs no longer using AFS*
  - CVMFS was already providing access to software releases and related files for ATLAS jobs at BNL
  - CVMFS Stratum Zero service created to host frequently used job files, including file transfer clients that were hosted in AFS
- Remaining AFS functionality being replaced by a multi-tiered solution based around CVMFS — phase out remaining dependencies this year
  - New CVMFS repo will cater to US ATLAS facility sites, with read-only access to files customized and collected specifically for US ATLAS job requests
  - Current CVMFS Stratum One service will continue to serve OSG and WLCG sites, with the additional functionality of hosting the new BNL repository
  - Writeable AFS functionality to be replaced by an ownCloud based, Ceph-backed, distributed storage solution (see Hiro Ito's BNL Box talk on Wed.)

CEPH Cluster

- Major Ceph cluster migration started in March
  - Installing new CephFS cluster as a storage backend for BNL Box
    - 60 x 8 TB SAS HDD JBOD arrays
    - Currently 3.8 PB raw/1.3 PB usable; 7.5 PB/ 2.5 PB by end of 2017
    - First Ceph cluster at RACF deployed entirely on new hardware w/o HW RAID

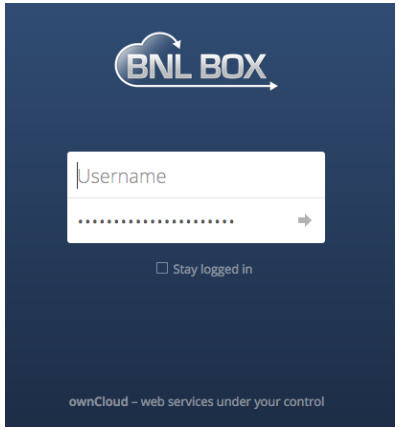

  - Repurposing “old” ATLAS Ceph cluster (1.2 PB/0.4 PB) as ATLAS "test" cluster (new OSD nodes, upgraded public network layout)
  - Replacing storage backend of ATLAS “new” cluster with newly retired dCache storage — will be 5.7 PB/1.9 PB once the process is finished
- Targeting Ceph 11.2.x (Kraken) on RHEL7 for all 3 clusters — currently Ceph 9.2.1 (Infernalis)
- Migration will require two reassignments of 'cephgw.usatlas.bnl.gov' alias used by ATLAS Event Service — expected in May.
  - Current plan assumes no need to preserve content of 'atlas*' buckets during migration sequence

Linux Farm

- New ATLAS Systems brought online this month (identical to last reported purchase)
  - 48 new Dell R430 systems
  - 2 Broadwell E5-2690v4 2.6 GHz CPUs (56 logical cores total)
  - 128 GB DDR4 2400 MHz RAM
  - 4 3.5" 2 TB 7200 RPM 6Gbps SATA drives
- RACF total now 59K cores
- Considering migration of RHIC hosts to SL7 before the next run (late Fall)

Monitoring

- Rolling out collectd deployment
  - Evaluated diamond, but process-per-metric too cumbersome
  - collectd easy, low resource usage, stable
- Grafana working out very well — deployed public and private instances
  - Per-job metrics and advanced node statistics from per-disk iops to slab cache usage

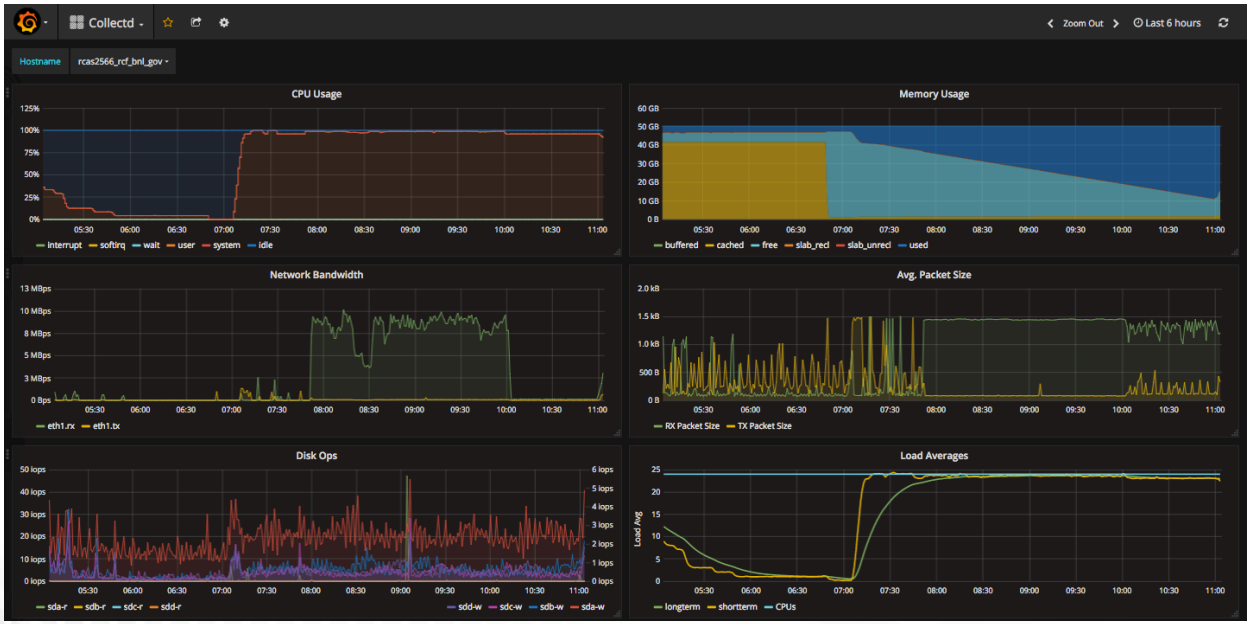

Singularity

- Linux container management software developed by LBL and SLAC
- Like Docker, LXC, etc., utilizes Linux namespaces to implement process containerization
- Developed primarily to solve the problem of HTC/HPC job mobility
    - Seamlessly redirects all I/O in and out of the container directly between environments — easy to handle MPI, X11 forwarding, etc.
- Run jobs on any Linux system, regardless of distribution or version, as long as namespaces are supported and Singularity installed
- Extremely lightweight — no associated daemons, one single package to install
- Does not support user escalation or context changes — all processes run as yourself in the container, not as root
- http://singularity.lbl.gov/

Example:

```
$ whoami
testuser
$ cat /etc/redhat-release
Scientific Linux release 7.2 (Nitrogen)
acas1801:~$ singularity shell /images/atlas_sl6.img
Singularity: Invoking an interactive shell within
container...

Singularity.atlas_sl6.img> cat /etc/redhat-release
Scientific Linux release 6.8 (Carbon)
Singularity.atlas_sl6.img> whoami
testuser
```

Singularity

- ATLAS SL6 test container for our environment created with necessary bindmounts (pnfs etc.)
  - Working to make this available via our CVMFS Stratum Zero server.
- Have set up a test SL7 system which forces execution of jobs in the ATLAS SL6 container via HTCondor's USER_JOB_WRAPPER configuration parameter
  - Demonstrated successful execution of HammerCloud jobs within the container
- In the process of expanding the test to ~30 hosts, and creating a new test Panda queue for real ATLAS jobs — update expected at Fall HEPiX

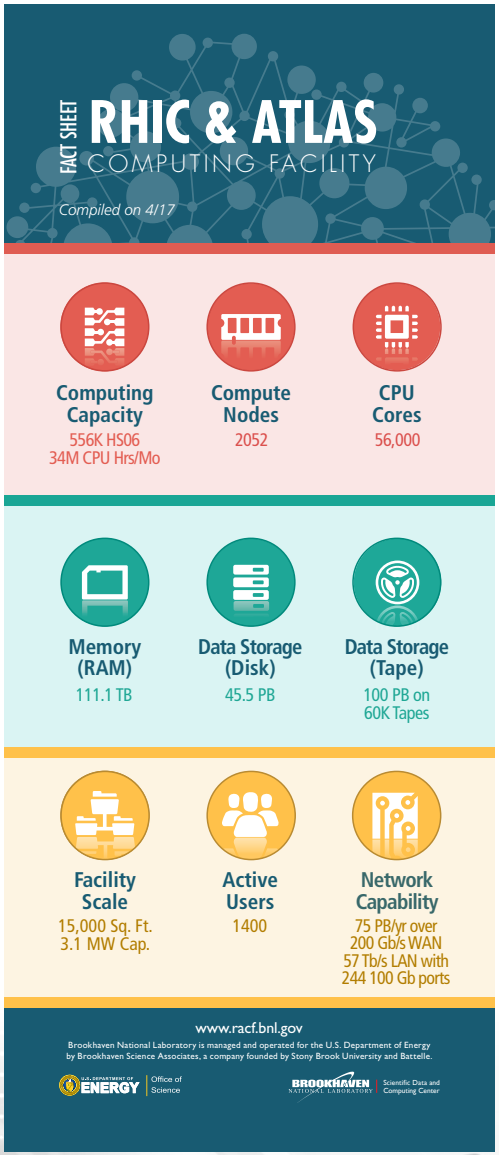

Many thanks to John DeStefano, Chris Hollowell, Hiro Ito, Shigeki Misawa, Jason Smith, Will Strecker-Kellogg, Tony Wong, David Yu, Alex Zaytsev and Xin Zhao for contributing the content of this presentation.