

BNL SDCC Site Report

David Yu

Spring HEPiX,

University of Wisconsin-Madison

May 14, 2018

BROOKHAVEN
NATIONAL LABORATORY

 U.S. DEPARTMENT OF
ENERGY


Scientific Data and Computing Center

- Service Operation For:
 - RHIC, LHC ATLAS, BELLE II, BER ARM, LQCD, RIKEN, BES Center for Functional Nano Materials, National Synchrotron Light Source II, National Nuclear Data Center, Simons Foundation,...
- ~1700 users from 20+ projects



SDCC Overview


- Located at Brookhaven National Laboratory on Long Island, NY
- Provides full service computing mainly for:
 - RHIC experiments (Tier-0): STAR, PHENIX
 - ATLAS (US Tier-1)
 - Belle-II (Tier-1)
 - Other groups: NSLS-II, CFN, LSST, Daya Bay, DUNE, EIC, etc.
- RHIC Run 18 expected to be completed in mid June



Belle II

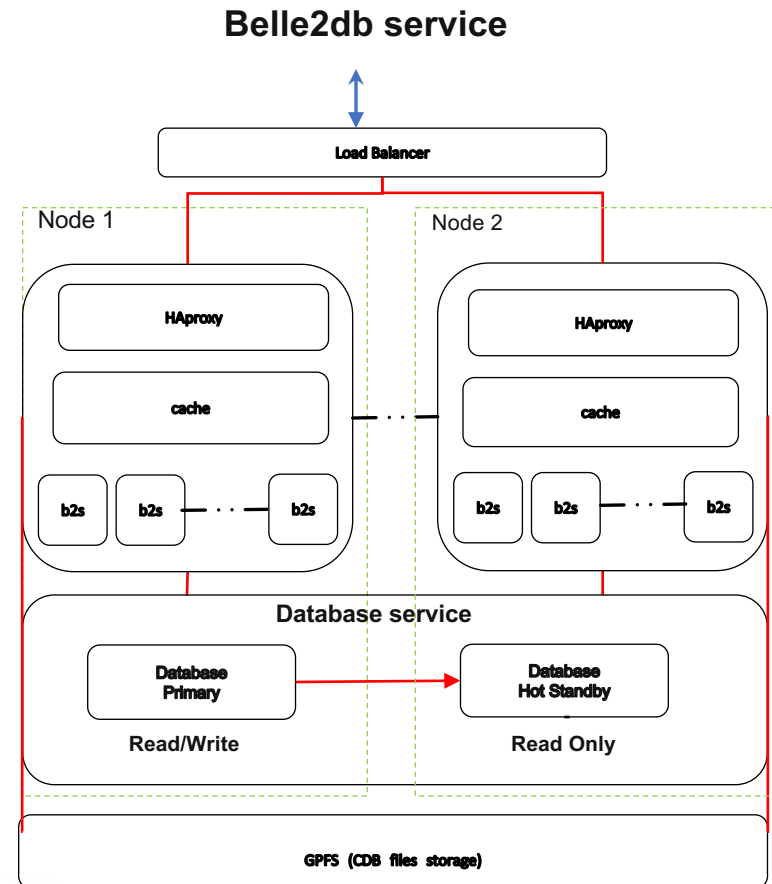
BNL is transitioning to become Belle II Tier1 computing center.

- Year-long transition of services from PNNL was started on September 2017.
- The compute and storage resources at BNL become the part of the Belle II Tier1 last year (2017).
- The production DDM DIRAC service was transitioned to BNL by the end of April, 2018.
- The conditions DB service will be transitioned to BNL by the end of May 2018. This will complete the transfer of all production services from PNNL to BNL.
- The transfers of the knowledge associated with the Belle II computing operations are on-going. They will be completed by the end of Sept, 2018.



Belle II Conditions Database (CDB) Metadata service

- Architecture deployed using Kubernetes / Docker framework and controlled via puppet
 - kubelet-1.9
 - docker-1.12
 - b2s: Java application on Payara micro 4.1.2
- Database replicated for reliability purposes
 - Postgres 9.6 hot standby replication
- Hardware:
 - 2 Nodes Dell R730xd:
 - Two Intel(R) Xeon(R) CPU E5-2667 v4 @ 3.20GHz -> total CPU thread 32
 - Memory 256 GB
 - Disk for Database Intel DC P3700
 - NVMe SSD on PCI Express 800 GB
 - 20Gb/s Channel Bounded connectivity

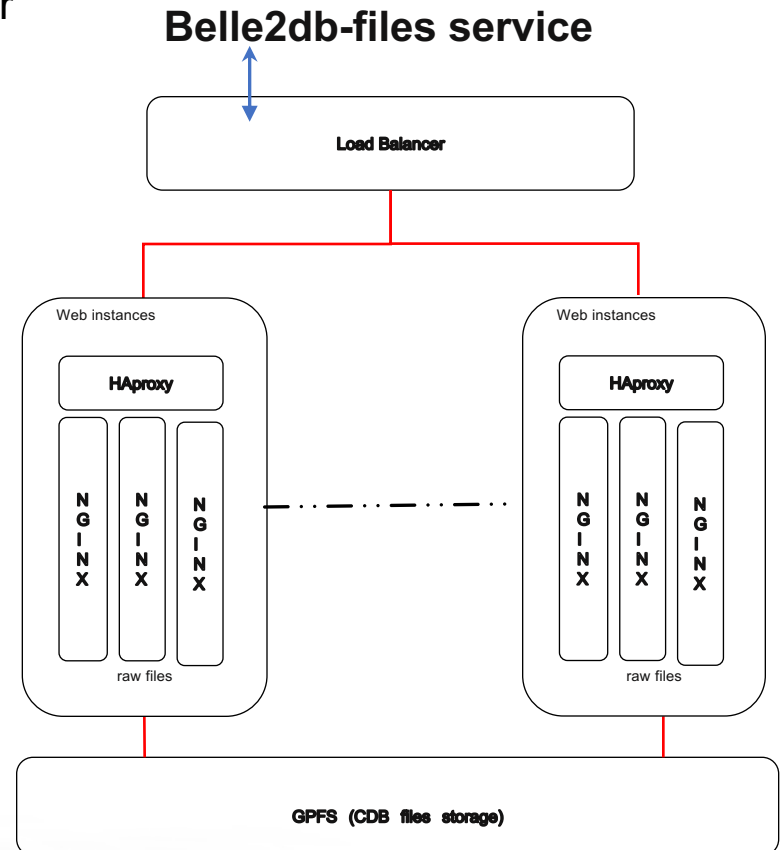


Belle II Conditions Database (CDB) Payload service

- Architecture deployed using Kubernetes / Docker framework and controlled via puppet
- General Parallel File System (GPFS) used to store CDB payloads

Hardware:

- 2 Nodes Dell R430
Two Intel(R) Xeon(R) CPU E5-2650 v4
- @ 2.20GHz total CPU threads 48
Memory 256GB
- 20Gb/s Channel Bounded connectivity



CDB services expected to be integrated into production at the end May 2018

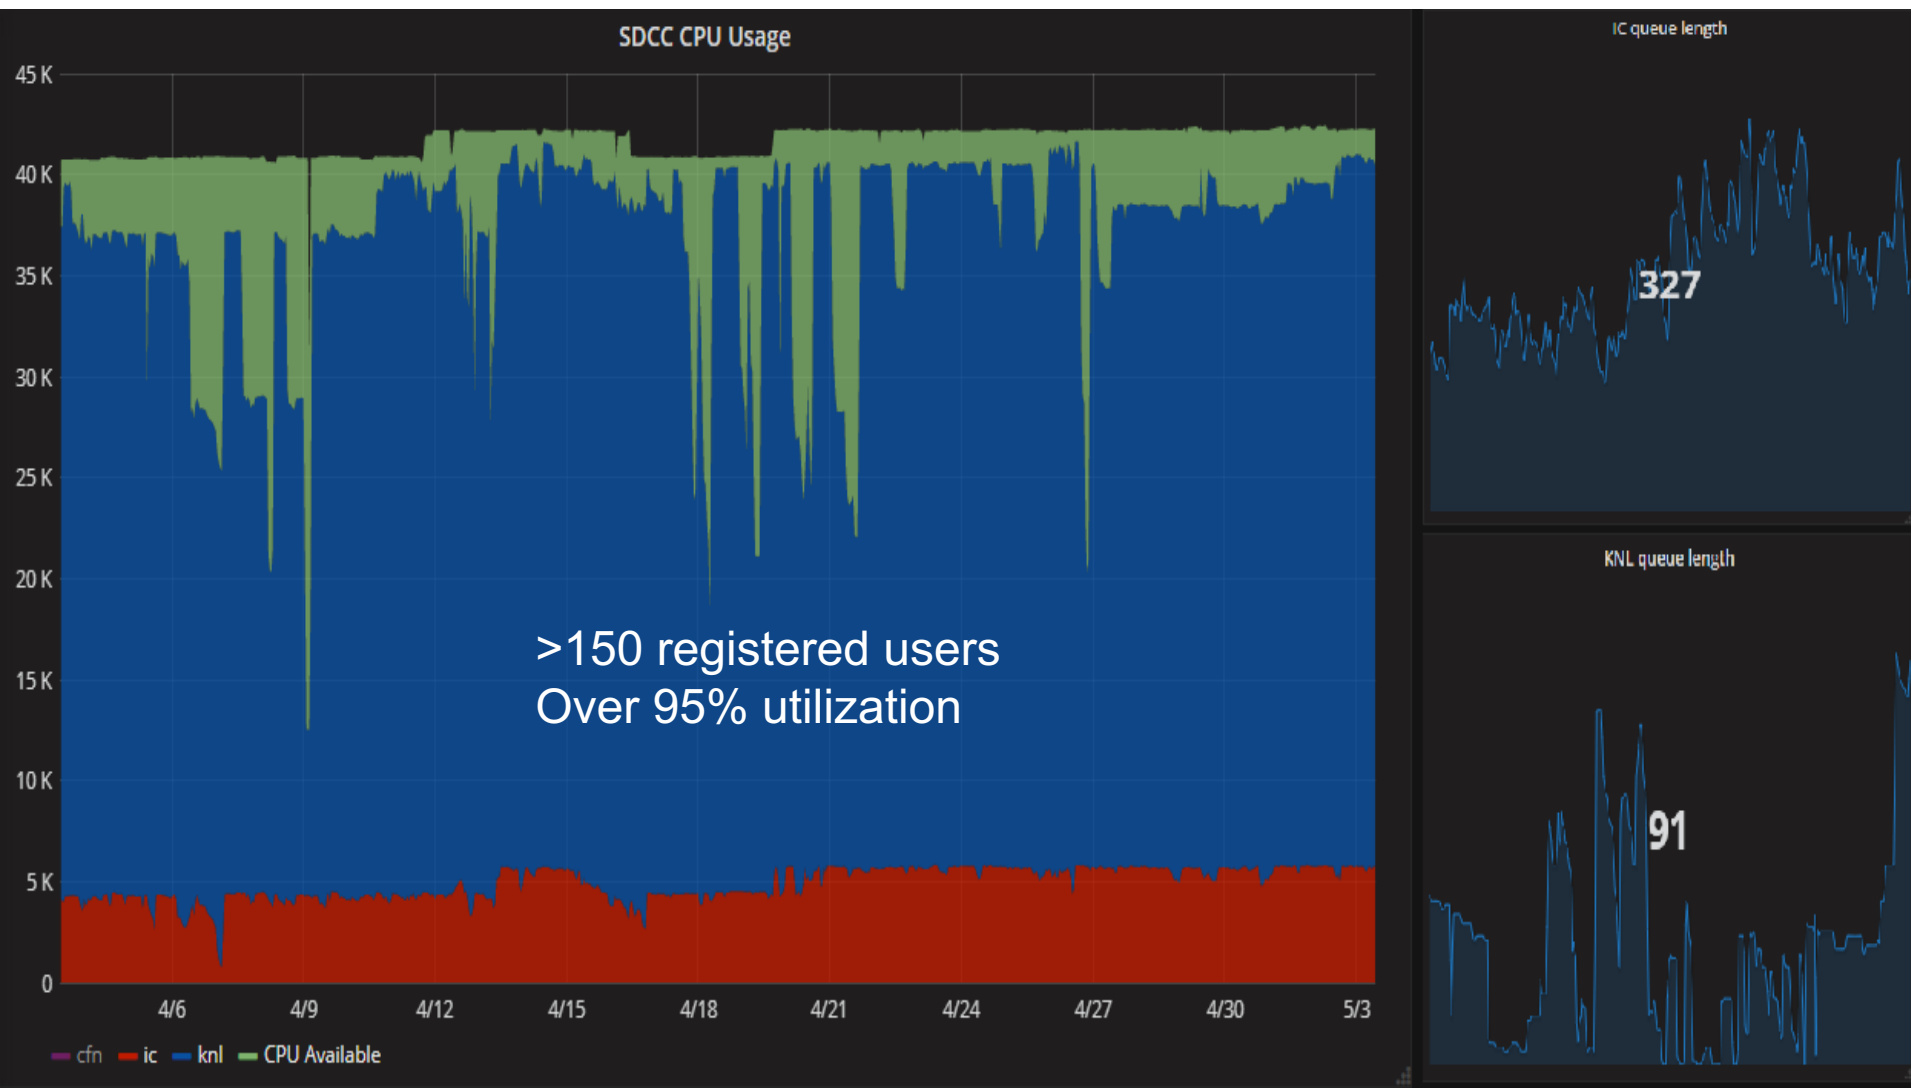
HPC Efforts

Three HPC components:

- **Institutional Cluster** - 162 x 36 core Broadwell nodes; EDR Infiniband
 - 108 with 2 x Nvidia K80 GPU
 - 54 with 2 x Nvidia P100 GPU, increasing to 108 (last 3 racks to be shipped in summer 2018)
- **KNL Cluster** - 144 x 64 core (256 HT) Xeon Phi 7230 1.3 GHz nodes; OPA interconnect
 - Stability and performance issues, mostly due to OPA, now under control
 - HS06 Benchmarking of KNL indicates factor 3 lower performance compared to Skylake
- **Sky Cluster (LQCD)** - 64 x 36 core Skylake; 192 GB; EDR Infiniband. To be delivered next month




HPC Cluster Usage




HPC Cluster Usage

- KNL Cluster



- Institutional Cluster (IC)



Linux Farm

New equipment arrived in Feb 2018 and in production

- RHIC - 34 Dell PowerEdge R740xd Servers, 2448 Job Slots, ~35 kHS06
 - 10 x 3.5" 8TB 7.2k RPM 6 Gbps SATA HW RAID5
 - 2 x 240 GB SSDs (for OS - HW RAID1)
 - PERC H730P+ RAID Card
 - 10 Gbps ethernet
- ATLAS - 90 Dell PowerEdge R640 Servers, 6480 Job Slots, ~93 kHS06
 - 2 Xeon Gold 6150 (Skylake) 2.7 GHz CPUs (72 logical cores total)
 - 192 GB DDR4 2666 MHz RAM
 - 4 x 3.5" 4TB 7.2k RPM 6 Gbps SATA
 - PERC HBA330 Storage controller
 - 1 Gbps ethernet
 - Issue with all of the Broadcom 5720 NICs: NMIs generated randomly. During use, resulting in reboots. Occurred in both OS and PXE ROM
 - Dell determined we had a bad batch of NICs. Sent a technician to replace all 90 daughter-cards with Intel I350s
- Belle II – one rack of 27 Dell R640 servers expected to arrive this month
 - Same configuration as the R640s purchased for ATLAS



Linux Farm


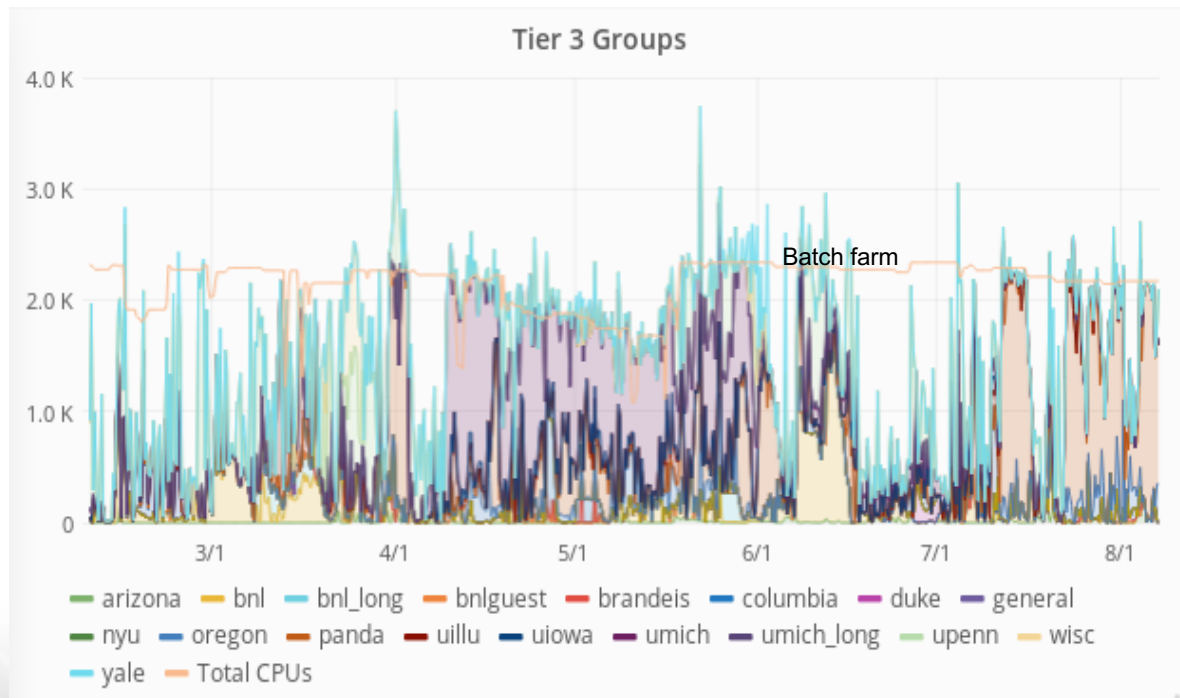
- Some older equipment moved to a “shared HTCondor pool”, usable by all of our HTC customers
 - ~2,800 job slots
 - ~30 kHS06
 - Guaranteed job runtime of 2 days
 - Near 100% utilization
- SL7 Upgrade Status and Plans
 - RHIC
 - STAR and small experiments upgraded to SL7 in Dec/Jan
 - PHENIX SL7 upgrade expected in June
 - ATLAS
 - Phased migration to SL7 in progress
 - Expected completion by the end of May
 - Jobs on SL7 hosts forced to run in SL6 Singularity containers to simplify migration



ATLAS Analysis facility at BNL

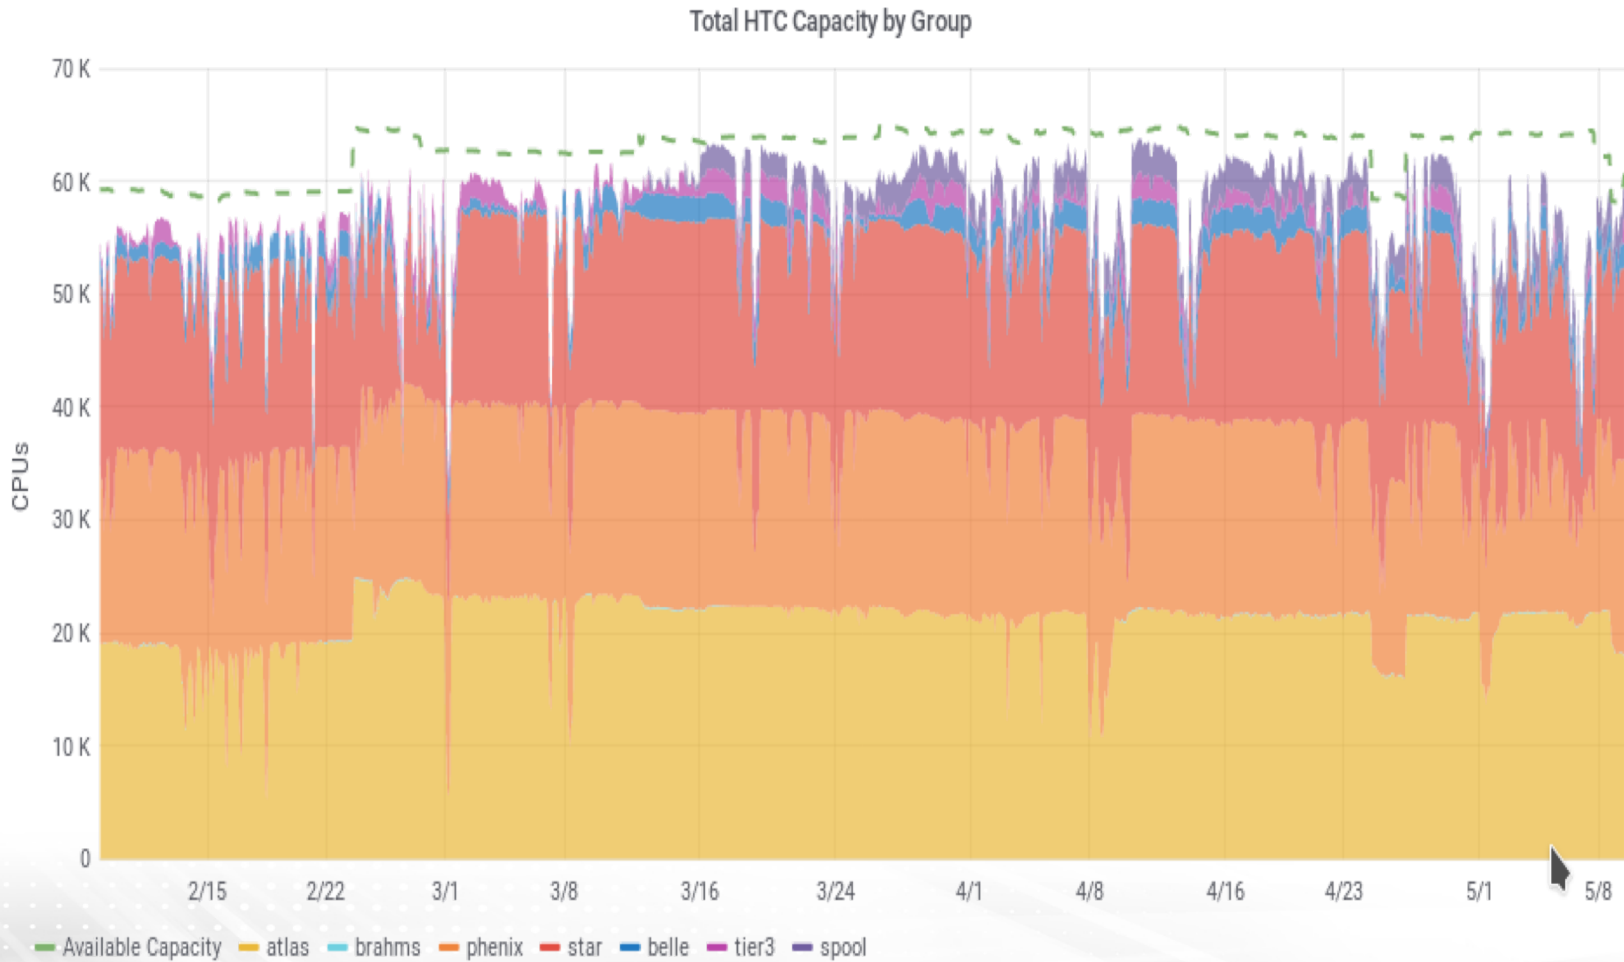
- Interactive nodes + batch farm
- ~300 users from a dozen of institutions

Groups past 6 months



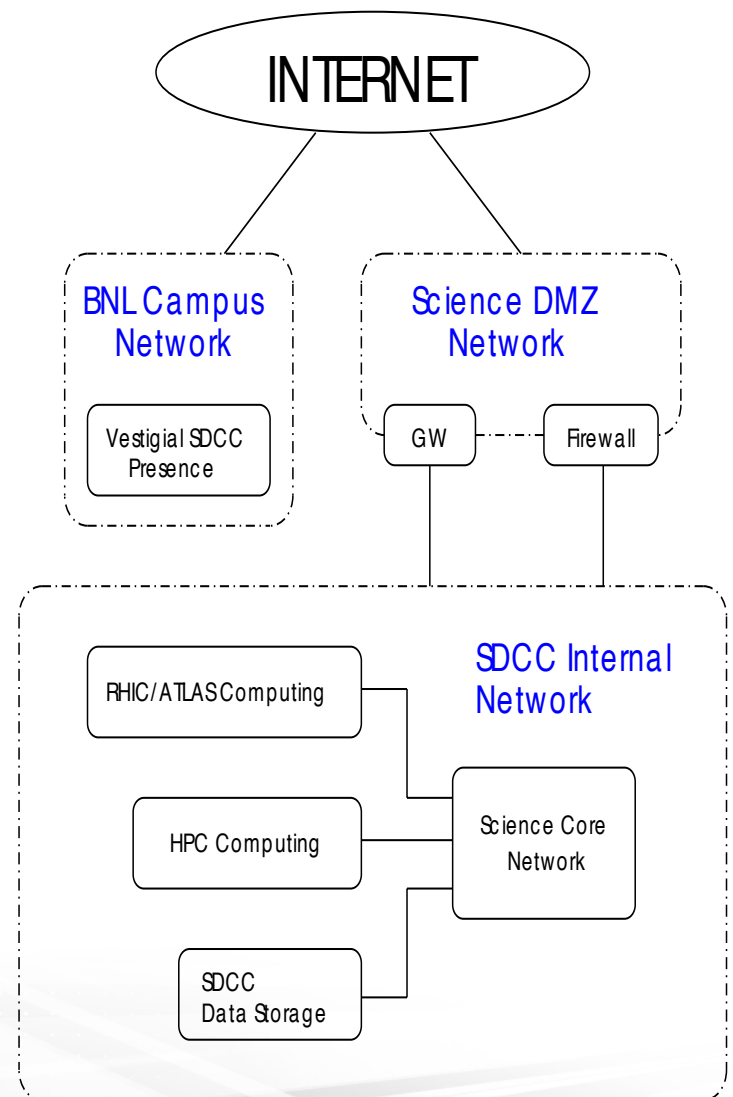
Linux Farm

- Last 3 months of farm utilization by experiment.



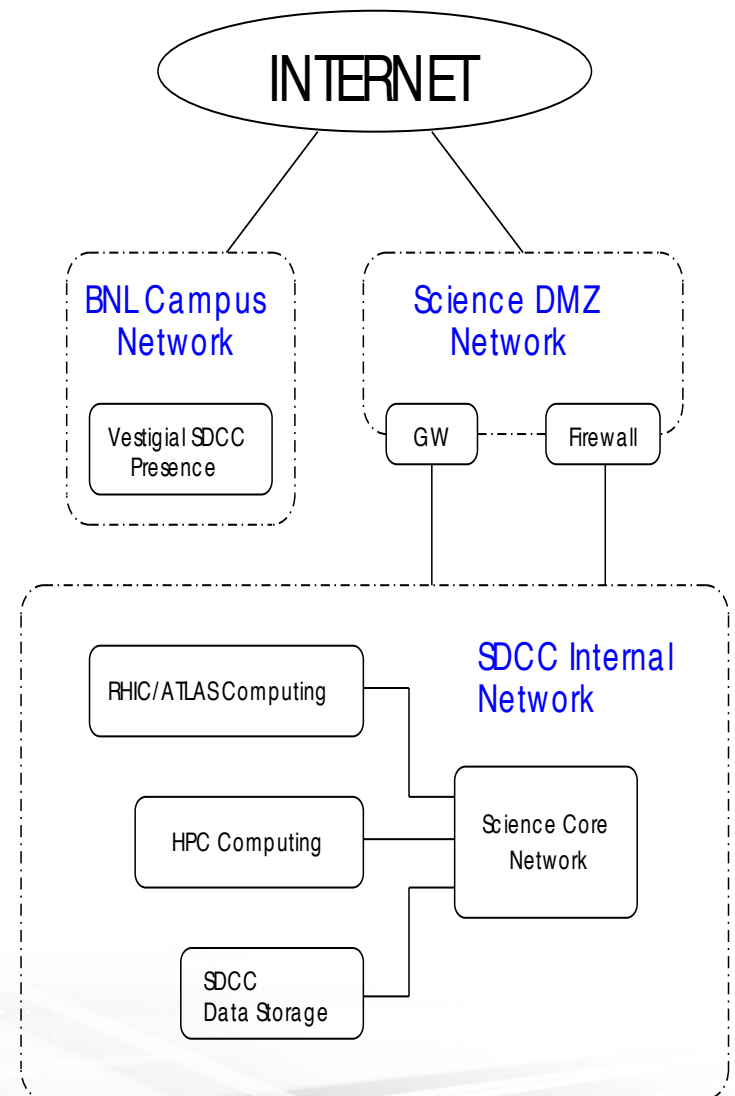
Network Reconfiguration Update

- RHIC facility was migrated outside of BNL campus network and a merge of RHIC and ATLAS network front-ends was performed in Dec'17
- Only a limited set of systems supported by the SDCC do still remaining on BNL Campus now
- The construction of 25/50/100 GbE enabled Storage Block of the Science Core was completed in Feb'18: migration of BNL ATLAS dCache storage system into it was completed in Mar'18, followed by migration of all other storage systems of the Facility (CVMFS, Belle II dCache, Ceph/CephFS clusters, RHIC GPFS, and eventually HPSS)



Network Reconfiguration Update

- All new Ethernet connected servers that require high bandwidth procured by the Facility in FY18 are configured with either 2x 25 GbE (LACP) or 4x 25 GbE (LACP) interfaces (configurations based on 2x 50 GbE are under evaluation, yet they seem to have limited usefulness for us so far)
- 1 GbE attached Farm compute nodes are migrating to Equal Cost Multi Path (ECMP) spine-and-leaf (ToR switch based) network infrastructure
 - The migration of BNL ATLAS Tier-1, BNL Belle II Tier-1 and BNL ATLAS Tier-3 compute node racks to spine-and-leaf layout is expected to be completed by the end of May'18
 - The conversion of all the RHIC Farm compute nodes that are 1 GbE attached is scheduled to happen before the end of CY18.
- Future consolidation of the majority of our Facility network distribution infrastructure on an upgraded and adequately scaled spine group is foreseen for the new B725 based data center

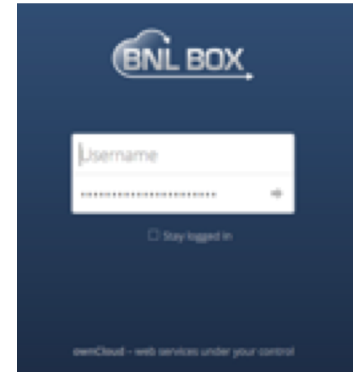


Distributed Storage

- ATLAS Tier-1 dCache
 - core server version: 3.0.11, NFS version: 3.0.38, pool version: 3.0.43
 - 17.5 PB of unique data
 - core server in SL7, phased migration to SL7 for pool hosts in progress
 - Deployed a newly written non-polling HSM backend component to reduce resource usage on pools in production
 - Pool version upgraded from 3.0.11 -> 3.0.43 to fix login handshake to support xrootd clients (> 4.7.0) after the ATLAS software was updated
 - Xrootd write functionality tested through.
 - GSI proxy needed
 - Await for production deployment
- Belle II Tier-1 dCache
 - core server version: 3.0.11, 1.7 PB deployed
- Simon's Foundation dCache
 - core server version: 2.1.6, 0.26 PB deployed in Ceph storage
- STAR XrootD
 - Major upgrade to v4.7.1 in November, just ahead of farm upgrade to SL7.
 - Had to rewrite STAR's custom Name2Name (N2N) library to compile against new XrootD codebase.
 - Deployment had to be set up to support STAR's 32-bit and 64-bit clients on both SL6 and SL7
 - Upgrade under SL6 was seamless with no impact on users. All data on farm (~8PB) was preserved successfully through upgrade to SL7.

Status of CEPH Clusters

- CephFS/BNLBox cluster: Provide CephFS capacity for BNL Box (no major changes since the last report):
 - 5.1 PB RAW, Luminous 12.2.5
 - 12 Gb/s SAS attached RAIDInc JBOD arrays
 - Public Network: dual attached 10/40 GbE; OSD network: dual attached 4x FDR IB (56 Gbps) with 1+1 redundancy
- Ceph ATLAS Prod cluster is being re-designed to provide the 3 PB scale test for Erasure code based pools to be used as a backend for the production BNL ATLAS dCache and to support the AMZ/S3 clients (ATLAS Event Service) at the same time:
 - 4.5 PB RAW, [to be deployed under Luminous 12.2.x \(currently not serving the users\)](#)
 - 8 GB/s FC and HW RAID based Nexsan arrays
 - Public Network: 2x 25 GbE; OSD Network: 4x 25 GbE
- Ceph ATLAS Test cluster:
 - 0.45 PB RAW, Kraken 11.2
 - 4 Gb/s FC and HW RAID Nexsan arrays
 - Public Network: 10 GbE; OSD network: 4X FDR IB (56 Gbps)



Centralized Storage

Items to be done in one month

■ IBM Spectrum scale (GPFS)

- Will be updating to GPFS version 4.2.3.8 from 4.1.1.13.
- Hardware refresh for STAR GPFS.
 - Six new Dell R740 servers with 100Gb connectivity on each server.
 - Replacing Hitachi unified storage (HUS130) with 2 X Netapp E5700's.
- Merging STAR filesystems into a single GPFS filesystem of capacity 4PB and performance 42GB/sec.

■ Home directory storage

- Replacing Hitachi HNAS 4100 with Netapp all flash AFF 300.
- Migrating all user data from NFS to GPFS and reserving NFS exclusively for home directories.


Centralized Services

- **RHEV Cluster**

- Currently hosts 300+ VMs, mostly infrastructure services — OpenAFS, auth, web, ssh gateways, ATLAS grid gatekeepers, etc.
- Running latest RHEV version 4.1, patched for Meltdown and Spectre vulnerabilities.
- ATLAS and RHIC virtual clusters will be merged into a single cluster.
- Non-production hosts will be isolated from production hosts.
- RHEV VMs will run on Netapp A300 all flash storage over NFSv3. (in a month)

Tape Storage (HPSS)

- Running HPSS v7.4.3 since Nov 2016
- ~120 PiB on 66K tapes
- Ready to accept Belle-II data
- Taking data for RHIC Run 18
- Evaluating HPSS 7.5.1
- Researching new tape libraries, and tape drives.
- Evaluating RAO - Recommended Access Order



■ Phenix Raw (15,202 TB)	■ Phenix DST (6,509 TB)	■ Phenix Archive (492 TB)
■ Star Raw (33,958 TB)	■ Star DST (21,483 TB)	■ Star Archive (793 TB)
■ Atlas (33,965 TB)	■ Star Raw 2nd Copy LTO-7 (1,502 TB)	■ Star Raw 2nd Copy T10KD (1,642 TB)
■ Phenix Raw 2nd Copy T10KD (793 TB)	■ Phenix Raw 2nd Copy LTO-7 (4,160 TB)	■ Phobos Raw (140 TB)
■ Star 2nd Archive (793 TB)	■ RACF Archive (514 TB)	■ RACF Archive 2nd copy (514 TB)
■ Star 2nd Archive (793 TB)	■ Phenix 2nd Archive (491 TB)	

Tape Storage (HPSS)



- Both STAR and ATLAS currently using LTO-7
- PHENIX DST remaining on LTO-6 due to lower volume
- Reclaimed 772 cartridge slots since October through migration of old LTO-4 (800 GB) data to dual-copy LTO-7 (6 TB) cartridges
- Completed migration of PHENIX RAW
- Processing STAR DST, stopped at the end of December, preparing for RHIC run.




Data volume in last 12 months

- Data import : 42 PB
- Data export : 45 PB

BNL WAN Traffic



BROOKHAVEN
NATIONAL LABORATORY



Other BNL SDCC presentations

- **Changing Compute Landscape at Brookhaven**
William Edward Strecker-Kellogg, 5/16 14:00
- **HPL and HPCG Benchmark on BNL linux farm and SDCC**
Zhihua Dong, 5/16 16:50
- **BNL New Data Center - Status and Plans**
Tony Wong, 5/17 09:20
- **Teraflops of Jupyter: A Notebook Based Analysis Portal at BNL**
Ofer Rind, 5/17 12:10

Other BNL SDCC News

BROOKHAVEN
NATIONAL LABORATORY

Scientific Data and
Computing Center

Since last Fall HEPiX meeting at KEK,

- 3 new staff members have joined the SDCC.
- 2 more will join this summer
- Still looking for more effort to meet programmatic needs. If interested, please contact Eric Lancon (elancon@bnl.gov).

Many thanks to Costin Caramarcu, Carlos Gamboa, Chris Hollowell, Hiro Ito, Jane Liu, Will Strecker-Kellogg, Tejas Rao, Ofer Rind, Iris Wu and Alex Zaytsev for contributing the content of this presentation.

