BNL Scientific Data and Computing Center (SDCC) Site Report

Ofer Rind
On behalf of SDCC, BNL
October 16, 2023
HEPiX Fall 2023, Univ. of Victoria, Canada

@BrookhavenLab

Scientific Data and Computing Center Overview

- Located at Brookhaven National Laboratory (BNL) on Long Island, New York
- Tier-0 computing center for the RHIC experiments
  - sPHENIX – began taking data in 2023
  - BNL is host site for the future Electron-Ion Collider (EIC)
- US Tier-1 Computing facility for the ATLAS experiment at the LHC
  - Also one of the ATLAS shared analysis (Tier-3) facilities in the US
- US Data center for the Belle II experiment at KEK
- Computing facility for NSLS-II photon science
- Providing computing and storage for proto-DUNE/DUNE along w/ FNAL serving data to all DUNE OSG sites
- Providing computing resources for a number of smaller experiments in NP and HEP
- Serving more than 2,000 users from >20 projects

PHENIX

sPHENIX

ATLAS EXPERIMENT

Belle II

DUNE DEEP UNDERGROUND NEUTRINO EXPERIMENT

National Nuclear Data Center

USQCD US Lattice Quantum Chromodynamics

ARM CLIMATE RESEARCH FACILITY

RBRC RIKEN BNL Research

We are hiring!

Brookhaven National Laboratory

2023Q4 (proj.)

National Synchrotron Light Source II, CryoEM

0.8 Tbps

Center for Functional Nanomaterials

0.4 Tbps

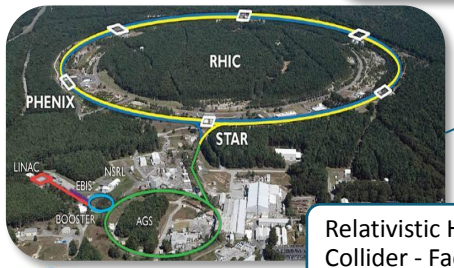

Relativistic Heavy Ion Collider - Facilities

0.7 Tbps

Brookhaven National Laboratory

Science DMZ

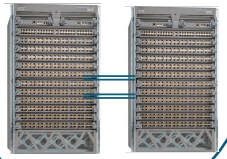

2x 400 Gbps

0.2 Tbps

1.2 Tbps

DTNs
(52 systems)

High Throughput Science Network

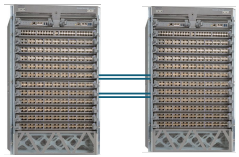

400 Tbps capability
(60 Tbps deployed)

1.2 Tbps

12 Tbps

Disk Storage
125+ PB

0.5 Tbps

18 Tbps

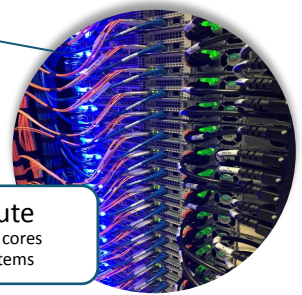

Compute
125k CPU cores
5 HPC systems

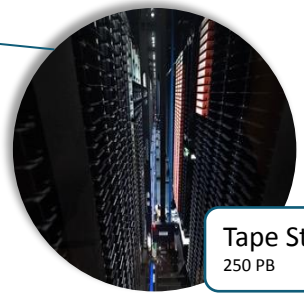

Tape Storage
250 PB

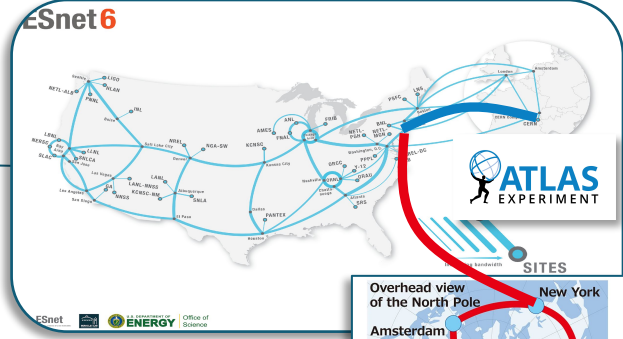

BNL Core Facility Revitalization (CFR) Project: New Data Center

New Data Center (Building 725) — 2023Q4: 2+ Years of Production Operations

- CFR project completed construction at end of FY21 with functional occupancy in 2021Q4
- Occupancy ramped up in 2022Q1-2023Q4 to the level of 104 racks populated with equipment in the B725 Main Data Hall (MDH)
  - 14 more storage / infrastructure racks are to be commissioned 2023Q4-2024Q1
  - Up to 27 more new HTC CPU racks and up to 4 more HPC CPU racks can be added 2023Q4-2024Q4
- Two library rows in B725 Tape Room are populated with IBM TS4500 tape libraries to serve ATLAS and sPHENIX (4 libraries, 128 tape drives in total).
  - One more library row (2 more 8-frame IBM TS4500 sPHENIX libraries) is expected to be added in FY24 and populated with FY25 sPHENIX data shortly after
- The transition of the majority of CPU and DISK to the new datacenter and retirement of the oldest parts of the old (B515) datacenter is now expected by early Nov 2023
  - Delayed due to late delivery of FY23 CPU purchase in late Sep 2023; this was needed to retire the rest of RHIC CPU located in the old datacenter
  - The retirement of B515 based HPC systems is currently well underway

Network Systems of B515 and B725 Data Centers (since Sep'23)

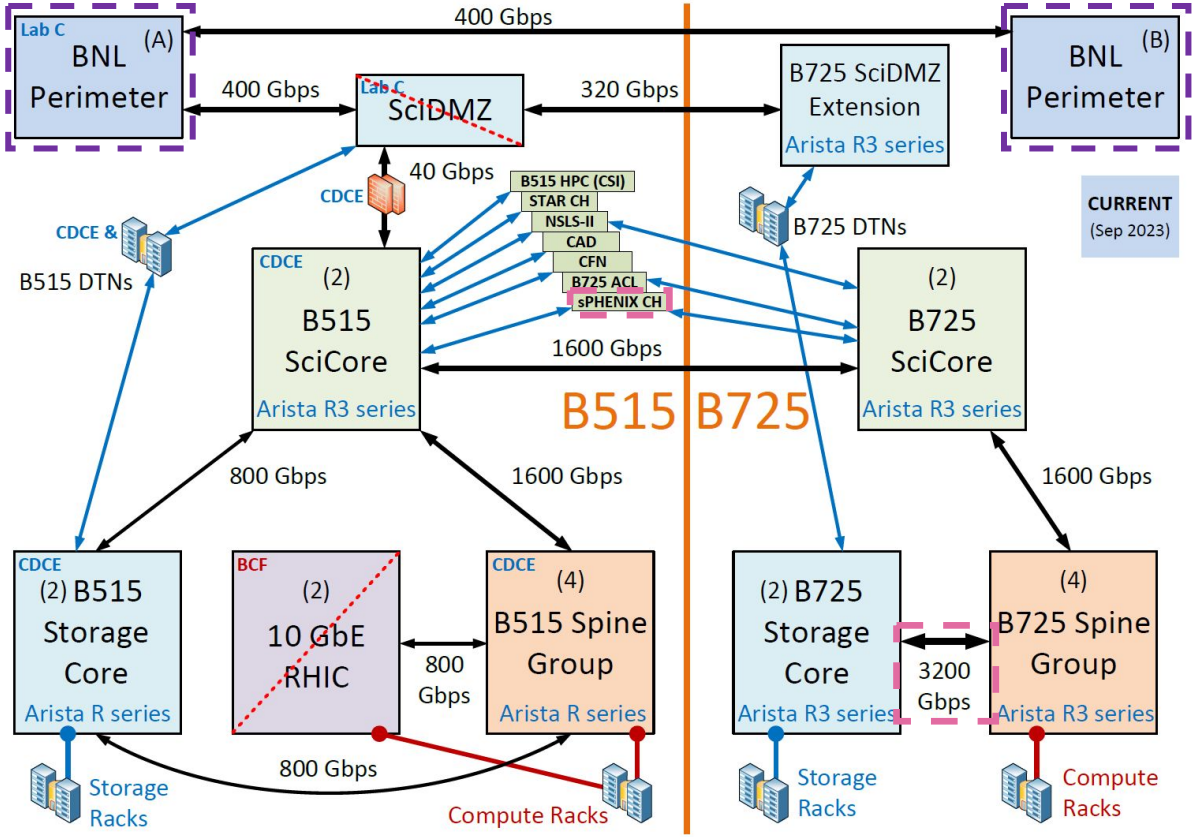

- BNL Perimeter (ESnet) was hardware refreshed with (400 GbE enabled) Nokia routers and now physically split between two different buildings on site for maximum redundancy and survivability
- sPHENIX Counting House (CH) uplinks to both B515 & B725 data centres are active at 600 Gbps combined
- The interswitch link between B725 Storage Core and the Spine Group is upgraded to 3.2 Tbps due to sPHENIX needs

Network Systems of B515 and B725 Data Centers (Dec'23 Proj.)

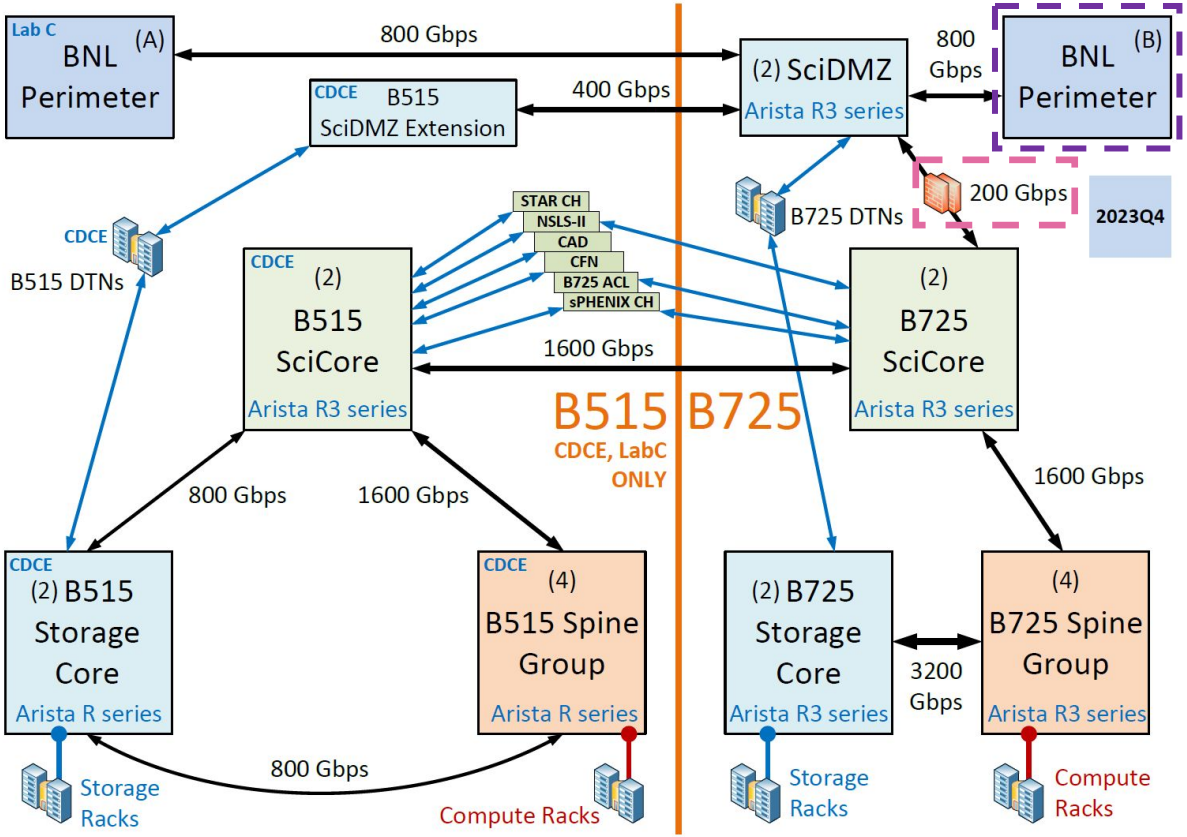

*The state expected to be achieved after **Dec 19, 2023** 15h long SDCC Facility-wide disruptive network intervention*

- Default routes for SciZone and SciDMZ are to be moved to B725 side
- Migration to the new 400 GbE enabled WAN infrastructure is to be completed (2x 400G with guaranteed BW of 400 Gbps)
- SciZone Perimeter Firewall is to be physically moved to B725 and upgraded to 200 Gbps of pass-through BW capability

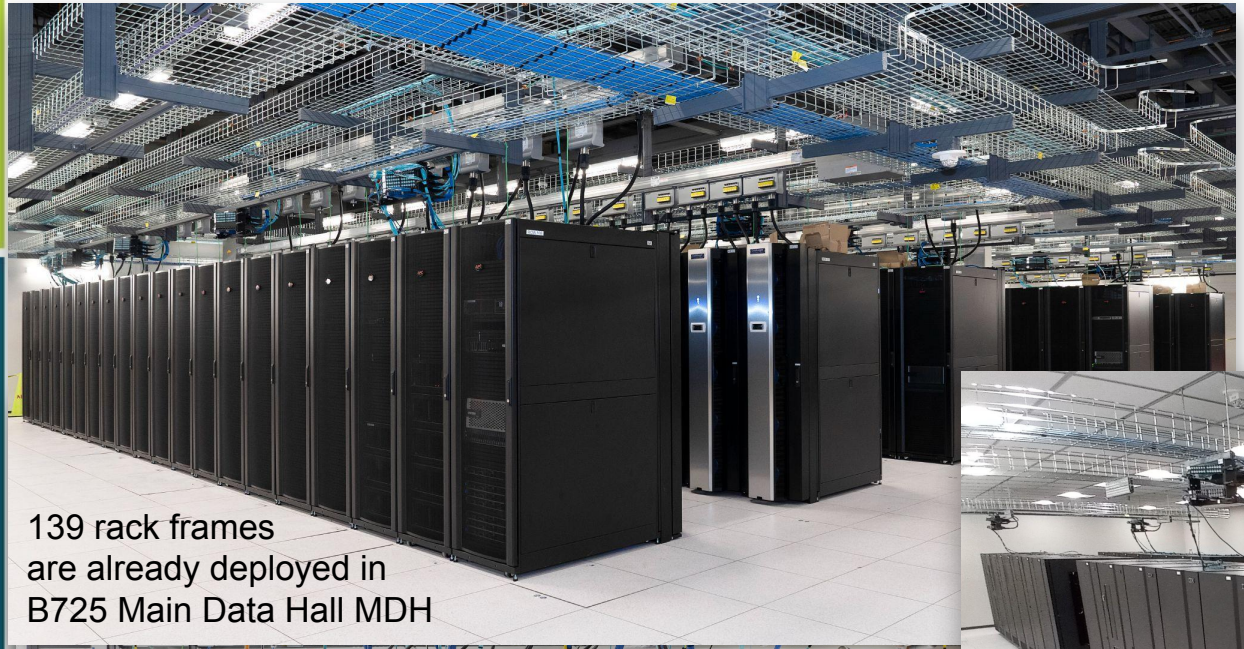

139 rack frames
are already deployed in
B725 Main Data Hall MDH

121 RDHx units are deployed in B725 MDH, out of which 104 are on racks with equipment while 17 are deployed for future growth

B725 Central Network Equipment Is Deployed & Active
(10x 400 GbE ready Arista modular chassis with 48x line cards slots in total)

4x 8-frame IBM TS4500
tape libraries are installed
in the B725 Tape Room

7

New Data Center Infrastructure

- SDCC received approval for power upgrade of the new data center (B725)
  - New (3rd) generator - increasing the total IT power that can be operated in B725 with N+1 redundancy in the generator group from 1.2 MW to 2.4 MW (ETA is summer 2025)
  - New (4th) power system - increasing the total amount of IT power available in B725 from 3.6 MW to 4.8 MW (ETA is summer 2025)
- Expansion needed to support sPHENIX computing and storage needs
  - Will directly benefit other programs (DUNE, ATLAS and EIC) later this decade

sPHENIX Commissioning

- Last major experiment for the RHIC program, scheduled to run from 2023 to 2025
  - Planned a 12-week commissioning run in Spring/Summer 2023 to validate the software and computing model before a short (few weeks) Au-Au run in the Fall
  - On August 1st, a magnet quenched due to cryo equipment failure that resulted in the loss of helium cooling at RHIC. Subsequent investigation discovered substantial damage. Because of the lengthy repair timeline, it was decided to end the 2023 RHIC run prematurely, and to plan for an extended 2024 to compensate for the loss of beam time.
- The cryo incident created a ripple effect on software and computing preparations at the SDCC in support of sPHENIX
  - While this bought more time to finish commissioning activities, it also shortened the timeline for procurement of computing/storage resources
  - Timeline challenges – supply chain delays, power/cooling capacity in new data center

Lustre for sPHENIX

- Upgrade from 2.12.X + RHEL7 to Lustre 2.15.2 + RHEL 8.7
  - 25 new OSS added, currently have: 1 MDS, 55 OSS (**53.3PB** capacity)
- Lustre OSS deadlock issue on Lustre 2.15.2 + RHEL 8.7 for MDRAID impacted stability and performance
  - Solutions (done live without bringing Lustre down):
    - Conversion from MDRAID to ZFS, node by node, and system-wide.
    - Migration of data from MDRAID OSTs to ZFS OSTs
- Robinhood policy engine
- Whamcloud Level 3 support
- Monitoring for Lustre 2.15.X + RHEL 8
  - Barreleye, Lustre Exporter, Node Exporter, LOKI syslog
- ML Anomaly detection using Lustre Jobstats
  - Based on Isolation Forest model developed by Amy Chen (HS Summer Intern)

S3 storage for the EIC

- Reliable, scalable, easy-access storage service in support of EIC detector simulation needs
- Ceph/RADOS solution replaces previous MinIO prototype
  - 4 storage hosts (48 Core, 96 GB RAM, 4x25 GbE, 8x16 TB OSD)
  - ~450 TB RAW, 4+2 Erasure Coding (300 TB usable)
  - 5.5 GB/s write performance inline with expectation and 2.6x previous setup
- Access via Federated ID (no BNL account required)

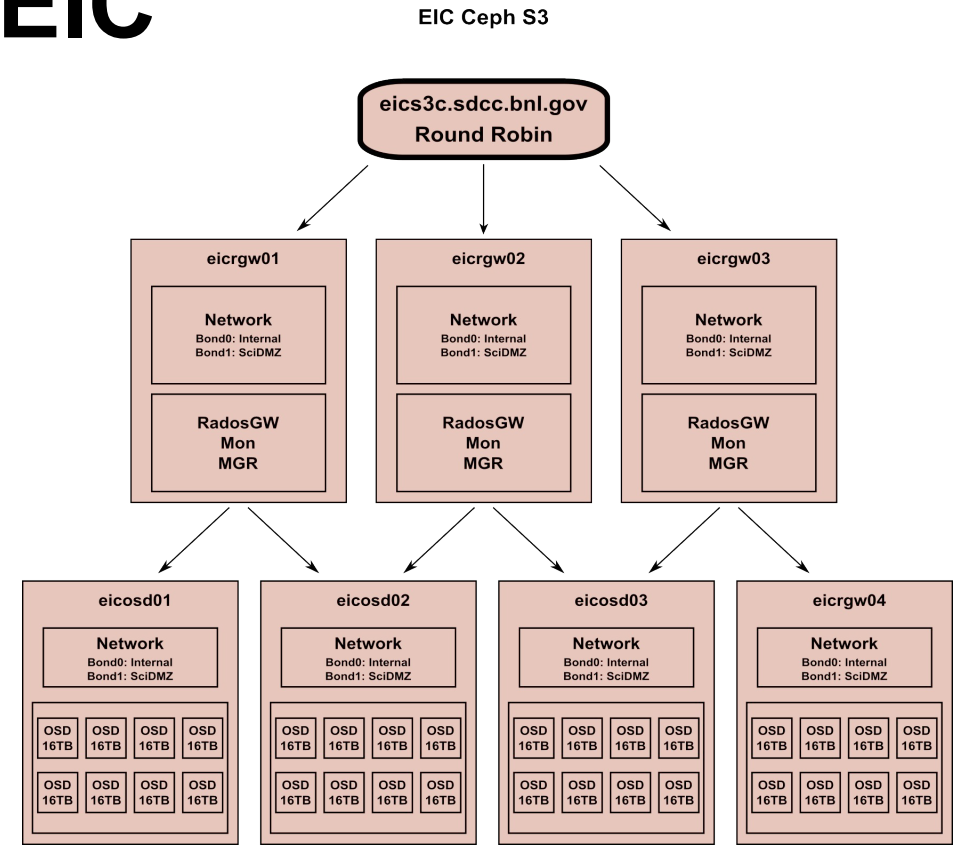

High Throughput Computing

- Providing our users with ~2,300 HTC nodes:
  - ~125,000 logical cores, ~1850 kHS06
  - Managed by HTCondor
- Purchased 120 Supermicro SYS-610C-TR nodes for Belle-II and the RHIC experiments (~12k logical cores total)
  - 6 racks, just delivered
  - System specs:
    - Dual Intel Ice Lake Xeon Gold 6336Y 24-core processors
    - 12x32 GB 3200 MHz ECC DDR4 RAM (384 GB total)
    - 4x2 TB SSD drives
    - 1U form factor
    - 10 Gbps NIC
- All nodes still running Scientific Linux (SL) 7
  - Preparations for an OS upgrade to Alma Linux 9 in progress
- HTCondor 10.0 fully tested, and a rolling upgrade mainly complete
- Purchased 2 Ampere Altra based systems from Supermicro
  - The servers arrived, and have been benchmarked
  - 2 Nvidia Grace-Grace based systems also purchased - expect arrival in the next few months

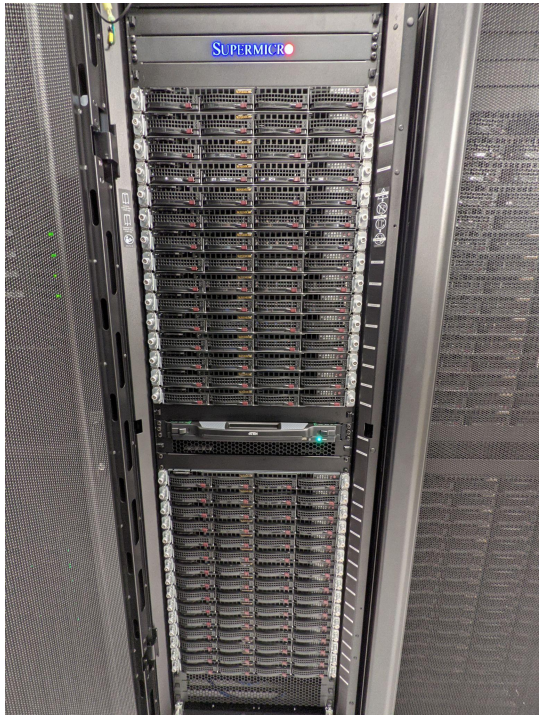

Supermicro SYS-6019U-TR4 Servers

ARM System Specs

ARM Ampere Altra Max	ARM Ampere Altra	x86_64 Xeon
Supermicro ARS-110M-NR	Supermicro ARS-110M-NR	Supermicro PIO-610C-TR-1-MT037
CPU - Ampere Altra Max 128 core 2.8GHz	CPU - Ampere Altra 64 core 2.6GHz	CPU - 2x Intel Xeon Gold 6336Y CPU 24 core 2.4GHz
Physical cores: 128 Threads (Job slots): 128	Physical cores: 64 Threads: 64	Physical cores: 48 Threads: 96
Memory – 512 GB (8x 64GB DDR4)	Memory – 256 GB (8x 32GB DDR4)	Memory – 384 GB (12x 32GB DDR4)
4GB per thread	4GB per thread	4GB per thread

ARM System Benchmarks

ARM 128 core	ARM 64 core	X86_64 Xeon
HEPscore – 2030	HEPscore – 1042	HEPscore – 1278
Max power – 370W	Max power – 252W	Max power – 601W
Idle – 125W	Idle – 107W	Idle – 179W

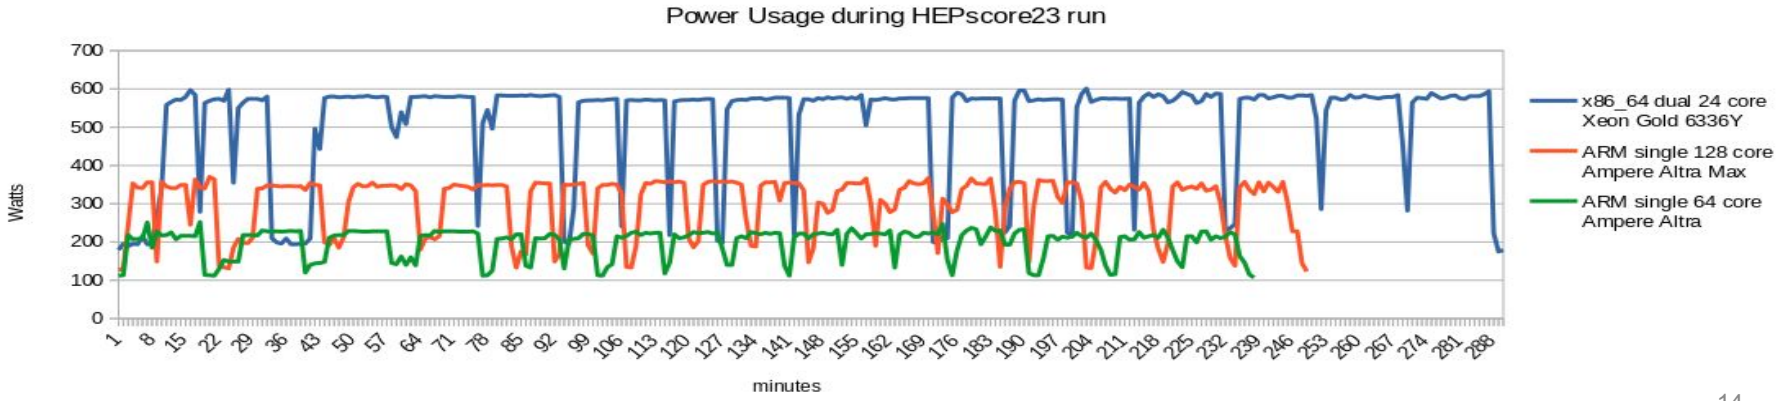

ARM Cost Analysis

Relative purchase cost in dollars per HS23:

- Ampere Altra was 89% cost of Xeon to purchase
- **Altra Max** was **61%** the cost of Xeon to purchase

Relative operating cost (electricity only):

Watts per KHS23:

Ampere Altra Max128 – **182.3 (39% of Xeon)**

Ampere Altra Q64 – 241.9 (51% of Xeon)

Xeon 6336Y – 470.3

High Performance Computing

Institutional Cluster (IC gen1)

Retired on October 2nd after ~7 years of service

Was 216 servers (mix of K80 and P100 GPU's) with IB EDR
144 node KNL cluster also retired

Institutional Cluster (IC) Gen2

System now fully online and in production

Specs:

- 2x Intel Xeon 6336Y 2.4GHz CPUs (Ice Lake)
- NDR200 IB (200 Gbps)
- 39 CPU-only nodes
  - 512GB DDR4-3200
- 12 GPU nodes
  - 4x Nvidia A100 80GB
  - 1TB DDR4-3200
- Planning to add another 20 GPU nodes this year, configured as above

New IC Gen2 Cluster

Other General Infrastructure Updates

- GPFS hardware refresh for IC Gen2: 2.4PB
- NSLS2 Lustre
  - New Lustre FS (3.5PB) + NetApp NFS (for small files)
  - Two of the existing Lustre systems upgraded to RHEL8 (14 servers)
- CVMFS upgrade
  - All old server hardware replaced by VMs + NetApp NFS storage for Stratum-1 cache: 44TB
- Mattermost upgrade: v7.8 + RHEL8
  - Added CILogon/COmanage login option for federated logins
    - Plan to migrate all users to CILogon/COmanage logins eventually

Questions?

Linux shipment entering Puget Sound

Thanks to the following people at BNL for contributing to this presentation:
Costin Caramarcu, Tim Chou, Vincent Garonne, Chris Hollowell, Jerome Lauret, Jason Smith, Tom Smith, Tony Wong, Alex Zaytsev

Extra Slides

Storage: Lustre, dCache & XROOTD

dCache.org

Total ~74 PB

- ATLAS (v8.2.*), Belle II (v8.2.*), DUNE (v8.2.*), PHENIX (v5.2.9), Pre-production/integration (v8.2.*)
- ATLAS:
  - Major refurbishment cycle involving pool/door/core replacements and migration to the new data center, necessitating an internal move of over 50 PB of data
  - Cell components' deployment model has been thoroughly reviewed and redefined to minimize single points of failure
- Pre-production services utilized by ATLAS, BELLE, and EPiC before any changes to production (e.g. major upgrade, tape interface changes, integration, etc)

XROOTD

~11 PB total storage for STAR

- Mix of central and farm node storage

lustre®

Total ~50 PB in Lustre

- ATLAS, EIC, LQCD, NSLS-II, sPHENIX, STAR
- Expanding the Lustre footprint with extensive operational experience, including major upgrades (2.15.*), storage addition (+ 25PB to sPHENIX) and finely-grained monitoring
- Excellent streaming sequential performance with aggregate throughput of 210 GB/s

Tape System

- Currently ~250 PB of data in HPSS with ~76k tapes
- 9 Oracle SL8500 and 5 IBM TS4500 tape libraries
- Total of 120K tape slots, 280 LTO tape drives

- New HPSS hardware in the newly commissioned data center for sPhenix experiment
  - 10 GB/sec data injection requirement
  - High performance/capacity disk cache (2.1PB, 330 HDDs)
  - Two new IBM TS4500 tape libraries
  - Total of 64 new LTO-9 tape drives in two TS4500 libraries

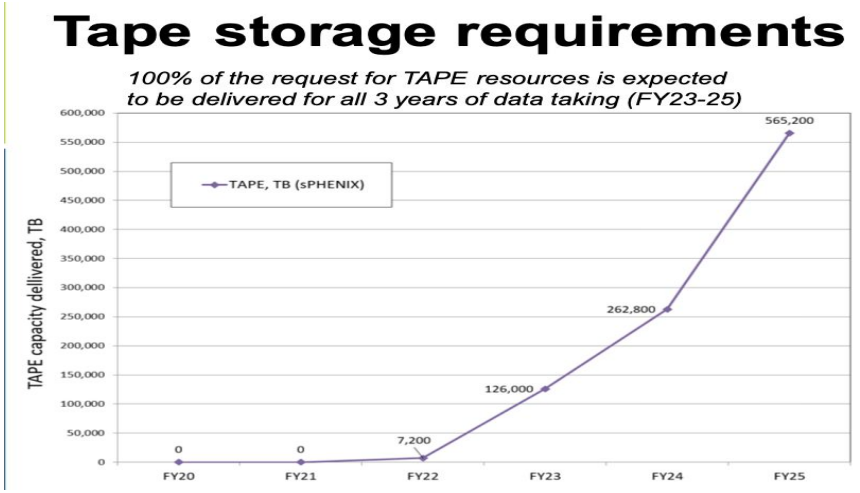

sPHENIX Tape Storage Space Requirements

Application Platforms

- **Two OKD clusters (sPHENIX and ATLAS), 7 nodes each, in production**

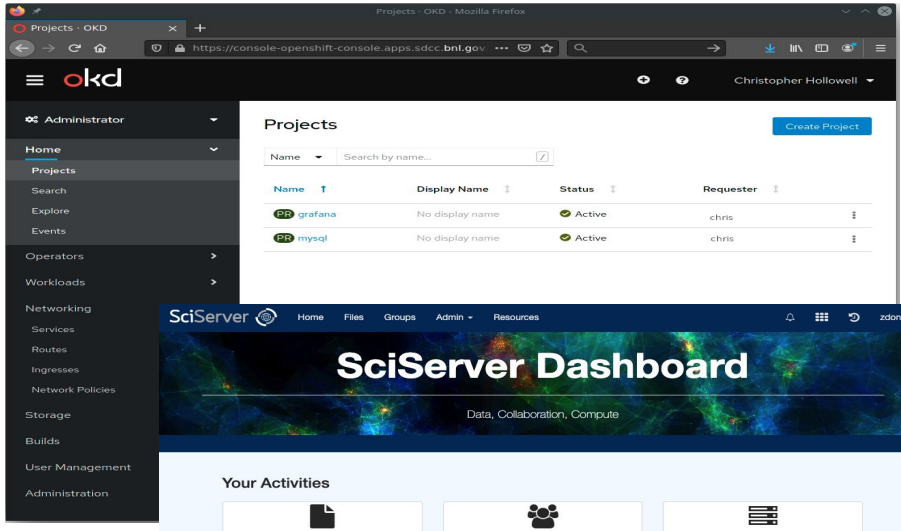

SciServer

The Science Platform

A collaborative environment for server-side analysis with large datasets

- Test instance: available for test users
- For production: 4 Compute nodes powered on, each with 8xH100 SXM GPUs