IT Fabric

Chris Hollowell <hollowec@bnl.gov>
Scientific Data and Computing Center (SDCC)

ATLAS Pre-Scrubbing Review - June 27, 2022

Brookhaven National Laboratory

U.S. DEPARTMENT OF ENERGY

@BrookhavenLab

IT Fabric - Overview

- 11 staff members
  - 8 of whom are partially funded by ATLAS
    - T. Chou (0.3), R. Hancock (0.9), C. Hollowell (0.4), O. Novakov (0.3), J. Spradley (0.35), T. Smith(0.2), I. Wu (0.3), X. Zhao (until FY22)
    - Full FTE expertise/cost shared by multiple programs
- *Responsible for:*
  - High Throughput Compute (HTC) Farms
    - HTCondor batch system
  - High Performance Compute (HPC) Clusters
    - Slurm batch system
  - HPSS mass tape storage system
    - BNL-developed ERADAT tape batch system
  - Distributed (dCache/XROOTD) storage hardware
    - ATLAS dCache software/service itself managed by Tools & Services group

IT Fabric

Chris Hollowell - Group Lead

HTC
Chris Hollowell*
Tom Smith
Xin Zhao

HPC
Costin Caramarcu*
Kevin Casella
Zhihua Dong

Central Storage
Tim Chou*
Robert Hancock
Ognian Novakov
Justin Spradley
Iris Wu

* Technical lead

High Throughput Computing

- *System administrators: C. Hollowell, T. Smith, M. Poat*

- Providing our users with ~1,900 HTC nodes
  - ~90,000 logical cores
    - About ~25k lcores of these for ATLAS T1
    - About ~2,400 cores for ATLAS analysis facility
  - ~1,050 kHS06
    - ~315 kHS06 of this for ATLAS T1
    - ~25 kHS06 for the ATLAS analysis facility

- All nodes running Scientific Linux (SL) 7
  - Have been evaluating RHEL8 and RHEL8-based distributions including Rocky Linux 8 due to SL discontinuation and CentOS 8 early EOL

- HTCondor batch system v9 (including token support) used to schedule jobs

- Dedicated subset of per-experiment interactive VMs for login and job submission. Also provide interactive access via Jupyter
  - Utilized by the shared ATLAS analysis facility at BNL

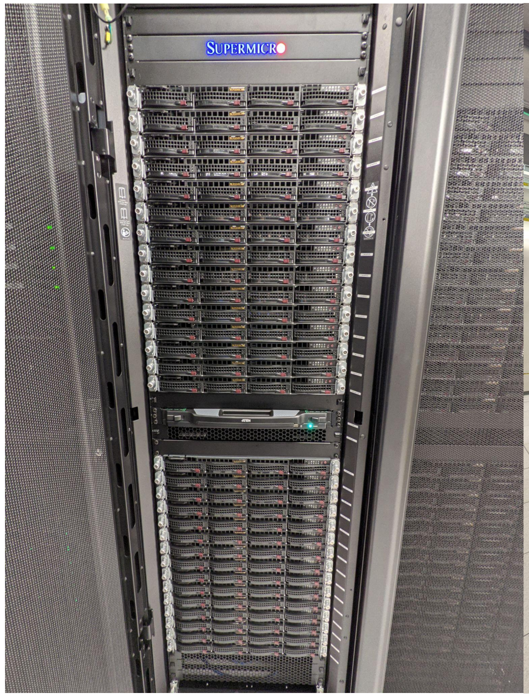

Supermicro SYS-6019U-TR4 Servers

HTC Racks with RDHx Cooling in our new B725 data center

Brookhaven National Laboratory

High Performance Computing

- *System administrators: C. Caramarcu, Z. Dong, K. Casella (all 3 not on ATLAS WBS)*

- **Currently supporting 5 HPC clusters**
  - Institutional Cluster (IC)
    - 216 HP XL190r Gen9 nodes with EDR IB
    - Each with 2xK80 or 2xP100 GPUs
    - ***ATLAS has an allocation here of 6 P100 nodes***
    - ***PanDA queue - ANALY_BNL_GPU_ARC***
- Skylake Cluster
    - 64 Dell PowerEdge R640 nodes with EDR IB
  - KNL Cluster
    - 142 KOI S7200AP nodes with dual rail Omnipath interconnect
    - Each with 2x Xeon Phi 7230 AVX-512 CPUs
  - ML Cluster
    - 5 HP XL270d Gen10 nodes with EDR IB
    - Each with 8x V100 GPUs
- NSLS2 Cluster
    - 30 Supermicro nodes with EDR IB
    - 13 of the nodes with 2xV100 GPUs

Institutional Cluster (IC)

- ***Requisition on for replacement next generation Institutional Cluster in process***
  - Mix of CPU-only (Xeon Ice Lake) nodes, and A100-80 HGX SXM4 based nodes

Distributed Storage Hardware

- *System administrators: R. Hancock, K. Casella (<u>not</u> on ATLAS WBS), J. Spradley*

- ~78 PB (on ~10,000 drives) total JBOD storage for dCache/XROOTD
  - ~59 PB raw for ATLAS (on ~7,000 drives)

- 92 JBOD/headnode pairs
  - 76 pairs for ATLAS

- Using JBODs for storage as considerably more performant, and less expensive that hardware RAID
  - Using a mix of MD and ZFS software RAID on the JBOD storage
  - MD/ZFS software RAID more complex to manage than hardware RAID, however
  - In-house developed custom management/monitoring software (R. Hancock) greatly assists us here

Western Digital Ultrastar Data 102-bay JBODs with Dell R740 headnodes

Brookhaven National Laboratory

HPSS

- *System administrators: T. Chou, O. Novakov, I. Wu*
- Multi-Hundred-PB scale HPSS mass tape storage system
  - ~215 PB accumulated data total for all experiments on ~75,000 tapes
  - ~70 PB for ATLAS on ~20,000 cartridges
- Currently running HPSS 8.3.10
- Tape libraries:
  - 9 Oracle SL8500 libraries in old B515 data center (CDCE)
    - 2 of these for ATLAS
  - 2 IBM TS4500 tape library pairs already in production
    - 1 pair for ATLAS
- In-house-developed ERADAT tape batch system used to optimize tape read access
  - Consolidates reads to minimize tape mounts

ATLAS IBM TS400 Library Assembly

Brookhaven National Laboratory

New ATLAS Compute Hardware

- Recently added 23 Supermicro SYS-6019U-TR4 1U servers to the T1 HTC farm
  - 2208 logical cores, 26.2 kHS06 total
  - Per system CPU specs:
    - 2x Intel Xeon Cascade Lake 6252 CPUs
    - 192 GB DDR4-2933 MHz RAM
    - 4 x 4 TB SSDs
    - 10 Gbps NIC
  - *FY22 ATLAS farm purchase in our requisition system for an additional 149 kHS06*
    - Delayed due to continuing resolution
    - Purchase combined with large sPHENIX order to maximize volume-discount pricing
- 2 new Supermicro SYS-120GQ-TNRT GPU compute nodes added to shared analysis facility
  - Accessible via HTCondor
    - Users added `request_gpus` line to their JDFs
  - Per system specs:
    - 2 x Intel Xeon Ice Lake 6336Y CPUs
    - 2 x A100-80 GPUs
    - 256 GB DDR4-3200 MHz RAM
    - 2 x 4 TB NVME drives
    - 10 Gbps NIC

New ATLAS Supermicro SYS-120GQ-TNRT GPU Nodes

- All existing 2017 and newer ATLAS T1 compute nodes moved to the new B725 data center with assistance of the Data Center Operations Team

New ATLAS dCache Hardware

- **New JBOD storage systems have been delivered**
  - In the process of being assembled

- **JBOD/headnode specs**
  - Western Digital Ultrastar Data102 JBOD
    - 4U chassis
    - 12 Gbps SAS connectivity
    - 102 14 TB SAS Drives
    - New systems all configured with ZFS
  - Dell R740 head node
    - 2x 6254 Xeon Cascade Lake CPUs
    - 384 TB DDR4-2933 MHz RAM
    - 2x 25 Gbps SFP NICs
    - 2x 12 Gbps SAS HBAs
    - 4x 1 TB SSDs for OS

1.7 Performance Specifications

Table 4: Performance Specifications

Specification	Value
Number of Drive Slots	102
Data Transfer Rates	12Gbps SAS / 6Gbps SATA
Max Raw Data Storage Capacity	2.0PB using 20TB Ultrastar HC650 drives
SAS Ports	12 x Mini-SAS HD (6 per IOM)
Management Ports	2 x 10/100/1G Ethernet

1.8 Ultrastar Data102 Layout

Figure 4: Front and Rear Product Layout

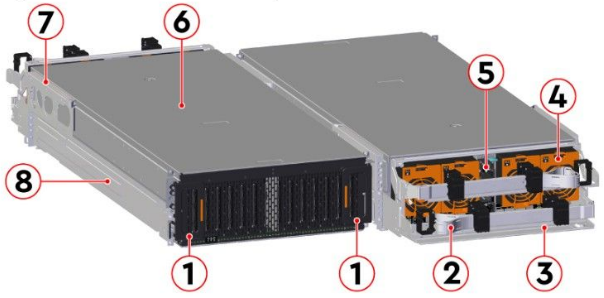

Ultrastar Data 102-bay JBOD Specs

New ATLAS Tape Library

- IBM TS4500 library pair installed in new B725 data center
  - 64 LT08 tape drives
  - 17,612 slots total
  - LTO8 capacity 12 TB = 211 PB total LTO8 potential future capacity added
  - Custom SCSI interface tape library management software developed by BNL staff (T. Chou)

- HPSS upgraded from 7.4.3.2 to 8.3.10 in August 2021
  - Done to support LTO8, and to enable migration by directory features
    - Important HPSS improvements required by ATLAS data carousel activity
  - First major upgrade in 5 years
  - Complex process, but due to careful planning and testing was completed successfully within the scheduled 4-day downtime window

ATLAS IBM TS4500 Library Pair

HEPscore Activities

- The HEPiX Benchmarking Working Group is developing the HEP Benchmarks Project, often referred to as "HEPscore", to replace HEPSPEC06
  - Based on HEP and other experiment workload benchmark containers
  - Includes multiple ATLAS workload containers (simulation, reconstruction, etc.)

- Identified HESPEC06 issues and improvements provided by HEPscore
  - HEPSPEC06 is getting old (currently over 12 years old)
    - Based on SPEC CPU2006, which is no longer supported by SPEC corporation
      - SPEC CPU2006 is no longer available directly for purchase on SPEC website
      - Potential issues compiling/running on RHEL8+
    - Moving to SPEC CPU2019 is an option, but it is still commercially licensed/paid software
      - HEPscore and its workloads are freely available and have open source (GPLv3 and other) licenses
  - HEPSPEC06 is a synthetic benchmark for HEP
    - Not based on real-world HEP applications
    - There have been concerns it may not accurately model real-world HEP/ATLAS application performance on new CPUs
      - Issues have already been seen in the HS06 performance modeling on Haswell and newer CPUs (so-called "Haswell Boost")
      - HEPscore developed to use real experiment workloads for benchmarking
        - Allows for more accurate hardware performance modeling
        - Can help sites to optimize their hardware purchases for real ATLAS software performance

Brookhaven National Laboratory

HEPscore Activities (Cont.)

- SDCC has been actively involved in several HEPscore related activities
  - SDCC staff member (C. Hollowell) is a participant in the HEPiX Benchmarking Working Group and an author/maintainer of the `hep-score` benchmark orchestration/execution utility which runs the benchmark workloads, and calculates an overall result/score
  - Released hep-score v1.4 in April. Made a number of changes this year:
    - Renamed default benchmark to HEPscoreBeta
      - Clarifies the benchmark set/configuration is in a development state and being reviewed by the WLCG HEPscore Task Force
    - Documented Podman Usage
      - Podman not yet officially supported, but various tests have shown it is usable with some modifications to OS settings
    - No longer contain PID namespace
      - Resolves issues running in k8s
    - Add support for oras:// and https:// registry URIs
      - Allows the use of SIF files in OCI container registries and webservers
  - Plan to test and add official support for Apptainer

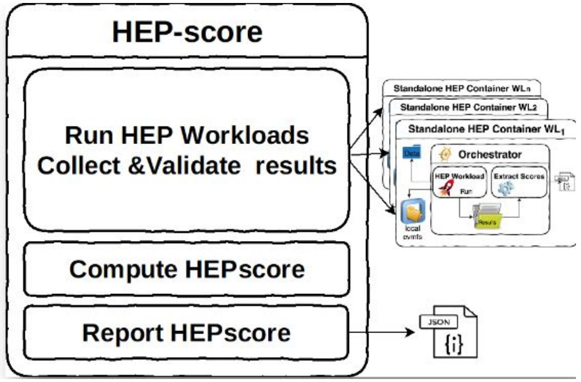

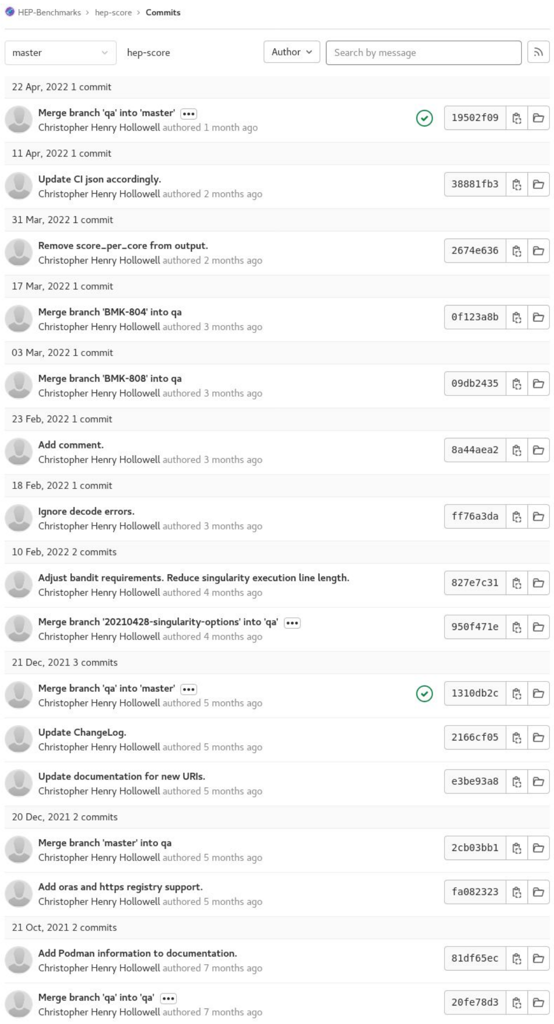

Brookhaven National Laboratory

HEPscore Activities (Cont.)

- SDCC staff member (A. Wong) a member of the WLCG HEPscore Task Force
  - Investigating the potential for HEPscore to replace HEPSPEC06
  - Will determine the final workload container set/configuration/parameters
    - Goal to have final benchmark set/configuration by September 2022
  - SDCC/BNL is one of a number of sites actively participating in running various benchmark sets and reporting results to CERN’s ELK instance for task force analysis and reproducibility studies
    - Investigating sqlite-related crashes in the `atlas-sim-mt` benchmark container
      - Due to other ATLAS-related activities, have not had time to dig deeper into the issue, but plan to do so soon

Progress Table from WLCG HEPscore Task Force June 1st Meeting

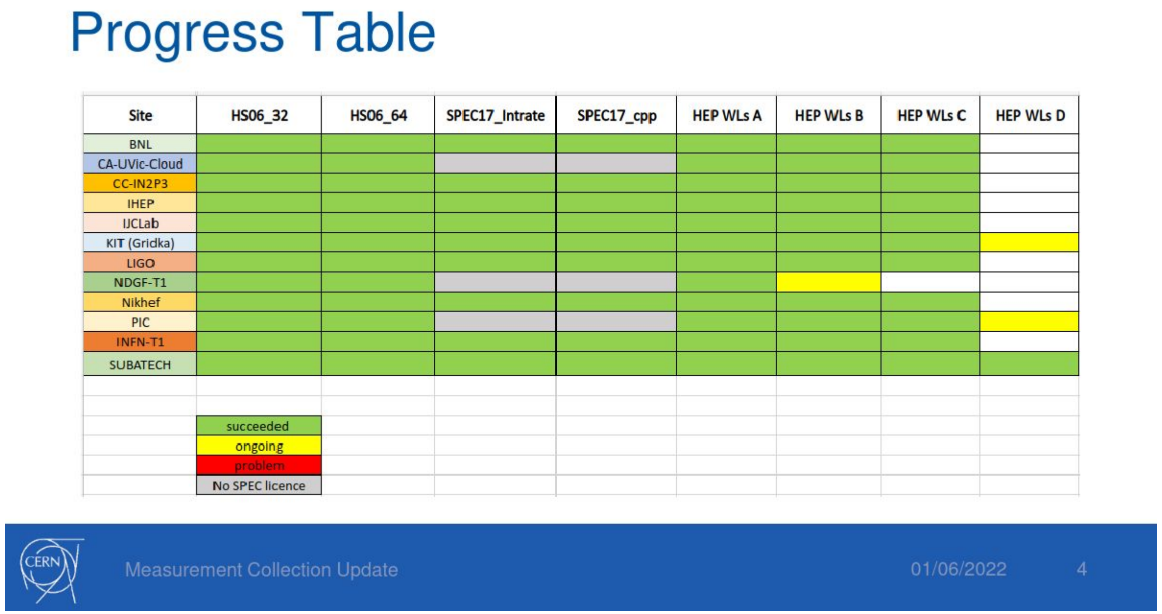

Datacenter Cost Model: Cloud Comparison

- SDCC staff (A. Zaytsev, C. Hollowell, S. Misawa) developed datacenter and cloud computing cost models to allow comparison with the cost of on-premises equipment purchase/operation
  - Used in determining if the next generation SDCC Institutional HPC Cluster (IC Gen2) and NSLS2 cluster should be purchased, or if cloud resources should be utilized instead
    - While primarily a comparison in the HPC context, HTC scenarios were also considered

- Developing a cost model for on-premises, and cloud computing IT resources is a complex multi-variable problem

- ***The conclusion was that for a 100% utilized system, the cost of locating IC Gen2 in the cloud would be 7-9 times more expensive over a 5 year period, even with considerably reduced DOE Google pricing***
  - **HTC compute was similarly ~10x the cost for 100% utilized resources**
  - **The cloud is still an attractive option for providing burstable resources to meet peak demand**

- It is important to note that cloud computing options do not significantly reduce staffing costs: high-level system/devops administrators/engineers are still required to run the system

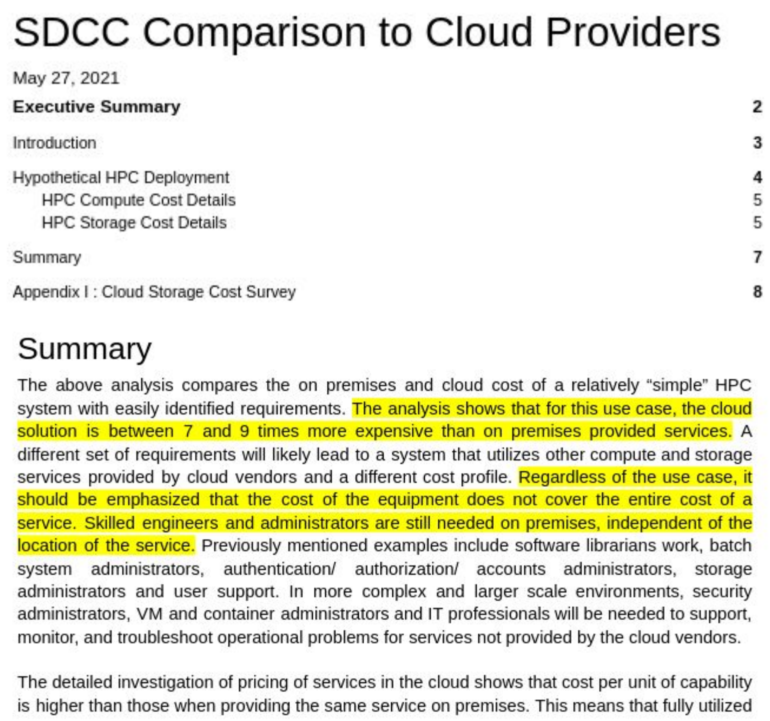

SDCC Comparison to Cloud Providers

May 27, 2021

Summary

The above analysis compares the on premises and cloud cost of a relatively "simple" HPC system with easily identified requirements. The analysis shows that for this use case, the cloud solution is between 7 and 9 times more expensive than on premises provided services. A different set of requirements will likely lead to a system that utilizes other compute and storage services provided by cloud vendors and a different cost profile. Regardless of the use case, it should be emphasized that the cost of the equipment does not cover the entire cost of a service. Skilled engineers and administrators are still needed on premises, independent of the location of the service. Previously mentioned examples include software librarians work, batch system administrators, authentication/ authorization/ accounts administrators, storage administrators and user support. In more complex and larger scale environments, security administrators, VM and container administrators and IT professionals will be needed to support, monitor, and troubleshoot operational problems for services not provided by the cloud vendors.

The detailed investigation of pricing of services in the cloud shows that cost per unit of capability is higher than those when providing the same service on premises. This means that fully utilized

Brookhaven National Laboratory

CE Token Transition

- OSG dropped support for X509 in CEs in the OSG 3.6 release

- Initial plan to EOL OSG 3.5 (which supported both X509 and tokens) in March 2022
  - Later delayed until May 1st 2022 to give sites and VOs more time to handle the transition

- T1 was one of the first US sites to successfully test token job submission in fall 2021

- *All Eight ATLAS T1 CEs upgraded to OSG 3.6 before the May 1st deadline*

Federated JupyterHub

- Technical implementation of a federated Jupyter instance for ATLAS complete
    - Allows users to login via BNL/SDCC, CERN, FNAL or SLAC credentials
        - Users without existing SDCC accounts are given “lightweight accounts” where their only access at the facility is Jupyter (i.e. cannot login to SSH gateways, etc.)
    - Implementation requires MFA to satisfy DOE requirements
    - Users must fill out a form and be approved for initial access after their ATLAS affiliation is confirmed
    - Account activation turn-around time goal of 1-2 business days

- EPPN OIDC token attribute mapping used in LDAP for authorization and to tie federated ID to a local UNIX account
    - Reverse proxies doing the authentication
    - Modifications to the Jupyter jhub_remote_user_plugin to support the mapping
    - Users have access to all SDCC ATLAS network filesystems

- Being tested by a number of users

Federated JupyterHub Login Screen

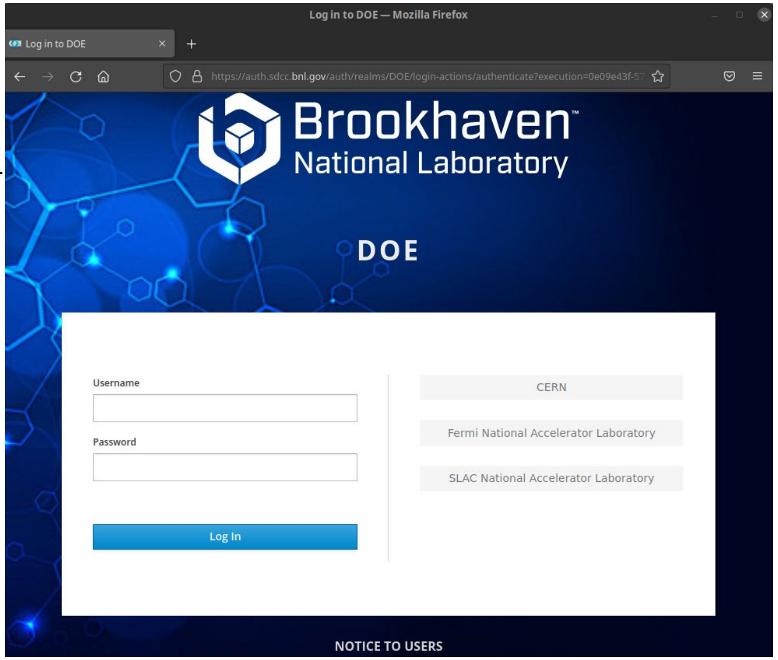

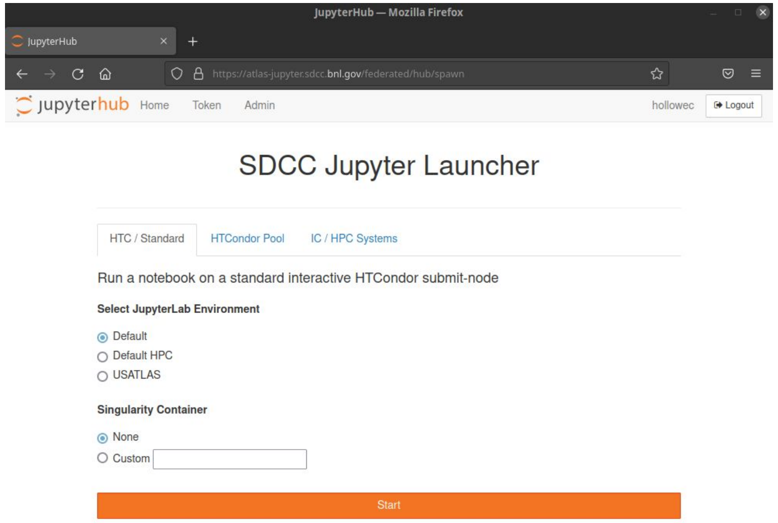

JupyterHub Web Interface

Brookhaven National Laboratory

ATLAS Analysis Facility

- Have test ServiceX and REANA instances setup in our staff kubernetes cluster, and our production ATLAS OKD cluster
  - ServiceX - Columnar data delivery service
  - REANA - framework for reusable analysis
    - Both may become important parts of future ATLAS/HEP analysis facilities

- JupyterHub service and eight interactive VMs provided for ATLAS user analysis

- SDCC shared pool utilized by ATLAS users for batch analysis
  - Users occasionally making opportunistic use of more than 10k cores

ServiceX Web interface

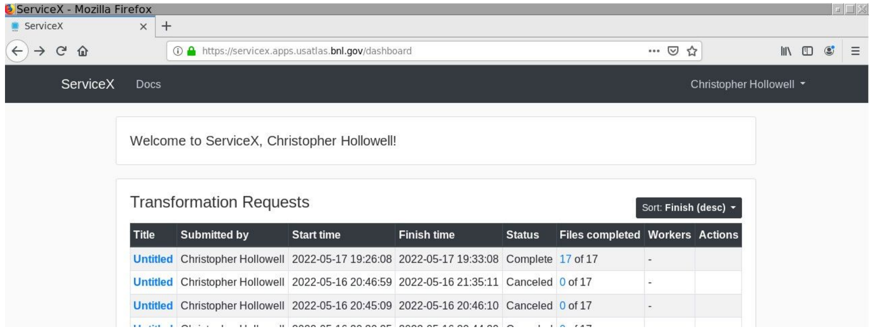

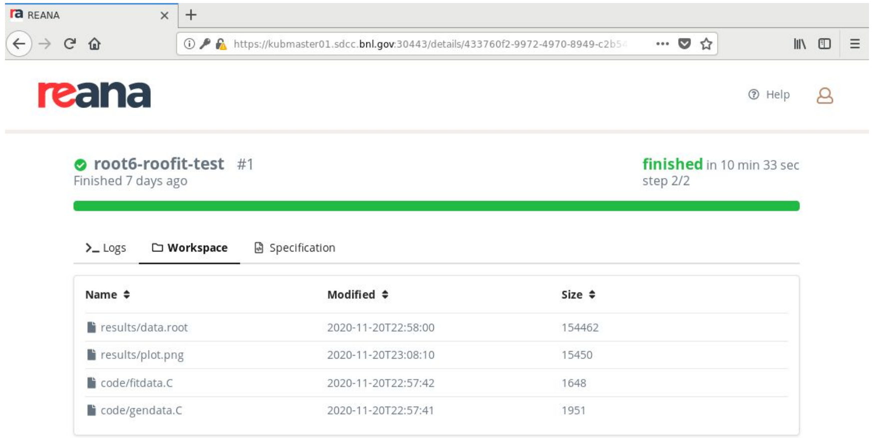

REANA Web interface

SDCC Shared HTCondor Pool Usage by the ATLAS Analysis Facility

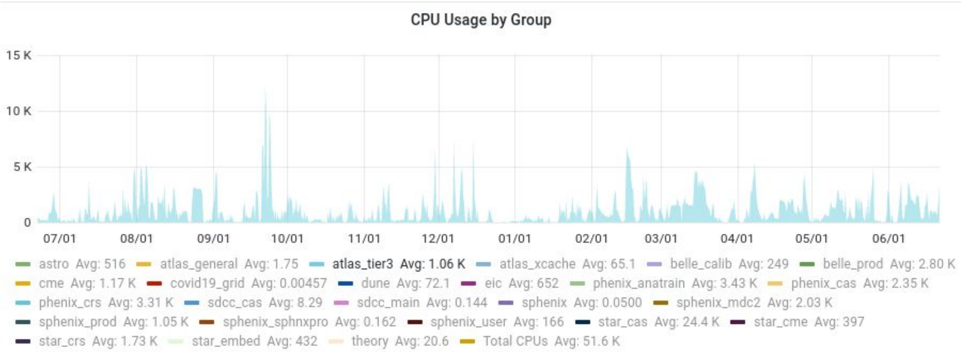

Brookhaven National Laboratory

IT Fabric Presentations and Publications

- Presentations
  - Tape Challenges - ATLAS SW&C Week, 6/17/2022 (X. Zhao)
  - Kubernetes/OKD at BNL - pre-GDB on Kubernetes, 6/7/2022 (C. Hollowell)
  - Tape Commissioning for Run-3 - ATLAS SW&C Week, 2/1/2022 (X. Zhao)
  - Migration of Archival Tape/Library Technology at SDCC - HEPiX Fall 2021 (I. Wu)
  - Linux OS Plans at The SDCC - HEPiX Fall 2021, 10/25/2021 (C. Hollowell)
  - Data Challenges - ATLAS SW&C week, 10/4/2021 (X. Zhao)
  - Worker Nodes at BNL - pre-GDB on Worker Nodes, 7/13/2021 (C. Hollowell)

- Publications
  - Giordano, D., Alef, M., Atzori, L., Barbet, J., Datskova, O., Girone, M., **Hollowell, C.**, *et al.* HEPiX Benchmarking Solution for WLCG Computing Resources. *Comput Softw Big Sci* **5,** 28 (2021). https://doi.org/10.1007/s41781-021-00074-y

ATLAS Challenges

- Increasingly difficult to compete with industry pay scales for system engineers/developers
  - Hiring process is long and difficult
  - Retention may become difficult as well

- COVID-19 related supply-chain issues have made it difficult to procure new equipment in a timely manner
  - Market volatility has also made it difficult to estimate/predict costs for planning purposes

- HL-LHC requirements
  - Vast amounts of computing/storage needed for HL-LHC
    - Multiple scenarios outlined
  - Would be good to refine the requirements
    - Choice of scenario to simplify planning process
    - Details such as local node IOPS performance, and desired memory and network bandwidth requirements, etc.

Conclusions

- SDCC's Fabric group consists of 11 staff members, with 8 partially funded by ATLAS

- Provide disk/tape storage and HTC/HPC compute hardware/sevices for ATLAS
  - Including management of batch systems, and HPSS tape system software
  - *Significant new hardware added for ATLAS in the past year*

- Token transition for our ATLAS CEs successfully implemented

- Active participant in HEPscore development and testing activities

- Developed (with A. Zaytsev and S. Misawa) cloud computing and local datacenter cost models/analysis

- Brought new k8s-based services online for testing at our ATLAS analysis facility
  - REANA
  - ServiceX