



Scientific Computing Division: Year-end Accomplishments

James Amundson

SCD All-Hands

December 15, 2021

Preliminaries

- **Welcome!**
- The division has continued to thrive in the remote environment
 - We are starting to get hints of a framework for a return to work
 - No schedule
 - Planning for a hybrid future
- The search for a new director is on...

Input on the Fermilab Director Search

From Fermilab members of the Director Search Committee (Maurice Ball, Brad Benson, Marcela Carena, Anna Grassellino, Jen Raaf):

On **January 5, 2022, from 10-10:30am CT**, we have arranged the first mini-townhall, to which we would like to invite a few representatives from each of the following organizations: CAO, CRO, CIO, CTO, NQIC, LBNF, PIP-II. In addition to ourselves, three other members of the search committee will be in attendance: Walter Massey (Committee Chair), Juan de Pablo, and John Mester.

Because the meeting will be very short in duration, we would like to focus the input to the following three questions:

- **What are the main characteristics that you are looking for in the next director?**
- **What should the priorities be for the next director to strengthen the lab (e.g., hires, infrastructure, etc.):**
 - **In the short-term (next 2-3 years)?**
 - **In the long-term (next 5-10 years)?**

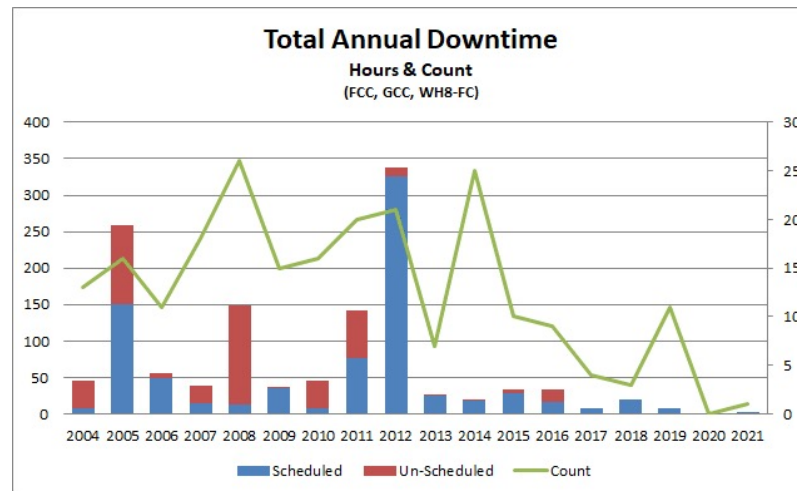
Input is also welcome on topics outside of these questions, but due to the time constraints of the meeting, these other comments should be sent via email to the Fermilab members of the search committee.

Please send your comments to Carla Lloyd, carlal@fnal.gov

Data Center Operations

Data Center Availability

- No major data center outages in 2020 or 2021
 - Requires continual preventative maintenance and mitigations to potential outages
 - Working closely with FESS, Procurement and subcontractors
- Much effort goes into planning mitigations for scheduled outages
 - Annual FESS ComEd Power Outages
 - Rental Generators at GCC
 - Annual FESS CUB Chilled Water Outage
 - Annual FESS Feeder Maintenance
- Rapid Response for Several Incidents
 - Water Leaks
 - Mechanical Failures (Cooling Equipment)
 - Power Disruptions
 - Thirty-Six After-hours Events in 2021
 - Nine Events Required Onsite Intervention



Data Center Improvements & Milestones

- Coordinating Upgrades of End-of-life IT Cabinets
 - Modern, Deeper Racks (typically 48” deep to accommodate new servers)
- Transitioning from 120v Electrical Distribution to 208v
 - Upgrading Toward the 208v Standard For All Racks (>90% complete)
 - Improves Energy Efficiency & Simplifies Rack Electrical Utilization for Servers
- Provisioned Space, Electrical Distribution, and Receiving Logistics for CMS Tape Library
 - CMS Tape Library Successfully Installed On FCC2 in May
- Portable Generators Deployed at GCC in September
 - Mitigated Two ComEd Site Power Outages
 - Also, Annual FESS Maintenance of High Voltage Substation
- Annual FESS CUB Maintenance
 - Impacted Primary Cooling for WH8-Fiber Central
 - Necessitated Deployment of Portable ‘backup’ Cooling Equipment
 - Saturday and Sunday (August 14-15)



Centralized Facility Management (CFM)


- CFM is the Consolidation of Conventional Facility Maintenance under FESS
 - Building envelope, roofs, windows, parking lots, office heating/cooling, lights, etc.
 - Computing onboarded in May for FCC, GCC, 18 Neuqua/Archives House
 - Wilson Hall Continues Largely Unchanged (John Kent, Sue Sprosty)
- Spoiler Alert! FESS was already performing most ‘building’ maintenance for CD
 - Conveniently, the CFM Transition Occurred While Most of Us Were Teleworking 😊
 - Several small subcontractor contracts were transferred to FESS
- Bruce Bieritz is the FESS Area Facility Manager (AFM) for Computing Sector
 - Bruce is Very Knowledgeable and Previously Worked in Technical Division
- Mark Thomas and the entire Data Center Operations team remain the ‘eyes and ears’ for FESS in the Tenant Liaison role. Many maintenance issues are noted during daily walkthroughs of the buildings.

Property Name	Landlord	Site	Zone	Area Facility Manager	Tenant Liaison/Building Manager
FEYNMAN COMPUTING CENTER	CD	CCA	AFM07	32156N - Bieritz, Bruce Phone: 6467 bbieritz@fnal.gov	MARK THOMAS
18 NEUQUA - HISTORICAL ARCHIVES	CD	VHA	AFM07	32156N - Bieritz, Bruce Phone: 6467 bbieritz@fnal.gov	MARK THOMAS
HDCF - GRID COMP CTR - WBL CTG HOUSE	CD	PAN	AFM07	32156N - Bieritz, Bruce Phone: 6467 bbieritz@fnal.gov	MARK THOMAS

Scientific Compute Services

Scientific Compute Services

- CCE
 - Darshan IO profiling on LArSoft to look for bottlenecks
 - HDF investigation for possible use as standard data format
- Continuous Integration
 - Jenkins/GitHub integration to trigger builds on PR and report status as comment to the PR and slack
 - improved support to monitor Physics Validation results (custom thresholds for individual metrics monitored by the experiment)
- CMS Production
 - Integrated ALCF Theta and generated HL-LHC MC samples
- Condor
 - Moving to tokens and retiring GSI
- FIFE
 - Jobs are running in containers by default
 - Retired SL6 for batch processing



Scientific Compute Services

• Elastic Analysis Facility

- We built a multi-VO, secure, integrated, EAF prototype in compliance with DOE cybersecurity requirements
- Started as a USCMS project but have grown to be an SCD-wide project
- Developed more than 15 environments for experiments with dedicated CVMFS mounts, shared storage and specific scientific software
- In Beta, give it a try <https://analytics-hub.fnal.gov>

• FermiCloud

- Moved Openstack to production in March, retired OpenNebula
- Working on IPv6 for developers (IPv6 only and dual stack)
- Working on a hardware refresh retiring old hardware

• Tokens/FedId

- Setup Token task force with experiments to explain/learn
- Storage layouts with tokens biggest issue

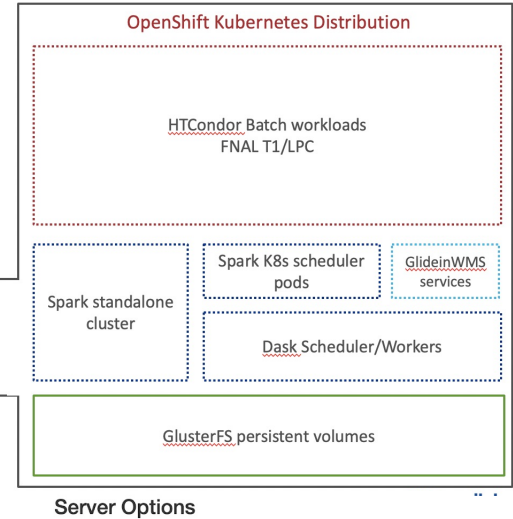
User interaction

HTCondor submission infrastructure (x509)

Direct submission (spark-submit, python scripts, etc.) (x509)

JupyterHub spawner Custom notebooks (FNAL LDAP)

* Dotted boxes can scale up/down dynamically depending on user demand



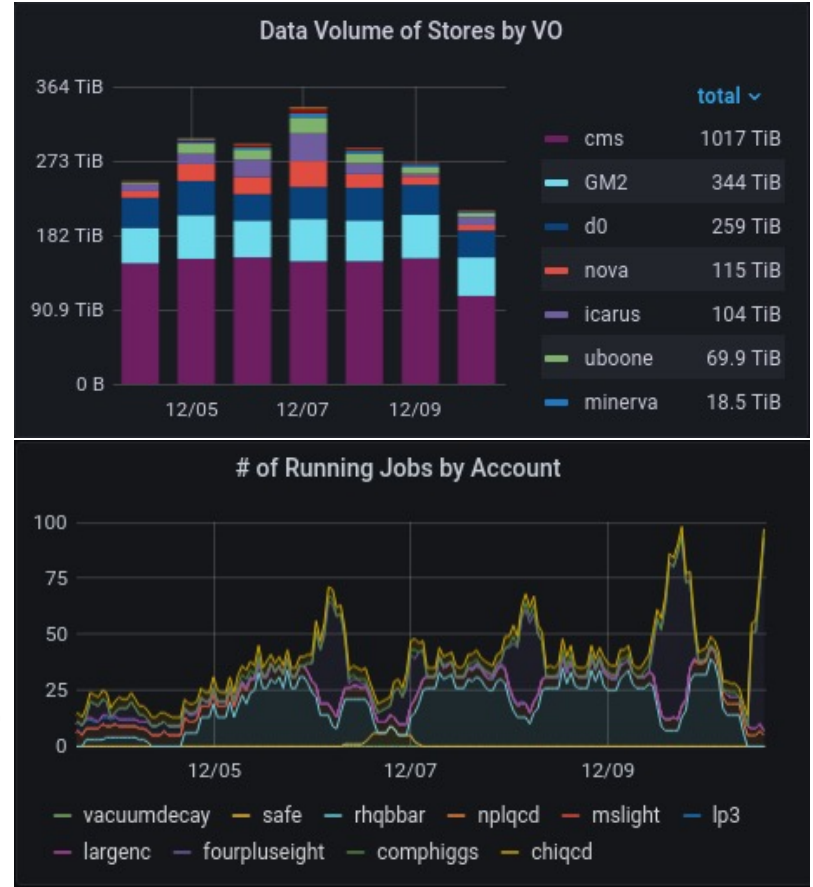
Server Options

<p>ACCEL-AI READS</p> <ul style="list-style-type: none"> ● AD-READS CPU Only SL7 Interactive ● [GPU] AD-READS NVIDIA Tesla K40m SL7 Tensor Interactive ● [GPU] AD-READS NVIDIA Tesla T4 SL7 Tensor Interactive 	<p>DES/LSST/ASTRO</p> <ul style="list-style-type: none"> ● DES/LSST - SL7 Interactive General Purpose Notebook ● [GPU] DES/LSST - NVIDIA Tesla K40m SL7 Interactive General Purpose Notebook ● [GPU] DES/LSST - NVIDIA Tesla T4 SL7 Interactive General Purpose Notebook 	<p>CMSLPC</p> <ul style="list-style-type: none"> ● CMSLPC - SL7 Interactive ● CMSLPC - COFFEA-DASK SL7 Interactive ● [GPU] CMSLPC - NVIDIA Tesla K40m SL7 Interactive ● [GPU] CMSLPC - NVIDIA Tesla T4 SL7 Interactive ● [GPU/GPUaaS] CMSLPC - NVIDIA Tesla T4 SL7 BDT Interactive
<p>LBNF/DUNE/ProtoDUNE</p> <ul style="list-style-type: none"> ● DUNE - SL7 Interactive General Purpose Notebook ● [GPU] DUNE - NVIDIA Tesla T4 SL7 Interactive General Purpose Notebook ● [GPU] DUNE - NVIDIA Tesla K40m SL7 Interactive General Purpose Notebook 	<p>FIFE/Neutrinos</p> <ul style="list-style-type: none"> ● FIFE - Interactive General Purpose Notebook ● [GPU] FIFE NVIDIA Tesla K40m SL7 Interactive General Purpose Notebook ● [GPU] FIFE NVIDIA Tesla T4 SL7 Interactive General Purpose Notebook 	<p>Fermi generic SL7/CC8</p> <ul style="list-style-type: none"> ● CSI DFC - SL7 Interactive


Start

Scientific Compute Services

- GlideinWMS
 - Improvements in scalability and Singularity support
 - Tokens fully integrated, GSI no longer needed
 - Python3 and CentOS8 now fully supported
- HepCloud (See also CPI)
 - Split-starter and Lumberjack HPC access methods
 - Enabling access to Rigetti Quantum computer(s)
- Jobsub
 - Streamlined jobsub client in testing
 - complete re-write is 1265 lines of code down from over 36k and should look the same to the end user
 - Moved to tokens for auth in new client
- Landscape (can we ever have enough stats?)
 - More Enstore statistics
 - More stats from Slurm for WC/LQ1 usage



Scientific Compute Services



• POMS/Project-py

- Deployed a ProductionShifter feature to enable production teams to more easily coordinate shift work
- Multiple users have greatly increased their job efficiency as best practices are defaults

• Spack

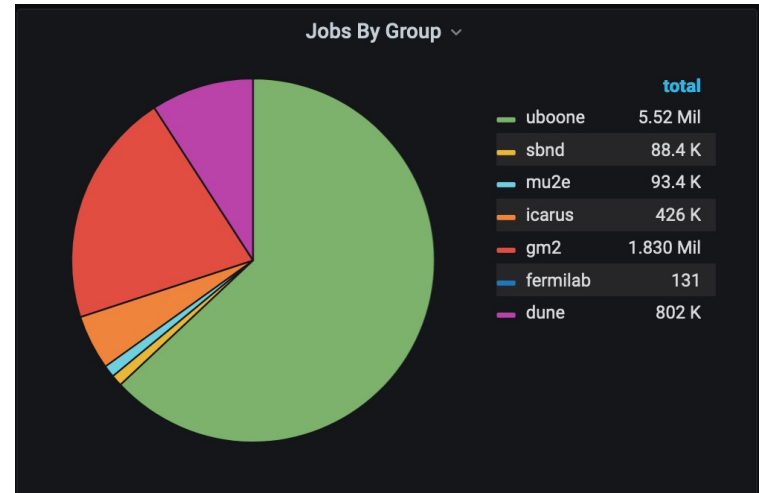
- Full build of Icaruscode for ALCF machines

• LQCD/Wilson

- Snow workflows setup for common Wilson tasks
- Integrated FERRY automation into admin flow
- Copying automation to LQ1 cluster next!

• LSST production

- Interim Data Facility a success. Two portals enabling test production and tooling.



Artificial Intelligence and Software for Physics Applications

AISP accomplishments: PDS

- Fermilab experiments and LQCD simulation efforts
 - Pythia8 development and support (8.304,5,6 releases)
 - Geant4 development and support (v10.7.p02,p03, v11 releases)
 - Optical photons in GPUs for liquid Argon detector applications
 - Example of Opticks integration to Geant4 completed and released
 - Completed example of how to decay taus and heavy quarks in Pythia8
 - Implemented in LArSoft
- Celeritas (G4 on GPUs)
 - Collaborating institutions: ORNL, FNAL, ANL and LBNL
 - Prototype includes Standard EM physics models for electron and gammas, a transportation loop, geometry navigation and support for uniform and non-uniform magnetic field
- Co-coordination of HEP Software Foundation Detector Simulation Working Group

AISP accomplishments: AS

• Publications

- Nagaitsev and **T. Zolkin**, Betatron frequency and Poincare Rotation Number, *Phys. Rev. Accel. Beams* 23 (2020) 5, 054001
- Mitchell, Ryne, Hwang, Nagaitsev, **T. Zolkin**, [Extracting dynamical frequencies from invariants of motion in finite-dimensional nonlinear integrable systems](#), *Phys.Rev.E* 103 (2021) 6, 062216
- **E. Stern**, et al., Self-consistent PIC simulations of ultimate space charge compensation with electron lenses, 2021 JINST **16** P03045

• Submissions and conference presentations

- **Q. Lu**, Upgrading Synergia for GPU Computing presented at SIAM March 2021
- **E. Stern**, et al., *Derivative-Free Optimization of a Rapid-Cycling Synchrotron* submitted to Optimization and Engineering

• Other notable accomplishments

- Upgraded Synergia build system
- Use WC Institutional Cluster resources with Spack based build
- Implemented more efficient Lie algebra calculations in Synergia for nonlinear normal forms
- Simulate PIP-II linac injection with phase space painting into the Booster

AISP accomplishments: ARA

Awards/recognition:

- **B. Nord**, Early Career Award, “Simulation-based inference for cosmological parameter estimation and discovery”
- **K. Pedro** et al, USCMS LHC S&C Award, “AI Denoising to Accelerate Detector Simulation”
- **K. Pedro**, **G. Perdue**, **B. Nord**, **A. Ćiprijanović**, DOE CompHEP “High Velocity AI: Generative Models”
- **N. Tran**, J. Ngadiuba et al, ASCR Data Reduction project, “Real-time data reduction codesign for the extreme edge”
- **N. Tran** et al, NP AI/ML project, “Intelligent experiments through real-time AI: Fast Data Processing and Autonomous Detector Control for sPHENIX and future EIC detectors”

AISP accomplishments: ARA

Papers and research milestones:

- "*DeepMerge II: Building Robust Deep Learning Algorithms for Merging Galaxy Identification Across Domains*": **A. Ćiprijanović, D. Kafkes**, K. Downey, S. Jenkins, **G. N. Perdue**, S. Madireddy, T. Johnston, G. F. Snyder, **B. Nord**, 2021, MNRAS, 605, 677-691
- Proceedings Paper from NeurIPS: "*Robustness of deep learning algorithms in astronomy-galaxy morphology studies*": **A. Ćiprijanović, D. Kafkes, G. N. Perdue, K. Pedro**, G. Snyder, F. J. Sánchez, S. Madireddy, S. M. Wild, B. Nord, arXiv:2111.00961. Workshop at the 35th Conference on Neural Information Processing Systems (NeurIPS), 13 December 2021, virtual conference.
- Community White Paper: "*Applications and Techniques for Fast Machine Learning in Science*", A. McCarn Deiana, **N. Tran** (editors), Submitted to Frontiers in Big Data, <https://arxiv.org/abs/2110.13041>
- Journal Paper: "*A reconfigurable neural network ASIC for detector front-end data compression at the HL-LHC*", G. Di Guglielmo, F. Fahim, C. Herwig, **N. Tran**, et al. IEEE Trans. Nucl. Sci. 68, 2179 (2021).
- NeurIPS benchmarks workshop: MLPerf Tiny Benchmarks (<https://arxiv.org/abs/2106.07597>) and submission to [MLCommons Tiny](#) (**B. Hawks, N. Tran**)
- Journal paper: "*Ps and Qs: Quantization-aware pruning for efficient low latency neural network inference*", **B. Hawks, N. Tran**, et al. Front. AI 4, 94 (2021).
- SONIC made large advances with production demonstrations for both CMS and ProtoDUNE and analysis facilities with Coffea (**K. Pedro, L. Gray, N. Tran**)
- **L. Gray** significant advances in *Graph NN for HGCal* (w/**G. Pradhan**) and ECal reconstruction
- **G. Pradhan** significant contributions to *L1 muon DL trigger algorithm*, pruning and quantization
- **P. Swamy** significant contributions to *AI for accelerator controls in READS* (real-time edge AI distributed systems)
- A lot of collaboration with other SCD groups!

AISP accomplishments: IES

- Publications

- 22+ papers published, 10 submitted and under review, 2 white papers authored (Aliaga-Soplin, Berkman, Cerati, Gardiner, Hobbs, Jena, Snider, Wang, Wospakrik, et al.)

- New leadership roles

- Co-convener of NOvA ND group (Aliaga-Soplin)
- GENIE Executive Board (Hobbs)

- Select technical accomplishments

- Deployment of parallelized track reconstruction for CMS Run 3 data (Cerati, Hall, et al.)
- Enable use of vectorized LArSoft code at HPC centers using Spack-based builds (Berkman, Cerati, et al)
 - Optimized hit finding algorithm 200x faster
- Development of GNN-based object reconstruction in LArTPC detectors and first application to DUNE and MicroBoone (Cerati, et al.)
- Extensible physics list factory adopted by Geant4, used in LArSoft for tau, charm decays (Hatcher, et al)
- Deployed GPUaaS capability in LArSoft (Wang)

AISP accomplishments: IES

- Select scientific accomplishments

- Leaders and major contributors to first results in search for MiniBooNE low-energy excess at MicroBooNE (Berkman, Cerati, Snider, Wospakrik)
- Delivered “MicroBooNE GENIE tune” used to interpret low-energy excess results (Gardiner)
- Measurement of muon neutrino, charged current double-differential cross section at NOvA (Aliaga-Soplin)
- Identified kinematic regions and hadronic interactions limiting precision of DUNE flux estimates, providing direction to hadron production experiments (Jena)
- Pioneering studies on treatment of nuclear effects in deep-inelastic scattering, and higher-order perturbative corrections for neutrino-nuclear cross sections (Hobbs)

- Awards

- **Mike Wang, LDRD award:** “Breaking the Big Data Bottleneck and Meeting the Realtime Constraints of Multi-messenger Astronomy in DUNE and LSST with Computational Storage and Machine Learning”

Scientific Data Services

New CMS Tape Library



- Spectra Logic TFinity Tape Library – SCD’s first from this vendor.
- Installation was finished on May 28th 2021
 - Procurement, installation, Enstore modifications and commissioning were the work of many – too many to thank individually

New CMS Tape Library

- 20 LTO-8 tape drives
- 150PB total capacity using LTO-8 tapes (12TB/tape)
- 12 frames (racks)
- Full redundancy for High Availability capacities (dual robots, isolation service frames, dual AC power modules, etc.
- First task for the new library is migration of remaining CMS data from Oracle T10k tapes.
 - Developed from scratch and running new, scalable, hands-off, migration software that uses dCache and delivers up to 170 TiB/day limited only by the number of drives
 - 56% complete (as of Tuesday)



Other storage accomplishments

- Public dCache upgraded to version 7.2 – with several of the new features contributed by us
 - Token support for xrootd
 - User/group quotas
 - Bulk REST API for file pinning and QoS transitions
- CDF & D0 media migration (very nearly) complete
 - Used the more scalable, less labor intensive, migration processes now extended to CMS
 - No more d0en and cdfen (for those that still remember them!)
- Total Enstore throughput (read+write) peaking to over 12 GiB/s, equivalent to 1 PiB/day. DC rate ~ 650 TiB/day

Data Management and database applications

- Rucio is in production on the OKD (Kubernetes) cluster for DUNE and ICARUS.
- Demonstrated the scalability of the new MetaCat metadata database by importing NOvA's existing metadata; 200 million files and 5.5 billion key-value pairs
- DUNE Hardware (Components) Database
 - First version completed and in production
 - Now integrating Parts DB which was originally separately developed at CERN
- FERRY
 - New APIs and features for working with Federated Identity and token authentication
 - FERRY is the source of truth for grid permissions and the data is now uploaded to LDAP to be used by the token issuer

Cross-cutting Projects and Initiatives

CPI: 2021 Accomplishments: VRO/LSST, CMS

VRO/LSST (Vera Rubin Observatory Legacy Survey of Space and Time)
Full operations start April 2024

1. Successful delivery of DP0.1 milestone (several FNAL EPRA awards given)
2. Network Architecture Design review and Support (on-going)
3. Data Curation leadership (inclusion of Rucio into Rubin operation plan)
4. Data Production operations (successful launch of DP0.2 production, due Jun '22)
5. Data Verification team involvement.
6. Coordination between Rubin/LSST and DESC groups (leverage DESSN/GW, GPU work).




Dome on Cerro Pachon, Chile
Photo credit: Rubin organization

CMS: Good utilization, lots of HEPCloud hours on HPC

New TFinity Tape Library Commissioned!
Nearly 10 PB of migrated data written since September

6 PB Storage added
5000 more CPU cores in process of delivery/burn in

Supply chain causing delays, but eventually getting here!



Fermilab

Fermilab

2013-01-01 2013-01-01 2013-01-01 2013-01-01 2013-01-01

CPI: 2021 Accomplishments: DUNE, LQCD/HPC

DUNE

- US Comp. R&D has **received funding** from the DOE (3 years: FNAL, ANL, BNL, 4 universities)
 - Databases, frameworks/event model, software infrastructure, training, new designs for ProtoDUNE II computing
- ProtoDUNE Prod 4/4a Campaigns complete: **3.8 PB, 15 M events** on OSG, WLCG, NERSC
- **Rucio** Data Management in production: **replicated 1.5 PB** of Prod4a files to UK storage element, consortium added NL, BR, IN, SP
- Successfully incorporated **SONIC-based GPU processing** into production workflow – and **saturated outbound GCC network** (sorry)
- Database instances in production for **Hardware DB, Metacat** (file catalog based on metadata)
- **Global Pool instance operating** and preparing for integration
- CCD Networking group has plan for **100 Gb/s path from SURF to FNAL** and multiple geographically distinct backup paths
- **Framework requirements taskforce finalizing** document after successful workshop with HSF Frameworks WG




LQCD/HPC

- Since July 1st, the **LQ institutional cluster** has delivered just over 29 million Sky-Lake core hours to the national USQCD lattice QCD collaborations
- The cluster is hosting ten large science projects, accounting for the bulk of the delivered time, plus several smaller projects having opportunistic access
- The SCS, SDS, and SCF departments together continue to operation the **Wilson institutional cluster** that has delivered small-scale HPC computing to the whole Fermilab community
- The cluster has delivered over 11 million core hours mainly for GPU and CPU development work as well as light production work.\

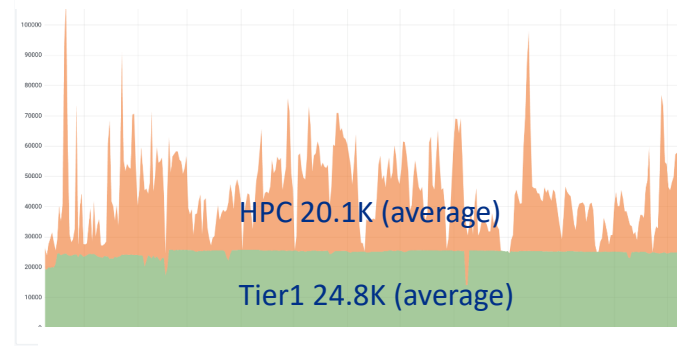


CPI: 2021 Accomplishments: HEPCloud

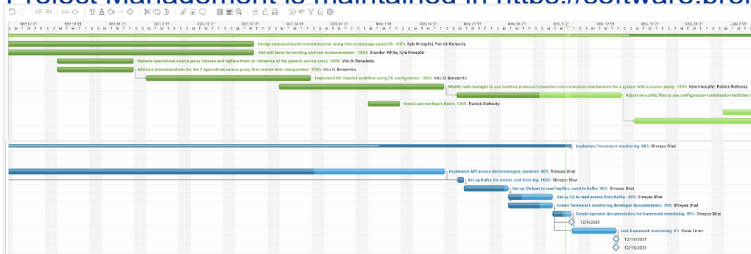

- Efforts contribution from 4 SCD departments
- Several **HPC** resources on-boarded by HEP Cloud and used by CMS during 2021 including PSC Bridges 2, TACC Frontera and **ALCF Theta**
- Coming in 2022: NERSC **Perlmutter**, Purdue **Anvil**
- Successful collaboration with **Fermilab Quantum Institute** and **Rigetti**
- **Decision Engine** Version 1.7 released in September 2021
- DE 2.0 planned mid-January, 2022
 - More scalable and **robust**, easy configuration, improved logging
 - Monitoring via Prometheus, **Grafana** and ES
 - Better architecture, **separation of GlideinWMS Frontend**
 - Export control approved, new **Apache 2.0** license
- **Project Management** is maintained in <https://software.proiectmanager.com>



CMS Claimed Cores on Tier1 and HPCs



CMS Claimed Cores by HPC resource



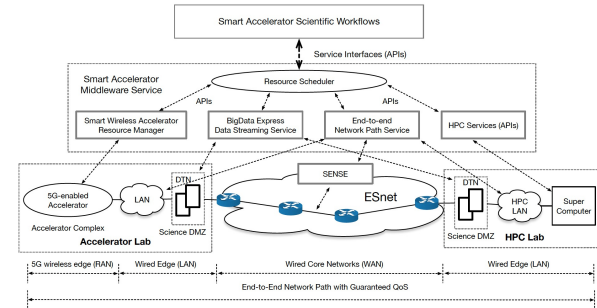
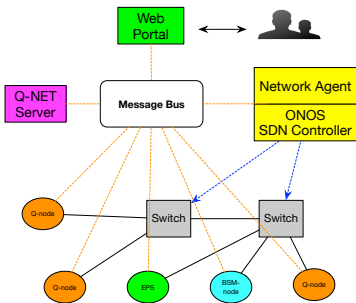
CPI: 2021 Accomplishments: R&D



- LDRD project developed **2 parallel algorithms for numerical integrations**; presented first at **SC21** and submitted second to ISC (European supercomputing). Currently in use for DES cluster cosmology; may scale up to use of 800 GPUs in a pipeline on Perlmutter. Versions in CUDA and Kokkos.
- Executed **C++ class through Zoom, with up to 70 students** – thanks to the many people who served at TAs. Reviews were positive – we may do it again next year.
- Completed first full year as DESC pipeline scientist: executing **major re-design of DESC tool for calculating DESC likelihood** for the Modeling and Combined Probes working group.
- HDF5** work gaining traction: invited to speak at The HDF Group Users Meeting, tools in use by multiple experiments, being looked into by others, now there is interest in HEP CCE. Parallel writing work with Northwestern developed scalable code that can **write in hours what previously took ~30 days** in a serial program.

5/6G R&D: Co-chaired Roundtable on Terahertz and 6G Wireless Communications in Science and Extreme Environments for DOE/ASCR

- Q-NetARCH: An **SDN-based Network Architecture for Quantum Networks** (presented at **SC21 Quantum Computing Workshop**)
 - Investigate, design, and develop a scalable, flexible, reliable, and platform-independent quantum network architecture that supports essential quantum network functionalities and capabilities.



CPI: 2021 Accomplishments: Projects



CTA (CERN Tape Archive)

- Making considerable progress evaluating CTA for Fermilab
- First accepted merge request in CTA code!

Storage R&D

- Setting up Ceph cluster for Object Store for Science evaluation

Appointed as **LPC Coordinator** for two year term, started September 2021

EAF

- See SCS slide!



AI

- Successfully demonstrated inference-as-a-service at scale in Cloud, on-prem, in EAF

Frameworks, DAQ and Electronics

Tools and Advanced Computing

artdaq and otsdaq

Provided 3 releases in 6 months of multi-experiment DAQ system

CMS

Migrated all essential CMSSW code to allow concurrent conditions loading

DUNE-DAQ

Migrated to GitHub software repository for issue tracking and integration building

Moving to Spack software packaging system


Provided 5 DAQ releases in last 6 months

Icarus

Provided 10 new DAQ releases in past 6 months

Mu2e

Installed first batch of nodes in the Mu2e DAQ room for integration tests



Cosmic Survey Support

- Vera Rubin Observatory

- Most of the group is working on preparation, processing, validation, and user support of Data Preview 0.1 and 0.2 in the Rubin Science Platform: submitting and monitoring jobs, validating output, diagnosing and reporting issues to the developers and educating users.
- Eric Neilsen is working on scheduler simulation analysis tools and plans for monitoring survey progress.
- Huan Lin began working as a DESC Photometric Redshift (PZ) Infrastructure Pipeline Scientist.
- Douglas Tucker began a term as co-convener of the DESC photometric calibration working group.




- DESI

- Steve Kent and Eric Neilsen continue support and refinement of PlateMaker, critical operations software for computing fiber positions. DESI recorded 41 tiles containing 141K redshifts in a single night last week. DESI now has more spectra than any other redshift survey.
- Liz Buckley-Geer supports the online database. Critical for all DESI operations.


- DES

- Huan Lin just completed a 4 year term as DES redshift working group co-coordinator.
- Several million gpGrid (DEgrid cluster nodes) core hours were used by the Transient and Moving Object group to find several hundred Trans-Neptunian Objects (& a comet).



Detector Electronics Group

- **In 2021, responsible for...**

- Physics Research Equipment Pool (PREP Counter at FCC1)
- New start! DUNE Far Detector 2 Photon Detector System
 - Level 2 project management, cold electronics design (-300° F in LAr)
- Mu2e Trigger & Data Acquisition 
 - Level 2 project management, artdaq/otsdaq, DAQ installation underway
- Test beam support (Silicon tracker, EMPAHTIC)




DUNE Cryogenic testing of digital and analog readout electronics at PAB

- **In 2021, collaborating with...**

- AD on PIP-II μ TCA readout and Machine Protection
- CMS on Correlator, Muon, and Global Track Trigger FPGA
- CMS Phase II upgrades (Tracker, Timing, HGCal)

- **In 2021, real-time AI with...**

- Xilinx Research, Booster, DUNE, hls4ml 
- MLPerf Tiny anomaly detection; Fast ML benchmarks



CMS FPGA card



Mu2e DAQ room

Framework and Software Technology

CMSSW

- Decided to adopt Alpaka for CPU/GPU software portability
- Working to extend concurrency above the luminosity block layer
- Improved static analysis tooling

Intensity frontier

- art has moved its repositories and issue management to GitHub
- LArSoft upgrading to Spack-supporting build system

HEPCloud

- Significant framework improvements

DUNE DAQ


- Many packages have been adjusted to support Spack

Collaborative efforts

- HSF conveners/contributors interacting with broader HEP community
- HEP-CCE involvement exploring code portability and I/O usage on HPC systems.

DAQ Controls & Detectors

- The QIS R&C system QICK paper is out. [arXiv:2110.00557](https://arxiv.org/abs/2110.00557) and accepted for publication in RSI (Review of Scientific Instruments).
- QICK is making a big splash in QIS. It seems like all DOE quantum centers want to adopt it. Several universities (UC, Princeton) have already done that. Even some companies (e.g. Rigatti) have put their eye on QICK. Fermilab QSC will also adopted QICK for their program of radiation studies in qubits.
- Dark matter and dark energy detectors
 - readout and control for MKIDs for CMB SPT4 (stakeholder Adam Anderson, Brad Benson).
 - FPGA for broadband axion search experiment BREAD (stakeholder Andrew Sonnenschein).
- GQuest proposal has been submitted and approved.
 - This is a gravitational wave interferometer in collaboration with several LIGO people (Caltech, MIT).
 - It is a \$10M project (total) for 5 years. It covers ~1.5FTEs for our department.
- We have submitted a full proposal for an LDRD grant in Quantum Networks using RFSoc.
- We have made improvements to skipper CCD readouts. In particular, to make the readout faster.
- We are designing cold electronics for QIS, at 4K and below.
- We designed the analog electronics for DUNE vertical drift 40 SIPM readout.
 - The readout is based on the hybrid 40 SIPM ARAPUCA configuration.



Scientific Computing Facilities

SCF in 2021

Installations:

- Disk storage: 16 PB public dCache, 11 PB CMS
- CPU farms: 80 CMS Tier 1 workers, 10k cores
- ~50 DAQ and other servers
- More ordered, in progress



DUNE DAQ:

- Support for ProtoDUNE DAQ computing
 - “Pocket” Kubernetes-based deployment
- Support for ICEBERG test stand
- Support for ArgonCube 2x2 demonstrator
- Networking plan for far detector (in progress for near detector)



SCF in 2021

Linux distribution:

- End of CentOS Linux announced late 2020
 - Compare options: RHEL, CentOS Stream, clones
 - Discussions with experiments, CERN, others
 - Recommend CentOS Stream
- SL 7 support continues; Fermilab resources for Stream 8 & 9



OKD/Kubernetes and container services:

- Elastic Analysis Facility
- "Production" services: Rucio, CTA, HTCondor
- Testing Ceph implementation to use with OKD
- Harbor image registry



SCF in 2021

IPv6:

- Collected information for DOE data calls
- Extensive work to implement support for dual stack, and test IPv6-only.



HPC: Completed Wilson Cluster transition; operations



CMS:

- Additions to CMS LPC interactive cluster:
 - Larger nodes for code builds
 - Test IPv6, test CentOS Stream 8
- Additions to Tier 1 farm, dCache, EOS




SCF in 2021

Shared services—examples:

- SCD Slack
- Jenkins
- Checkmk monitoring
- Developer effort for HEPCloud



Jenkins



Ongoing support for experiments—offline and online/DAQ

- Interactive nodes: GPCF (“gpvms”), CMS LPC
- DAQ: ICARUS, SBND, g-2, MicroBooNE, ANNIE, others
- csresearch cluster
- Database servers hosting ifbeam, experiment DBs, elogs, ...
- etc.



SCD Science

SCD Scientific Publications

SCD Scientists are contributing to Physics at Experiments and Projects

- Concentrating on 2021 publications for this all-hands meeting.
- Adam asked the SCD Scientific Community to send a list of publications for which they **had made a major contribution** (did analysis, wrote paper, reviewed, etc.)

- Remember, not all experiments are in the publishing phase
- Remember, important contributions may not lead directly to one publication (e.g., operations work)

- Next slides are the lists **that were sent**
Only SCD authors are shown (order is author order on paper)
- Conference proceedings are not included
- Adam owns any mistakes/omissions (sorry in advance!)

SCD Scientific Publications

CMS

- R. Harris, Measurement of the double-differential inclusive jet cross section in proton-proton collisions at $\sqrt{s} = 5.02$ TeV
- R. Harris, Search for high mass trijet resonances using final states with boosted dijet resonances in proton-proton collisions at $\sqrt{s} = 13$ TeV
- R. Harris, Search for heavy resonances decaying to b quarks in proton-proton collisions at $\sqrt{s} = 13$ TeV
- B. Jayatilaka, Search for new particles in events with energetic jets and large missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV

SCD Scientific Publications

Cosmology

- J. Annis, H. Lin, D. Tucker, B. Yanny, The Dark Energy Survey Data Release 2
- J. Annis, S. Kent, Reducing Ground-based Astrometric Errors with Gaia and Gaussian Processes
- E. Buckley-Geer, B. Nord, DeepZipper: A Novel Deep Learning Architecture for Lensed Supernovae Identification
- E. Buckley-Geer, H. Lin, D. Tucker, B. Yanny, B. Nord, The DES Bright Arcs Survey: Candidate Strongly Lensed Galaxy Systems from the Dark Energy Survey 5,000 Sq. Deg. Footprint
- A. Ćiprijanović, K. Pedro, B. Nord, DeepAdversaries: On Adversarial and Noise Robustness of Deep Learning Algorithms in Astronomy
- A. Ćiprijanović, B. Nord, M. Wang, DeepGhostBusters: Using Mask R-CNN to Detect and Remove Ghosting and Scattered-Light Artifacts from Optical Survey Images
- A. Ćiprijanović, Validating Synthetic Galaxy Catalogs for Dark Energy Science in the LSST Era
- A. Ćiprijanović, B. Nord, DeepMerge II: Building Robust Deep Learning Algorithms for Merging Galaxy Identification Across Domains
- A. Ćiprijanović, DeepShadows: Separating Low Surface Brightness Galaxies from Artifacts using Deep Learning
- K. Herner, M. Paterno, J. Annis, Dark Energy Survey Year 3 Results: Multi-Probe Modeling Strategy and Validation
- S. Kent, B. Nord, M. Wang, A machine learning approach to the detection of ghosting and scattered light artifacts in dark energy survey images
- H. Lin, B. Yanny, J. Annis, Dark Energy Survey Year 3 results: redshift calibration of the weak lensing source galaxies
- B. Nord, J. Annis, E. Buckley-Geer, H. Lin, Finding quadruply imaged quasars with machine learning. I. Methods
- D. Tucker, J. Annis, H. Lin, B. Yanny, SOAR/Goodman Spectroscopic Assessment of Candidate Counterparts of the LIGO--Virgo Event GW190814
- D. Tucker, Photometric cross-calibration of the SDSS Stripe 82 Standard Stars catalogue with Gaia EDR3, and comparison with Pan-STARRS1, DES, CFIS, and GALEX catalogues
- D. Tucker, B. Yanny, The DECam Local Volume Exploration Survey: Overview and First Data Release

SCD Scientific Publications

Neutrino Physics

MicroBooNE: (S. Berkman, G. Cerati, S. Gardiner, W. Ketchum, M. Kirby, E. Snider, M. Wospakrik)

Search for an Excess of Electron Neutrino Interactions in MicroBooNE Using Multiple Final State Topologies

Semantic Segmentation with a Sparse Convolutional Neural Network for Event Reconstruction in MicroBooNE

The Continuous Readout Stream of the MicroBooNE Liquid Argon Time Projection Chamber for Detection of Supernova Burst Neutrinos

Search for an anomalous excess of charged-current ν_e interactions without pions in the final state with the MicroBooNE experiment

Novel Approach for Evaluating Detector-Related Uncertainties in a LArTPC Using MicroBooNE Data

Measurement of the Longitudinal Diffusion of Ionization Electrons in the MicroBooNE Detector

DUNE: (S. Berkman, K. Biery, G. Cerati, P. Ding, J. Freeman, S. Fuess, S. Gardiner, R. Hatcher, K. Herner, D. Jena, W. Ketchum, M. Kirby, A. Norman, E. Snider, S. Timm, M. Wang)

Design, construction and operation of the ProtoDUNE-SP Liquid Argon TPC

First results on ProtoDUNE-SP liquid argon time projection chamber performance from a beam test at the CERN Neutrino Platform

SCD Scientific Publications

Muon Physics

Muon g-2: (A. Lyon, J. Stapleton, T. Walton)

Measurement of the Positive Muon Anomalous Magnetic Moment to 0.46 ppm

Measurement of the anomalous precession frequency of the muon in the Fermilab Muon g-2 Experiment

Magnetic-field measurement and analysis for the Muon g-2 Experiment at Fermilab

Beam dynamics corrections to the Run-1 measurement of the muon anomalous magnetic moment at Fermilab

Theory

S. Mrenna, A Study of QCD Radiation in VBF Higgs Production with Vincia and Pythia

S. Mrenna, BROOD: Bilevel and Robust Optimization and Outlier Detection for Efficient Tuning of High-Energy Physics Event Generators

COVID-19

G. Lukhanin, M. Paterno, M. Wang, The Novel Mechanical Ventilator Milano for the COVID-19 pandemic


SCD Scientific Publications

Computing

- S. Berkman, G. Cerati, K. Knoepfel, M. Mengel, A. Hall, M. Wang, Optimizing the Hit Finding Algorithm for Liquid Argon TPC Neutrino Detectors Using Parallel Architectures
- W. Ketchum, M. Paterno, S. Sehrish, M. Wospakrik, Grid-based minimization at scale: Feldman-Cousins corrections for SBN
- S. Mrenna, E. Sexton-Kennedy, Challenges in Monte Carlo event generator software for High-Luminosity LHC
- S. Mrenna, Multivariate Rational Approximation
- M. Paterno, PAGANI: A Parallel Adaptive GPU Algorithm for Numerical Integration
- S. Sehrish, M. Paterno, A Case Study on Parallel HDF5 Dataset Concatenation for High-Energy Physics Data Analysis
- M. Wang, B. Holzman, K. Knoepfel, K. Pedro, N. Tran, GPU-Accelerated Machine Learning Inference as a Service for Computing in Neutrino Experiments
- M. Wang, H. Wenzel, Extracting low energy signals from raw LArTPC waveforms using deep learning techniques – A proof of concept
- M. Wang (ArgoNeuT), A deep-learning based raw waveform region-of-interest finder for the liquid argon time projection chamber
- M. Wang, A systematic study of projection biases in weak lensing analysis

Quantum Computing

- A. Lyon, Impacts of Noise and Structure on Quantum Information Encoded in a Quantum Memory
- A. Lyon, Simulating Large PEPs Tensor Networks on Small Quantum Devices
- S. Mrenna, Bosonic field digitization for quantum computers
- L. Stefanazzi, K. Treptow, N. Wilcer, C. Stoughton, G. Cancelo, The QICK (Quantum Instrumentation Control Kit): Readout and control for qubits and detectors



Happy Hanukkah

