



Scientific Computing Division

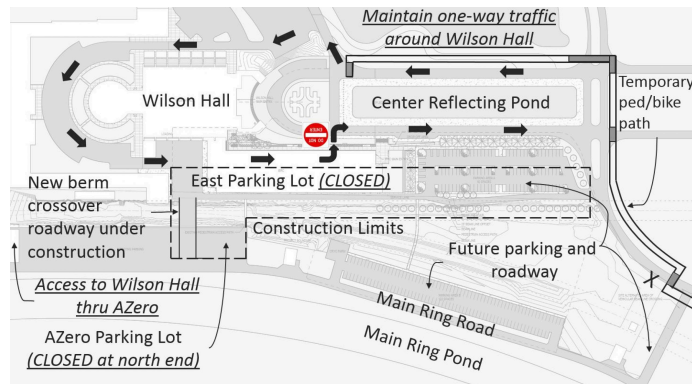
James Amundson
CIO All-hands Meeting
2019-10-28

Top 5 (done)

Safety


Safety

- Indications that the lab had seen a small decline in commitment to safety
 - Need to remind everyone that your safety is the highest priority
- SCD has been doing a very good job
 - Remember: complacency is a major error precursor!
- IERC construction
 - Foot trips between FCC and WH now require ~73 road crossings



Safety

- Indications that the lab had seen a small decline in commitment to safety
 - Need to remind everyone that your safety is the highest priority
- SCD has been doing a very good job
 - Remember: complacency is a major error precursor!
- IERC construction
 - Foot trips between FCC and WH now require ~73 road crossings



New People

Welcome to new people (transfers)

- Steven Gardiner
 - Postdoc
 - Transferred from Neutrino Division
 - Working on Neutrino Simulations
 - Joined Neutrino Simulation group
- Nhan Tran
 - Wilson Fellow
 - Transferred from Particle Physics Division
 - Working on AI and CMS
 - Leading the Algorithms for Reconstruction and Analysis group
 - Recently won a DOE Early Career Award



Welcome to new people (former postdocs)

- Kevin Pedro
 - Associate Scientist
 - Working on CMS, AI, new architectures
 - Joined Algorithms for Reconstruction and Analysis group
- Timofey Zolkin
 - Peoples Fellow
 - Working on accelerator simulation
 - Joined (still in) Accelerator Simulation group



Welcome to new people (actual new person)

- Gavesh Kirby Jayatilaka, born October 25, 2019; 6 lb., 14 oz.



Climate Survey

SCD Survey results

- Similar results to lab as a whole
 - 63rd percentile (SCD) vs. 61st percentile (lab)
- High point: “My supervisor treats me with respect”
 - 87th percentile (SCD) vs. 80th percentile (lab)
 - Lower ratings for “management”
 - Everyone loves their congressperson, but everyone hates congress
- Low points: diversity and sexual harassment
 - Also, career development

63. I understand Fermilab's sexual harassment policy.	Total Population	1,358	5.54	80
	Scientific Computing Division	117	5.56	82
62. I understand how I can anonymously report business abuses such as theft, fraud, or sexual harassment.	Total Population	1,342	5.19	76
	Scientific Computing Division	115	5.28	80
65. Sexual harassment is effectively addressed at Fermilab.	Total Population	1,311	4.77	46
	Scientific Computing Division	106	4.50	44
64. Sexual harassment does not occur at Fermilab.	Total Population	1,322	4.14	43
	Scientific Computing Division	106	3.89	40

Thanks to all who attended the SCD Climate Survey All-hands

Actions

- Diversity and sexual harassment
 - Recruiting at Grace Hopper
 - Conference for women in computer science
 - Recruiting at Society of Women Engineers
 - Diversity recruitment efforts to follow as positions become available
- Relationship with management/appreciation
 - Reorganizing the Division
 - Removed one layer of management
 - People closer to top management
 - Weekly division head/group lunches
- Career development
 - Identifying wider group of people who would like to develop programming skills
 - Pair programming model to share expert knowledge with developing programmers

Reorg

5 Core departments + 1 cross-cutting department

Scientific Data Services

Scientific Compute Services

Scientific Computing Facilities

AI and Software for Physics Applications

Frameworks, DAQ and Electronics

Cross-cutting Proj and Ini

Senior management and standalone groups

- Associate Head for Projects: Margaret Votava
- Associate Head for Science: Adam Lyon
- Associate Head for CMS: Oli Gutsche
- Associate Head for DUNE: (vacant!)
- Associate Head for Facilities: Adam Walters

- Associate heads report to the Deputy Division Head: Stu Fuess
- Deputy Division Head reports to the Division Head

- Two standalone groups
 - Data Center Operations
 - Science
 - (science that does not fit in departments)

Status

- Lab systems (Workday, Kronos, etc...) should be updated by the end of the week
 - Tremendous amount of work on the part of Carla Lloyd, Stu Fuess, Jackie Zolna, and the financial group
 - Thanks to all!
- Finally!
 - Note to self: summer is a terrible time to start a reorganization

Division Status

Community of SCD Scientists


- The re-organized SCD Science Organization
- Adam Lyon is the Associate Head for Science
- Big thank you to Robert Harris, Rob Kutschke, Daniel Elvira, Steve Kent for their management and hard work
- Features of the new organization:
 - Continue focus on mentoring, especially for all postdocs (Robert Harris is Head of Mentoring)
 - Postdoc seminars will continue. Next one is James Stapleton from Muon g-2 on Nov 5.
 - No internal structure by frontier
 - Increase opportunities to share, collaborate and interact
 - Organize more along topics
 - Host the “Joint Task Force Initiative” **Computational Science Seminar**. A rotating seminar among UChicago, Argonne and us. Inaugural seminar is Nhan Tran about AI at UChicago on Nov 12.
- More info on all of the above will be coming soon! Science accomplishments will be featured at next All-Hands

Standalone Group: Data Center Operations

Computer Room	Availability Goals		Planned and Unplanned Outages		Unplanned Outages Only	
	Availability (goal%)	Downtime (hours)	Availability (actual %)	Downtime (hours)	Availability (actual %)	Downtime (Hours)
FCC2	99.67	28.80	100.000	0.00	100.000	0.00
FCC3	99.67	28.80	100.000	0.00	100.000	0.00
GCC-CRA/NRA/TRR	99.50	43.92	99.824	15.44	100.000	0.00
GCC-CRB/NRB	99.50	43.92	99.915	7.42	100.000	0.00
GCC-CRC	99.50	43.92	99.941	5.17	100.000	0.00
WH8-FC	99.67	28.80	100.000	0.00	100.000	0.00

DOE Quarterly Data Calls only recognize unscheduled outages. Last unplanned outage in 2016

- Mitigated Scheduled Lab Outages 2019 - 4-Site, 2-GCC, 4-WH8-FC, 2-FCC, and 3-GCC DCO scheduled maintenance outages
 - Data Center Operations responds and mitigates unscheduled, off hours operational problems every 4.5 days
- Managed lots of equipment repairs, elevator upgrades, UPS battery replacements, 208V electrical distribution, bunch of water leaks, automation of energy readings for DOE reporting. WH8-FC Cooling Upgrade Design & Planning, Lab GPP/Infrastructure Planning (5 year plan), ESH and other sub-committees, working tirelessly with Procurement on maintenance contracts
- 2019 GCC Energy Star Award, GCC PUE ~1.6, FCC PUE ~1.6





Department: Scientific Data Services

- Maintains primary scientific data of **all Fermilab experiments**
 - Part of US CMS Tier 1; **largest store of CMS data** outside of CERN
- Seven tape robots (193 drives) with 0.7 EB capacity total
- **Over 180 PB of data** currently stored and managed
 - Expected to double by 2022
- Since 2018: **complete overhaul of tape technology**
 - Necessitated by industry changes
 - First new library: delivery to production in 3 months, **in time for protoDUNE beam data**
 - Migration of **110 PB** of data to new tape format underway



Department: Scientific Data Services

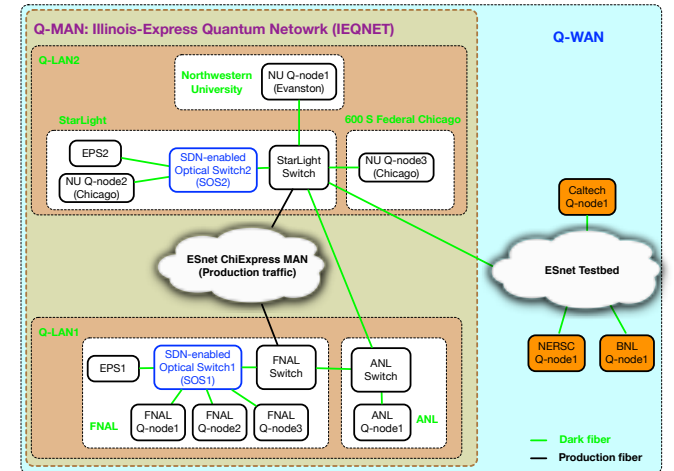

- dCache and EOS
 - Providing **~46 PB** of disk capacity for production, data distribution, and end-user analysis use
- Data management
 - CMS and DUNE both in the process of moving to Rucio, a common data management system
- SAM data management
 - Maintaining data management services for existing Fermilab experiments
 - DZero SAM finally turned off after **20 years** of service



Department: Scientific Data Services

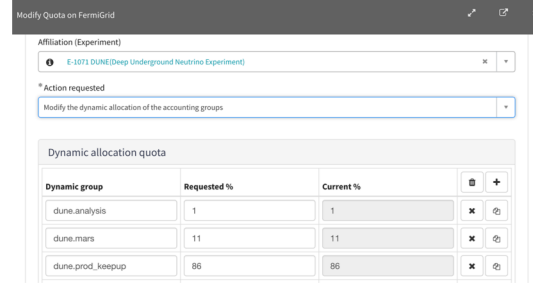
- Scientific Database Applications
 - Provides conditions, hardware, and other database applications to multiple experiments
 - Expanding the beams database to the ProtoDUNE beamlines at CERN was an important part of being able to process the test beam data
- Network Development
 - Illinois-Express Quantum Network (IEQNET)
 - Aim to build a metropolitan-scale quantum network testbed that demonstrates important advanced quantum networks capabilities beyond the lab.
 - Funded by DOE ASCR, \$3.2M, announced on Sep, 2019.

IFBeam quarterly growth

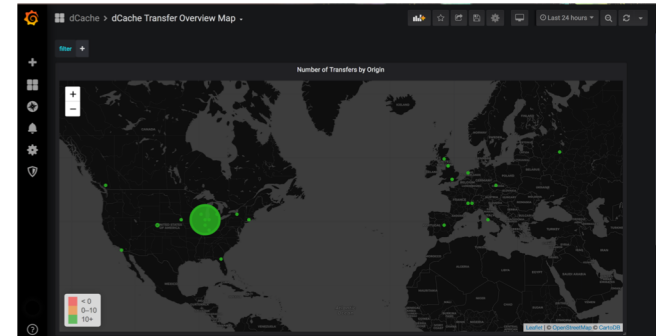


Department: Scientific Compute Services

- FERRY - Central registry of scientific computing administrative relationships
 - FERRY team released v2.0 which is being used by many scientific services to automate authorization needs
 - Integration of common tasks (e.g. experiment subgroup quota mgmt., adding users to experiments) with the Service Now team has streamlined operations and made us more efficient.
- Landscape - Common infrastructure for scientific services monitoring
 - Landscape team has upgraded Elasticsearch, Grafana, Kafka, and other tools to provide a valuable view into our computing
 - Service providers use it to monitor their services, Users get views of their computing activity
- Continuous Integration
 - Automatic builds of onboarded products has shortened the time needed to find problems introduced by code changes




Dynamic group	Requested %	Current %
dune.analysis	1	1
dune.mars	11	11
dune.prod_keepup	86	86




Department: Scientific Compute Services

- User and Experiment Support
 - Integrating and onboarding new sites
 - Jobs running in experiment-built containers
 - Working on art/g4 multithreading
- HEPCloud - access to commercial and exascale computing for HEP community
 - HEPCloud facility in production since March 2019 expanding our computing power
- HTCondor
 - Enabled user selection of Singularity containers on FermiGrid and CMS T1/LPC
 - Investigating token auth within HTCondor
- Fermicloud - essential VM resource for developers
 - Running stably since 2011
 - Converting to OpenStack as time permits
- CMS/OKD
 - Enabling new analysis methods (coffea) via SPARK in Kubernetes
 - Exploring HTCondor in Kubernetes




ProtoDUNE reprocessing campaign



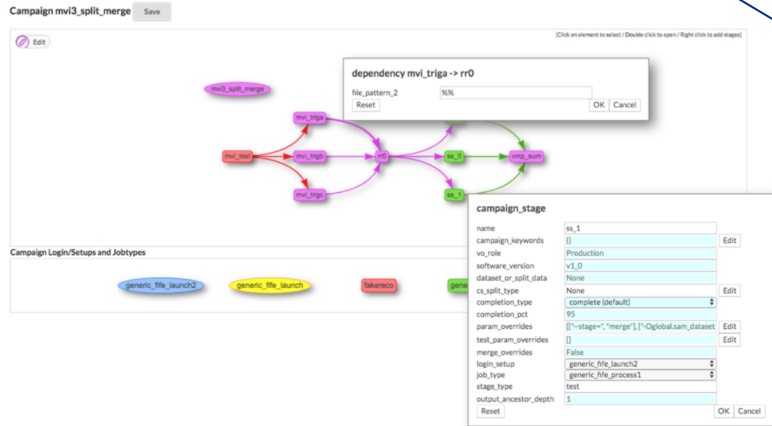
Department: Scientific Compute Services

- GlideinWMS
 - GWMS team facilitating access to millions of core hours of computing around the globe for Fermilab researchers
- Facility Interfaces and Workflow management (Jobsub/POMS)
 - New releases providing needed features via a common interface for job submission and mgmt
- CMS production
 - Produced 45 billion MC events in the last 12 months
 - Transitioning legacy to new infra for Run 3
- CVMFS - code and file distribution
 - CVMFS team released rapid code distribution feature enabling efficient transfer of custom user code to jobs
- LQCD
 - Wilson Cluster refresh in the works
 - Learning and planning for SLURM administration



"DUNE has a very positive interaction with the (POMS) developers and appreciates their willingness to incorporate DUNE's requirements."


"The Muon g-2 experiment uses POMS for data processing and MC production. We found it beneficial for submitting, tracking, bookkeeping monitoring grid jobs. In addition, with POMS we can automate our production and use the multistep-processing of our data which will be mandatory for the next round of data taking in October 2019. POMS is a great asset to the offline team of Muon g-2 experiment."



Department: Scientific Computing Facilities

What do we do?

- Stable operation of hardware, OS (architect, install, configure, monitor, retire):
 - CMS, FermiGrid batch systems
 - HPC clusters: LQCD, Wilson
 - Online/DAQ computers for MicroBooNE, g-2, NOvA, FTBF; soon for ICARUS, SBND, Mu2e, DUNE
 - Control rooms: ROC East/West, g-2, FTBF, NOvA far detector
 - GPCF and CMS interactive systems
 - System administration tools: Puppet, Check_MK, ...



Department: Scientific Computing Facilities

What do we do?

- Linux OS support
 - Package, distribute, support Scientific Linux 6 & 7
 - Preparing for CentOS 8 support
 - Help troubleshoot OS problems; help file bug reports
- Jenkins CI/CD system
- Software containers: Kubernetes (with OKD/OpenShift); platform to build & deploy service containers; registry



Department: Scientific Computing Facilities

What else have we been working on?

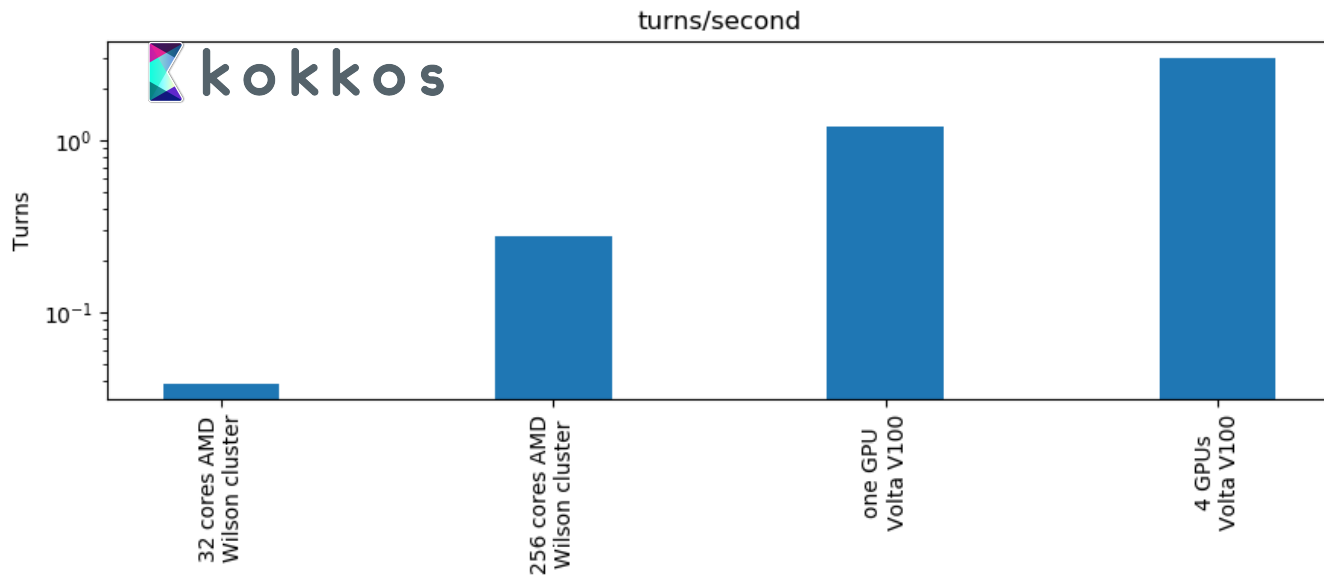
- Transition to IPv6
- Planning for DUNE, especially online/DAQ computing
- ICARUS, SBND starting up
- Testing new hardware: NVMe storage, fabrics; GPUs, FPGAs; new processor models
- LQ1 HPC cluster—first piece of a shared institutional cluster
- Working with CMS on other initiatives: containers, Apache Spark



Department: Artificial Intelligence and Software for Physics Applications (Accelerator Modeling)

Speedup with GPUs!

Large benchmark circular accelerator (protons – includes space-charge effects)



The Kokkos framework allows us to run the same code on both CPUs and GPUs. We achieved a huge speed up running simulations on GPUs (Qiming Lu, Eric Stern)



Department: Artificial Intelligence and Software for Physics Applications (Generators, detector, neutrino simulation)

- PYTHIA8 support of experiments (Steve Mrenna) – new release in preparation
- GeantV prototype (Ph. Canal – S. Banerjee, K. Genser, S. Jun, G. Lima, K. Pedro)
 - GeantV prototype achieves ~ 1.4 -2 speedup with respect to Geant4 (depends on memory cache size)
 - R&D also translated into performance improvements of Geant4 (>1.1)
- Geant4 support of experiments + profiling & benchmarking & hadronic validation of official and internal releases (S. Banerjee, K. Genser, S. Jun, J. Yarba)
 - For LAr experiments: Opticks/OptiX optical photons G4 package integration on GPU (ongoing – H. Wenzel)
- New LArSoft Geant4 interface (LArG4) released
 - usage examples created, starts being used (H. Wenzel)
- GeantX (HPC/GPUs): ECP and HEP collaboration (Ph. Canal leader + part of GeantV team)
 - Preparing for next generation of supercomputers with efficient use of GPUs
- Detector R&D, development of test infrastructure & studies of detector effects (Adam Para)

Department: Artificial Intelligence and Software for Physics Applications (Generators, detector, neutrino simulation)

- Fermilab GENIE team is working hard for GENIE to meet the needs of Fermilab experiments.
 - R. Hatcher completed project to add Geant4 nuclear models to GENIE
 - L. Fields organizing workshop to discuss streamlining process to get theoretical models to ν experiments
- NS group postdocs have been appointed to leadership positions on MicroBooNE (Steven Gardiner), NOvA (Leo Aliaga), and MINERvA (Deepika Jena)
- Papers submitted to conferences/journals:
 - "Geant4 parameter tuning using Professor". JINST, arXiv, and ACAT (L. Fields editor, S. Jun, Robert Hatcher, J. Yarba)

Identified ways to
improve Geant4 physics
performance



Highlights where more degrees
of freedom are needed.

- "Vectorization of RN generation and reproducibility of concurrent particle transport simulation" (Soon Jun)

Department: Artificial Intelligence and Software for Physics Applications (AI and Reconstruction)

- Neutrino and Mu2e (S. Berkman, G. Cerati, R. Kutschke, G. Perdue, E. Snider, M. Wospakrik, M. Wang)
 - LArSoft:
 - **Large-scale data production campaigns for ProtoDUNE, ICARUS, and MicroBooNE**
 - Making code thread-safe (large workflows and improved efficiency), migrate to GitHub
 - LArSoft workshop well-attended by SCD and neutrino community
 - SciDAC work: implemented improvements in hit finder with parallelization/vectorization
 - General development, support, coordination of reconstruction across neutrinos
 - e.g. First data-driven energy calibration of μBooNE with $\pi^0 \rightarrow \gamma\gamma$ (joint with ND)
<https://arxiv.org/abs/1910.02166>
- Cosmic (Nord, Caldeira) <https://arxiv.org/abs/1810.01483>
 - New study in deploying AI for CMB de-noising
Papers in review: galaxy mergers, strong lens regression, SZ signal cluster detection
 - LDRD for deep learning pipelines in HEP; JTFI grant for telescope scheduling/controls

Department: Artificial Intelligence and Software for Physics Applications (AI and Reconstruction)

- CMS/LHC (G. Cerati, L. Gray, A. Hall, K. Pedro, N. Tran)
 - mkFit parallelized track fitter – adapting Kalman Track fitter to SIMD architectures
<https://arxiv.org/abs/1906.11744>
 - Graph Neural Networks for Calorimetry LDRD: new ML architectures to learn representations from higher dimensional non-Euclidean detector inputs
 - Accelerated machine learning as a Service using FPGAs with Microsoft
<https://arxiv.org/abs/1904.08986>
- High velocity AI grant (DOE) to accelerate inference and training for uncertainty quantification
- ExaTrkX project kicking off: advanced track finding and pattern recognition with ML, building off of successful HepTrkX

Department: Frameworks, DAQ and Electronics

- Moving LArSoft to use GitHub
- Framework discussions between SCD and DUNE have begun
- Extending the open source build tool Spack to support HEP and neutrino experiment use, including LArSoft software stack support and parallel package development (new feature).
- Supporting CMS Offline software and improving its efficiency when using many-core systems
- HSF frameworks working group has restarted with CMS and neutrino representation from Fermilab
- Continuing to support HEP and neutrino experiments with new external software and compilers, and troubleshoot and resolve problems with external software (most recently, ROOT schema evolution).

Department: Frameworks, DAQ and Electronics

- Supporting the DAQ for NOvA, SBN, ProtoDUNE & DUNE
- Leading the Trigger & Data Acquisition subsystem for Mu2e and we are in the production phase - one year from taking cosmics!
- *artdaq* and *otsdaq* were chosen as the online DAQ software for Mu2e. Other customers include CMS Outer Tracker and the Fermilab Test Beam Facility.
- We are heavily subscribed to FPGA firmware projects around the lab including for Mu2e, PIP-II, USCMS Correlator Trigger, USCMS Endcap Muon Trigger, and real-time Machine Learning applications.
- PREP is supporting many experiments with equipment that is significantly reducing their costs. Keeping up with all the experimenters' requests, 10's per month.

Department: Frameworks, DAQ and Electronics

- DESI has first light!
- Playing an important role in getting cluster cosmology science out of DES.
- DES Y6A1 was released to the DES Collaboration
- Working on roles for FNAL in the LSST Operations
- LTA-QSM electronics picked as readout and DAQ for skipper CCD experiments (SENSEI, CONNIE, DM 10Kg).
- fMESSI readout being used for CMB-S4 R&D
- Next version of fMESSI is being developed and will cover needs for CMB-S4, ADMX, QC, and MKIDs.
- The ARAPUCA photon detector for DUNE has achieved 4% efficiency with active array of SIPMs
- Fermilab is leading an effort in readout and control electronics for fault tolerant quantum computers, with partners at UC and MIT. We are already controlling qubits.

Thank you!
Questions?