Updates in Scientific Computing at Fermilab

Shreyas Bhat, CSAID, Scientific Computing Systems and Services Division

FNAL Users Meeting

29 June 2023

Outline

- CSAID and the resources we provide to users
- Highlights from CSAID in the past year
- New plans for 2023-24 for CSAID
- The many ways to get access to a GPU

The Computational Science and AI Directorate

Last fall the Computing Sector split into the Information Technology Division (mostly what was Core Computing) and the Computational Science and AI Directorate (mostly what was Scientific Computing)

CSAID consists of three divisions and has extensive ties to the Lab AI Office

James Amundson
Associate Lab Director, Computational Science and AI Directorate

Keynote speaker

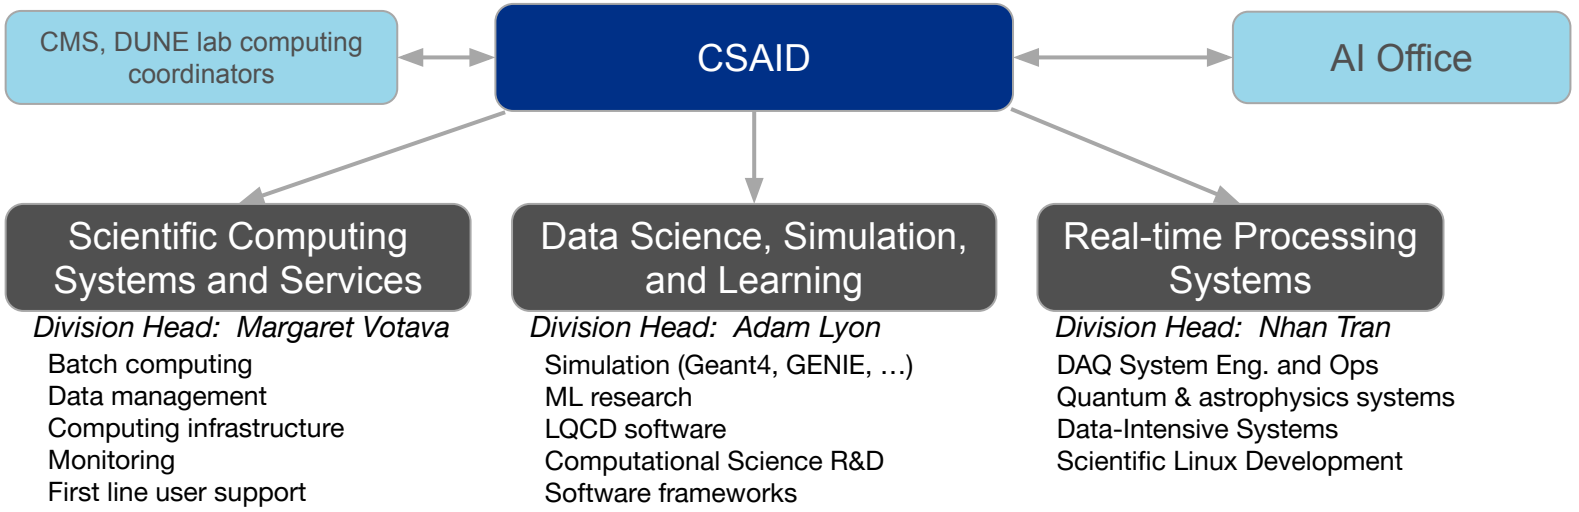

Division Head: Margaret Votava
- Batch computing
- Data management
- Computing infrastructure
- Monitoring
- First line user support

Division Head: Adam Lyon
- Simulation (Geant4, GENIE, …)
- ML research
- LQCD software
- Computational Science R&D
- Software frameworks

Division Head: Nhan Tran
- DAQ System Eng. and Ops
- Quantum & astrophysics systems
- Data-Intensive Systems
- Scientific Linux Development

All divisions have personnel who participate in experiment science activities

Some of the Resources We Provide at Fermilab

GPGrid Cluster

Data Storage and
Handling/Active Archive Facility

CMS Tier1 Center

Elastic Analysis Facility

Wilson Cluster-Institutional Cluster/Lattice QCD Cluster

HEPCloud

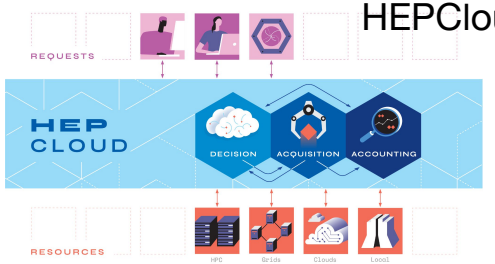

Highlights from the Scientific Computing Systems and Services Division

- Moved job submission authorization to tokens, via jobsub_lite (see next slide)
- Carried out experiment production workflows at HPC centers via HEPCloud
- Formed CREST committee to plan future of facility
- Started the migrations off of Scientific Linux 7
- Elastic Analysis Facility (EAF) in production with full 8x5 support (more on this later)

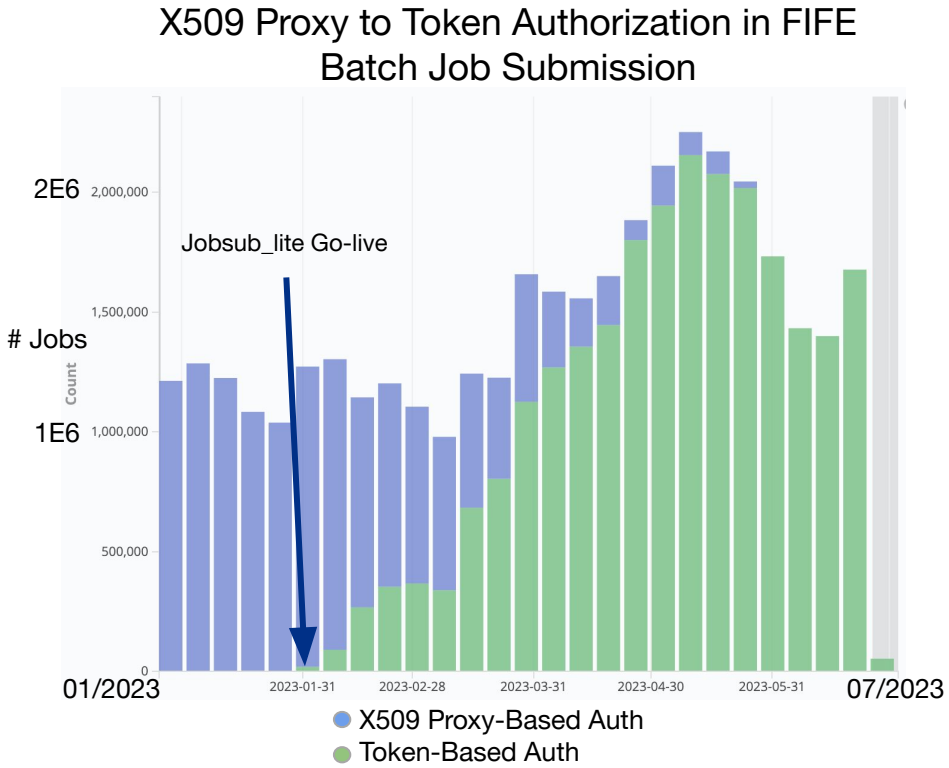

Fermilab

Transition to Token-Based Authorization

- This year, we rolled out major change to grid authorization: X509 Proxies → Tokens
- SciTokens (https://scitokens.org)/HTCSS IDTOKENs are form of JSON Web Token (JWT), which implements OAuth2 standard (think “Sign in with your Google Account”)
- Shift from “Who are you” (X509 Proxies) to “What can you do?” (SciTokens)
- Has allowed experiments to enforce more fine-grained access control to resources, especially storage
- Used/supported by all FIFE services: jobsub_lite, POMS, SAMWEB client, dCache
- Tokens used for internal authorization as well (e.g. GlideinWMS): Fine-grained access control within infrastructure too!

Transition to Token-Based Authorization (2)

- Before - many of these components used copies of the same X509 proxy.
- Now - disparate, appropriately-scoped tokens for different services

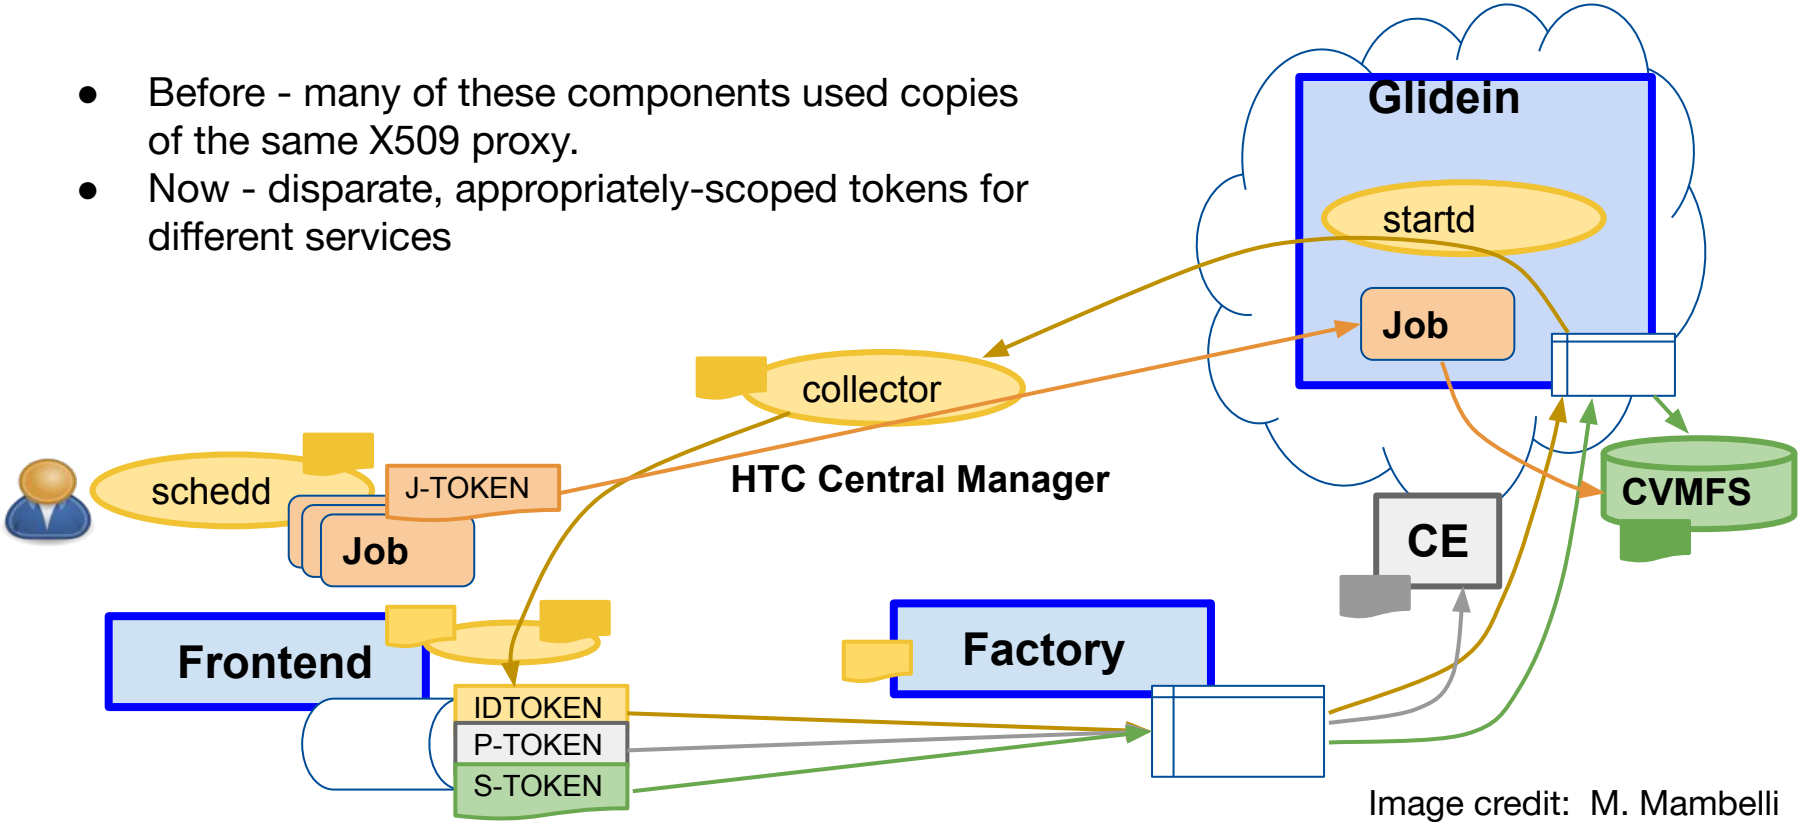

Image credit: M. Mambelli

Highlights from the Data Science, Simulation and Learning Division

Supporting experiments with Scientists who are computing experts and Computing Professionals

Cutting edge research and development

- AI for cosmology, CMS, neutrino experiments (eg, tracking)
- Computer science and physics publications
- Snowmass leadership roles and report contributors

Developing and supporting community tools

- GENIE, GEANT4, CMSSW, art, LArSoft, Marley, Root and more

Highlights from the Data Science, Simulation and Learning Division

Listening to the experiments for software needs

- DUNE framework requirements → R&D / LDRD on frameworks
- Hosted LArSoft Multithreading and Acceleration Workshop
  - Experiences, problems, solutions related to resource / throughput optimization and efficient high performance computing (HPC) / supercomputer utilization to lead common tool development
- Contributed to AI in Production Workshop
  - Working with experiments to improve toolkit integration
- Hosted Frameworks Workshop
  - Framework design concepts, art development, experiment requirements
  - Highlighted needs for HPC, improved documentation
- HPC for LArTPC reconstruction

Highlights from the Real-Time Processing Systems Division - Data-Intensive Systems Department

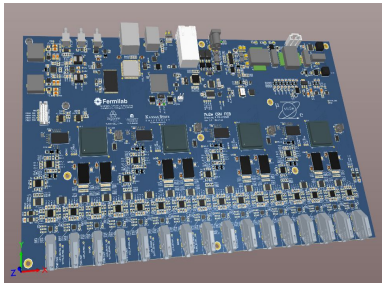

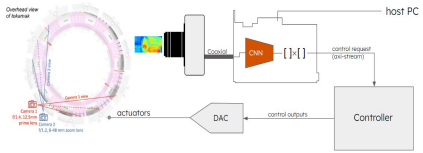

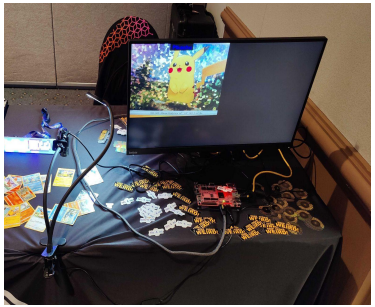

- Leading Accelerator Data Acquisition and Controls hardware modernization effort (ACORN)
- Developing readout hardware and firmware for multiple experiments (including DUNE, PIP-II, Mu2e, CMS)
- Supporting particle tracking for multiple Test Beam users.
- AI/ML: Leading low-latency ML in hardware package (hls4ml)
  - Leading research on ML architectures optimized for low-power FPGA applications.
  - Tech transfer for other sciences (fusion, microscopy, etc.)
  - Outreach to broader communities (DEFCON): real-time, FPGA-based, image identification demos
- Physics Research Equipment Pool (PREP)

Highlights from the Real-Time Processing Systems Division - DAQ Systems Engineering and Operations

- **DAQ Software Development**
  - DUNEDAQ
    - coordination roles in DAQ consortium, several major releases/milestones this year
  - ArtDAQ/OTSDAQ
    - development and support for mu2e, SBN, various test beam efforts
- **Providing online support** for running and upcoming experiments and test stands
- **DUNE**
  - DUNEDAQ and US Project management
  - ProtoDUNE-II, coldboxes, global test stands
- **Linux Distribution**
  - Support SL7 until End of Life (June 30, 2024)
  - Linux Steering

Highlights from the Real-Time Processing Systems Division - Quantum & Astrophysics Systems Department

- **The LTA (Low Threshold Architecture), supports all CCD and skipper-CCD instrumentation at Fermilab and collaborators.**
  - Dark matter experiments: DAMIC, SENSEI and Oscura 10Kg.
  - Space-LTA for future space mission.
  - Used for readout of quantum imaging with skippers and Coherent Neutrino Nucleus Scattering experiments.

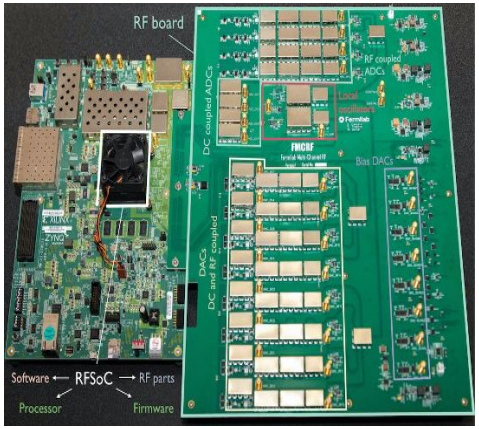

- **QICK (Quantum Instrumentation Control Kit):** A comprehensive, control and readout system for QIS (Quantum Information Science), including quantum computing, quantum networks, and quantum sensors
  - Adopted by and supporting a growing community of over 200 institutions (Labs, academia and industry).
  - Types of qubits: superconducting, AMO, NV-centers, trapped ions, spin.
  - Sensors, e.g. MKIDs, RF broadband, SNSPDs, quantum capacitor, etc.

New plans for 2023-24 from the various divisions

- Continue transition from proxies to tokens
- Help manage transition away from SL7
  - Monitoring industry activity
  - Steering Linux development in light of recent events
- Continue to investigate replacing NAS with CEPH
- Expanding GPU expertise for scientific computing
- Focusing support on DUNE & Mu2e following recent reviews
- Continue supporting experiments in DAQ/Real-time processing
- Hardware purchases (More on this later)

Highlights from AI Office

- Lots of interesting stuff going on in the AI office
  - See Tia Miceli's upcoming talk
- Let's say you need a GPU…

So you need a GPU…

- AI Practitioners have ready* access to the following resources
  - Elastic Analysis Facility (EAF)
  - Wilson Cluster, through experiments and/or dedicated projects
  - GPU clusters at remote sites via usual job submission tools through experiments
  - HPC centers such as NERSC through experiments and HEPCloud
  - Other HPC centers with individual/university/experiment allocations (CAN be integrated with HEPCloud!)
- Each type has slightly different strengths and weaknesses, along with different interfaces

"When you come to a fork in the road, take it." - L. P. "Yogi" Berra

*commercial clouds are also possible, though access is less ready (and no large common pool available)

Slide Credit: K. Herner

Elastic Analysis Facility

- Jupyter Hub deployment with general CPU and GPU-enabled notebooks available. Highly scalable, customizable, and replicable elsewhere
- GPUs are available through https://analytics-hub.fnal.gov (must be on-site or on VPN for now)
  - They are finite of course, but A100s now deployed
- Latest documentation is at https://eafjupyter.readthedocs.io/en/latest/. In particular, look at the server and notebook options section to see what's in each flavor (CVMFS also available)

Slide Credit: K. Herner

EAF (2)

- How do I get an account?
  - Anyone with a services account can log in, but follow your experiment's usual instructions
- Can I customize the environment and packages?
  - Yes. Several things are possible with preamble scripts, mamba, pip, spack etc. See the instructions about customizing: https://eafjupyter.readthedocs.io/en/latest/02_customization.html
- What storage is available?
  - 24 GiB work area in Ceph as home area, usual /nashome areas (FIFE/DUNE/Cosmic), usual LPC /uscms NFS areas also available
  - Streaming with xrootd also works to access larger storage elements
  - Just added 10 TB of scratch space - admins reserve right to delete anything older than 7 days if it gets full (*/scratch/7DayLifetime*)
  - https://eafjupyter.readthedocs.io/en/latest/01_storage.html

Slide Credit: K. Herner

GPUs via Wilson Cluster

- The Wilson Cluster is available to all of FNAL either through experiments/departments or specific projects (focused on a particular area; could be sub-project of an experiment, but do *not* have to be part of something larger)
- Available GPUs range from K80s (100) to A100s (4); base OS is SL7
- More information available at https://computing.fnal.gov/wilsoncluster/
  - There is a SNOW form to be added to a project (automatically added to base experiment projects)
- Rumors persist that queue times are interminable for the GPUs: generally not true

SLURM Partition	SLURM GPU Resource:Type:Numberof	Processor	Nodes	Total Threads	Memory/Core	Memory/Host	GPU
cpu_gce	No GPU	2.6GHz Dual CPU Eight Core Intel	108	16	8GB	128GB	None
cpu_gce_test	No GPU	2.6GHz Dual CPU Eight Core Intel	7	16	8GB	128GB	None
gpu_gce	gpu:k40:4	2.6GHz Dual CPU Eight Core Intel	4	16	8GB	128GB	4x NVIDIA Kepler K40
	gpu:p100nvlink:2	2.4GHz Dual CPU Fourteen Core Intel	1	56	2GB	112GB	2x NVIDIA Pascal P100 (w/NVLINK)
	gpu:p100:8	1.7GHz Dual CPU Eight Core Intel	1	16	48GB	768GB	8x NVIDIA P100
	gpu:v100:2	2.5GHz Dual CPU Twenty Core Intel	4	40	4.7GB	188GB	2x NVIDIA Volta V100
	gpu:a100:4	2.8GHz Dual CPU Thirty-Two Core EPYC 7543	1	64	8GB	512GB	4x NVIDIA Ampere A100-80 80GB HBM2e

cpu_gce | cpu_gce_test | gpu_gce | gpu_gce_ppc | gpu_ose | knl_gce | Project Usage

Refresh Rate=5 mins. Last updated on: 04-05-2023 1:05 AM

SLURM job status for partition gpu_gce

JobID	Job Name	User	Account	# of Nodes	Partition	QoS	List of Nodes -or- (Reason for job queued)	Priority	GPU Resources	Start Time	Requested Walltime (DAY-HH:MM:SS)	Time Remaining (DAY-HH:MM:SS)	Job State
408879	ecd_gpu	xinyuan	fwk	1	gpu_gce	regular	wcgpu04	67492	gres:gpu:1	2023-04-04T06:10:52	1-00:00:00	5:05:47	RUNNING
408880	ecd_gpu	xinyuan	fwk	1	gpu_gce	regular	wcgpu04	67491	gres:gpu:1	2023-04-04T06:10:52	1-00:00:00	5:05:47	RUNNING
408958	bash	Soon yung Jun	g4p	1	gpu_gce	opp	wcgpu05	339	gres:gpu:1	2023-04-04T17:06:55	8:00:00	1:50	RUNNING
407393	diffu_dataset3_mar21_hybrid_n400	oamram	cms	1	gpu_gce	opp	wcgpu03	44752	gres:gpu:1	2023-04-04T17:35:20	8:00:00	30:15	RUNNING
408966	train	calcuttj	dune	1	gpu_gce	opp	wcgpu06	19260	gres:gpu:v100:2	2023-04-05T00:17:21	8:00:00	7:12:16	RUNNING
407392	diffu_dataset3_mar21_hybrid_weight_n400	oamram	cms	1	gpu_gce	opp	wcgpu05	46830	gres:gpu:1	2023-04-05T00:31:22	8:00:00	7:46:17	RUNNING

Slide Credit: K. Herner

Fermilab

GPUs on External Clusters

- Several sites have GPUs available; accessible via usual job submission commands for IF/CF expts (separate for CMS)
  - **As of now, user is responsible for specifying appropriate container**
- Hardware ranges from a few years old up to A100s depending on the site. Contact FIFE for details.
- Several advantages to using these resources
  - Very little competition for them (i.e. no long queues a la NERSC)
  - No finite allocations!
- GPGrid will be getting GPU nodes; could consider more if demand (more later)
- NERSC Perlmutter (both CPU and GPU queues) available via HEPCloud for currently onboarded expts (CMS, DUNE, g-2, Microboone, Mu2e, Nova)
  - Contact your experiment's production group

Slide Credit: K. Herner

HEPCloud and other HPC centers

- It is also possible to integrate other HPC center or commercial cloud allocations into HEPCloud; contact HEPCloud team
  - Advantages include not having to learn separate infrastructures at every different site
- Some HPC centers offer testbed platforms for AI/ML apps; see ANL as one example
- GlideinWMS and HEPCloud are working to provide GPUs as a Service (GaaS)
  - GPUs will be shared between all jobs running on a node (fractional use possible)

Slide Credit: K. Herner

Future Hardware Additions at FNAL

- CSAID is planning to buy roughly $500k of GPU workers equipped with A100s
  - Form factor not yet decided (2 or 4 GPUs per node)
  - Expect this will provide up to a few dozen GPUs
  - Looked at Hopper cards but decided better bang for the buck with A100s given our typical models
- Machines will have NVLink, Infiniband, ethernet
- Distribution among the various clusters not decided yet. Depends a bit on needs, demand, and stakeholder requirements
  - *One could imagine* shifting machines between clusters as demand ebbs and flows between different areas (requires sufficient lead time)

Slide Credit: K. Herner

Summary

- Lots of exciting activity going on in Scientific Computing
- Wide variety of resources now available from interactive containers to grid slots to large blocks on HPC centers
- Different tools for different jobs: don’t hesitate to contact division personnel even if it’s just to consult about what’s the best option
- How? Contact your experiment liaison/spokesperson or open an “Ask a Question” ticket to Distributed Computing Support in ServiceNow

WE WANT TO HEAR FROM YOU!!!

Thank you!

Thanks to all who contributed to this talk: J. Boyd, G. Cancelo, B. Coimbra, V. Di Benedetto, J. Eisch, L. Goodenough, K. Herner, B. King, A. Lyon, M. Mambelli, R. Rechenmacher, K. Retzke, E. Snider, N. Tran, M. Votava

Backup Slides

Data-Intensive Systems Department

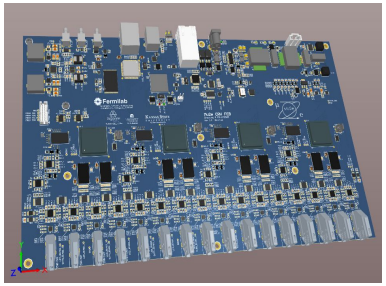

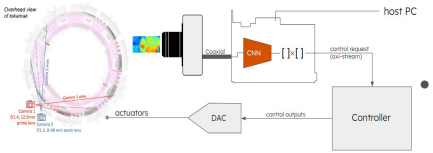

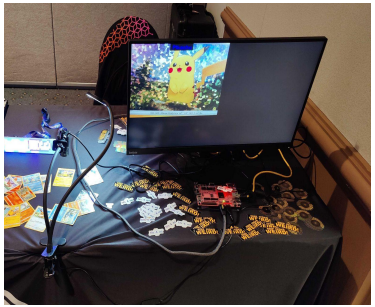

- Accelerator Controls Operations Research Network (ACORN):
  - Leading Accelerator Data Acquisition and Controls hardware modernization effort.
- DUNE:
  - Photon Detection System readout hardware: DAPHNE for FD-1 and Analog and Digital signal over fiber R&D for FD-2.
- PIP-II:
  - High-speed digitization and low-latency processing system for machine protection.
- Mu2e:
  - Data acquisition system hardware, firmware and software development, teststands and installation.
- CMS:
  - Outer Tracker production testing and correlator trigger firmware.
- Test Beam:
  - Particle tracking support for multiple users.
- AI/ML
  - Presented a real-time, FPGA-based, image identification demo at DEFCON in Las Vegas
  - Major contributions to HLS4ML low-latency ML software.
  - Leading research on ML architectures optimized for low-power FPGA applications.
  - Enabled real-time inference on frame grabber FPGA for optical instability tracking and suppression in Fusion experiments.
- Physics Research Equipment Pool (PREP)

Slide Credit: J. Eisch

More details on using GPUs on HTC clusters

- Simply add `--lines='+RequestGPUs=1'` to your submission if using jobsub.
- **As of now, user is responsible for specifying appropriate container** (use the --singularity-image option and make sure the image is available to the worker, usually via CVMFS). One could imagine the division maintaining a container including CUDA, IF there is demand for it. Must also be sure data transfer in/out is possible

HEPCloud as a gateway to HPC and LCF

- NERSC Perlmutter (both CPU and GPU queues) available via HEPCloud for currently onboarded expts (CMS, DUNE, g-2, Microboone, Mu2e, Nova)
  - Contact your experiment's production group
- Demonstrations of Inference server setups
- No true "co-scheduling" of GPU servers w/ other workflows on non-GPU queues yet
  - Ad-hoc or requires partitioning between workflow components
- No support for MPI workflows currently within HEPCloud
  - Need to rely on local slurm batch queue