



Scientific Computing at Fermilab

Kyle J. Knoepfel

12 June 2019

52nd Fermilab Users Meeting

Scientific Computing Division at Fermilab

- **What we do**

- Meet scientific-computing needs for DAQ, production teams and analyzers for all Fermilab-supported experiments
 - Software for data processing
 - Interactive and build environments
 - Large-scale batch system processing
 - Data storage and access
 - Simulations and physics modeling
 - DAQ systems, etc.



```
ngSwitchMatchController {  
  constructor(attrs, element, attr, ngSwitchController) {  
    this.$inject = ['attrs', 'ngSwitchController'];  
    this.$scope = $scope;  
    this.$element = element;  
    this.$attrs = attrs;  
    this.$ctrl = ngSwitchController;  
  }  
  $onInit() {  
    this.$ctrl.$on('ngSwitchMatchAction', function(value) {  
      // ...  
    });  
  }  
}
```



- **Who we are**

- We are computing professionals (many with particle-physics PhDs), physicists and computing science researchers, postdocs, engineers, etc.
- Each experiment has a CS liaison, who is a bridge between the experiment and SCD

Ultimately, we are here to help you achieve your physics goals.

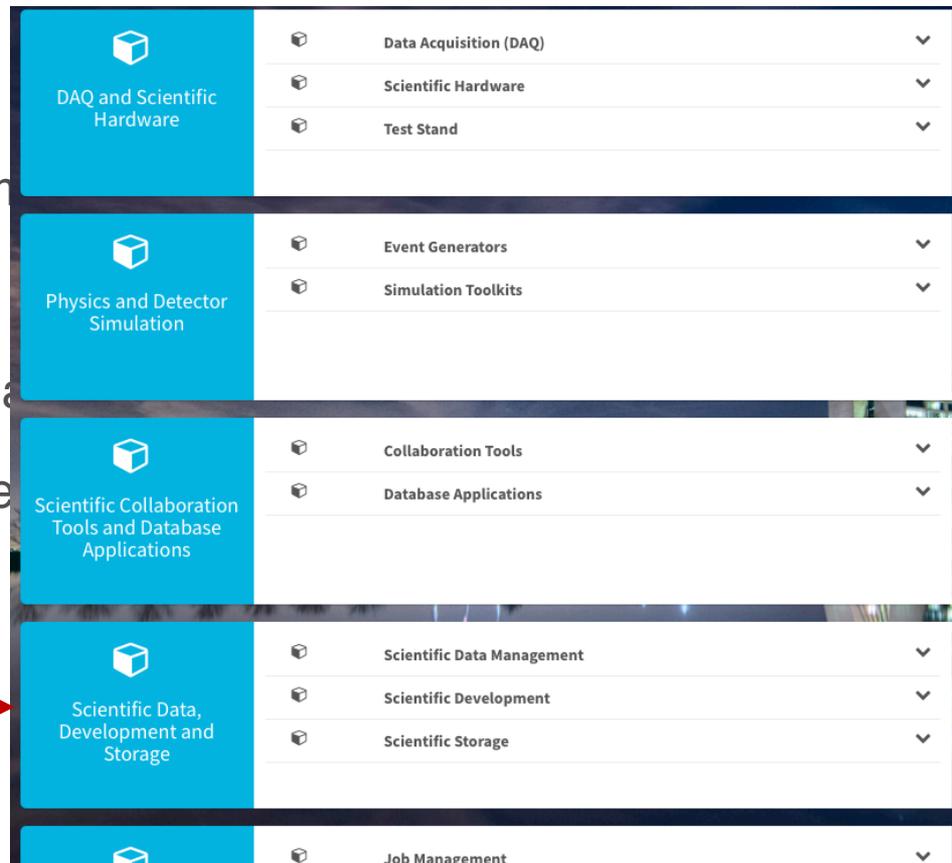
What do you as a user care about?

- Will the experiment software work?
- Will the community software we use be maintained and supported?
- How do I submit and monitor my jobs?
- Where do I store my data, and how do I access it?
- If something is wrong, how do I get it fixed?

What do you as a user care about?

- Will the experiment software work?
- Will the community software we use be maintained?
- How do I submit and monitor my jobs?
- Where do I store my data, and how do I access it?
- If something is wrong, how do I get it fixed?

You can always consult the service desk at servicedesk.fnal.gov



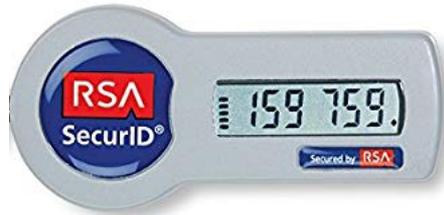
DAQ and Scientific Hardware	Data Acquisition (DAQ)	▼
	Scientific Hardware	▼
	Test Stand	▼
Physics and Detector Simulation	Event Generators	▼
	Simulation Toolkits	▼
Scientific Collaboration Tools and Database Applications	Collaboration Tools	▼
	Database Applications	▼
Scientific Data, Development and Storage	Scientific Data Management	▼
	Scientific Development	▼
	Scientific Storage	▼
	Job Management	▼

Multifactor authentication (MFA) at Fermilab

Public service announcement

New threats to cybersecurity require us to **expand use of MFA**

- By July 30 **VPN access** will require use of **RSA** (hard or soft token) or **YubiKey**
- Non-employees will typically receive RSA (hard or soft token)
- Soft tokens can be used on phones and other mobile devices
- YubiKeys are typically used only by lab employees who access financial or HR data



Multifactor authentication (MFA) at Fermilab

Public service announcement

- **Planned rollout:**

- **Visit the Service Desk (WH ground floor, north side across from credit union) to pick up your RSA token during the User's Meeting**

OR

- Service Desk will distribute (150 YubiKeys/tokens per week)
- VPN users have received or will receive email notice from Service Desk when it is their week to pick up/request their token.

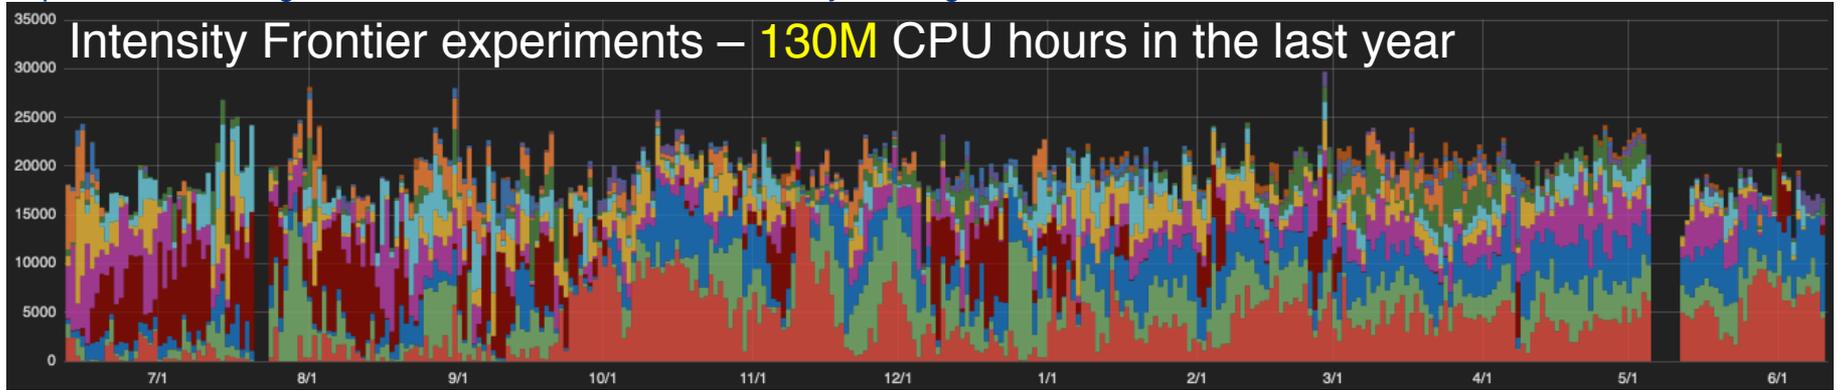
Remote users: The above email notice will explain how to obtain your device by post (if hard token) or by email (if soft token).

Questions: mfa-questions@fnal.gov

How many jobs have you submitted?

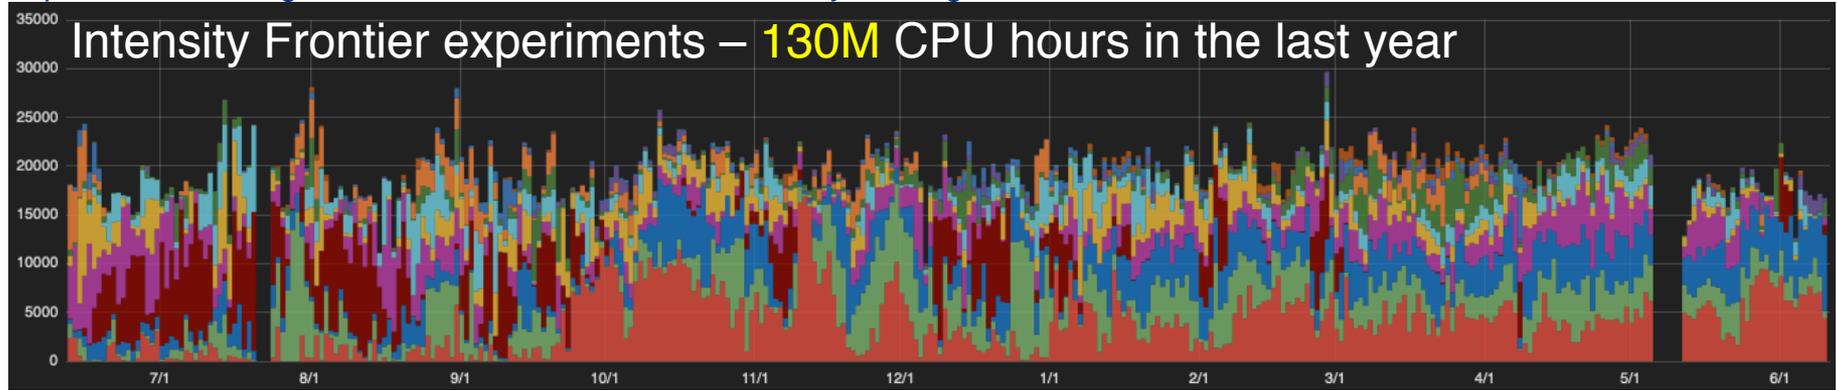
How many jobs have you submitted?

<https://fifemon.fnal.gov/monitor/d/000000180/scd-summary-fife?orgId=1>

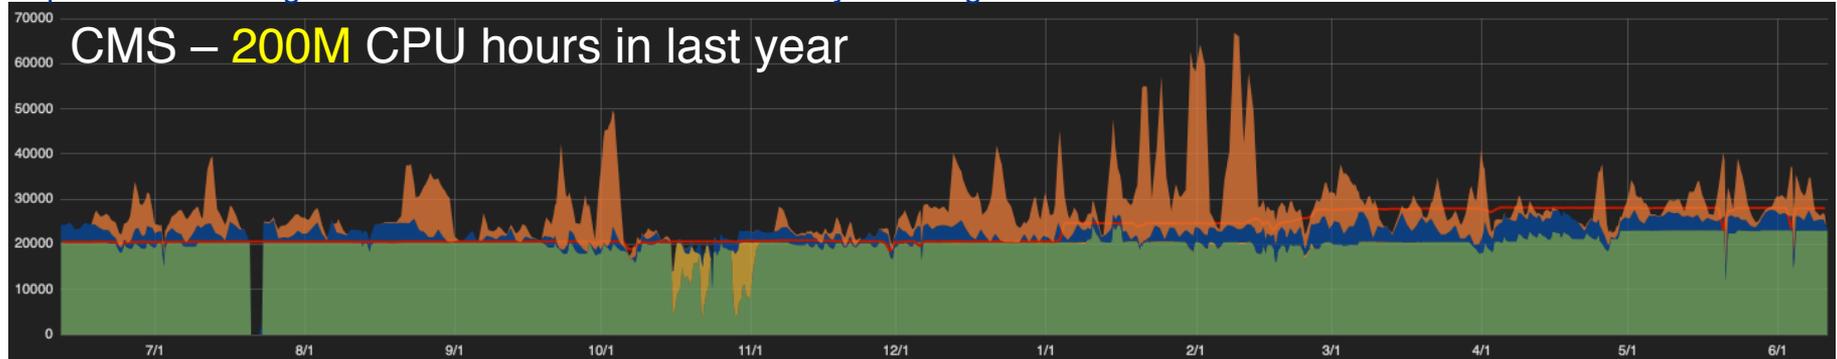


How many jobs have you submitted?

<https://fifemon.fnal.gov/monitor/d/000000180/scd-summary-fife?orgId=1>

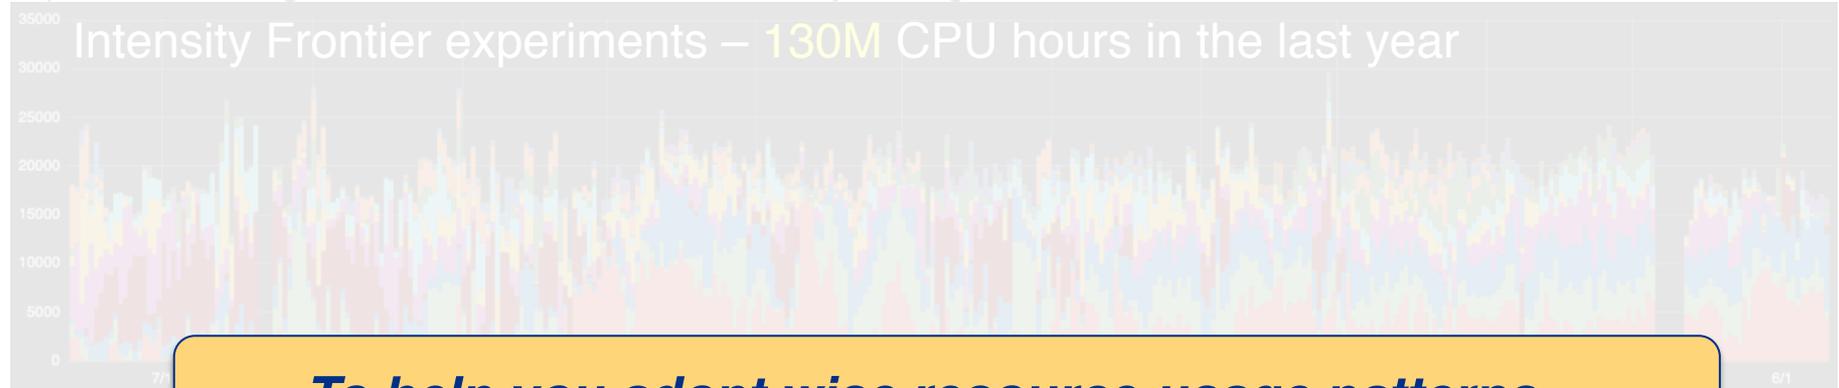


<https://fifemon.fnal.gov/monitor/d/000000130/scd-summary-cms?orgId=1>



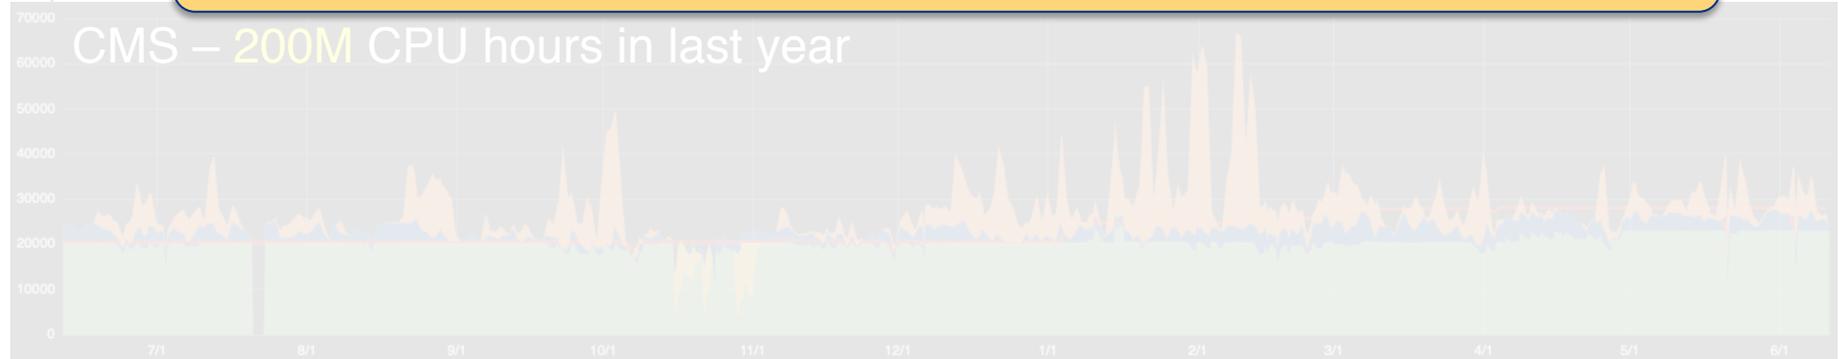
How many jobs have you submitted?

<https://fifemon.fnal.gov/monitor/d/000000180/scd-summary-fife?orgId=1>



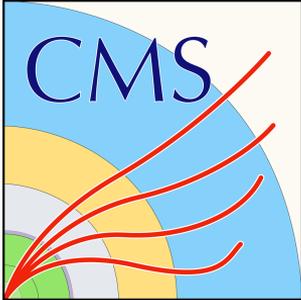
To help you adopt wise resource-usage patterns...

[https://fifemon](https://fifemon.fnal.gov/monitor/d/000000180/scd-summary-fife?orgId=1)



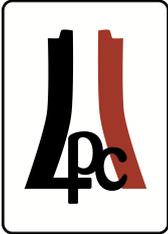
Software training

Software training



CMSDAS

10 years!



Software training



CMSDAS



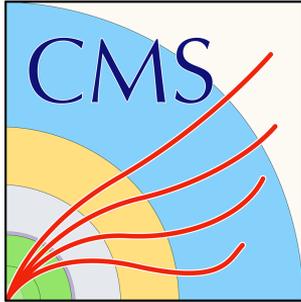
C++/STL (May 20-24, 2019)

Instructor: Glenn Downing (UT Austin)

Over 50 attendees!

- Fermilab is re-energizing the software training that been provided in the past

Software training



CMSDAS



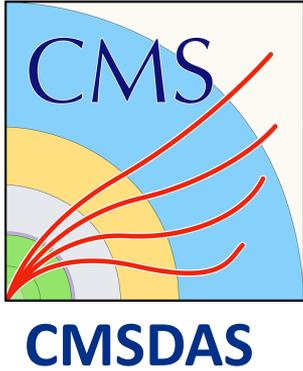
- Fermilab is re-energizing the software training that been provided in the past

Courses available

<https://hr.fnal.gov/training/curriculum/>

- **C++ / Standard Template Library** (5 days) – Waitlist
- **WordPress Training** – Waitlist
- **Python Programming Basics** (3 days) – Waitlist
- **Python Programming Advanced** (3 days) – Waitlist (*Previous Linux experience not required for either course.*)

Software training

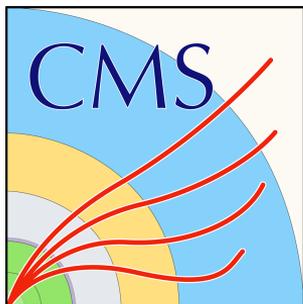
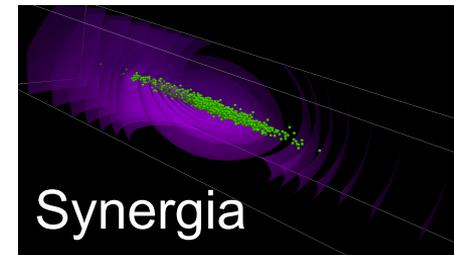


- Fermilab is re-energizing the software training that been provided in the past
- Why training?
 - Prudent use of resources necessitates well-designed, well-tested, efficient code
 - Methods and languages evolve
 - C++17 is the current standard (it looks a lot different wrt C++ from 20 years ago)
 - Python 2 end-of-life is January 1, 2020

Some projects supported by SCD



Pythia



CMSSW



artdaq
otsdaq

Enstore

dCache.org 

Common themes

- We struggle with the following tensions:
 - Current experiments must be supported
 - New experiments often require new ideas/R&D
 - We need to keep up with new architectures and technologies



Common themes

- We struggle with the following tensions:
 - Current experiments must be supported
 - New experiments often require new ideas/R&D
 - We need to keep up with new architectures and technologies
- **How do we manage these tensions and the additional work they bring while meeting your needs?**



Common themes

- We struggle with the following tensions:
 - Current experiments must be supported
 - New experiments often require new ideas/R&D
 - We need to keep up with new architectures and technologies

- How do we manage these tensions and the additional work they bring while meeting your needs?
 - **Use tools supported by a broader community**

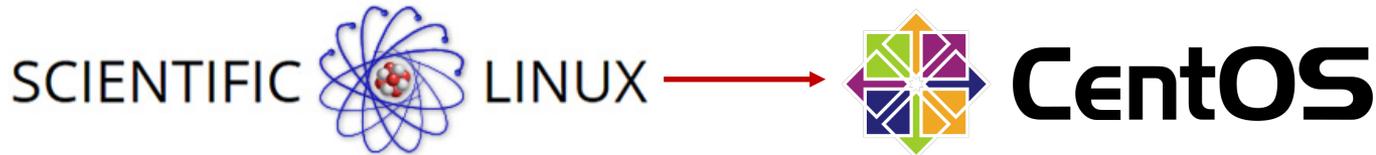


Common themes

- We struggle with the following tensions:
 - Current experiments must be supported
 - New experiments often require new ideas/R&D
 - We need to keep up with new architectures and technologies
- How do we manage these tensions and the additional work they bring while meeting your needs?
 - **Use tools supported by a broader community**



Scientific Linux has served us well; after SL7, the lab will support CentOS.



Common themes

- We struggle with the following tensions:
 - Current experiments must be supported
 - New experiments often require new ideas/R&D
 - We need to keep up with new architectures and technologies
- How do we manage these tensions and the additional work they bring while meeting your needs?
 - **Use tools supported by a broader community**



Improving the code-development and -distribution system

- Some Fermilab-supported code being moved from Redmine to GitHub.
- Replacing decades-old, Fermilab-based UPS.

Common themes

- We struggle with the following tensions:
 - Current experiments must be supported
 - New experiments often require new ideas/R&D
 - We need to keep up with new architectures and technologies

- How do we manage these tensions and the additional work they bring while meeting your needs?
 - Use tools supported by a broader community
 - **Engage communities with similar challenges to pursue common solutions...**

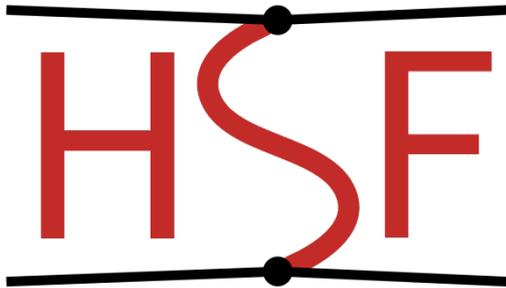


Interactions within HEP



- **HEP Software Foundation**
 - International collaboration becoming the “Apache foundation” of HEP
 - Large CERN presence that welcomes Fermilab’s collaboration (since the beginning)
- **Its goals are:**
 - sharing expertise
 - raising awareness of existing software and solutions
 - catalyzing new common projects
 - promoting commonality and collaboration in new developments to make the most of limited resources
 - aiding developers and users in creating, discovering, using and sustaining common software
 - supporting career development for software and computing specialists

Interactions within HEP



HEP Software Foundation

- **Companion white papers**

- Conditions Data
- Data Processing Frameworks
- Data Organization, Management and Access
- Visualization
- Data and Software Preservation
- Training, Staffing and Careers
- Data Analysis and Interpretation
- Detector Simulation
- Software deployment

Community White Paper – arXiv:1712.06982 [physics.comp-ph]

A Roadmap for HEP Software and Computing R&D for the 2020s

HEP Software Foundation¹

ABSTRACT: Particle physics has an ambitious and broad experimental programme for the coming decades. This programme requires large investments in detector hardware, either to build new facilities and experiments, or to upgrade existing ones. Similarly, it requires commensurate investment in the R&D of software to acquire, manage, process, and analyse the sheer amounts of data to be recorded. In planning for the HL-LHC in particular, it is critical that all of the collaborating stakeholders agree on the software goals and priorities, and that the efforts complement each other. In this spirit, this white paper describes the R&D activities required to prepare for this software upgrade.

Scientific computing R&D at Fermilab

- Largely geared toward running on HPC and heterogeneous systems
 - Utilizing vectorization where possible
 - Offloading to GPUs
 - Rethinking I/O
 - Investigating multiple I/O formats
 - Exploring *non*-file formats like object stores
 - How far can/should multi-threading be pushed?
 - Exploring the balance between multi-threaded and multi-process jobs
 - etc.
- How do these new concepts influence our computing paradigms?



HPC use in HEP

- More and more data processing is happening on HPC systems.
 - Experiments' computing needs are getting more intensive
 - Physics problems of interest are usually naturally parallel
 - Significant DOE funding has gone into supporting HPC systems

HPC use in HEP

- More and more data processing is happening on HPC systems.
 - Experiments' computing needs are getting more intensive
 - Physics problems of interest are usually naturally parallel
 - Significant DOE funding has gone into supporting HPC systems

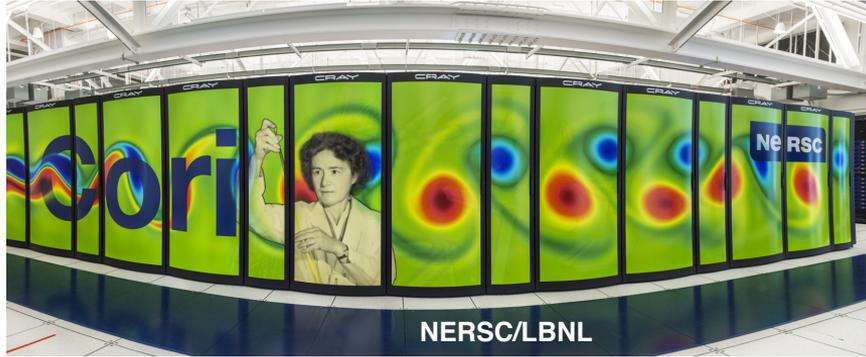
Center for Computational Excellence

HEP-CCE

Argonne, Brookhaven, Fermilab, Lawrence Berkeley, SLAC

“The HEP-CCE is a cross-cutting initiative to promote excellence in high performance computing (HPC) including data-intensive applications, scientific simulations, and data movement and storage.”

HPC use in HEP



Fermilab computing experts bolster NOvA evidence, 1 million cores consumed

July 3, 2018 | [Marcia Teckenbrock](#)



How do you arrive at the physical laws of the universe when you're given experimental data on a renegade particle that interacts so rarely with matter, it can cruise through light-years of lead? You call on the power of advanced computing.

The NOvA neutrino experiment, in collaboration with the Department of Energy's Scientific Discovery through Advanced Computing (SciDAC-4) program and the HEPCloud program at DOE's Fermi National Accelerator Laboratory, was able to

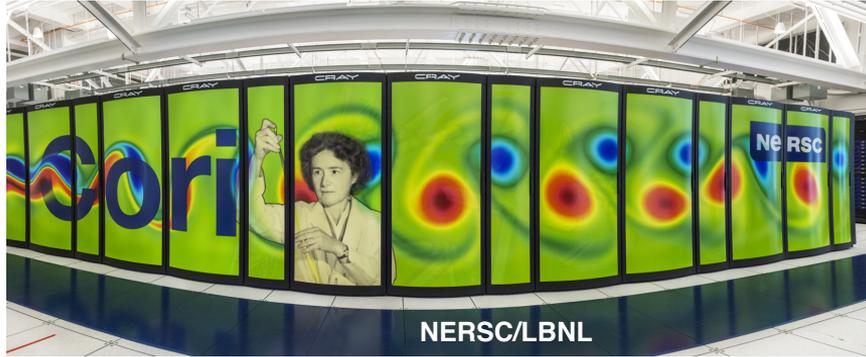
These machines are not your typical CPU-style computers.

More systems coming online in the next few years with architectures tailored for AI

- Figuring out how to use them for HEP is a challenge.



HPC use in HEP



Fermilab computing experts bolster NOvA evidence, 1 million cores consumed

July 3, 2018 | Marcia Teckenbrock

[Share](#) [Tweet](#) [Email](#)

How do you arrive at the physical laws of the universe when you're given experimental data on a renegade particle that interacts so rarely with matter, it can cruise through light-years of lead? You call on the power of advanced computing.

The NOvA neutrino experiment, in collaboration with the Department of Energy's Scientific Discovery through Advanced Computing (SciDAC-4) program and the HEPCloud program at DOE's Fermi National Accelerator Laboratory, was able to

These machines are not your typical CPU-style computers.

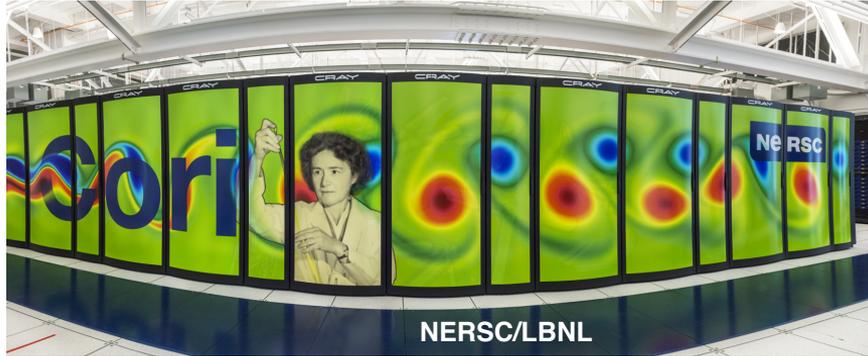
More systems coming online at Fermilab, including the Aurora supercomputer at ANL.

- Figuring out how to use these machines

How can my experiment use these types of machines?



HPC use in HEP



Fermilab computing experts bolster NOvA evidence, 1 million cores consumed

July 3, 2018 | Marcia Teckenbrock

[Share](#) [Tweet](#) [Email](#)

How do you arrive at the physical laws of the universe when you're given experimental data on a renegade particle that interacts so rarely with matter, it can only be seen on the power of advanced computing.

Scientific Discovery
HEPCloud program
the largest-scale an

These machines are not your typical CPU-style computers.

More systems coming...

- Figuring out how...

How can my experiment use these types of machines?



ORNL (already online)



ANL



HEPCloud

“A scientific gateway to resources beyond local worker nodes and grids”

- **Onboarded** experiments have access to resources that expand **dynamically**, based on the jobs’ requirements.
- Expanded resources include:
 - HPC (NERSC)
 - Amazon Elastic Compute Cloud
 - Google Cloud Platform
- Designed to mesh cleanly with current job-submitting tools
 - Minimal user-facing changes required to enable HEPCloud for onboarded experiments
- HEPCloud is authorized for experiments’ production jobs
 - **Ask your experiment about their intentions for using HEPCloud.**



Takeaways

- We are tackling HEP's computing challenges, but not in isolation:
 - We are interacting with a broad array of HEP, computing, and scientific communities.
 - We continue to grow our common software and commodity hardware solutions.
- Fermilab's Scientific Computing Division exists to support **you**.
 - If you have technical questions or problems, file a service-desk ticket. 😊
servicedesk.fnal.gov
 - If you have strategic questions or concerns, talk to us.

Takeaways

- We are tackling HEP's computing challenges, but not in isolation:
 - We are interacting with a broad array of HEP, computing, and scientific communities.
 - We continue to grow our common software and commodity hardware solutions.
- Fermilab's Scientific Computing Division exists to support **you**.
 - If you have technical questions or problems, file a service-desk ticket. 😊
servicedesk.fnal.gov
 - If you have strategic questions or concerns, talk to us.

Thank you.