Data, Analytics and AI Services for Science

Wahid Bhimji

Data Day 2022

NERSC has a rich data ecosystem!

globus online
jupyter

data transfer and access

PyTorch
scikit learn

machine learning

netCDF
mongoDB®
MySQL™

data management

julia

data analytics

visit
ParaView
Parallel Visualization Application

visualization

SHIFTER
Spin

containers

GNUparallel
Parsl
papermill
FireWorks

workflows

Data, Analytics and AI Services for Science

Modified from WB et. al O'Reilly Blog *(2016)*

Scientific Instruments

Scientist Interaction

Data transfer
- Globus
- GridFTP
- Xrootd

Portals
- Newt/SFAPI
- Django
- Jupyter

Analytics
- Python
- ROOT
- R
- Matlab

File Format
- HDF5
- NetCDF
- ROOT

Storage
- Lustre
- GPFS
- DataWarp
- HPSS

Workflow
- Fireworks
- Taskfarmer
- Dask

Processing
- SLURM realtime queues
- Spark
- Shifter

Machine Learning
- TensorFlow
- PyTorch
- Scikit-Learn

Databases
- MongoDB
- PostgreSQL
- MySQL

Edge Services
- Spin

Visualization
- Visit
- ParaView
- Matplotlib

BACK TO THE FUTURE

NERSC

Powering Scientific D

HOME ABOUT SCIENCE SYSTEMS FOR USERS NEWS R & D EV

FOR USERS

- Getting Help
- NERSC Code of Conduct
- Live Status
- Getting Started
- Accounts & Allocations
- Documentation

Home » For Users » Training & Tutorials » D

DATA DAY 2016

August 22nd - 23rd, 2016

Agenda Overview

Monday, August 22: Data Day

Remote connection available via zoom.

Some notebooks that will be shown are available at **NERSC's data-day examples** git repo.

Time	Topic	Presenter(s)	Room
	Data talks		Bldg. 50 Aud
8:30 am	Welcome (video, slides)	Katie Antypas, Head, NERSC Data Dept	Bldg. 50 Aud
8.45	Intro to Machine Learning (video, slides)	Prabhat, NERSC Data and Analytics Services Group Lead	Bldg. 50 Aud
9.15	Machine Learning tutorial (video, slides)	Evan Rac	
9:30	Science with Machine Learning (video, slides)	Marcus S	
9.45	**Break**		
10:10	Python Tutorial (video, slides)	Rollin Th	
10:40	Science with Python (video, slides)	Ben Bow	
11:10	Spark tutorial (video, slides)	Lisa Gerh	
11:40	Science with Spark (video, slides)	Zhong W Genome	
12.10 - 1.30	**Lunch and poster preview**	Lunch wi registere	
1:45	Visualization tutorial and discussion (video, slides)	Annette C	
2:30	Burst Buffer Tutorial (video, slides)	Debbie Bard, NERSC	Bldg. 50 Aud
3:00	Science with the Burst Buffer	Andrey Ovsyannikov	Bldg. 50

Available Tools

Deep Learning Frameworks

- **Theano** - flexibility, not for beginners (good for research)
- **Keras / Lasagne** - Theano-based but higher-level for ease of use
- **TensorFlow** - ease of use and flexibility, large, growing community, some *multi-node support*
- **Caffe** - high performance (IntelCaffe with performance highly optimised for KNL), *multinode (no programming necessary)*

General Machine Learning:

- **Scikit-Learn -** great for non-image based machine learning, easy to use, support for wide range of algorithms
- **Spark -** *multinode*, great for data parallel, relatively easy to use, support for only a

TensorFlow

theano

Spark

scikit learn

Many data services are now ubiquitous - others rapidly growing

Jupyter Usage at NERSC

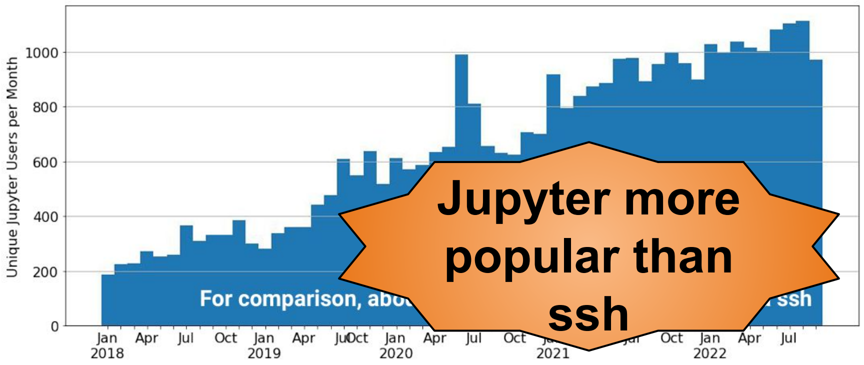

Jupyter more popular than ssh

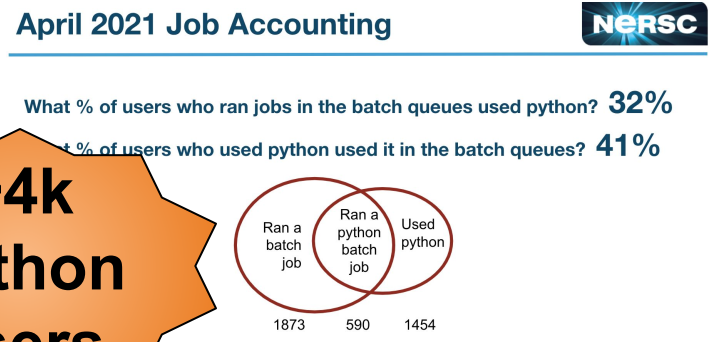

~4k python users

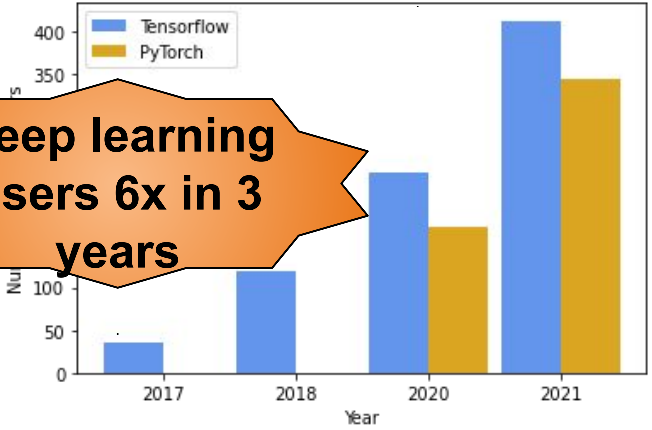

Deep learning users 6x in 3 years

Top 500 HPL run in shifter container

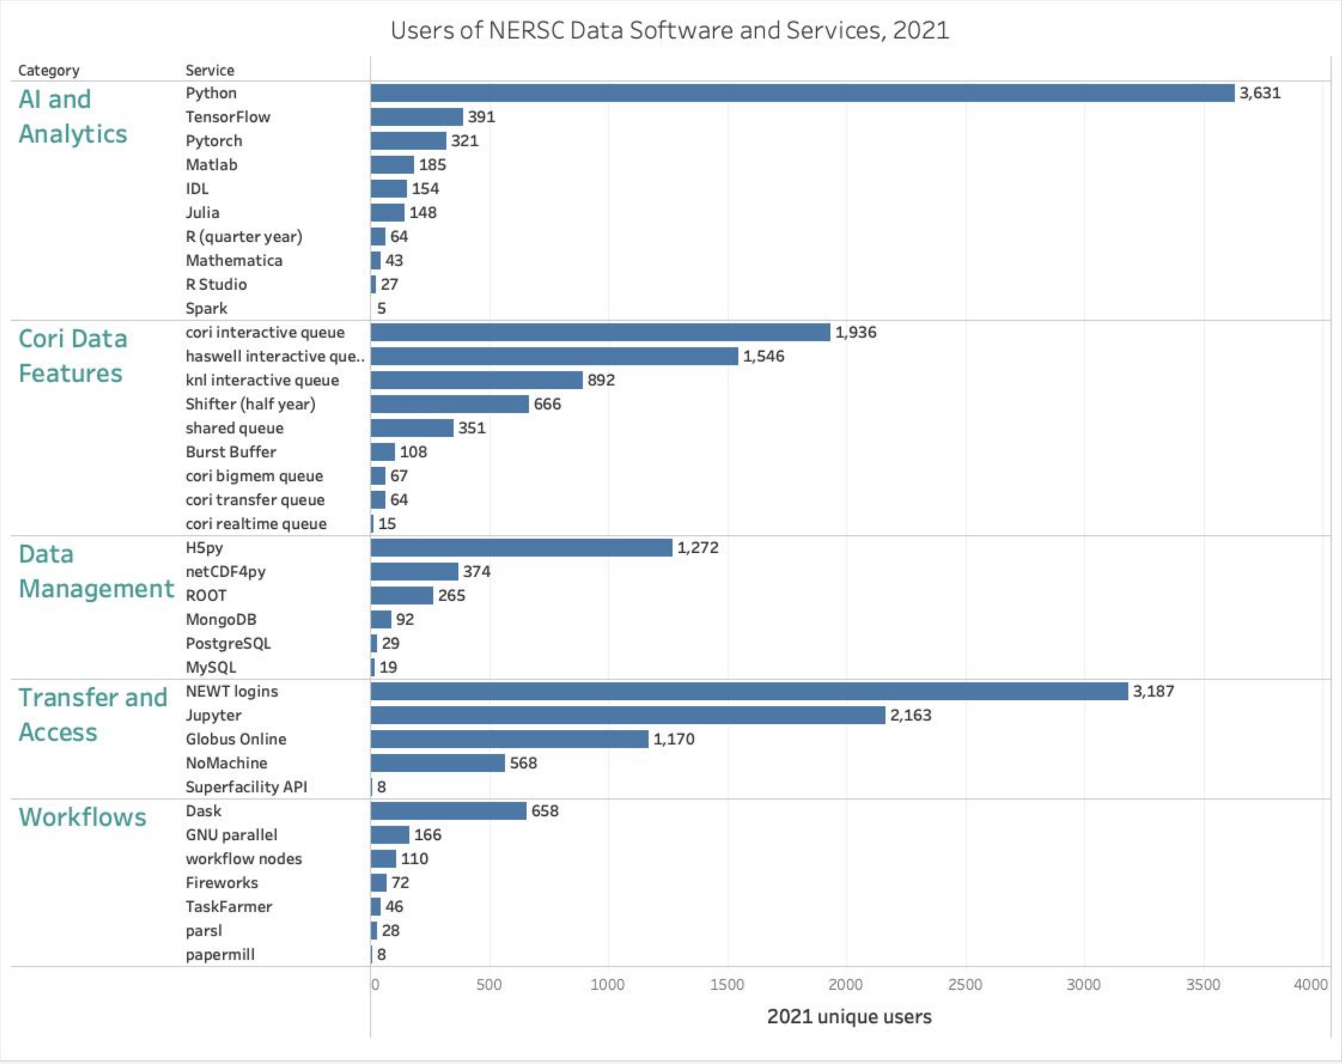

Users of NERSC Data Software and Services, 2021
Category
Service
AI and Analytics
Python
3,631
TensorFlow
391
Pytorch
321
Matlab
185
IDL
154
Julia
148
R (quarter year)
64
Mathematica
43
R Studio
27
Spark
5
Cori Data Features
cori interactive queue
1,936
haswell interactive que..
1,546
knl interactive queue
892
Shifter (half year)
666
shared queue
351
Burst Buffer
108
cori bigmem queue
67
cori transfer queue
64
cori realtime queue
15
Data Management
H5py
1,272
netCDF4py
374
ROOT
265
MongoDB
92
PostgreSQL
29
MySQL
19
Transfer and Access
NEWT logins
3,187
Jupyter
2,163
Globus Online
1,170
NoMachine
568
Superfacility API
8
Workflows
Dask
658
GNU parallel
166
workflow nodes
110
Fireworks
72
TaskFarmer
46
parsl
28
papermill
8
0
500
1000
1500
2000
2500
3000
3500
4000
2021 unique users

Perlmutter a Data and AI supercomputer

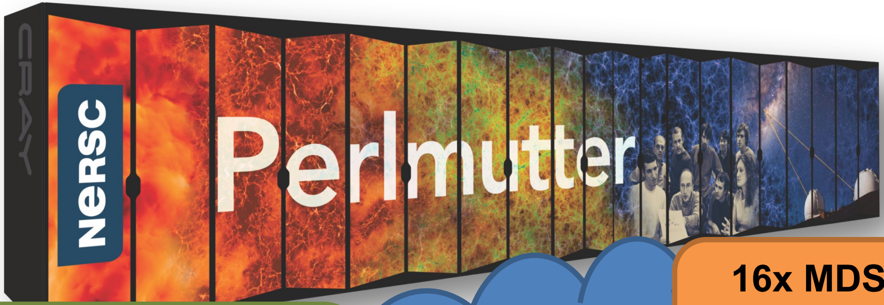

1,536 GPU nodes
1x AMD Epyc 7763
4x NVIDIA A100
4x Slingshot NICs

3,072 CPU nodes
2x AMD Epyc 7763
1x Slingshot NIC

Slingshot
200 Gb/s
2-level dragonfly

16x MDS + 274 OSS
1x AMD Epyc 7502P
2x Slingshot NICs
24x 15.36 TB NVMe

24x Gateway nodes
2x Slingshot NICs
2x 200G HCAs

2x Arista 7804 routers
400 Gb/s/port
> 10 Tb/s routing

SAN

SAN

Image: G. Lockwood

WAN

Data services impact entire workflow

From Lockwood, WB, … NERSC-9 Technical Design Review 2018

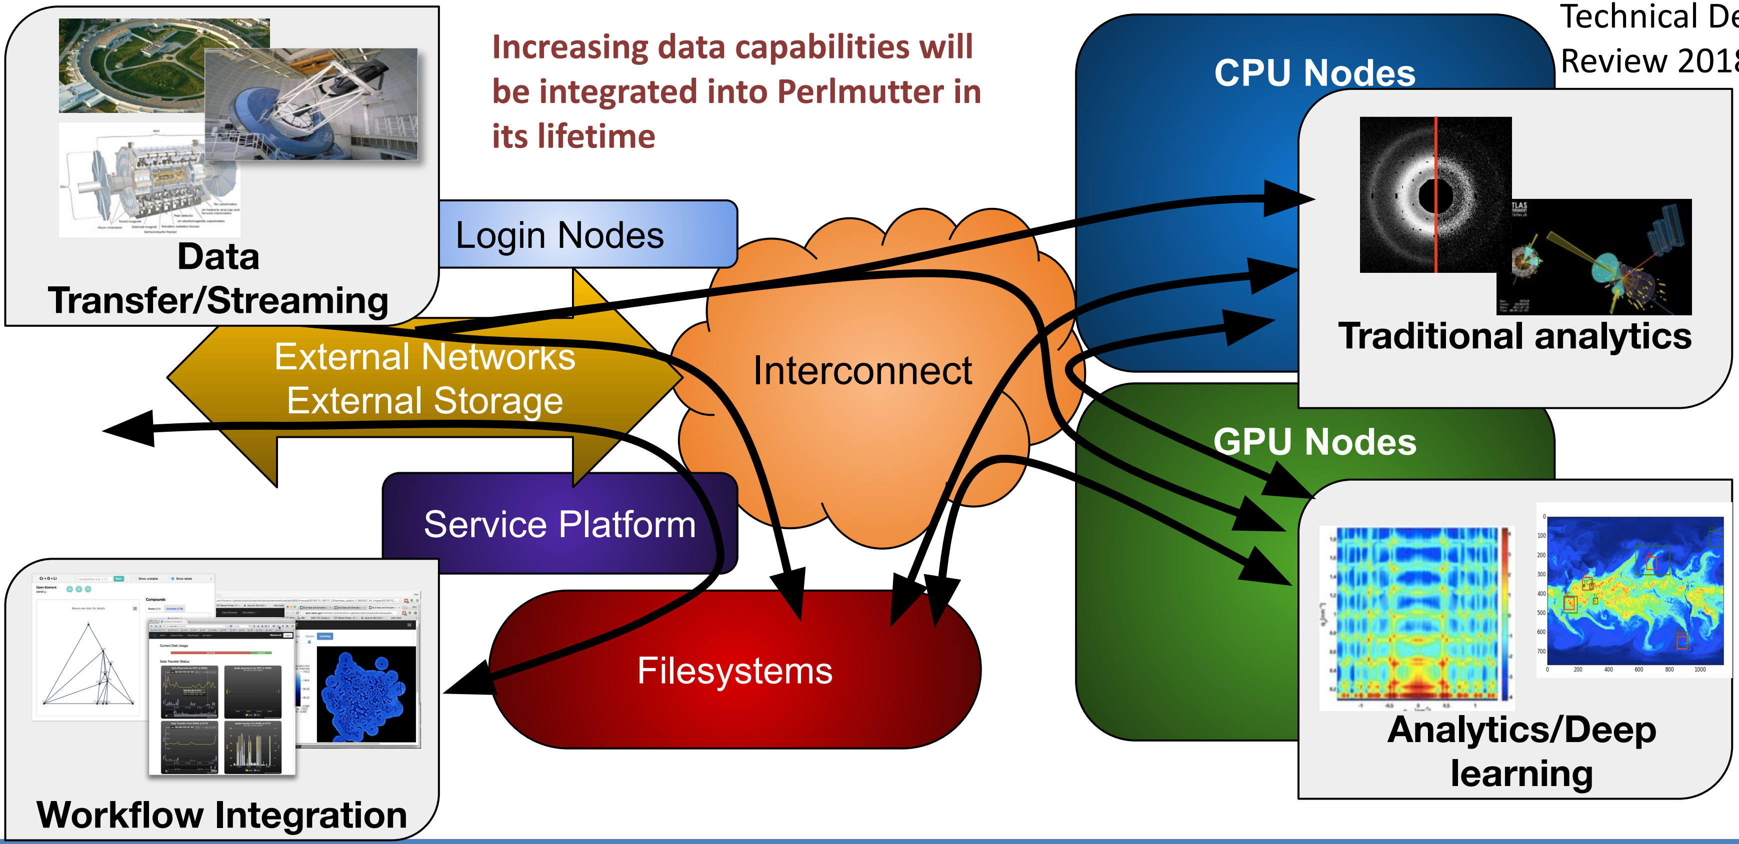

NERSC-10 Architecture: Designed to support complex simulation and data analysis workflows at high performance

NERSC-10 will provide on-demand, dynamically composable, and resilient workflows across heterogeneous elements within NERSC and extending to the edge of experimental facilities and other user endpoints

Complexity and heterogeneity managed using complementary technologies

- **Programmable infrastructure**: avoid downfalls of one-size-fits-all, monolithic architecture
- **AI and automation**: sensible selection of default behaviours to reduce complexity for users

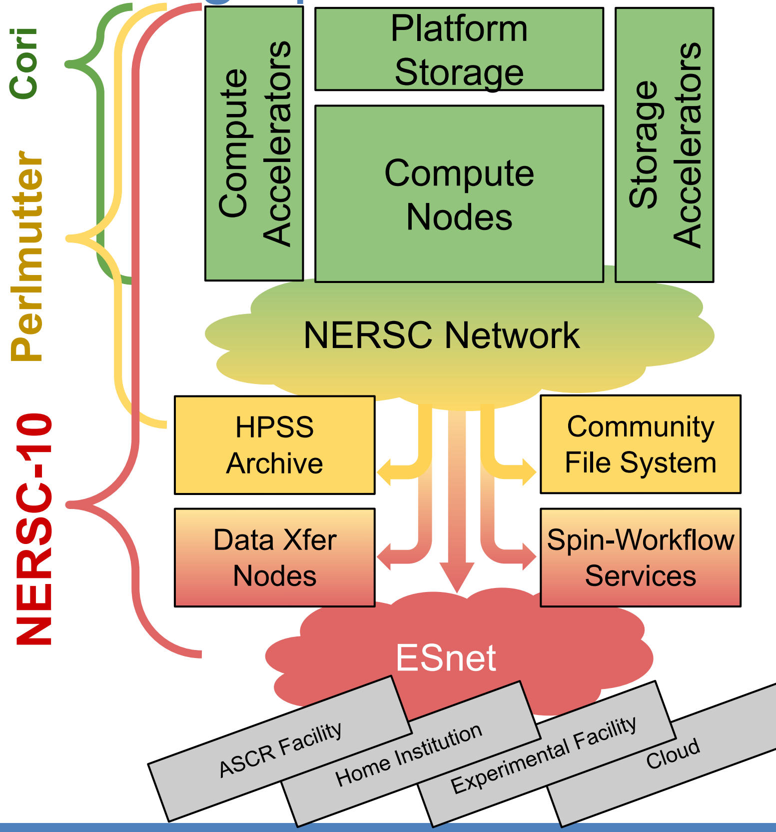

Sophisticated data services still evolving

- Managed, multi-stream data transfer
- I/O libraries and flash filesystems
- Sophisticated deep learning, and software frameworks
- Containerised services - portable and resilient
- Rich ecosystem of libraries to build portal and workflow tools
- Python ecosystem - productive language with performant libraries

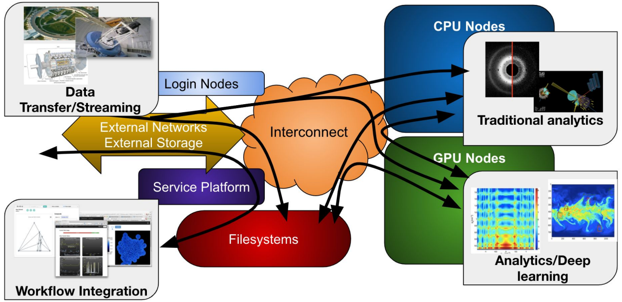

From WB. Data, Analytics and AI on Supercomputers for Science https://sites.google.com/lbl.gov/data-talks/

But remaining challenges include:

- Workflow services don't extend into compute and data infrastructure
- Divergence between HPC and cloud workflow and data tools and approaches
- Lack of widely accessible tooling to support FAIR data principles
- Interactive user interfaces for HPC compute are still quite limited
- Productive languages are difficult to scale to large HPC systems
- Data volumes still outpace I/O so batch processing and filtering needed (and inefficient)
- Deep learning methods can be opaque, need heavy tuning and further tuning at scale
- Ad-hoc inference on experimental data based on modelling and simulation

11:00 AM	I/O Profiling on Perlmutter with Darshan	Alberto Chiusole & Jean Luca Bez

But remaining challenges include:

- Workflow services don't extend into compute and data infrastructure
- Divergence between HPC and cloud workflow and data tools and approaches
- Lack of widely accessible tooling to support FAIR data principles
- Interactive user interfaces for HPC compute are still quite limited
- Productive languages are difficult to scale to large HPC systems
- Data volumes still outpace I/O so batch processing and filtering needed (and inefficient)
- Deep learning methods can be opaque, need heavy tuning and further tuning at scale
- Ad-hoc inference on experimental data based on modelling and simulation

1:00 PM	Containers-as-a-Service: Spin	Cory Snavely
1:30 PM	Workflows: Pegasus Workflow Manager	Nicholas Tyler
2:00 PM	Superfacility API	Bjoern Enders

But remaining challenges include:

- Workflow services don't extend into compute and data infrastructure
- Divergence between HPC and cloud workflow and data tools and approaches
- Lack of widely accessible tooling to support FAIR data principles
- Interactive user interfaces for HPC compute are still quite limited
- Productive languages are difficult to scale to large HPC systems
- Data volumes still outpace I/O so batch processing and filtering needed (and inefficient)
- Deep learning methods can be opaque, need heavy tuning and further tuning at scale
- Ad-hoc inference on experimental data based on modelling and simulation

10:30 AM	Containers for HPC: Shifter and Podman	Daniel Fulton
11:00 AM	Scaling Python Applications	Daniel Margala
11:30 AM	Julia	Johannes Blaschke

But remaining challenges include:

- Workflow services don't extend into compute and data infrastructure
- Divergence between HPC and cloud workflow and data tools and approaches
- Lack of widely accessible tooling to support FAIR data principles
- Interactive user interfaces for HPC compute are still quite limited
- **Productive languages are difficult to scale to large HPC systems**
- Data volumes still outpace I/O so batch processing and filtering needed (and inefficient)
- Deep learning methods can be opaque, need heavy tuning and further tuning at scale
- Ad-hoc inference on experimental data based on modelling and simulation

1:00 PM	Data Visualization: Altair Demo	Annette Greiner
1:30 PM	Deep Learning at Scale on Perlmutter	Steven Farrell
2:00 PM	Python on GPUs: JAX	Nestor Demeure

But remaining challenges include:

- Workflow services don't extend into compute and data infrastructure
- Divergence between HPC and cloud workflow and data tools and approaches
- Lack of widely accessible tooling to support FAIR data principles
- Interactive user interfaces for HPC compute are still quite limited
- Productive languages are difficult to scale to large HPC systems
- Data volumes still outpace I/O so batch processing and filtering needed (and inefficient)
- **Deep learning methods can be opaque, need heavy tuning and further tuning at scale**
- Ad-hoc inference on experimental data based on modelling and simulation

2:00 PM	Python on GPUs: JAX	Nestor Demeure

But remaining challenges include:

- Workflow services don't extend into compute and data infrastructure
- Divergence between HPC and cloud workflow and data tools and approaches
- Lack of widely accessible tooling to support FAIR data principles
- Interactive user interfaces for HPC compute are still quite limited
- Productive languages are difficult to scale to large HPC systems
- Data volumes still outpace I/O so batch processing and filtering needed (and inefficient)
- Deep learning methods can be opaque, need heavy tuning and further tuning at scale
- Ad-hoc inference on experimental data based on modelling and simulation

Evolution of data services

Compose services and compute seamlessly

Experiment with and apply performant, productive analytics at scale

Leverage large AI models, fine-tune to new problems, apply to new data pipelines

Discover through robust science-informed AI and inference approaches

Curate and re-use data through FAIR management services

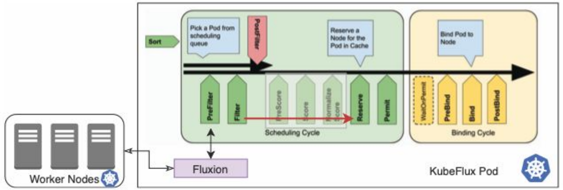

e.g. KubeFlux Misale et. al

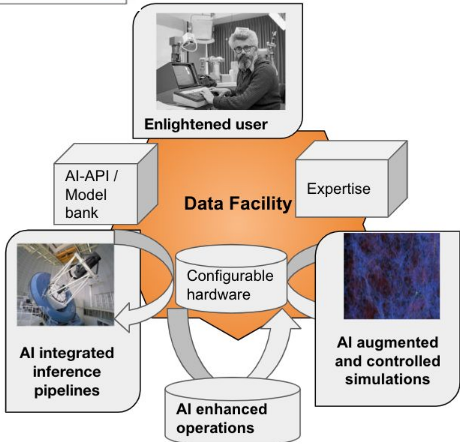

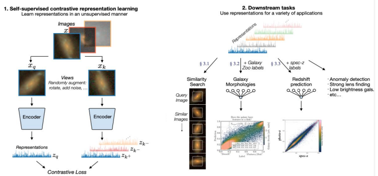

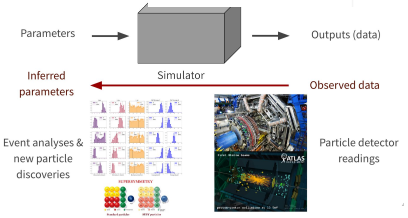

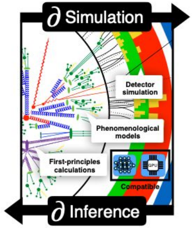

Etalumis: Baydin et. al . SC19: arXiv:1907.03382
See also "Surrogating" project PI Seljak (ML Forum talk)

Unifying HEP Simulation and Inference
Nachman et. al.

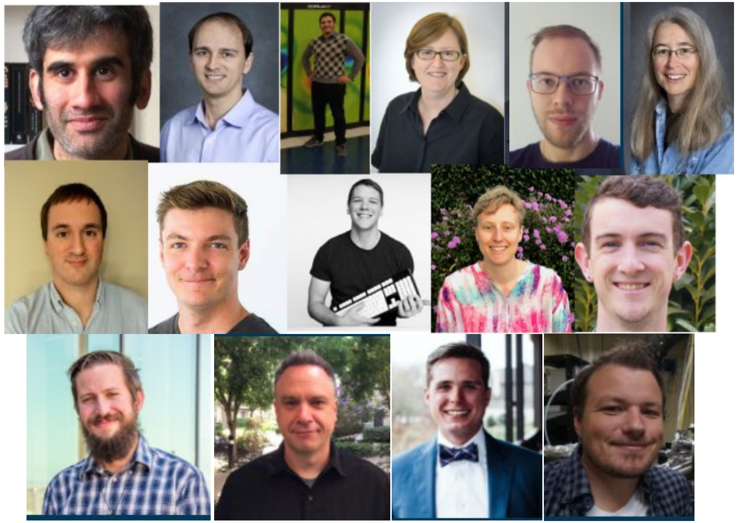

Questions?
Collaboration?

Wahid Bhimji
wbhimji@lbl.gov

https://docs.nersc.gov/analytics/analytics/
https://docs.nersc.gov/machinelearning/
https://docs.nersc.gov/services/