IRI: the Integrated Research Infrastructure

Update to ASCAC

29th May 2024

Ryan Adamson, Debbie Bard, Amber Boehnlein, Ben Brown, Tiffany Connors, Eli Dart, Sudip Dosanjh, Rafael Ferreira da Silva, Hal Finkel, Chin Guok, Graham Hayes, Inder Monga, Sarp Oral, Mike Papka, Arjun Shankar, Nhan Tran, Tom Uram

Standup of the IRI Program is a DOE FY24-25 Agency Priority Goal

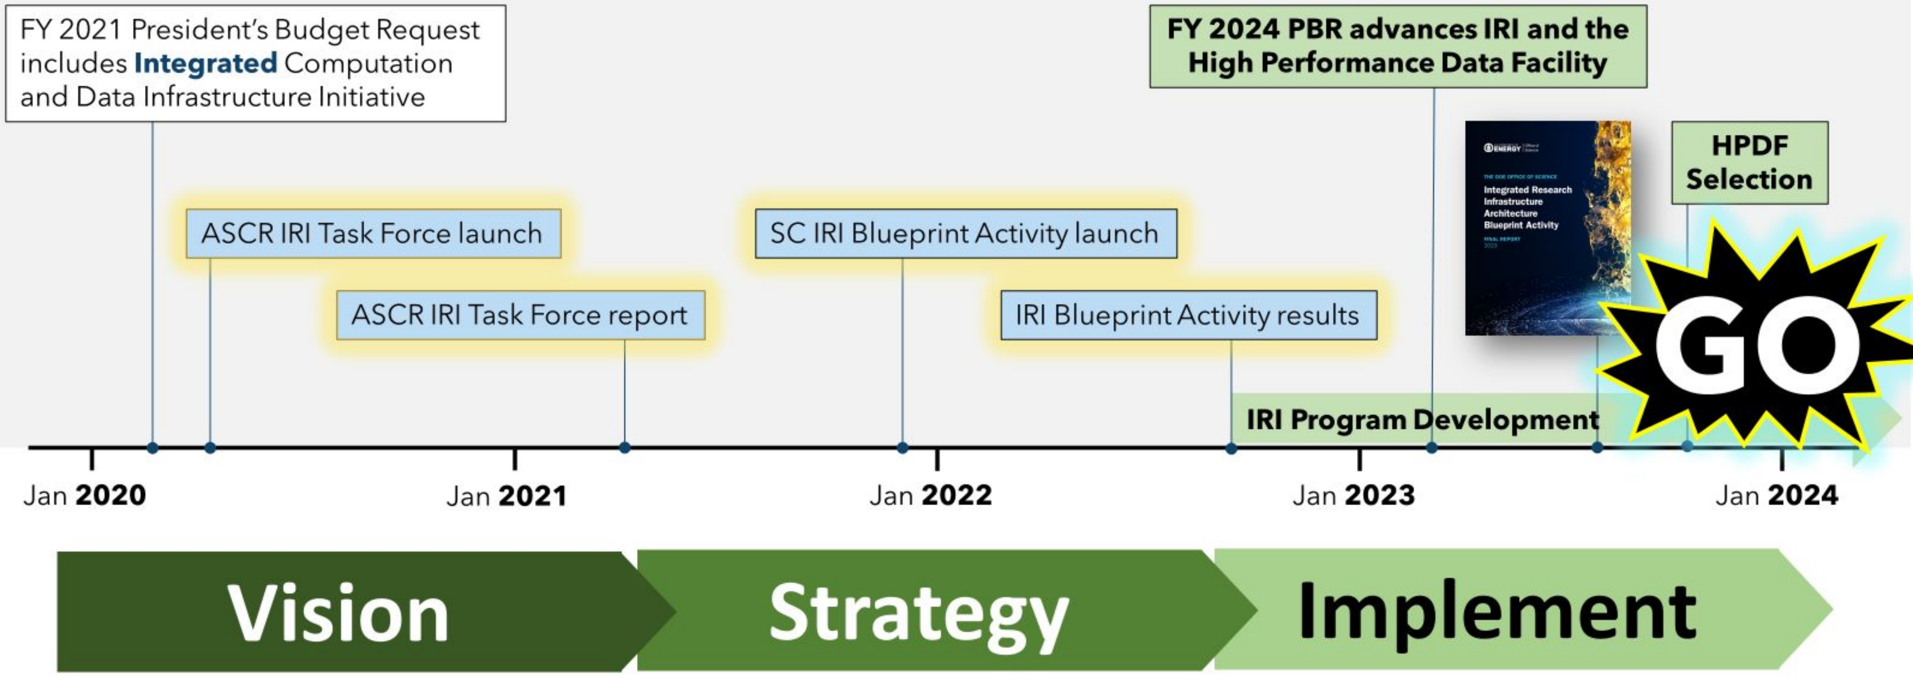

IRI Program launch is a DOE FY24-25 Agency Priority Goal. ASCR is implementing IRI through these four major elements.

1. Invest in IRI foundational infrastructure
2. Stand up the IRI Program governance and technical activities
3. Bring IRI projects into formal coordination
4. Deploy an IRI Science Testbed across the ASCR Facilities

These are all connected.
These are each essential.

The ASCR Facilities are firmly embedded in the IRI ecosystem

IRI is permeating everything we do - multi-facility workflows are an integral part of our major infrastructure design and strategic planning.

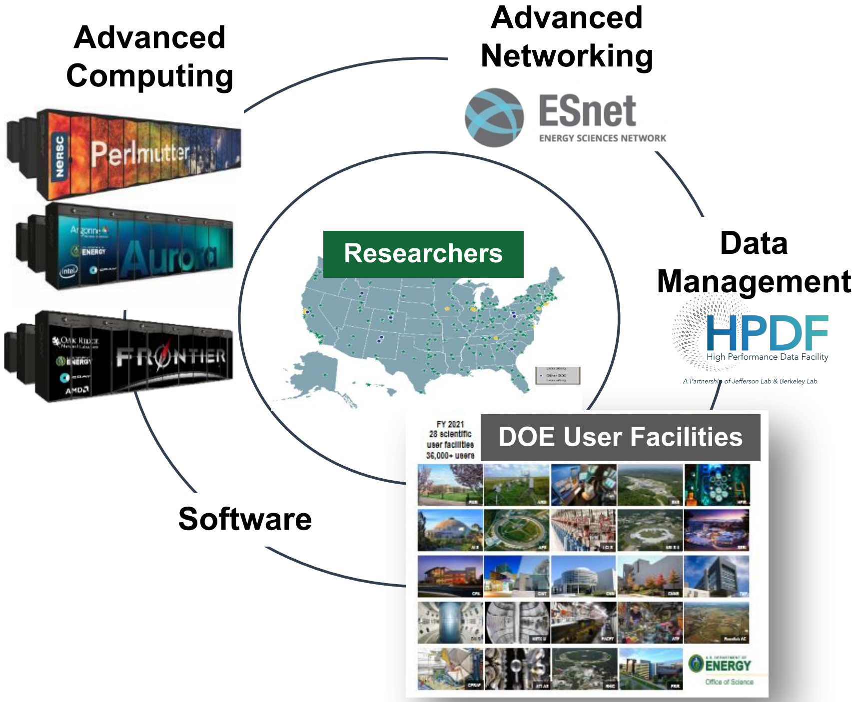

- IRI requires an exponentially growing amount of collaboration across ASCR facilities - it's changing how we work together
  - Discussed every week at the ASCR Facility Directors meeting
  - Catalysing cross-facility events on everything from training to API design
- ASCR facilities have made a lot of progress in supporting complex workflows
- Ongoing projects, testbeds, and major acquisitions are targeting IRI (ESnet6, NERSC-10, OLCF-6, HPDF and ALCF-4)

Standup of the IRI Program is a DOE FY24-25 Agency Priority Goal

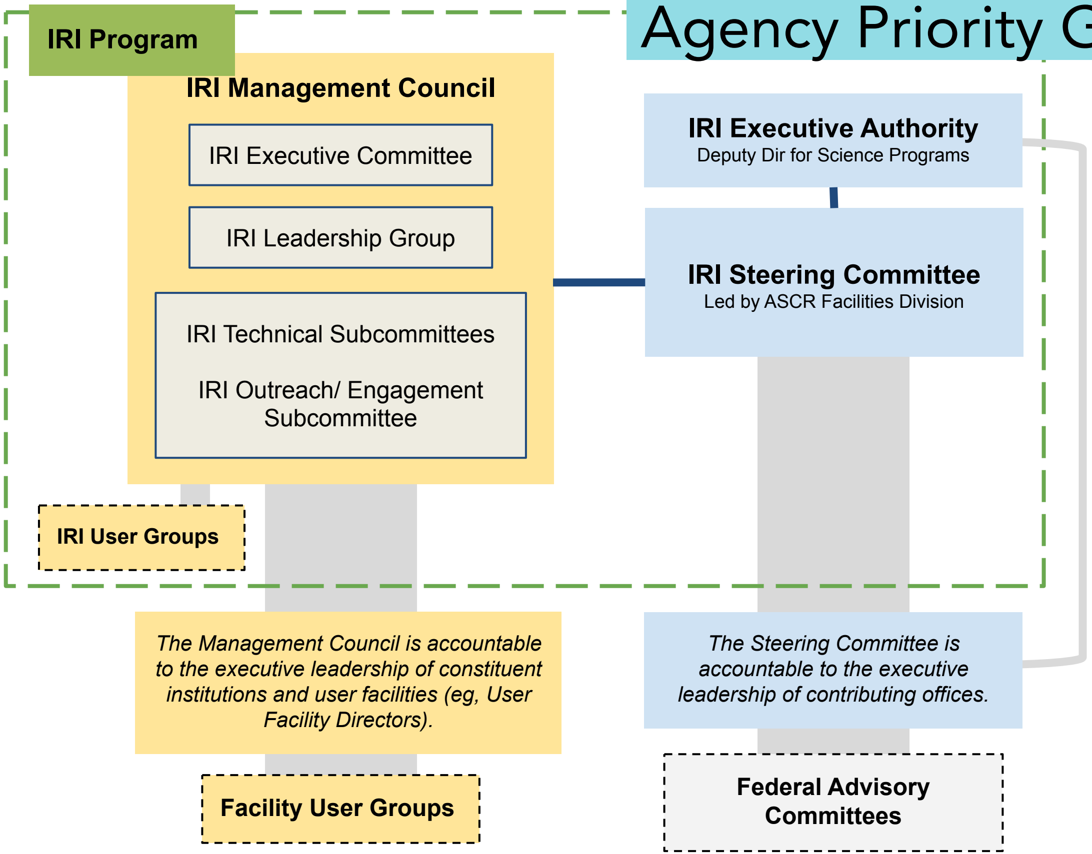

In the past 6 months, the IRI Management Council has launched the IRI program, building on work and experience from across the DOE.

We're excited to share with you the science IRI will enable, the initial areas of technical work including Pathfinder Projects, how IRI will engage with the DOE community, and how all this work will be coordinated and managed.

IRI timeline: Where we're heading

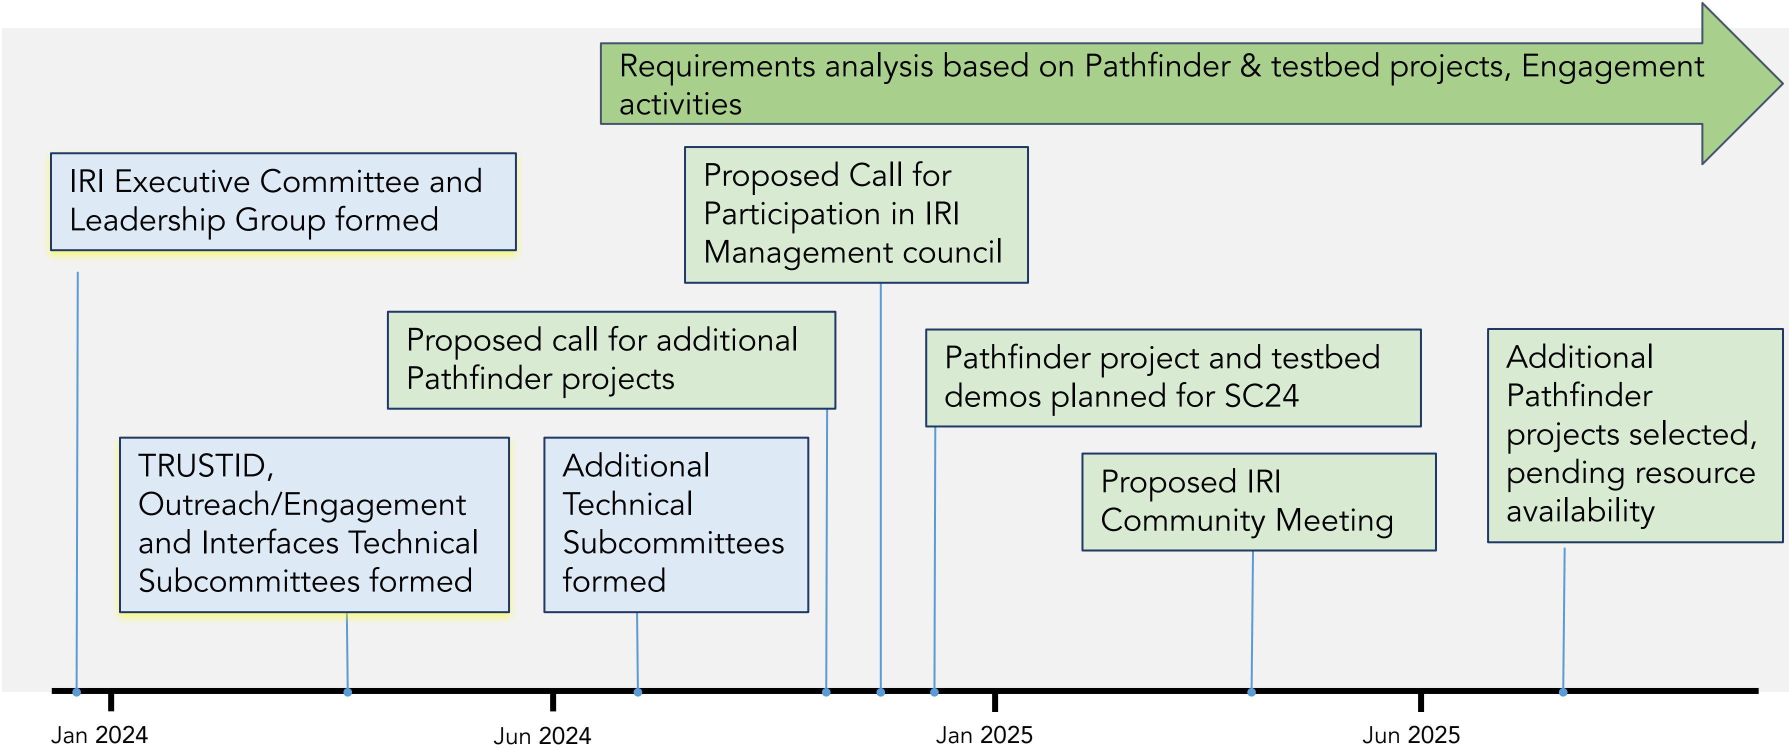

Outline

- IRI activities happening at the ASCR facilities now [Debbie Bard/Tom Uram/Chin Guok/Rafael Ferreira da Silva]
- How the Architectural Blueprint Activity is driving our work [Debbie Bard]
- IRI Partner pathfinder projects [Tom Uram]
- IRI program areas and technical work [Tom Uram]
  - TRUSTID Technical Subcommittee [Tiffany Connors/Ryan Adamson]
  - Engagement/Outreach Technical Subcommittee [Eli Dart]
- ASCR major infrastructure + testbeds [Debbie Bard]
- IRI Management structure [Debbie Bard]
- IRI Research [Hal Finkel]
  - Fast ML for science and the *extreme* edge [Nhan Tran]
- Summary [Debbie Bard, ASCR Facility Directors]

The IRI program builds on long-term support for multi-facility workflows

National Energy Research Scientific Computing Center

NERSC supports over 10,000 users each year, from all across the DOE.

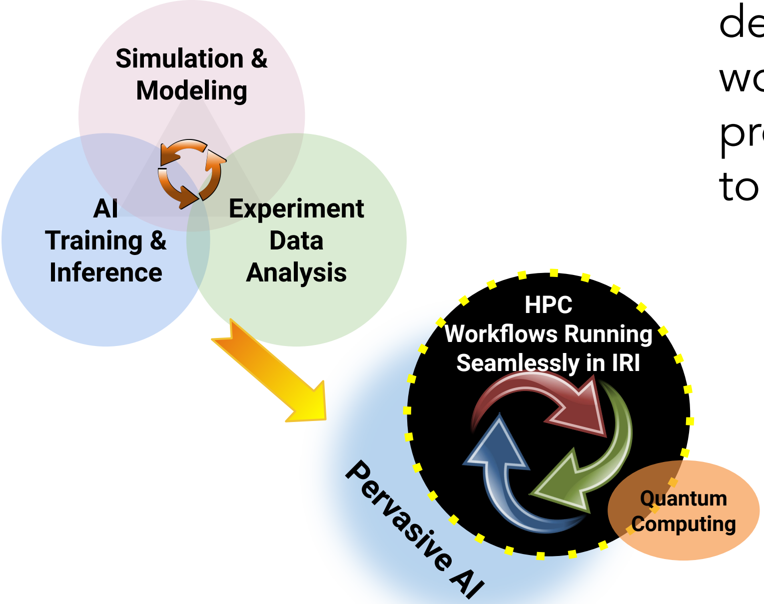

The Superfacility project [2019-2022] kick-started our work on the tools needed for multi-site science. It was designed to enable autonomous workflows from a set of 8 partner projects, and the impact has scaled to benefit all NERSC users.

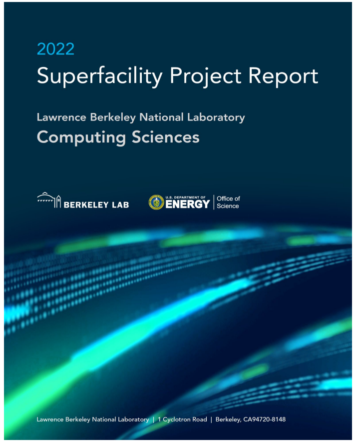

- > 20 projects use the realtime qos
- > 250 users, >85 projects use Spin
- > 1500 unique Jupyter users per month
- > 40 projects use the NERSC API, one request every 2 sec
- > 1400 users are now logging in with a home lab identity via Federated ID

Bring IRI projects into formal coordination

The IRI program builds on long-term support for multi-facility workflows

Argonne NATIONAL LABORATORY | Argonne Leadership Computing Facility

ALCF works closely with experimental and observational facilities for mod/sim, near real-time analysis, and data-intensive science and integrated deep learning. Argonne's Nexus project has demonstrated automated cross-facility workflows using a combination of ALCF-deployed services and Globus tools.

HEP - ATLAS

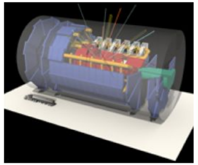

2015-16

Hundreds of millions of core hours of simulation and analysis for ATLAS (DD, ADSP, ESP)

FES - DIII-D

2017

Near real-time analysis of DIII-D fusion experiment data, powered by Balsam workflows (DD); deep learning for fusion (ESP)

HEP - LSST

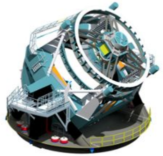

2017-18

Simulation and analysis of telescope images for LSST-DESC (ADSP)

HEP - LIGO

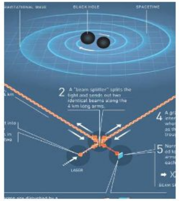

2018-19

Deep learning for gravitational wave detection with LIGO (ADSP)

BES/OTHER

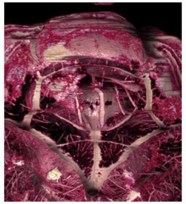

2018-

3D Reconstruction of mouse brain from APS imagery at Argonne (DD, ADSP, ESP)

FES

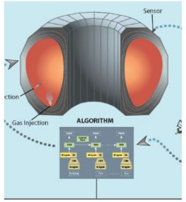

2018-

Hyperparameter tuning at exascale to predict and mitigate disruption events (ESP)

BES - LS

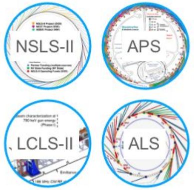

2019-

Online analysis of experiment data from multiple light sources using near-real-time queue on Theta (ADSP)

HEP - SBND

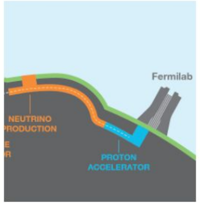

2019-

Large-scale simulation and deep learning for SBND/DUNE (DD, ADSP)

Bring IRI projects into formal coordination

The IRI program builds on long-term support for multi-facility workflows

ESnet interconnects all the labs and facilities in DOE

- Multi-facility IRI workflows *depend* on existence of a high-performance science network
- ESnet supports IRI workflows between all DOE user facilities and facilitates data/instrument connections from outside

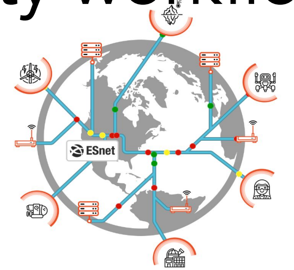

ESnet is the data circulatory system for the DOE Office of Science

ESnet requirement reviews meta-analysis shows IRI workflow patterns permeate across SC

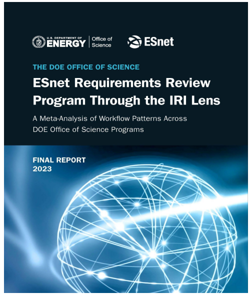

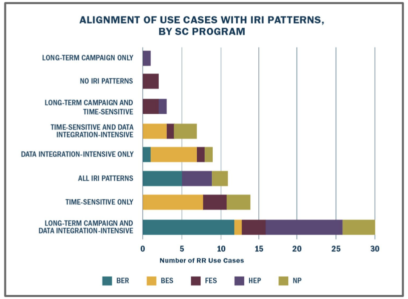

Developing advanced services to tackle IRI workflow patterns

- Science-grade networking supports current needs well
- Advanced services co-designed to match with partner facility new IRI needs

Successful collaboration examples of prototyping IRI workflows

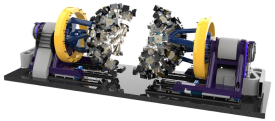

Co-design high performance streaming for GRETA

EJFAT
ESnet-Jefferson Lab FPGA Accelerated Transport

Co-develop ESnet/JLab real-time streaming load balancer

ECP EXASCALE COMPUTING PROJECT
Exascale for Free Electron Lasers

Co-deploy dynamic resource provisioning for fast feedback

Bring IRI projects into formal coordination

The IRI program builds on long-term support for multi-facility workflows

OLCF supports 1700+ users annually from 25+ countries

Collaborations with INTERSECT and ESnet led to the development of capabilities supporting IRI science patterns, multi-facility workflows, AI, quantum computing, and energy efficiency

The Advanced Computing Ecosystem (ACE), Frontier, and Summit+ are foundational resources for enabling IRI science

Program	Workflows Project Example
BES	CNMS, SLAC, SNS Beamlines…
BER	ARM, ESGF, APPL, JGI…
NP	GRETA/Deleria, ALICE…
HEP	Celeritas, Panda…
FES	WDMApp Coupling…

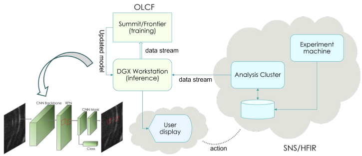

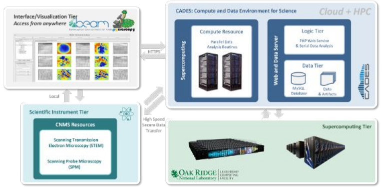

ACE/IRI Testbed

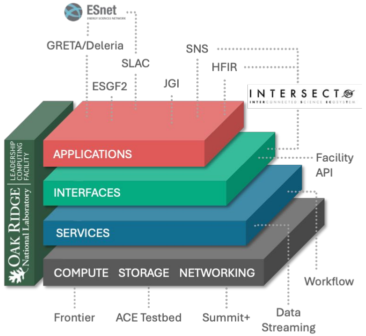

Summit+ Allocation for IRI

End-to-End pipelines from facility-to-facility

Autonomous HPC-AI Workflows

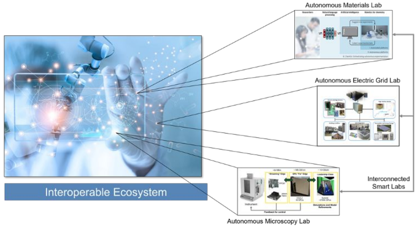

INTERSECT Initiative: Lab-wide integration of research infrastructure

INTERSECT
INTERCONNECTED SCIENCE ECOSYSTEM

Bring IRI projects into formal coordination

The IRI ABA gathered requirements from 24 science teams, covering the span of DOE science

- **NP** Large-scale simulations integrating experimental data from DOE and international nuclear physics facilities.
- **BES** Autonomous AI/ML-driven experiments such as for ALS and NSLS-II
- **NP** AI/ML incorporated into simulations to drive exploration of parameter space using codes like LAMMPS in the workflow.
- **BER** Applying AI/ML to EMSL data for analysis
- **BER** Combining multi-modal data from simulations and experiments for molecular scale imaging workflows (instrument + compute + storage)
- **HEP** Astronomical spectroscopic observations at DESI/KPNO, survey data daily streaming to HPC, target selection pipeline
- **FES** Heterogeneous data handling and real-time analysis for fusion power experiment steering at DIII-D and JET for fusion.
- **HEP** Astronomy data management - DESI workflows - processing 100s of TBs of data utilizing HPC.
- **BER** Crystallography microscopy and synchrotron light sources workflows and local and HPC computing.
- **FES** Fusion workflows: 10s ~45GB per pulse, about 500TB total using custom SQL and NoSQL databases, computing with local and institutional resources, data exposed via APIs also for future ITER workflows
- **BER** HPC-enabled high throughput sequencing, large scale sequence data analysis, sample and data life cycle, data product development (e.g., at JGI)
- **BES** Beamline data transfer at scale to HPC for real-time analysis and experiment steering, data transfer at scale to HPC (e.g., at SSRL)

- **FES** DIII-D Fusion experiments + simulations, near-real-time analysis to determine/predict plasma cond
- **HEP** LHC simulation workflows; eve
- **NP** GRETA spectrometer online co campaigns.
- **FES** DIII-D Tokamak plasma physics Integrating ASCR HPC facilities
- **HEP** Event generation for of comple but not data-intensive.
- **BES** APS 8-ID-I Small-angle XPCS management workflow includin
- **BES** Metallurgy, MIDAS x-ray analys microscopy beamline data. Hig and ALCF for on-demand comp
- **BES** Materials science. High data th analysis
- **BES** Light source data processing w ptychography; using AI/ML to h
- **FES** Data management and real-tim
- **BES** Data pipelines for high-speed detectors used at lights sources and NSRCs. Workflow development using NERSC.
- **BES** ML Autonomous materials characterization workflows using data (100s GBs per day) collected at light sources, neutron sources, and NSRCs.

The IRI-ABA collated and categorized the many challenges scientists face in building workflows integrated across DOE resources.

Need insights into how workflows/resources are performing, for troubleshooting or performance tuning

Need common/compatible AuthN/AuthZ models

Need to move data across facilities and use different systems at different steps of data processing chain

Need for some level of workflow/software standardization (including APIs to access resources, data formats, etc) and associated support model

Need a joint proposal mechanism so that campaigns can request resources for longer duration resources across the facilities

Facility cooperation requires a systematic ability to ask for reservations or resources in a timely manner

Allocations and accounts that span resources

The IRI-ABA collated and categorized the many challenges scientists face in building workflows integrated across DOE resources.

The IRI Framework comprises:

> 3 IRI Science Patterns represent integrated science use cases across DOE science domains.

Time-Sensitive Patterns

Data-Integration Patterns | Long Campaign Patterns

> 6 IRI Practice Areas represent critical topics that require close coordination to realize and sustain a thriving IRI ecosystem across DOE institutions.

Workflows, Interfaces & Automation | Scientific Data Lifecycle | User Experience

Portable/Scalable Solutions | Cybersecurity & Federated Access | Resource Co-Operations

IRI is now implementing the frameworks proposed in the ABA

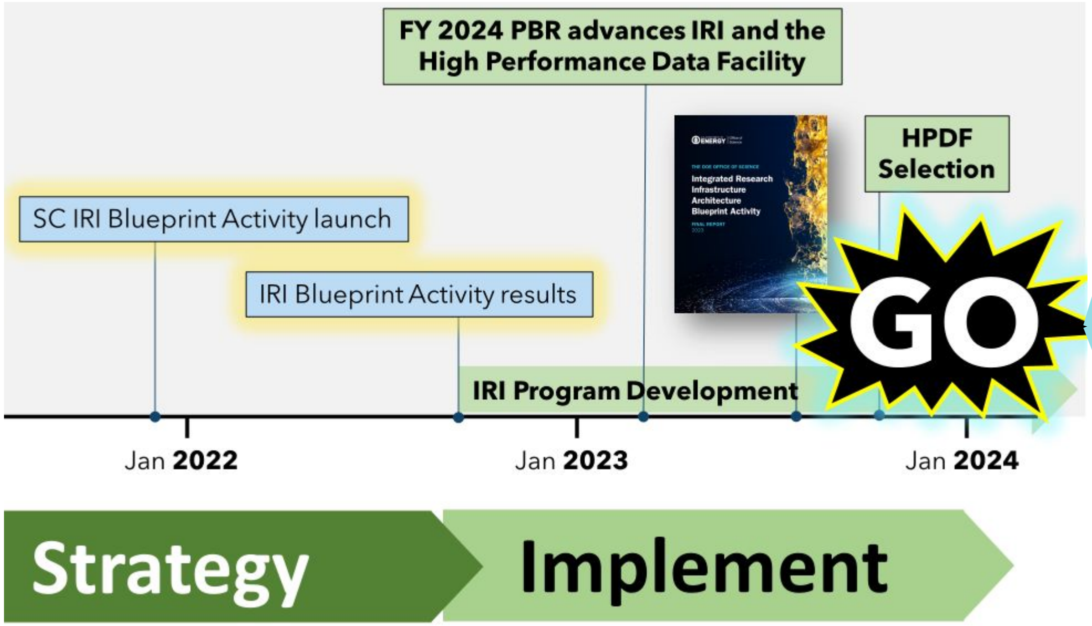

- Engaged with Pathfinder projects to partner with IRI in initial stages
- Launched Technical Subcommittees to do the work proposed in the ABA
- Developed the IRI Management Council structures necessary to coordinate multi-institutional work

IRI Pathfinder projects are partnering with us in the first stages of developing IRI-enabling technologies

Many science teams are already doing cross-site science. We heard about the challenges they face during the ABA exercises, and in our everyday interactions with ASCR users.

The IRI program has identified "pathfinder projects", nominated by their program managers as ready to partner with us in the first stages of developing IRI technologies.

Initial pathfinder projects:
- Lightsources: ALS, APS, LCLS, NSLS-II (*BES*)
- DIII-D (*FES*)
- ESGF (*BER*)

All have dedicated resources to working with the emerging IRI technical work areas, providing requirements, feedback on initial plans and implementations

ASCR facilities are also working with several other projects that are IRI-related.

We know other science teams are interested in becoming IRI pathfinder partners; we will talk about how we plan to expand our partnerships later.

DIII-D pathfinder project

- The DIII-D fusion science team need to make decisions about how to tune their experiment parameters between plasma shots.
  - Working towards a computing model for ITER, which will produce 1PB/day
- DIII-D is partnering with ALCF, ESnet, and NERSC to run analyses robustly during experiment.
- IRI goal: DIII-D will have automated, resilient, fast turnaround for large-scale data analysis coupled with simulation (“digital twin”).

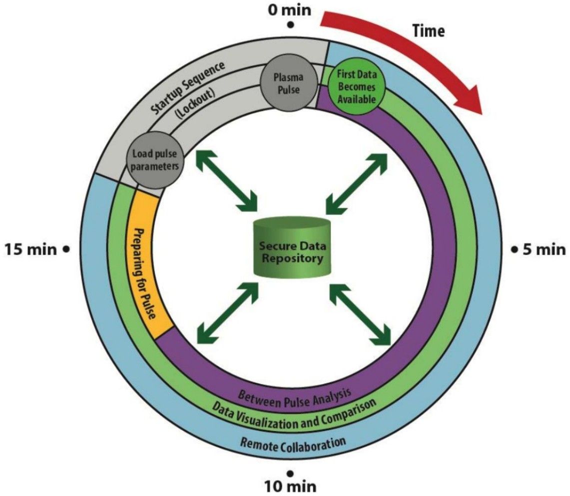

ESGF Pathfinder project

- The Earth System Grid Federation (ESGF) provides access to world's largest climate model data archive
  - With a growing multi-PB archive across the labs, and worldwide, the ESGF project is building tailored services to empower thousands of scientists to leverage climate data in their research.
- The ESGF2-US project is collaborating with ALCF, OLCF and Globus to support cross facility workflows, including resource selection flexibility and execution optimization
- IRI goal: ESGF will have resilient, server-side computing distributed data services to deliver value-added services to make the archive more accessible to a broader set of users.

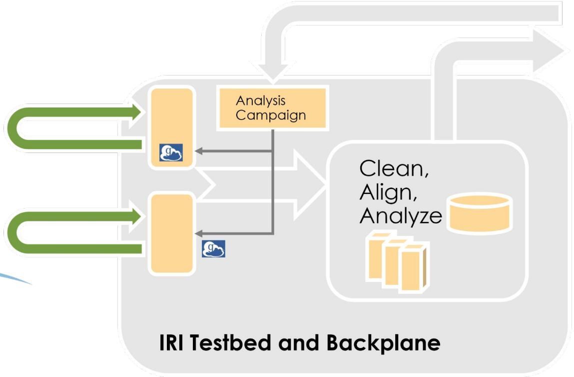

User-facing value added products accessed through

- Notebooks
- Viz tools
- Portals

BES Light Sources Pathfinder projects

- The BES light sources need a scalable and adaptive end-to-end architecture for data management through the full lifecycle to accelerate time to science and insights
  - Multiple order-of-magnitude increase in demand for computing resources over the next decade
- APS, ALS, LCLS and NSLS-II are collaborating with ALCF, ESnet, NERSC, and OLCF to demonstrate multi-facility resource coordination
- IRI goal: BES light sources will have access to a multi-tiered architecture with varying types and scales of computing, networking, and storage capabilities

Advanced Photon Source

ALS ADVANCED LIGHT SOURCE

LCLS

NSLS-II

Pathfinder Projects

NP	Large-scale simulations integrating experimental data from DOE and international nuclear physics facilities.
BES	Autonomous AI/ML-driven experiments such as for ALS and NSLS-II
NP	AI/ML incorporated into simulations to drive exploration of parameter space using codes like LAMMPS in the workflow.
BER	Applying AI/ML to EMSL data for analysis
BER	Combining multi-modal data from simulations and experiments for molecular scale imaging workflows (instrument + compute + storage)
HEP	Astronomical spectroscopic observations at DESI/KPNO, survey data daily streaming to HPC, target selection pipeline
FES	Heterogeneous data handling and real-time analysis for fusion power experiment steering at DIII-D and JET for fusion.
HEP	Astronomy data management - DESI workflows - processing 100s of TBs of data utilizing HPC.
BER	Crystallography microscopy and synchrotron light sources workflows and local and HPC computing.
FES	Fusion workflows: 10s ~45GB per pulse, about 500TB total using custom SQL and NoSQL databases, computing with local and institutional resources, data exposed via APIs also for future ITER workflows
BER	HPC-enabled high throughput sequencing, large scale sequence data analysis, sample and data life cycle, data product development (e.g., at JGI)
BES	Beamline data transfer at scale to HPC for real-time analysis and experiment steering, data transfer at scale to HPC (e.g., at SSRL)

FES	DIII-D Fusion experiments + simulations, near-real-time analysis to determine/predict plasma conditions.
HEP	LHC simulation workflows; event generation for integration of WLCG with HPC
NP	GRETA spectrometer online computing pipeline construction - for 1-5 day campaigns.
FES	DIII-D Tokamak plasma physics experiments - diagnostic data collection Integrating ASCR HPC facilities with the WLCG, for LHC computing.
HEP	Event generation for of complex simulation pipeline. Distributed, compute-intensive but not data-intensive.
BES	APS 8-ID-I Small-angle XPCS x-ray spectroscopy, high-frame rate camera, data management workflow including HPC. 0.2 PB of unsparsified data generated/day.
BES	Metallurgy, MIDAS x-ray analysis software for APS high-energy X-ray diffraction microscopy beamline data. High data throughput. Pipeline development to NERSC and ALCF for on-demand computing at scale
BES	Materials science. High data throughput, facility data transfer from APS to PNNL for analysis
BES	Light source data processing workflows for large datasets (TBs to PBs), particularly ptychography; using AI/ML to help reduce data quickly.
FES	Data management and real-time analysis for fusion experiments at PPPL.
BES	Data pipelines for high-speed detectors used at lights sources and NSRCs. Workflow development using NERSC.
BES	ML Autonomous materials characterization workflows using data (100s GBs per day) collected at light sources, neutron sources, and NSRCs.

Architectural Blueprint Activity

IRI technical work is driven by the Architectural Blueprint activities, requirements from Pathfinder projects, and many collaborations with science teams doing IRI-relevant work.

Initial IRI program areas are drawn from the ABA Practice Areas

IRI-ABA Practice Area	IRI Program Area	Status
Resource Co-Operations	IRI Allocations Program/ Scheduling/Preemption Technical Subcommittee	Under development; Under development
Cybersecurity and Federated Access	TRUSTID Technical Subcommittee	Launched!
User Experience	Outreach/Engagement Technical Subcommittee	Launched!
Workflows, Interfaces & Automation	Interfaces Technical Subcommittee	Launched!
Scientific Data Lifecycle	Data Movement Technical Subcommittee	Under development
Portable / Scalable Solutions	Software Deployment and Portability	Under development

IRI technical work will be managed via Technical Subcommittees

The role of the IRI Technical Subcommittees is to coordinate work on a particular technical area, identify goals and estimate resources required to perform the work.

- Each subcommittee has two leads, who are responsible for reporting, voted for by the subcommittee members.
- Membership open to DOE labs and DOE user facilities
- Report to IRI Leadership Group, who will help prioritize work and identify requirements

Each subcommittee will stand up "working groups" to tackle specific technical topics, answer charge questions, define implementation plans etc.

- Membership may include non-DOE consultants

IRI program areas have goals for FY24/25

IRI Program Area	Sample Goals for FY24/25
IRI Allocations Program	Develop cross-facility IRI allocations test program
Scheduling/Preemption Technical Subcommittee	• Identify possible common policies for scheduling urgent compute needs • Optimize preemption capabilities across scheduling technologies
TRUSTID Technical Subcommittee	• Create design patterns to enable trusted interoperable cross-facility workflows • Deploy and pilot federated identity infrastructure across multiple facilities
Outreach/Engagement Technical Subcommittee	• Bring initial Pathfinder Projects under common coordination • Design and execute plan for broader community engagement
Interfaces Technical Subcommittee	• Design a minimal functional API and deploy it on the testbed systems • User interface: Jupyter already deployed at sites; explore how to make this uniform across testbeds
Data Movement Technical Subcommittee	• Deploy Science DMZ and Data Transfer Node infrastructure for IRI workflows • Deploy Globus; consider extensions and alternatives
Software Deployment and Portability	• Containers: Agree on a container technology and deploy on testbed systems

IRI program areas have goals for FY24/25

IRI Program Area	Sample Goals for FY24/25
IRI Allocations Program	Develop cross-facility IRI allocations test program
Scheduling/Preemption Technical Subcommittee	● Identify possible common policies for scheduling urgent compute needs ● Optimize preemption capabilities across scheduling technologies
TRUSTID Technical Subcommittee	● Create design patterns to enable trusted interoperable cross-facility workflows ● Deploy and pilot federated identity infrastructure across multiple facilities
Outreach/Engagement Technical Subcommittee	● Bring initial Pathfinder Projects under common coordination ● Design and execute plan for broader community engagement
Interfaces Technical Subcommittee	● Design a minimal functional API and deploy it on the testbed systems ● User interface: Jupyter already deployed at sites; explore how to make this uniform across testbeds
Data Movement Technical Subcommittee	● Deploy Science DMZ and Data Transfer Node infrastructure for IRI workflows ● Deploy Globus; consider extensions and alternatives
Software Deployment and Portability	● Containers: Agree on a container technology and deploy on testbed systems

Trusted IRI Designs (TRUSTID) Technical Activity

TRUSTID was established to assist in enabling secure and performant execution of cross-facility workflows within the IRI ecosystem.

Objectives of TRUSTID:

- Define design patterns to enable trusted interoperable cross-facility workflows within the IRI ecosystem
- Provide a path to incremental adoption of design patterns to enable IRI workflows to access a subset of services via federated authentication
- Promote interoperability and use of existing standards for integration of facility resources into diverse research environments
- Create documentation, best practices, and references to assist stakeholders in implementing the trusted design patterns
- Foster collaboration and knowledge sharing within the IRI and DOE community

TRUSTID: IRI Federated ID Demonstrations

Problem: IRI facilities and laboratories maintain independent identity access management systems which means more work for scientific end users

- The TRUSTID FedID working group is working to understand and solve this problem
- Our strategy is to:
  - Prioritize the scalability and repeatability of solutions
  - Ensure inclusivity of ASCR facilities through Show-and-Tell roundtable discussions
  - Develop blueprints and shareable best practices
  - Allow facilities to participate to a level commensurate with the security posture required by their laboratory

globus oneID InCommon LOGIN.GOV

Potential non-laboratory identity providers (IdP). Participating facilities will utilize external IdPs such as these or other external laboratory IdPs in demonstration of a Federated ID Pilot

The Outreach/Engagement Technical Activity

Bring initial Pathfinder Projects under common coordination

- Build momentum through early successes
- Identify and develop requirements and science requirements strategy for IRI

Broader community engagement

- Serve as contact interface for IRI constituents
- Identify additional projects, synergies, common issues, etc. as IRI develops

Initial pathfinder projects:

- Lightsources: ALS, APS, LCLS, NSLS-II (*BES*)
- DIII-D (*FES*)
- ESGF (*BER*)

Outreach/Engagement: early engagements

Readiness is a key aspect of early engagements

- Need to be able to make progress without significant new funding
- Experimental facility/project must be able to use multi-facility workflow

Different teams have different needs and capabilities

- Endeavor to cover multiple programs, workflow types early (learn quickly)
- Identify synergies and gaps
- Discover resource needs of different workflows/programs/facilities
- Collect requirements as we go - build on what we have

Human contact is important - conversations with constituents are vital

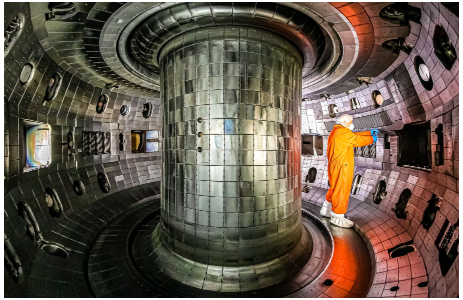

ESGF2 US

Advanced Photon Source

NSLS-II

LCLS

ALS ADVANCED LIGHT SOURCE

Nominal timeline for engagement activities

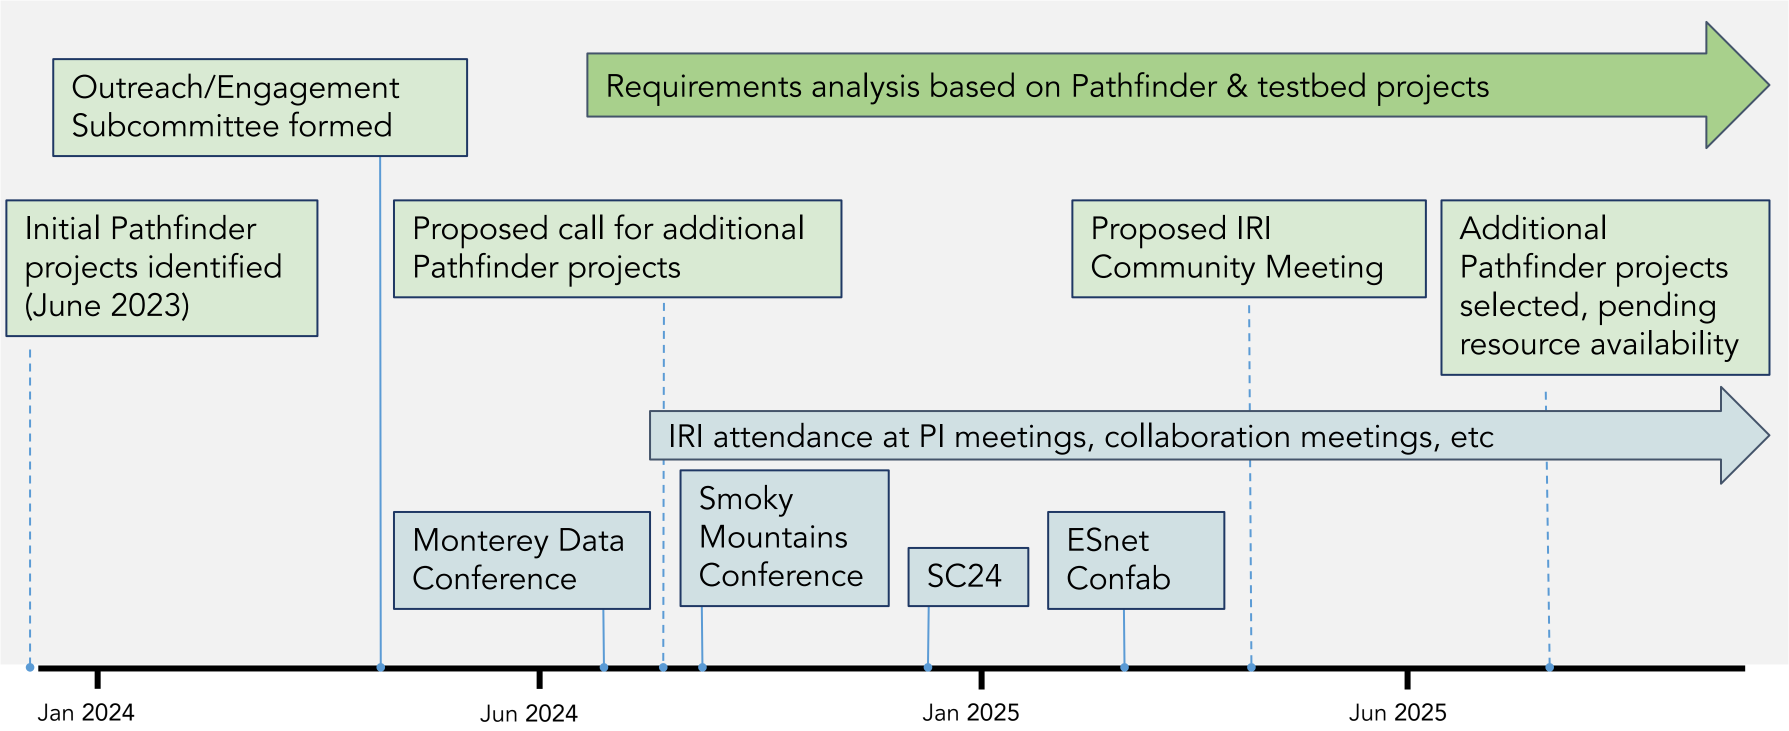

IRI is driving our major Infrastructure investments

- The NERSC-10 RFP explicitly asks for IRI-relevant features (eg end-to-end QoS, QoS Storage system, APIs, workflow portability)
- The OLCF-6 RFP asks for increased bandwidth and connectivity in and around the system to extend Leadership capabilities to enable new workflows between facilities
- The ALCF4 RFP will also feature IRI prominently

} KPPs , benchmarks and readiness programs are explicitly designed for IRI

- HPDF was conceived as a data focused component of the IRI ecosystem; we intend to provide a full range of services to support data driven IRI use cases.
- ESnet 6/7 is designed with IRI in mind
  - Fully integrated network automation
  - High fidelity traffic monitoring at scale
  - Network service composability
  - In-network compute and storage

Invest in IRI foundational infrastructure

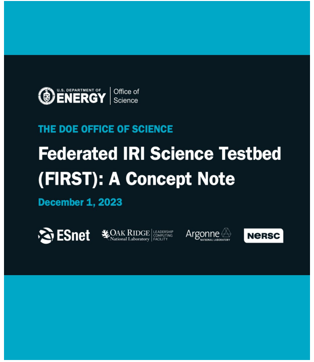

The IRI Testbed is a shared playground for IRI developers and pilot applications, designed to accelerate transition to operations

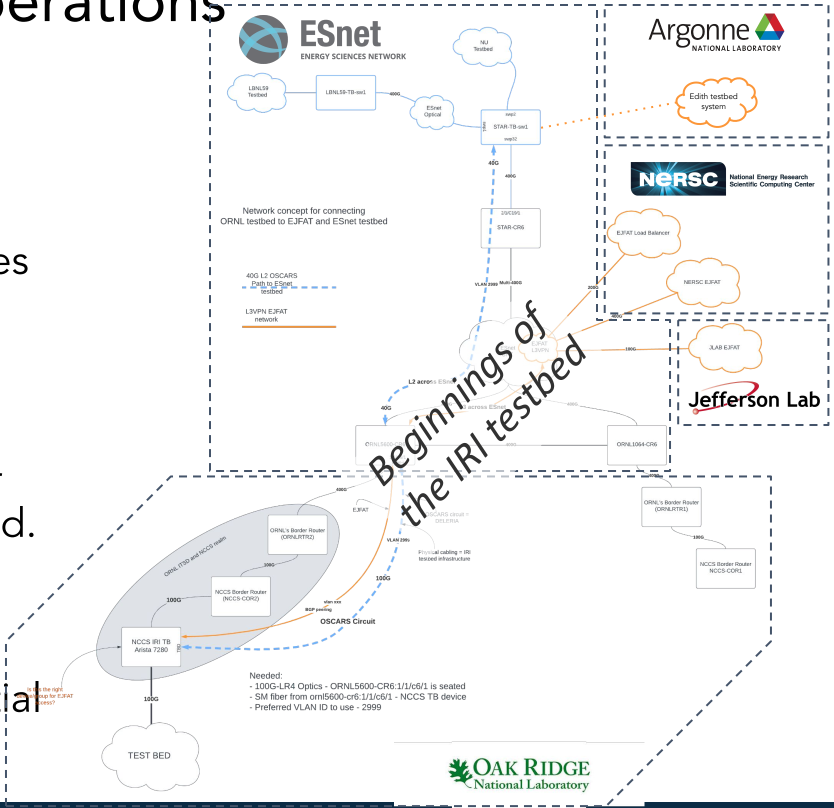

Key Testbed activities in FY24/25:

- Catalog existing testbeds at ASCR compute facilities
- Co-design and implement plan to connect existing ASCR testbeds using ESnet testbed network
- Develop a best current practice (BCP) connection service template that can be used to connect other (e.g. non-ASCR) test environments to the IRI testbed.
- Publish facility-specific testbed access procedures, AUPs, etc.
- Identify and execute early test cases to validate initial IRI testbed design, e.g., GRETA/Deleria, EJFAT

Deploy an IRI Science Testbed across the ASCR Facilities

ASCR Facilities are standing up the IRI Testbed

ESnet Testbed-NextGen upgrades driven by IRI testbed requirements

OLCF ACE testbed effort to explore gaps and develop novel capabilities for OLCF-6 and OLCF-7

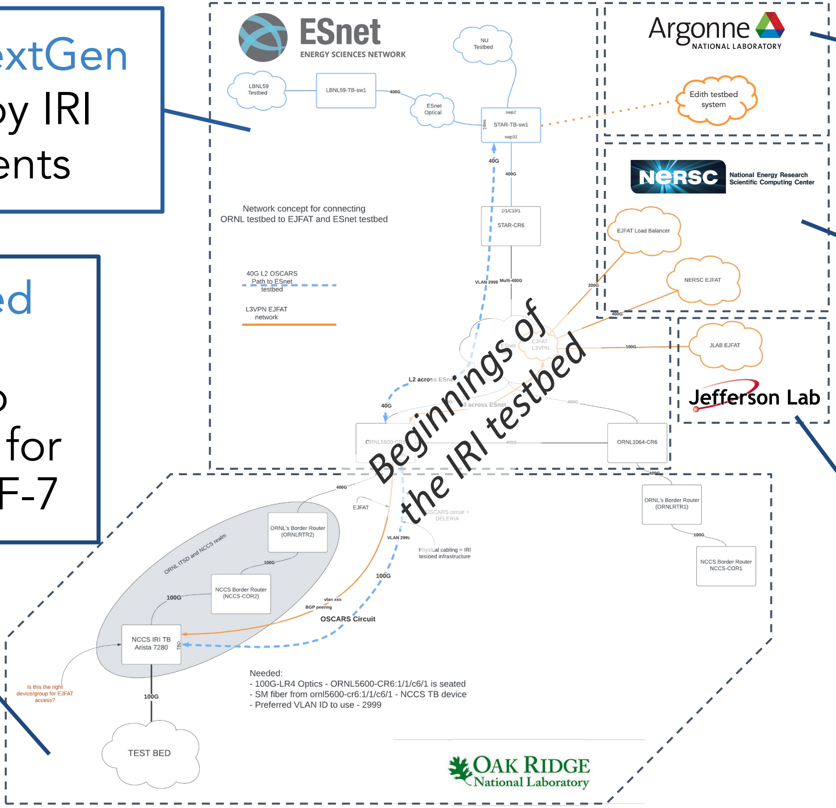

Edith IRI testbed resource is available for exploratory IRI use cases

The "POD" testbed is used for IRI experiments

HPDF will have a phased approach, starting with a testbed, followed by a small scale production facility that will be extended into the full scale HPDF

IRI Governance

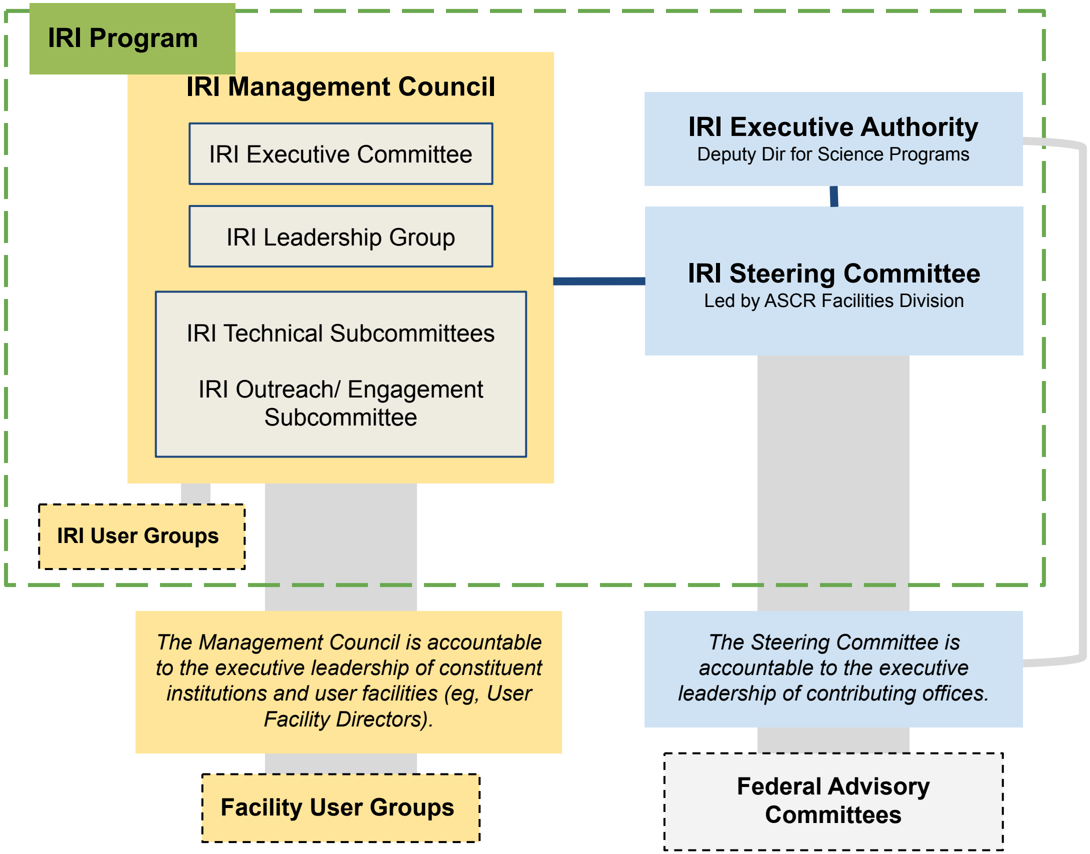

This structure was developed by the SC Headquarters IRI Coordination Group in 2022-23.

In late 2023 ASCR directed the ASCR Facilities to jumpstart the IRI Management Council.

Since early 2024, the Executive Committee and Leadership Group have been working to:

- Define roles, responsibilities and reporting structures
- Initial staffing of committees
- Early focus area: coordinating work of Technical Subcommittees and the IRI Pathfinder Projects.

The roles of the Executive Committee and the Leadership Group

The IRI Executive Committee:
- ASCR Facility Directors
- DOE Leadership

The role of the EC:
- Ongoing oversight of the program
- Allocate resources strategically
- Identify and manage potential risks
- Conduct annual reviews of IRI activities

ASCR Facilities are providing the initial leadership. The fully constituted Management Council will include senior personnel from partner SC user facilities and other DOE personnel.

IRI Leadership Group

NERSC — National Energy Research Scientific Computing Center
- Debbie Bard (Chair)
- Wahid Bhimji

Argonne National Laboratory — Argonne Leadership Computing Facility
- Tom Uram (Vice Chair)
- Rachana Ananthakrishnan

ESnet — Energy Sciences Network
- Chin Guok
- John MacAuley

HPDF — High Performance Data Facility — A Partnership of Jefferson Lab & Berkeley Lab
- Graham Heyes
- Shane Canon

Oak Ridge National Laboratory — Leadership Computing Facility
- Sarp Oral
- Rafael Ferreira da Silva

The role of the LG:
- Prioritize and organize the work
- Translate strategic goals into tactical plans
- Outreach to DOE Program Offices
- Work closely with Technical Subcommittees
- Coordinate closely with Executive Committee

The IRI Management Council will be expanding membership as it becomes more established

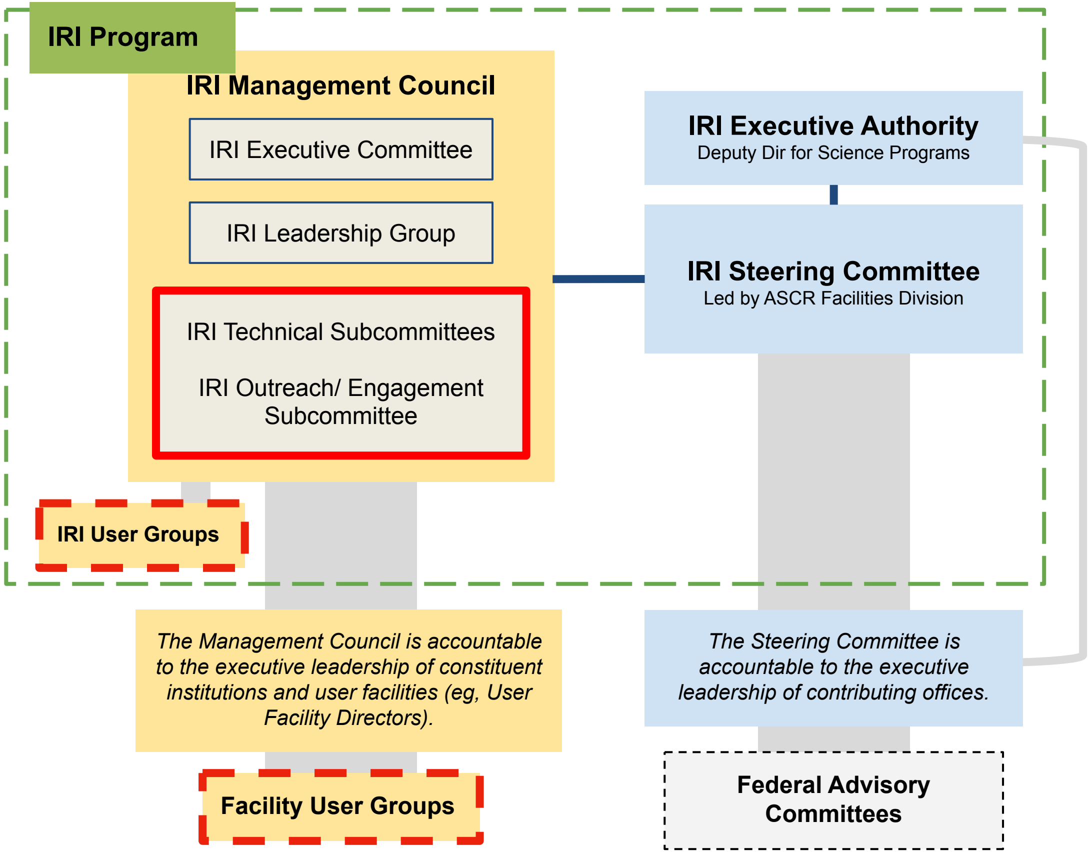

In the next year, the IRI Program will be seeking to expand its membership (in addition to engaging with the scientific community).

- A call for participation in the Technical Subcommittees and User Groups is expected in August 2024

ASCR Research and the Integrated Research Infrastructure (IRI)

- ASCR Research supports the IRI by aligning basic research priorities with future needs of IRI objectives.
- The following recent portfolios are strongly aligned with IRI objectives:
  - FY 2021: Integrated Computational and Data Infrastructure (ICDI)
    - Experimental/Computational/Computer Science collaborations
    - Intelligent Distributed Infrastructure Simulation Capabilities
  - FY 2023: Distributed Resilient Systems
    - Scalable System Modeling
    - Adaptive Management and Partitioning of Resources

Other portfolios have individual projects which are strongly aligned with IRI objectives, for example:

- FY 2021: Data Reduction for Science – *The next presentation will discuss one such project in detail.*

Fast ML for science and the *extreme* edge

Nhan Tran (Fermilab) on behalf of our "extreme edge" data reduction team
Team PIs: J. Agar (Drexel), J. Duarte (UCSD), A. Gholami (ICSI/UCB), P. Harris (MIT), R. Kastner (UCSD), M. Mahoney (ICSI/LBNL), J. Ngadiuba (Fermilab) and the Fast ML for Science community

"Scientific discoveries come from groundbreaking ideas and the capability to validate those ideas by testing nature at new scales - finer and more precise temporal and spatial resolution. This is leading to an explosion of data that must be interpreted, and ML is proving a powerful approach. The more efficiently we can test our hypotheses, the faster we can achieve discovery. To fully unleash the power of ML and accelerate discoveries, it is necessary to embed it into our scientific process, into our instruments and detectors."

Applications and Techniques for Fast Machine Learning in Science
https://doi.org/10.3389/fdata.2022.787421

Edge capabilities and interface/synergy with HPC critical to develop within IRI objectives

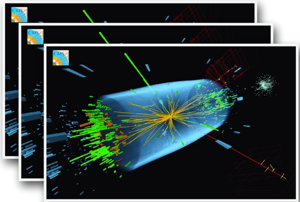

CMS Experiment

40MHz collision rate

~1B detector channels

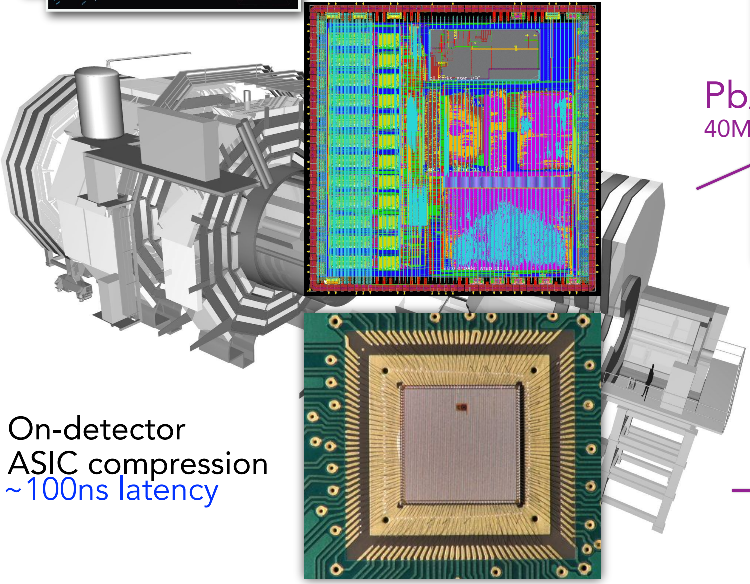

On-detector
ASIC compression
~100ns latency

Pb/
40MH

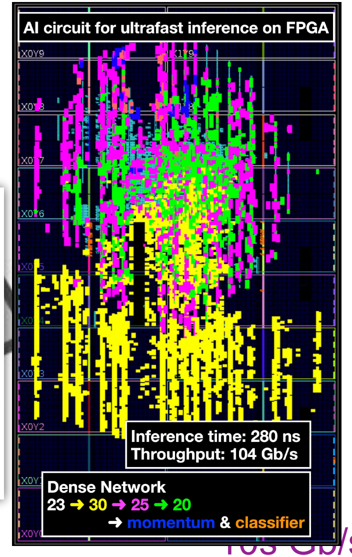

10s Gb/s
~5 kHz

10s Tb/s
100s kHz

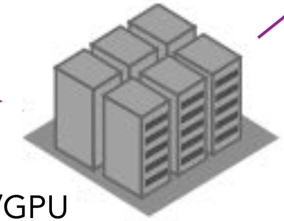

On-prem CPU/GPU
filter farm
~100 ms latency

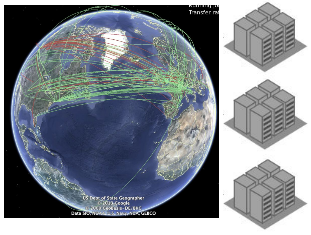

Worldwide
computing grid
Exabyte-scale
datasets

Comparable to global internet traffic bandwidth

The need for speed

<u>Aligned with IRI Science Patterns:</u>
Time-sensitive pattern
Data integration-intensive pattern

↓

Benchmarks bring innovation

The Fast ML for Science community aims to bring **seemingly different domains** together to develop **techniques, tools, and platforms** for challenges that far outpace industry.

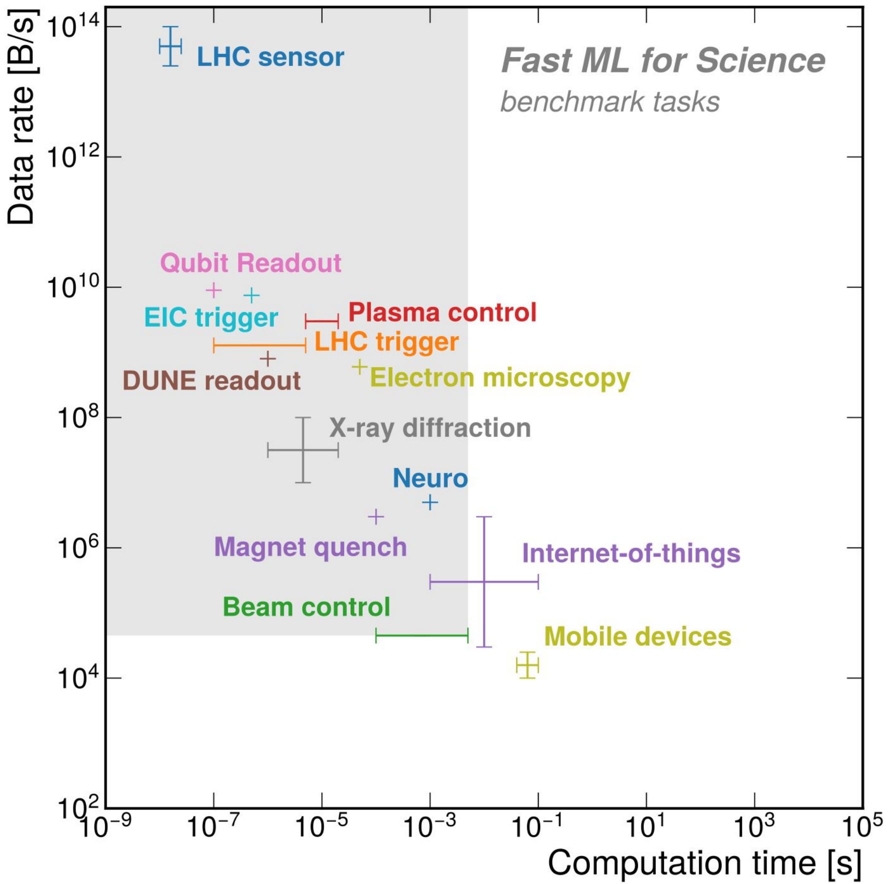

The need for speed

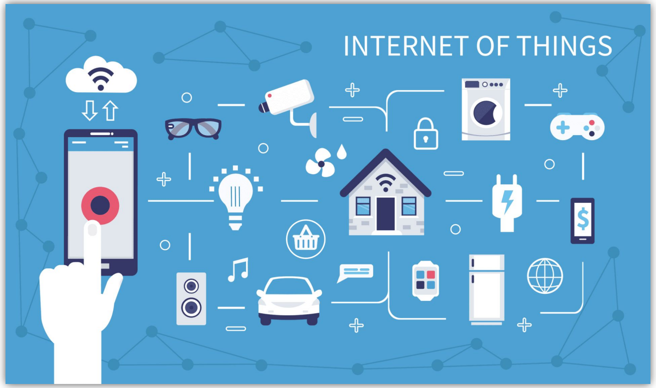

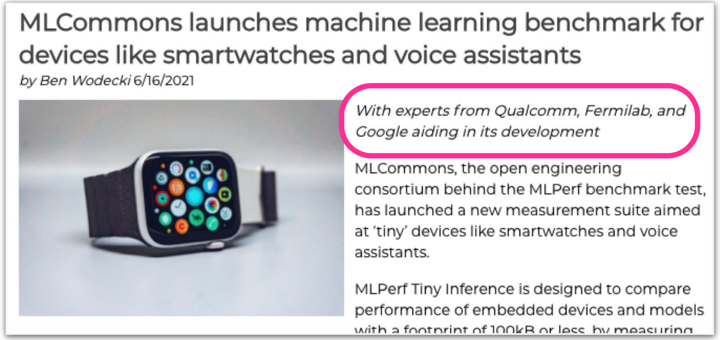
MLCommons launches machine learning benchmark for devices like smartwatches and voice assistants

by Ben Wodecki 6/16/2021

With experts from Qualcomm, Fermilab, and Google aiding in its development

MLCommons, the open engineering consortium behind the MLPerf benchmark test, has launched a new measurement suite aimed at 'tiny' devices like smartwatches and voice assistants.

MLPerf Tiny Inference is designed to compare performance of embedded devices and models with a footprint of 100kB or less, by measuring

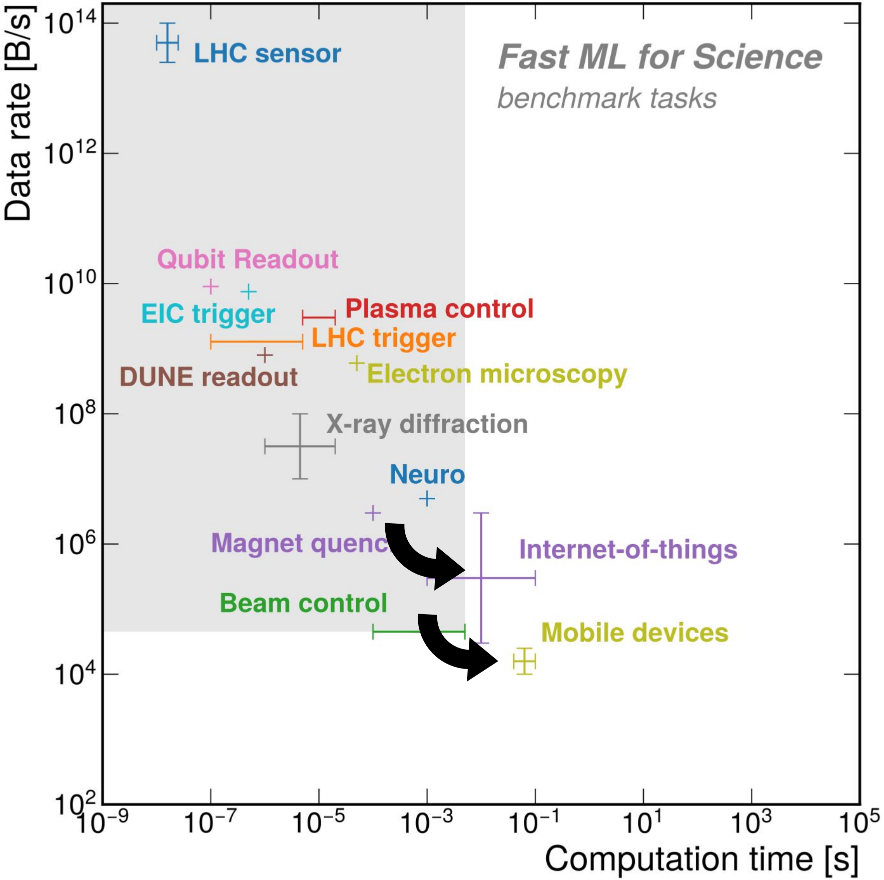

The need for speed

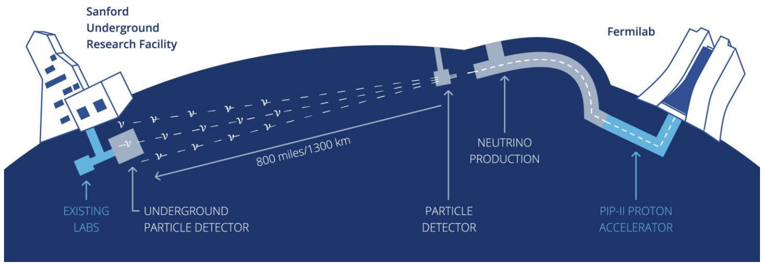

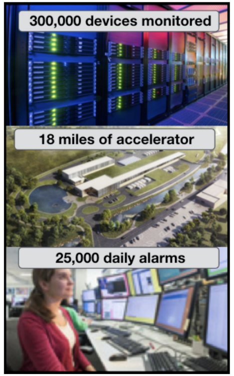

Fermilab accelerator complex

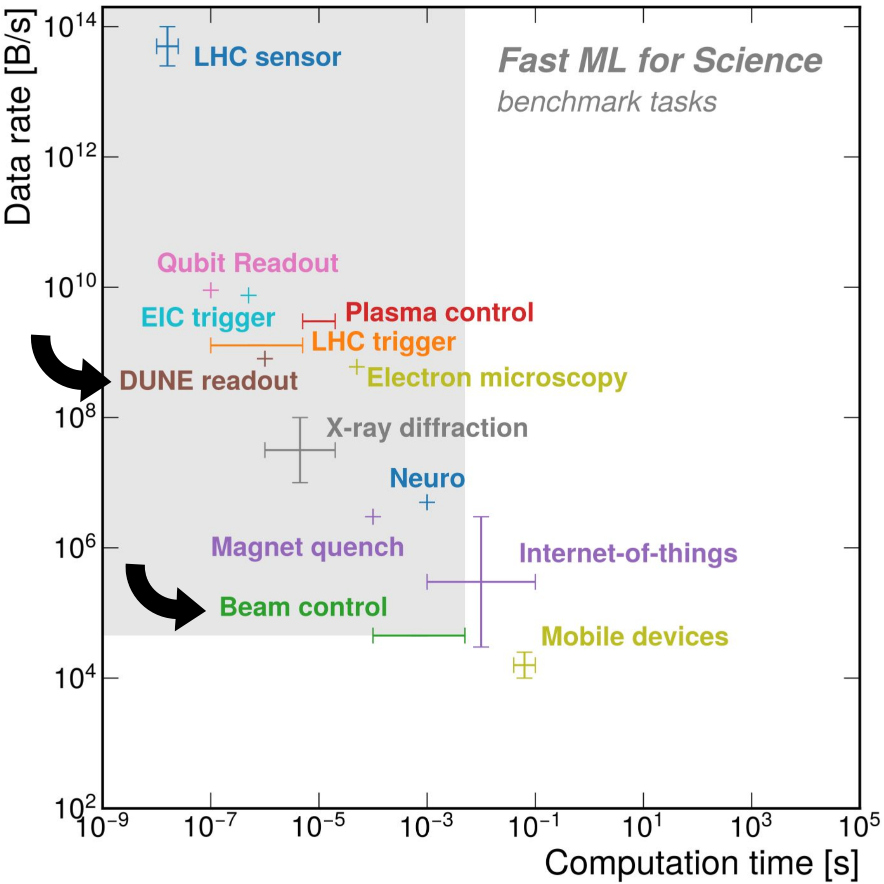

The need for speed

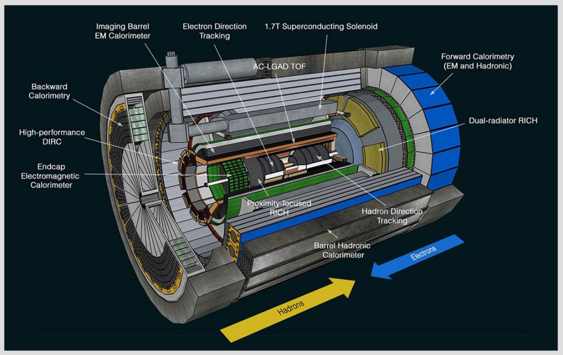

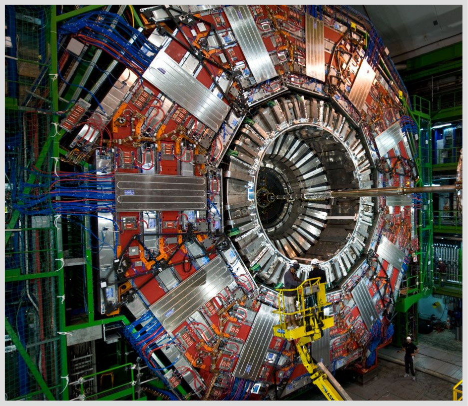

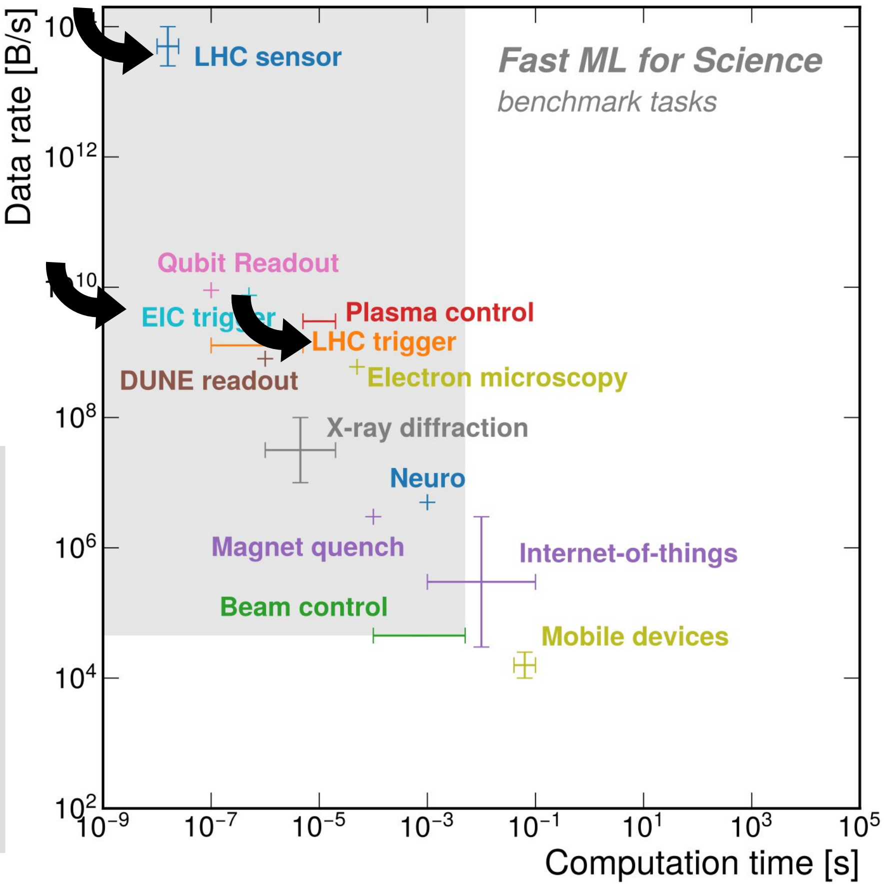

The need for speed

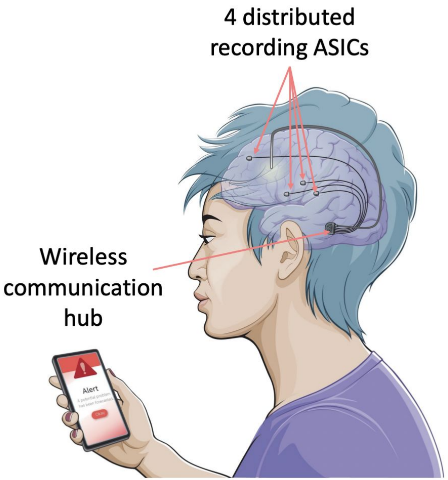

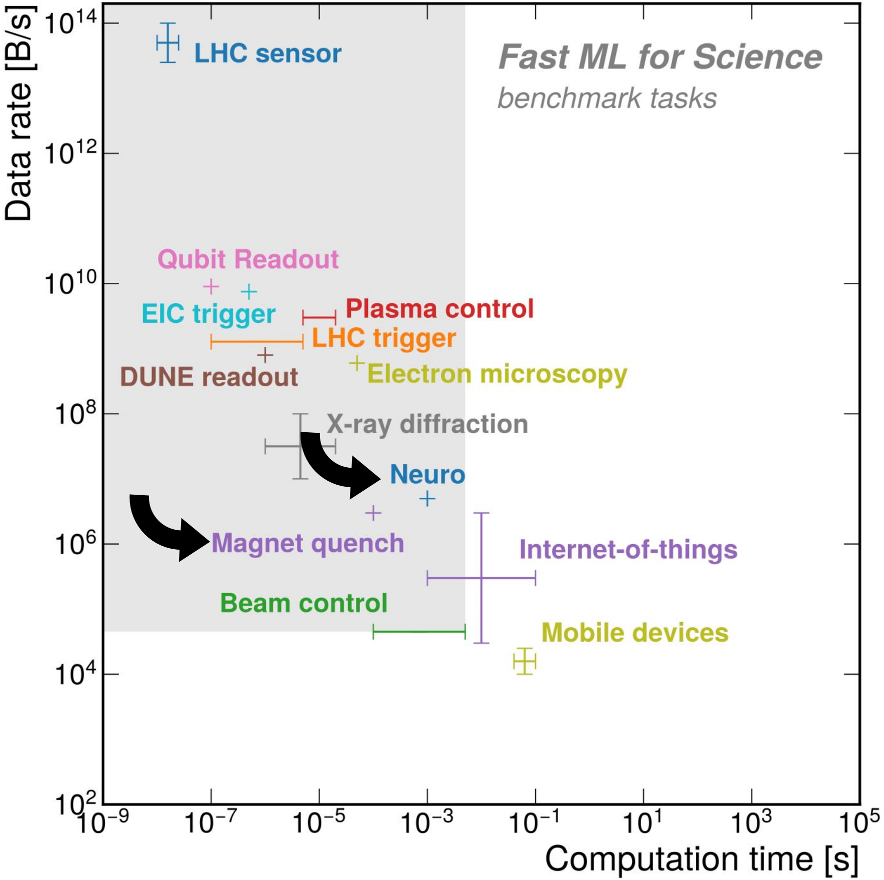

The need for speed

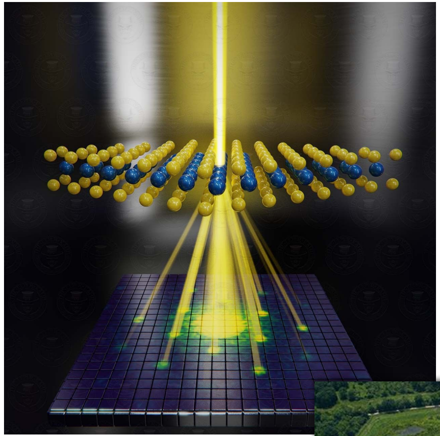

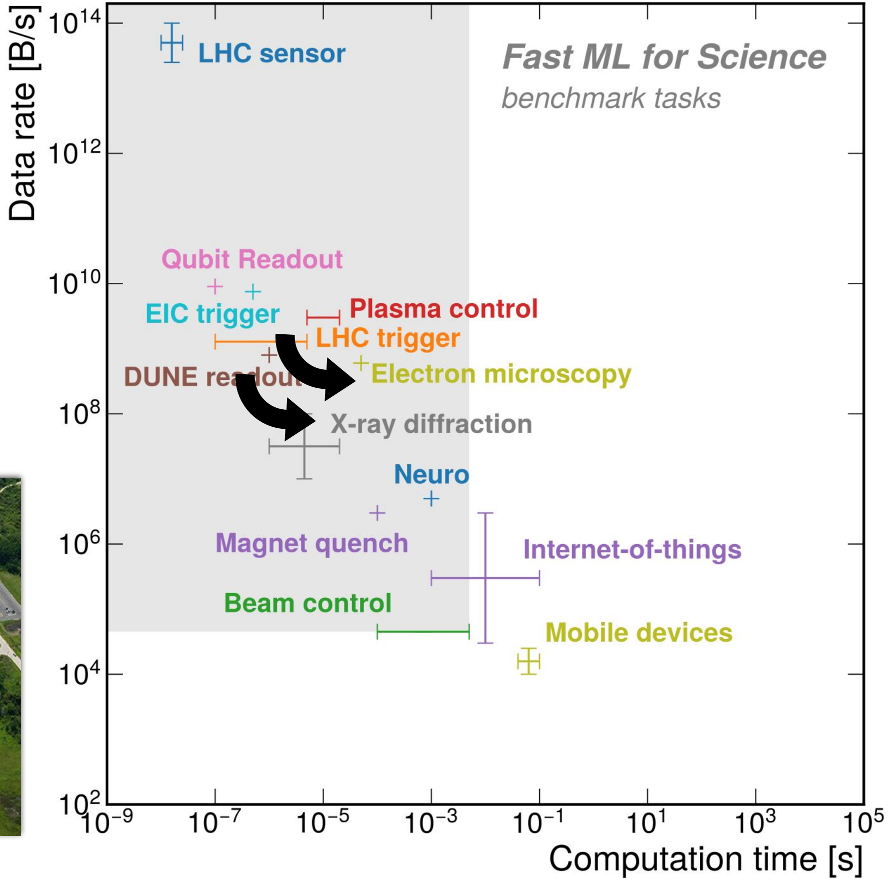

The need for speed

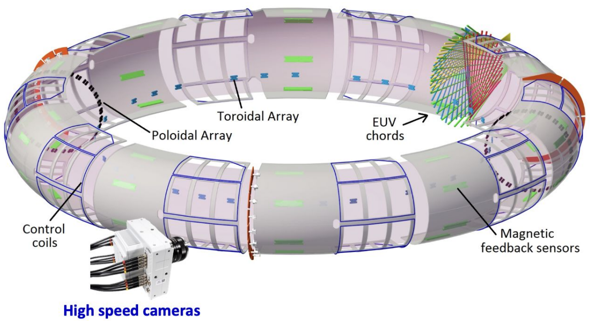

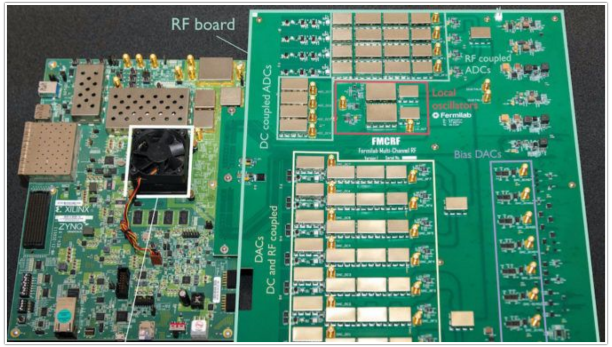

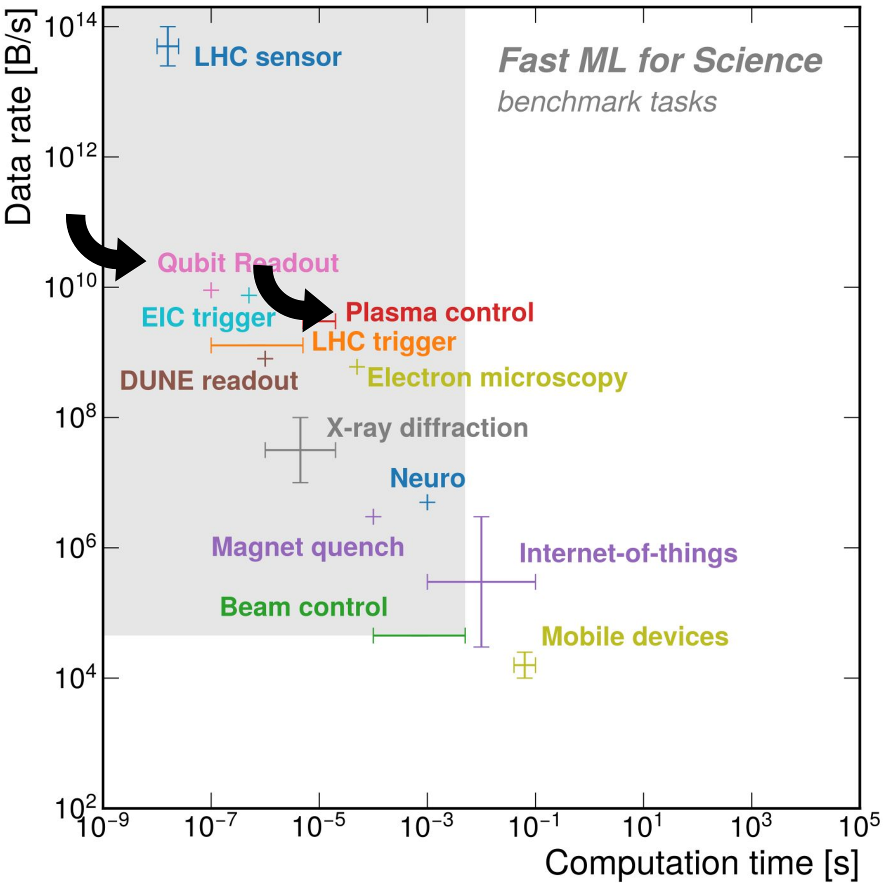

Tools and techniques for the extreme edge

<u>Aligned with IRI Practice Areas:</u>
User-experience practice
Resource co-operations practice
Workflows, interfaces, & automation practice
Portable/scalable solutions practice

↓

Efficient codesign
^
methods, tools, workflows, collabs, and support for

https://arxiv.org/abs/2304.06745
https://arxiv.org/abs/2306.07629

Efficient codesign
methods, tools, workflows, collabs, and support for

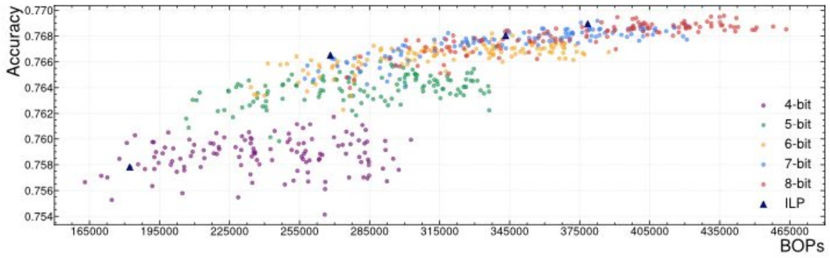

- Hessian-aware quantization with ILP solver provides optimal solutions without brute force scans
- Circuit-level deployment with hls4ml and QONNX toolflow

- New methods for sparse/dense matrix decomposition of greatly reduce compute resources

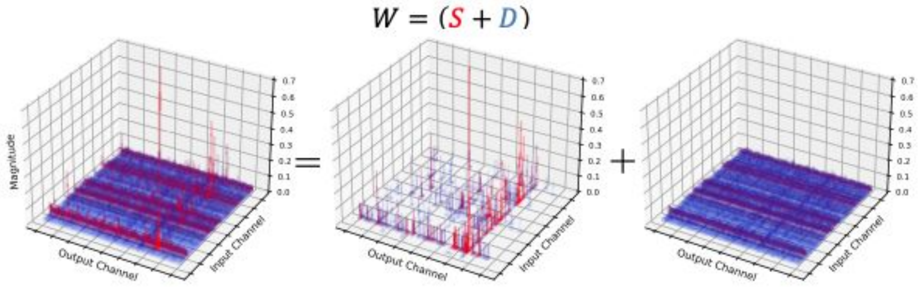

Efficient codesign
methods, tools, workflows, collabs, and support for

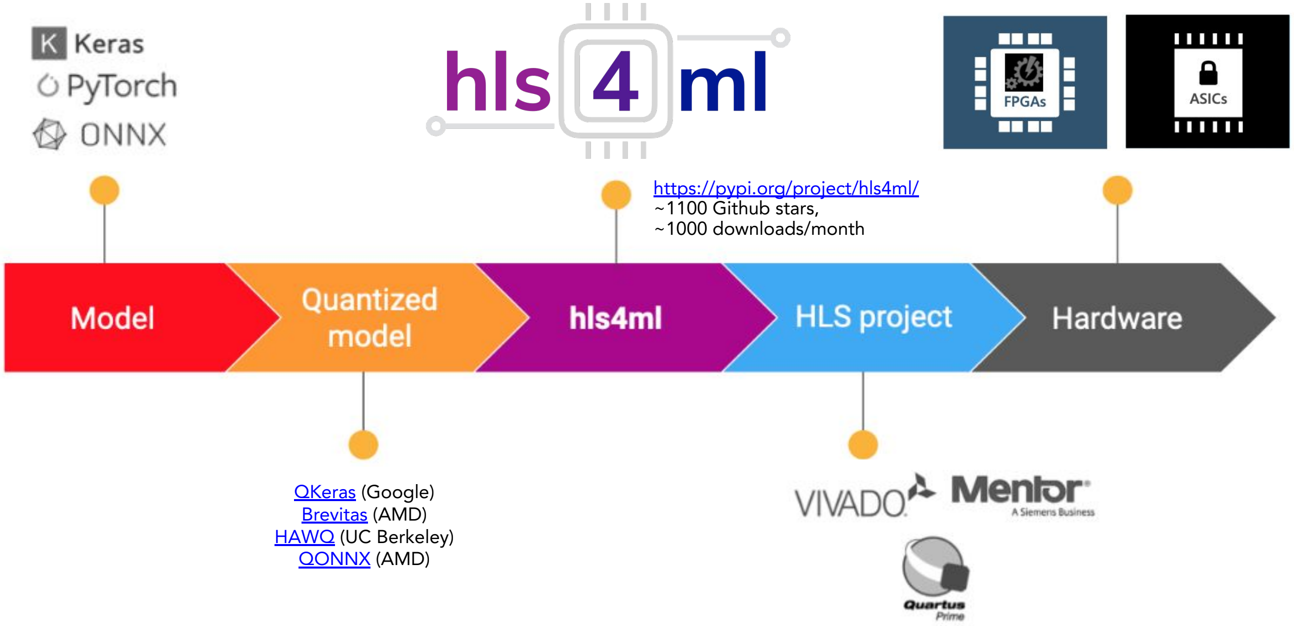

https://newsroom.sw.siemens.com/en-US/siemens-catapult-ai-nn/

Efficient codesign

^

methods, tools, workflows, collabs, and support for

yahoo!finance

Search for news, symbols or companies

My Portfolio News Markets Sectors Screeners Personal Finance Videos

Siemens simplifies development of AI accelerators for advanced system-on-chip designs with Catapult AI NN

PR Newswire
Tue, May 21, 2024, 8:00 AM CDT • 5 min read

In This Article:

SIEGY -0.78% SMAWF +0.35%

- **Catapult AI NN offers software e… solution to synthesize AI Neural…**
- **Enables software development t… models designed in Python into… facilitating faster and more pow… to standard processors**

PLANO, Texas, May 21, 2024 /PRNew… Industries Software today announce… High-Level Synthesis (HLS) of neural… Application-Specific Integrated Circu… (SoCs). Catapult AI NN is a complete… network description from an AI frame… synthesizes it into an RTL accelerator… implementation in silicon.

Catapult AI NN brings together hls4ml, an open-source package for machine learning hardware acceleration, and Siemens' Catapult™ HLS software for High-Level Synthesis. Developed in close collaboration with Fermilab, a U.S. Department of Energy Laboratory, and other leading contributors to hls4ml, Catapult AI NN addresses the unique requirements of machine learning accelerator design for power, performance, and area on custom silicon.

Efficient codesign
methods, tools, workflows, collabs, and <u>support</u> for

- Open-source
- Community-supported
- User-driven
- Accessible and usable

Check performance

Check the accuracy and make a ROC curve

In [ ]:
```python
import plotting
import matplotlib.pyplot as plt
from sklearn.metrics import accuracy_score

y_keras = model.predict(X_test)
print("Accuracy: {}".format(accuracy_score(np.argmax(y_test, axis=1), np.argmax(y_keras, axis=1))))
plt.figure(figsize=(9, 9))
_ = plotting.makeRoc(y_test, y_keras, le.classes_)
```

Convert the model to FPGA firmware with hls4ml

Now we will go through the steps to convert the model we trained to a low-latency optimized FPGA firmware with hls4ml. First, we will evaluate its classification performance to make sure we haven't lost accuracy using the fixed-point data types. Then we will synthesize the model with Vivado HLS and check the metrics of latency and FPGA resource usage.

Make an hls4ml config & model

The hls4ml Neural Network inference library is controlled through a configuration dictionary. In this example we'll use the most simple variation, later exercises will look at more advanced configuration.

In [ ]:
```python
import hls4ml

config = hls4ml.utils.config_from_keras_model(model, granularity='model')
print("-----------------------------------")
print("Configuration")
plotting.print_dict(config)
print("-----------------------------------")
hls_model = hls4ml.converters.convert_from_keras_model(
    model, hls_config=config, output_dir='model_1/hls4ml_prj', part='xcu250-figd2104-2L-e'
)
```

Let's visualise what we created. The model architecture is shown, annotated with the shape and data types

In [ ]:
```python
hls4ml.utils.plot_model(hls_model, show_shapes=True, show_precision=True, to_file=None)
```

Compile, predict

Now we need to check that this model performance is still good. We compile the hls_model, and then use `hls_model.predict` to execute the FPGA firmware with bit-accurate emulation on the CPU.

In [ ]:
```python
hls_model.compile()
X_test = np.ascontiguousarray(X_test)
y_hls = hls_model.predict(X_test)
```

Accessible through high-level Python APIs

Tutorials, demos, and resources for non-experts

Community of developers and experts to help https://dl.acm.org/doi/10.1145/3665334

Robust & Efficient codesign

^ methods, tools, workflows, collabs, and support for

Extreme edge applications dealing with raw & unfiltered data requires **robust data processing**

Questions being answered:

- Which bits should be protected from single event upsets (due to radiation)?
- How does quantization effect model robustness?
- How can we make models insensitive to distribution shift (e.g. noise) a priori?
- Can we visualize important properties of ensembles for performance?

Also in collaboration with Visualization (VisSciML) team!

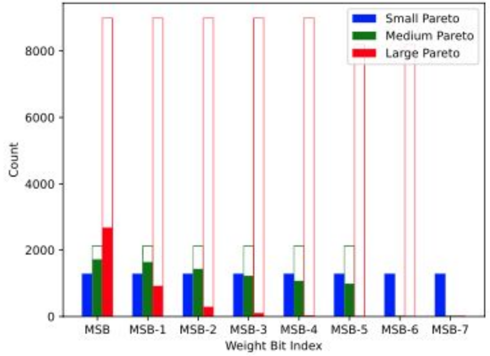

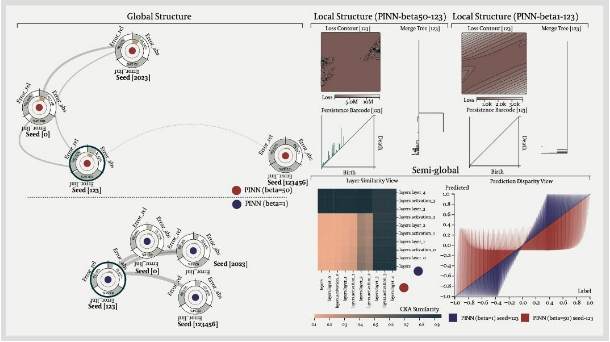

https://arxiv.org/abs/2312.00128

Cross-domain impact: plasma control

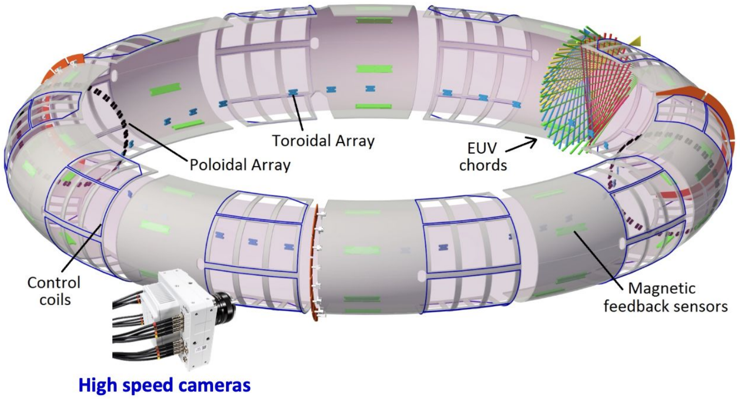

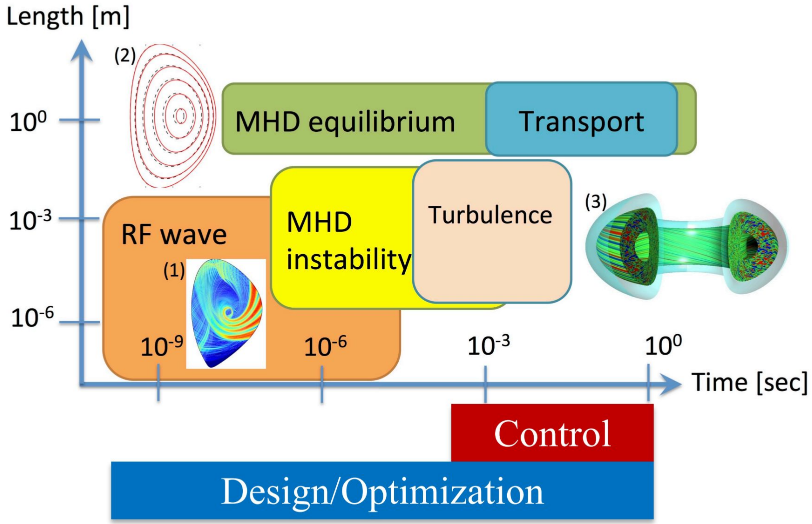

Cross-domain impact: plasma control

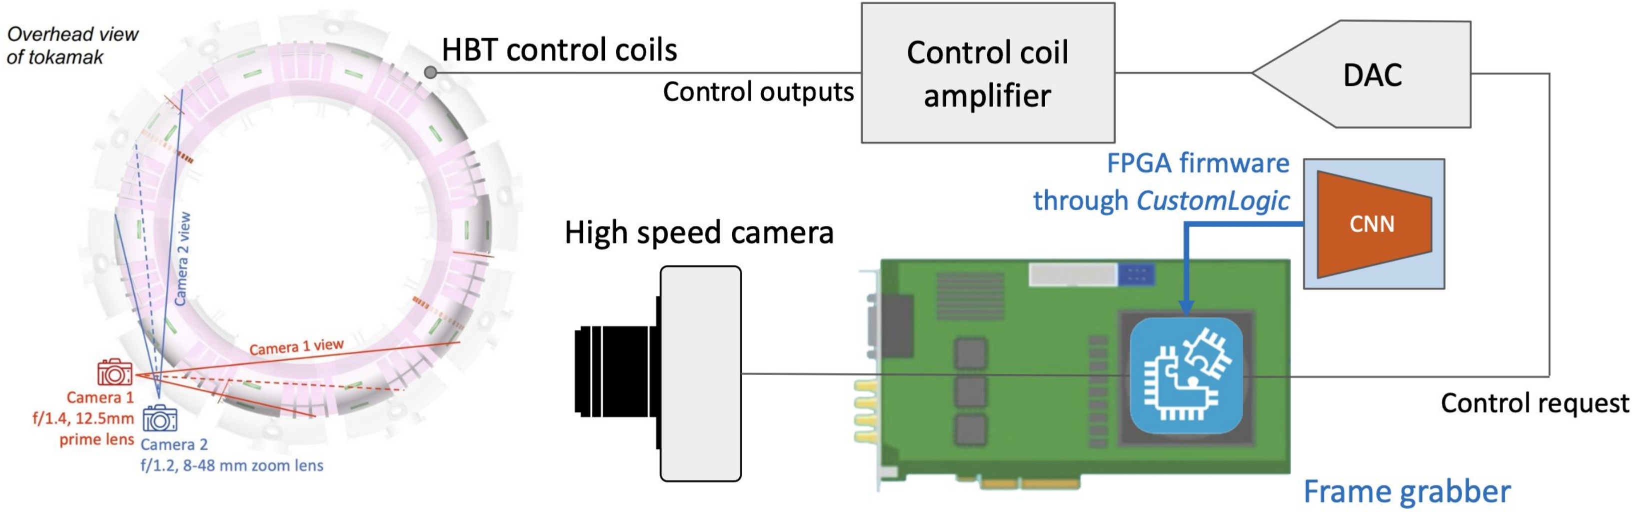

Cross-domain impact: plasma control

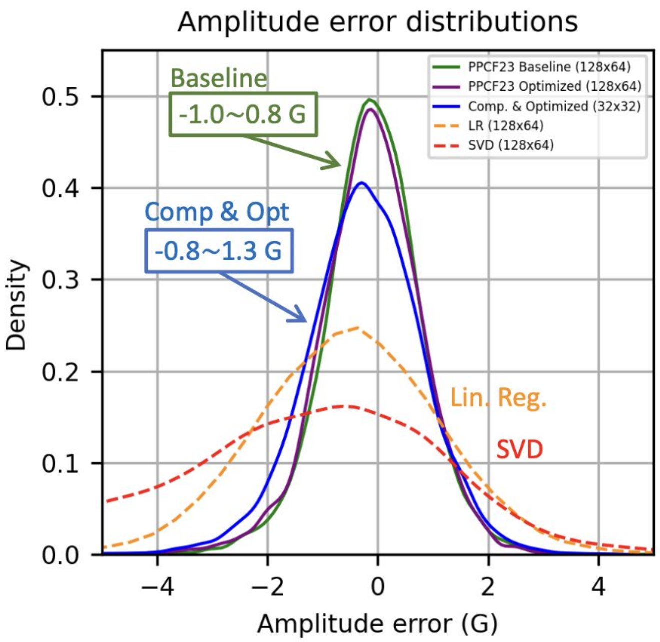

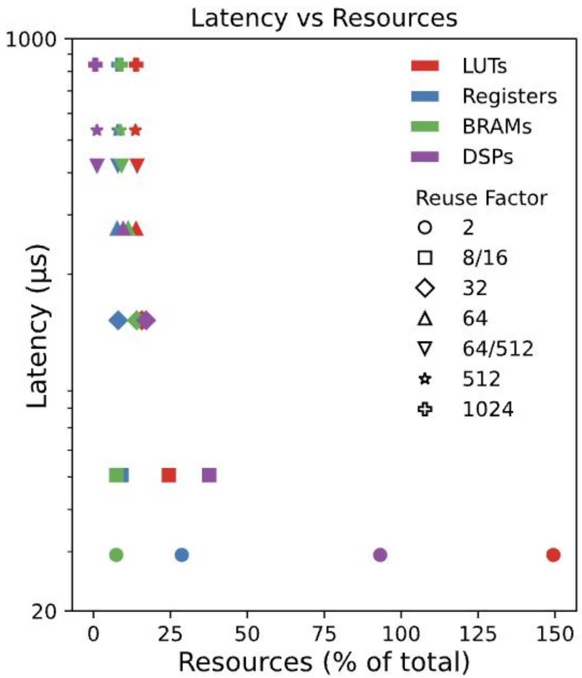

Cross-domain impact: plasma control

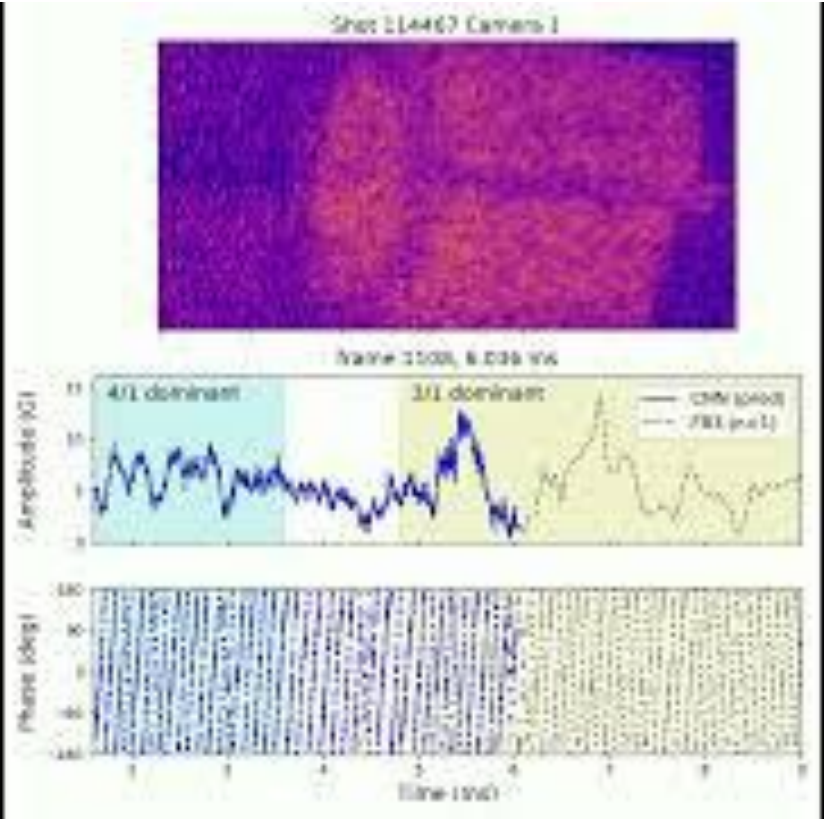

120 kfps throughput
17.6 μs latency

– enabling new capabilities for fusion experiments

IRI Program launch is a DOE FY24-25 Agency Priority Goal
ASCR is implementing IRI through these four major elements

1 Invest in IRI foundational infrastructure

ASCR Facilities are architected to support IRI

- KPPs , benchmarks and readiness programs are explicitly designed for IRI

2 Stand up the IRI Program governance and technical activities

The IRI Management council has been established, and is working to:

- Define roles, responsibilities and reporting structures
- Launch committees to coordinate technical and outreach work

3 Bring IRI projects into formal coordination

We are working closely to coordinate existing work and forge new collaborations

- Roadmap for partnering with science teams
- Coordinating with CS research programs

4 Deploy an IRI Science Testbed across the ASCR Facilities

Shared development environment for IRI developers and pilot applications, designed to accelerate transition to operations

Summary

The IRI program is moving out of the initial stages of development and process-setting

- Launching first phases of science engagements and technical work
- Building organizational structure to manage IRI efforts and integrate across program offices
- Connecting to research programs
- Developing partnerships with pathfinders, and laying out the pathways to engage more broadly

IRI will provide a foundational infrastructure for next-generational experimental science, post-exascale computing, AI/FASST and other initiatives.

IRI is moving!

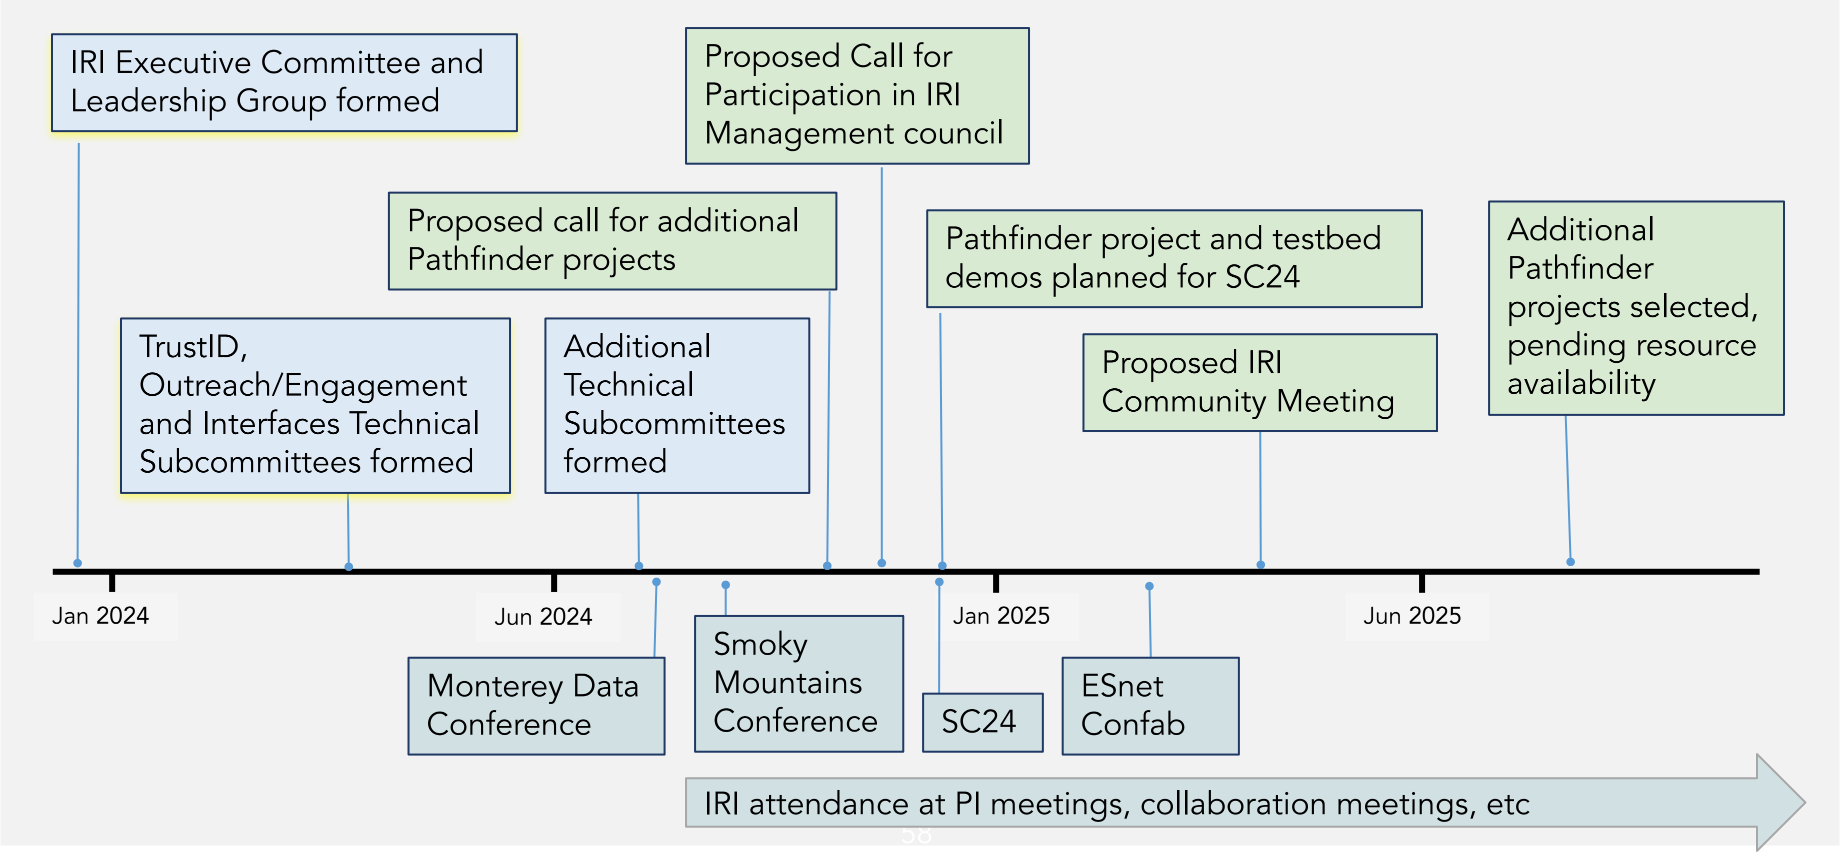

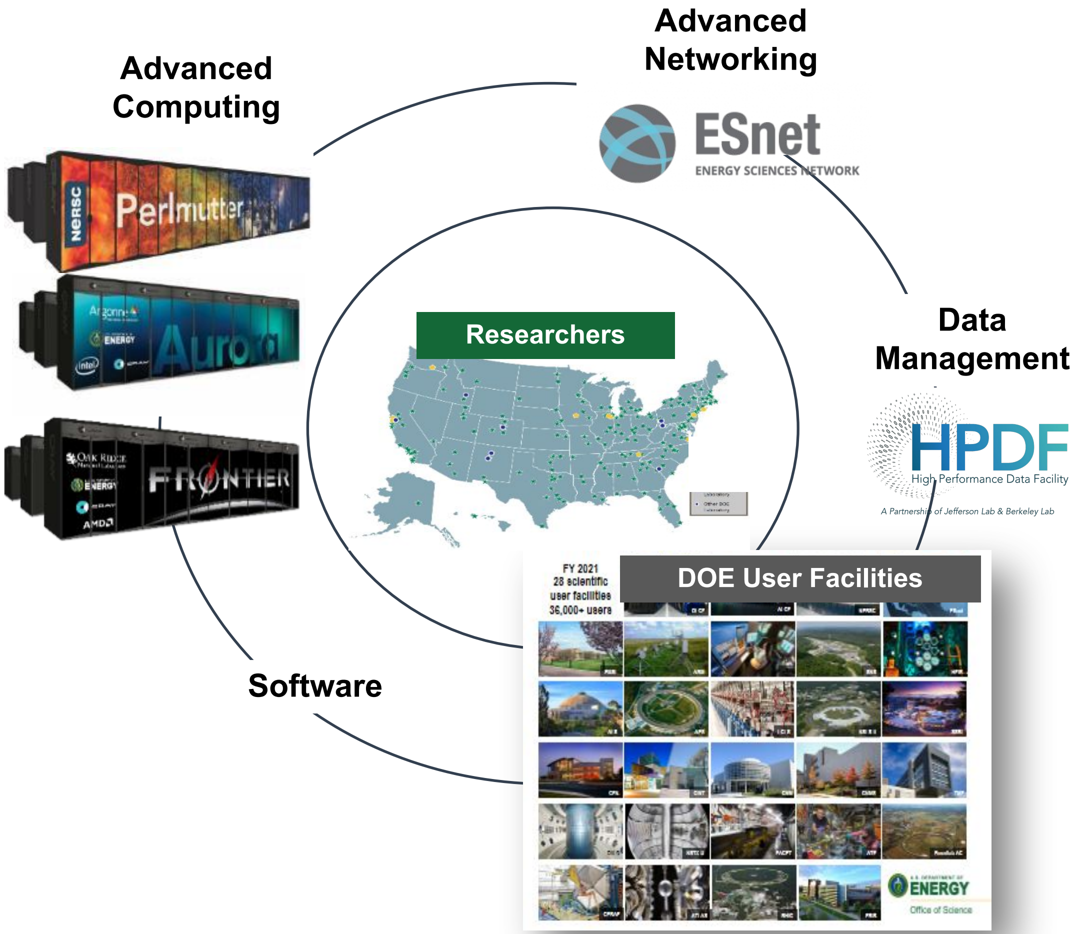
Advanced Computing
Advanced Networking
ESnet
ENERGY SCIENCES NETWORK
NERSC
Perlmutter
Aurora
FRONTIER
Researchers
Data Management
HPDF
High Performance Data Facility
A Partnership of Jefferson Lab & Berkeley Lab
FY 2021
28 scientific
user facilities
36,000+ users
DOE User Facilities
ENERGY
Office of Science
Software

Thanks to our speakers!

Argonne National Laboratory | Argonne Leadership Computing Facility

Mike Papka
Tom Uram

ESnet Energy Sciences Network

Eli Dart
Chin Guok
Inder Monga

Fermilab

Nhan Tran

HPDF High Performance Data Facility — A Partnership of Jefferson Lab & Berkeley Lab

Amber Boehnlein
Graham Hayes

NERSC National Energy Research Scientific Computing Center

Debbie Bard
Tiffany Connors
Sudip Dosanjh

Oak Ridge National Laboratory | Leadership Computing Facility

Ryan Adamson
Rafael Ferreira da Silva
Sarp Oral
Arjun Shankar

U.S. Department of Energy | Office of Science

Ben Brown
Hal Finkel