
# HEP Community Organization for AI

**Tulika Bose, University of Wisconsin-Madison**
**June 29th, 2026**

# “Building an AI-Native Research Ecosystem” Whitepaper

arXiv > hep-ex > arXiv:2602.17582

Search... | All fields | Search

Help | Advanced Search

**High Energy Physics – Experiment**

*[Submitted on 19 Feb 2026]*

## Building an AI–native Research Ecosystem for Experimental Particle Physics: A Community Vision

Thea Klaeboe Aarrestad, Alaa Abdelhamid, Haider Abidi, Jahred Adelman, Jennifer Adelman-McCarthy, Shuchin Aeron, Garvita Agarwal, Usman Ali, Cristiano Alpigiani, Omar Alterkait, Mohamed Aly, Oz Amram, Saeed Ansari Fard, Aram Apyan, John Arrington, Marvin Ascencio-Sosa, Mohammad Atif, Aneesha Avasthi, Muhammad Bilal Azam, Bhim Bam, Joshua Barrow, Rainer Bartoldus, Amit Bashyal, Aashwin Basnet, Ayse Bat, Lothar A. T. Bauerdick, John Beacom, Chris Bee, Michael Begel, Matthew Bellis, Rene Bellwied, Rakitha Beminiwattha, Gabriele Benelli, Douglas Benjamin, Catrin Bernius, Binod Bhandari, Avinay Bhat, Meghna Bhattacharya, Saptaparna Bhattacharya, Prajita Bhattarai, Sudip Bhattarai, Wahid Bhimji, Jianming Bian, Burak Bilki, Mary Bishai, Kevin Black, Kenneth Bloom, Brian Bockelman, Johan Sebastian Bonilla Castro, Tulika Bose, Nilay Bostan, Othmane Bouhali, Dimitri Bourilkov, Dominic Brailsford, Gustaaf Brooijmans, Elizabeth Brost, Maria Brigida Brunetti, Quentin Buat, Brendon Bullard, Jackson Burzynski, Paolo Calafiura, Rodolfo Capdevilla, Fabian Andres Castaño Usuga, Raquel Castillo Fernandez, Fabio Catalano, Viviana Cavaliere, Flavio Cavanna, Giuseppe Cerati, Aidan Chambers, Maria Chamizo-Llatas, Philip Chang, Andrew Chappell, Arghya Chattopadhyay, Sergei Chekanov, Jian-ping Chen, Yi Chen, Zhengyang Chen, J. Taylor Childers, Hector Chinchay, Yuan-Tang Chou, Tasnuva Chowdhury, Neil Christensen, Wonyong Chung, Rafael Coelho Lopes de Sa, Simon Corrodi, Kyle Cranmer, Matteo Cremonesi, Roy Cruz, Mate Csanad, Mariarosaria D'Alfonso, Carlo Dallapiccola, Daine Danielson, Sridhara Dasu, Gavin Davies, Kaushik De, Patrick de Perio, Klaus Dehmelt, Marco Del Tutto, Carlos Ruben Dell'Aquila, Sarah Demers et al. (359 additional authors not shown)

Experimental particle physics seeks to understand the universe by probing its fundamental particles and forces and exploring how they govern the large-scale processes that shape cosmic evolution. This whitepaper presents a vision for how Artificial Intelligence (AI) can accelerate discovery in this field. We outline grand challenges that must be addressed to enable transformative breakthroughs and describe how current and planned experimental facilities can implement this vision to advance our understanding of the vast and complex physical world from the smallest to the largest scales. We show how facilities currently under construction, such as the HL-LHC, DUNE and soon EIC, can both benefit from and serve as proving grounds for this vision, while also enabling a longer-term goal for how future experiments -- like FCC-ee at CERN, IceCube-Gen2, a Muon Collider in the U.S., and smaller to mid-scale projects -- can be fully AI-native. We describe how a truly national-scale collaboration, jointly managed across large funding partners, and involving both DOE laboratories and universities, can make this happen.

**Access Paper:**

- View PDF
- HTML (experimental)
- TeX Source



Current browse context:
hep-ex
< prev | next >
new | recent | 2026-02

References & Citations
- INSPIRE HEP
- NASA ADS
- Google Scholar
- Semantic Scholar

Export BibTeX Citation

Bookmark

**More than 500 people from 105 universities and 7 national labs from the U.S. signed the paper (40 U.S. states/territories, including 16 EPSCoR jurisdictions)**

https://arxiv.org/abs/2602.17582

Community vision put together in close collaboration with the new DPF Coordinating Panel for Software & Computing (CPSC)

# “AI-Native”: From Design to Discovery

**Four inter-connected Grand Challenges****:**

- **Accelerated design informs intelligent instrumentation**
- Intelligent sensing enables autonomy
- Autonomous operation feeds higher-quality data into inference
- From data to discovery closes the loop back to design
- *Together, these challenges define an end-to-end **AI-native experimental paradigm**.*

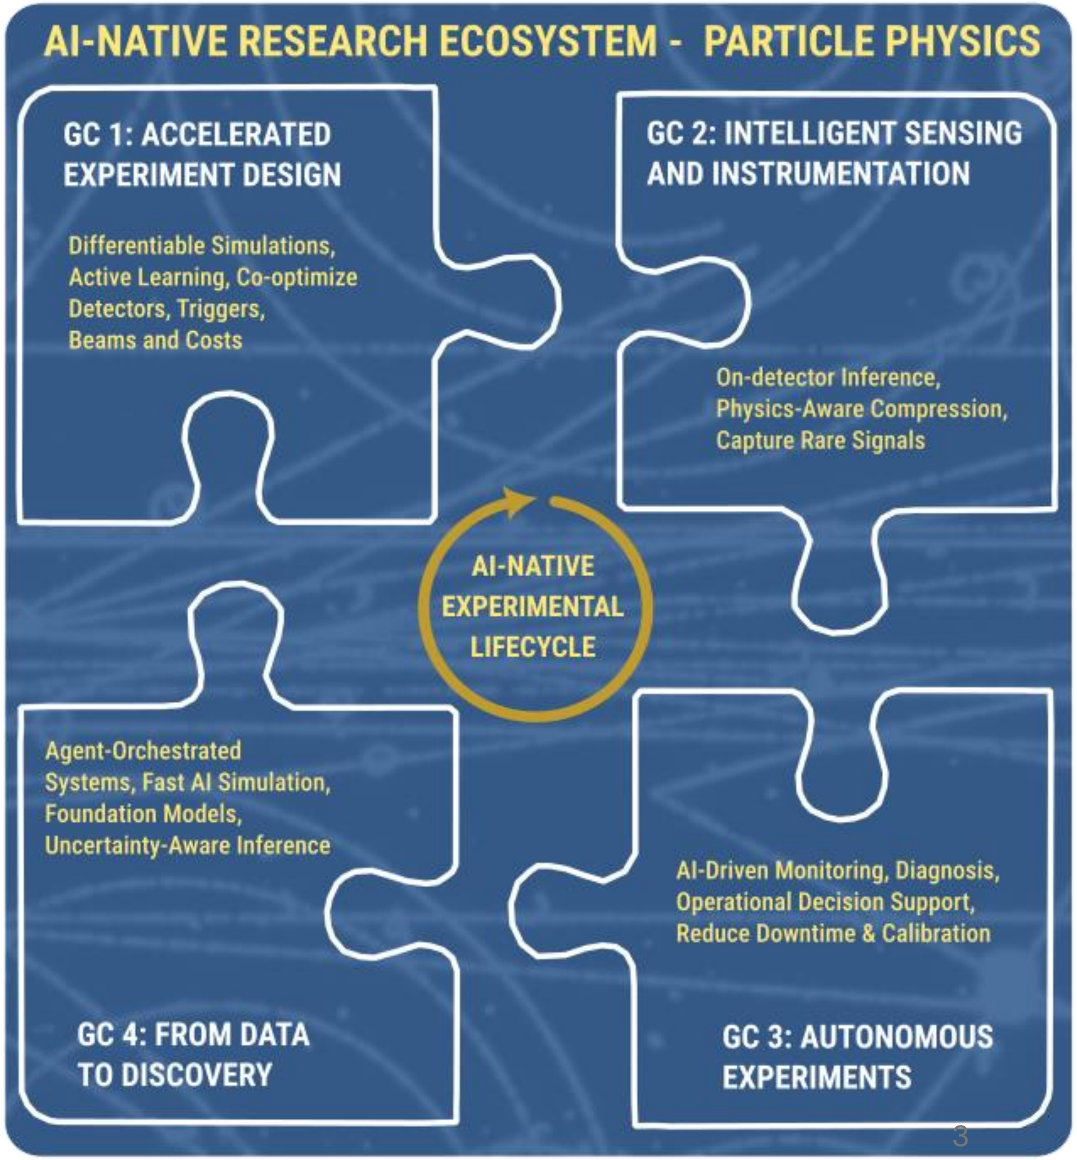


**Plus some cross-cutting themes and emerging opportunities

# The case for collaboration

Fulfilling this ambitious vision requires a **national-scale collaboration** that can:

- Bring together expertise across labs and universities,
- pool computing and data services
- align with emerging national infrastructure
- scale workforce development

enabling each experiment/project to benefit from shared tools, sustained support, and the ability to marshal greater collective effort on every project.

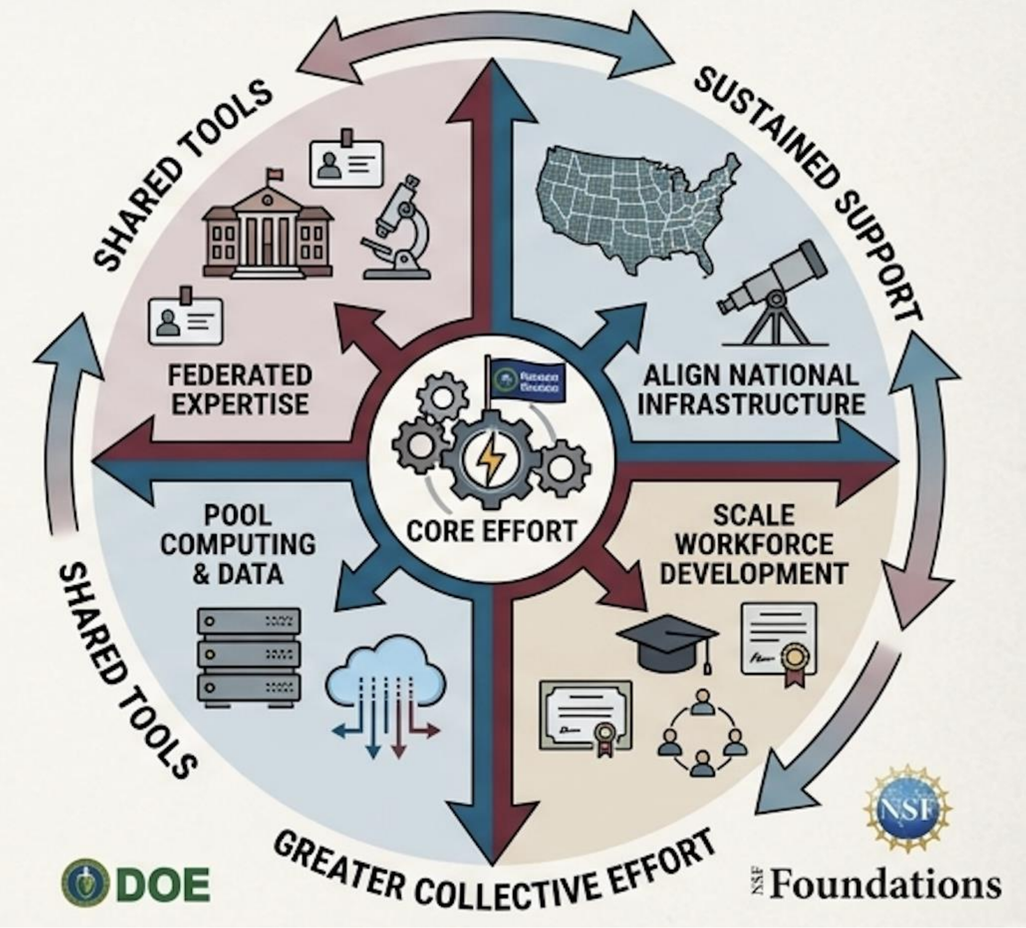


*We proposed a* **national-scale, multi-institutional collaboration that unites universities, laboratories, and industry partners**. Benefiting from *multiple large funding partners -- DOE (Genesis Mission), NSF, and private foundations.*
*The successful U.S. LHC operations programs can be used as a model*

# Blueprint Meeting

Meeting organized on May 18-19 in Washington, DC

~60 in-person participants

~120 Remote attendees

Strong international participation (CERN, INFN, UK)

Wide-ranging discussion of potential areas of collaboration

Indico page: [https://indico.flatironinstitute.org/event/4120/](https://indico.flatironinstitute.org/event/4120/)

**Organizers: P. Elmer, I. Fisk, V. Martinez Outschoorn & T.B.**

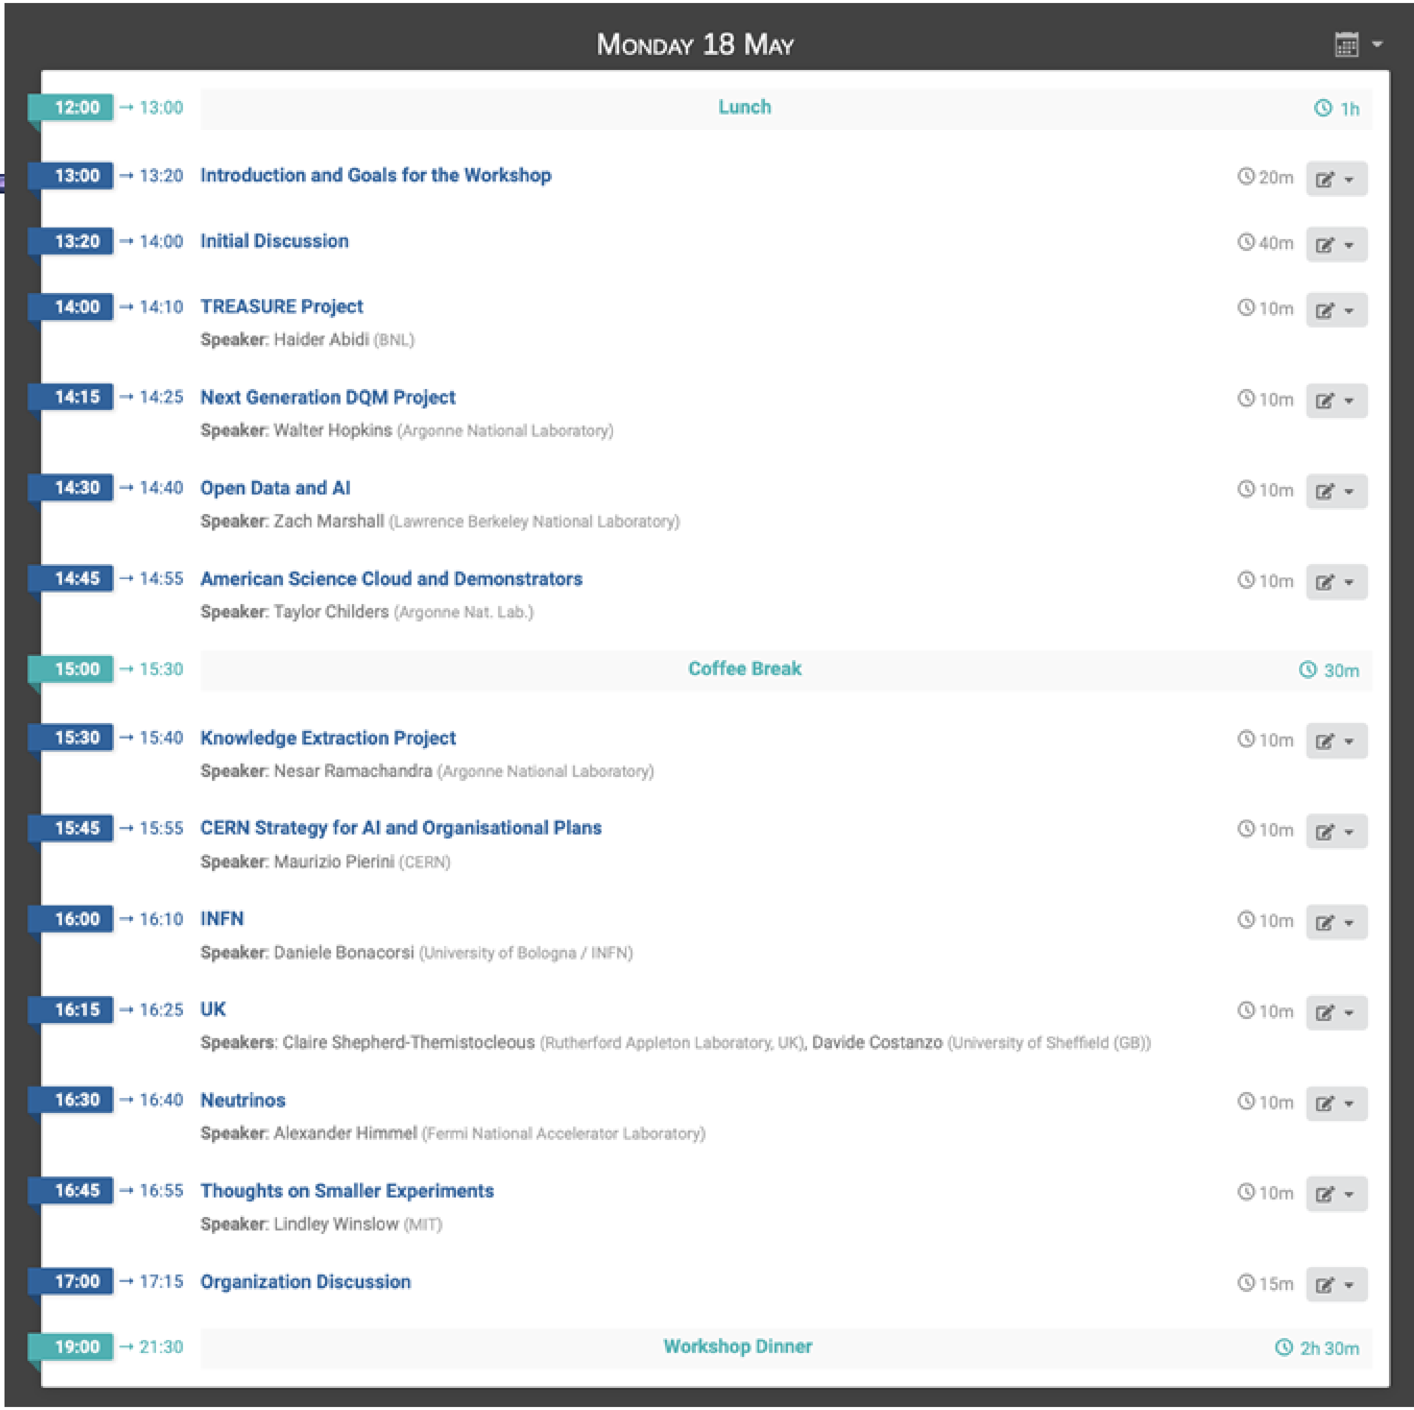

Monday 18 May

| | | | |
|---|---|---|---|
| 12:00 | 13:00 | Lunch | 1h |
| 13:00 | 13:20 | Introduction and Goals for the Workshop | 20m |
| 13:20 | 14:00 | Initial Discussion | 40m |
| 14:00 | 14:10 | TREASURE Project<br>Speaker: Haider Abidi (BNL) | 10m |
| 14:15 | 14:25 | Next Generation DQM Project<br>Speaker: Walter Hopkins (Argonne National Laboratory) | 10m |
| 14:30 | 14:40 | Open Data and AI<br>Speaker: Zach Marshall (Lawrence Berkeley National Laboratory) | 10m |
| 14:45 | 14:55 | American Science Cloud and Demonstrators<br>Speaker: Taylor Childers (Argonne Nat. Lab.) | 10m |
| 15:00 | 15:30 | Coffee Break | 30m |
| 15:30 | 15:40 | Knowledge Extraction Project<br>Speaker: Nesar Ramachandra (Argonne National Laboratory) | 10m |
| 15:45 | 15:55 | CERN Strategy for AI and Organisational Plans<br>Speaker: Maurizio Pierini (CERN) | 10m |
| 16:00 | 16:10 | INFN<br>Speaker: Daniele Bonacorsi (University of Bologna / INFN) | 10m |
| 16:15 | 16:25 | UK<br>Speakers: Claire Shepherd-Themistocleous (Rutherford Appleton Laboratory, UK), Davide Costanzo (University of Sheffield (GB)) | 10m |
| 16:30 | 16:40 | Neutrinos<br>Speaker: Alexander Himmel (Fermi National Accelerator Laboratory) | 10m |
| 16:45 | 16:55 | Thoughts on Smaller Experiments<br>Speaker: Lindley Winslow (MIT) | 10m |
| 17:00 | 17:15 | Organization Discussion | 15m |
| 19:00 | 21:30 | Workshop Dinner | 2h 30m |

# Goals of the workshop

**A central focus of the workshop will be discussions on how such a national-scale collaboration could be effectively bootstrapped and organized, as well as identifying priorities over the near and medium term. Such a collaboration, including both universities and DOE laboratories, has a number of potential benefits and opportunities, including:**

- a forum for **develop**ing **a shared roadmap** toward an AI-native research ecosystem
- a platform for enabling larger-scale efforts while **integrating expertise and resources from across the community**
- clear pathways to scale and transition promising **R&D activities towards deployment within experiments**
- a structure to develop and provide the services needed for the experiments to fully **leverage evolving national infrastructure** (e.g., AmSC, NAIRR) and large-scale data resources
- a means for **fostering both collaboration and healthy competition** through the development of shared benchmarks and the development and curation of the supporting datasets
- **accessible entry points for individual university groups** to engage and contribute
- a coherent framework for **cross-experiment activities**, enabling solutions with impact to deliver benefit to multiple experiments
- a clear structure for **US engagement with international efforts**, including those at CERN, and other national and regional collaborations
- a community structure for **engagement with industry**
- a structure for **large-scale workforce development activities** leveraging both lab and university expertise

This first workshop will consider the potential impact of these opportunities, identify near-term priorities, and explore ways to enable additional initiatives as resources become available. **We expect this discussion to continue in subsequent community workshops in 2026.**

# Computing Resource Strategy in the AI era

**What's changing:**

- AI infrastructure ≠ traditional Worldwide LHC Grid (WLCG)
- GPUs / accelerators / external facilities dominate
- Facility priorities and access models shifting (HPC, cloud, AI centers)

**Implications/Key questions:**

- Dedicated CPU-based infrastructure won't scale
- Standard data/network/storage assumptions break
- Need to understand how training, tuning, inference, HPC, AI facilities, networks, storage, and data access will change HEP computing needs

**Near-term Actions:**

- Convene a community workshop and develop a white paper on future computing models, AI workflows, access to HPC/AI facilities, inference facilities, on-prem resources
- Pursue common interfaces + allocation strategies, shared development resources, and coordinated access etc

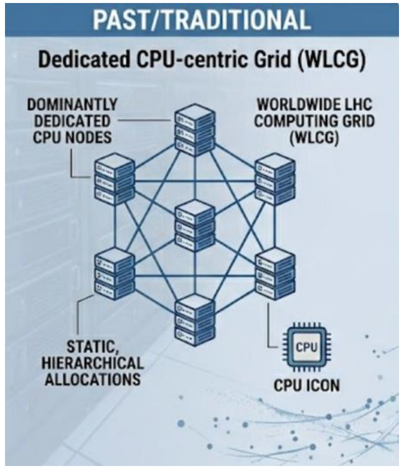


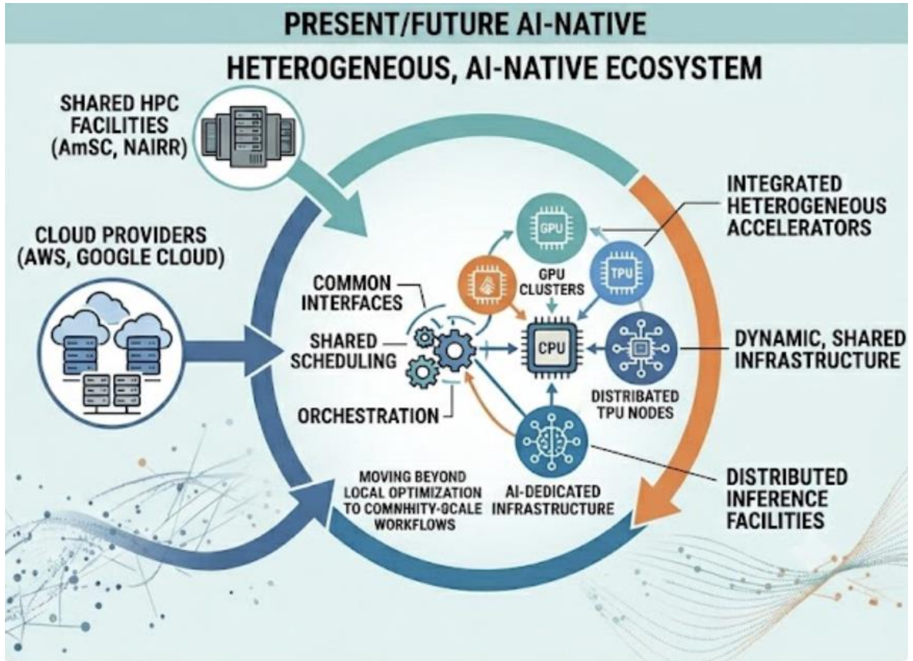


***Evolve in the direction of ecosystem-level computing***

# Agentic AI

**What's changing:**

- Automated agent-driven workflows being explored for Analysis, Operations, Coding, Hardware Design…

**Implications/key questions:**

- Need trusted benchmarks, validation procedures, reusable agents, common skills, and AI-ready knowledge repositories.
- Need to understand realistic use cases beyond demonstrations

**Near-term Actions:**

- Build an agentic tutorial repository, curate standard agents and skills,
- Organize tutorials, hackathons around already-defined problems,
- Develop benchmarks & validation procedures


*Agentic workflows should become reusable infrastructure*

# Open Data & Foundation Models in the AI era

**Open Data**

- Open data must evolve for the AI era, with better formats, metadata, documentation, provenance, and hosted access.
  - Need machine-readable documentation, integrated hardware/software/model-serving environments, workflow capture, usage monitoring, and harmonized experiment policies.

**Foundation Models**

- Scale (data + compute) exceeds individual experiments
  - We can collaborate to build shared datasets + benchmarks; Physics-aware guardrails/constraints/validation; common platforms for models + tools

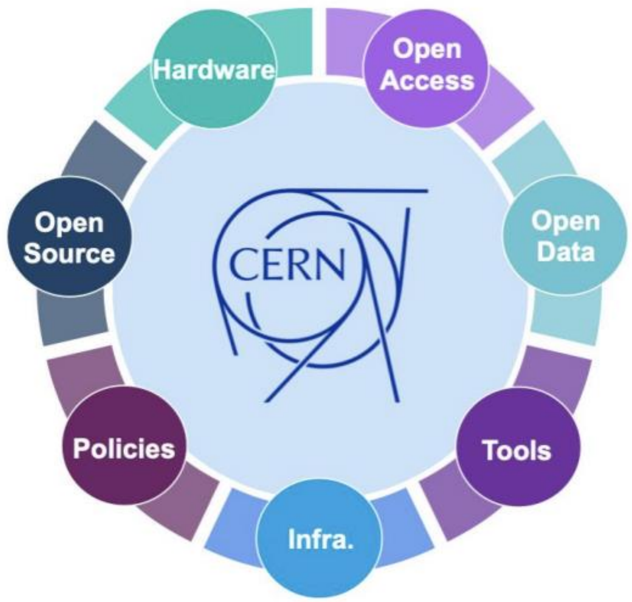


***Open data must evolve from "released" → "AI-ready and usable"***

**Near-term Actions:**

- Aim for a hosted open dataset on a cloud provider
- Collect/share best practices
- Identify resources for large-scale training
- Develop benchmarks and common software solutions…

***Community-scale effort is crucial for foundation models***

# Workforce: Training AI-native physicists

**The challenge:**
- We need training to use AI tools effectively while <u>retaining physics intuition and responsibility</u>

**What we need:**
- Best practices (validation, ethics, responsibility)
- AI-integrated curriculum
- Training resources + starter kits
- Lab–university coordination

**Near-term Actions:**
- Best-practice modules (via IRIS-HEP etc.)
- Starter kit + training sessions
- Access to tools (LLM/agent platforms)
- Build contributor community

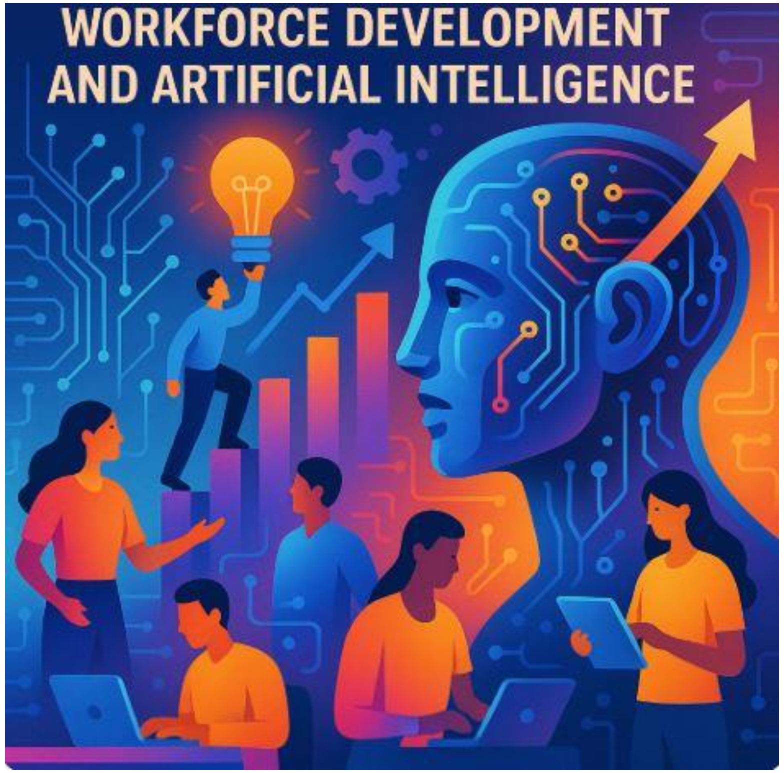


**We must train physicists to *direct, validate, and trust - but verify - AI.***

# “Proto-collaboration” Near Term Activities

- Engage with existing funded projects (and emerging funded projects in the rest of FY26) and create links between them and to the community
- Identify a set of activities and/or demonstrators doable on the time scale of the next six months
- Begin to create environments to catalyze both working collaborations and connections for eventual proposals
  - Slack space created: Join link
  - Channels:
    - #agentic-analysis
    - #workforce-development
    - #marketplace (for skills, MCPs, tools etc)
  - Please engage and start/join discussions!
- Checkpoint and engage the community regarding the proto-collaboration at APS DPF @FNAL in July (Computing parallel session)
- Continue discussions w/ international partners regarding collaborative efforts
- **Aim for a first “collaboration” meeting in the fall.**

# Outlook & Conclusions

- Blueprint meeting was an important step to establish goals and foster connections across sub-communities
  - Plan to organize smaller follow-up workshops to focus on the technical details of specific priority areas
- Encourage interested USMCC members to join the “proto-collaboration”
  - Effort is explicitly structured to provide low-barrier entry points for university groups to contribute meaningfully without duplicating effort
- Lots of opportunities for collaboration across universities and DOE labs
  - Could be used to bring people together for future proposals for funding opportunities
  - Brainstorm about developing industry connections
  - Leverage expertise in the community for the benefit of all
- Start engaging with discussions on Slack and join upcoming meetings
  - Join Slack space!

# Additional Material

# What does it mean to be AI-native ?

**End-to-End AI Integration**

AI-native systems embed AI methods from facility and detector design to intelligent sensing, data acquisition, autonomous operations and calibration, and accelerated analysis for discovery

**Rigorous Validation and Interpretability**

AI-native requires strict validation, interpretability, and physics-informed constraints for trustworthy scientific results.

**Organizational and Infrastructure Changes**

AI-native research demands shared software, standardized data, reproducibility, and model lifecycle management.

**Reshaping Scientific Process**

The AI-native paradigm reshapes science by fully exploiting modern computational intelligence; less time on mechanical steps and more on scientific interpretation (where physicists are needed!)

***AI becomes a strategic accelerator of scientific impact,***
***opening pathways to new transformative scientific capabilities***

# GC1: Accelerated Experimental Design

**Objective:**

- Leverage AI to dramatically shorten and expand the design space for experiments, detectors, facilities.

**Key elements:**

- AI-driven co-design of detectors, accelerators, computing
- Fast exploration of large, high-dimensional parameter spaces
- Integration of simulation, optimization, and engineering constraints
- Earlier feedback b/w physics goals, design choices

**Why it matters:**

- Traditional design cycles are slow and serial;
- AI can enable *iterative, data-driven design* at a scale commensurate with next-generation experiments.

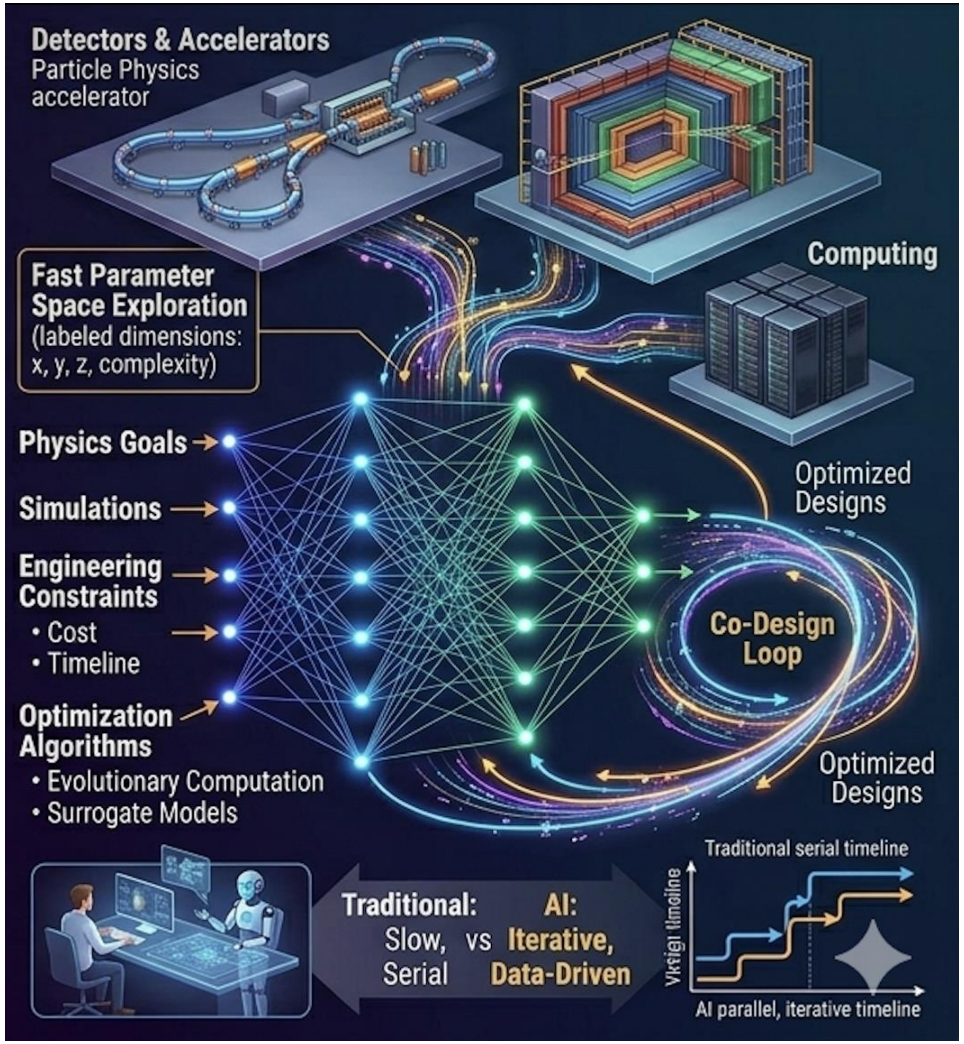


*Co-design accelerators and detectors so that ambitious ideas become buildable, higher-impact experiments on faster timescales with reduced technical risk and costs*

# GC2: Intelligent Sensing & Instrumentation

**Objective:**

- Move intelligence closer to the detector to enable more selective, adaptive, and informative data acquisition.

**Key elements:**

- AI-enabled sensors, front-end electronics, and readout
- Real-time feature extraction and pattern recognition
- Tight coupling b/w hardware capabilities & inference models
- Robust performance under extreme rates and radiation environments

**Why it matters:**

- Future experiments cannot record everything;
- intelligent sensing enables *physics-aware data reduction* without sacrificing discovery potential.

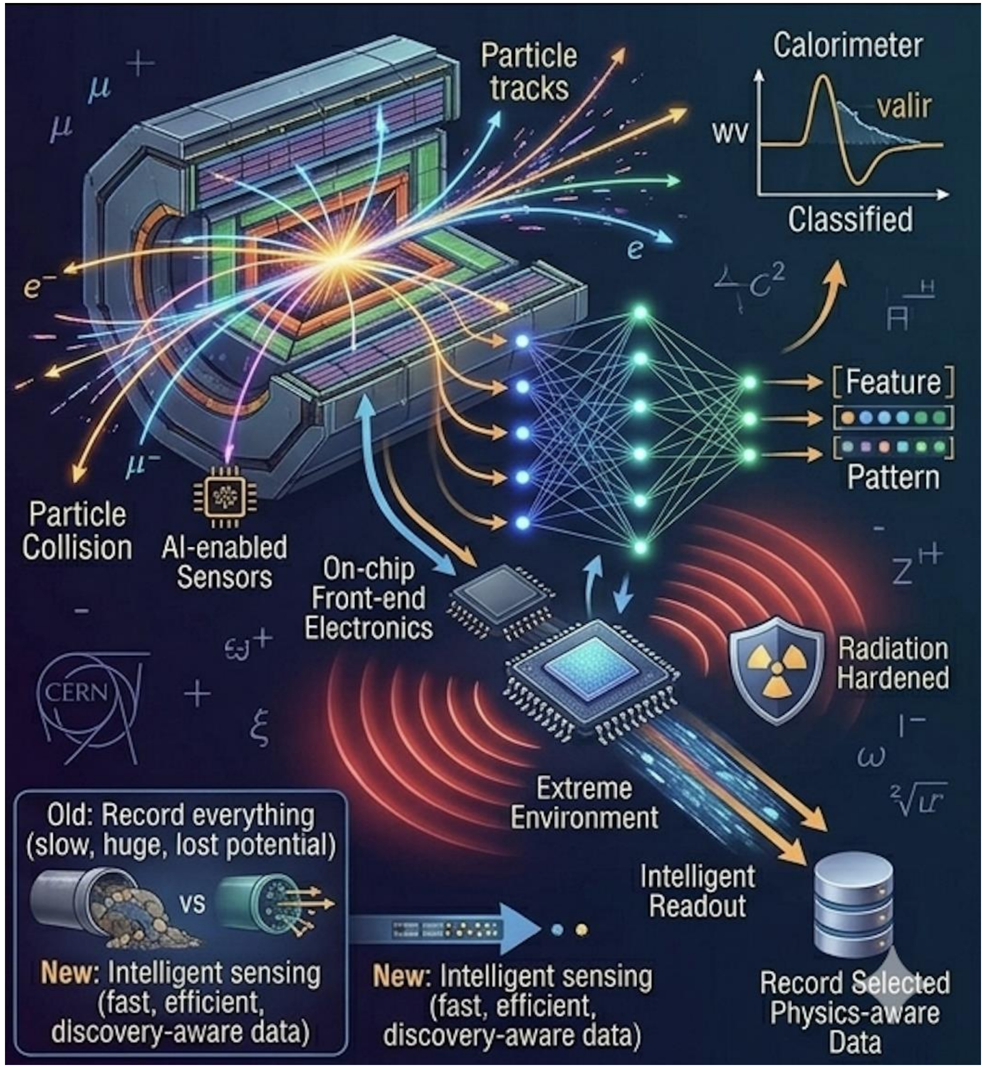


*Move intelligence upstream to preserve rare, unexpected, or time-critical signals while operating within bandwidth and storage constraints*

# GC3: Autonomous Experiments

**Objective:**

- Enable experiments that can monitor, adapt, and optimize themselves with minimal human intervention.

**Key elements:**

- AI-assisted calibration, alignment, and monitoring
- Closed-loop control of experimental conditions
- Detection of anomalies and novel behavior in real time
- Human oversight with AI handling routine decision-making

**Why it matters:**

- Autonomy increases efficiency, resilience, and scientific uptime while allowing human expertise to focus on interpretation and strategy.

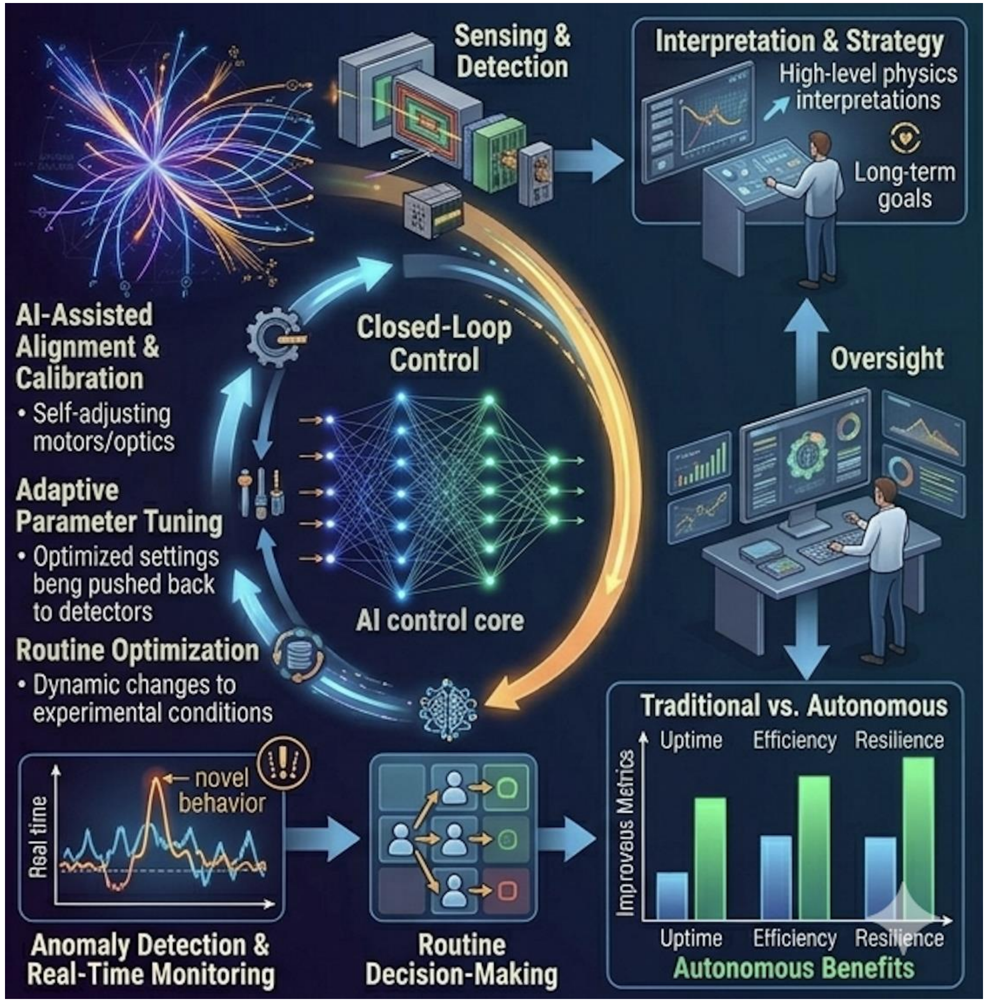


*Transform labor-intensive, reactive operations into proactive, resilient, and continuously calibrated systems, capturing more high-quality data with less downtime and preserving institutional knowledge over decades-long experiment lifetimes.*

# GC4: From Data to Discovery

**Objective:**

- Transform the full analysis chain into an integrated, AI-native inference pipeline.

**Key elements:**

- End-to-end workflows from raw data to physics results
- AI-accelerated reconstruction, simulation, and statistical inference
- Models, data, and uncertainties treated as co-evolving objects
- Rigorous validation, interpretability, and reproducibility

**Why it matters:**

- Fragmented workflows limit scientific reach; AI-native pipelines enable *faster, more reliable extraction of knowledge* from complex data.

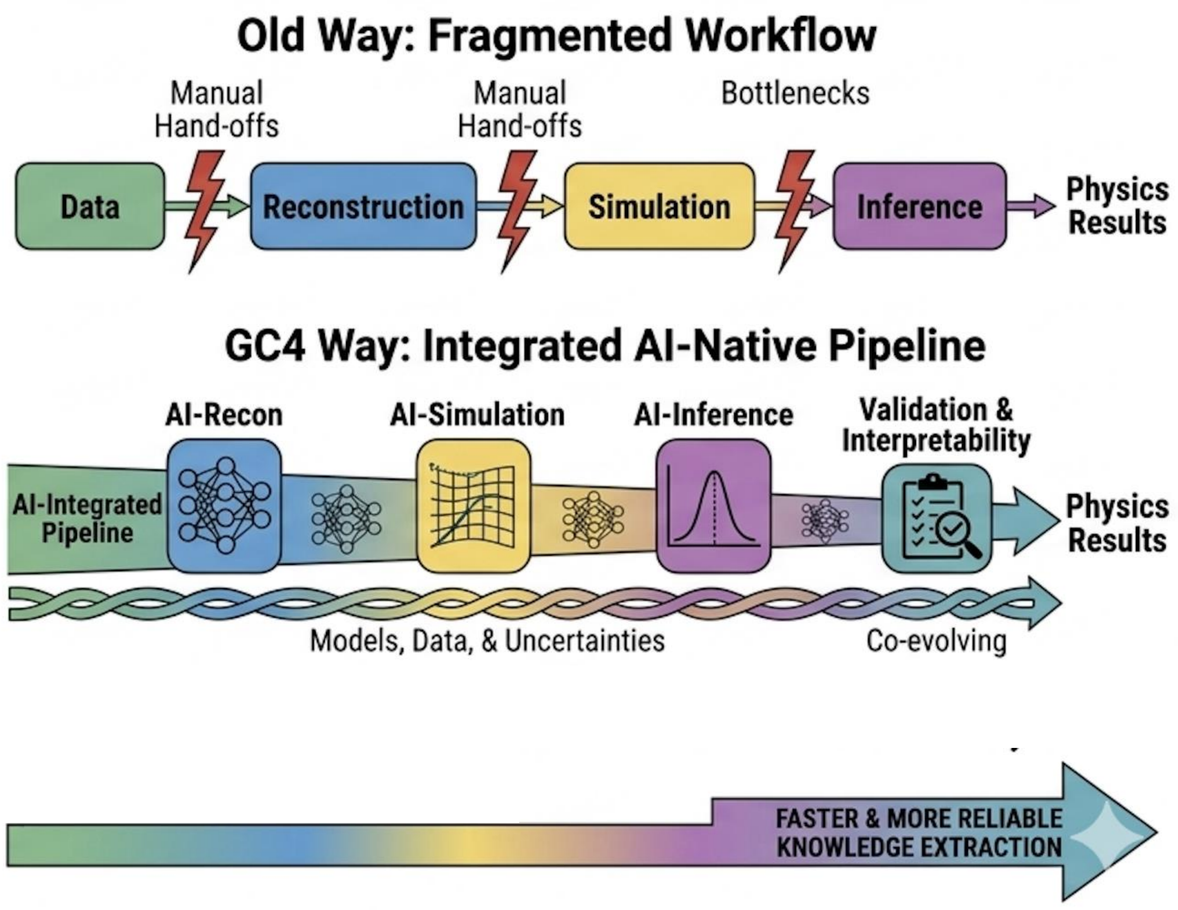


*Integrate foundation models, fast AI-enabled reconstruction and simulation, and agent-orchestrated workflows to compress analysis cycles by orders of magnitude and open new regions of theory space to exploration*

# National Scale Collaboration?

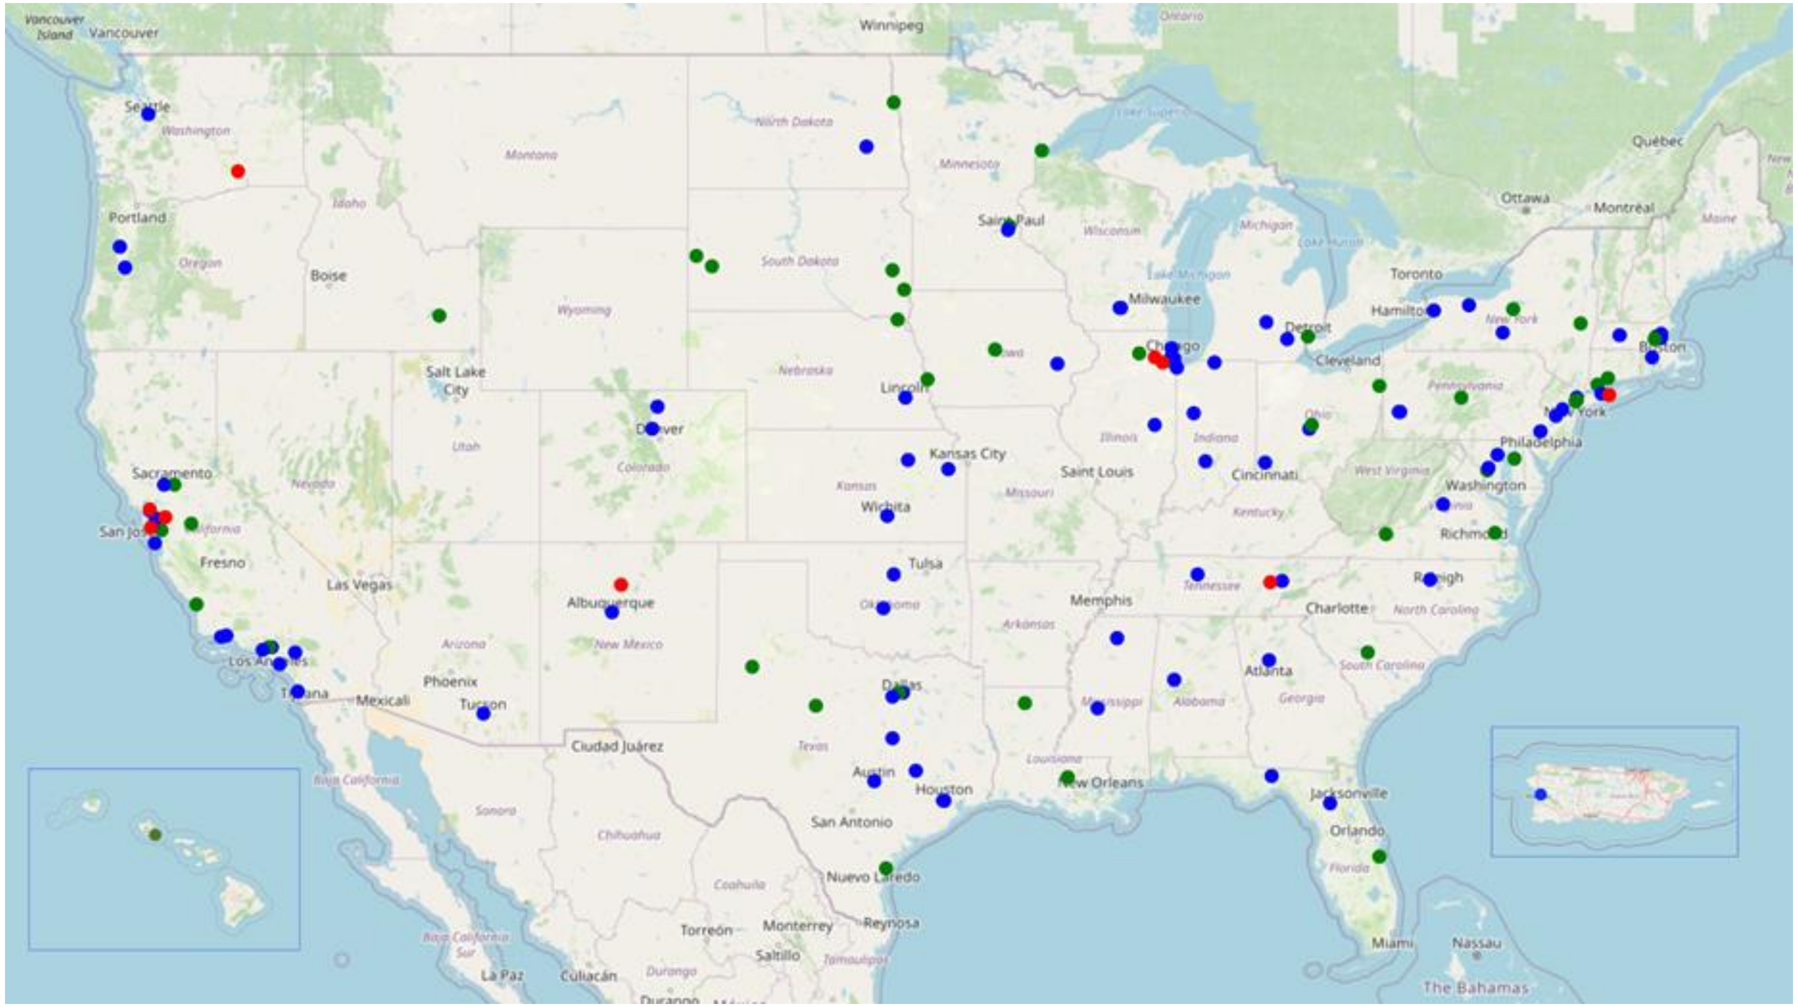


U.S. Institutions (blue dots = universities, red dots = DOE labs) that are currently part of large national-scale funded projects:

- U.S. ATLAS Operations (DOE/NSF, 33 institutions)
- U.S. CMS Operations (DOE/NSF, 44 institutions)
- DUNE Operations (DOE, 16 institutions)
- HEP-CCE (DOE, 4 labs)
- IRIS-HEP (NSF, 16 universities)
- A3D3 (NSF, 12 universities)

Additional universities that are part of the U.S. ALICE, U.S. ATLAS, U.S. CMS and (U.S.) DUNE scientific collaborations are shown in green.

Taken together, these overlapping experiments and projects tightly connect 9 DOE labs and 124 different universities and colleges.

# Supporting ecosystem

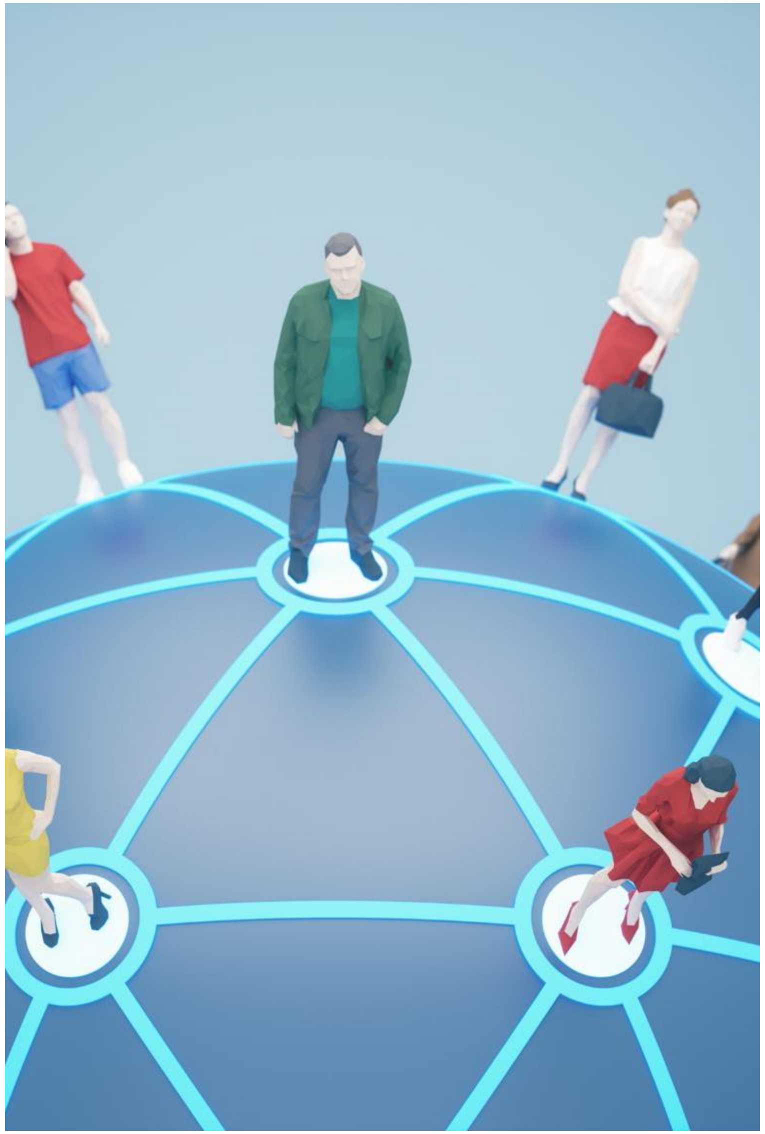

**Interdisciplinary Collaboration**

Success requires **cooperation** among experimental physicists, theorists, computer scientists, machine learning experts, industry

**Community-Level Coordination**

Workshops, shared roadmaps, and platforms enable effective coordination and progress.

**Global Collaboration**

International partnerships are essential due to the worldwide nature of particle physics projects.

**Workforce Development**

Train a new generation fluent in physics, AI, and software

**Incentive Alignment**

**Recognizing contributions** to infrastructure, software, and training motivates sustained involvement. **Career paths** that value software and ML contributions

# (Some) Topics of discussion at Blueprint meeting

**Big-picture talks on (among others) American Science Cloud, TREASURE, Next-Generation DQM…**

**Breakout sessions:**

1. Computing Resource Strategy in the age of AI
2. Agentic AI
3. Open Data & Benchmarks
4. Foundation Models
5. Workforce Development

Summary slides next

# AmSC/Genesis

- **American Science Cloud/Genesis**
  - See talks by T. Childers at the workshop
  - Fermi Data Platform is an Infrastructure Partner
  - U.S. CMS (and U.S. ATLAS) have been invited to become “Infrastructure Partners” for the emerging Genesis Platform
    - Opportunity to initiate new AI projects on top of existing ones

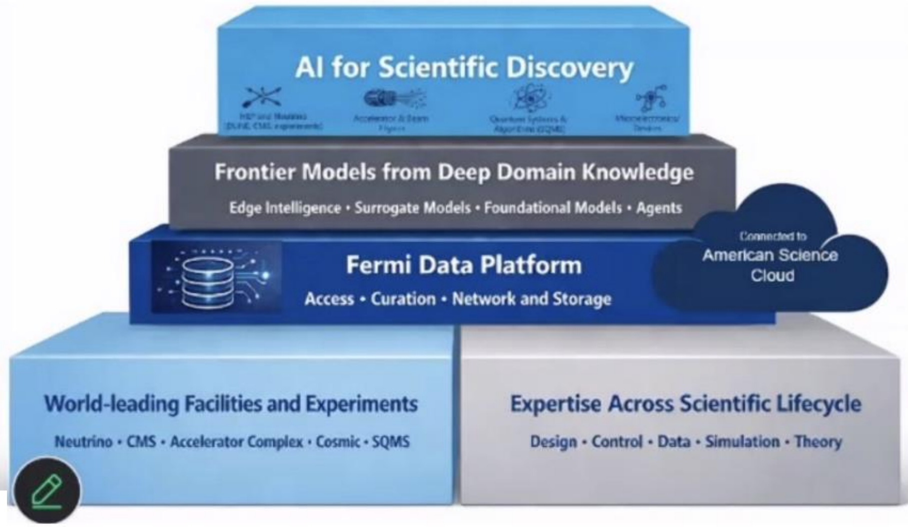


- **<u>TREASURE</u>**
  - Transform heterogeneous collider data into standardized, tokenized, AI-ready representations that enable foundation models to learn across experiments, detectors, and data-taking eras
    - Plan to bring together people from Phase 1 + collaborators for a potential Phase 2 proposal

- **Next-Generation DQM**
  - Building a knowledge base of how to apply and combine AI/ML models for HEP
    - Framework for cross-experiment activities exists w/ hackathons at FNAL/ANL

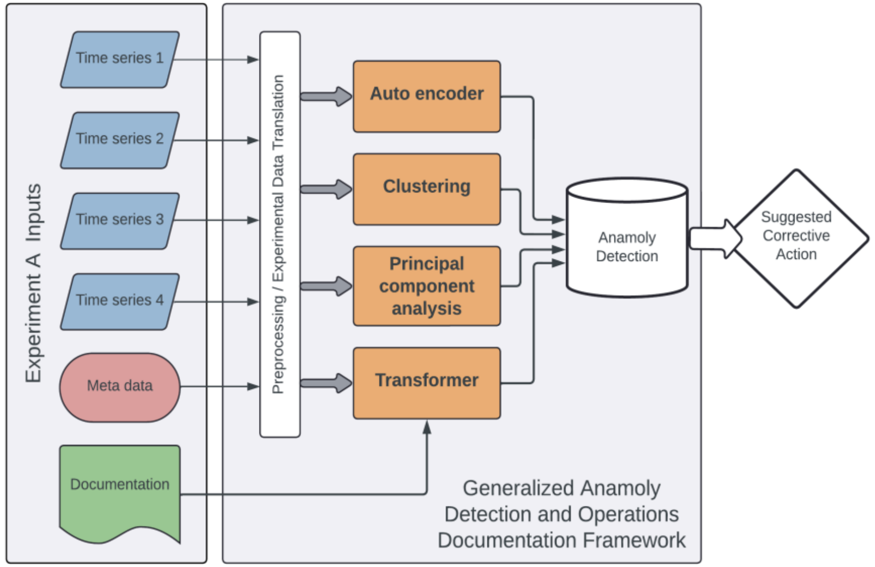

# Open Data in the AI era

**What's changing:**

- Open data must evolve for the AI era, with better formats, metadata, documentation, provenance, and hosted access.

**Implications/key questions:**

- Need machine-readable documentation, integrated hardware/software/model-serving environments, workflow capture, usage monitoring, and harmonized experiment policies.

**Near-term Actions:**

- Produce a white paper with a new, larger vision for open data [in progress]
- Aim for a hosted open dataset on a cloud provider
- Integrate benchmarks and common software solutions,
- (and include neutrino, EIC, and small-experiment communities)

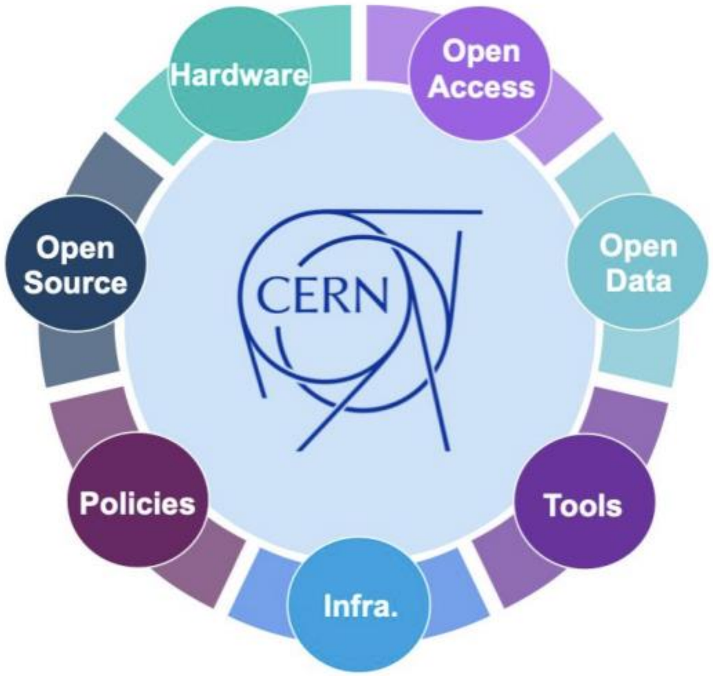


***Open data must evolve from "released" → "AI-ready and usable"***

# Foundation Models

**Why collaborate:**
- Scale (data + compute) exceeds individual experiments
- Common needs (encoders, reconstruction, multimodal models)
- Opportunity to influence broader AI

**What we can build together:**
- Shared datasets + benchmarks
- Physics-aware guardrails/constraints/validation
- Common platforms for models + tools
- Links to existing efforts (e.g., TREASURE)

**Near-term Actions:**
- Define scopes of interest across sub-communities
- Collect/share best practices
- Identify resources for large-scale training
- Develop downstream physics and computational benchmarks.

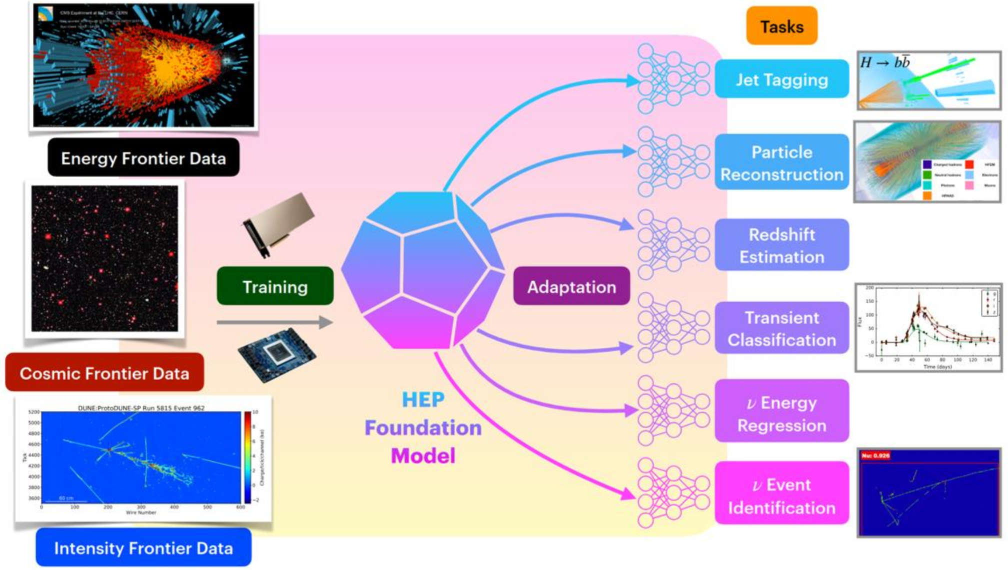


***Community-scale effort is crucial for foundation models***