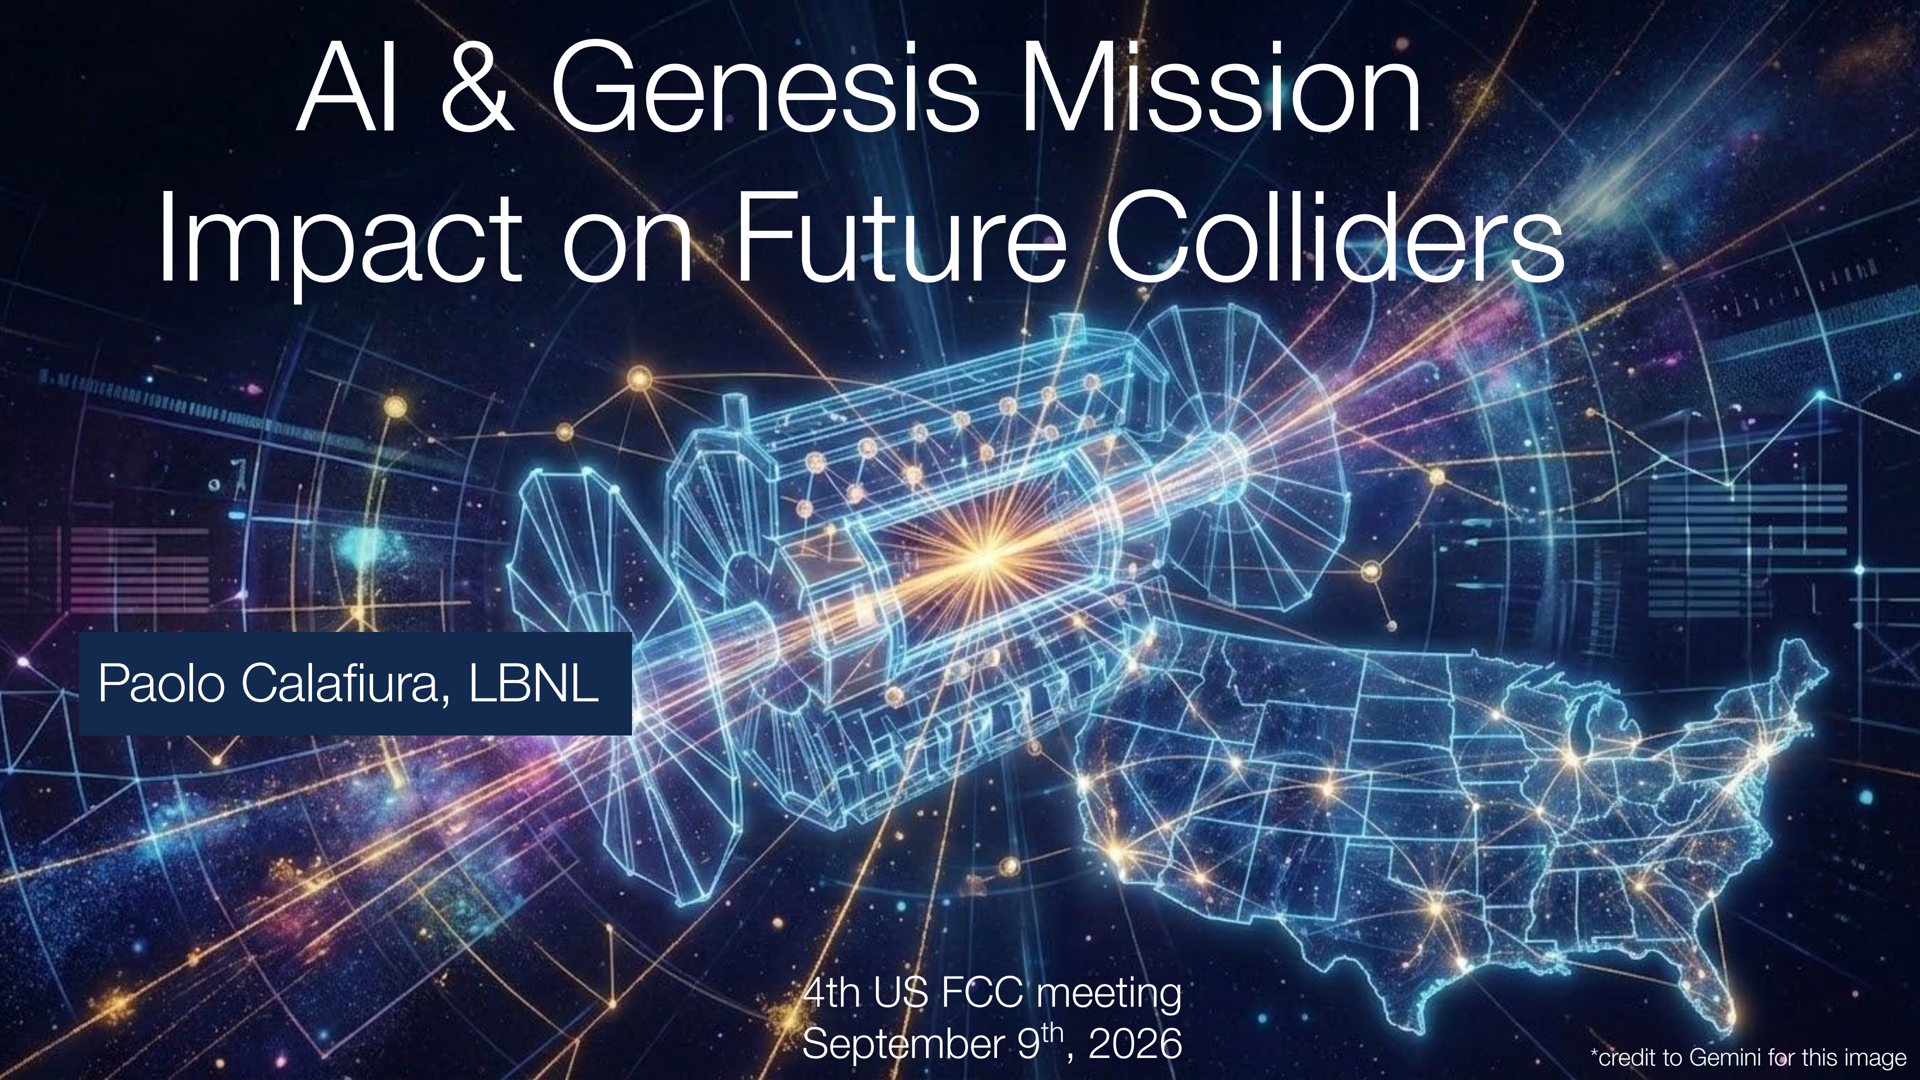

AI & Genesis Mission
Impact on Future Colliders

Paolo Calafiura, LBNL

4th US FCC meeting
September 9th, 2026

*credit to Gemini for this image

~~AI &~~ Genesis Mission
Impact on ~~Future~~ Colliders

Paolo Calafiura, LBNL

4th US FCC meeting
September 9th, 2026

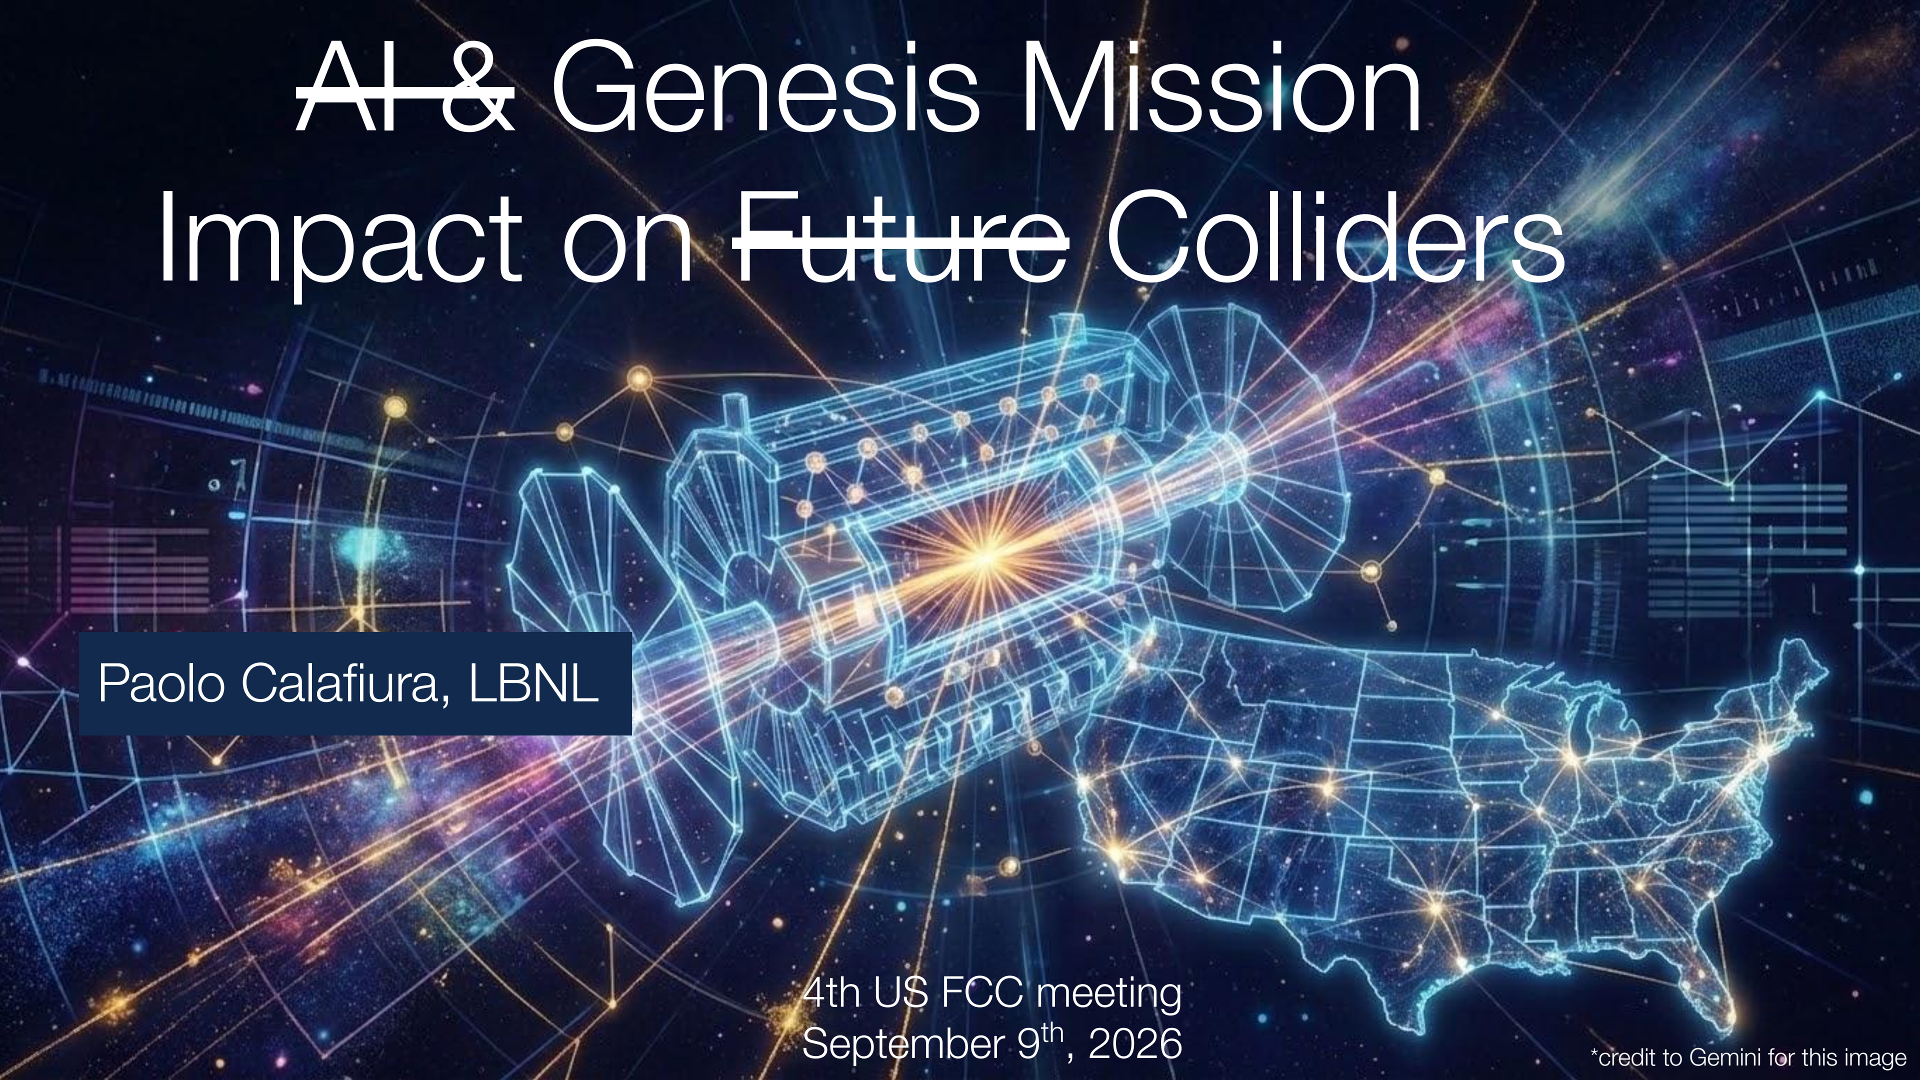

*credit to Gemini for this image

Outline

1) Biased Status Report on Genesis Mission — focusing on HEP Energy Frontier
    a) Model Consortium and American Science Cloud
    b) HEP-related AmSC Infrastructure Partnerships
    c) HEP Pilot Project TREASURE

2) Random comments on the impact of Genesis on HEP experiments

Genesis Mission

[PR page]

A US national initiative led by DOE to build the world's most powerful scientific **platform** to harness **AI** for breakthroughs in **discovery science**, national security, and energy innovation.

17 DOE labs, five NNSA plants and sites, and 41 industry, nonprofit, and philanthropic organizations.

Science & Technology Challenges

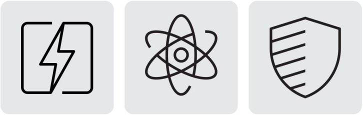

– 278 "Phase 1" science teams, including
– [~15 HEP-related]
– 3 HEP-funded Pilot Projects
– 4 AmSC-funded "Infrastructure Partnerships"
–2 ModCon-funded seed teams

American Science Platform

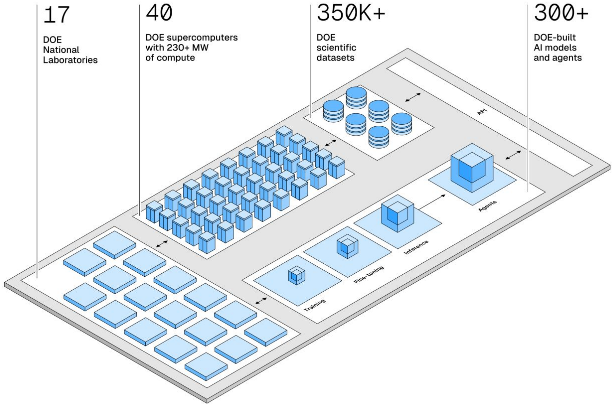

Federated compute/data platform for AI in the DOE national labs

Collaboration & Consortium

Transformational Model Consortium

PR Page →
GitHub Repository →

Baseline AI R&D Capabilities (BASE)

Reusable methods, frameworks, and tools that underpin accelerated AI.

Scientific Workflows (BPSW)

Guidance, patterns, and tools for integrating AI into scientific workflows.

Data Brokers & Standards (DBS)

Centralized access to FAIR datasets with standards and governance tools.

IP & Partnerships (IPPF)

Streamlined frameworks for collaboration across labs, industry, and academia.

Two HEP-related seed teams:
1) Q2C: From Quarks to Cosmos through the lens of AI
2) MOAT: Multi-Office particle Accelerator Team

Led by Argonne with major roles for Pacific Northwest, Berkeley, Oak Ridge, Jefferson, SLAC, NREL, and Idaho National Laboratories, along with supporting roles from Brookhaven, Ames, Savannah River, and Fermi National Laboratories in coordination with NNSA labs.

American Science Cloud (AmSC)

PR Page →
Documentation →

A national resource covering data, model & AI services, and compute, storage & networking infrastructure for science.

Cohesive AI-Ready Platform

Shared Infrastructure: workflows run everywhere

- **Common identity layer**
- **Interoperable service interfaces**
- AI-enabled orchestration
- Commercial technology leverage

Collaborative Ecosystem

Involve future partners from the get-go

- **Infrastructure partners** contribute and integrate core capabilities into the AmSC - key role for DOE labs

Data Services

FAIR, AI-ready datasets across DOE

Model Services

State-of-the-art models for discovery

AI Services

Extreme-scale training & inference

Infrastructure

Secure compute, storage, networking

American Science Cloud - Status & Plans

Demonstrated federated access, next phase focused on onboarding science teams

Currently Available (to early adopters):

Federated Login

– SSO + White List authorization (currently limited to DOE labs)
– In the process of switching to ping identity

Data Management

– OpenMetaData as central register for all DOE data
– Globus/EJFAT for data transfer

Compute

– IRI REST API enables remote job submission at NERSC/OLCF/ALCF and limited filesystem operations

Inference

– **MAG (Model Access Gateway)**
+ Industry cloud resources and models
+ Open weight models on DOE HPCs

AI Tools

– MLFlow, ClearML, MCP, REST & Python API

Genesis Open Model Initiative

– Genesis Science 1 by Arcee
– Call for participation

HEP priority should be to contribute to AmSC infrastructure and integrate our workflows into it.
AmSC poised to become the main source of DOE computing resources, especially for AI

Outline

1) Biased Status Report on Genesis Mission — focusing on HEP Energy Frontier
   a) Model Consortium and American Science Cloud
   b) HEP-related AmSC Infrastructure Partnerships
   c) HEP Pilot Project TREASURE

2) Random comments on the impact of Genesis on HEP experiments

Scientific User Facilities Infrastructure Partnership (IP)

BES/HEP/NP AmSC IP representing tens of thousands of facility users, researchers, and engineers, and decades of achievement and experience with unique scientific instruments including colliders.

High-Impact R&D

Scaled ***Inference as a Service (IaaS)*** to accelerate ATLAS workflows, focusing on high-throughput ML tracking (GNN4ITk) and analysis derivations (DAOD Prod.)

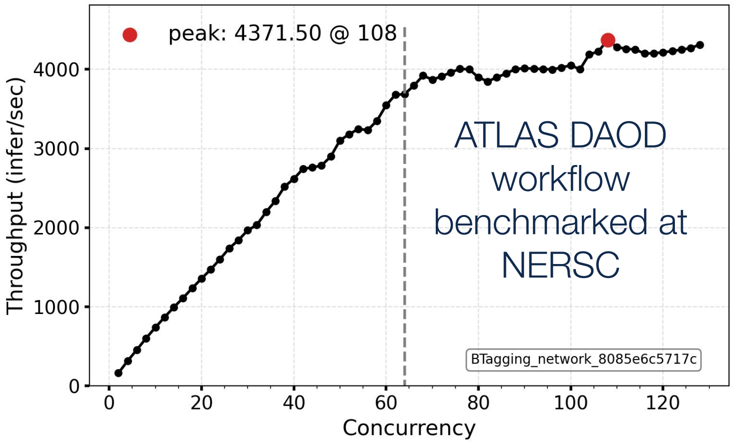

Argonne NATIONAL LABORATORY

BERKELEY LAB

Brookhaven National Laboratory

Fermilab

Jefferson Lab

SLAC NATIONAL ACCELERATOR LABORATORY

AmSC Infrastructure Advancement

Developed a Python client for ***IRI/SuperFacility APIs*** and ***re-architected the IaaS server*** to improve system flexibility and resource optimization

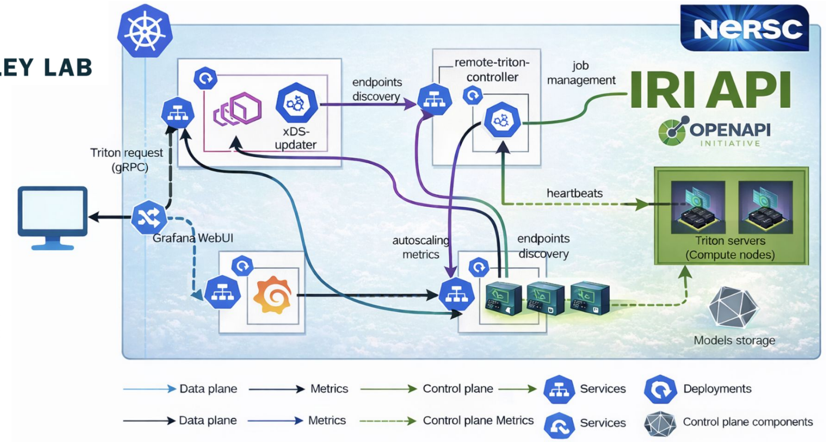

Fermi Data Platform IP

- Created to **host and serve curated scientific datasets**
  - Initially those related to the Genesis mission
  - Expanding to include non-Genesis data
- Provide starter data storage allocations and metadata tools for scientific users to store and manage their AI-ready datasets
- Provide **access** to AI-ready datasets from the AmSC infrastructure **via** the **AmSC metadata catalog and client tools**

Fermi Data Platform is both **storage hardware and infrastructure software** for AI-ready datasets

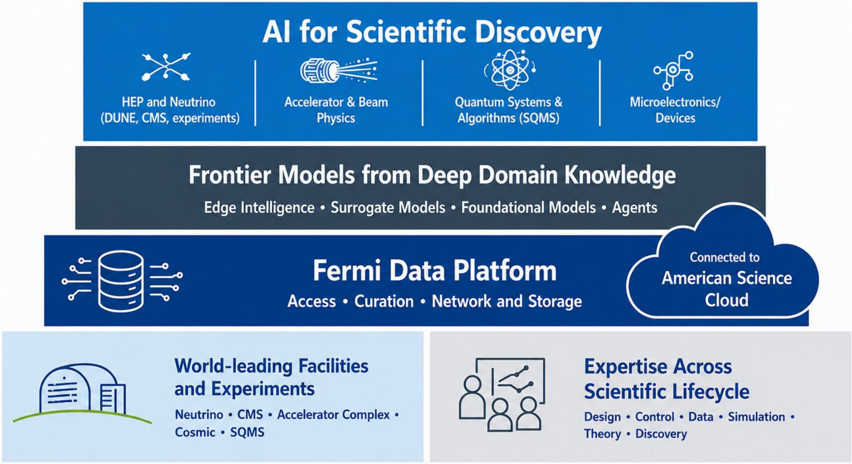

Credit: Robert Harris (FNAL)

FDP: areas, users, datasets, usage, write date

- axess: Hawks, Tran
  - fpga-hpc-blocks: 0.2 MB, Dec. 2025
  - sonos_rad: 0.1 GB, May 2026
  - wa-hls4ml: 5 GB, Dec. 2026
  - wa-hls4ml-projects: 9 TB, April 2026
- cosmiq: Bowring, nexus_run_22: 1 GB, Mar. 2026
- dune: Wang, Bhattacharya, Cerati, Timm
  - dune_testdataset*, 10 GB, May 2026
  - global, 6 GB, Mar 2026
  - jobsub, 6 MB, Jul 2026
  - mwgen*_rbkg_g4_detsim*.root, 30 TB, Mar 2026
- lqcd: Wagman, Fleming
  - demo, 8.3 GB, May 2026
  - misc. files, 0.4 GB, Apr 2026
- rhic: Lancon, tpcpp-1m: 78 MB, Jun. 2026
- treasure: Gandrakota, Pedro
  - aoj: 0.2 TB, Nov. 2025
  - aoj_tokenized_Feb2026: 23 GB, Apr. 2026
  - collide-1m: 32 GB, Nov. 2025
  - PCDF Test: 41 GB, Jun. 2026
- uboone: Wang, Bhattacharya, Cerati,
  - open-bnb-iaas: 2 GB, May 2026
- Areas initiated but not yet used:
  - accelai: Ratnikova
  - moat / sc-magnet: Velev
  - opencosmo: Wells
- Areas planned but not yet created:
  - SQMS: van Zanten, Murthy
  - CODs-CL-pilot: Voetberg, Ratnikova
  - ARIADNE: Wang

Credit: Robert Harris, status as of Aug 26

Federated Orchestration of Heterogeneous Workflows IP

BNL Multi-agent Integration of Multiple Experimental & Computing Facilities

AI-Driven Discovery

Foundation Model for Nuclear & Particle Physics (FM4NPP) to standardize and accelerate HEP & nuclear physics pattern recognition

Operational Agility

Implement federated orchestration to manage complex workflows across diverse, hybrid computing environments

Real-Time Scientific Feedback

Develop a distributed infrastructure prototype that connects offline HPC/Cloud resources to DAQ for near real-time experimental steering

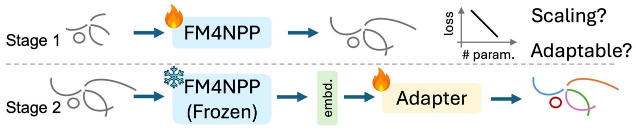

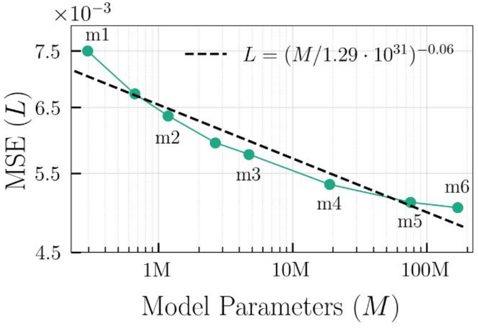

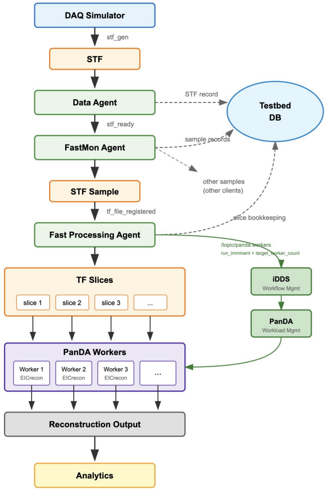

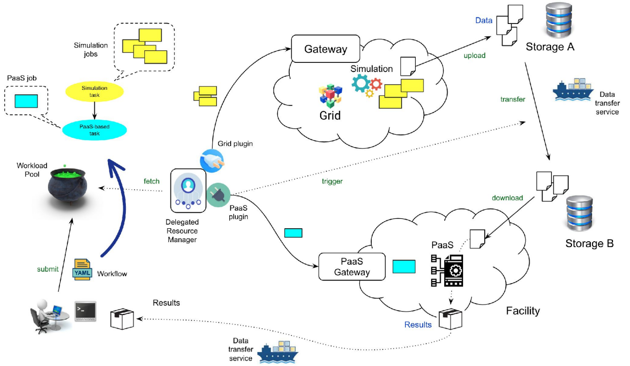

SLAC's S3DF and S3AI Resources

- The SLAC Shared Scientific Data Facility (S3DF) supports over two dozen projects and experiments at SLAC - a range of BES, BER, HEP, and FES experiments.
- S3DF also serves SLAC's Edge computing needs - and supports the SLAC Streaming Sandbox for AI (S3AI)
- S3DF is indirect funded with contributions from projects for their computing and storage needs

Platform Services: User defined services; Kubernetes Platform; Cloud interfaces

Compute: >540 CPU/GPU accelerated nodes; >62k cores

SLAC-ESNET 2x100Gbps (2x400Gbps in FY26)

Core services: Identity/ Access; Batch; Workflow control and Data Transfer

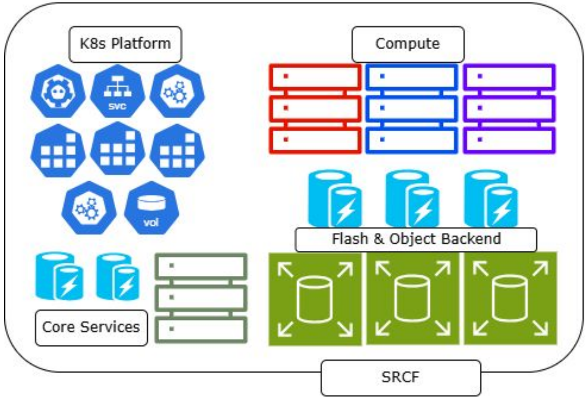

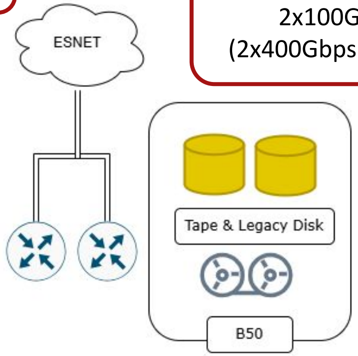

Archival: >100PB science data capacity

Science data storage capacity: >70PB; Posix, S3 interfaces

- **S3AI:** Heterogeneous hardware testbed for streaming inference
  - Access to S3DF infrastructure (storage, computing & networking)
  - Detectors (optical cameras, digitizers, etc); SLAC developed high speed x-ray imagers (ePIX family)
  - Streaming inference and event processing (FPGA boards, commercial accelerators - NVIDIA, GROQ, Cornami, Alveo)
  - Fast and secure encryption & memory pools (Kove, Cornami)
  - SLAC Neural Network Library (SNL) for deploying FPGA based inference

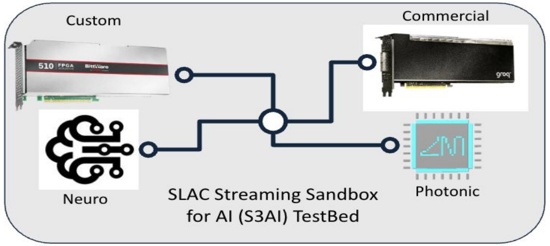

SLAC TID TECHNOLOGY INNOVATION DIRECTORATE

AmSC AMERICAN SCIENCE CLOUD

Outline

1) Biased Status Report on Genesis Mission — focusing on HEP Energy Frontier
    a) Model Consortium and American Science Cloud
    b) HEP-related AmSC Infrastructure Partnerships
    c) **HEP Pilot Project TREASURE**

2) Random comments on the impact of Genesis on HEP experiments

What is TREASURE?

*Tokenized **R**epresentations for **E**nergy-frontier **A**I **S**earches via **U**nderstanding and **RE**asoning*

AmSC Intelligent Data Activities HEP Pilot

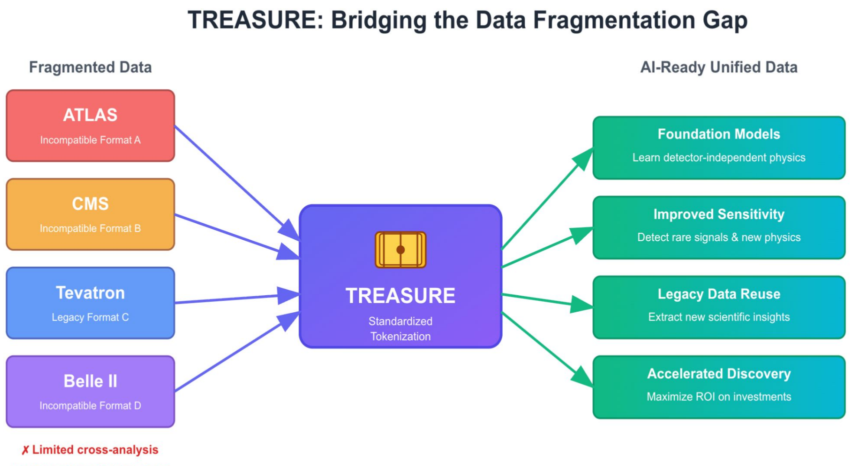

Why this matters?

Collider experiments generate enormous data volumes in incompatible formats.

This fragmentation limits cross-experiment analysis and reuse of legacy data.

Multi-Experiment Training

Learn detector-independent physics features.

- Improved sensitivity to rare signals.
- More data!

Legacy Data Revaluation

Legacy datasets underutilized due to technical obstacles.

⇒ Extract new information and extend scientific impact.

U.S. DEPARTMENT of ENERGY

Genesis Mission

Brookhaven National Laboratory

Argonne NATIONAL LABORATORY

BERKELEY LAB

Fermilab

SLAC NATIONAL ACCELERATOR LABORATORY

the TREASURE Pipeline

GitHub Repository →

Public

To be public framework

To be public

Experiment Data (ROOT)
Object-level Data
(e.g., ATLAS PHYSLITE)

→

AI-READY (HDF5, PARQUET)
per-event arrays,
masked, padded

→

TOKENIZER
VQ-VAE / FSQ /
PQ / RVQ per type

→

PARQUET
Integer token
sequences

→

FOUNDATION MODEL
Pretrain → fine-tune
→ transfer

PCDF: The Particle Cloud Data Format

Stores events and physics objects in linked **flat** tables. The relationships are preserved using **indices**.

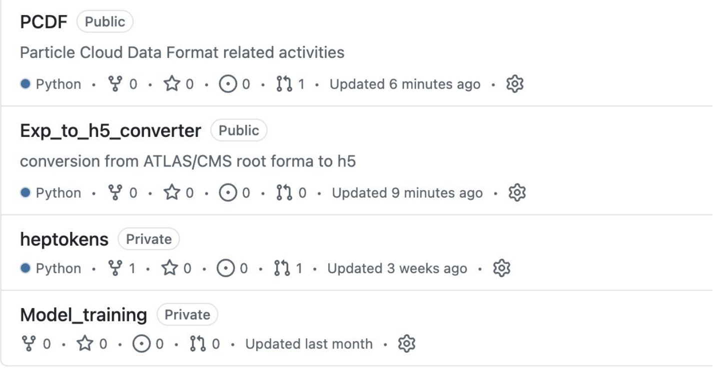

Credit: Merve Nazlim Agaras, EuCAT 2026

TREASURE High-Level Object Tokenizers

Vector quantized variational autoencoder (VQ-VAE) compresses high-dimensional inputs into a learned codebook
Each object type has its **own separate VQ-VAE** with its own codebook, because each object type has completely different features:

- Electron: pT, η, φ, isolation variables
- Jet: pT, η, φ, mass, b-tagging etc
- Track: qOverP, η, φ, d0, z0, pT, χ^2, nDoF

TREASURE tokenizer builds on substantial existing jet-level work (like 2401.13537 and 2412.10504) extending it towards full-event representations. The $1M question is: **how much information do we lose by tokenizing**?

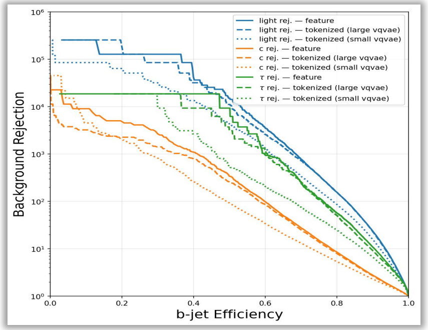

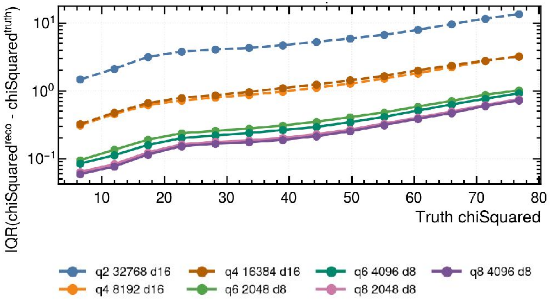

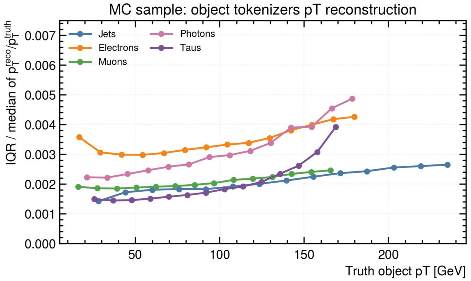

ATLAS 2015/16 Open Data used in this study

Credit: Merve Nazlim Agaras, EuCAT 2026

Genesis & HEP Summary

HEP EF has actively contributed to the 1st year of Genesis.

- Our contributions were recognized by AmSC and DOE program managers.
  - Opportunity to expand our role

What's Next?

- Genesis moves fast and evolves rapidly. Wait for official announcements!
  - We expect **more award announcements** in the coming weeks.
  - There will be a **review of AmSC and ModCon** in the Fall.
  - After that, we expect a **new RFA (call for proposals) early next year**.

Outline

1) Biased Status Report on Genesis Mission — focusing on HEP Energy Frontier
    a) Model Consortium American Science Cloud
    b) HEP AmSC Infrastructure Partnerships
    c) HEP Pilot Project TREASURE

2) **Random comments on the impact of Genesis on HEP experiments**

Genesis Challenges: AI for Science and Science for AI

HEP experiments develop and control their hardware and software platform focusing on **minimizing cost and time to result**.
Fundamental AI research is (at best) a secondary goal.

Genesis vision is to transform US science through a **unified hardware and software platform.**
Advancing AI state-of-the-art very much a goal.

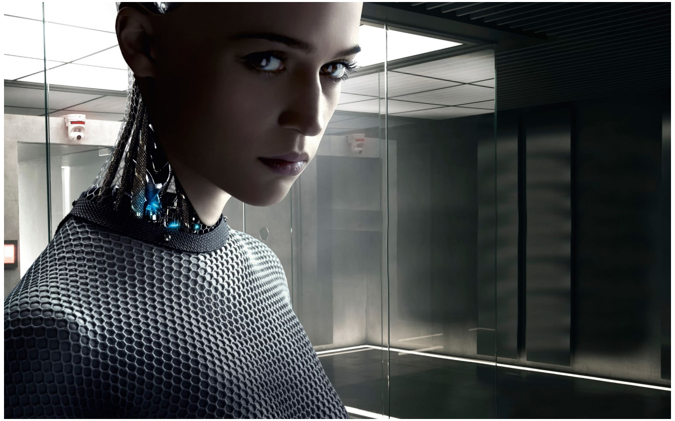

Genesis Challenges: Time-Horizon and Policies

- Future colliders have a longer time horizon than Genesis (5-10 years)
  - This does not mean that FCC applications are out of scope
    - Detector optimization, ideally through a full-physics differentiable simulation+reco+analysis chain would be a (very) transformative project if we could have it in five years.
- Genesis emphasizes the role of industry partnership, and the impact of using their frontier models to accelerate and eventually transform US Science
  - HEP experiments currently favor open source software and commodity hardware
    - OTH the experiences of ATLAS analysis groups at NERSC have been very positive
      - Good support
      - Flexibility
      - Quick account approval
      - Access to hard-to-obtain GPU resources

The Funding Challenge

- We did not do well in the FY26 RCA for “Phase 1” projects
  - No one seem to know why, there are probably many reasons
    - My favorite: we are a large, mature AI community with years of successful R&D behind us.
      - To be “transformative” (as required by Genesis Mission) we may well need a moonshot project requiring AGI-scale resources.
- We are doing better in the “AI infrastructure” projects
  - Probably because of our experience working at scale in a distributed production environment.
    - We should leverage that.
      - Do FCC applications need AI services (IaaS, model repositories, MCP, etc)?
- One area of growth are industrial partnerships.
  - We have a lot to learn from their R&D teams.
  - In my experience, industrial partners are excited to work with “big science” projects.
    - Can we make a good business case?

In any case, **we can’t sit this one out:**
we need Genesis resources, models, and tools that meet our needs.
Some funding too!

Thanks for your attention!

I am grateful to:

D. Bard,
V. Martinez Outschoorn,
H. Abidi, G. D'Amen,
and the Treasure team,
O. Gutsche and the FDP team,

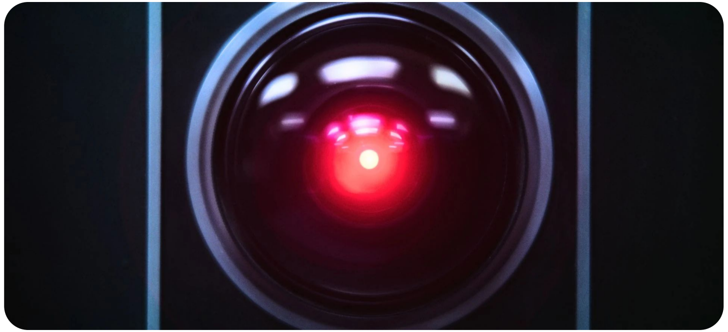

A. Klimentov and the BNL IP team,
Pete Elmer,
T. Childers,
C. Leggett,
A. Rastogi, X. Ju, and
S. Pagan Griso

and others for all their inputs – all errors are mine – credit also to Gemini for slide design.

Community & Collaboration in AI for Experimental HEP

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

Interested in building a collaboration around AI for experimental HEP

See recent meeting for a discussion of potential areas of collaboration at this link

Whitepaper

More than 500 people from 105 universities and 7 national labs from the US & other international institutes signed the paper

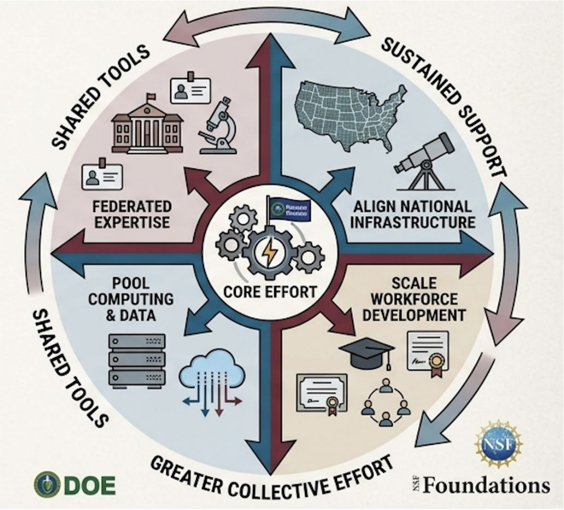

In 2015 I gave a similar talk about ML and HL-LHC

Mercifully accessible only by ATLAS members

These were my “predictions” at the time:

Wager #1: ML will be predominant in Run 4 analysis

Wager #2: Deep neural networks in tracking, jet reco, and clustering will allow us to exploit GPUs and FPGAs, and possible neuromorphic architectures

Wager #3: 2025 grad-students will wonder why we wrote all that C++ junk instead of training a few good networks

Right. Remember that Run 4 was scheduled to start ~now.

Maybe, the jury is still out on this one (no neuromorphic though!)

Mostly wrong, students mainly wonder how was life before Claude

TREASURE Project Timeline

Year 1

Year 2

Foundations

- AI-ready datasets from LHC Open Data
- Common data models, metadata, and tokenization schemes
- Assess AI-readiness and quantify tokenization impact
- **Demonstrate cross-experiment learning with prototype AI models**

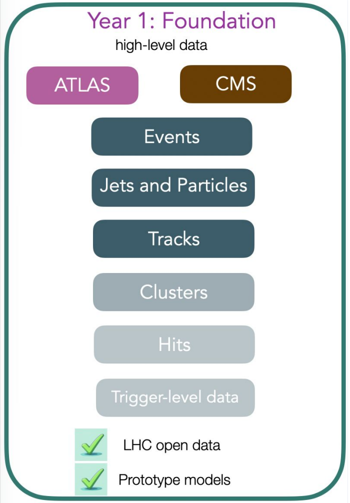

Expansion and Integration

- **Detector data:** hits, clusters, streaming compression
- **Beyond the LHC:** Tevatron (ppbar), Belle II (e+e-), future experiments
- **Scale up cross-experiment foundation model training**

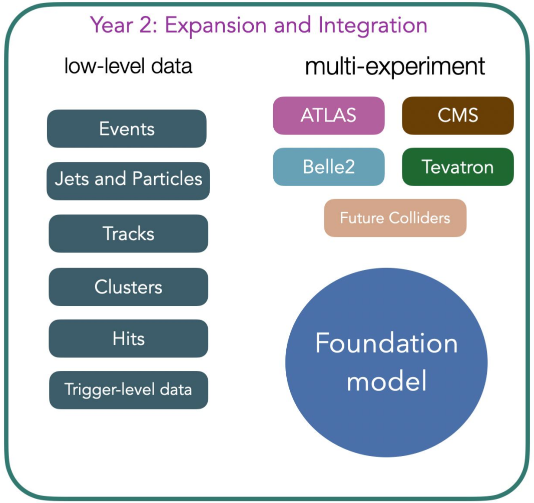

New scope added: SSO + Model Access Gateway for RFA Phase 1 teams delivered early!

- Supports access to frontier models on CSPs (AWS, Microsoft now; GCP, OpenAI ~Sept)
  - (Based on Chat UI prototype already tested by AmSC Early Users)

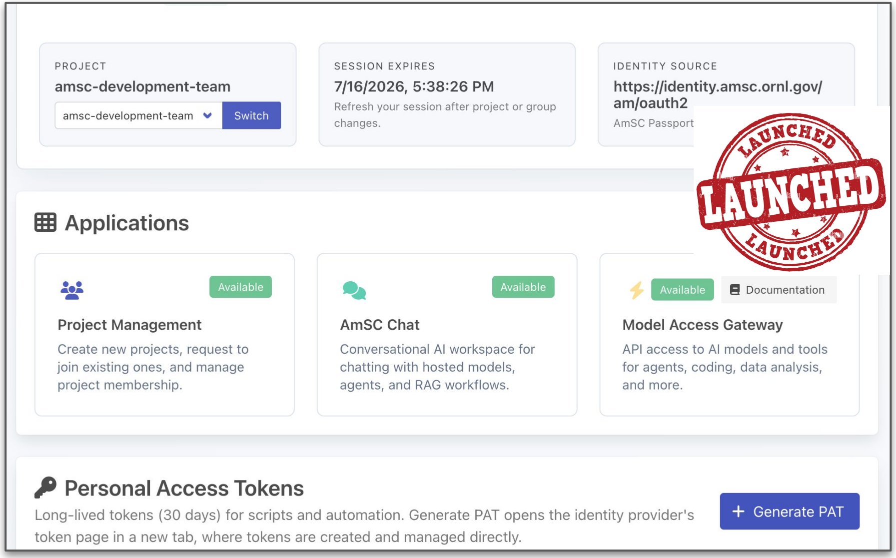

- Launched at the July 22nd Genesis Mission event
  - Our staff put in a lot of time and effort to pull this together on a short timeline
- Extremely useful "forcing function":
  - Obtained IATO (Interim Authority To Operate)
  - SSO: users can register with AmSC using their home identity from 12+ labs

Federated ID (🥳)

Many users can now use their home lab ID to authenticate via SSO with the AmSC Platform

- FedID enabled by reciprocity agreements with other labs. Working on adding other ID providers.
- Currently MAG is available via SSO
- In progress: integrating additional services into SSO

IP teams, ModCon, vendors… want to integrate their own services and infrastructure into AmSC's SSO and therefore make them available in the platform

- Expose their resources to AmSC users via SSO
- Expect to integrate most IP team services/infrastructure into AmSC SSO by Q3 FY27

Data Services

Goal: Users can find, register and move datasets, setting appropriate access controls.

Data Catalog	• Centralized metadata store for **discovering** AmSC datasets • Catalog of catalogs synthesizing dataset registries from IPs and facilities
Data Lakehouse	• Columnar data **storage** supporting highly scalable SQL-like data querying • Hosting solution for AI-ready curated datasets generated from raw data
Data Curation	• **Pipelines** for processing and generating AI-ready data from raw data • Facilitating ingestion into the Data Lakehouse
Data Movement	• **Transferring** registered AmSC datasets from IPs to compute facilities • **Streaming** data transfer from experimental facilities to compute

AI Services

Goal: Users can leverage industry-standard tools to run automated training, fine-tuning and optimization pipelines, including agentic workflows.

At-Scale services	• Provide **curated software stacks** across AmSC resources • **Benchmark** AI applications to validate software stack functionality and track performance (with ModCon)
Model Services	• MLFlow, ClearML and other frameworks for **experiment tracking** and model catalog • **DBaaS** for AI
Inference Services	• **Serving LLM and non-LLM models** hosted at DOE sites and CSPs • Accessible OpenAI-compliant **API**
Intelligent Interfaces	• Deploying **Agentic stack** in AmSC - including LLM API access, ChatUI, Agent Builder and MCP Server hosting • Interface to **LLM inference** services at facilities and commercial providers

The IRI-API: Cross-DOE and CSP Compute Access

Goal: Enable automated workflows across all AmSC sites

The IRI-API can be used across AmSC sites (and Cloud Service Providers) by humans and agents to manage:

- Resource status
- Job management
- Local file management
- Allocation information
- Resource orchestration
- WAN management

Operation ID	Demo	NERSC	ESnet	ALCF	OLCF
GET /api/v1/facility	PASS	PASS	PASS	PASS	PASS
GET /api/v1/facility/sites	PASS	PASS	PASS	PASS	PASS
GET /api/v1/facility/sites/{site_id}	PASS	PASS	PASS	PASS	PASS
GET /api/v1/status/events	PASS	PASS	PASS	PASS	PASS
GET /api/v1/status/events/{event_id}	PASS	PASS	PASS	PASS	FAIL
GET /api/v1/status/incidents	PASS	PASS	PASS	PASS	PASS
GET /api/v1/status/incidents/{incident_id}	PASS	PASS	PASS	PASS	PASS
GET /api/v1/status/resources	PASS	PASS	PASS	PASS	PASS
GET /api/v1/status/resources/{resource_id}	PASS	PASS	PASS	PASS	PASS
GET /api/v1/account/capabilities	PASS	PASS	PASS	FAIL	MISSING
GET /api/v1/account/capabilities/{capability_id}	PASS	PASS	PASS	FAIL	MISSING
GET /api/v1/account/projects	PASS	PASS	PASS	PASS	MISSING
GET /api/v1/account/projects/{project_id}	PASS	PASS	PASS	PASS	MISSING
GET /api/v1/account/projects/{project_id}/project_allocations	PASS	PASS	PASS	PASS	MISSING

Other labs/CSPs in progress!

AmSC is moving from build/MVP mode to operations

As we onboard new users from RFA teams, we need to offer them reliable, secure, supported services. This means we have to change the way we plan and manage work, and how we allocate staff time. (standard practice for start-ups!)

"Velocity built the MVP. It cannot run the platform."

	Build mode	**Operate mode**
SW Release management	Sandbox: move fast, break things	Staged, reviewed, reversible releases
Reliability and SLAs	No commitment	Uptime is a contract
Incident support	No commitment	On-call rotation with named POCs
Security & compliance	Restricted access for development	Vital: High scrutiny and increased attack surface
Support obligations	Tickets answered at best-effort level	Committed ticket response time
Docs & playbooks	System explained by devs	Aim for self-sufficiency
Technical debt	Debt is the price of building an MVP	Debt is scheduled work with a named owner

AmSC priorities for FY27

Launch MVP Platform

- **Release** MVP services to users in monthly cadence
- **Integrate** AmSC, IP and ModCon services and resources into SSO and the Platform
- **Harden** services: scale and address technical debt

Onboard new science teams

- RFA teams
- Additional IP teams (DOE-wide)

Deploy new services

- CUI data support
- AI developments

Expand partnership model

- External partners (e.g., US agencies, Japan Collaboration, CSPs)
- Vendor engagement
- Governance structure