What is ModCon?
Transformational AI Model Consortium (ModCon)

April 3, 2026

ModCon
Transformational AI Models Consortium

Mission

- Establish a consortium to accelerate the technical development and scientific discovery of the Model Teams
- Develop and deliver domain cross-cutting services as an engine for transformational AI model development
- Convene partners from industry, academia, and internationally to accelerate AI development and adoption

Four Core Teams

deliver support to Genesis Teams

IPPF
Partnerships and IP Management

DBS
Data Brokers & Standards

BPSW
Best Practices for Scientific Workflows

BASE
Cross-Cutting AI Capabilities

What We Deliver

AI-Ready Data Pipelines
Transform raw scientific data into training-ready datasets

Scientific Workflows
Leaderboards, workflows, and upskilling for continuous discovery

Transformational Capabilities

- Core Agentic Framework
- Self-Improving Models Harness
- Multimodal Reasoning Frontends
- Safety, Security, Assurance
- Evaluation

Deployed on Genesis Mission Infrastructure
The American Science Cloud

Partners
17 DOE Labs + Industry + Academia

Genesis Mission Mapping
Data and Model Teams

Genesis Mission

Governance and DOE Engagement
Lab Council, ALDs, DOE Lab Reps
Partnership Council, External Executive Voice

ModCon Office

Director
Rick Stevens, Argonne

Deputy Director
Court Corley, PNNL

Project Management
Emily Dietrich, Argonne

Model Team Coordinator
Tom Brettin, Argonne

AmSC Coordination
Neeraj Kumar, PNNL

Cross-Cutting Capabilities

IPPF: IP & Partnership Formation
Megan Clifford, Argonne
Daniel Ratner, SLAC
Jason Hales, INL

BPSW: Best Practices for Scientific Workflows
Franck Cappello, Argonne
Rafael Ferreira Da Silva, ORNL

DBS: Data Broker Standards
Lauren Biven, JLAB;
Jeren Browning, INL
Ray Grout, NLR

BASE: Baseline AI R&D Capabilities
Stefan Wild, LBNL
Nathan Hodas, PNNL

CAF: Core Agentic Framework
Ian Foster, Argonne

EVAL: Evaluation
Stefan Wild, LBNL

SAFE: AI Security & Safety
Nathan Hodas, PNNL

SIM: Self-Improving Models
Ana Gainaru, ORNL

DATA: AI-Ready Data
Shreyas Cholia, LBNL

MM: Multimodal Reasoning
Sandeep Madireddy, Argonne

American Science Cloud (AmSC)

A Unified Platform for Transformational AI and Science Across DOE

- AmSC is a secure, federated, and science-optimized cloud environment
  - *Not centralized servers but an environment supporting contributed tools, services, and ModCon core capabilities*
- Integrates DOE's world-leading computing, experimental facilities, data resources, and HPC networks
  - *A flexible structure to support integration of capabilities*
- Enables AI-ready end-to-end workflows across the complex
  - *Contributions from across the DOE lab complex, not predefined workflows*

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

Genesis Mission defines ambition and urgency.
AmSC scaffolds a shared platform to execute.

DOE Labs and Infrastructure Partners

ANL · LBNL · BNL · FNAL · INL · JLab · NETL · NLR · ORNL · PNNL · PPPL · SLAC

https://amsc.energy.gov/

ModCon Core
Baseline AI Research & Development
BASE

Genesis Mission

U.S. DEPARTMENT *of* ENERGY

Base R&D is Delivering Baseline AI Capabilities to the Genesis Mission

- Providing reusable methods, frameworks, and tools that underpin every Model Team
- Responsive to the requirements set by the BPSW, DBS, IPPF crosscuts
- Collaborating directly with Model Teams and AmSC

We are also:

- Genesis Mission's crosscutting AI dev team
- Filling in what is missing from the current AI ecosystem
- Helping inspire others to use the command line

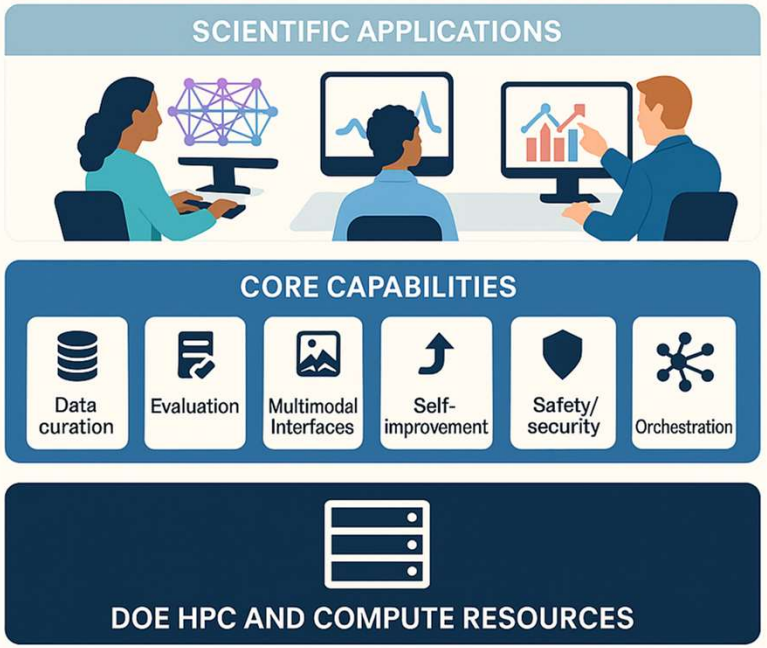

FASTER SCIENCE

SAFER AI

SCALABLE ORCHESTRATION

Stefan Wild (LBNL)

Nathan Hodas (PNNL)

Base/CAF: Building a Common Agentic Framework


Ian Foster
(ANL)


Yadu Babuji
(ANL)


Matt Baughman
(PPPL)


Kyle Chard
(ANL)


Tanwi Malllick
(ANL)


Fred Suter
(ORNL)


Bella Yang
(ANL)

Agents enable **closed-loop, autonomous scientific workflows** across complex, distributed systems

To enable agentic AI applications across AmSC at unprecedented scales, CAF is:

- **Developing scalable agentic orchestration** (1000s of concurrent LLMs, agents, tool calls; flexible inter-agent comms)
- **Promulgating best practices** for agent use in DOE environments

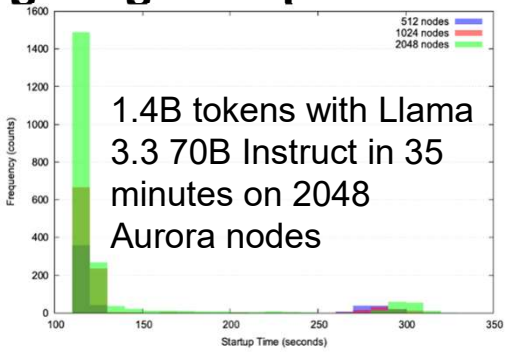

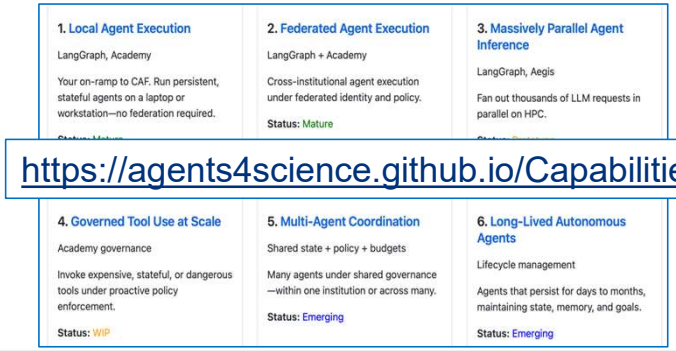

https://agents4science.github.io/Capabilities

Base/Data: Automating the Data Preparation Pipeline

Lead: Shreyas Cholia
(LBNL)

Jean Luca Bez
(LBNL)

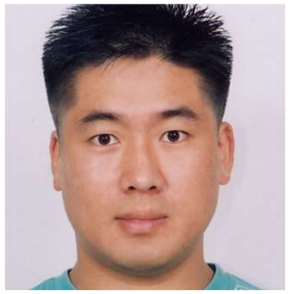

Jong Choi
(ORNL)

Rohith Varikoti
(PNNL)

Kyle Parfrey
(PPPL)

Andrew Tritt
(LBNL)

Aaron Tuor
(PNNL)

Developing an AI-agent driven data curation pipeline that autonomously transforms raw scientific data into AI-ready form

Contact:

SCholia@lbl.gov

Base/Data: Composable Data Pipeline Framework

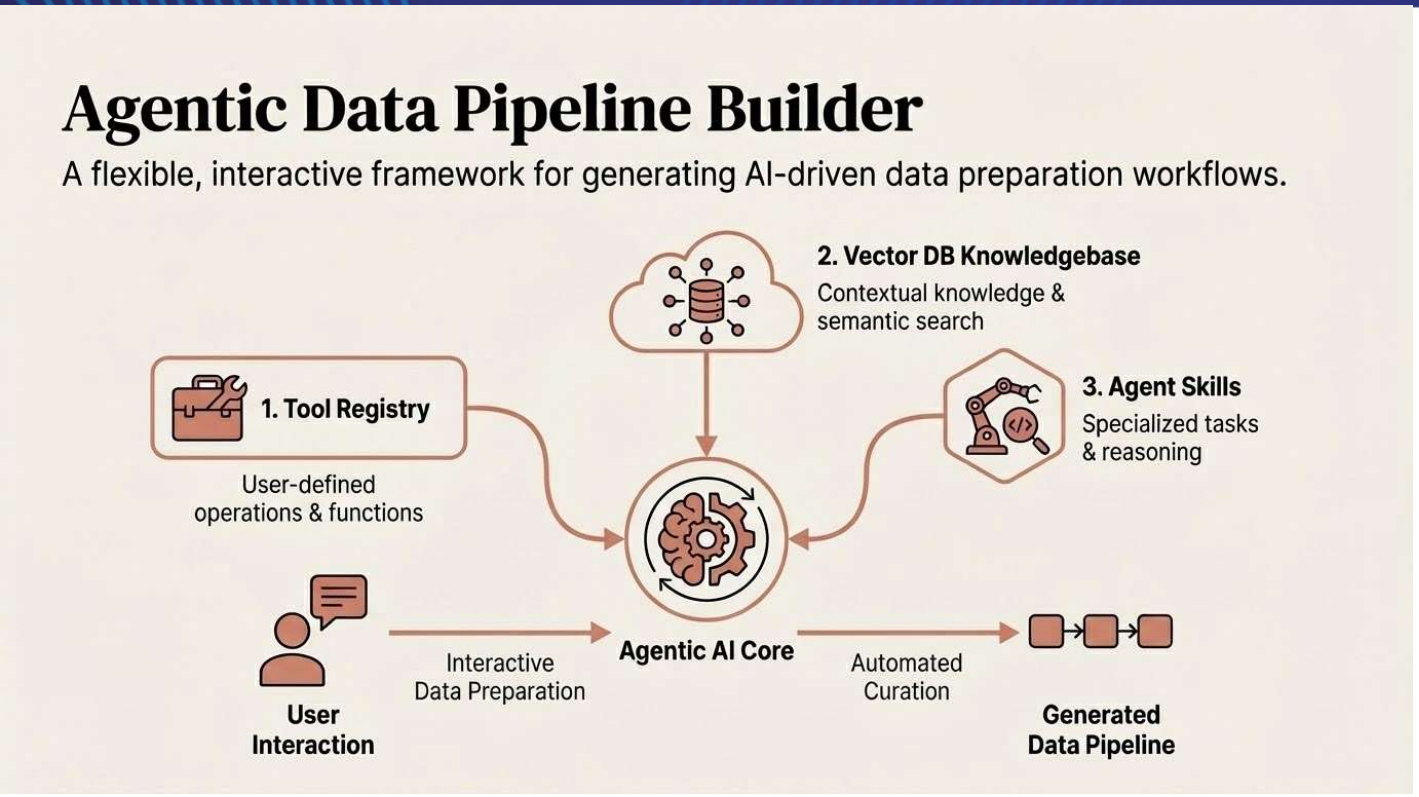

Data Science Agent Toolkit (DSAGT): Agentic Harness to manage data preparation pipeline

- Use your favorite interactive LLM agent (Claude, Codex, OpenCode, Goose etc.)
- Support for data tool execution, RAG knowledge-bases, Agent Skills
- Use of "Skills" for agentic tasks: Data Card Generator, HDF5 data ...

Base/Data: AI Readiness

AI Readiness

- Working with ModCon DBS on AI readiness tool list / metrics / evaluations.
- Practical definition of AI-readiness: Getting to **model.load(data)** + additional metadata
- Example pipelines / steps / definitions vary across model teams

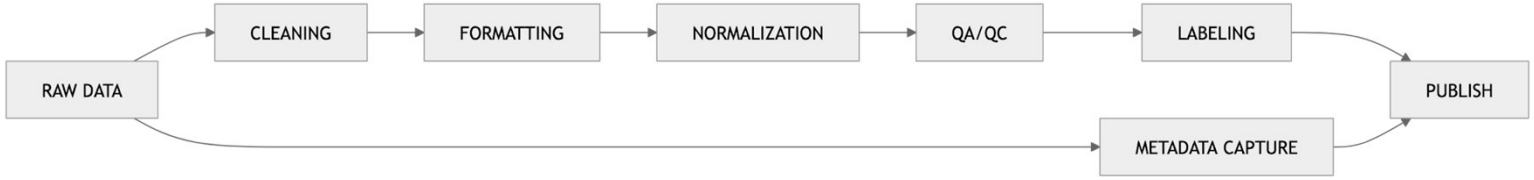

- Data card generation and visualization tools
- Work with model teams on preparation pipeline to fit their needs

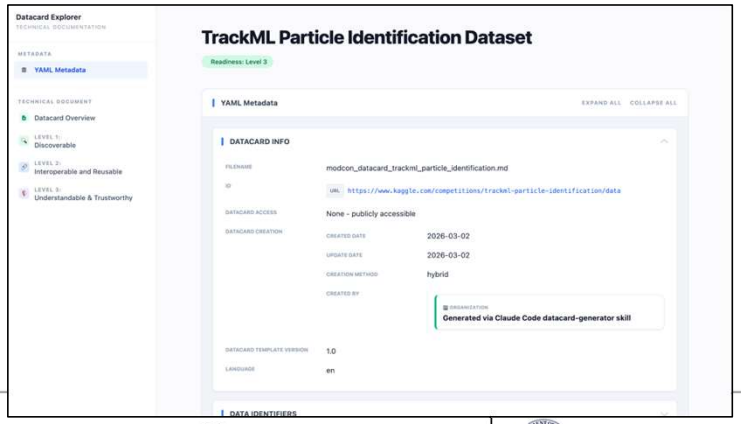

Base/Eval: Model Assessment and Evaluation


Stefan Wild
(LBNL)


Eric Chagnon
(LBNL)


Neil Getty
(ANL)


Fernando Llorente
(BNL)


Zarija Lukic
(LBNL)


Emily Saldanha
(PNNL)


Shashank
Subramanian (LBNL)


Fred Suter
(ORNL)


Angel Yanguas-Gil
(ANL)

+ AmSC & BPSW embedded liaisons
(Wahid Bhimji, Nhan Tran, Rafael Zamora-Resendiz)

Developing tools, benchmarks, & protocols that assess robustness, generalizability, reproducibility, & scientific validity under evolving data and uncertain ground truths

- Track multi-metric progress over time
- Serve as peer reviewer, ultimately with SME-HitL
- Integrating with AmSC UI/Infrastructure

Key principles

- Open-source framework, multiple techlanes
- Address non-language modes, domain-specific info
- Build to evolve and scale

Base/Eval: Delivering a Comprehensive Evaluation Suite

Model
(or agentic pipeline)

Benchmark
Tasks, data, metrics, …

Registered on AmSC AMERICAN SCIENCE CLOUD

GM Proving Ground
Eval Harness

`MLFlow, NeMo, lm-evaluation-harness`

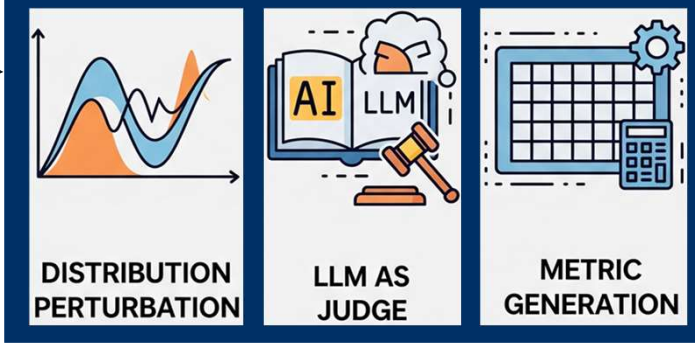

Performance Metrics
accuracy, precision, recall, F1, ROC-AUC, confusion matrix, …

Visualizations

Logs

Error Analysis & Auditing
anomalous responses, tail instances, constraint violations, …

SME-HitL review | Leader board | … | S.I.M

Base/MM: Multimodal Scientific Reasoning Models

Sandeep Madireddy
Argonne National Lab

Adarsha Balaji
Argonne National Lab

Guojing Cong
Oak Ridge National Lab

Patrick Emami
National Lab of the Rockies

Ben Erichson
Lawrence Berkeley National Lab

Neil Getty
Argonne National Lab

Sam Wheeler
Argonne National Lab

Fangfang Xia
Argonne National Lab

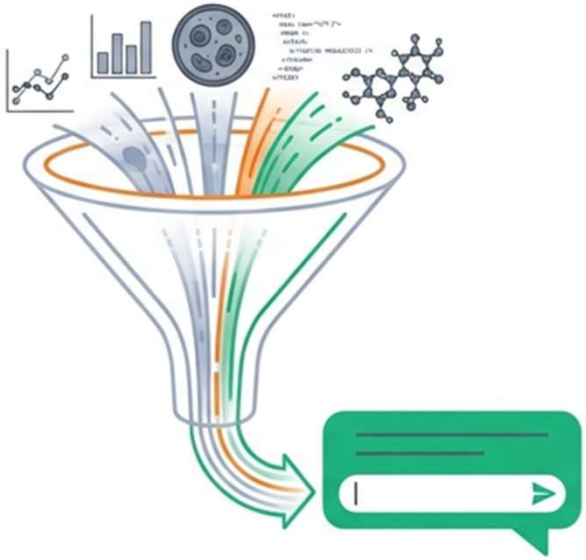

Goal: *Develop multimodal scientific reasoning models and an intelligent conversational interface that accepts inputs across modalities, performs cross-modal reasoning, to answer complex scientific questions*

AmSC AMERICAN SCIENCE CLOUD

At-Scale Services / MLOps

Software: https://github.com/AI-ModCon/BaseMM_PRISM/

Contact: Sandeep Madireddy (smadireddy@anl.gov)

Base/MM: Multimodal Scientific Reasoning Models

PROJECT GOALS & CHALLENGES

Goal: Cross-modal reasoning across text, images and other structured data for scientific discovery

Capability: Utilize the best-in-class modality/domain encoders and enhance multimodal reasoning - tool-free capabilities and tool-based discovery and tasks.

Challenge: Bridging the gap between all diverse modalities, data quality, and enable flexible dialogue.

PRISM HUB & SPOKE ARCHITECTURE

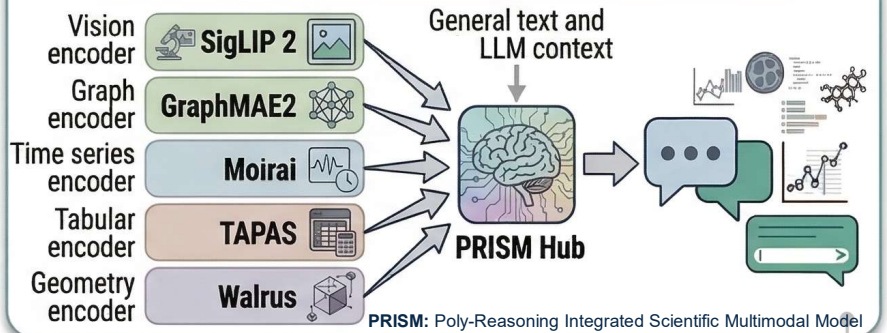

PRISM: Poly-Reasoning Integrated Scientific Multimodal Model

CURRENT STATUS

- **Scaling Training on Aurora, Perlmutter & Frontier**
  - 5x Throughput increase, efficient communication and memory optimization using FSDP, 85 samp/s → 420 samp/s via Native DDP
- **Projector Training & Modality Alignment** (Stage 1 & 2)
- **Instruction Tuning:** Joint Multimodal SFT (Stage 3)
- **RL & Safety Pipelines with Reward/Policy Models** (Stage 4)
- **Early results on reasoning:** Image & Time Series modalities

ONGOING RESEARCH & PLANS

- **Data Curation...**
  - ✓ Data Curation
  - ✓ Joint SFT
  - ✓ Ablating Core
  - ✓ Joint SFT
- Standardize multimodal data ingestion, evaluation/safety, SIM interfaces
- Model-generated Synthetic Data for future training stages (feedback loop from SFT)
- **Evaluate multimodal/cross-modal reasoning**
- **Prototype** end-to-end 'GM Proving Ground' working on AmSC

Base/Safe: Model Safety and Security

Nathan Hodas
(PNNL)

Robin Cosbey
(PNNL)

David Florey
(PNNL)

Natalie Isenberg
(PNNL)

Steven Goldenberg
(JLAB)

Olivera Kotevska
(ORNL)

Reilly Raab
(PNNL)

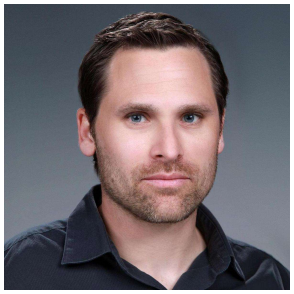

Malachi Schram
(JLAB)

Developing rigorous frameworks, scalable methods, and institutional safeguards that guarantee safety and security across the scientific AI lifecycle

Base/Safe: Agentic Framework with Flexible Execution Strategy for Model Teams

On Model Team (MT) laptop: Either through command-line or embedded within VS Code IDE. Provides rich analysis tools directly in coding environment.

As a Continuous Integration (CI): Enables MT's to focus on functionality while the Base/Safe Agents focus on providing safety, uncertainty, and explainability insights.

Within testing environment: This fully integrated approach enables more real-world analysis and insights.

Execution instructions

Design, training methodology, purpose, etc.

Model architecture, weights

Source code

Input / output pairs

Fingerprinting Agent

Utilizing LLMs on the model artifacts provided, create a structured dossier for use by downstream agents.

Uncertainty Quantification Agent

Perform baseline uncertainty quantification to provide insight to the model team.

Red and Blue Teaming Agents

Find vulnerabilities and exploit or defend against them to help better secure the model.

Explainability Agent

Enable the model team to better understand how the model is working with given inputs/outputs.

Report w/ Action Suggestions

Base/SIM: Enabling Self-Improving Models

APEIRON
SELF-IMPROVING MODELS

Self-improving AI Models

Agentic framework for keeping domain science models up to date with data from new experimental and simulations

Software: https://github.com/AI-ModCon/BaseSIM_APEIRON
Contact: gainarua@ornl.gov


Ana Gainaru
ORNL


Steffen Schotthoefer
ORNL


Anna Quach
INL


Zilinghan Li
ANL


Krishnan Raghavan
ANL


Rafael Zamora-Resendiz
LBNL


Auralee Edelen
SLAC


Andrew Ayres
ORNL


Alvaro Sanchez-Villar
PPPL

Base/SIM: APEIRON's Continuous Self-Improvement

Goals

Keep domain-specific AI models synced to new experimental data and methods that emerge

Monitor workflow runs
Detect data and model drift
Trigger continual learning
Monitor performance

Status

APEIRON framework has been developed and can connect to domain scientists through API calls

In progress: Deployment at scale
In progress: SKILLS / MCP tools

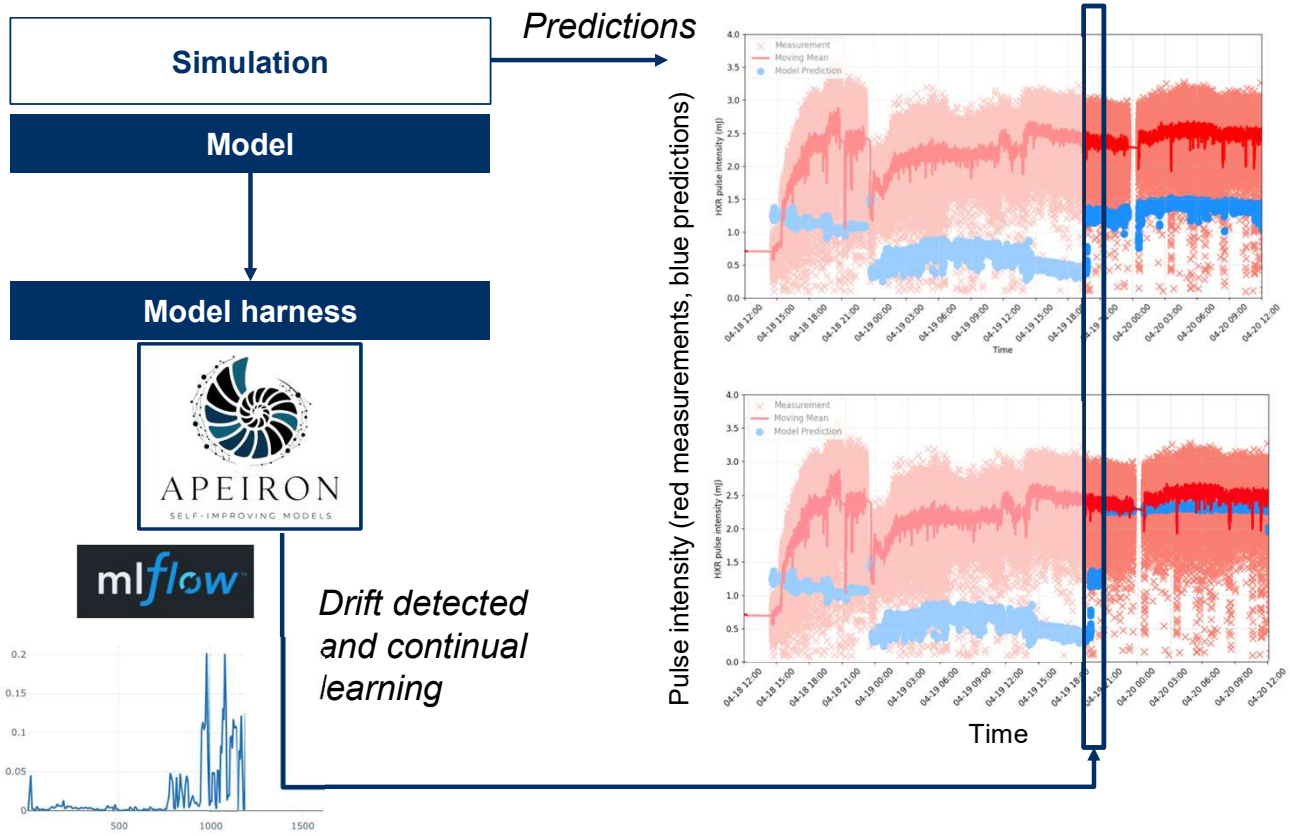

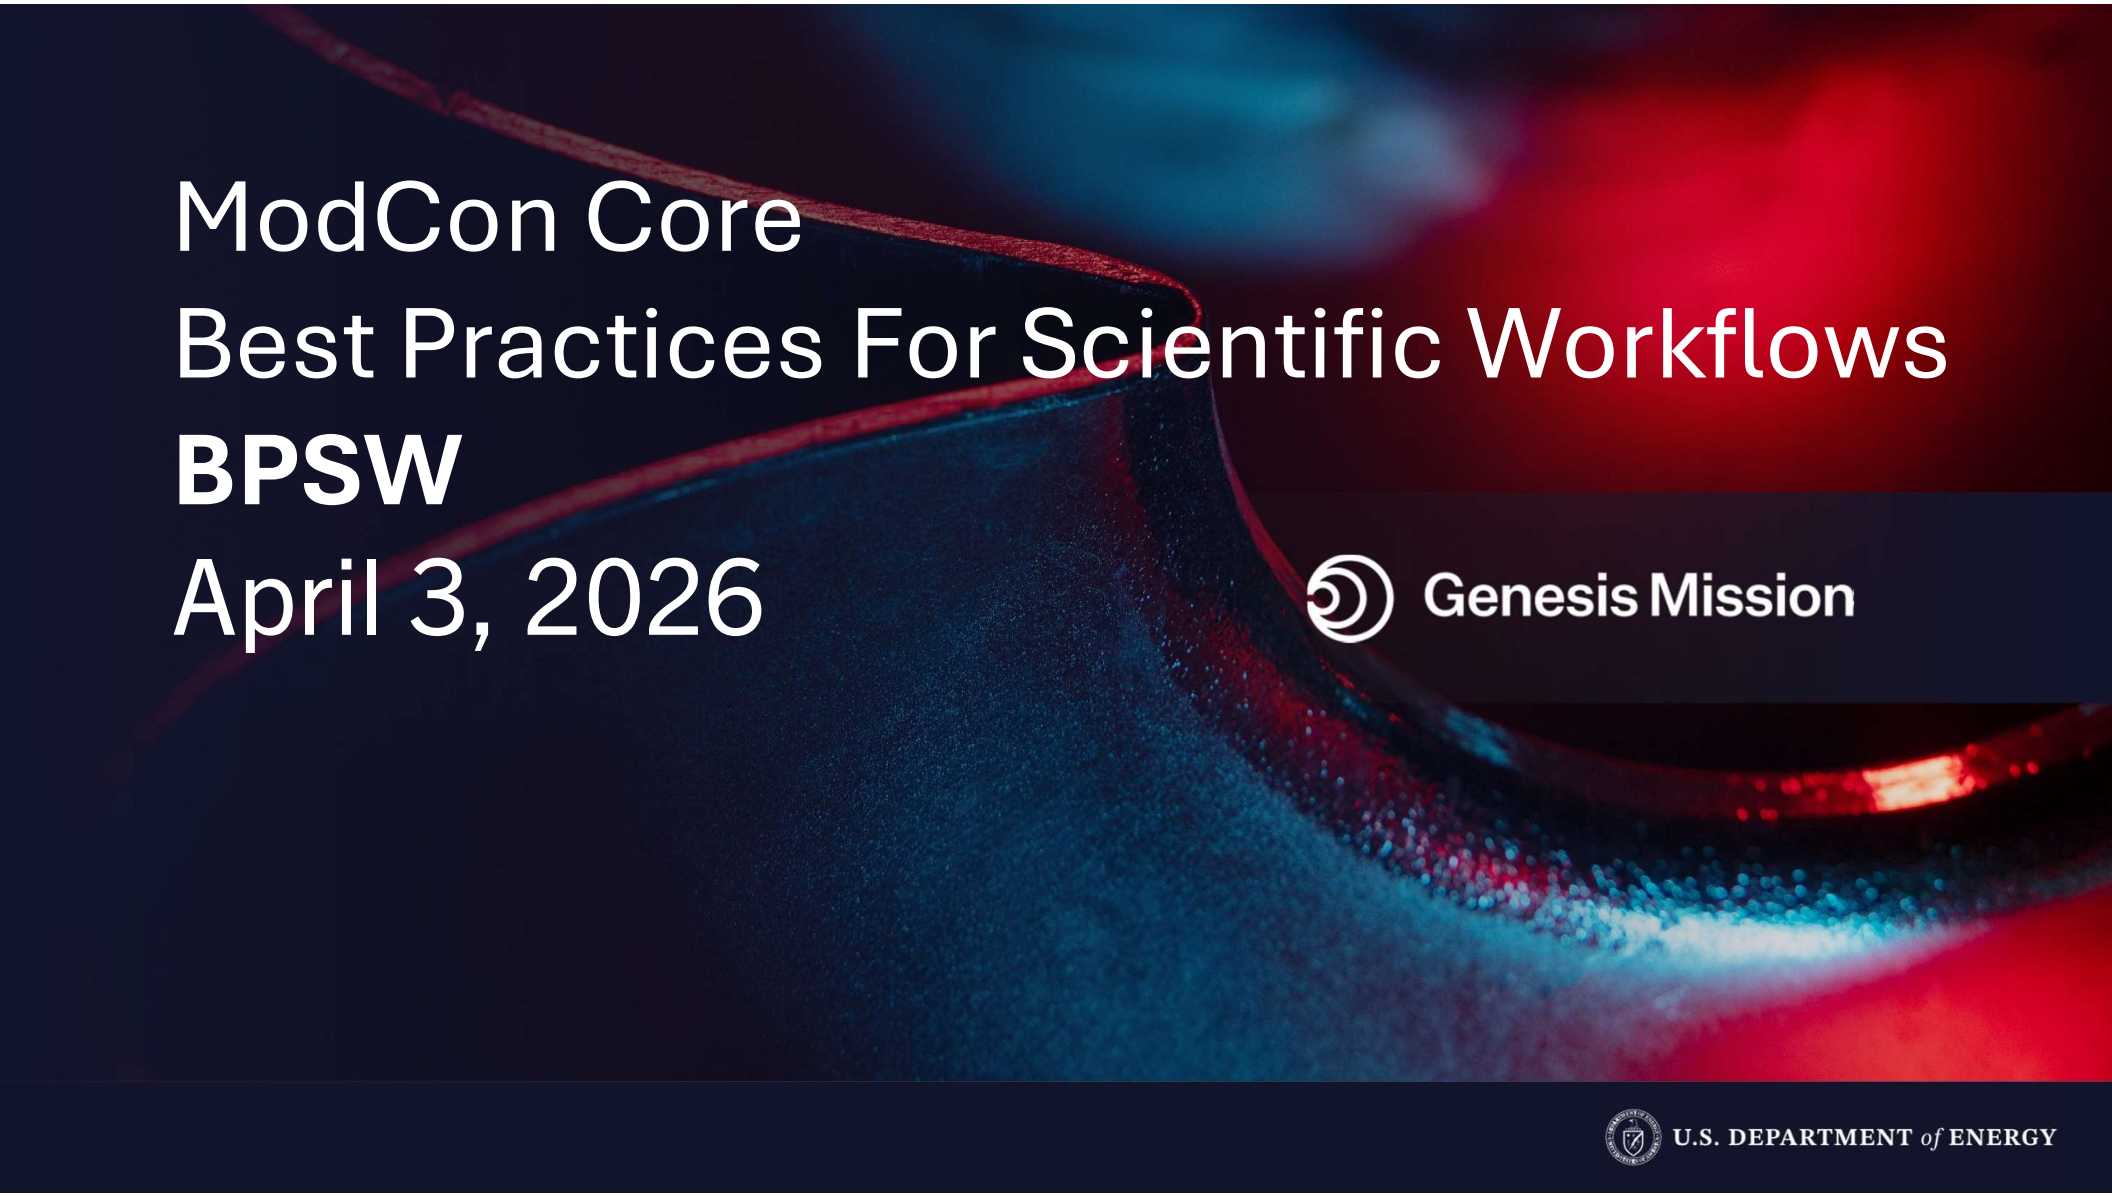
ModCon Core
Best Practices For Scientific Workflows
BPSW
April 3, 2026
Genesis Mission
U.S. DEPARTMENT of ENERGY

BPSW 23 Members from 11 DOE Labs


Franck Cappello, **Co-lead,** ANL


Rafael Ferreira da Silva, **Co-lead**, ORNL


Robert Underwood, **Extended Team Lead,** ANL


Massimiliano Lupo Pasini, **Extended Team Lead,** ORNL


Rinku Gupta, **Extended Team Lead**, ANL


Neil Getty, ANL


Murat Keceli, ANL

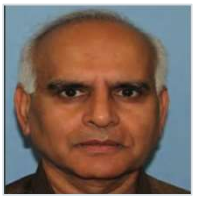
Nagi Rao, ORNL


Jong Youl Choi, ORNL


Helia Zandi, ORNL


Nathan Tallent, PNNL


Natalie Isenberg, PNNL


Anna Giannakou, LBNL


Dan Gunter, LBNL


Mohammad Abdo, INL


Long Qi, AMES


Mikhail Titov, BNL


Joseph Smith, NLR


Nhan Tran, FNAL

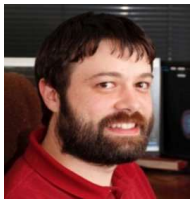
Thomas Britton, JLAB


Shantenu Jha, PPPL


Thorsten Hellert, LBNL


Amith Murthy, SLAC

Scientific Workflows

BPSW considers every steps of the whole research workflow from the research questions to report generation.

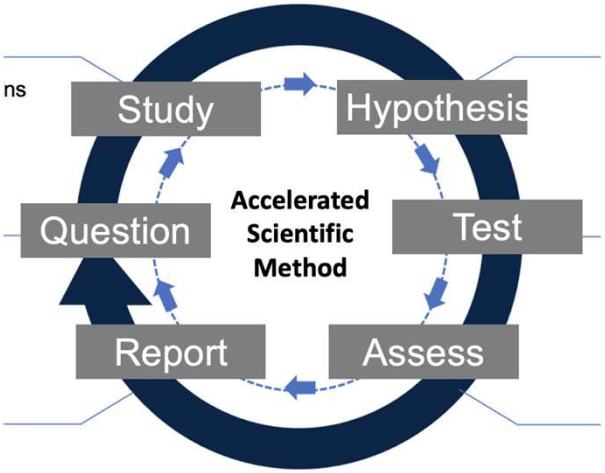

https://doi.org/10.1038/s41524-021-00765-z

Some teams implement the entire cycle.
Most teams focus on a subset of steps
(e.g. AI4HPC: LLM-based code translation workflow)

MOAT March Demo

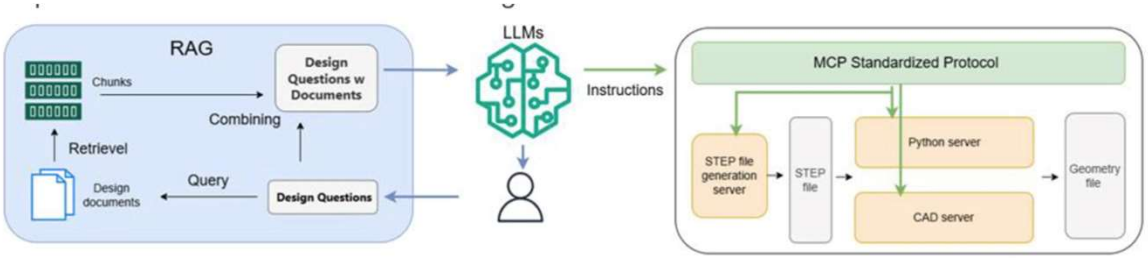

End-to-End Loop: couples an LLMS to a RAG, MCP servers.

1. User's query is passed to the RAG to retrieve relevant documents (Study).
2. The LLM uses this context to generate a plan, which involves calling simulation tools via the MCP. (Plan Generation)
3. The MCP executes the simulation and returns results (e.g., a STEP file for a CAD model). (Test)
4. The process is iterated with human oversight. (Assess)

BPSW Specific Goals

1) Catalog the practices of AI integration in scientific workflows at DOE

2) From the catalog, Identify, catalog and develop training about best practices, the integration methods that provide the most acceleration

3) Identify the 20 most important tasks to accelerate with AI in scientific workflow

4) Identify where AI models break and define research directions to mitigate this
(e.g. LLMs, Hypothesis Generation Multi-Agents)

5) Define AI integration templates that DOE workflows not using AI can use for acceleration

Examples of Workflows for March demos

Microelectronics

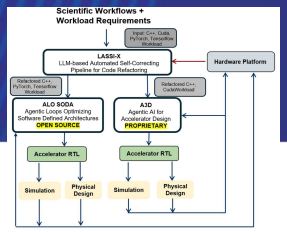

Q2C

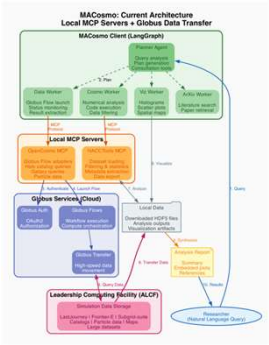

AXESS

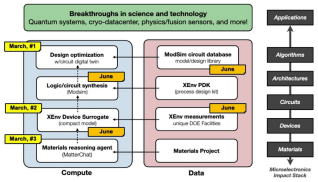

Synapse-I

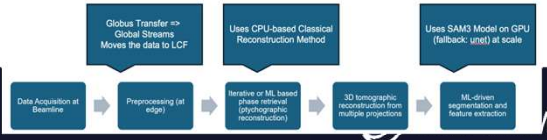

MOAT

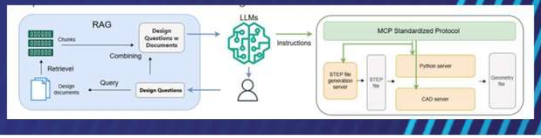

Prometheus

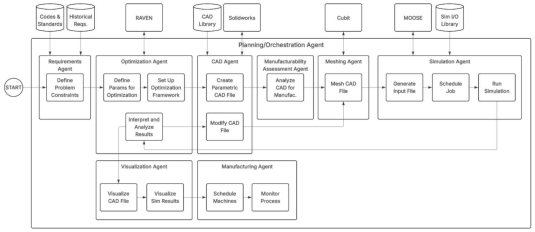

COMB-Flow

MAIQMAG

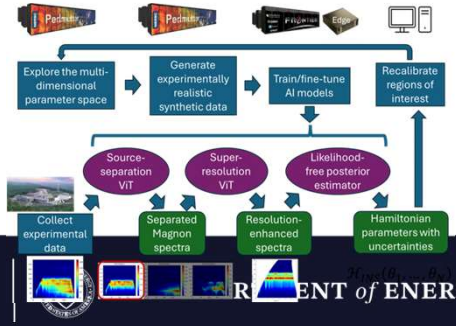

Catalog of Current Practices

Ran a survey collecting current practices within the DOE complex

144 workflow responses so far including AI

- We are targeting ~180-200 responses, but expect more
- Several labs are just now distributing the survey

This includes detailed Reponses for **12 model teams** (not included: AI4Quantum, AI4HPC, and new model teams added Jan'25)

Developed a dashboard from the workflow survey responses, and working to deploy onto AmSc next week.

Chart shows AI model type usage across the surveyed scientific workflows

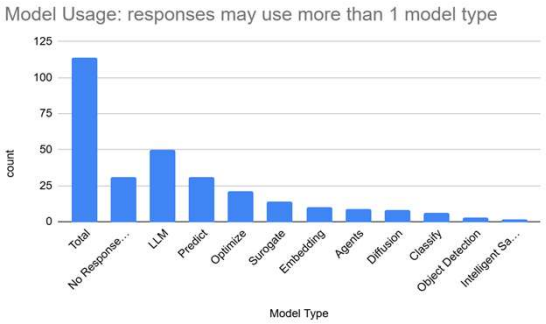

Responses only for the Model type question

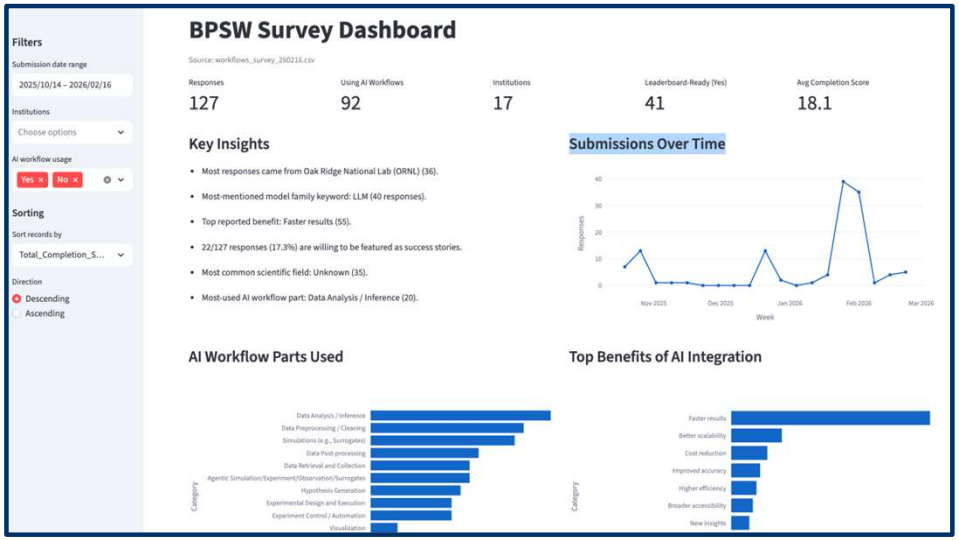

Best Practices

Proven, evidence-based techniques or procedures that consistently deliver superior results and efficiency. Act as standards or benchmarks designed to minimize risk and optimize performance.

For AI-Based Workflows: "**Patterns**" of reusable AI integration techniques that have four elements:

- **Problem** (steps in the workflow) where AI-integration is it useful
- **Techniques** how to solve the problem (what AI, what training, what integration)
- **Consequences** of using the pattern (what benefits. e.g. speed, what risks)

Include policies (e.g. safety, security)

Goal in our case: **Provide guidance** on tasks such as AI design, model training, evaluation, integration, and deployment

Artifacts: user-guide, templates, trainings, and other resources

Adapted from Gamma et al, "Design Patterns for Object Oriented Software"

Some **Patterns** Identified in MT March Demo Workflows:

- LLM establishing a plan
- LLM calling multiple specialized agents
- Globus transfer to acquire, transfer data
- RAG requests to retrieve data
- Generation/Correction loops
- Calling surrogates for optimization.
- AI for data cleaning
- LLM as a judge for validation outputs

Prioritized AI Research Needs from the Model Teams

How was this list developed?

- Each team provided up to their top 5 needs
- Similar responses were grouped
- Ranked by the most top needs, most second needs etc.

Observations

- Currently based on a small subset of teams
- We expect that some tasks (e.g. #12) will become more important as teams gain experience
- We expect that more tasks will be identified as we expand to more teams

Next

- Collect more responses (send reminders, etc.)
- Participation in the International Forum on Detectors for Photon Science in Kyoto (April 2026)

Priority Research Tasks

1. Development of Surrogates (7 teams)
2. Code Development (5 teams)
3. Adversarial Evaluation (5 teams)
4. Optimization (4 teams)
5. AI+Modsim Orchestration (2 teams)
6. Data Sharing (2 teams)
7. Segmentation (1 team)
8. Data Integration (2 teams)
9. Hypothesis Generation (2 teams)
10. Real Time Visualization (2 teams)
11. Integrated Research Infrastructure (1 team)
12. **Data Extraction and Cleaning (4 teams)**
13. Scientific/Technical Data Search (1 team)

Model/Agent Cards

https://gitlab.osti.gov/genesis/modcon-bpsw

Model/Agent Cards capture the capabilities, constraints, and intended use of DOE and community AI models & agents

Co-designed with stakeholders (AmSC, BASE-CAF, DBS, and leadership) to ensure consistency and adoption

- Scale is real: **330+ cards** already assembled and curated
- Broad coverage: spanning **30+ mission-relevant topics/domains**

Living inventory: continuously expanding as new models and agents emerge

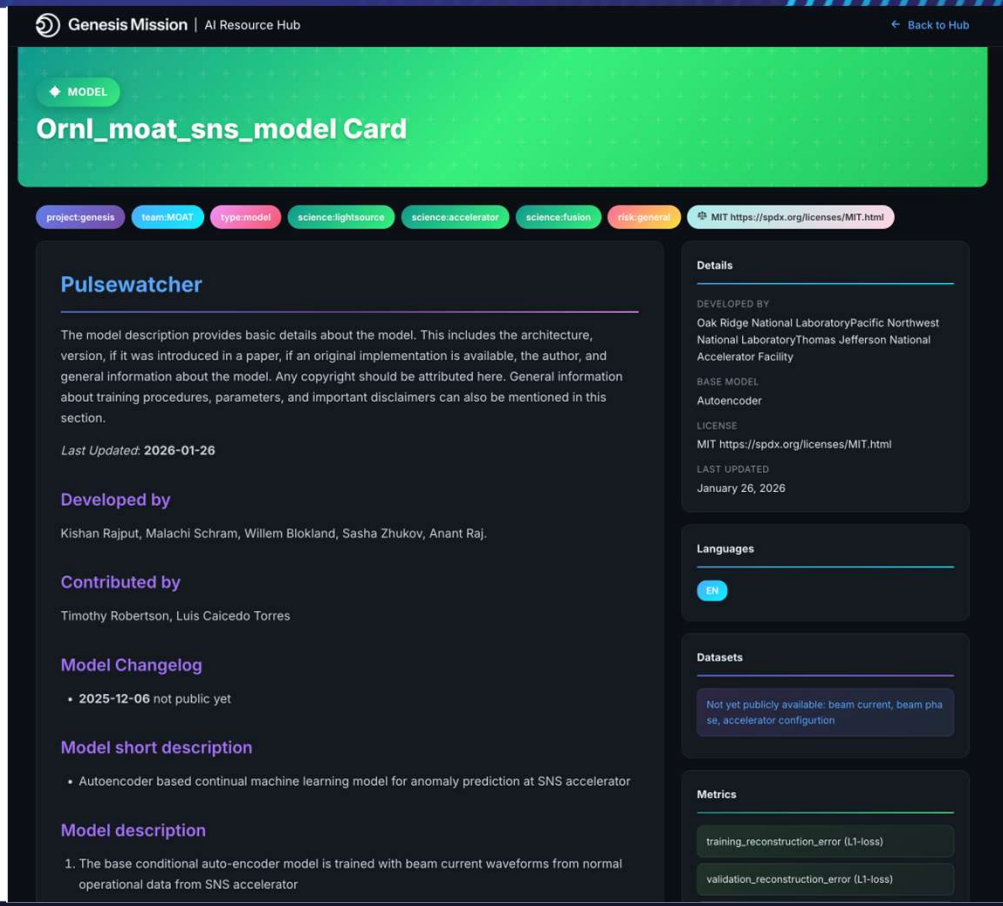

Leaderboard

The leaderboard will expose the performance of different AI models on a selection of scientific workflow **tasks** identified by the MTs and other DEO teams as critical for their scientific workflow accelerations.

The leaderboard will follow the general format of the benchmarks defined by ML common: https://mlcommons-science.github.io/benchmark

Rigorous specifications of benchmark quality:

- Software environment
- Problem Specification and Constraints
- Dataset (FAIR)
- Performance Metrics
- Reference Solution
- Documentation

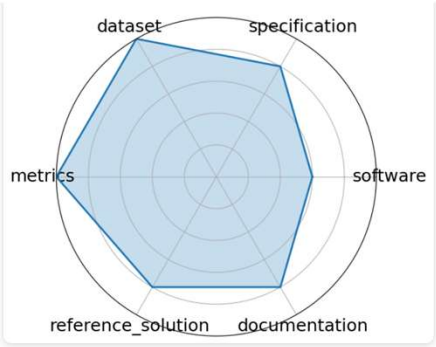

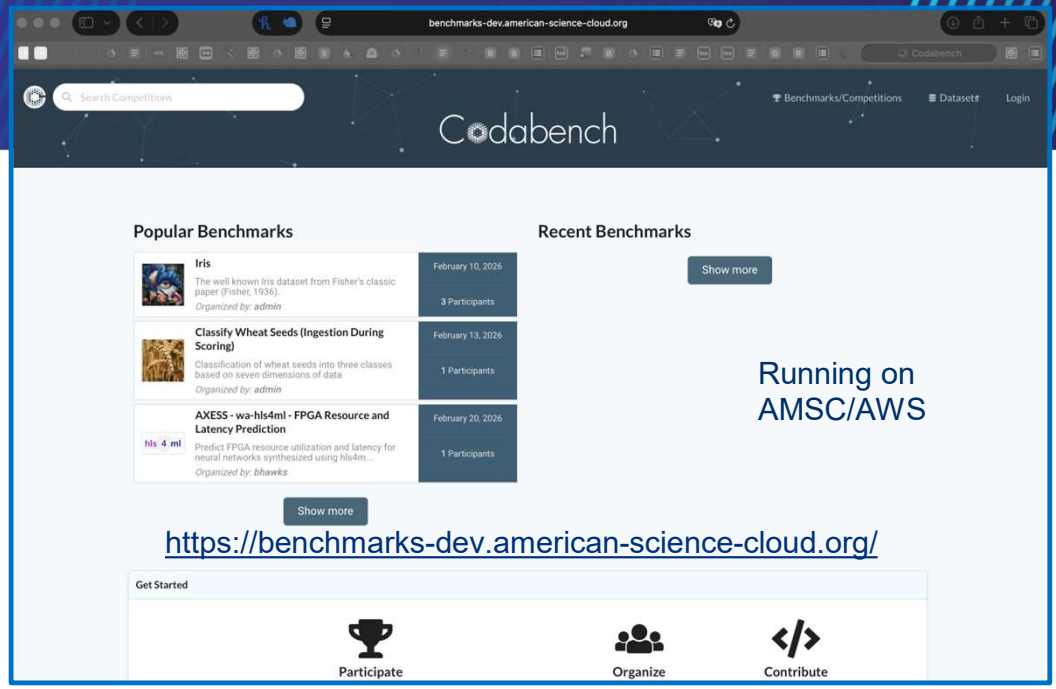

Hosted on the American Science Cloud (AWS)

Examples of benchmarks

- [LLMs] SciCode, CritPT
- [Agentic Evaluation] MOAT team
- [Surrogate Models] AXESS team
- [Segmentation] SYNAPS-I team

AI Discovery Engine

**Developed in partnership with PNNL

Natural language interface to compose and interrogate workflows of DOE and NNSA AI assets to solve national S&T challenges

~300 DOE and NNSA model and agent cards inform that possible workflows

Data and compute resources cards will be incorporated soon

AMSC deployment: to be hosted and operated within the AMSC environment

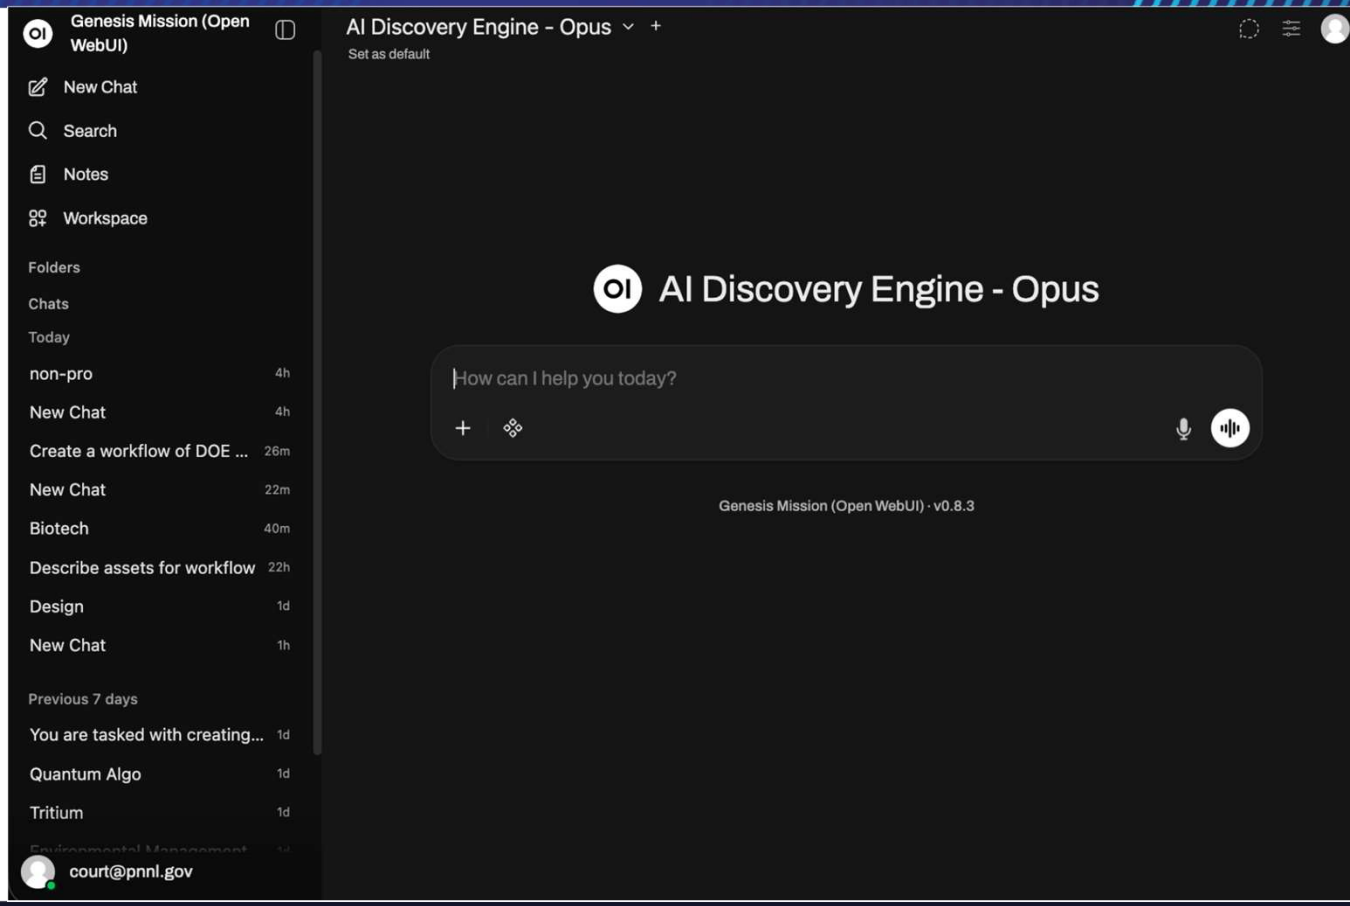

Establishing Scaling Laws Best Practices for MTs

- MTs are developing AI models where efficient allocation of compute and data is critical.
- Scaling laws predict the performance relationship between model size, dataset size, and compute.
- BPSW Neil Getty and BASE Nathan Hodas are establishing standardized best practices for Scaling Law analysis.
- Conducted initial viability audit of MTs using technical reports, model cards, and architecture documentation.
- Will produce Scaling laws documentation and training materials.

Tier	Name	Criteria
Tier 1	Scaling-Ready	Custom architecture trained from scratch; large dataset (>1B tokens); parametric code; well-defined loss; baselines available
Tier 2	Data-Scaling Only	Fine-tuning pre-trained models; cannot modify base model size; data scaling viable
Tier 3	Remediation Required / Not Applicable	Unclear metrics, no baselines, hard-coded architecture, workflow team, or not a trainable model

Metric	MTs
Tier 1 (Scaling-Ready)	7 (50%)
Tier 2 (Data-Scaling Only)	5 (36%)
Tier 3 (Remediation/N/A)	2 (14%)
High Confidence	8
Teams with >1B Tokens	6
Teams with De Novo Training	9
Teams Reporting FLOPs	1
Workflow/Agent Teams	2

What to consider in RFA proposals with respect to BPSW integration

Workflow patterns: Proposal can leverage, if relevant, one or more of the workflow patterns identified by BPSW.

Priority tasks: Proposal can focus, if relevant, on one (or more) of the priority tasks to accelerate with AI

Model/Agent Cards: Proposal can produce, if relevant and authorized, Models/Agent cards to include in the resource hub

Leaderboard: Proposal can produce benchmarks that will be included in the leaderboard

Scaling laws: Proposal can leverage the scaling law training material to include projections.

ModCon Core
Data Brokers and Standards
DBS

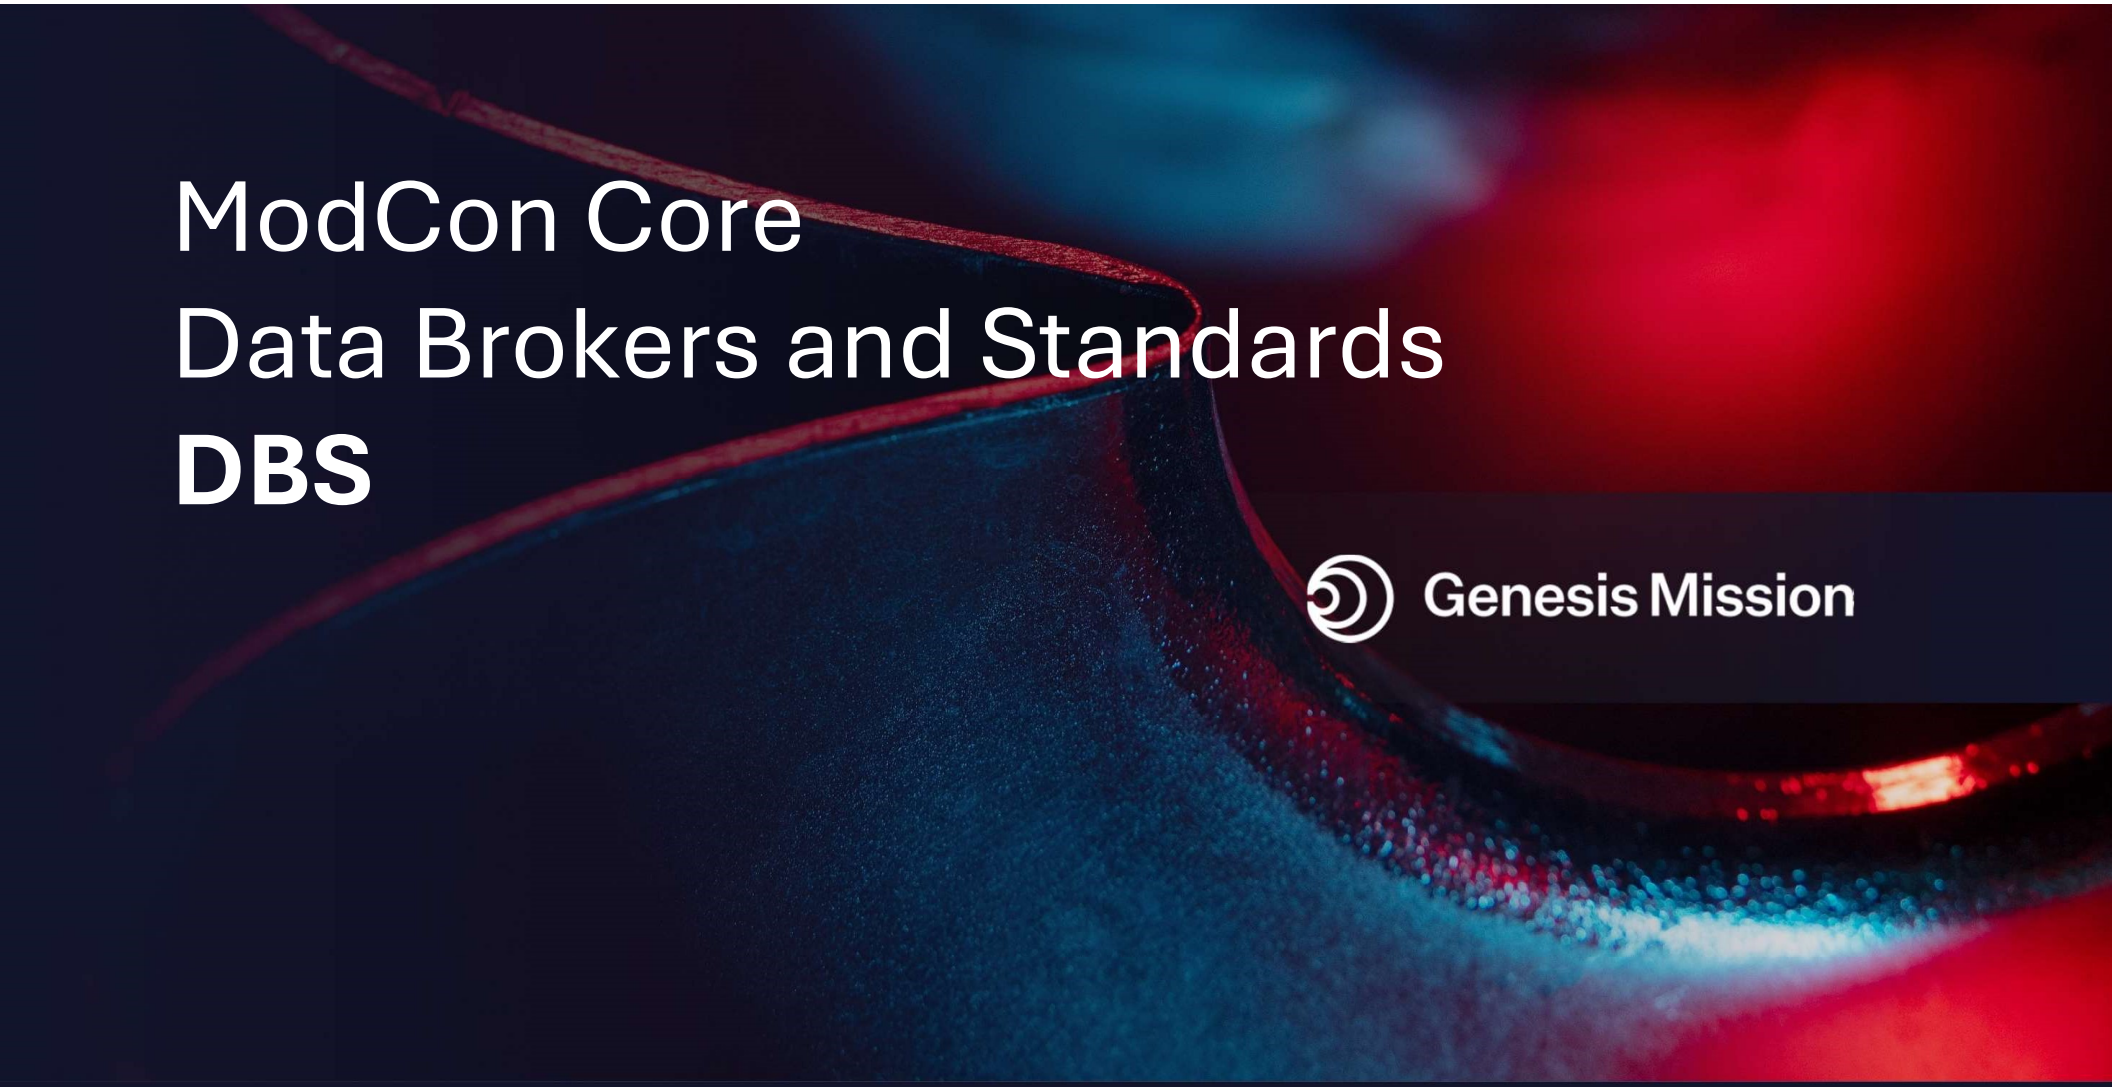

Genesis Mission

U.S. DEPARTMENT of ENERGY

DBS Enhances the Productivity of Science

DBS addresses current rate-limiters to scientific productivity and security by...

- improving the findability of data and related artifacts for people and agents
- developing secure, trustworthy frameworks for data sharing and publication
- automating data curation for AI-readiness

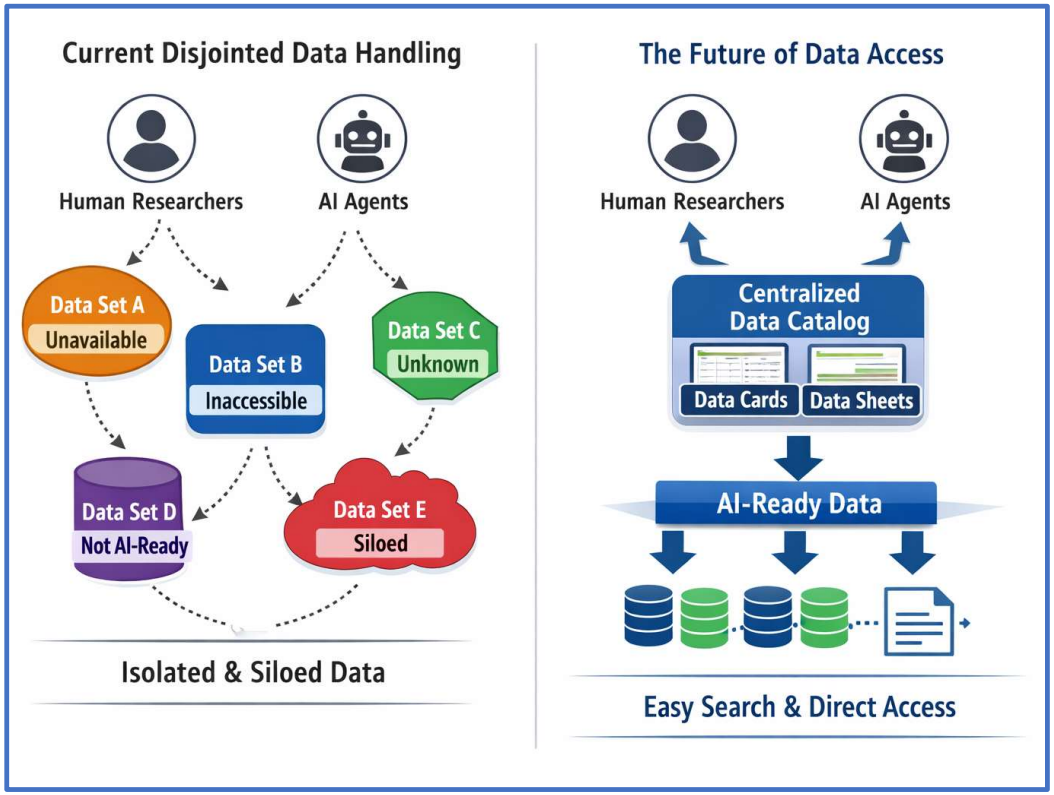

DBS Tasking

Current DBS work is organized into 5 Tasks:

- Task 1: Register Open, Curated Data in AmSC Catalog
- Task 2: Data Documentation
- Task 3: Curate, Publish, and Register New Data with AmSC
- Task 4: Automate Data Curation
- Task 5: Governance

Task 1: Register Open, Curated Data in AmSC Catalog

- 8 catalogs currently accessible through Genesis Central Catalog containing petabytes of data
- Catalogs and data repositories identified across all scientific domains

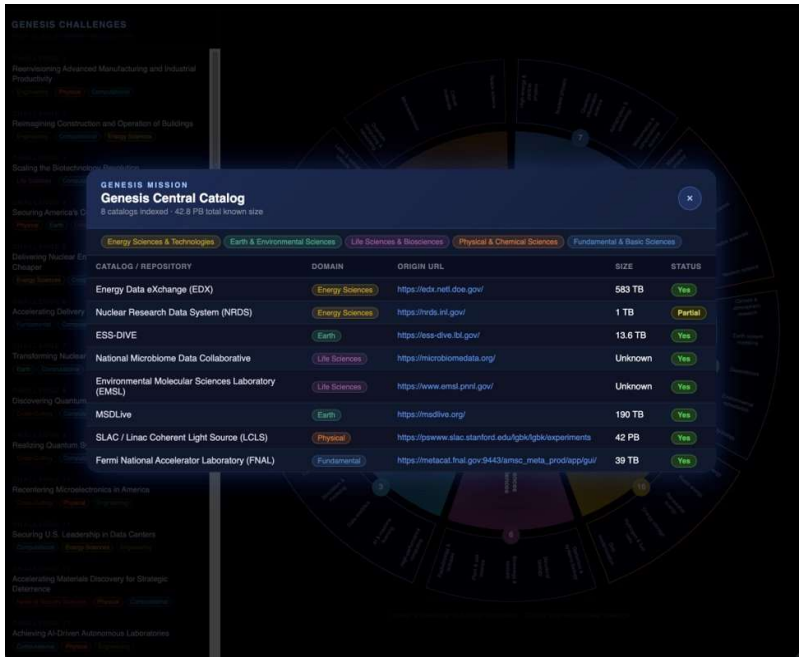

Task 2: Data Documentation

- V1 of metadata, data card, and data sheets available for use
- Tools and skills created for data card generation and viewing

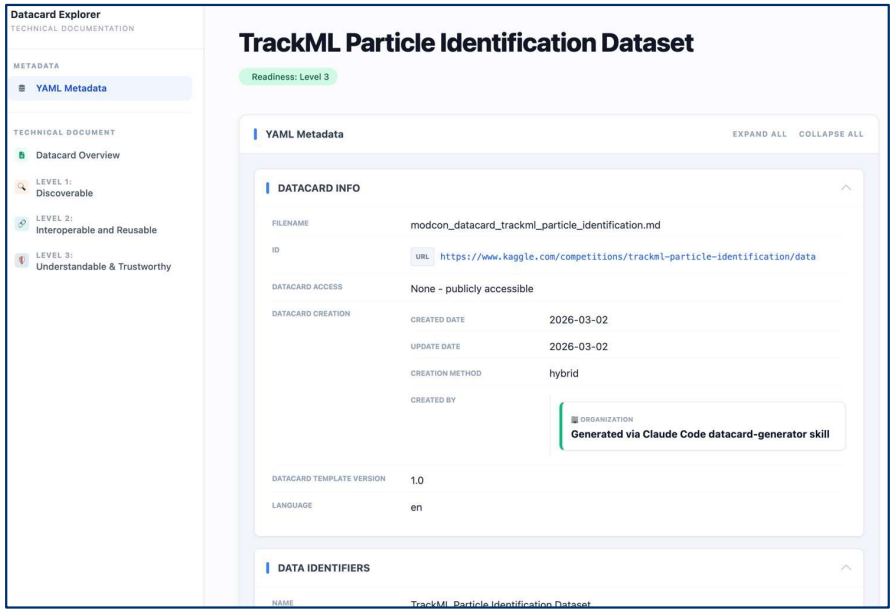

Data Card Viewer

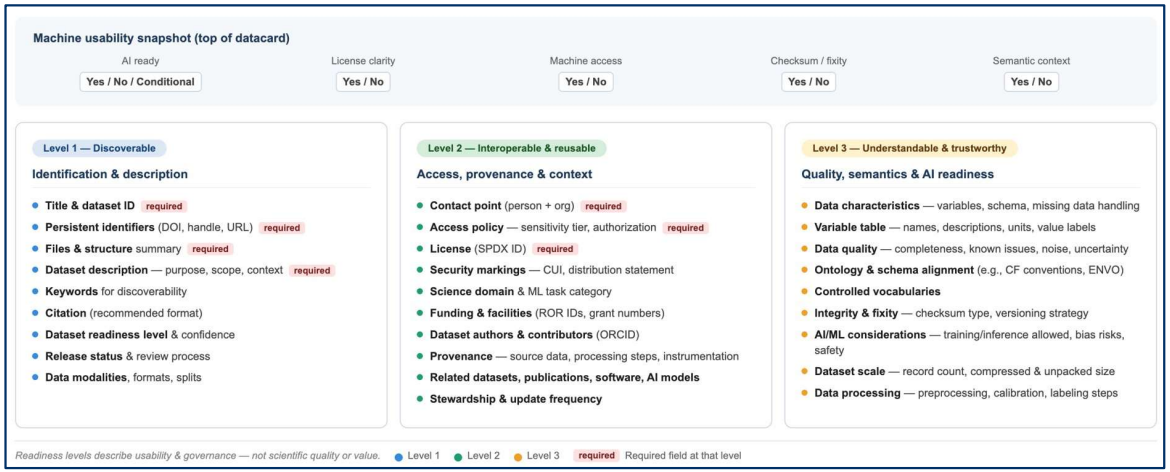

Metadata v1 Summary

Task 3: Curate, Publish, and Register New Data with AmSC

- AI-Readiness workshop series launched
  - Identify datasets, assist in curation, certify for Genesis Mission use

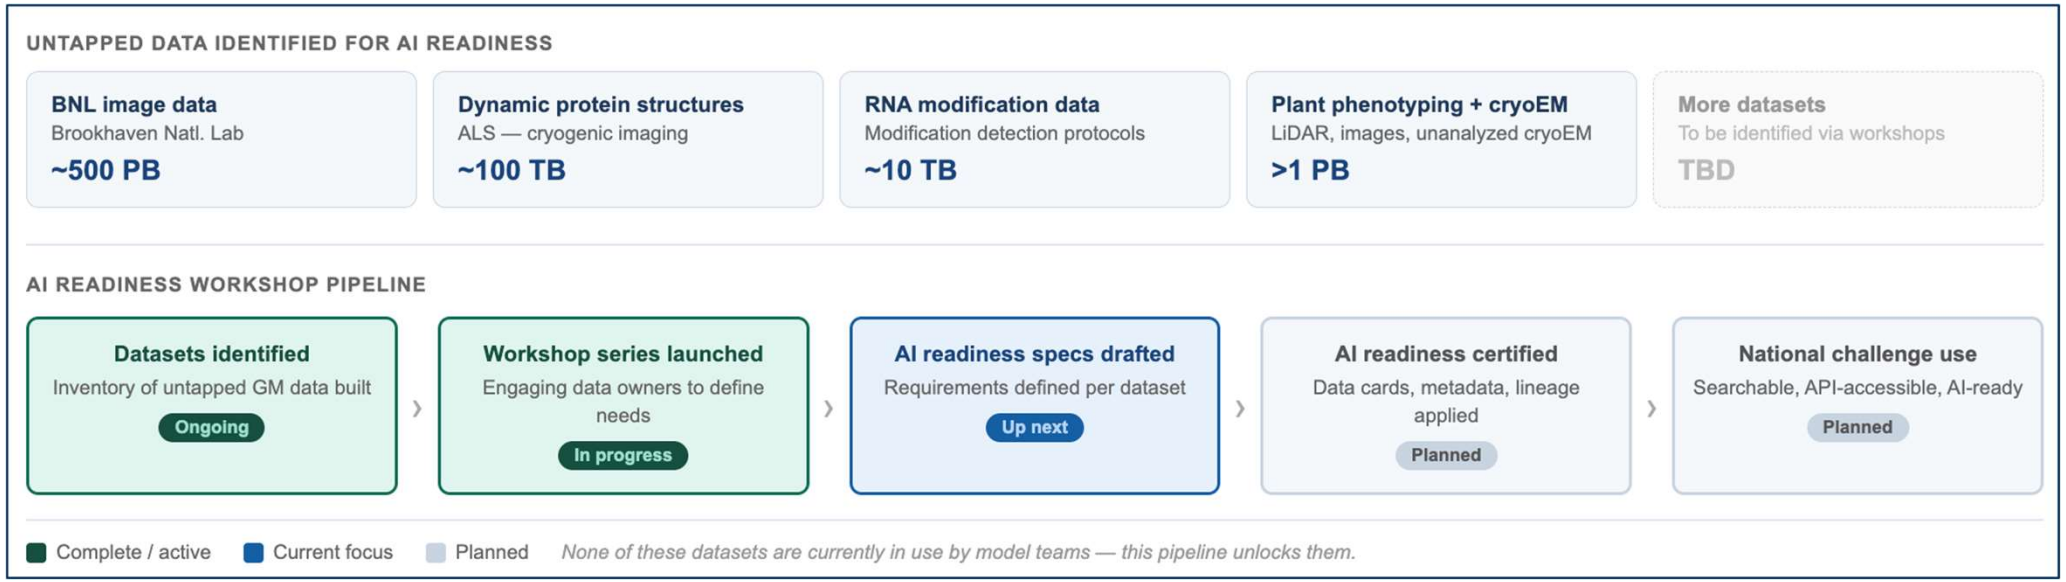

Task 4: Automate Data Curation

- Testing AI-readiness assessment tools on registered data to inform tool choice and usage

Tool	Python API	CLI Interface	Tabular Data	Scientific / Domain Formats	Distributed Processing	Data Quality Metrics	FAIR Compliance Assessment	Fairness / Bias Metrics	Privacy Metrics	Custom Metrics / Extensions	Automated Data Transformation	LLM / AI-Assisted Features	Metadata / Schema Management	Multi-dataset / Batch Support	GPU Acceleration	Domain Agnostic
YData Profiling	Yes	Yes	Yes	No	Partial	Yes	No	Partial	Yes	Partial	No	No	No	No	No	Yes
Evidently	Yes	No	Yes	No	No	Yes	No	No	No	Yes	Partial	Yes	No	Yes	No	Yes
Anemoi-datasets	Yes	Yes	No	Yes	Yes	No	No	No	No	Partial	Yes	No	Yes	Yes	No	No
PhysicsNeMo Curator	Yes	Yes	No	Yes	Partial	Partial	No	No	No	Yes	Yes	No	No	Yes	Yes	No
NVIDIA NeMo Curator	Yes	Partial	No	No	Yes	Partial	No	No	Yes	Yes	Yes	Yes	No	Yes	Yes	Partial
draviz	No	Yes	N/A	No	N/A	No	No	No	No	Partial	No	No	No	Partial	No	Partial
DREAMER	Yes	No	Yes	No	Partial	Yes	No	No	No	No	Yes	No	No	No	No	Partial
Great Expectations	Yes	Yes	Yes	No	Partial	Yes	No	No	No	Yes	No	No	Yes	Yes	No	Yes
Soda Core	Yes	Yes	Yes	No	Yes	Yes	No	No	No	Yes	No	No	Yes	Yes	No	Yes
F-UJI	Yes	No	No	No	No	No	Yes	No	No	Yes	No	No	Yes	No	No	Yes
AIDRIN	Yes	Yes	Yes	Yes	Partial	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Partial	No	No	Yes
DSAGT	No	Yes	Yes	Yes	No	No	No	No	No	Yes	Yes	Yes	Partial	Yes	No	Yes
Bridge2AI	No	No	No	No	No	No	Yes	Yes	Partial	Partial	No	No	Yes	N/A	No	No
REDI	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	Partial	Yes	Yes	Yes	Yes	Yes	No	Partial

* Table generated by Sonnet based on collected report from the evaluation by the team. Missing or undocumented information might introduce mismatching results. Some tools are under heavy development (e.g., DSAGT, AIDRIN, REDI) or have not been published yet (e.g. REDI).

Task 5: Data Governance

- Established a Technical Data Working Group, in conjunction with partnerships teams
- TDWG is establishing federated metadata harvesting, AI-centric ontologies, FAIR-aligned infrastructure, best practices for data management plans, and specialized training and documentation to support AI-driven research

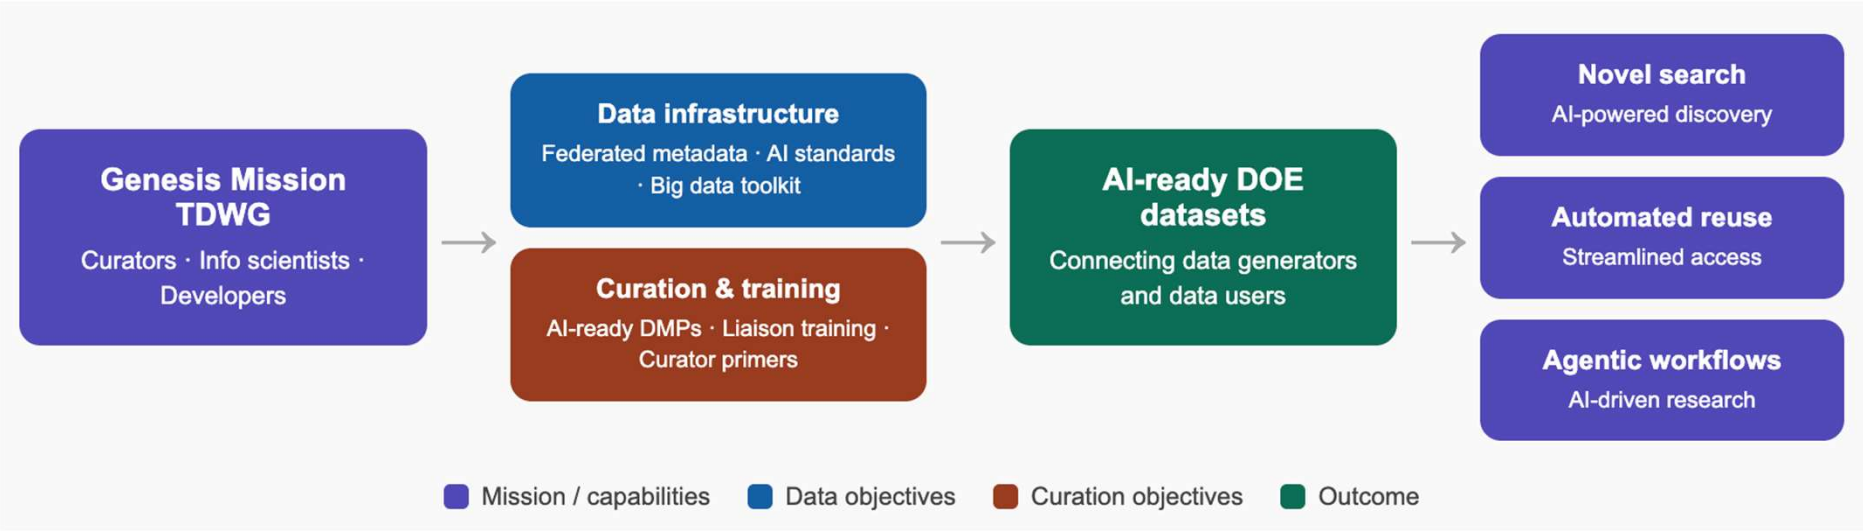

Putting it all together: Document-based search, access, reuse

Transforms DOE datasets into AI-ready resources by **centralizing access to metadata, datacards, and datasheets** and establishing clear standards and policies for data usage.

Bridges data generators and data users and enables:

- Novel search capabilities
- Automated access and reuse
- Integration into agentic workflows

metadata
- Highly structured (JSON, XML, RDF)
- Often standardized fields
- Possible to automated
- Enables data to be found, indexed, managed, and understood by software systems/AI

agent cards

model cards

datacard / README
- Introduces and explains dataset background, context, organization & file structure, data details, methodological info, access & reuse
- Possible to automated
- Human & AI-readable
- Adjacent to files in the data repo
- Published & unpublished datasets

datasheet
- Comprehensive documentation of creation, scope, intended use, limitations, & ethical/technical considerations
- Intentionally human-written (possibilities to automate an initial draft)
- Capture tacit information & community knowledge
- Enables frontier models to leverage rich, unstructured dataset information
- For published, **elite** datasets

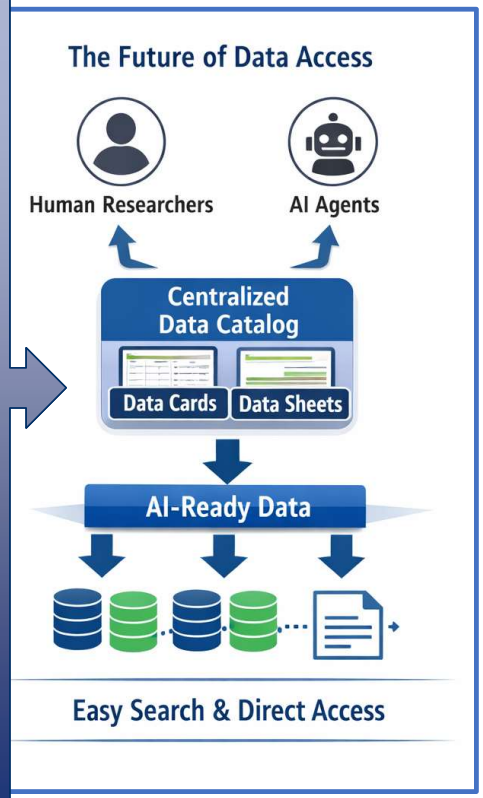

DBS Members

DBS Leads

Laura Biven JLab
Jeren Browing INL
Ray Grout NLR

Genesis Data Pillar Leads

Ana Kupresanin LBNL
Dana Grisham SNL

Le	Chen	ANL
Nicholas	Schwarz	ANL
Rajeev	Thakur	ANL
Yi	Huang	BNL
Yihui (Ray)	Ren	BNL
Adam	Lyon	FNAL
Oliver	Gutsche	FNAL
Victor	Walker	INL
Brad	Sawatzky	JLAB
Diana	McSpadden	JLAB
Jai	Sachdev	PPPL
Jean Luca	Bez	LBNL

Kjiersten	Fagnan	LBNL
Sam	Blau	LBNL
Kelly	Rose	NETL
Nick	Wunder	NLR
Ray	Grout	NLR
Steven	Spurgeon	NLR
Struan	Clark	NLR
Ana	Gainarua	ORNL
Christian	Engelmann	ORNL
Mayanka	Chandra Shekar	ORNL
Olga	Kuchar	ORNL
Chitra	Sivaraman	PNNL
Terao	Kazuhiro	SLAC
Lindsay	Roy	SRNL

DBS Liaisons for each Model Team

Accelerating Scientific Discovery through AI-Driven Code Development	Diana McSpadden
AI Driven Co-Design for Microelectronics	YQ Chen
AI for Efficient Quantum Algorithms	Le Chen
AI for Planning and Operation of the US Power Grid	Chitra Sivaraman
AI Foundational Model Transforming Combustion Science and Technology for Energy and National Security	Lindsay Roy
America's AI Revolution Transforms Scientific Discovery at the Speed of Light	Nicholas Schwartz
CM²US: Critical Minerals and Materials to Unlock Supply	Steven Spurgeon
FAMOUS: Foundational AI Models for Optimizing and Understanding Biological Systems	Kjiersten Fagnan
From Quarks to Cosmos through the Lens of AI	Brad Sawatzky and Terao Kazuhiro
Fusion-FM: A Foundational Model for Magnetic Confinement Fusion	Jai Sachdev
MAIQMag: Multimodal AI for 2D Quantum Magnets	Yihui Ren
MOAT: Multi-Office Particle Accelerator Team	Yi Huang
PROMETHEUS: A Revolution in Design and Operation of Nuclear Power Plants	Victor Walker
Specs-to-Silicon Microelectronics in Extreme Environments for Transformative Science	Adam Lyon

ModCon Core Intellectual Property & Partnership Formation

IPPF

April 3, 2026

Genesis Mission

U.S. DEPARTMENT of ENERGY

Background

Intellectual Property and Partnerships Team (IPPF or P3)

- Partnerships (INL-led)
- IP and Agreements (ANL-led)

Partnerships Team Focus

- Partner alignment with labs
- Cross-lab coordination

Intellectual Property (IP) and Agreements Team Focus

- Responsive to EO section 5(c)(i) and (ii)
- Model Agreements for industry collaborators
- Inter-lab alignments (“Virtual Lab”)
- Requirements for non-open data protections
- Master Scope of Work

LAUNCHING THE GENESIS MISSION

Executive Orders | November 24, 2025

(c) The Secretary, in coordination with the APST and the Special Advisor for AI and Crypto, shall establish mechanisms for agency collaboration with external partners possessing advanced AI, data, or computing capabilities or scientific domain expertise, including through cooperative research and development agreements, user facility partnerships, or other appropriate arrangements with external entities to support and enhance the activities of the Mission, and shall ensure that such partnerships are structured to preserve the security of Federal research assets and maximize public benefit. To facilitate these collaborations, the Secretary shall:

(i) develop standardized partnership frameworks, including cooperative research and development or other appropriate agreements, and data-use and model-sharing agreements;

(ii) establish clear policies for ownership, licensing, trade-secret protections, and commercialization of intellectual property developed under the Mission, including innovations arising from AI-directed experiments;

(iii) implement uniform and stringent data access and management processes and cybersecurity standards for non-Federal collaborators accessing datasets, models, and computing environments, including measures requiring compliance with classification, privacy, and export-control requirements, as well as other applicable laws; and

DOE/Genesis-level Partnerships

- Genesis Mission teams are strongly encouraged to work with industry, academia, and public-sector partners
- Partnerships can take many forms: data/model sharing, collaborative R&D, in-kind contributions on shared tasks, etc., and range from Genesis-wide to individual PIs.
- PIs are encouraged to work with their local lab tech-transfer office to understand partnering process and scope of agreements

DOE/Genesis-level Consortium

- The Genesis mission consortium established to manage public-private partnerships
- DOE has signed MOUs with 26 partners to be initial consortium members.
- Non-members can also participate in consortium, including partner-matching

Anthropic	Dell Technologies
OpenAI	Accenture Federal Services
Google	Cerebras Systems
Microsoft	Groq
Amazon Web Services	CoreWeave
Oracle	DrivenData
NVIDIA	Armada Systems, Inc.
Intel	Periodic Labs, Inc.
Advanced Micro Devices	Radical AI, Inc.
IBM	XPRIZE Foundation, Inc.
Hewlett Packard Enterprise	Project Prometheus

IP and Agreements – Accomplishments

External - DOE approved agreements

- SOW template (short 1.5 pager) – first step toward an agreement between MT and industry (contact your MT P3 POC)
- AI Agreement for Open Science Projects
- Agreement for AI Projects (non-open data / science)
  - Project Data Use Addendum (to handle proprietary / sensitive project data)
- 1-page NDA for initial project discussions between labs and companies
- Bridge Agreement (for anticipated PIA funded projects)
- Lab's existing agreements may be used

Inter-lab Collaboration Mechanisms

- Virtual Lab Teaming Framework – all 17 labs aligned
- Intellectual Property (IP) Advisory Committee to identify, coordinate, and support tech transfer and IP matters
- Data Use Plan (DUP) for MT projects (in coordination with Data team)
  - Ensure project work's deployment per mission objective
  - Lower risks of unauthorized data use / release
- Genesis Master Scope of Work (MSW) - eliminates local site office approvals for domestic Genesis agreements
- Blanket approval on Project Data protection (requested, with DOE IPL)
- Inter-lab project 1-pager (links Project Lead Lab to other Project Labs)

IP and Agreements – What's next

Current developments

- Internal and partner discussions started
- Records of partnerships and potential partnerships being developed
- MTs developing statements of work to establish agreements
- Data tracking and protections discussions started; joining DBS and Data Team discussions

Next steps

- Labs to sign Genesis consortium agreement
- Convert aligned SOWs to agreements
- Proliferate MSW across labs
- Continue identifying ways to reduce duplicative efforts (e.g., DOE reviews, foreign vetting)

P3 POCS IN MODEL TEAMS

Seed Model Teams (MTs)	P3 POC	P3 POC Email Address
AI-Driven Code Development	Charlie Catlett, Demetra Adrahtas, Christina Lomasney	catlett@anl.gov, dadrahtas@anl.gov, christina.lomasney@pnnl.gov
AI Driven Co-Design for Microelectronics	Charlie Catlett, Phil Smith	catlett@anl.gov, pesmith@anl.gov
AI for Efficient Quantum Algorithms	Charlie Catlett, Phil Smith	catlett@anl.gov, pesmith@anl.gov
AI for the US Power Grid	Ray Grout, Eric Payne	ray.grout@nrel.gov, eric.payne@nrel.gov
Combustion Science and Technology	Charlie Catlett, Hemant Bhimnathwala	catlett@anl.gov, hbhimnathwala@anl.gov
Scientific Discovery at the Speed of Light	David Martin, Peter Nugent, Gail Chen, Sebastian Ainslie	demartin@lbl.gov, penugent@lbl.gov, gailchen@lbl.gov, sainslie@lbl.gov
CM²US	Charlie Catlett, Maricely Ramirez, Julienne Krennrich	catlett@anl.gov, mrhernandez@anl.gov, jmkrenn@ameslab.gov
FAMOUS	Christina Lomasney-Mattis, Paul Rigor	christina.lomasney@pnnl.gov, paul.rigor@pnnl.gov
From Quarks to Cosmos	Charlie Catlett, Peter Nugent, Ilya Kats	catlett@anl.gov, penugent@lbl.gov, ikats@anl.gov
Fusion-FM:	Arturo Dominguez, Chris Wright	adomingu@pppl.gov, cwright@pppl.gov
MAIQMag:	Shaun Gleason	gleasonss@ornl.gov
MOAT:	David Martin, Peter Nugent, Gail Chen, Sebastian Ainslie	demartin@lbl.gov, penugent@lbl.gov, gailchen@lbl.gov, sainslie@lbl.gov
Prometheus:	Jason Hales, Jason Stolworthy	jason.hales@inl.gov, jason.stolworthy@inl.gov
Specs-to-Silicon Microelectronics	Whitney Hastings, James Amundson	whasting@fnal.gov, amundson@fnal.gov

U.S. DEPARTMENT of ENERGY

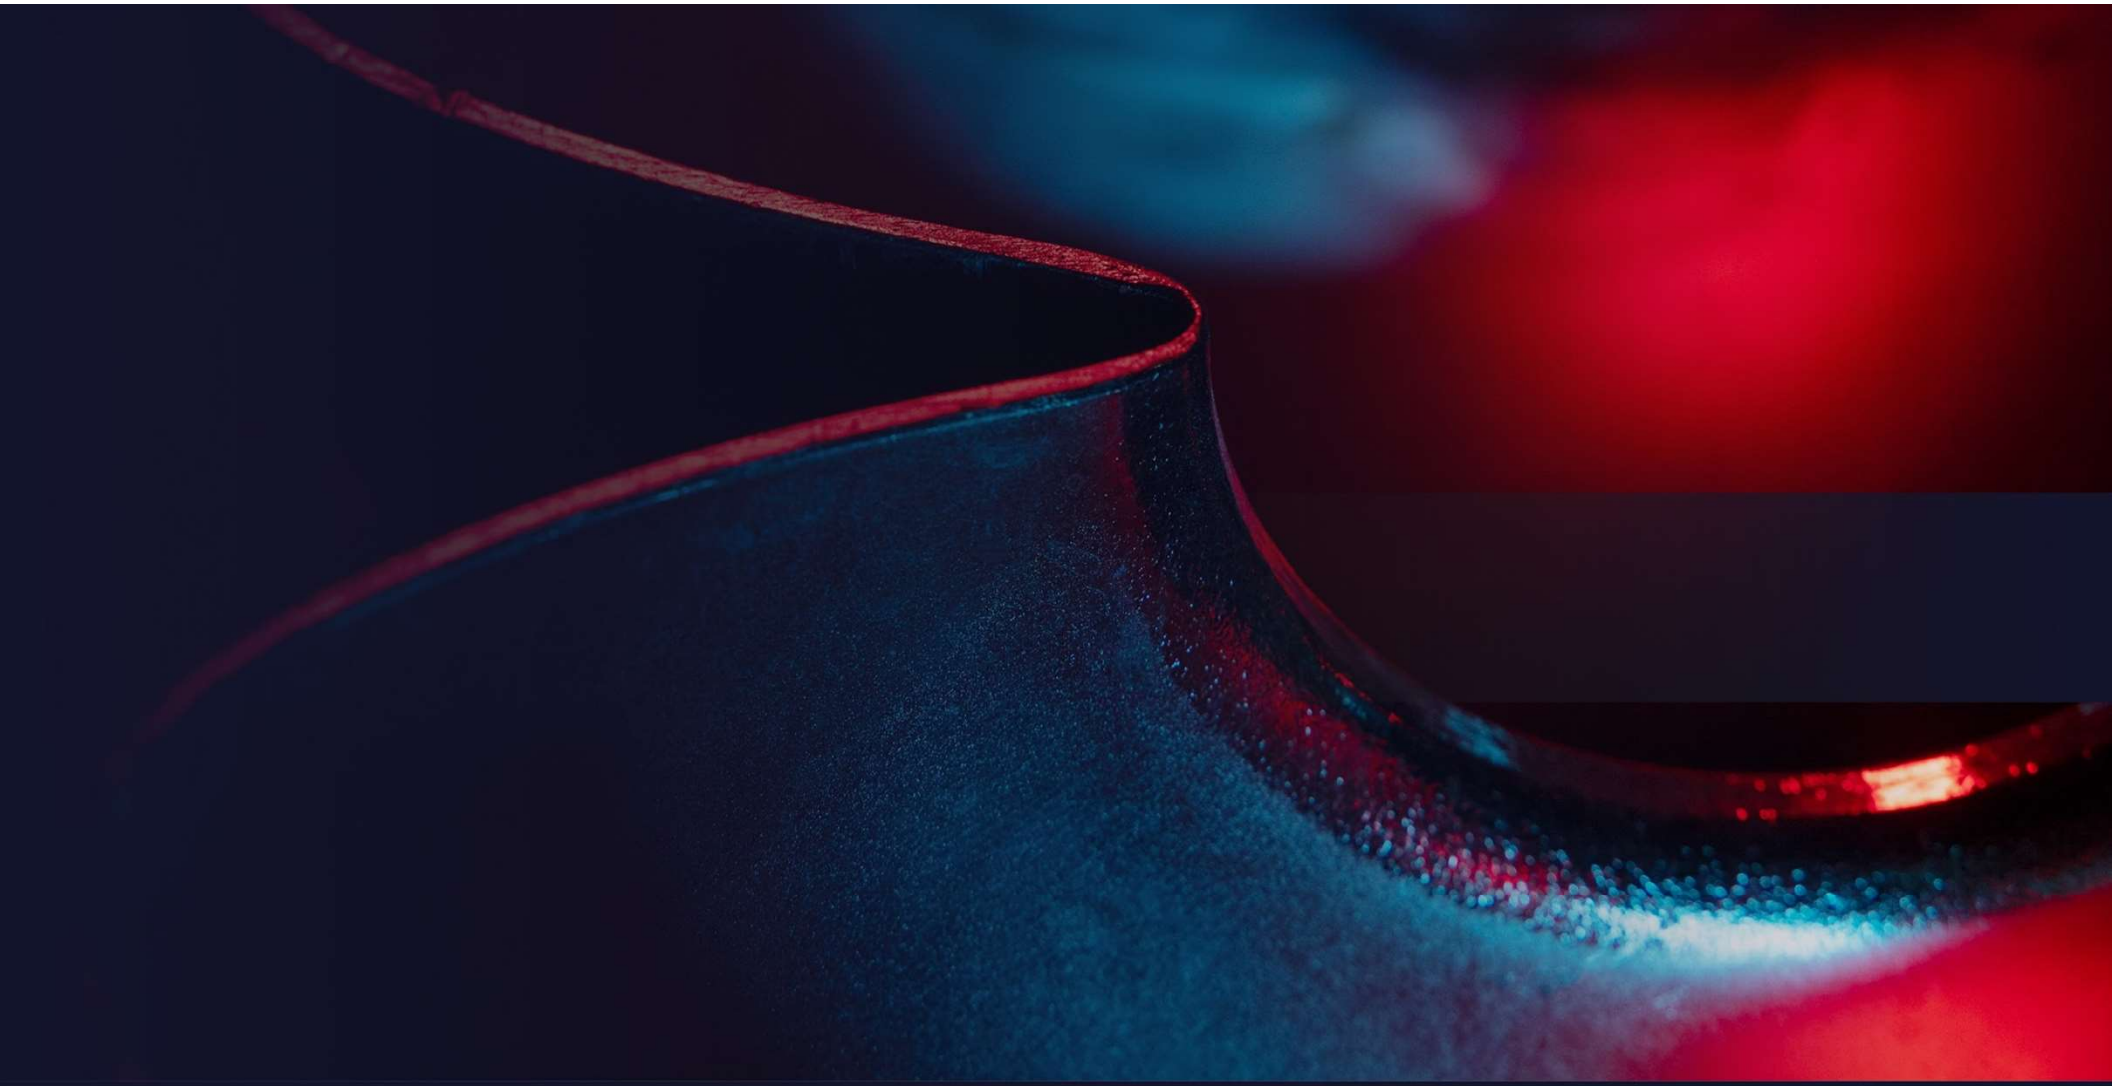

U.S. DEPARTMENT *of* ENERGY