Welcome!

- **Audio & Video**: All participant microphones and cameras are muted and will remain disabled during the presentation
- **Sound Check**: Please verify your speaker volume is turned up and not muted
- **Questions**: The chat function is disabled, but you can submit questions through the Q&A box. We will address questions at the end during the Q&A session
- **Recording**: This webinar is being recorded. The recording and a compilation of Q&As will be made available to all attendees afterwards

Genesis Mission

AMERICAN SCIENCE CLOUD

January 20, 2026

Introduction to the American Science Cloud (AmSC): Part 2

U.S. DEPARTMENT of ENERGY

What is the American Science Cloud (AmSC)

- AmSC is a secure, federated, and science-optimized cloud environment
  - *Not centralized servers or services but rather an environment that will support contributed tools and capabilities*
- AmSC integrates the DOE's world-leading computing and experimental facilities, data resources, and high-performance networks
  - *A flexible structure to support integration of capabilities*
- AmSC enables AI-ready end-to-end workflows
  - *Not a single set of predefined workflows but inclusive of contributions across the complex*

Genesis Mission defines ambition and urgency.

AmSC is a cornerstone building a shared platform to execute the Genesis Mission

AmSC Concepts and Connections to ModCon and GM

- **Infrastructure Partners (IPs):** DOE labs that contribute and integrate core capabilities into AmSC
- **Model Teams (MTs):** Multi-lab teams developing domain AI models and workflows
- **Seed Projects:** Early, mission-driven use cases that pressure-test data, models, and workflows
- **Lighthouse Projects (LHPs):** High-impact scientific efforts that built on MTs/Seed projects to address Genesis Mission priorities in energy, science, and national security.
- **Lighthouse Challenges (LHCs):** Grand challenges composed of multiple coordinated LHPs

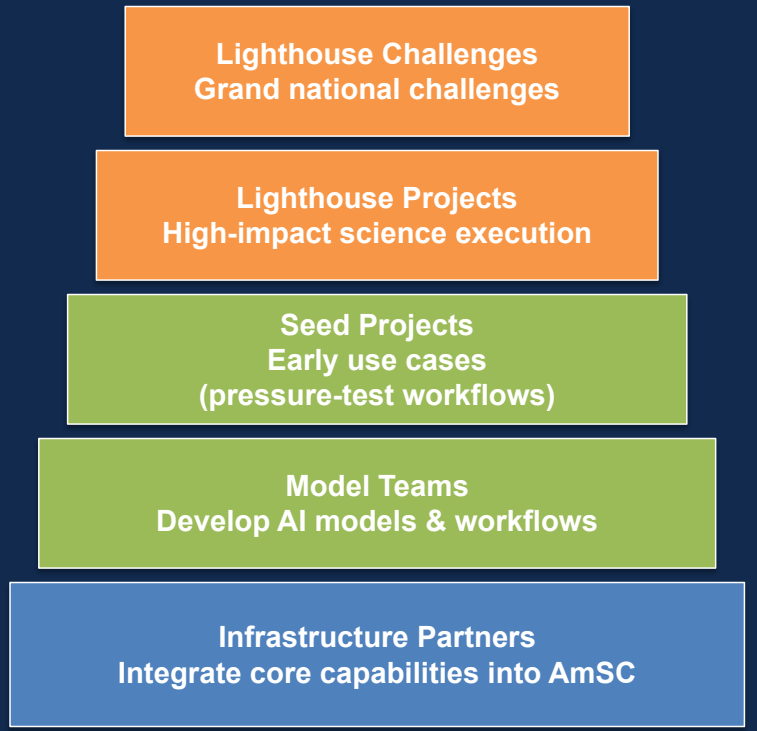

Infrastructure Partners

Infrastructure Partner(s)	Lab(s)	Key Contacts (POC/PI, Leads)
OLCF	ORNL	Arjun Shankar
ALCF	ANL	Mike Papka
ESnet	LBNL	Inder Monga
NERSC	LBNL	Sudip Dosanjh
HPDF	TJNAF, LBNL	Graham Heyes
C3	PNNL	Robert Rallo
Scientific User Facilities	LBNL, ANL, TJNAF, BNL, FNAL, SLAC, ORNL	Paolo Calafiura, Nicholas Schwarz

Infrastructure Partner(s)	Lab(s)	Key Contacts (POC/PI, Leads)
DeepLynx	INL	Peter Suyderhoud
DIII-D Digital Twin	GA	Sterling Smith
EDX	NETL	Kelly Rose, Chad Rowan, Jack Searle
FemtoMind	TJNAF, FNAL	Robert Edwards
Fermi Data Platform	FNAL	James Amundson
HAIDIS	TJNAF	Ilya Baldin
SCDF	BNL	Adolfy Hoisie
S3DF	SLAC	Jay Srinivasan
Stellar-AI	PPPL	Shantenu Jha
VEE-ARIES/HERO	NLR	Kristi Potter

Multi-Lab and Public-Private Partnership

Science Council

Chair:Andreas Kronfeld (FNAL)
Kevin Yager (BNL)
Michael Begel (BNL)
Kjiersten Fagnan (LBNL)
Oliver Gutsche (FNAL)
Joel England (SLAC)
Todd Satagota (JLAB)
Jana B. Thayer (SLAC)
Neeraj Kumar (PNNL)

Industry Council

AWS
AMD
Cisco
Dell
Google
HPE
Microsoft
NVIDIA
Nokia
CORNELIS Networks

AmSC Project Leadership Office

DIRECTOR: **Gina Tourassi** (ORNL)

DEPUTY: **Arjun Shankar** (ORNL)
DEPUTY: **Inder Monga** (LBNL)

CTO: **Sarp Oral** (ORNL)
Product Manager: **Deborah Bard** (LBNL)
PMO: **Denise Hoomes** (ORNL)
CISO: **Ryan Adamson** (ORNL)

ModCon CoDesign Liaisons

Wahid Bhimji (LBNL)
Feiyi Wang (ORNL)
Venkat Vishwanath (ANL)

IRI Liaisons

Deborah Bard (LBNL)
Thomas Uram (ANL)

OPERATIONS

L1: **Ashley Barker** (ORNL)

Virtual Organization
Veronica Vergara (ORNL)
Adam Slagell (LBNL)

Intelligent Operations
Ed Balas (LBNL)
Eric Pershey (ANL)

User Engagement
Haritha Siddabathuni Som (ANL)
Chris Fuson (ORNL)

INTERFACES AND SERVICES

L1: **Mike Brim** (ORNL)

AmSC Interfaces
Taylor Childers (ANL)
John MacAuley (LBNL)

Infrastructure Services
Zach Mayes (ORNL) *
Shane Canon (LBNL)
P. Shyamshankar (ANL)

Data Services
Ilya Baldin (TJNAF)
Rajesh Kalyanam (ORNL)

AI SERVICES

L1: **Thomas Uram** (ANL)

Model Services
John Gounley (ORNL)
Huihuo Zheng (ANL)

At-Scale Services
Murali Emani (ANL)
Steven Farrell (LBNL)

Intelligent Interfaces
Wahid Bhimji (LBNL)
Shreyas Cholia (LBNL)

AmSC PARTNER INTEGRATION

L1: **Chin Guok** (LBNL)

IRI Integration
John MacAuley (LBNL)
Paul Rich (ANL)

CSP Integration
Roger Cass (PNNL)
Jack Deslippe (LBNL)

Industry & Gov IP Integration
David Martin (LBNL)
Rebecca Hartman-Baker (LBNL)
Ben Mintz (ORNL)

** Integration Architect*

✓ Federated leadership with **Science Council** and **Industry Council**

The American Science Cloud (AmSC)

A Unified Cloud Platform for Transformational AI and Science Across DOE
A cornerstone of the Genesis Mission

Mission

Deliver a seamless integration of DOE science instruments, computing, experimental, and networking capabilities, AI, data, modeling and simulation software tools in a single advanced programming interface – the **AmSC API**

What We Deliver

✓ Unified APIs for data, model, and workflow access
✓ Secure, federated identity (Virtual Organization model)
✓ Seamless interoperability between DOE facilities and CSPs

Reduce “time to insight” from months to days through composable AI-driven workflows

AmSC Detailed Architecture

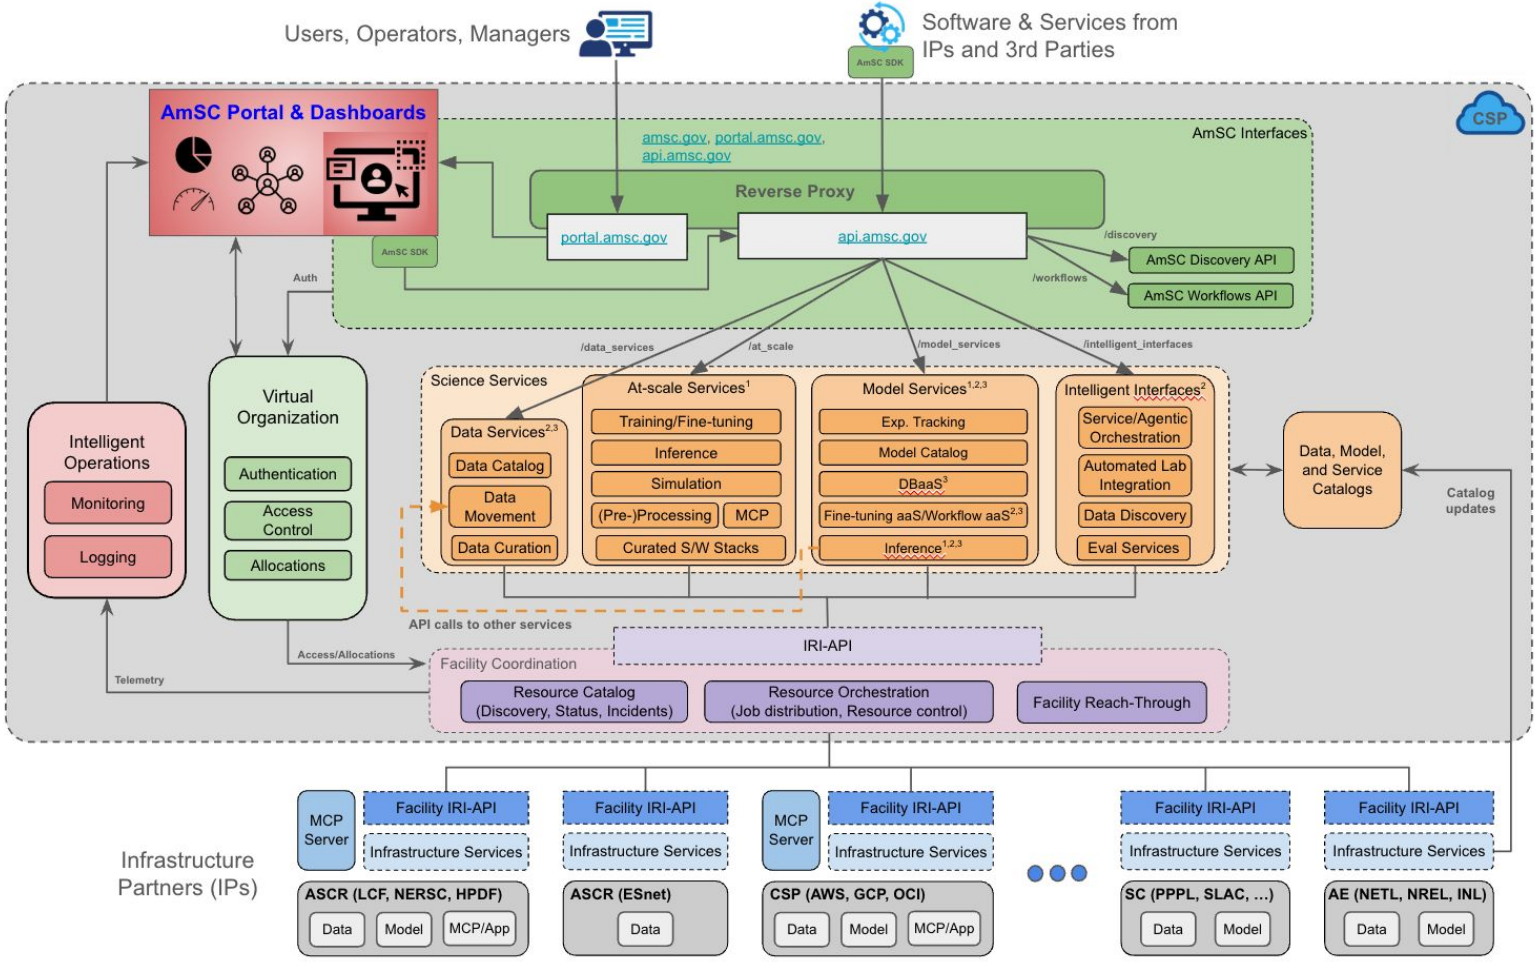

Genesis Mission

Our milestones will build towards our MVP

3-MONTHS
FEDERATED ACCESS DEMONSTRATION

✓ Unified login across multiple DOE HPC sites; first operational **REST API and Python client** enabling authentication, dataset discovery, and workflow execution

6-MONTHS
SCIENCE TEAM ONBOARDING

✓ Two to three research teams execute **end-to-end workflows** (data search → transfer → model fine-tuning) on **Frontier**, **Perlmutter**, or **Aurora** using the AmSC Data Catalog

9-MONTHS
INTELLIGENT WORKFLOW MILESTONE

✓ First early **agentic workflow framework** autonomously discovers data, trains, fine-tunes, and performs **reasoning-based inference** across federated DOE and cloud resources

12-MONTHS
MINIMUM VIABLE PRODUCT

✓ Link IPs and CSPs into a **unified pre-production environment** to support AI-ready data services, agentic frameworks, large-scale model training/inference and advanced simulation. Establish a **foundation for future expansion** to additional partners and services.

LONG-TERM VISION

Accelerate discovery cycles by building the world's most integrated scientific platform

In Month 3, AmSC is delivering the first capabilities to support end-to-end scientific AI workflows

3-MONTHS
FEDERATED ACCESS DEMONSTRATION

✓ Unified login across multiple DOE HPC sites; first operational **REST API and Python client** enabling authentication, dataset discovery, and workflow execution

Unified login across multiple DOE sites:

- The AmSC federated identity system
- The AmSC User Dashboard
- The AmSC staff hub and myAmSC user portal

First operational APIs and Services:

- AmSC Science AI Analysis Support with ChatUI
- AmSC Data Catalog and Data Movement API
- AI Software Stack, including prototype MLOps platform and Inference services
- The AmSC REST API with Python Client Module
- CSP Integration via IRI API
- Our AmSC development AWS environment

The AmSC Month 3 deliverables are the result of new efforts, underway as a direct result of the Genesis Mission.

What does the AmSC MVP look like for a simple workflow for data discovery and model training?

1. Register & Discover Data
- Register datasets in Data Catalog
- Search for similar relevant datasets

2. Find Base Model
- Search Model Catalog for appropriate model

3. Train Model
- Move data to appropriate site
- Train model with selected data
- Run on AmSC infrastructure

This set of basic AI capabilities forms part of multiple science use cases:

- **GridAI:** foundation model development for electrical grid planning and decision-support.

- **DIII-D:** surrogate model development for fusion reactor facility digital twin.

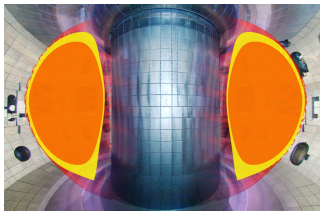

- **HAIDIS**: extracting physics information from GLUEX nuclear physics data.

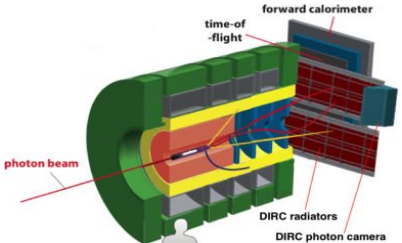

What does the AmSC MVP look like for a simple inference workflow?

1. Data produced at experiment facility
- May be file-based or streaming

2. Data processed with AI-enhanced analysis
- May be time-sensitive
- May include simulations

3. Results inform next steps
- Action may be directed by user, or performed automatically as part of an agentic workflow.

This is a common workflow for facilities who will use AmSC to accelerate experimental science:

- **ALS:** Image analysis, used to give lightsource users feedback on experiment progress in real time.

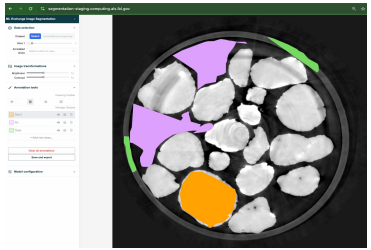

- **Genomics** (FAMOUS/OPAL): agent-driven feedback for steering autonomous lab experiments.

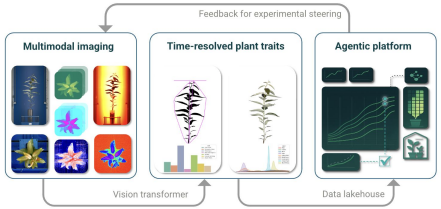

- **SUF + MOAT:** build framework of AI-powered real-time feedback for particle accelerator operators.

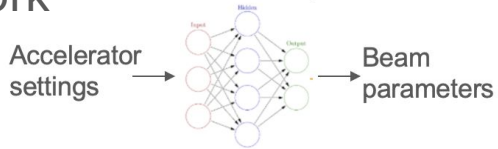

AmSC Identity Management

MVP Goal: User can use a single identity to access all AmSC services and interfaces, running across multiple DOE sites and CSPs.

Month 3 Deliverable: User can login via ESnet and OLCF accounts (other identity providers in progress) and gain access to the Staff Hub and User Portal.

AmSC User Dashboard

MVP Goal: User can access all AmSC Services via a Dashboard. APIs enable service execution. User can also view and track their work executed within the AmSC ecosystem.

Month 3 Deliverable: User Dashboard design mockup with stubbed backend services.

Note: All AmSC services shown hereafter will be accessible via the Dashboard in the MVP

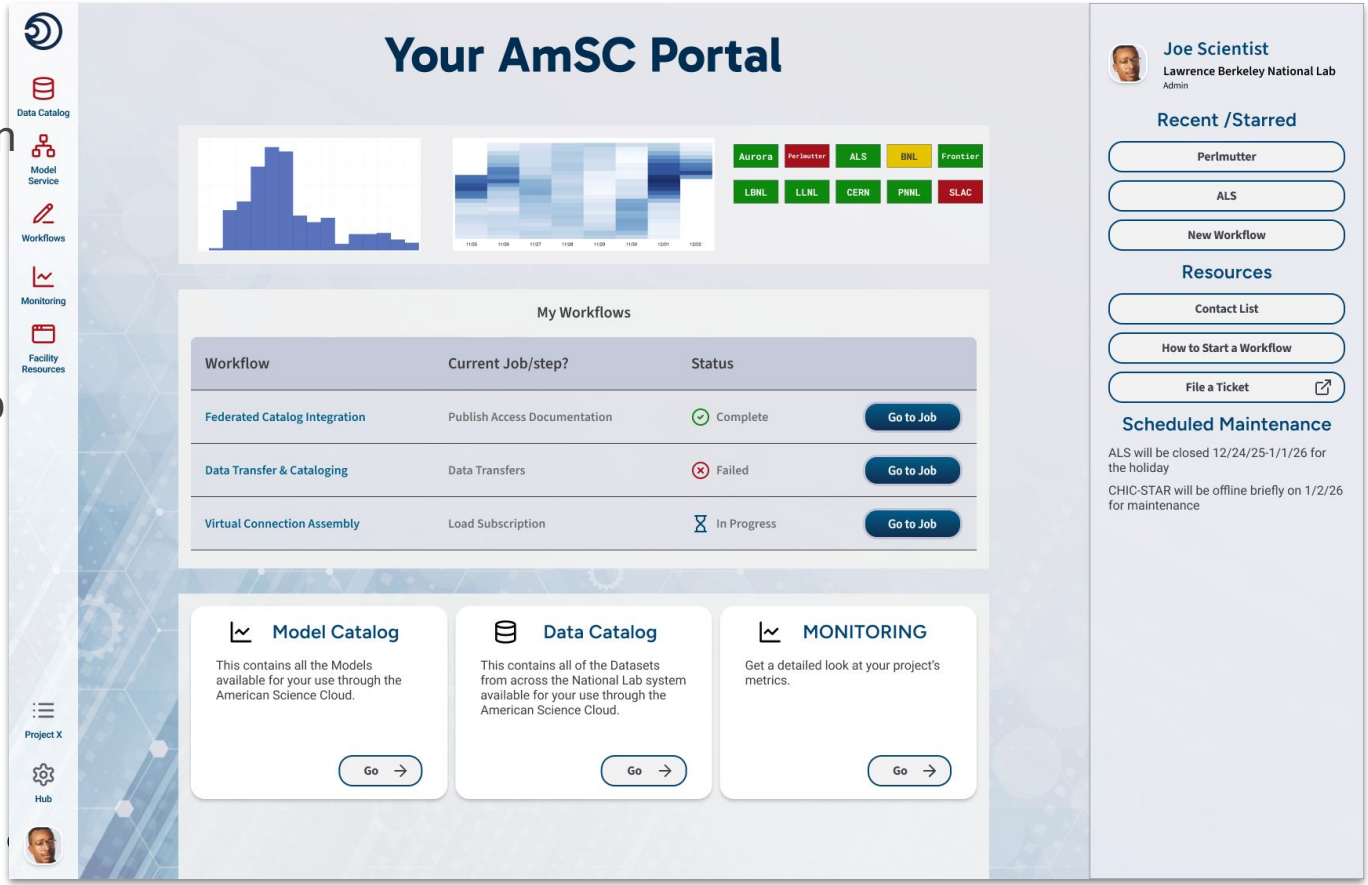

AmSC User Dashboard

MVP Goal: User can access all AmSC Services via a Dashboard. APIs enable service execution. User can also view and track their work executed within the AmSC ecosystem.

Month 3 Deliverable: User Dashboard design mockup with stubbed backend services.

Example services:

1. Graphs for quick assessment.
2. Recent work presented with current status.
3. Current Site Status
4. Data/Model Catalog

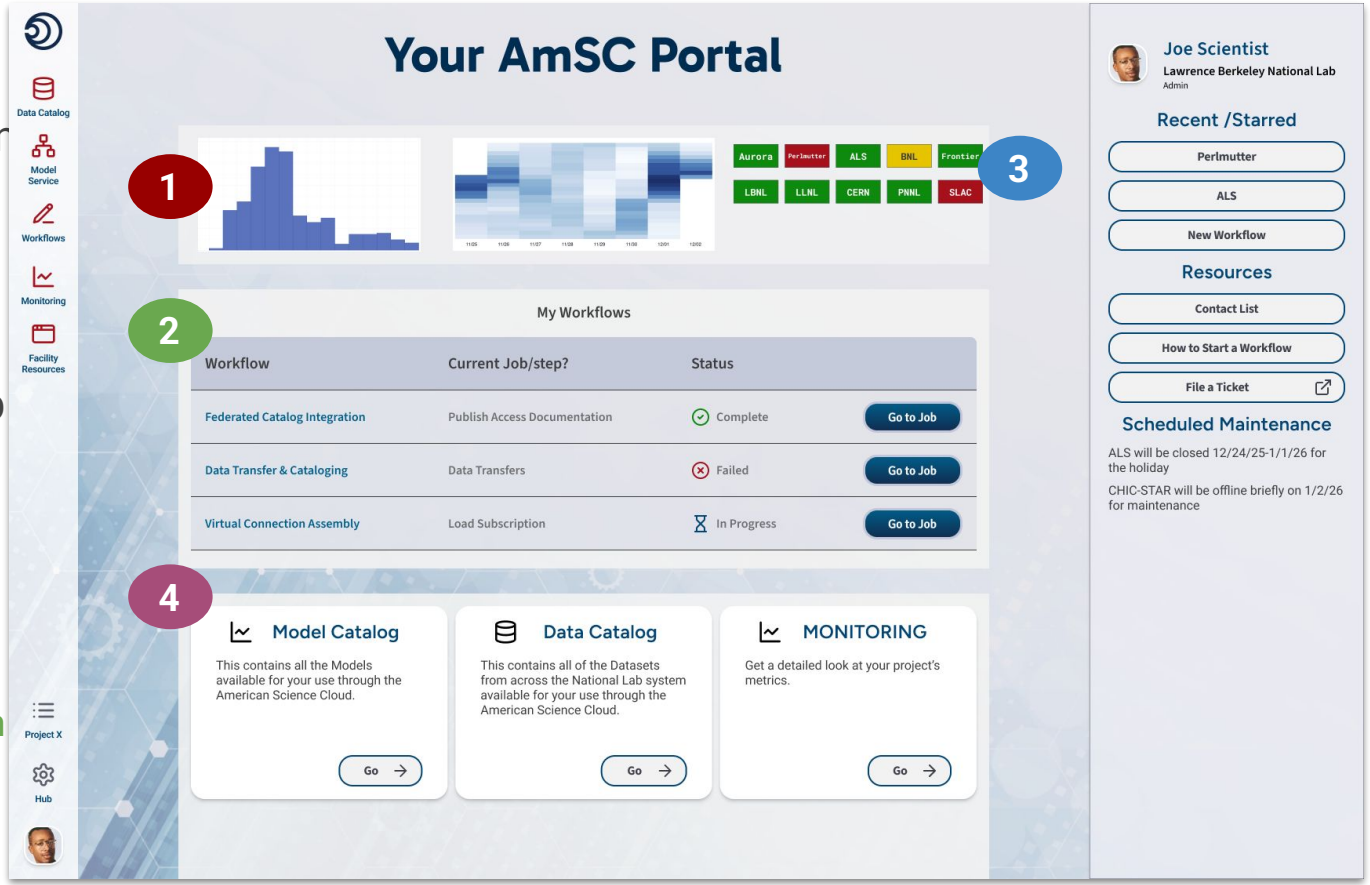

AmSC PI Hub

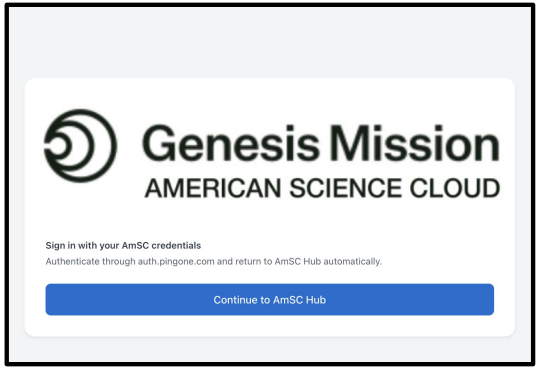

MVP Goal: PIs can manage AmSC projects, resources, services and users. Hub is customizable for all infrastructure partners.

Month 3 Deliverable: Functional prototype hub with view of projects and allocations, accessible via AmSC Fed ID.

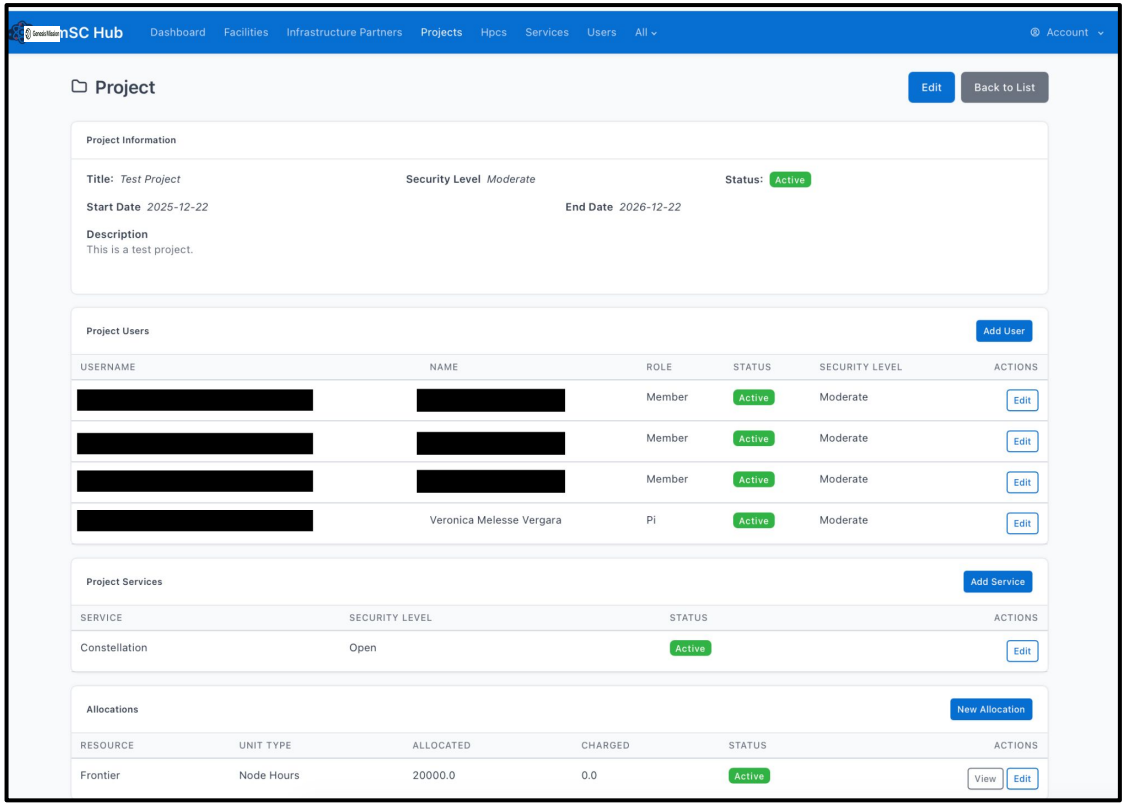

myAmSC User Portal

MVP Goal: Users can access the myAmSC user portal to: apply for accounts, view account and project information, and manage their projects.

Month 3 Deliverable: Functional initial portal with user profiles and account application, accessible via AmSC Fed ID.

Account Application

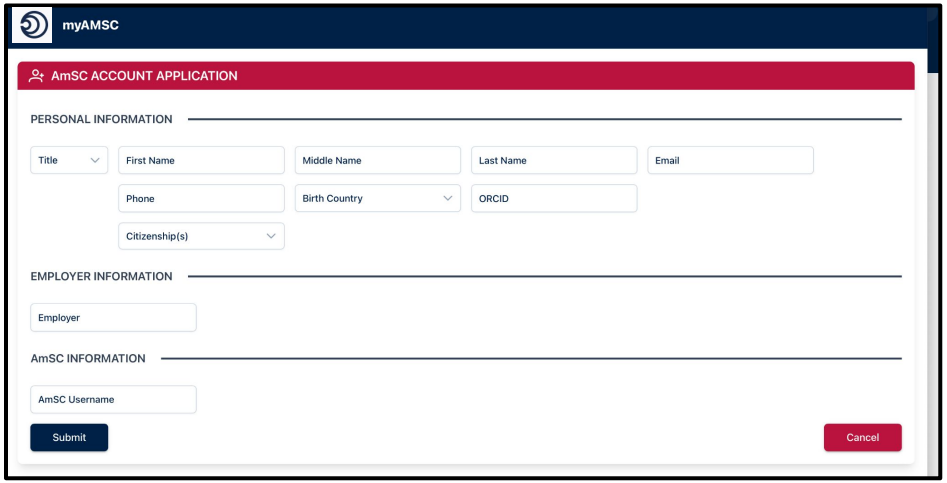

Project Profile

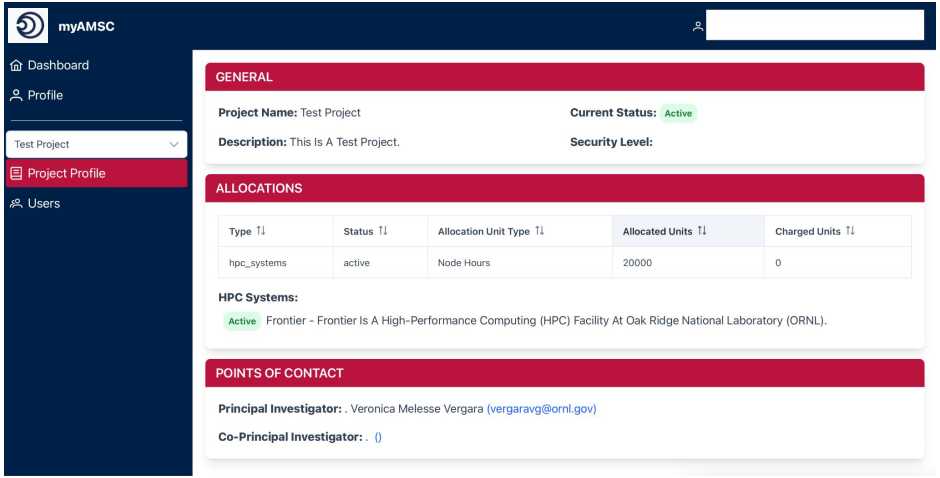

AmSC Interfaces: Python & REST API

MVP Goal: Users, workflow engines and AmSC services can access and manage AmSC services via a fully-featured Python API client.

Month 3 Deliverable: use the AmSC Python Client API running in a Jupyter Notebook to communicate with AmSC Data services via the REST API.

This example shows listing datasets and registering a new dataset in the catalog. APIs now exist for model catalogs and workflow launching.

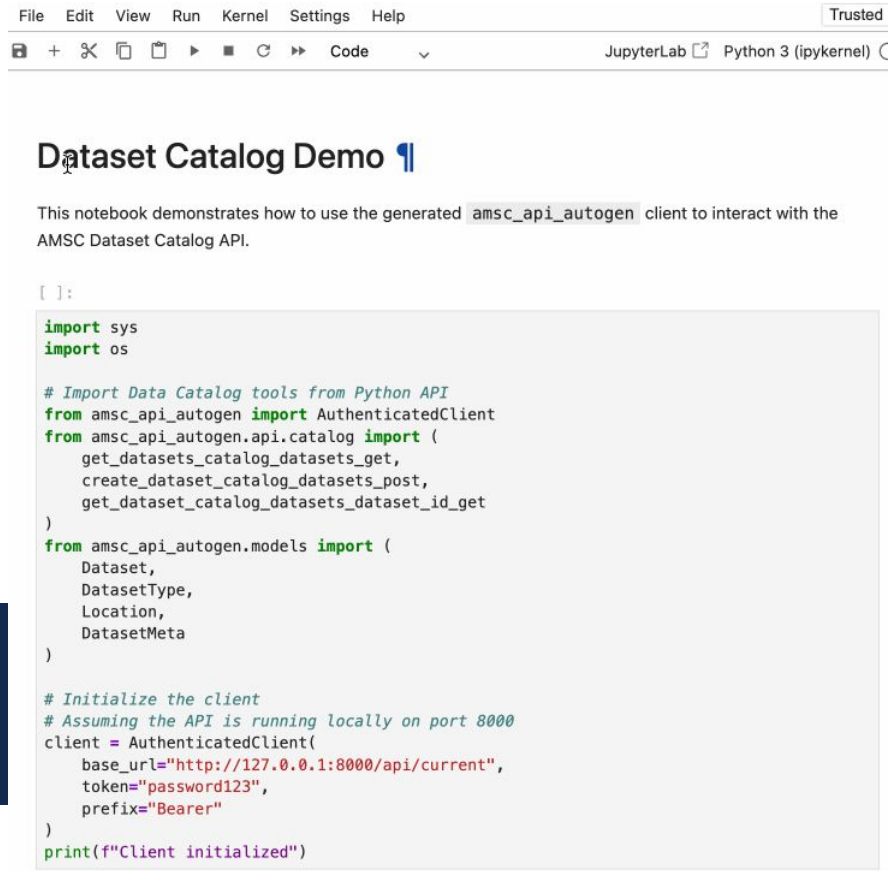

AmSC Agentic Stack

MVP Goal: User can interact with all AmSC services via an agentic interface, and use agents to discover data, train/ fine-tune models, and perform reasoning-based inference across federated DOE and cloud resources.

Month 3 Deliverable: AmSC staff have prototyped and evaluated the first full AmSC agentic stack, deployed on AmSC AWS platform.

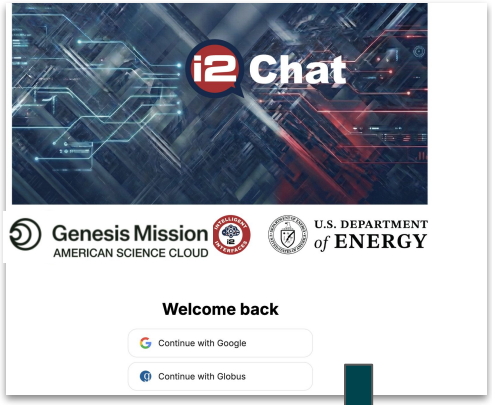

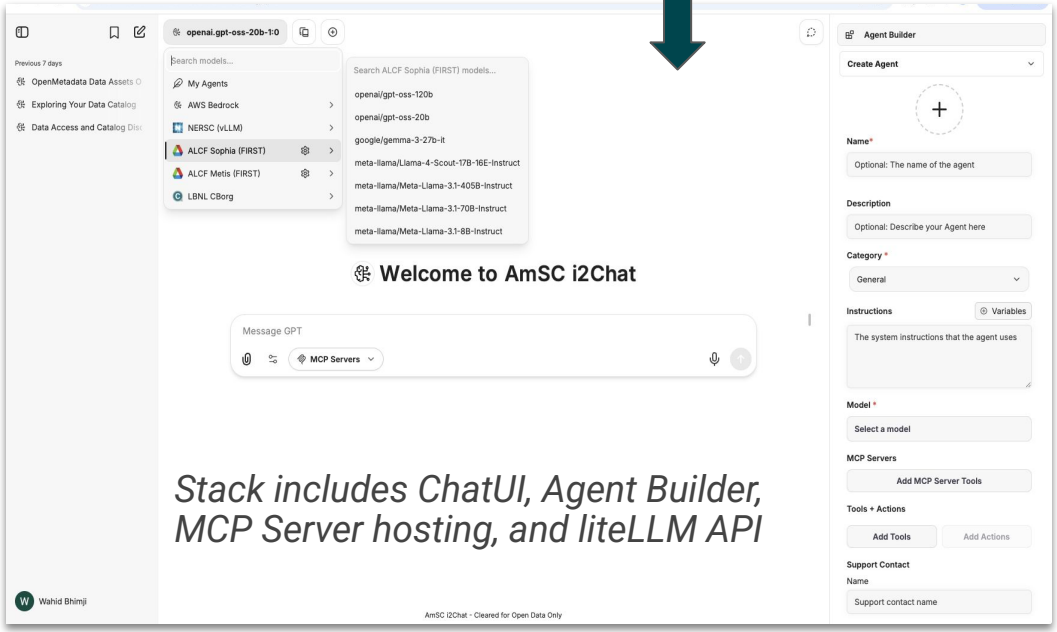

Stack includes ChatUI, Agent Builder, MCP Server hosting, and liteLLM API

AmSC Agentic Stack: Agents for Data Discovery

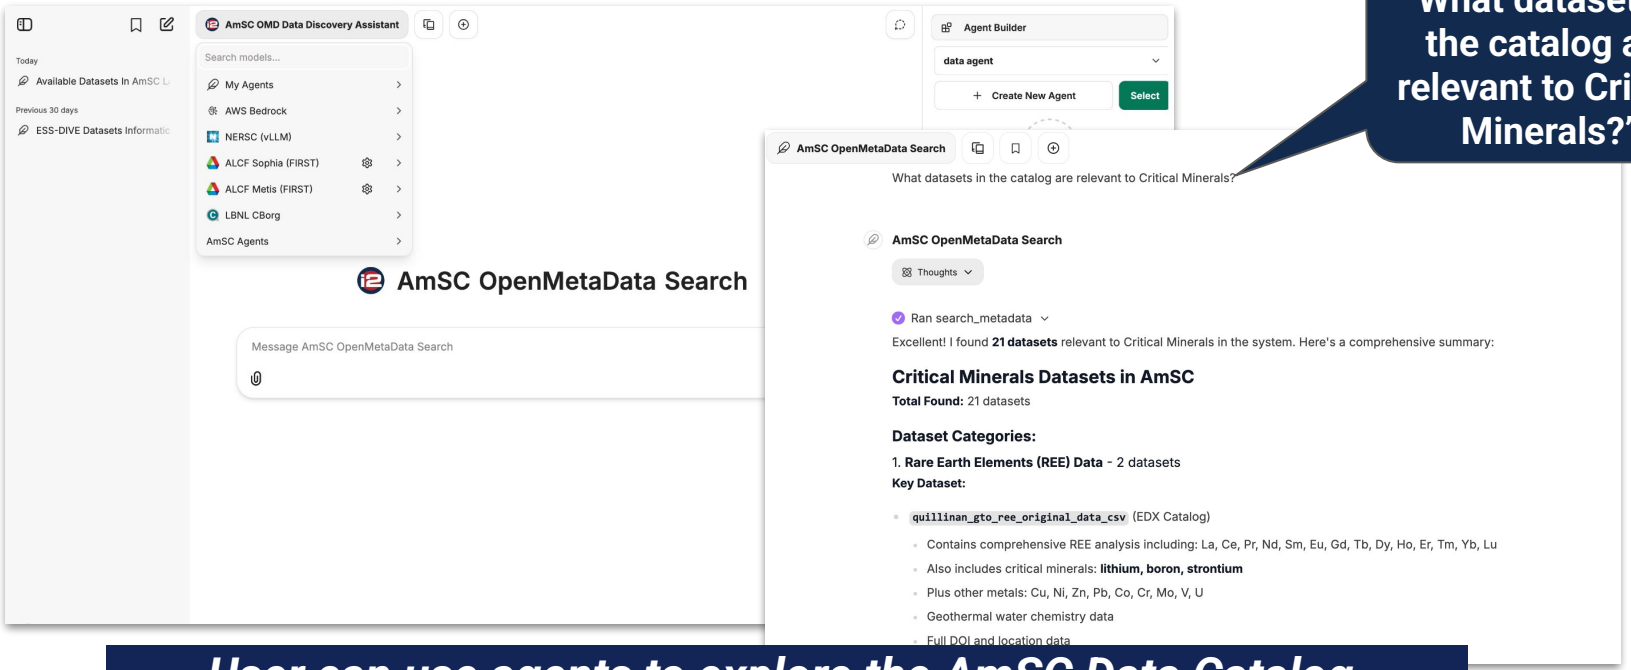

User can use agents to explore the AmSC Data Catalog, integrated via an MCP service.

AmSC Data Catalog

MVP Goal: Users can register new datasets and find and/or update existing datasets, setting appropriate access controls. Datasets can be accessed by users and all AmSC services and interfaces.

Month 3 Deliverable: Functional prototype Data Catalog based on OpenMetadata and populated with datasets from AmSC IPs.

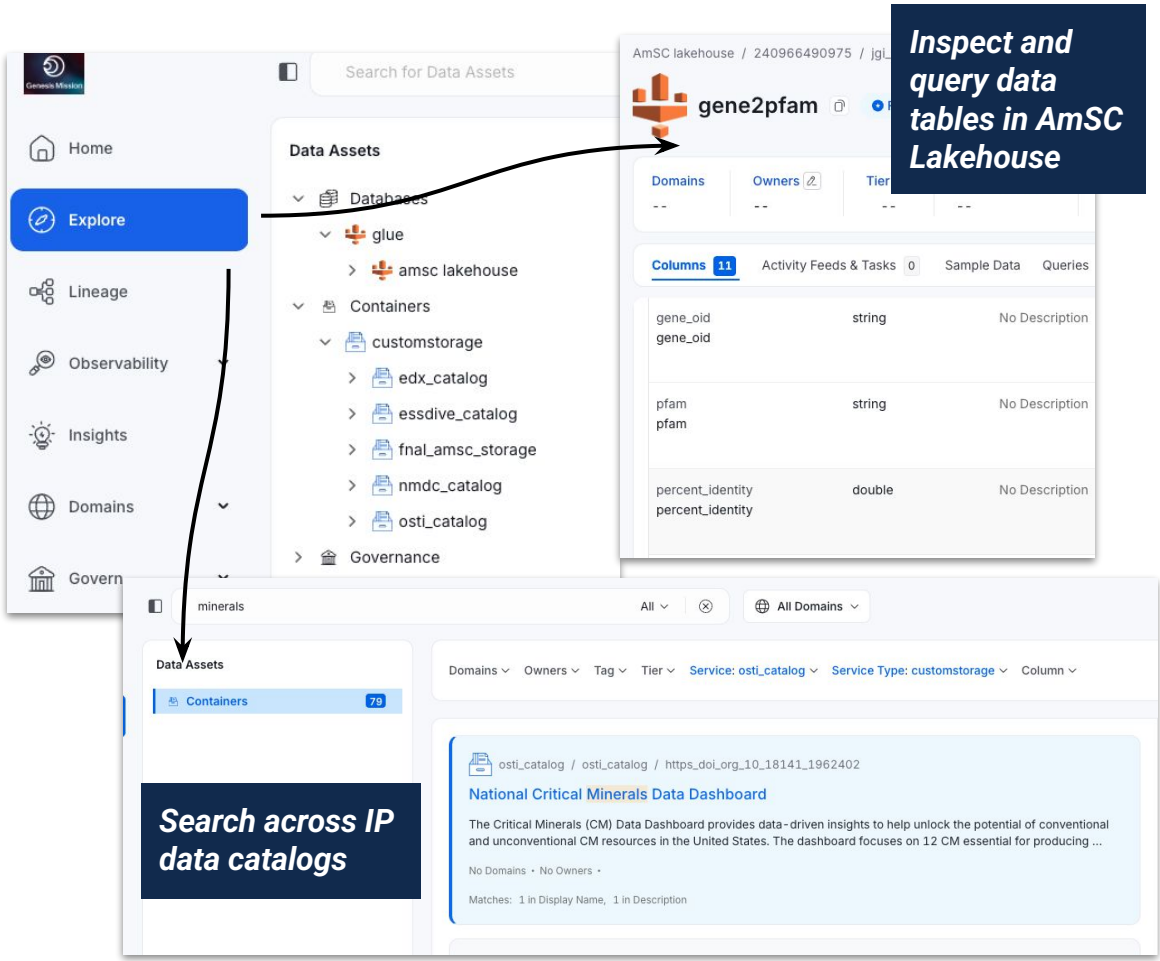

AmSC Data Movement API

MVP Goal: User can move data between AmSC sites (DOE and CSP) via a unified service; AmSC services can transfer data to appropriate compute resources as necessary.

Month 3 Deliverable: Globus-based data movement between NERSC and OLCF, using AmSC API-minted tokens.

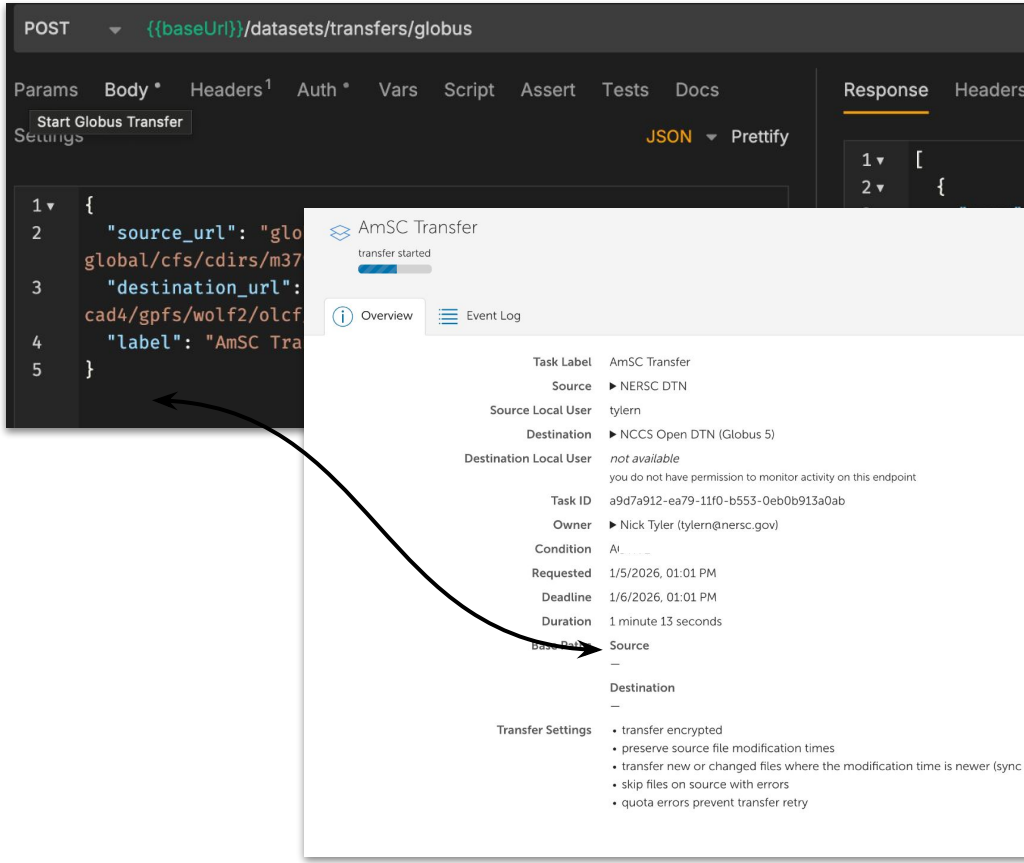

AmSC MLOps Platform: Experiment Tracking and Model Registry

MVP Goal: Users can leverage industry-standard tools to track experiments and run automated training, fine-tuning and hyperparameter optimization.

Month 3 Deliverable: Deployment of MLFlow in AWS, with connection tested from ALCF, BNL, NERSC, OLCF. Evaluating ClearML and ZenML across AmSC sites.

AmSC deploys support for model development at AmSC sites, and for cataloging and sharing AI models

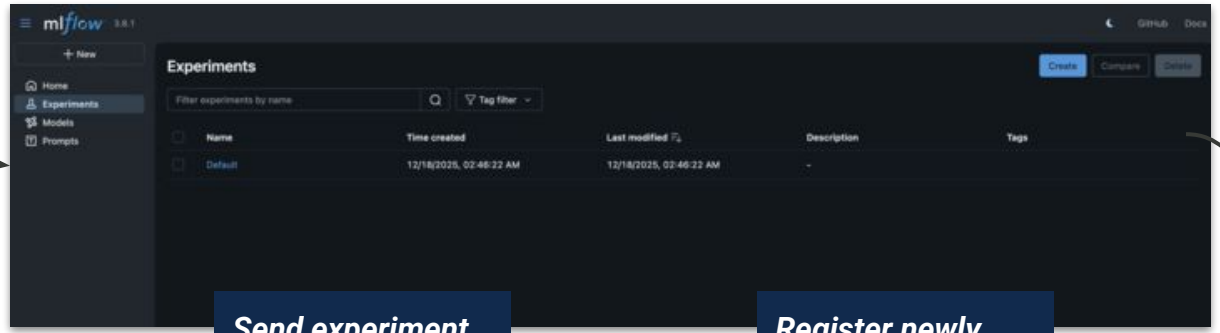

Send experiment data to MLOps

Register newly trained models

Models pulled from catalog for use at AmSC sites

AI Inference Services

MVP Goal: Users can leverage high-performance inference servers (e.g. vLLM, SGLang) hosted at AmSC sites.

Month 3 Deliverable: Prototype deployment of inference services at ALCF, OLCF, and NERSC.

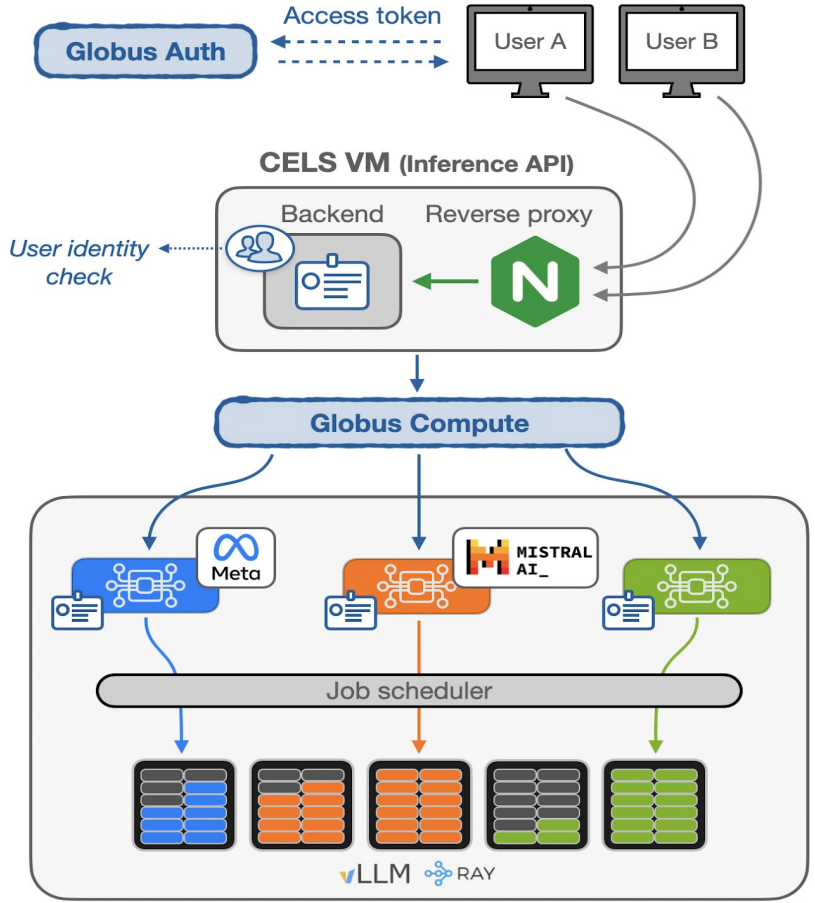

AmSC At-Scale Services: AI Software stacks and benchmarks

MVP Goal: Users can access the same software stack across all AmSC sites; software stack performs efficiently at each site.

Month 3 Deliverable: Completed survey of AI software at ASCR facilities; developed baseline benchmarks with initial performance measurements across sites.

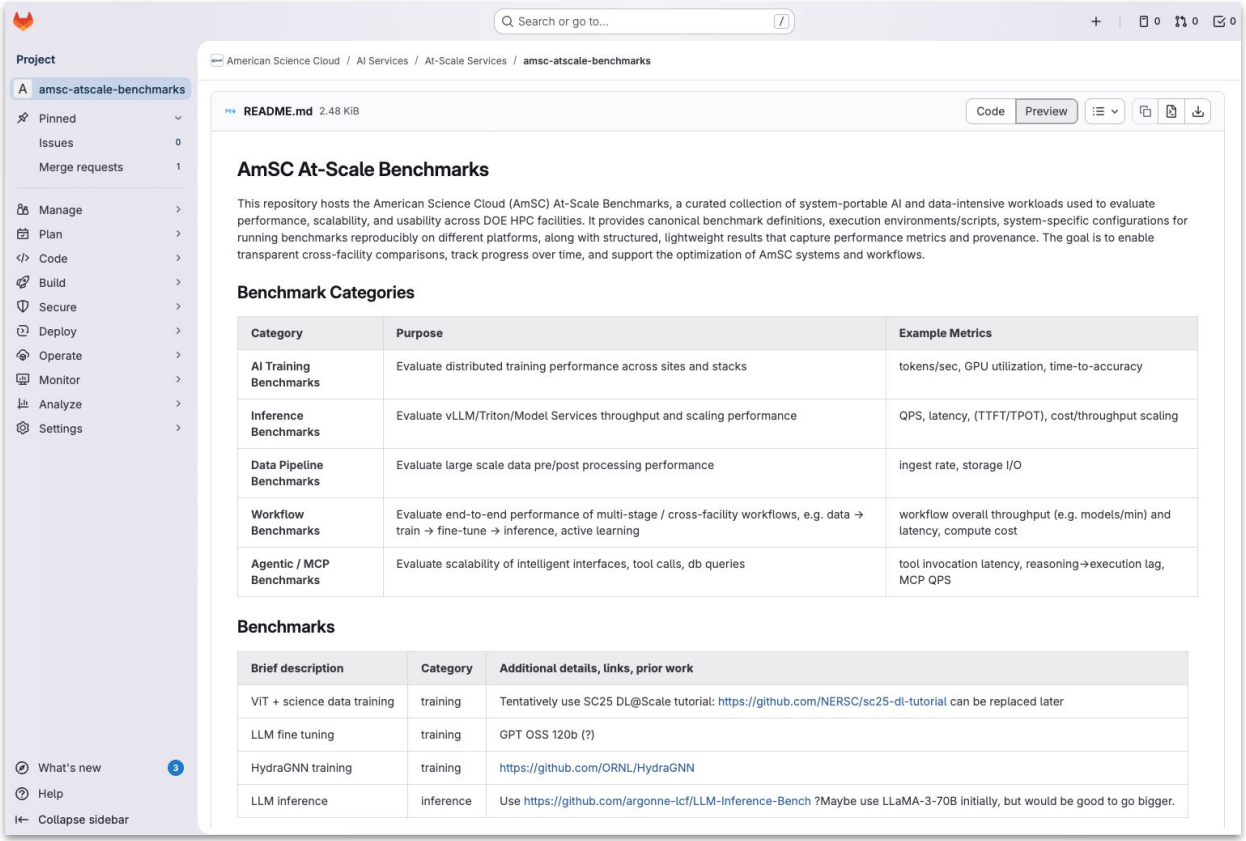

IRI API: Integrating with AmSC sites and facilities

MVP Goal: AmSC services can use the IRI API to interact with AmSC IP sites. Endpoints include resource status, authenticated endpoints for job management, local file management, allocation information, resource orchestration and WAN management.

Month 3 Deliverable: completed API design spec and stubs for compute and filesystem control, as well as resource/site status endpoints. Connected AWS to AmSC APIs and services via an extension to the IRI API.

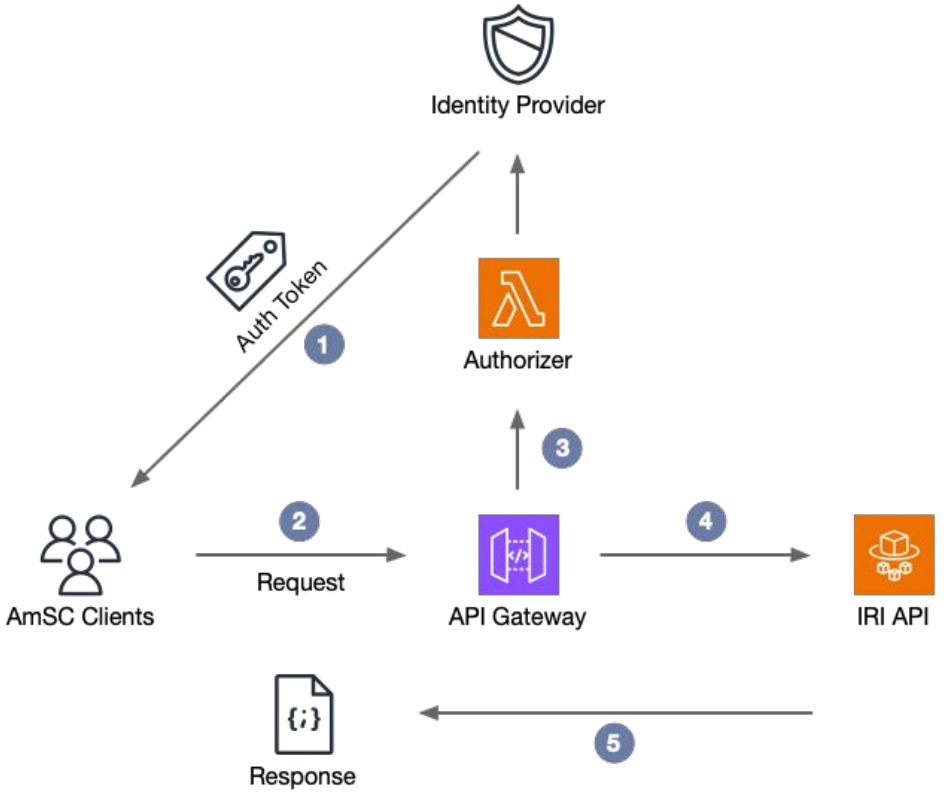

AmSC is developing capabilities via a process of iterative co-design

Next step: iterate with key AmSC users to evaluate how well capabilities meet science requirements and recommend changes to better align with science needs.

There are many AmSC stakeholders. Iteration will be managed via:

- Deeper dives into AmSC services via hackathons and workshops for current early adopters
- AmSC Liaisons to our early adopters are playing a pivotal role by giving feedback on individual use cases
- Coordinated process for scope changes

We will learn a huge amount as we work together to prepare for the Month 6 deliverables!

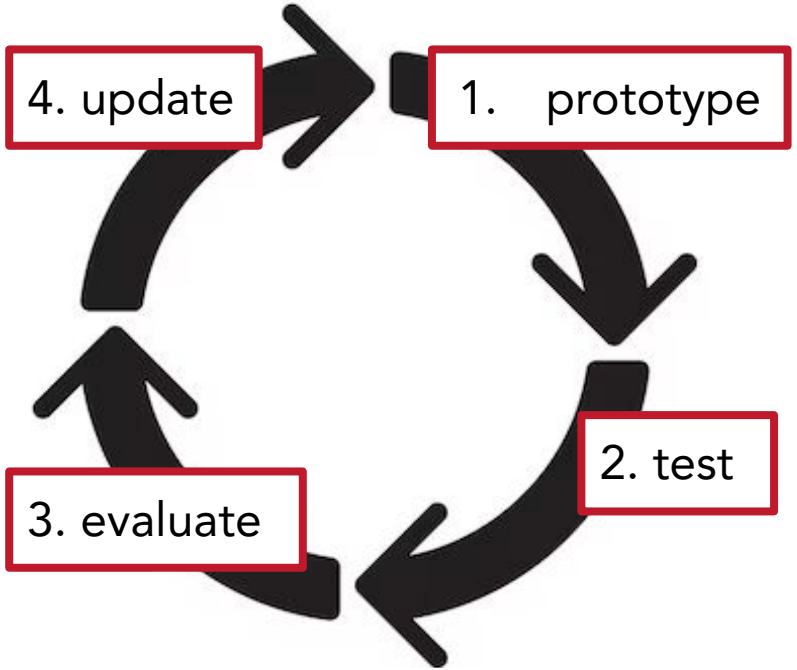

Infrastructure Partners (IPs) are a necessary component in the success of AmSC

IPs leverage AmSC to provide enhanced capabilities for the scientific workflows they currently support.

IPs integrate resources into the AmSC ecosystem to offer additional capabilities for all AmSC scientific workflows. Resources include physical infrastructure (e.g., compute, storage, network), services (e.g., data curation, inference), and data.

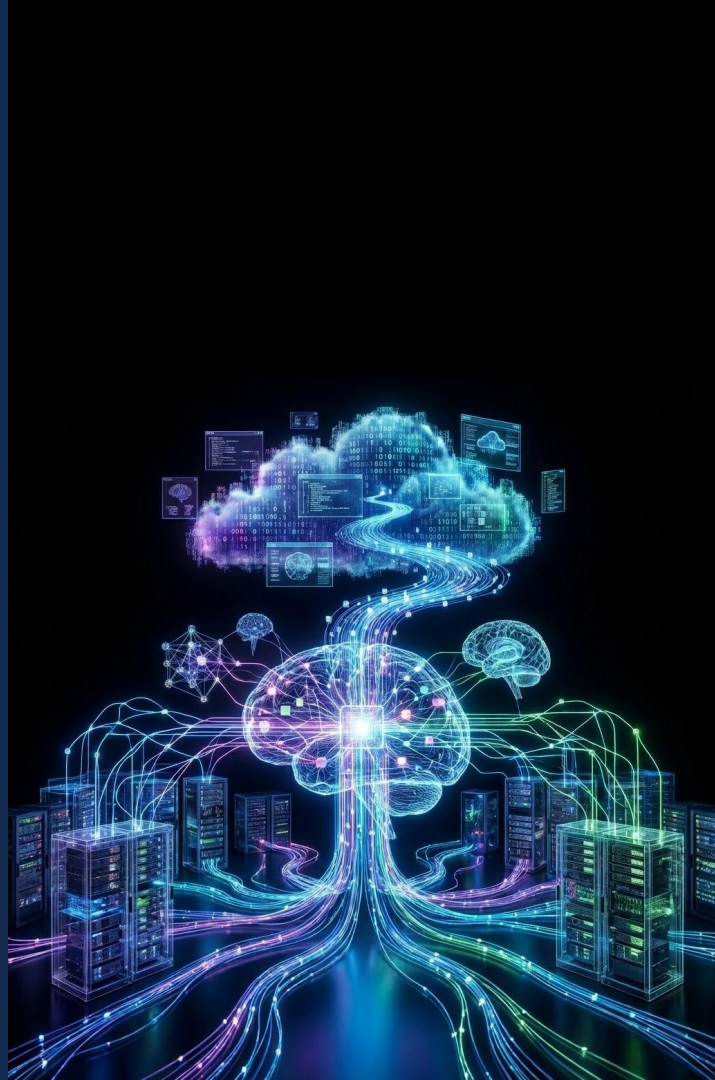

BES/HEP/NP Scientific User Facilities Infrastructure Partnership (AmSC SUF IP)

Establish **AmSC** as a **shared platform for the research communities** of multiple (accelerator-based) **BES**, **HEP**, and **NP Scientific User Facilities (SUFs)**, along with their industry and research partners, to **reduce the time to science** and **unlock challenging measurements and experiments**

SUFs power the DOE SC mission by advancing discovery

- >> 10,000 users per year, producing thousands of publication annually

Pose complex computing challenges and opportunities

- Exabytes of high-quality scientific data
- Multiple petabytes of curated data
- Hundreds of production workflows and AI models

AmSC Capabilities

- Federated identity and virtual organizations
- AI services: IaaS, training, model catalogs, agentic AI
- Robust data services
- Operations services / IRI foundation

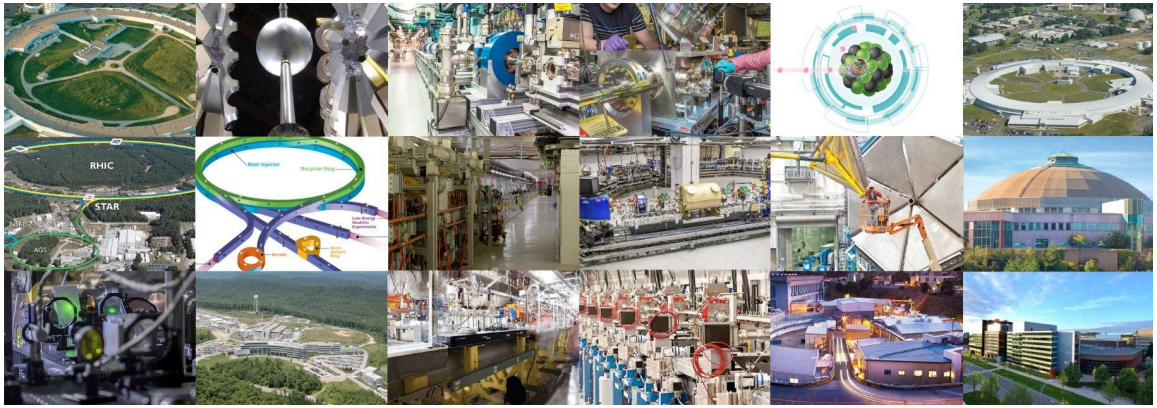

1. Agentic Accelerator Controls

Real-time guidance and optimization of particle accelerators
(LBNL, SLAC, ANL, BNL, …)

2. Cross-Facility Data Ingestion and AI Workflows for NP

Automated AI pipelines
(JLAB, BNL)

3. Inference as a Service for HEP and NP Data Analysis

Distributed IaaS for data analysis
(FNAL, LBNL)

4. Light and Neutron Source AmSC Demonstrator

AI-enabled workflows across 6 facilities
(ANL, BNL, LBNL, ORNL, SLAC)

5. OpenCosmo on AmSC

Seamless access to peta-scale cosmological data
(DOE Computing Facilities)

Aligned with: MOAT, SYNAPS-I, MAIQMag, NARAD, ISAAC, LAMBDA

APS, ATLAS, AWA at ANL, ATF, EIC, NSLS-II, RHIC at BNL, the Accelerator Complex, FAST, IOTA at FNAL, CEBAF at JLAB, ALS, BELLA at LBNL, SNS at ORNL, FACET-II, LCLS, SSRL, NLCTA at SLAC, and FRIB at MSU

Paolo Calafiura (LBNL, PI), Nicholas Schwarz (ANL, co-PI) - POCs: Salman Habib (ANL), David Lawrence (JLAB), Meifeng Lin (BNL), Auralee Edelen (SLAC), Jonathan Taylor (ORNL), Nhan Tran (FNAL), Remi Lehe (LBNL)

Genesis Mission | AMERICAN SCIENCE CLOUD

Argonne NATIONAL LABORATORY · BERKELEY LAB · Brookhaven National Laboratory · Fermilab · Jefferson Lab Thomas Jefferson National Accelerator Facility · OAK RIDGE National Laboratory · SLAC

STELLAR-AI: Contribution to AmSC

Deployment, Development, and Integration/Usage

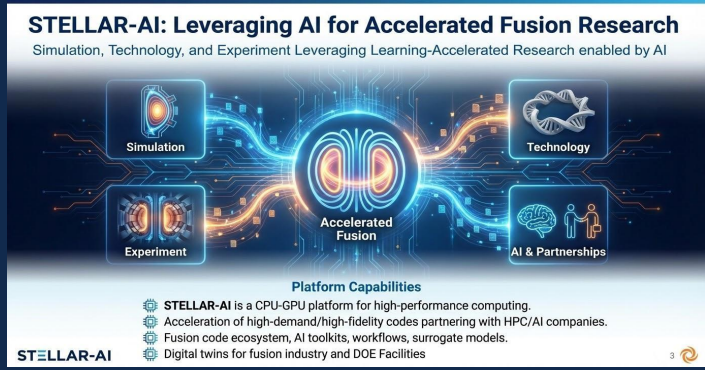

- **Deployment** of AmSC Software and Services
  - Accelerated Fusion requires Learning Everywhere: AI toolkits, workflows, surrogate models
- **Development** of AmSC Services
  - Consumers and producers of scalable software and services
  - Provides baseline ECP ExaWorks SDK and associated capabilities
- **Integration** into production computational campaigns
  - NSTX-U Digital Twin: A FES AmSC Project
  - Stellaris: Seed Fusion Foundation Model

STELLAR-AI: Contribution to AmSC Development

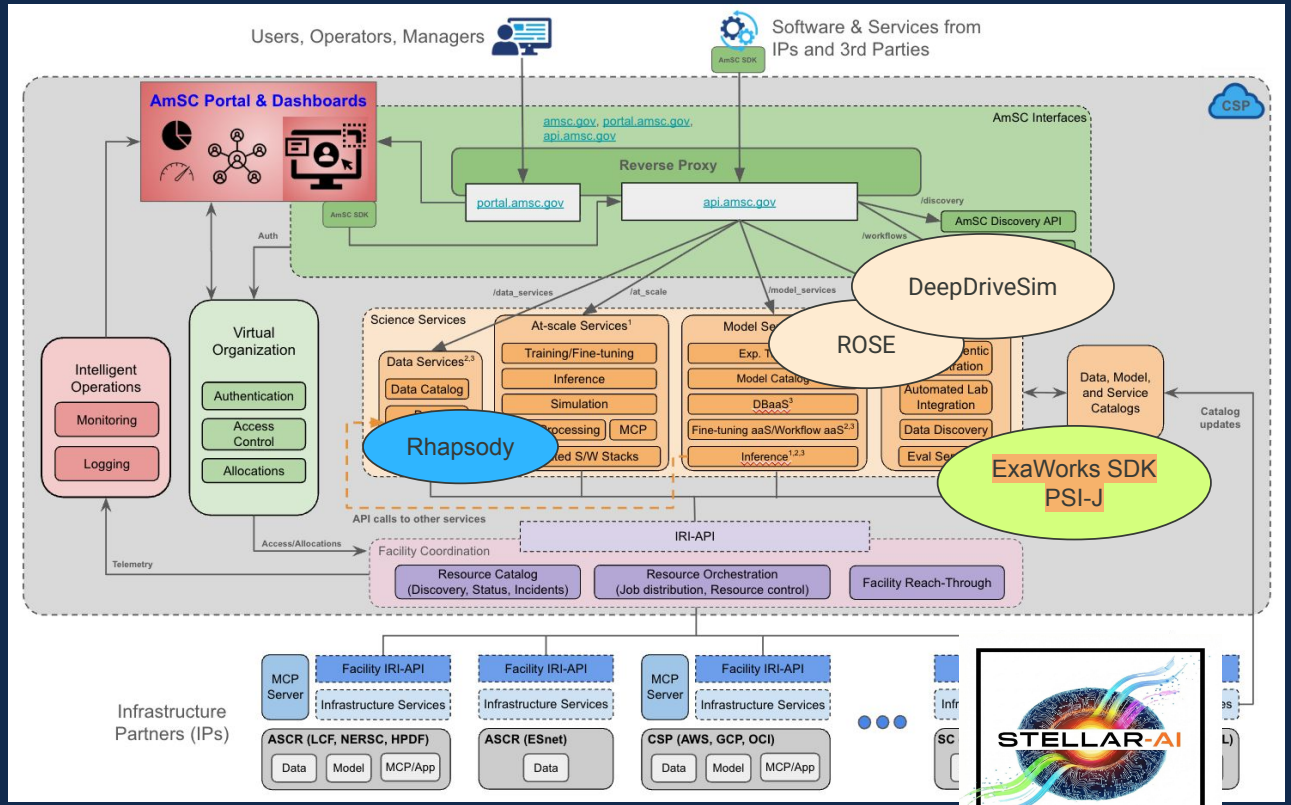

Getting Help

Questions, Problems, Requests

SEED PROJECTS

- Questions about using the AmSC environment?

Engage with your UE Liaison

AmSC INFRASTRUCTURE

- Feature Requests or requirements (routed to Prod Mgt)
- Bug reports (routed to appropriate AmSC team)
- Anything else; questions, comments, problems

help@amscproject.atlassian.net

COMMUNICATION TOOLS

- Questions about communication/ collaboration tools used by AmSC teams such as Slack, Jira, Confluence, Google Cloud, GitLab?

amsc-admins@ornl.gov

https://american-science-cloud.github.io/amsc-site/

Future Events

- We will continue to have webinars for the user community as we move forward
- Deeper dives into AmSC services, updates via hackathons and workshops for current early adopters
- https://american-science-cloud.github.io/amsc-site/events

Q&A Session