

An update on the Genesis Mission Platform in three contexts: AI, Quantum, HPC

September 14, 2026

Ben Brown, Ph.D.

Director, Facilities Division

Acting Director, Advanced Computing Technology Division

Advanced Scientific Computing Research



U.S. DEPARTMENT
of **ENERGY**

Office of
Science

[Energy.gov/science](https://energy.gov/science)

SYNAPS-I: Speeding up X-ray data analysis with artificial intelligence



Genesis Mission Summit, July 22, 2026




Dr. Dario Gil
DOE Under Secretary for Science
Director, Genesis Mission



<https://youtu.be/d5CSBRORcuw?t=4040>



Genesis Mission Platform should make complex infrastructure simple for users & empower complex agentic workflows



The Genesis Mission Platform encompasses all connected infrastructure, data, models, etc.



Users

Goal: make complex infrastructure simple for users & empower complex agentic workflows

The **American Science Cloud** is akin to the “operating system” that allows users to compose complex research workflows.



Data Services

Orchestration Services

AI Services

Data

Models

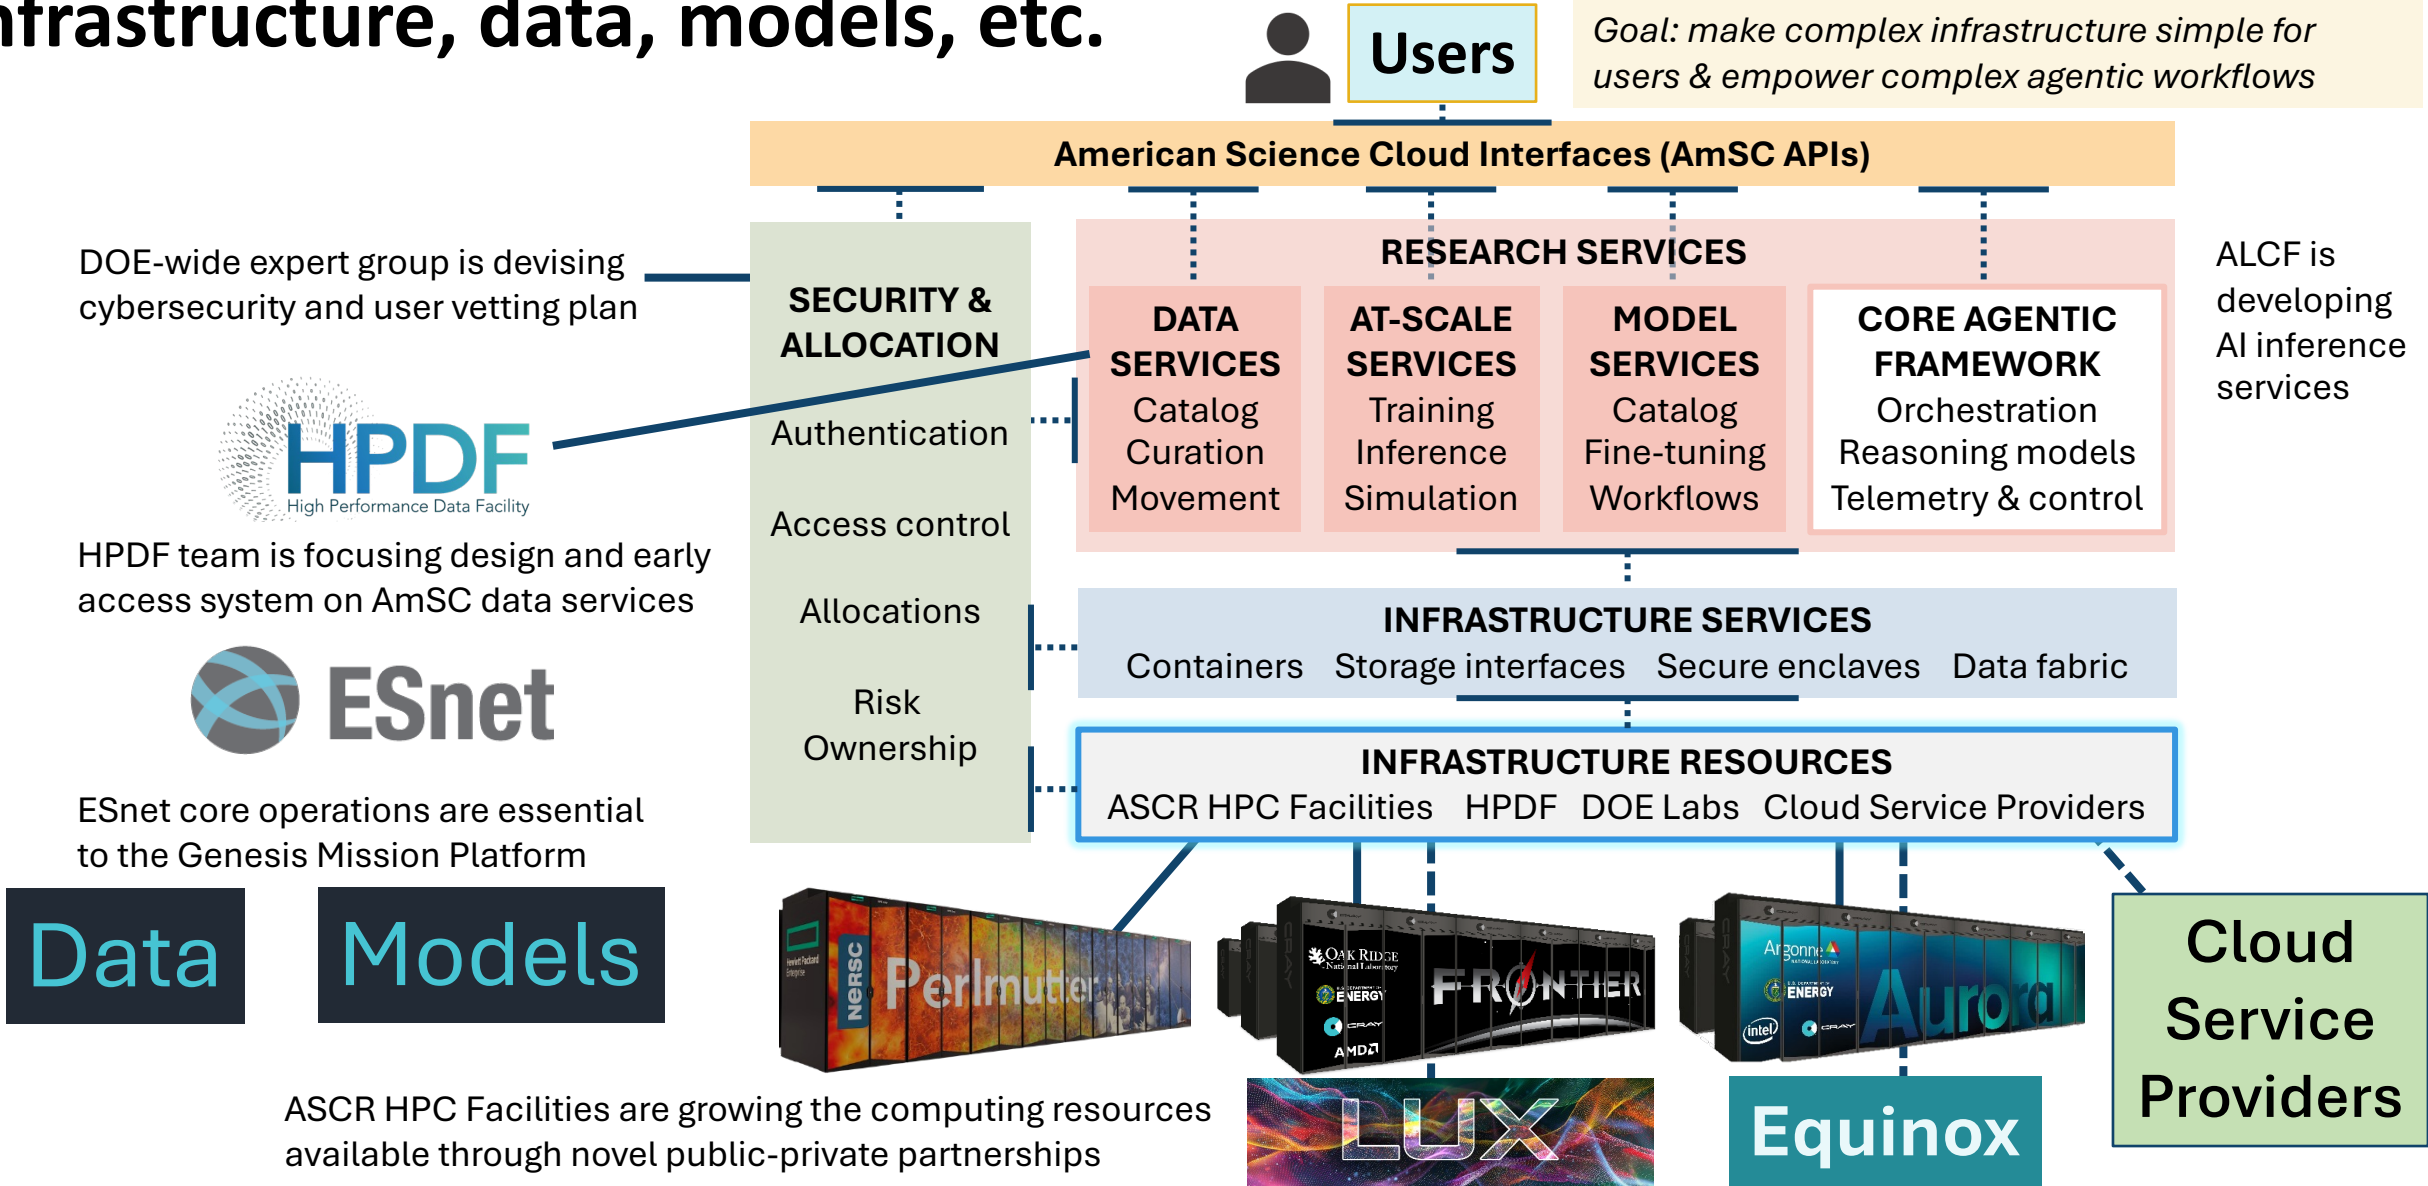


Equinox

Cloud Service Providers



The Genesis Mission Platform encompasses all connected infrastructure, data, models, etc.



The Genesis Mission Platform encompasses all connected infrastructure, data, models, etc.



Users

Goal: make complex infrastructure simple for users & empower complex agentic workflows



Data

Models



Equinox

Cloud Service Providers



Defining the Platform: A Requirements Framework

Federated Governance

- Multi-institutional registry: capability advertisement, trust at discovery, versioning
 - Skills and workflows as governed, versioned, OCI-compatible artifacts
- Cost attribution across agent-to-agent chains spanning projects and institutions

Operational Continuity

- Agent memory governance: isolation, lifecycle, auditable deletion
- Structured Agent Run Records supporting replay, audit, reproducibility
- Graceful degradation when delegating human sessions are revoked

Autonomy & Human Oversight

- Risk-adapted autonomy ladder with explicit controls per band
- Human approval gates propagate constraints across autonomous hops
 - Independent, agent-proof telemetry for high-impact actions

Security & Compliance

- ATO-compatible within existing DOE SSP inheritance
- NIST 800-53 coverage for non-human agent identities and agent-generated audit events
 - AI safety threat model: prompt injection, tool poisoning, supply chain integrity

Identity & Trust

- Every agent traceable to an owning human principal — no anonymous pooling or account sharing
 - Agent identity as a distinct platform primitive with full lifecycle
- Multi-hop delegation chains with cryptographic lineage and mid-chain revocation
 - Sender-constrained tokens at every institutional boundary

The American Science Cloud: Where we are today

3-MONTHS FEDERATED ACCESS

✓ 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.

January 2026

April 2026

July 2026

September 2026



Concept and architecture
Foundational platform

Initial capabilities
Prototype services
tested

Initial operational
AmSC platform for
Genesis teams


Integration of all
services into a
coherent platform

DOE announced Quantum Genesis on June 24, 2026

The Quantum Genesis will serve as a foundational pillar of the broader Genesis Mission, designed to usher in a new era of quantum computational power to accelerate scientific discovery and innovation


DOE Q Competition

A bold competition to demonstrate fault-tolerant quantum systems by 2028, targeting hundreds of logical qubits and hundreds of millions of operations.




National Quantum Computing User Facility

A first-of-its-kind facility providing U.S. researchers with access to advanced quantum computers, integrating HPC and AI to create a unified discovery platform



Quantum Computing Algorithms & Applications

Targeted research to identify and implement high-impact use cases for quantum computers that address DOE's Quantum Computing National Challenges.



Other Quantum updates

- **Request for Information on Scientifically Relevant Fault Tolerant Quantum Computers closed on June 9th** (see <https://tinyurl.com/26cr9pr4>)
 - Focused on potential deployment of a first-generation, scientifically relevant fault-tolerant quantum computer by 2028.
 - The RFI sought views on the viability of achieving technical targets, an overview of a technical approach, and the specific scientific applications such a system would enable.
- **ASCR has merged its two quantum computing user access activities at NERSC and OLCF under a combined QUEST program:**
 - QUEST = Quantum User Expansion for Science and Technology
 - Combines the Quantum Computing User Program (QCUP) at Oak Ridge National Laboratory and the Quantum Computing Access for NERSC (QCAN) program at Lawrence Berkeley National Laboratory.
 - Includes resources from IBM, IonQ, IQM, Quantinuum, QuEra.



HPC context: DOE is pursuing novel public-private partnerships for accelerated deployment of AI compute capacity


The Lux supercomputer at ORNL is a co-investment of AMD and DOE, built by HPE at ORNL.

AMD MI355X GPUs

5 PF (FP8) and 78 TF (FP64)

288 GB of HBM3E at 8 TBps


24 TB of NVMe per node



<https://www.olcf.ornl.gov/olcf-resources/compute-systems/lux/>




ASCR Announced HPC Systems and Projects Timeline




The evolving practice of science, in an HPC context

Responding to the DOE Mission



The evolving practice of science, in a high performance networking context: Load balancing and streaming of large data across the wide area network


The ESnet JLab FPGA Accelerated Transport (EJFAT) project has demonstrated load-balanced compute across any HPC Facility — simultaneously:



While the experiment was running...

... the EJFAT team added compute nodes at ORNL
... and still more compute nodes in FABRIC;
... with option to add compute from LabIT, Cloud, or multi-cloud environments **with NO event loss.**

Where we are today



Special Thanks

ASCR PARTNERS

AmSC Program Manager Xiaofeng Guo

ASCR Facilities Director Ben Brown

ASCR Advanced Computing Director Ben Brown (Interim)
and Pavel Lougovski

AmSC LEADERSHIP TEAM

Director Gina Tourassi

Deputy Directors Inder Monga · Arjun Shankar

CTO Sarp Oral

Product Manager Debbie Bard

Integration Architect Zach Mayes

CISO Ryan Adamson

L1 Leads Ashley Barker · Mike Brim · Chin Guok · Tom Uram


PMO Denise Hoomes

AmSC SCIENCE COUNCIL

Andreas Kronfeld · Michael Begel · Cami Collins · Joel England · Kjersten Fagnan · Oliver Gutsche · Joseph Ilon · Neeraj Kumar · Todd Satogata · Jana Thayer · Kevin Yager ·



Risk management must mature as new architectures mature, and this is a distinct issue from obtaining funding



Credit to Prof. Ilkay Altintas, UC San Diego, for “Useable, Useful, Used”

