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

AmSC
AMERICAN SCIENCE CLOUD

AmSC
AMERICAN SCIENCE CLOUD

November 13, 2025

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

U.S. DEPARTMENT of ENERGY

ORNL IS MANAGED BY UT-BATTELLE LLC
FOR THE US DEPARTMENT OF ENERGY

A Platform for Transformative Science

AmSC is a *secure*, *federated*, and *science-optimized* cloud environment that integrates the DOE's world-leading computing and experimental facilities, data resources, and high-performance networks

AmSC enables DOE scientists to create, access, and integrate world-class AI-ready datasets, run scalable model training on leadership-class systems, perform distributed simulations, control instruments, and move data efficiently across sites.

Start: Oct 1st, 2025, for 12 months

AmSC
AMERICAN SCIENCE CLOUD

Our Unique Infrastructure Enables Transformational AI Models

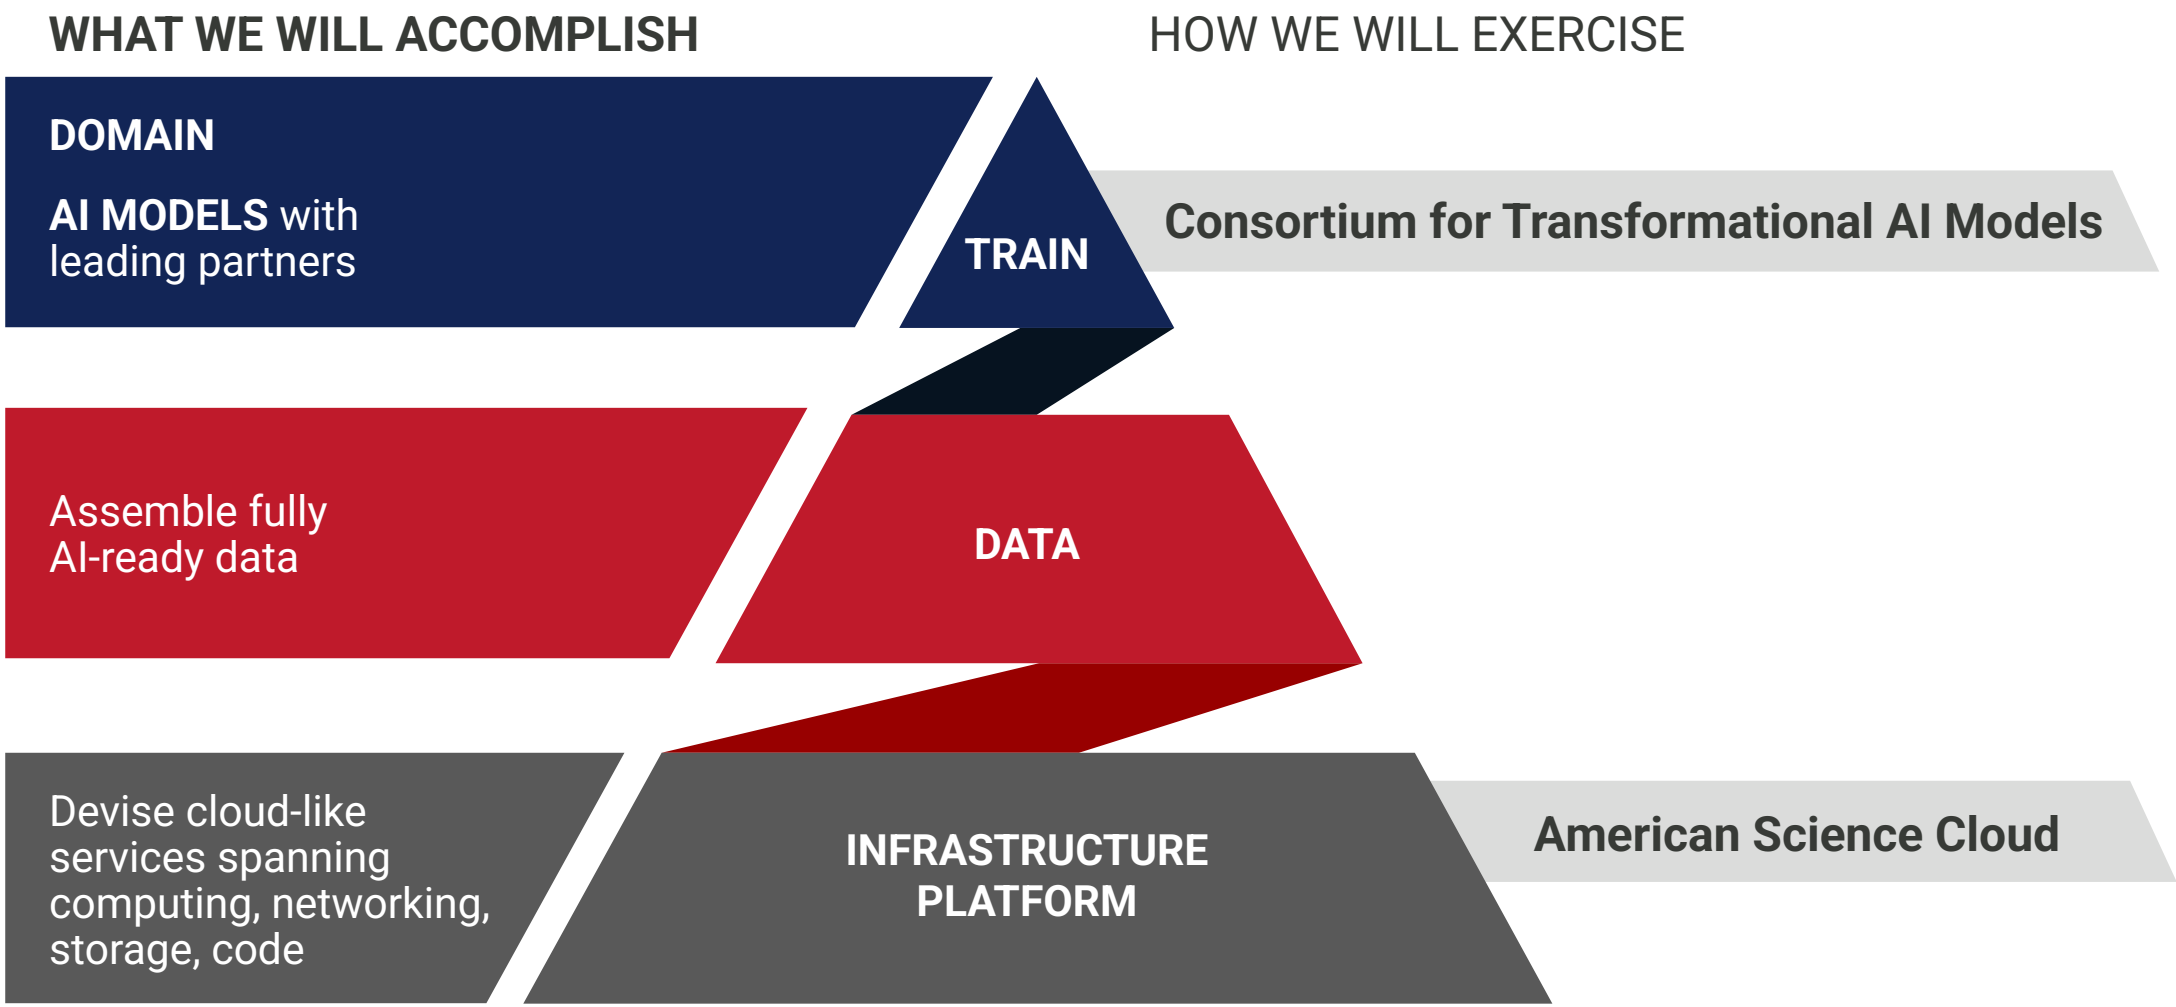

Terms/Definitions

INFRASTRUCTURE PARTNERS (IP)

✓ DOE labs that contribute, use, and integrate **resources and capabilities** which, federated under the Lead and Host structure, form the shared, integrated scientific cloud of the AmSC

MODEL TEAMS (MTs)

✓ Multi-lab groups within ModCon that develop **self-improving AI models** to advance science and engineering

Create integrated AI technologies that reason within their domains, analyze and predict data, enable autonomous experimentation, and build modular, multimodal architectures that make data AI-ready for others

SEED/FLASHLIGHT PROJECTS

✓ Mission-driven, high-impact scientific challenges that AI Model Teams address as **proof-of-concept use cases**, requiring both data and model readiness

LIGHTHOUSE CHALLENGES

✓ Mission-driven, high-impact scientific challenges that **showcase and guide the application** of AI across domains

Discovery Science

- Understanding the universe, from quarks to cosmos
- Seeing molecules in motion
- Discovering new quantum algorithms

Energy Dominance

- An intelligent, resilient grid
- Fusion you can plug into
- Advanced nuclear, faster and safer

National Security

- Discovering mission-ready materials
- Securing critical materials
- Accelerating advanced manufacturing

AmSC AMERICAN SCIENCE CLOUD

Infrastructure Partners

Infrastructure Partner(s)	Lab(s)	Key Contacts (POC/PI, Leads)
OLCF	ORNL (Lead Host)	Gina Tourassi (Lead PI); Arjun Shankar (OLCF); Sarp Oral (CTO)
ALCF	ANL	Mike Papka
ESnet	LBNL	Inder Monga
NERSC	LBNL	Sudip Dosanjh
HPDF	TJNAF, LBNL	Kate Mace, Lavanya Ramakrishnan
C3	PNNL	Robert Rallo
BES/HEP/NP Scientific User Facility	LBNL, ANL, TJNAF	Paolo Calafiura

Infrastructure Partner(s)	Lab(s)	Key Contacts (POC/PI, Leads)
DeepLynx	INL	Jeren Browning
EDX/MACH	NETL	Kevin Kuhn
Fermi Data Platform	FNAL	Jim Amundson
HADIS	TJNAF	Ilya Baldin
QCD	FNAL, TJNAF	Michael Wagman, Robert Edwards
S3DF	SLAC	Jay Srinivasan
SDCF	BNL	Adolfy Hoisie
Stellar-AI	PPPL	Shantenu Jha
VEE-ARIES	NREL	Kristi Potter

Multi-Lab and Public-Private Partnership

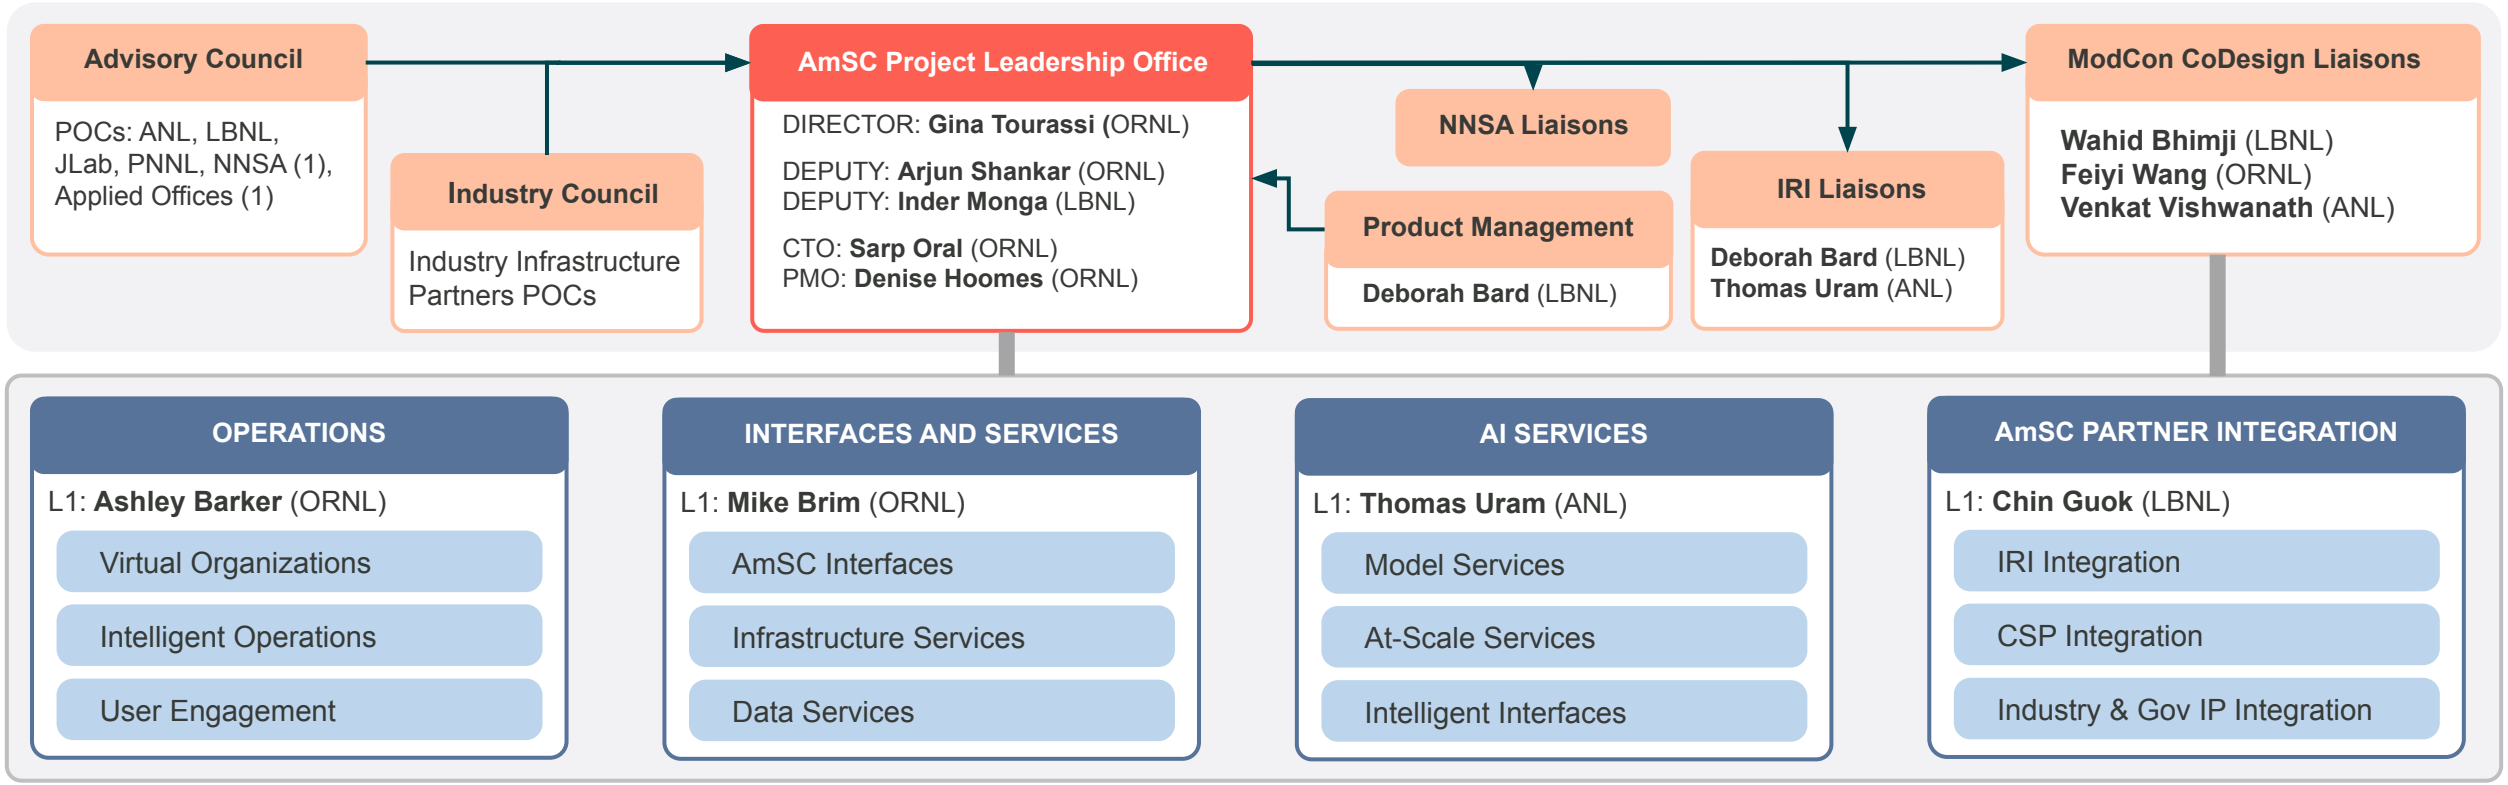

✓ **Industry Collaborators:** AWS, Microsoft, Google, NVIDIA, AMD, Dell, HPE, Cisco, Nokia

✓ Federated leadership with **Laboratory Advisory Council** and **Industry Council**

AmSC AMERICAN SCIENCE CLOUD

The American Science Cloud (AmSC)

A Unified Cloud Platform for Transformational AI and Science Across DOE

Mission

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

Core Teams

Operations

Interfaces & Services

AI Services

AmSC Partner Integration

Product Management

What We Deliver

✓ Unified APIs for data, model, and workflow access
✓ Secure, federated identity (Virtual Organization model)
✓ High-speed data fabric powered by ESnet
✓ Seamless interoperability between DOE facilities and CSPs

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

AmSC is a *secure*, *federated*, and *science-optimized* cloud environment that integrates the DOE's world-leading computing and experimental facilities, data resources, and high-performance networks

AmSC
AMERICAN SCIENCE CLOUD

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 Model 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

Complements capability delivered by
The American Science Cloud

Partners
17 DOE Labs + Industry + Academia

AmSC
AMERICAN SCIENCE CLOUD

ModCon Co-Design Liaisons

Facilitating Productive AmSC Infrastructure for ModCon Science Teams

- **Translate requirements from ModCon** to AmSC technical teams and leadership
- **Communicate** AmSC plans and roadmaps back to ModCon teams
- **Advise on feature priorities** to achieve ModCon science goals
- **Develop synergies** between AmSC and ModCon to advance technical and science objectives.

Liaisons align ModCon priorities with AmSC development to advance scalable and impactful scientific innovations

AmSC
AMERICAN SCIENCE CLOUD

AmSC–ModCon

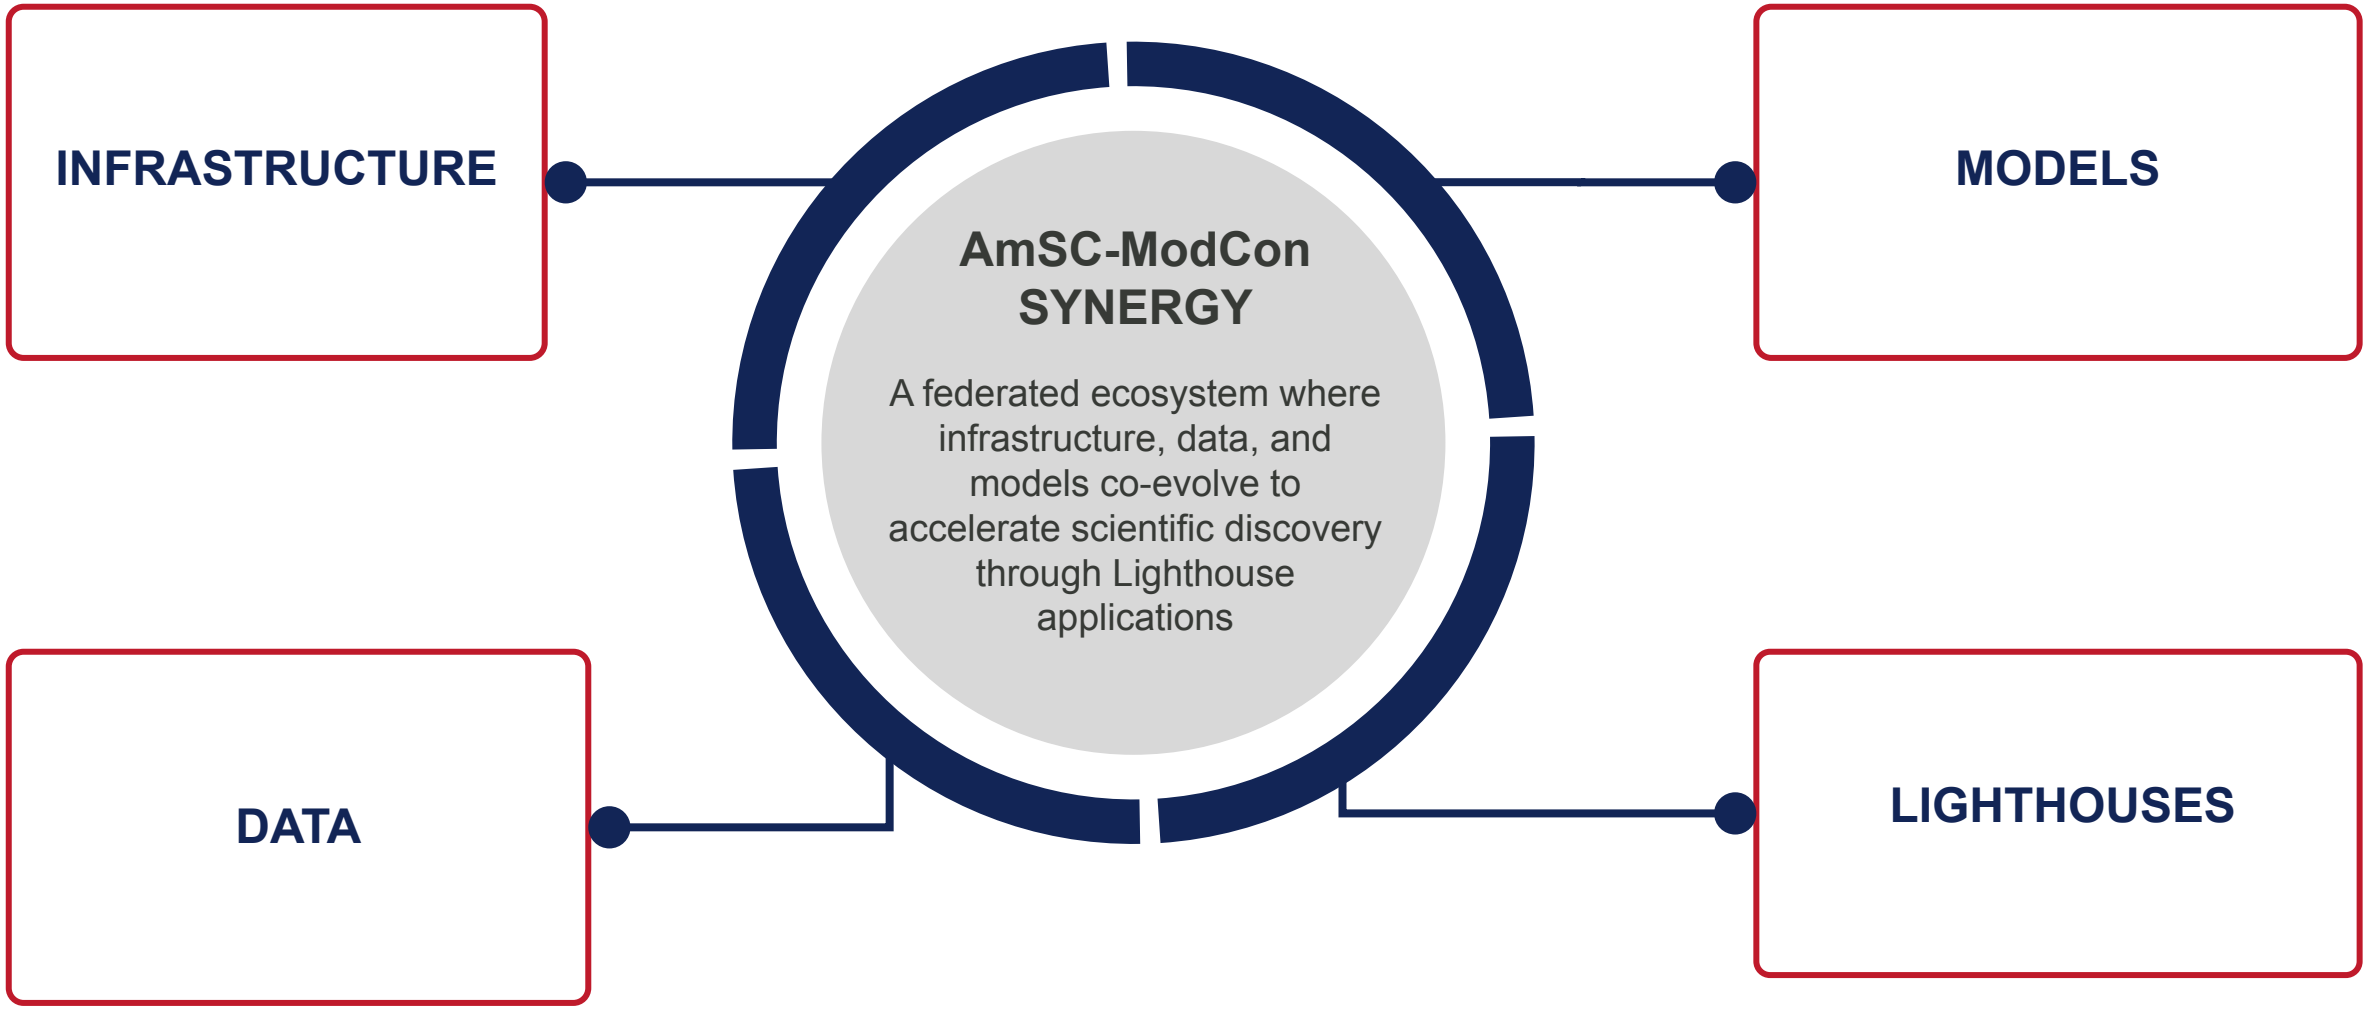

AmSC
AMERICAN SCIENCE CLOUD

AmSC Architecture Overview

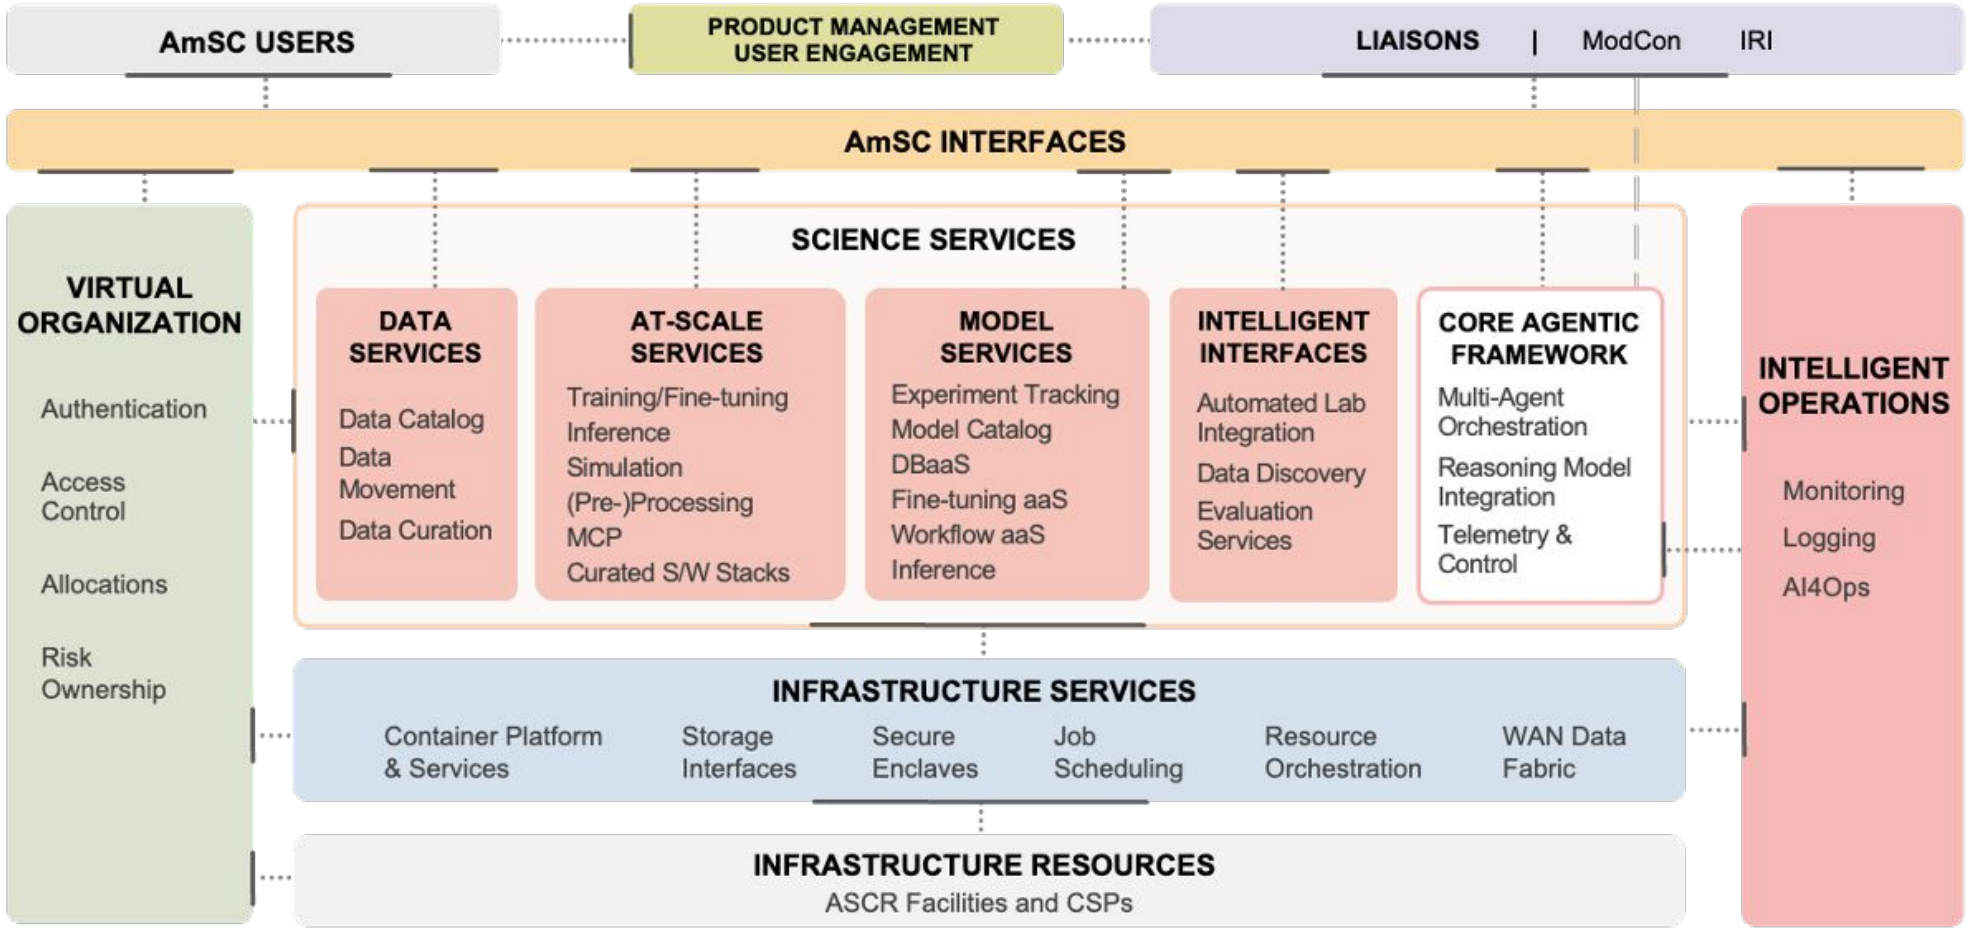

AmSC Detailed Architecture

Science Services

- There is only a single instance of the science services, i.e., a single point of entry for each service, but services can distribute work to multiple facilities.
- Each service (e.g., Data Catalog, Inference, Simulation, …) must expose a Northbound-API that can be proxied through the AmSC Interfaces API end-points.

Infrastructure Backbone

- HPC Systems (Aurora, Frontier, Perlmutter)
- Kubernetes-based container federation
- CSP Partners
- Infrastructure Partners
- Object Storage
- WAN Data Fabric (ESnet L3VPN)
- Secure enclaves

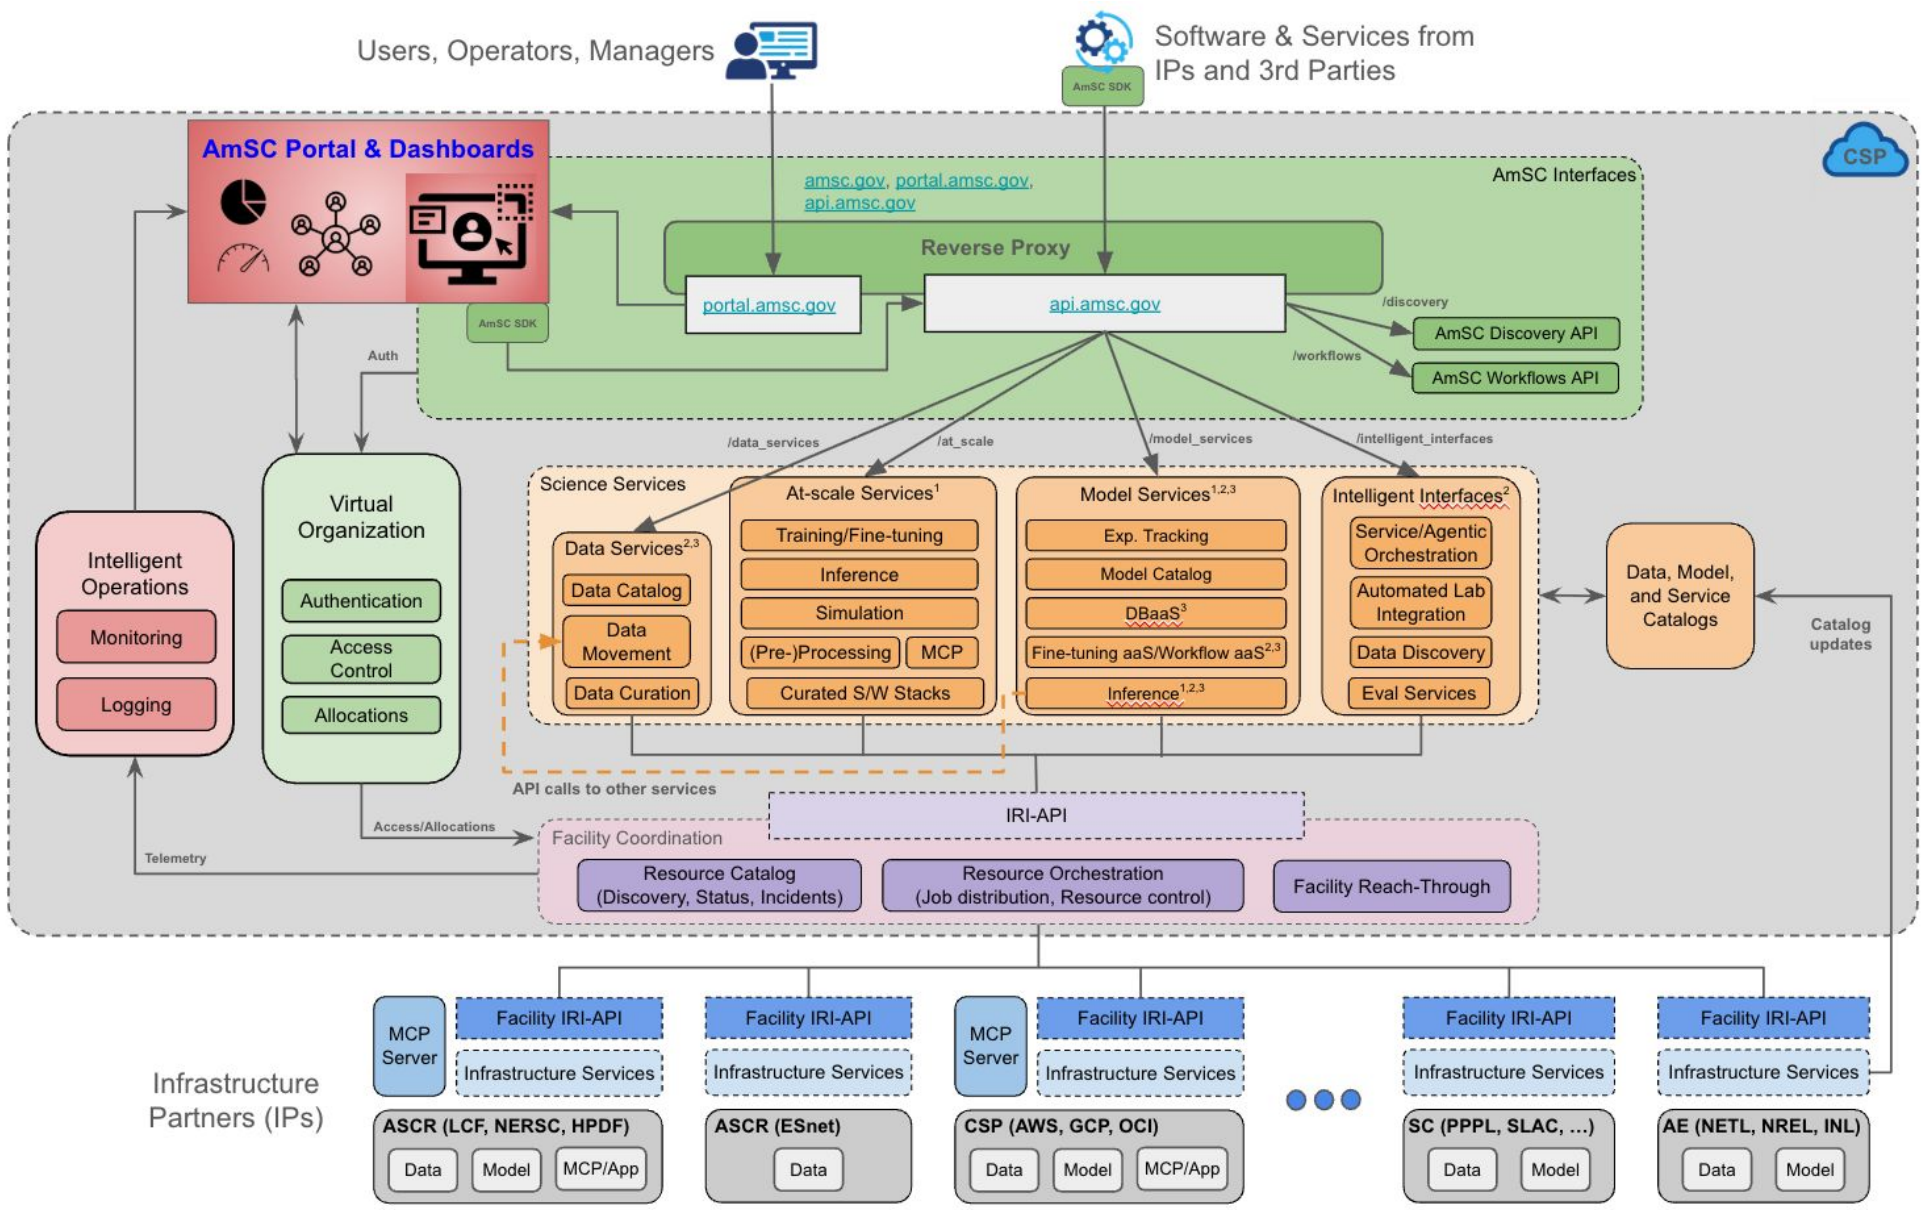

Deliverables and Impact

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

LONG-TERM VISION
Foundation for a National Strategic Computing and AI Reserve, ensuring resilience and leadership in AI-driven science

AmSC
AMERICAN SCIENCE CLOUD

AmSC project scope in a nutshell

We are working on solving three major problems

- Simplify account creation and authentication via common processes
- Improve automation for users and programmatic agents via a common API
- Support AI and simulation scientific workflows supporting flexible training and inference pipelines

The list is not exhaustive; there are multitude of other small and large problems

Operations

Unified Access, Smarter Operations, and Engaged Users for Seamless Science

INTELLIGENT OPS

- **Monitor system health** through unified cross-facility metrics & dashboards
- **Enable debugging & performance insights** to support users & improve workflow reliability
- **Automate fault detection & optimization** using AI-driven telemetry and anomaly detection
-

AI-driven observability to monitor, debug, and optimize AmSC end-to-end.

USER ENGAGEMENT

- **Onboard & train users** to effectively leverage AmSC
- **Support & collaborate** with SEED/FLASHLIGHT teams to resolve issues & improve workflows
- **Align user needs** with platform capabilities & science goals

Connecting researchers with the tools, training, and support they need to use AmSC effectively.

VIRTUAL ORGANIZATION

- **Integrate support for federated identities** to simplify user experience
- **Streamline access processes and policies** for improved, cross-facility entry
- **Coordinate allocations** across available resources

Enable more seamless integration of user access and allocation management across AmSC resources.

AmSC
AMERICAN SCIENCE CLOUD

Interfaces and Services

Seamless User Interactions, Federated FAIR Data, and Unified Science Infrastructure

AmSC INTERFACES

- **User-focused web interfaces, APIs, and SDK** for using Science Services
- **Consistent and seamless user experience** for interactive use of AmSC capabilities

Expose AmSC platform capabilities to users and workflows in a consistent, reusable manner.

DATA SERVICES

- **Data catalogs and services** to publish, find and access AI-ready scientific datasets
- **Data curation technologies and pipelines** to prepare data for use with AI Models
- **Data movement** throughout the AmSC IP ecosystem

Provide a robust data management environment that promotes FAIR principles.

INFRASTRUCTURE SERVICES

- **Integrate IP compute, data, and service resources** into the AmSC platform
- **Uniform infrastructure capabilities** for use by Science Services
- **Resource controls** in coordination with Virtual Organization

Unify the capabilities provided by Infrastructure Partners to Science Services.

AmSC
AMERICAN SCIENCE CLOUD

AI Services

Scalable AI Model Training, Fine-Tuning, and Inference, SOTA Tools, and Agentic Orchestration

AT-SCALE SERVICES

- **Large-scale AI model training and fine-tuning** guidance
- **Deployment of AI Models** for scalable inference
- **Catalog of simulations** with execution guidance

Support large-scale simulation and model training and inference

MODEL SERVICES

- **Searchable catalog** of language and science models
- **Database hosting** in support of traditional projects and for AI vector embeddings
- **Array of model development tools** to support rapid refinement

Industry-standard support for model development, tuning, and inference

INTELLIGENT INTERFACES

- **Agentic framework development** in collaboration with ModCon
- **Active data discovery** for model training and evaluation
- **Automated integration of AmSC models** with physical labs

Integration of intelligent services within and for AmSC

18

Partner Integration

Connecting Visions, Enhancing Capabilities, and Unlocking Potential

IRI INTEGRATION

- **Facility APIs** for using Infrastructure Services
- **Consistent and well defined interface** for infrastructure partners to participate in AmSC capabilities

Providing a framework to seamlessly integrate user facilities, data assets, and advanced computing resources

CSP INTEGRATION

- **Unified interface** for managing CSP capabilities that is aligned with and enhances the IRI API
- **Coordinate interaction with CSPs** to reduce friction when accessing cloud resources

Enable CSP resource integration with AmSC

INDUSTRY & GOVERNMENT IP INTEGRATION

- **Collaborate with non-ASCR IPs** on workplans, requirements, integration, and AmSC progress
- **Building bridges with industry partners** to enhance the AmSC platform

Create an open, diverse ecosystem of IPs for AmSC, and integrating unique resources

AmSC
AMERICAN SCIENCE CLOUD

Product Management

Delivering science capabilities for AmSC users

Collect and prioritize science requirements from stakeholders

- **Gather, maintain and communicate up-to-date requirements** from the AmSC user communities
  - including IP teams, ModCon SEED/FLASHLIGHT projects, Lighthouse challenges and also broad community requirements
- **Advise on feature priorities,** based on requirements to complete science demos
- **Track delivery schedule** and work with CTO to prioritize and advocate for critical features
- **Help define the long-term scientific vision** in collaboration with the CTO and the project leadership (beyond the year 1 demos) and keep work aligned with the long-term vision as well as short-term milestones and demos.

The Product Management Team will ensure the platform meets the needs of AmSC users

AmSC AMERICAN SCIENCE CLOUD

Integration Models Across IPs

Three Models Of AmSC Integration

INFRASTRUCTURE DEVELOPERS

✓ Build and host AmSC services
✓ Provide compute, data, and network foundations
✓ Enable federation with DOE IRI

Develop the AmSC platform

INFRASTRUCTURE USERS

✓ Deploy and use AmSC APIs locally
✓ Integrate AmSC services into workflows and instruments
✓ Contribute models

Demonstrate science services

HYBRID INTEGRATORS

✓ Combine DOE and CSP infrastructures
✓ Validate secure, federated execution
✓ Support protected and cross-domain workflows

Bridge DOE and industry environments

AmSC
AMERICAN SCIENCE CLOUD

An IP "Infrastructure User" of AmSC might:

- Move data from an experiment facility to an AmSC facility and register with the Data Catalog.
- Add a model to the Model Catalog.
- Train model on the experiment data.
- Report back results to science team.
- *Initially performed interactively via AmSC portal, eventually all automated via AmSC APIs.*

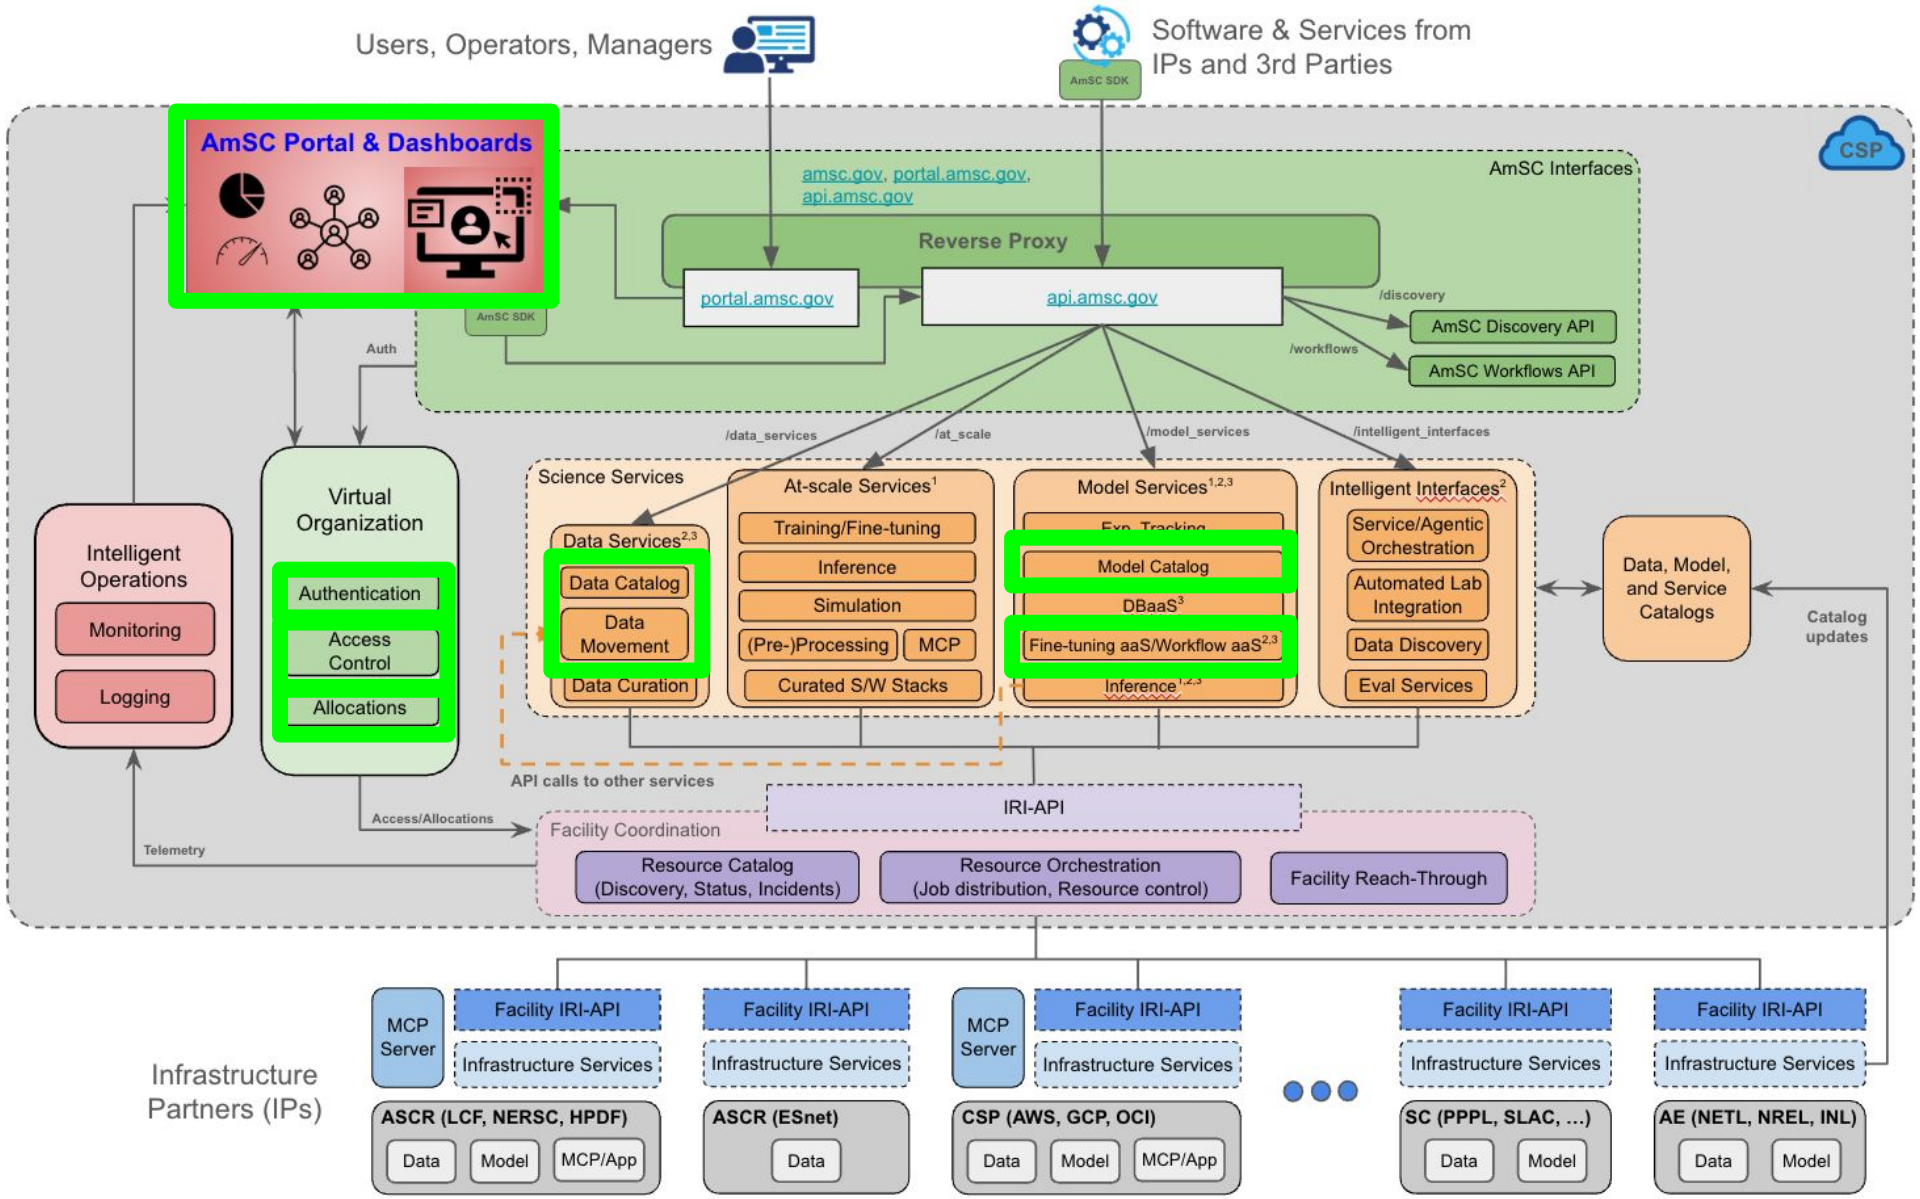

An IP "Infrastructure Developer" might: Integrate existing data infrastructure into AmSC

- Deploy IRI-API and facility coordination services onto facility infrastructure.
- Deploy AmSC data services at facility and make datasets accessible to Data Catalog (with access restricted to designated users).
- Provide monitoring information, eg facility status.

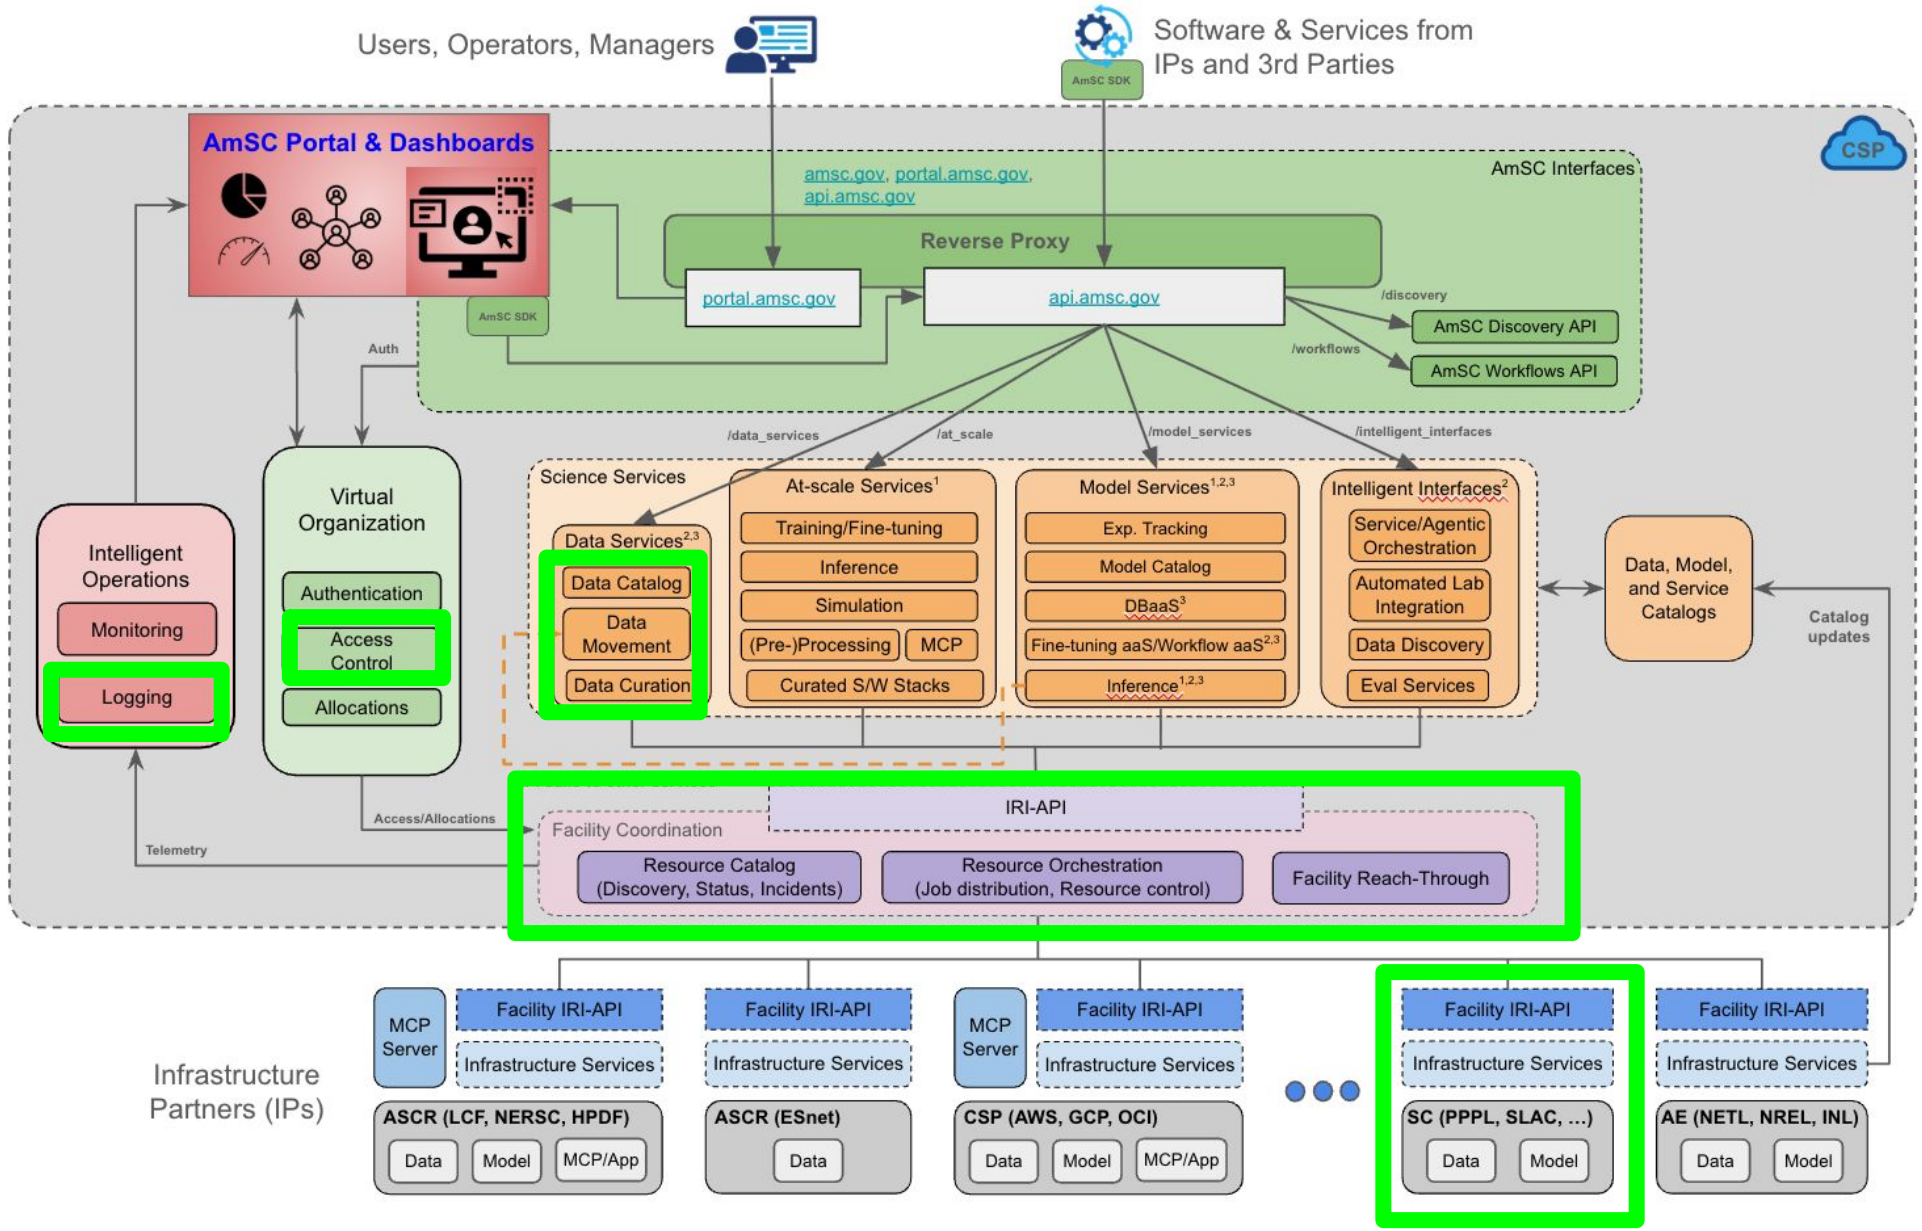

Rules of Engagement

AmSC and ModCon are in a Sprint

- Unprecedented **pace and visibility**;
- Operating in a **national spotlight**, with performance and outcomes under close review.
- **Targets and priorities may shift quickly**; responsiveness will define our success.

Our Operating Principles

- **Stay Aligned:** Communicate early, often, and clearly. Be prepared for multiple channels.
- **Move Fast, Together:** Favor action and iteration over perfection. Adapt as priorities evolve.
- **Own Outcomes:** Everyone is accountable for quality, speed, and mission impact.
- **Decide and Deliver:** Be ready to pivot and make hard choices when called upon.
- **Support Each Other:** High tempo, high trust. Professionalism and respect are non-negotiable.

Urgency, Speed, and Excellence = Our Foundation of Success in this Sprint

AmSC
AMERICAN SCIENCE CLOUD

What Success Looks Like for AmSC

1
Technical Success

Technical Success — Build the Federated Ecosystem for AI and Accelerated Science
Deliver a **secure, interoperable multi-site cloud** linking DOE HPCs, user facilities, and public clouds.

- **Milestones:** Federation architecture, identity management, AI workflows (training, inference, agentic, …), AmSC API - leading to a Minimum Viable Product (MVP)

2
Programmatic Success

Programmatic Success — Integration & Coordination

- **Infrastructure Partners:** Demonstrated workflows that span IPs to enable AI model development (with IP data) and inference (agentic and otherwise). Compute IPs run major AI campaigns.
- **Model Consortium:** Demonstrated coordination via SEED/FLASHLIGHT projects and *model teams* driving flagship science applications.

3
Strategic Success

Strategic Success — National Impact

- **SC-Wide:** Supports *Lighthouse Problems* showcasing cross-SC discovery.
- **DOE-Wide:** Coordinated with NNSA's AI Innovation Platform to advance shared architectures, security, sovereign AI/HPC leadership

AmSC
AMERICAN SCIENCE CLOUD

Success = Technical Excellence + Programmatic Integration + Mission Impact

Getting Help

Questions, Problems, Requests

SEED/FLASHLIGHT 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)
- Any other queries

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/

AmSC AMERICAN SCIENCE CLOUD

SEED/FLASHLIGHT Team -> User Engagement Liaison

Project Name	User Engagement Liaison
PROMETHEUS	Sam Foreman
CM2US	Brian Etz
Fusion-FM	Kyle Felker
GRID-UNI	In-Saeng Suh
SYNAPS-I	Marshall McDonnell
FAMOUS	Brian Etz
Q2C Science at Scale	Gustav Jansen
Specs-to-Silicon	William Godoy
MAIQMag	Marshall McDonnell
CombustionFM	Phil Roth
AI-Microelectronics	Kyle Felker
MOAT	Jeff D'Ambrogia
AI-CodeDevelopment	Helen He
AI-QuantumAlgo	In-Saeng Suh

AmSC
AMERICAN SCIENCE CLOUD

Next Webinar – January

- Follow-on part-2 webinar will be scheduled for January
- A deeper dive into AmSC services, including updates on timelines and the availability of various infrastructure/platform components

Q&A Session