ORNL Town Hall

Genesis RFA and AmSC Connections

Frequently Asked Questions | April 6, 2026

Q. What are the opportunities for expert groups with long experience to engage with AmSC developers and help. One example is the Global Network Advancement Group (GNA-G) and its SENSE and Data Intensive Sciences Working Groups (H. Newman, Caltech)

This webinar is for potential respondents to the Genesis RFA, and that is one vehicle to engage with Genesis mission funded projects like AmSC and ModCon. We will introduce the leads of infrastructure partners involved in AmSC later in the webinar, and we hope that the list of those partners will expand as we make progress as the Genesis mission expands, providing more opportunities for the larger community to engage and contribute.

Q. How will system monitoring be architected across the AMSEC multi-lab ecosystem to ensure operational visibility, reliability, and cross-site interoperability for telemetry? Has that been fully-scoped?

We do have an Intelligent Operations planned for AmSC that will not only monitor and collect data but will be the platform over which AIOps will be developed. For the MVP we are putting a framework and data collection techniques together, and the RFA has options for people to propose AIOps Techniques.

Q. Will a list of available hardware and capabilities from AmSC and the Platform be provided for the July through December timeframe?

This question was discussed in the townhall Q&A, please see the recording.

Q. Where and when will the recording of this webinar be available?

The recording (and slides) are available: https://amsc.energy.gov/#events

Q. What exactly is AI-ready data and what is the process? Is there a standards-based approach, such as MLOps?

There is no single standard for “AI-ready”. Domain scientists need to define the necessary transformations to take raw scientific data and make it suitable for use with an existing AI model.

Q. Is the AmSC architecture built on an established architecture framework (Zachman, DoDAF, TOGAF, etc.), and is it composable for future needs?

TOGAF, with some elements of ArchiMate. Composability, extensibility, and utilizing open-source software are three of our core design goals.

Q. Is integration with AmSC a technical requirement for responsive Phase I applications, or is it aspirational? And for intensive workflows is there an expectation for teams to budget independent compute for Phase I, given that AmSC will still be in prototype stage at project start in July 2026?

We are providing this informational webinar for you to have enough knowledge to determine if AmSC integration is a requirement or optional in your RFA response. Please do read the RFA call to determine what is a requirement or not.

Regarding allocations for compute, these will be determined by DOE program managers as needed after the awards have been made.

Q. What is the difference between using AmSC and using commercial cloud services such as Azure and AWS? Why use AmSC rather than AWS, which already manages a lot of DOE data?

AmSC integrates DOE compute, storage, data and network resources in addition to and along with public Cloud resources, Public cloud resources. DOE resources are available based on allocation and need as determined by DOE program managers and aligned with the mission. AWS and other cloud providers are surely managing some DOE data but there is a lot of data and data sources that are still within the DOE compute ecosystem. AmSC enables the glue with the complex's significant scientific instruments.

Q. What about cybersecurity and interoperability?

Our platform architecture is designed using zero-trust principles. Interoperability is achieved through an API-first approach, specifically through the IRI and AmSC APIs.

Q. How can we collaborate to get our data catalogs, agents, models, etc. on AmSC and vice versa?

This question was discussed in the townhall Q&A, please see the recording.

Q. For infrastructure partners seeking to provide value to AmSC capabilities, are there proactive measures we can take? For example, are there any established architectural patterns or reference models we could leverage?

As an IP, the best way to prepare for integrating your infrastructure and offered resources into the platform, is to reach out to User Engagement via the support mechanisms shared in the webinar. They can put you in contact with members of the AmSC architecture team, who can provide access to our AmSC RFC repository, which contains technical specification of the platform, identity, data, etc.

We are working on a shareable, Kubernetes-based reference implementation that IPs can deploy on their infrastructure.

Q. Do we need to write proposals to request computing time at individual HPC sites?

This question was discussed in the townhall Q&A, please see the recording.

Q. Can NETL's EDX OAuth provider services be included to support our user base's easy access as part of AmSC's Federated Identity?

Our goal is to work with all the laboratories to Federate with their IDP of choice. Please reach out to us through the contact form to work together.

Q. Will AmSC infrastructure be open source? Can researcher provide improvements for workflows that are not supported? How is it different than, let’s say AWS or Google Cloud?

Software released as part of the AmSC ecosystem must use an open-source license and allow reuse and collaboration across the AmSC platform, DOE facilities, and partner institutions. A specific license under which to release software has not yet been chosen. Researchers will be able to collaborate by submitting improvements and feature enhancements to workflows and infrastructure.

Q. How much scientific data is currently housed in this facility? How is the data curated and validated?

Most data is not stored in the platform. It lives within storage at infrastructure partners and is catalogued within IP catalogs that are then federated with the AmSC Data Catalog.

Q. Where can we request access to an american science cloud account?

This question was discussed in the townhall Q&A, please see the recording.

Q. Like the earlier question, our current ML workflow is built on Amazon Bedrock, and the phase 1 call requires a demo or artifact in six months. From a schedule and risk perspective, it would make the most sense to continue using Bedrock through this initial phase . Will AmSC be able to provide access to Bedrock for us during this phase, so we can focus on demonstrating capabilities in this tight timeframe?

Access to Amazon Bedrock, and other CSP and vendor model providers, will be provided though the Model Access Gateway (MAG). Access to Bedrock will be provided to users of the platform, subject to a given project’s token allocation.

Q. Will AmSc be ready to use for A work?

This question was discussed in the townhall Q&A, please see the recording.

Q. will it be possible to integrate domain-specific MCP servers that run on user/experiment infrastructure to the AmSC-based agentic workflows? Is there (or will there) be a specific authentification protocol to manage access to these domain-speciifc MCP servers?

Yes. Via OIDC, using issued tokens tied to human user or service account that initiated the agentic workflow.

Q. Could one propose to use commercial products (such as LLMs on the market) in combination with AmSC capabilities?

Yes, but AmSC will not be in position to acquire or service commercial models beyond what has been planned (exception being MAG and frontier models GM providing). If one suggests commercial products outside these boxes, AmSC can't acquire it but perhaps service it.

Q. Will AmSC be capable of handling CUI?

In FY27 and beyond, AmSC plans to begin incorporating capabilities to handle CUI/FISMA High data, likely based on the Scalable Protected Infrastructure (SPI) / CITADEL framework, also incorporating CSP secure computing capabilities.

Q. How is AmSC related to the American Science and Security Platform and the Genesis Mission Platform?

live answered

AmSC is the core part of Genesis Mission Platform and American Science and Security platform. Since this is one platform for DOE, there are components that might be targeted towards secure workflows that are not covered by the presentation today.

Q. What is the support model for this multi-lab platform, and how will it be updated over time?

This question was discussed in the townhall Q&A, please see the recording.

Q. Is there any public documentation for these AmSC capabilities?

This question was discussed in the townhall Q&A, please see the recording.

Q. Is it possible to deploy customized MCP servers on AmSC and then call those services via AmSC?

This question was discussed in the townhall Q&A, please see the recording.