

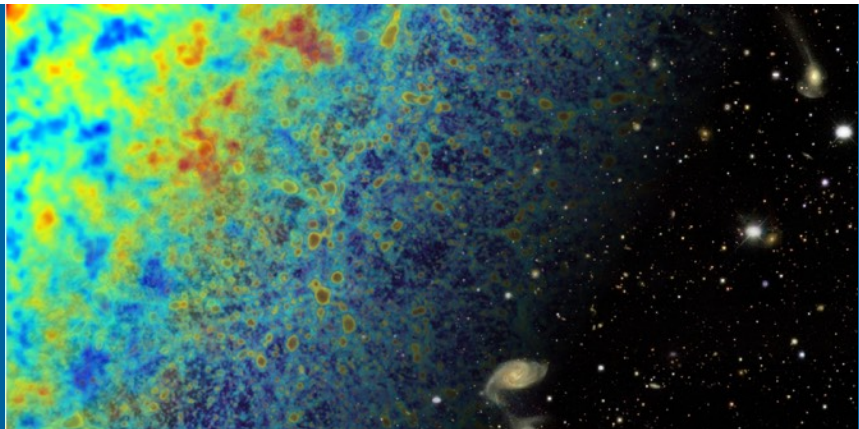
MARCH 11, 2026

ALCF SERVICE-ENABLED SCIENCE SERIES



OpenCosmo

Sharing and Analyzing
Cosmological Simulations Across
DOE Computing Facilities



PATRICK WELLS

High Energy Physics Division
Argonne National Laboratory

MICHAEL BUEHLMANN

Computational Science Division
Argonne National Laboratory

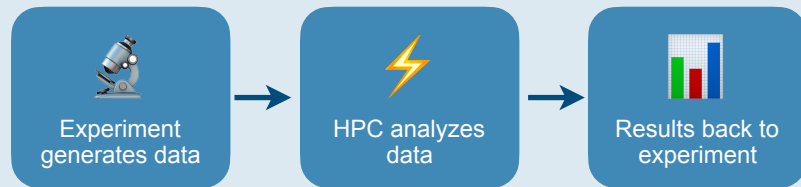


ALCF SERVICE-ENABLED SCIENCE SERIES

Two complementary perspectives on the same infrastructure

WEBINAR 1 Oct 2025

ALCF on-Demand APS + ALCF



Real-time pipeline · single facility · seconds per dataset

WEBINAR 2 This talk

OpenCosmo Cosmology + ALCF/NERSC/OLCF



Interactive queries · multi-facility · PB-scale simulations

Shared Infrastructure:

GLOBUS AUTH

GLOBUS COMPUTE

GLOBUS FLOWS

GLOBUS TRANSFERS

ALCF SERVICE ACCOUNTS

ALCF DEMAND QUEUES

EAGLE SHARED FILESYSTEM

The Team



Michael Buehlmann



Salman Habib



Benoit Cote



Idu Adeniji



Tom Uram



Nick Frontiere



Will Hicks



Patrick Wells



Patricia Larsen



Mike Papka



Katrin Heitmann



Manpreet Dhillon

Thanks to:
LCFs & NERSC
Globus



Supported by
DOE HEP Computing



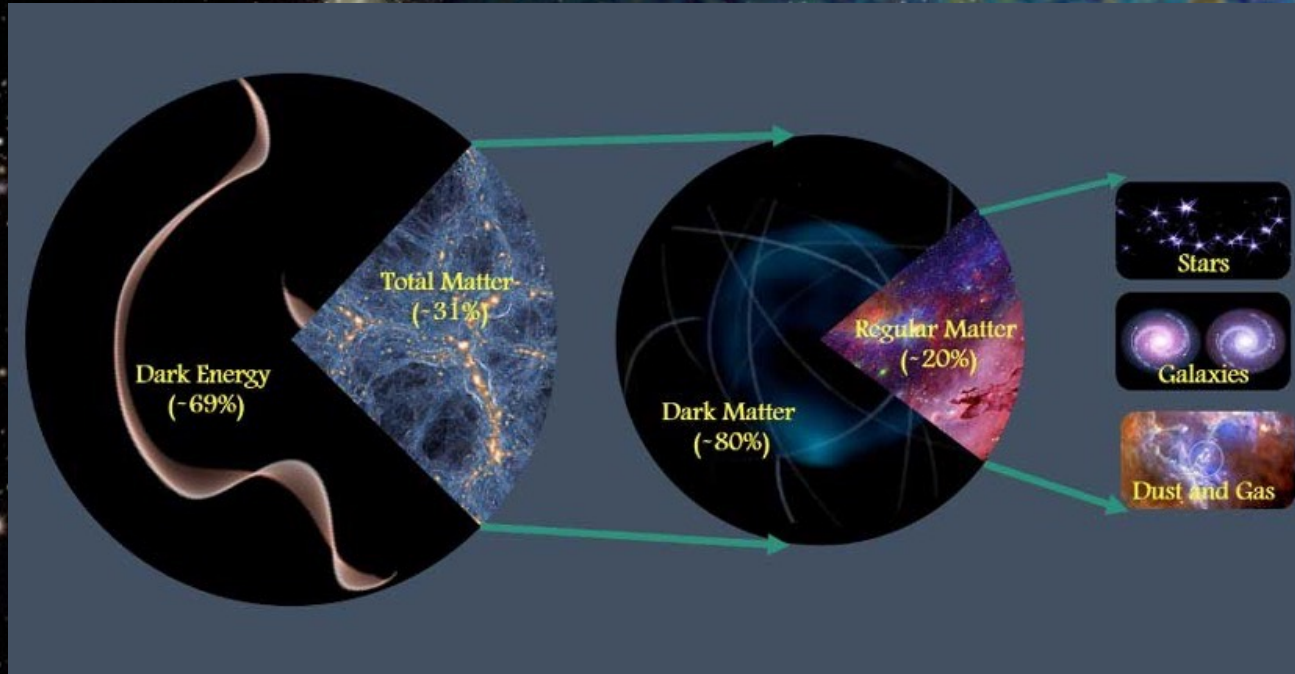
Simulating the Universe at Exascale

The Universe is Dark!

95% of the Universe is dark!

- 69% of the matter-energy content is a mysterious agent that drives the accelerated expansion of the Universe
- 26% is a dark matter component that does not emit or absorb light
- Only 5% of the Universe is made up of "regular matter"

Our quest: Understanding the secrets of the dark Universe!



Exploring the Physics of the Dark Universe

A wide-angle photograph of the Vera C. Rubin Observatory on Cerro Pachon, Chile, during twilight. The observatory's large, white, multi-faceted dome is the central focus, perched on a rocky ridge. In the background, another smaller observatory with a blue dome is visible on a higher peak. The sky transitions from a deep blue to a soft orange near the horizon, with a bright star or planet visible in the upper left. The foreground shows a dirt road and some construction equipment, indicating ongoing work at the site.

Vera C. Rubin Observatory on Cerro Pachon at nearly 9,000 feet in the foothills of the Chilean Andes

- The ten-year Rubin Observatory Legacy Survey of Space and Time (LSST) will observe half of the sky and record 37 billion stars and galaxies
- 10 million alerts, 20 TB of data, every night!
- The Rubin Observatory LSST 3200-megapixel camera weighs more than 3 tons and will produce images so large that 1500 high-definition TV screens would be required to view each of them
- First Look Event: June 23 2025; Rubin Observatory started operations in October 2025



Two Flagship Simulations at Exascale

Simulating the Universe on the world's fastest supercomputers – producing PBs of data

Frontier-E

SC25 GORDON BELL FINALIST

OLCF
Frontier

Largest cosmological hydrodynamics simulation

~4T
particles

4.7 Gpc
box size

~100PB
raw data written
during run

O(PB)
processed data

CRK-HACC: gravity + smoothed-particle hydrodynamics

Physics: Gravity, gas dynamics, radiative cooling, star formation, supernova feedback, AGN feedback, chemical enrichment

Hardware: 9000 nodes, AMD Instinct MI250X GPUs

Impact: 15x larger than previous hydro simulations. First to match survey-scale volumes with full baryonic physics, enabling direct comparison with Rubin/LSST, DESI, Euclid observations

Frontiere, et al. SC '25 (arXiv:2510.03557)

SPHEREx Run

ALCF
Aurora

Record-breaking gravity-only cosmological simulation

~13.4T
particles

7.8 Gpc
box size

~40PB
raw data written
during run

O(PB)
processed data

HACC: pure N-body gravitational solver

Physics: gravity only to maximize volume and resolution for large-scale structure statistics

Hardware: 8100 nodes, Intel Data Center GPU Max (Ponte Vecchio)

Impact: Will support cosmological interpretation of SPHEREx full-sky observations. Unprecedented volume for precision clustering statistics at scales matching next-generation galaxy surveys

recent run; publications forthcoming

ALCF Webinar (Developer Sessions): Nick Frontiere & Esteban Rangel ([Link](#))

64 h^{-1} Mpc



<https://youtube.com/watch?v=RPqcANuUvNs>

$z = 0.34$

$M_{200} = 1.1 \times 10^{15} h^{-1} M_{\odot}$

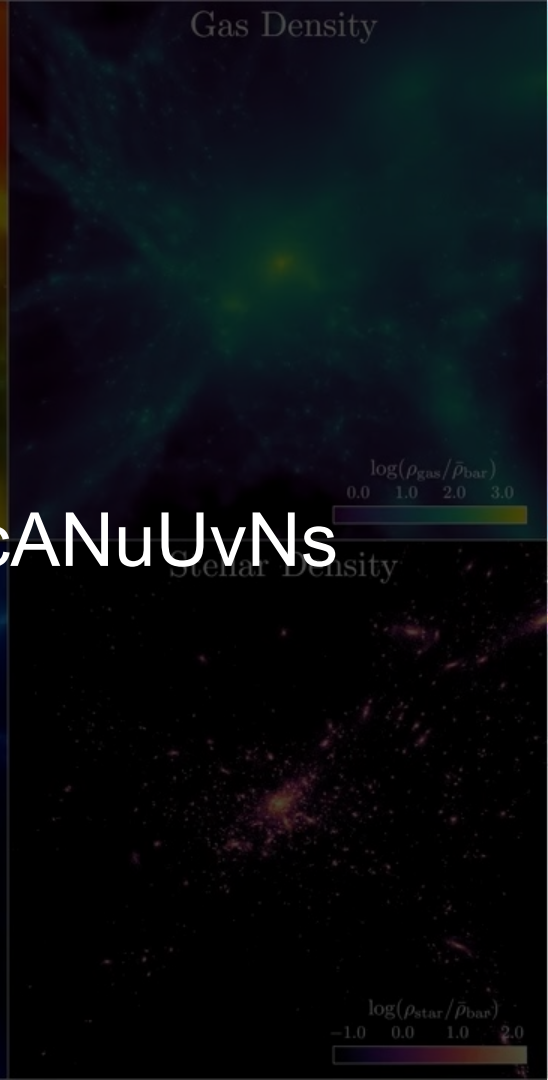


Gas Temperature

$\log(T_{\text{gas}}/K)$
4.0 5.0 6.0 7.0 8.0

Dark Matter Density

$\log(\rho_{\text{cdm}}/\bar{\rho}_{\text{cdm}})$
0.0 1.0 2.0 3.0



Gas Density

$\log(\rho_{\text{gas}}/\bar{\rho}_{\text{bar}})$
0.0 1.0 2.0 3.0

Stellar Density

$\log(\rho_{\text{star}}/\bar{\rho}_{\text{bar}})$
-1.0 0.0 1.0 2.0

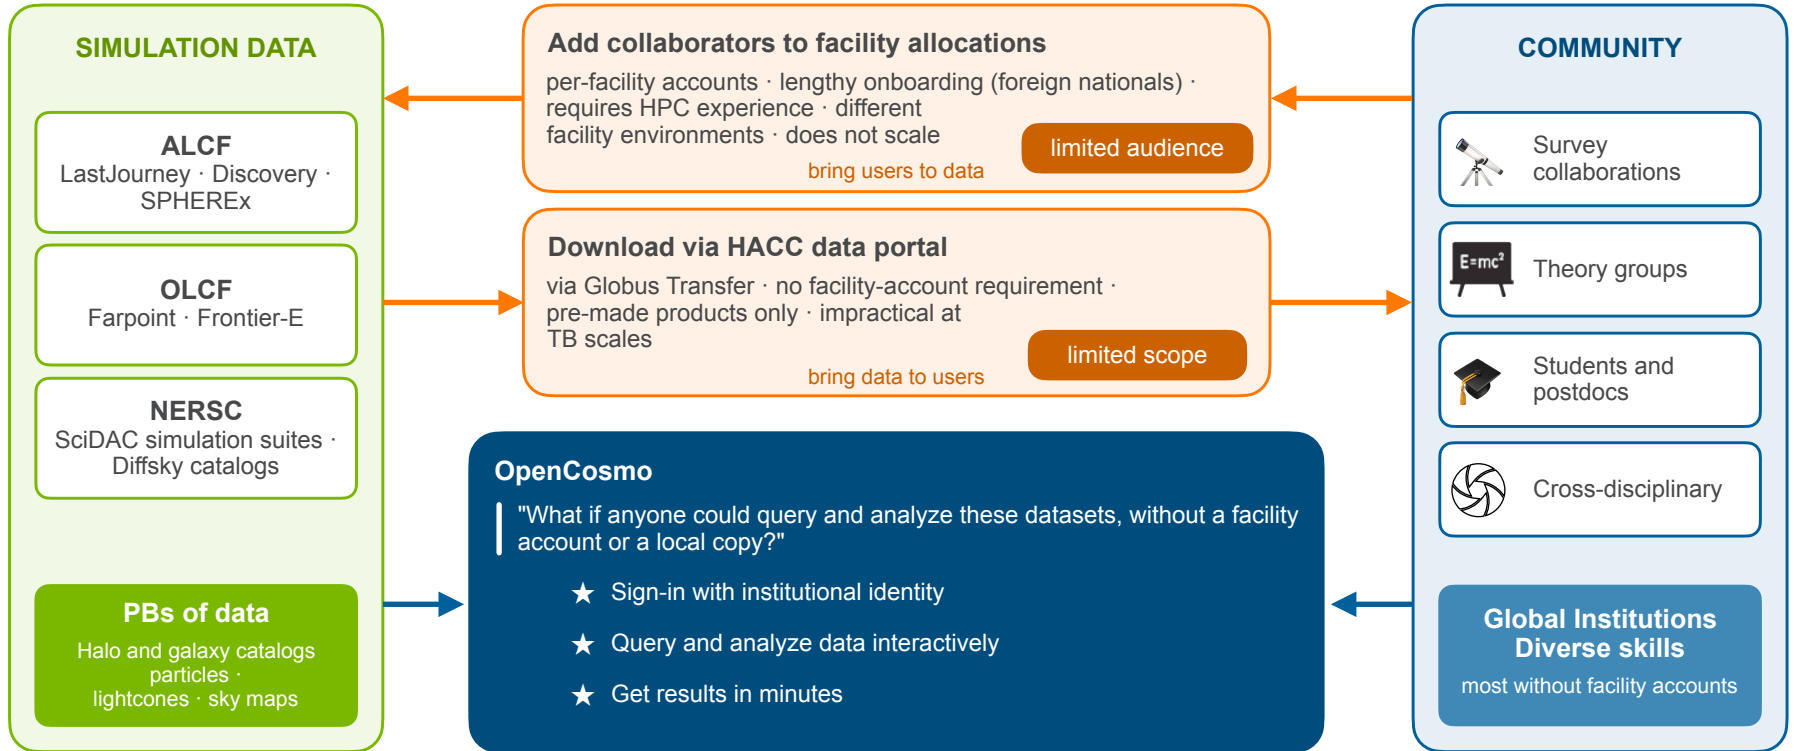
https://youtube.com/watch?v=JAyrpJCC_dw

$z = 0.000$

106 Mpc/h

The Data Sharing Challenge

PBs of simulation data across DOE facilities – a world-wide research community to collaborate



OpenCosmo

A platform for sharing and analyzing cosmological simulation datasets across DOE Leadership Computing Facilities

The core idea

Don't ask users to download petabytes or navigate HPC systems. Instead, let them **describe what they want**; the system figures out where the data lives, runs the analysis on facility resources, and returns the results.



Curated data & Toolkit

Simulation outputs converted to efficient and standard formats with a Python library for reading, filtering, and transforming the data



Interactive access

Web portal and API for querying data in place. Analysis runs on HPC, results returned to the user



AI-driven discovery

MCP servers let AI agents query simulation data with natural language → a new way to explore

Built on Globus services

Globus Auth: Federated identity; institutional login, no facility accounts

Globus Groups: fine-grained permissions for accessing tasks, facilities, tools

Globus Flows: workflow orchestration and routing

Globus Compute: remote execution

Globus Transfer: data movement from facilities to users

ALCF

Polaris ·
Eagle

NERSC

Perlmutter ·
CFS

OLCF

Odo ·
Wolf



Data Curation and Toolkit



Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

Foundation: Data Curation and Toolkit

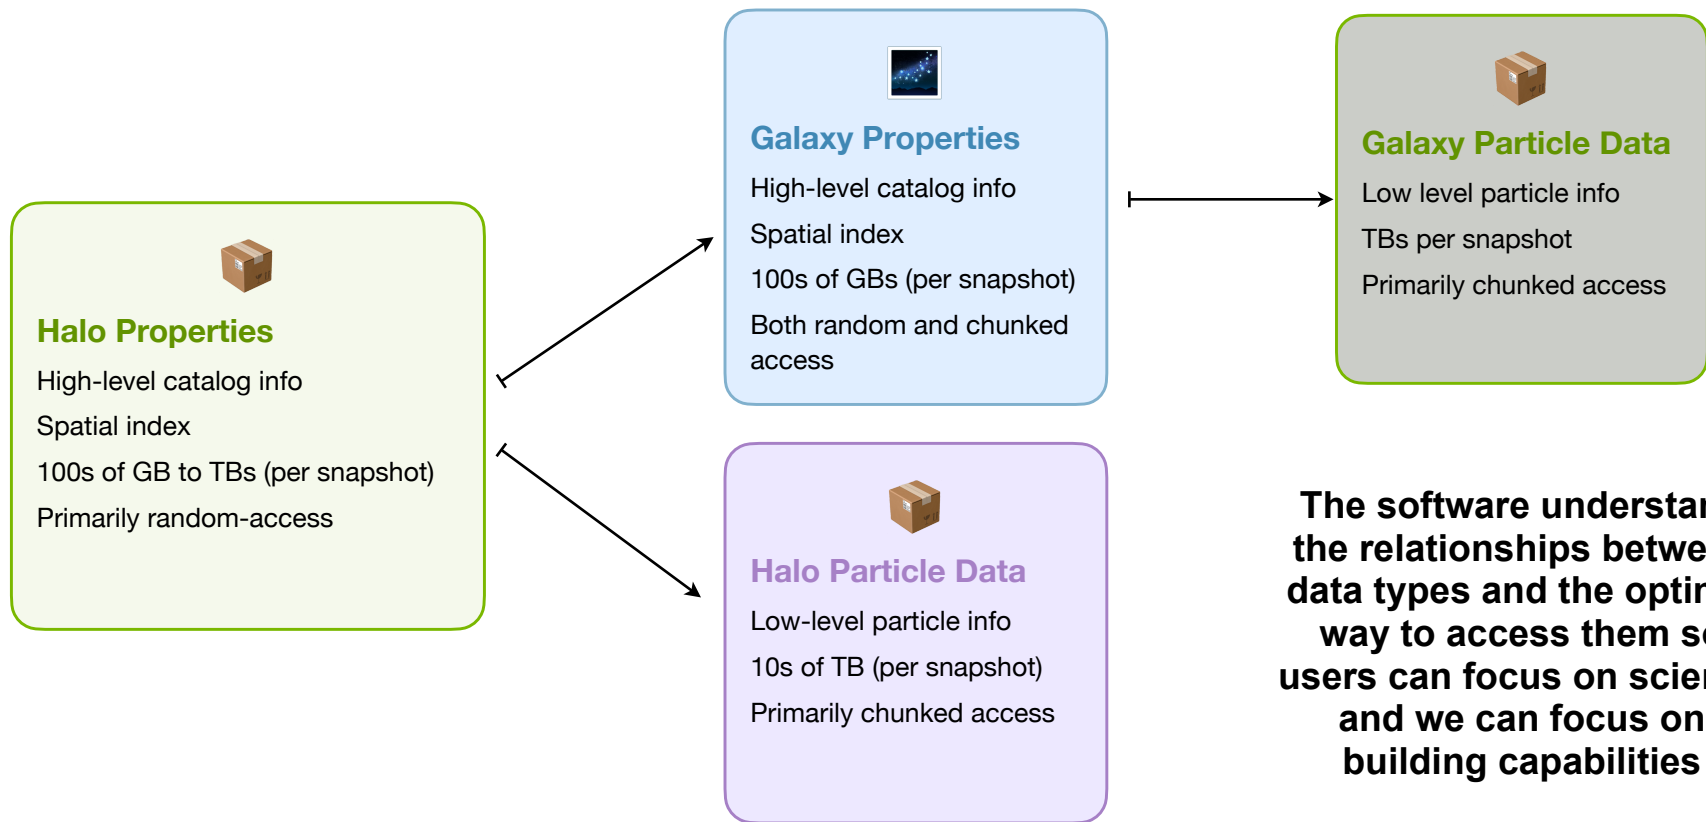
```
import opencosmo as oc

ds = oc.open("haloproperties.hdf5")
high_mass_filter = oc.col("fof_halo_mass") > 1e14
ds = ds.filter(high_mass_filter)
oc.write("high_mass_halos.hdf5", ds)
```

- Domain-specific analysis tool built top of a generic lazy query engine
- Common data format for both large and small-scale analysis
- Scales down to single-core, small-data work or up to multi-node, petabyte-scale analysis.

Result: Our backend talks to our data with the same software our users do

Foundation: Data Curation and Toolkit



The software understand the relationships between data types and the optimal way to access them so users can focus on science and we can focus on building capabilities



Interacting with datasets through OpenCosmo

Live Demo



Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.



Hydro Simulations

Galaxy Query

Halo-Particles Query (Hydro)


Halo Query (Hydro)

Gravity-Only Simulations

Halo-Particles Query (Gravity)

Halo Query (Gravity-Only)

Map Query (Gravity-Only)

 Help Contact Us

Welcome to the OpenCosmo Portal!

Select a task to get started, or check out [the documentation](#)

Galaxy Query

This flow retrieves and filters galaxies from HACC simulations

Halo-Particles Query (Gravity)

This flow retrieves halo properties and particles from HACC gravity-only simulations

Halo-Particles Query (Hydro)

This flow queries the particles of a given halo (run the "Halo Query" first to get a halo ID!). Particles and surface density maps, divided by particle species, are returned in an HDF5 file.

Halo Query (Gravity-Only)

This flow retrieves and filters halos from HACC gravity-only simulations

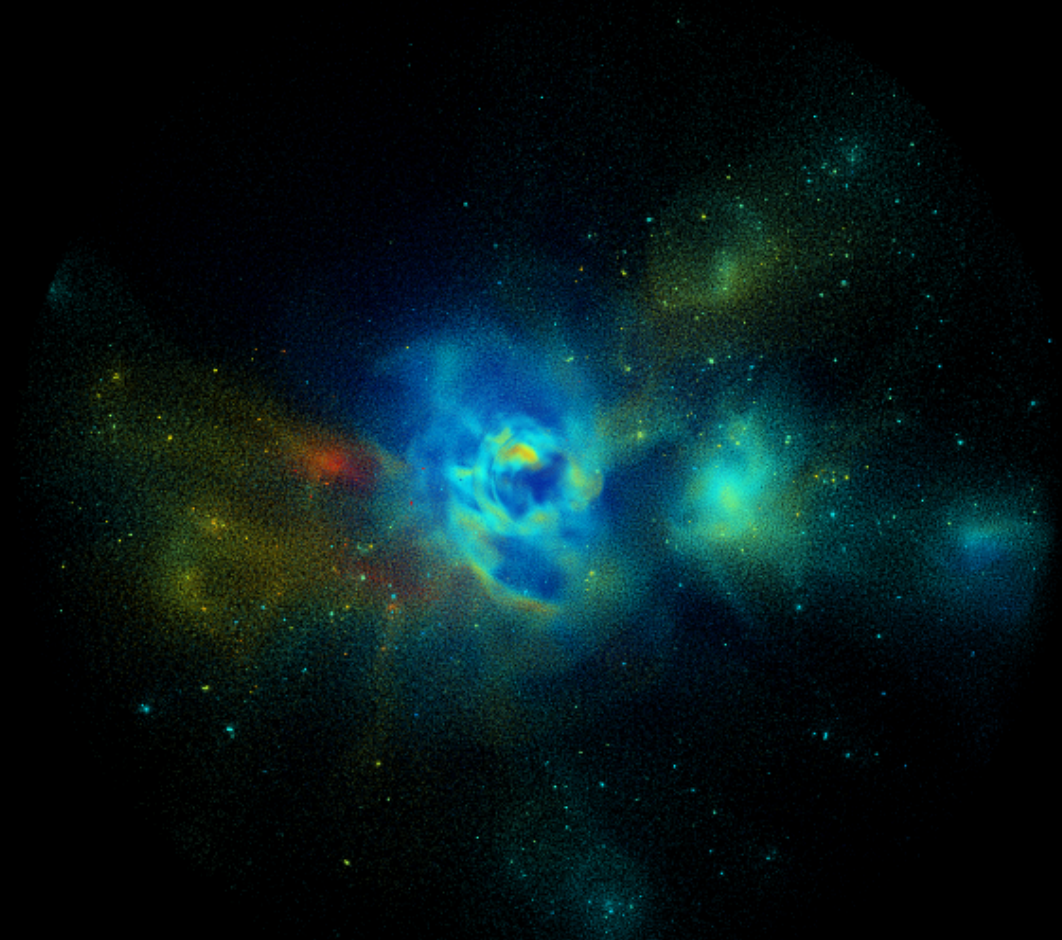
Halo Query (Hydro)

This flow retrieves and filters halos from HACC hydro/subgrid simulations.

Map Query (Gravity-Only)

This flow retrieves and filters lightcone maps from HACC gravity-only simulations

Recent Runs



uet, galaxy_particles.parquet, gas_particles.parquet, moon.parquet,

Camera
Presets:

Top

Front

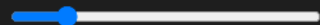
Side

Reset

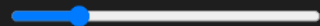
?

Visuals

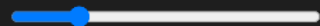
Point Size: 4



Opacity: 0.2



Scale: 10



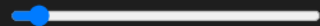
Show Colormap Bars

Show Axes Viewer

Show Vector Arrows

Hide Points with Arrows

Arrow Size: 0.0004



Data Filtering

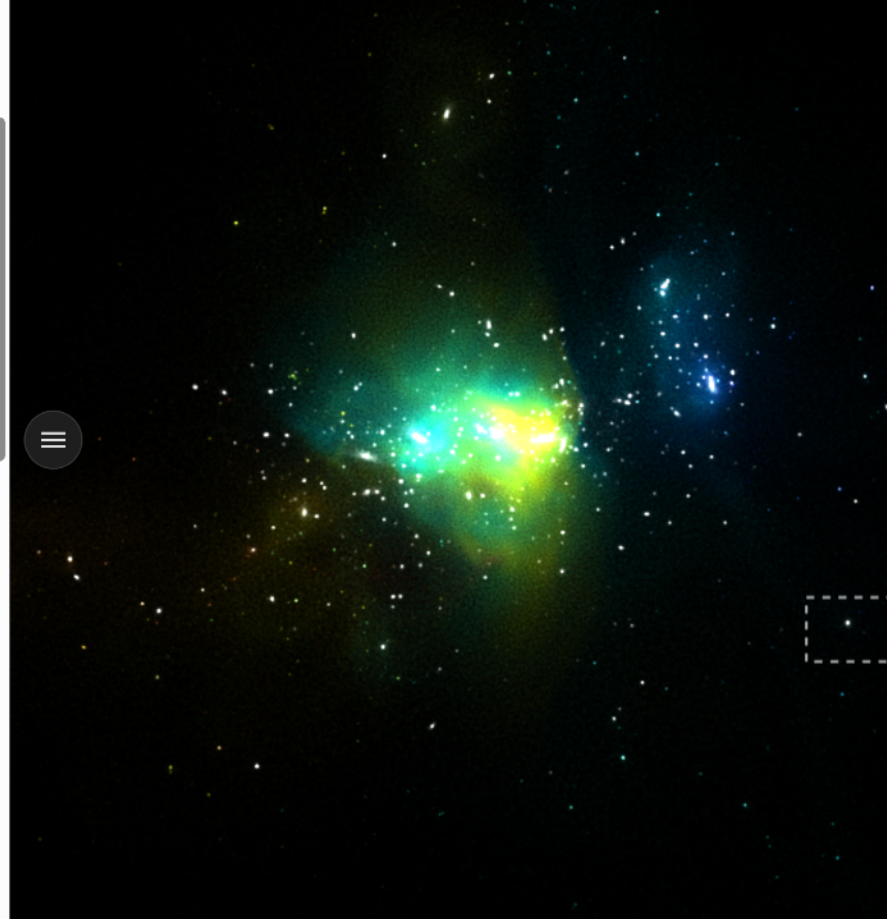
Limit Rendered Points:

e.g., 1,000,000 1,644,366

Apply Limit

Reset Limit

Filter by Scalar Range:



Data Legend

Scalar: vx

Min: -3286.25

Max: 4491.74

Total Rendered: 1,644,366



Gas (1,495,439)

Scalar: vx

Reset

Colormap: Jet

Reset

Size: 4

Reset

gas - vx

Selected Points Summary:

3,746

Type: gas (3,087 points)

Scalar Statistics:

mass:

Min: 3.1389e+8, Max: 1.2391e+9

Avg: 3.2817e+8, StdDev: 3.9815e+7

Median: 3.1389e+8

temperature:

Min: 1.7294e+3, Max: 1.3181e+8

Avg: 1.7060e+7, StdDev: 2.1733e+7

Median: 6.0119e+6

MCP: Letting AI access the tools

Example Conversation

What is MCP?

Model Context Protocol: Open standard that lets AI agents use external tools.

- AI agent can browse available tools and datasets
- Users can describe their queries in natural language
- Same queries and analysis as the web portal
- Agent translates text to tool call, can ask clarifying questions

Why this matters

- Explore tools and data conversationally; AI can help finding the optimal query and dataset given the user's question
- Towards automated science:
 - Ask scientific question
 - AI decomposes question into steps, uses tools as appropriate

User

I want to study radial profiles of cool-core galaxy clusters, what tool and dataset should I use?

AI

For cool-core galaxy clusters, I recommend `halo_catalog_query_hydro`. Cool-core clusters require hydrodynamic simulations to capture gas physics and cooling. This tool retrieves density profiles via `include_profiles`, and you can filter for cluster-mass halos ($M > 10^{14} M_{\odot}/h$).

Available hydro datasets: `SCIDAC_SUITE` and `Frontier-E` (single large volume simulation).

Would you like me to start a query?

User

Use the Frontier-E simulation, redshift 0.5, $M > 1e14 M_{\text{sun}}/h$

AI

Flow submitted! Run ID: 5b89fb6d-28ca-449a-9dfc-73488e263fa3

Query: Frontier-E, $z = 0.5$, $M200c > 10^{14} M_{\odot}/h$, profiles included, 1000 halos.

You can check status with `get_flow_status`.



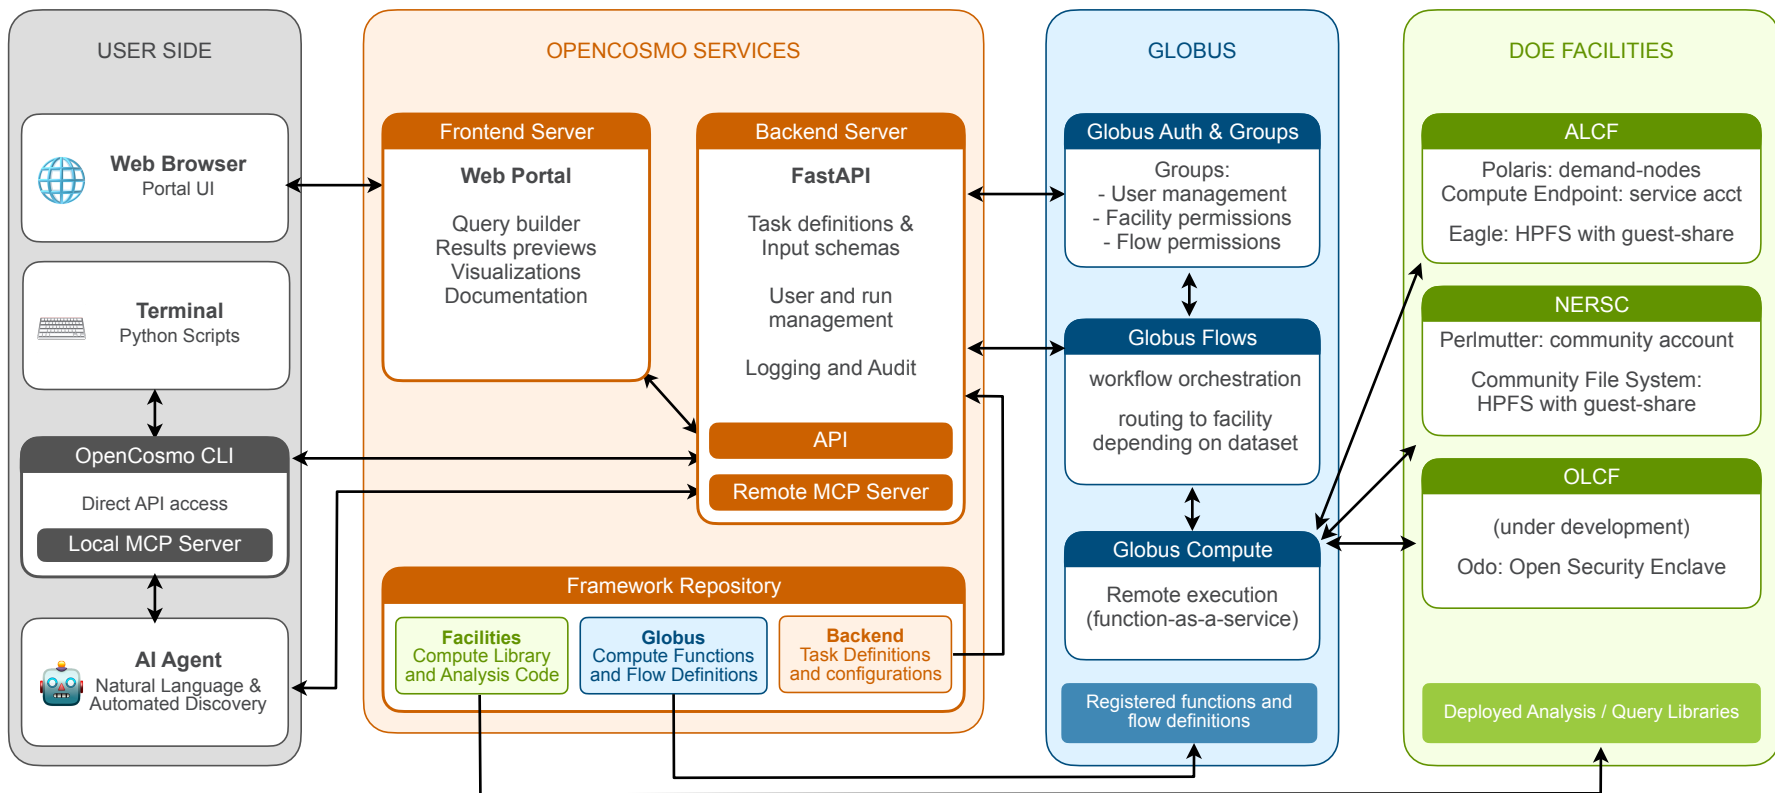
Under the Hood: Architecture and Globus Services



Argonne National Laboratory is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC.



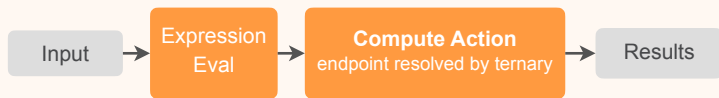
Framework Overview



Globus Flows: Routing

Challenge: Datasets are distributed among facilities – flow must pick the right compute endpoint at runtime

Expression Eval



```
// In the flow definition:
"SelectEndpoint": {
  "Type": "ExpressionEval",
  "Parameters": {
    "endpoint_id.": "'uuid-for-sim1' if $.input.dataset ==
      'sim1' else ('uuid-for-sim2' if $.input.dataset ==
      'sim2' else ..."
  },
  "ResultPath": "$.computed",
  "Next": "Compute"
},
"Compute": {
  "Type": "Action",
  "endpoint.$": "$.computed.endpoint_id",
  ...
}
```

- Difficult to maintain and scale
- Limited on what you can do

OUR APPROACH

Routing Function



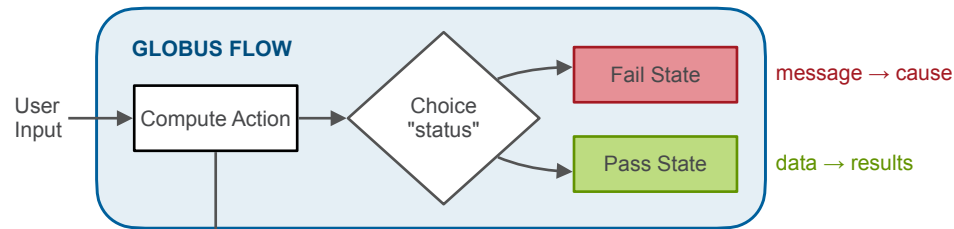
```
// Action 1: routing function
def route(task, params, user):
  validate(task, params)
  endpoint = lookup_endpoint(params["dataset"])
  function = lookup_function(task)
  log_request(user, task, params)
  return {"endpoint": endpoint, "function": function}

// Action 2: dynamic endpoint reference
"Compute": {
  "Type": "Action",
  "endpoint.$": "$.R.details.result[0].endpoint",
  "function.$": "$.R.details.result[0].function",
  ...
}
```

- Runs on lightweight VM
- Additional validation and logging outside HPC
- Update compute functions without redeploying flow

Globus Flows: Error Handling

Challenge: Full tracebacks from uncaught exceptions in Compute – not user-friendly and hard to decipher



Compute Function catches exceptions, always returns a dict

```
return {
  "status": "failed",
  "message": "no objects found
for query",
  "data": None
}
```

```
return {
  "status": "success",
  "message": None,
  "data": {...}
}
```

Without Pattern

```
Compute - FlowFailed
1 {
2   "state_name": "Compute",
3   "exception": "ActionFailedException",
4   "cause": {
5     "action_id": "tg_2877ca16-b6a8-4cfd-8538-4a7e69e2cec3",
6     "status": "FAILED",
7     "details": {
8       "result": [
9         "Traceback (most recent call last):\n File
\\opt/compute_action_provider/compute_action_provider/provider.py", line 534, in
get_status\n task_result_raw = fxc.get_result(task_uid)\n
~~~~~\n File \\usr/local/lib/python3.12/site-
packages/globus_compute_sdk/sdk/client.py", line 399, in get_result\n task =
self.get_task(task_id)\n
~~~~~\n File
\\usr/local/lib/python3.12/site-packages/globus_compute_sdk/sdk/client.py", line 62,
in _wrapper\n return gare_handler(login, f, self, *args, **kwargs)\n
~~~~~\n File \\usr/local/lib/python3.12/site-
packages/globus_compute_sdk/sdk/utils/gare.py", line 13, in gare_handler\n return
f(*args, **kwargs)\n
~~~~~\n File
\\usr/local/lib/python3.12/site-packages/globus_compute_sdk/sdk/client.py", line 383,
in get_task\n return self._update_task_table(r.text, tid)\n
~~~~~\n File \\usr/local/lib/python3.12/site-
packages/globus_compute_sdk/sdk/client.py", line 354, in _update_task_table\n raise
TaskExecutionFailed(fr_dict["exception"]]
```

With Pattern

```
ComputeFail - FlowFailed
1 {
2   "state_name": "ComputeFail",
3   "cause": "No halos found matching tags [123451234] at step 415"
4 }
```

Facility Policies and Cyber Security

External users trigger (predefined) compute on facility – defense in depth at every layer

OpenCosmo Server

- ☑ Globus Group membership for authentication
- ☑ Parameter validation against task schemas
- ☑ Request logging and audit trail

Globus Flows

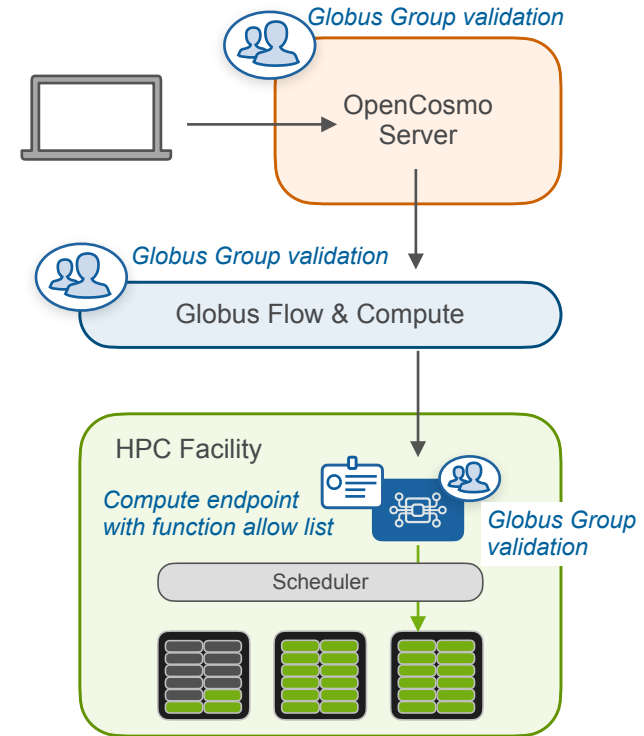
- ☑ Globus Group permissions on flows
- ☑ Flow input schema validation
- ☑ Globus internal logging and audit trail

Facility Compute Endpoint

- ☑ Globus Group permissions on Compute endpoint
- ☑ Compute endpoint with function allow list
- ☑ Parameter re-validation at execution time

Facility-specific access policies: mapped to Globus subgroups

ALCF:	OpenCosmoALCF	Collaborators with DOE lab credentials or ALCF accounts
NERSC:	OpenCosmoNERSC	Collaborators with NERSC accounts
OLCF:	OpenCosmoOLCF	Part of our OLCF allocation



IRI and the American Science Cloud

LCFs are good at

- Large-scale storage
- Large-scale compute

"Clouds" are good at

- Multi-facility orchestration
- Small-scale compute for persistent services

Building "Science as a Service" platforms requires an orchestration layer that looks more like a commercial software stack than a science workflow

Goal: LCFs and NERCS provide *storage* and *compute*, IRI and AmSC provide the *orchestration*

Success stories

Real Science Uses:

- Early dark energy and physics beyond the standard cosmological model (S. Hodgson)
- Formation of galaxies within large clusters (G. McFarland)
- Variation of galaxy and structure formation as a function of subgrid parameters in hydrodynamic simulations (E. Shinbaro)
- Relationships between observational data and synthetic galaxy catalogs (DESC collaboration)
- Concentration of mass in simulated halos (W. Hicks)
- Agentic discovery (J. Coburn, N. Ramachandra, A. Wells)
- And many more...



"The system is excellent! It would have saved me so much time during my last project to have the tools available."

F. Keruzore

Level 3: Collaboration Access

- Allow user to deploy their own analysis tools, will require accounts at LCFs/NERSC
- First tests/demos with LSST DESC members
- Collected feedback about further steps

Simulation data sharing exploration #6

New Issue



patricialarsen opened this issue 2 weeks ago · 0 comments



patricialarsen commented 2 weeks ago

Member ...

Simulation data sharing exploration

At Argonne we have a new project for simulation data sharing, we'd like to give you a preview of how it will work, what sort of data we'll have and how you can interact with it. We'd also love to get ideas on what would make this most useful for your science!

Contacts: Patricia Larsen [@patricialarsen](#) , Katrin Heitmann [@katrinheitmann](#)
(plarsen and heitmann on slack)

Time: Morning

Assignees

heather999

Labels

General Sprint **New Contributor Friendly**
Tutorial

Projects

None yet

The Vision (for users)

An integrated compute environment where users (human or otherwise) can run complex astrophysics analyses across one or more full-scale dataset without thinking about data management or even logging into a machine directly.

The Vision (for builders)

Common software and infrastructure patterns for building domain-specific, cross-facility analysis tools on top of domain-agnostic backends and orchestration layers

Ongoing Tasks

- Further development of analysis tools (*Based on user feedback!*)
- Integration of new simulations (*in-progress!*)
- Visualization capabilities (*integrated!*)
- Documentation, testing (*exists, but can be improved*)
- Workshops with test users (*several done*)
- First public release (*almost ready!*)
- User-specified analysis (*Early stages, but proofs of concept exist!*)
- Integration into AmSC and IRI backends (*Early days*)

Summary and Outlook

- Exascale supercomputers and ongoing and upcoming experiments generate A LOT of data
- Sharing data seamlessly and providing analysis tools to interact with the data in place will become ever more important
- We have built a robust framework that can be easily ported to other science areas as well
- Cyber security issues still need to be resolved to make the experience for the users fully satisfactory
- Another big thanks to DOE HEP, LCFs/NERSC and Globus for their support!

