JGI JOINT GENOME INSTITUTE
A DOE OFFICE OF SCIENCE USER FACILITY

JAWS - Scientific Workflow Execution Service

Seung-Jin Sul
ssul@lbl.gov
June 10, 2026

The DOE Joint Genome Institute (JGI)

JGI JOINT GENOME INSTITUTE

About us

- High-throughput AI-centric genomic science user facility located at Lawrence Berkeley National Laboratory
- Provides the genomic capabilities, data, and expertise that supports the global research community in studying complex biological and environmental systems

Advanced genomic capabilities

Large-scale sequencing and synthesis

AI-ready scientific data

Curated data products for reuse

Professional expertise

Support from domain and workflow experts

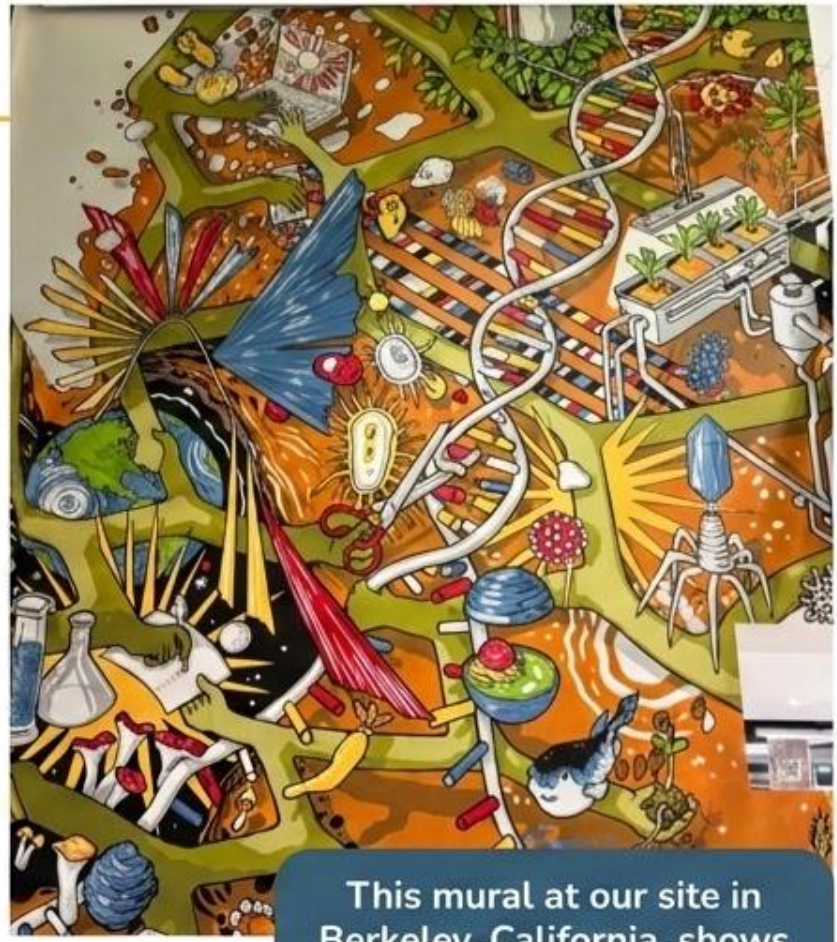

This mural at our site in Berkeley, California, shows some of the species that JGI users have studied.

Big Data, Bigger Science, and Workflows

JGI JOINT GENOME INSTITUTE

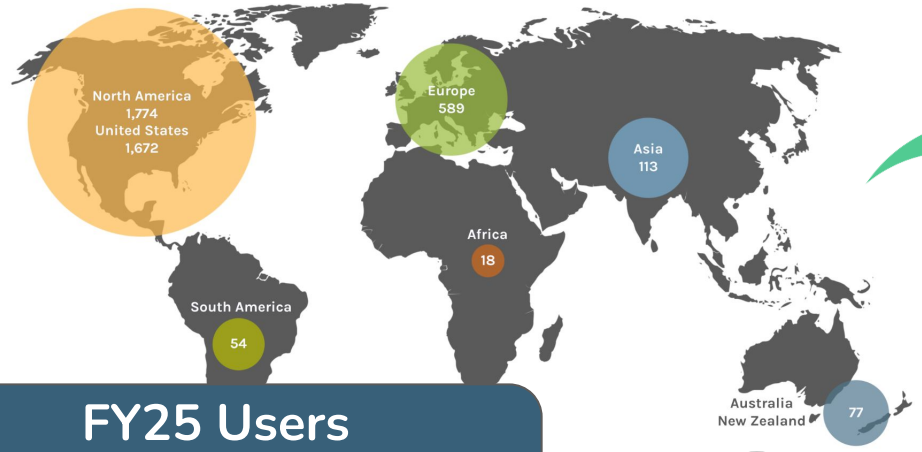

FY25 Users
Primary: 2,627
Secondary: > 18,000

Primary Users provide unique samples, from fungi, plants and microbiomes as part of their studies

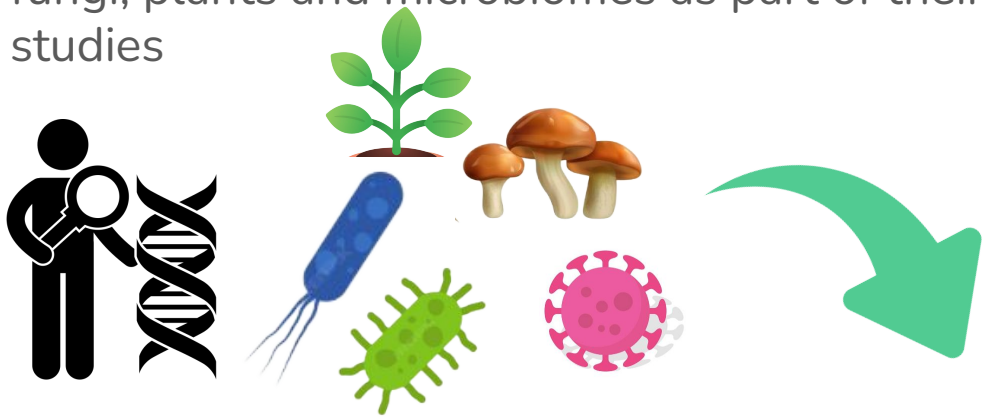

FY25
- Generated **1.14 Petabases** of sequence
- Synthesized **8.53 Megabases** of DNA
- Processed **17,179 metabolomics** samples

The sample becomes data

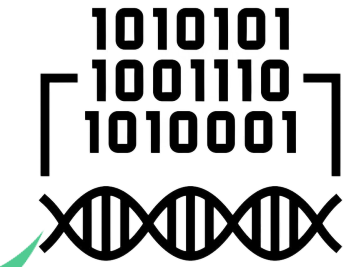

~22M Files

~20 PB Data

- JGI Data Portal
- JGI MycoCosm — THE FUNGAL GENOMICS RESOURCE
- JGI PhycoCosm
- JGI Phytozome 14 — THE PLANT GENOMICS RESOURCE
- JGI IMG/M — INTEGRATED MICROBIAL GENOMES & MICROBIOMES
- JGI SMC — SECONDARY METABOLISM COLLABORATORY

Primary and Secondary Users leverage data through JGI Portals

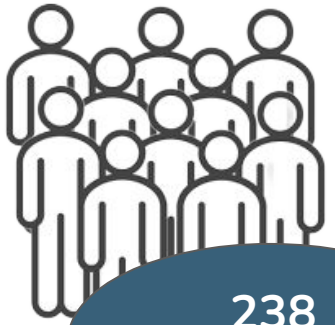

238 publications

Fragmented Science: Porting is Hard, Availability is Necessary

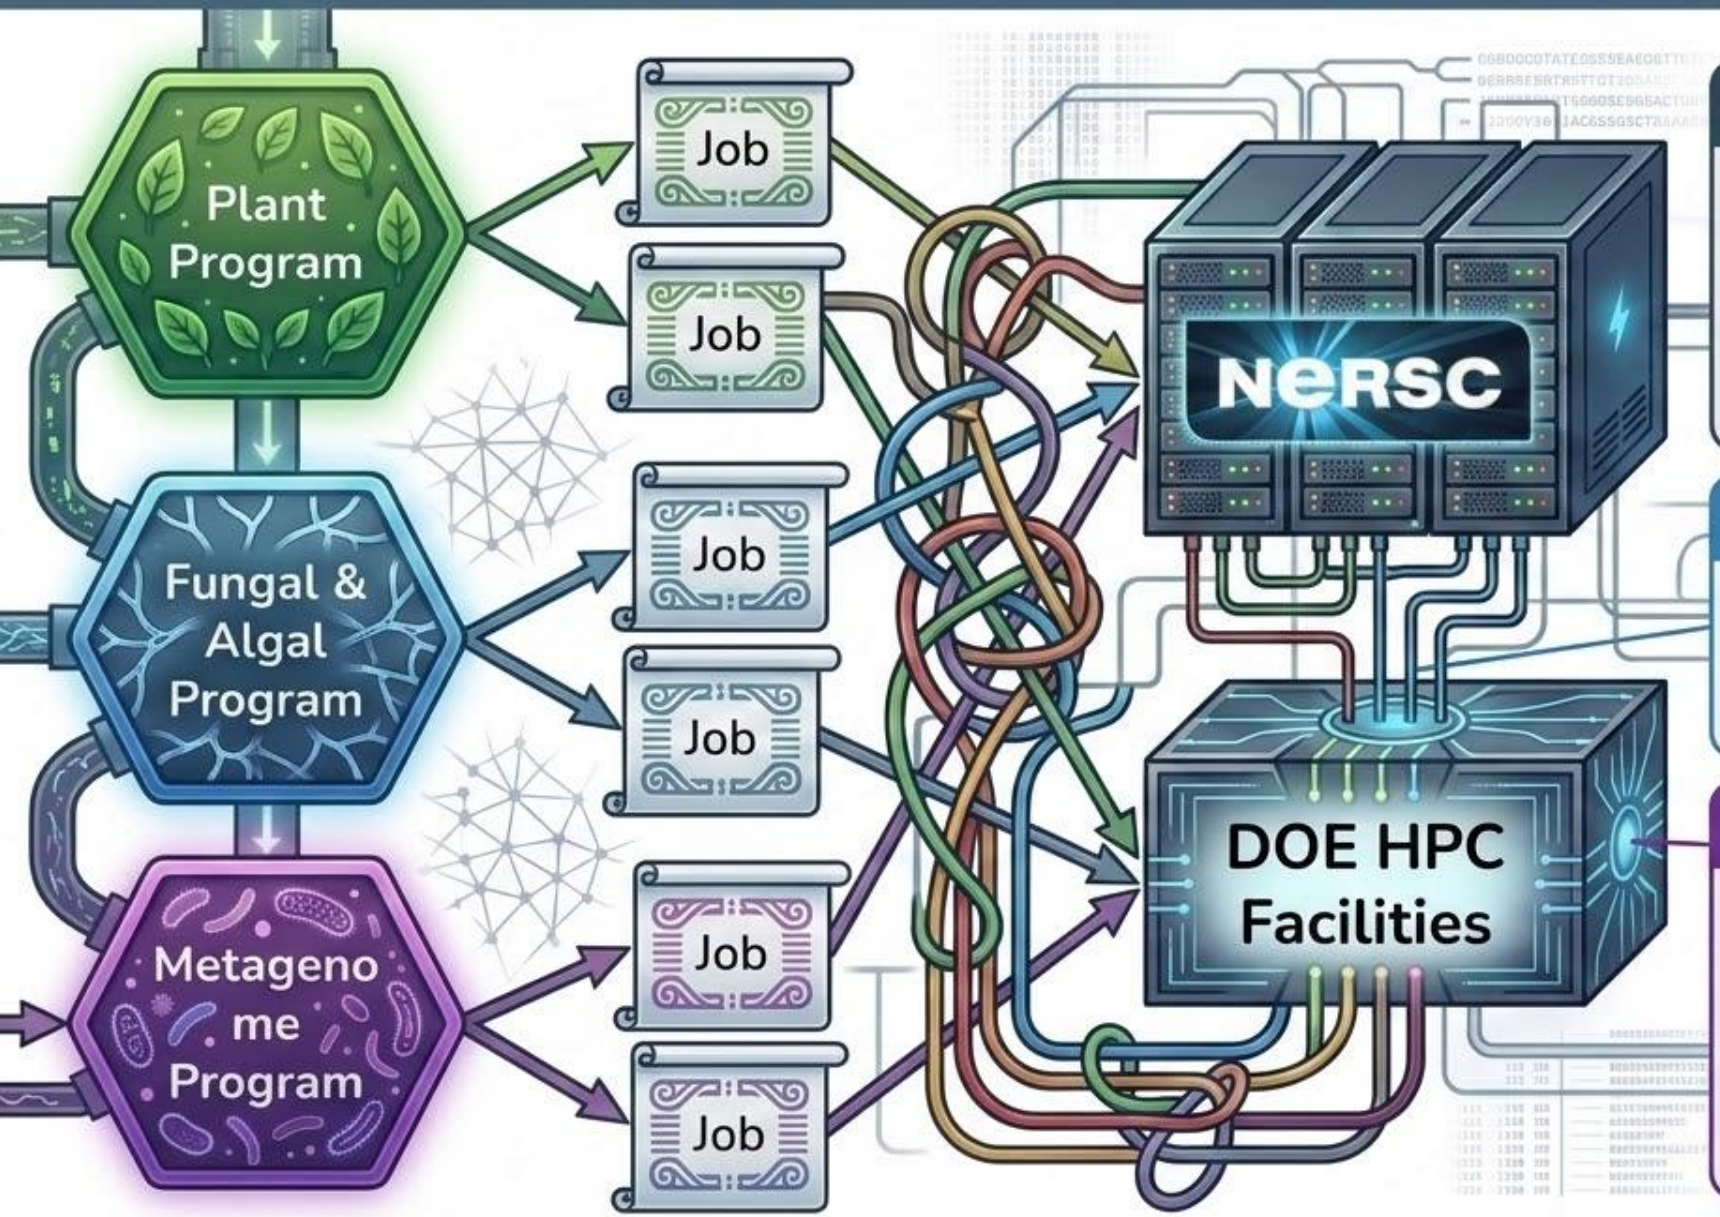

Development & Maintenance Issues

- Duplicate effort
- Many languages/glue codes
- Low maintainability
- Susceptibilty to downtime

Data Movement

Data movement is managed by each team

Portability & Environment

- Low portability
- Different HPC environments
  - Batch systems
  - Container systems
  - Account & security

* NERSC: DOE HPC center at LBNL

User provides the Science; We handle the Execution

- Developed a workflow management service called JGI Analysis Workflow Service (**JAWS**)
- Unifies workflows across JGI groups and DOE facilities
- Users write their workflows in Workflow Description Language (**WDL**)
- **JAWS** takes care of moving data, handling compute backends of multiple facilities, and executing workflows under the hood
- Write once, run anywhere

{wdl}

JAWS
JGI Analysis
Workflow Service

JAWS Components

User writes WDL; JAWS handles the rest

JAWS Client

- CLI / API / REST
- WDL workflow submit, status, cancel, logs, ...

JAWS Central & Sites

- File Transfer: Globus
- Protocols: HTTP and AMQP

Cromwell

- WDL execution engine
- Creates task DAG and HTCondor submit files
- Submits jobs to HTCondor

* DAG: directed acyclic graph

HTCondor

- Job scheduling and execution system between Cromwell and HPC clusters

JAWS Pool Manager

- HTCondor glidein pool manager
- Dynamically scales site compute pools
- Communicates with SLURM, PBS (AWS, k8s)

WDL Example: Hello!

JGI JOINT GENOME INSTITUTE

```
version 1.0

workflow hello {
  call say_hello { }
  output {
    File greeting = say_hello.out
  }
}

task say_hello {
  input {
    String message= "Hello!"
  }
  command <<<
    echo ~{message} > final_message.txt
  >>>
  output {
    File out = "final_message.txt"
  }
  runtime {
    docker: "ubuntu:22.04"
    cpu: 1
    memory: "2G"
    runtime_minutes: 3
  }
}
```

Workflow Description Language "Recipe"

- *func* Workflow Call
- Define task
- Define inputs type
- Series of commands
- Save results → outputs
- Compute resources requirement

JAWS: You Bring the Science, We Manage the Complexity

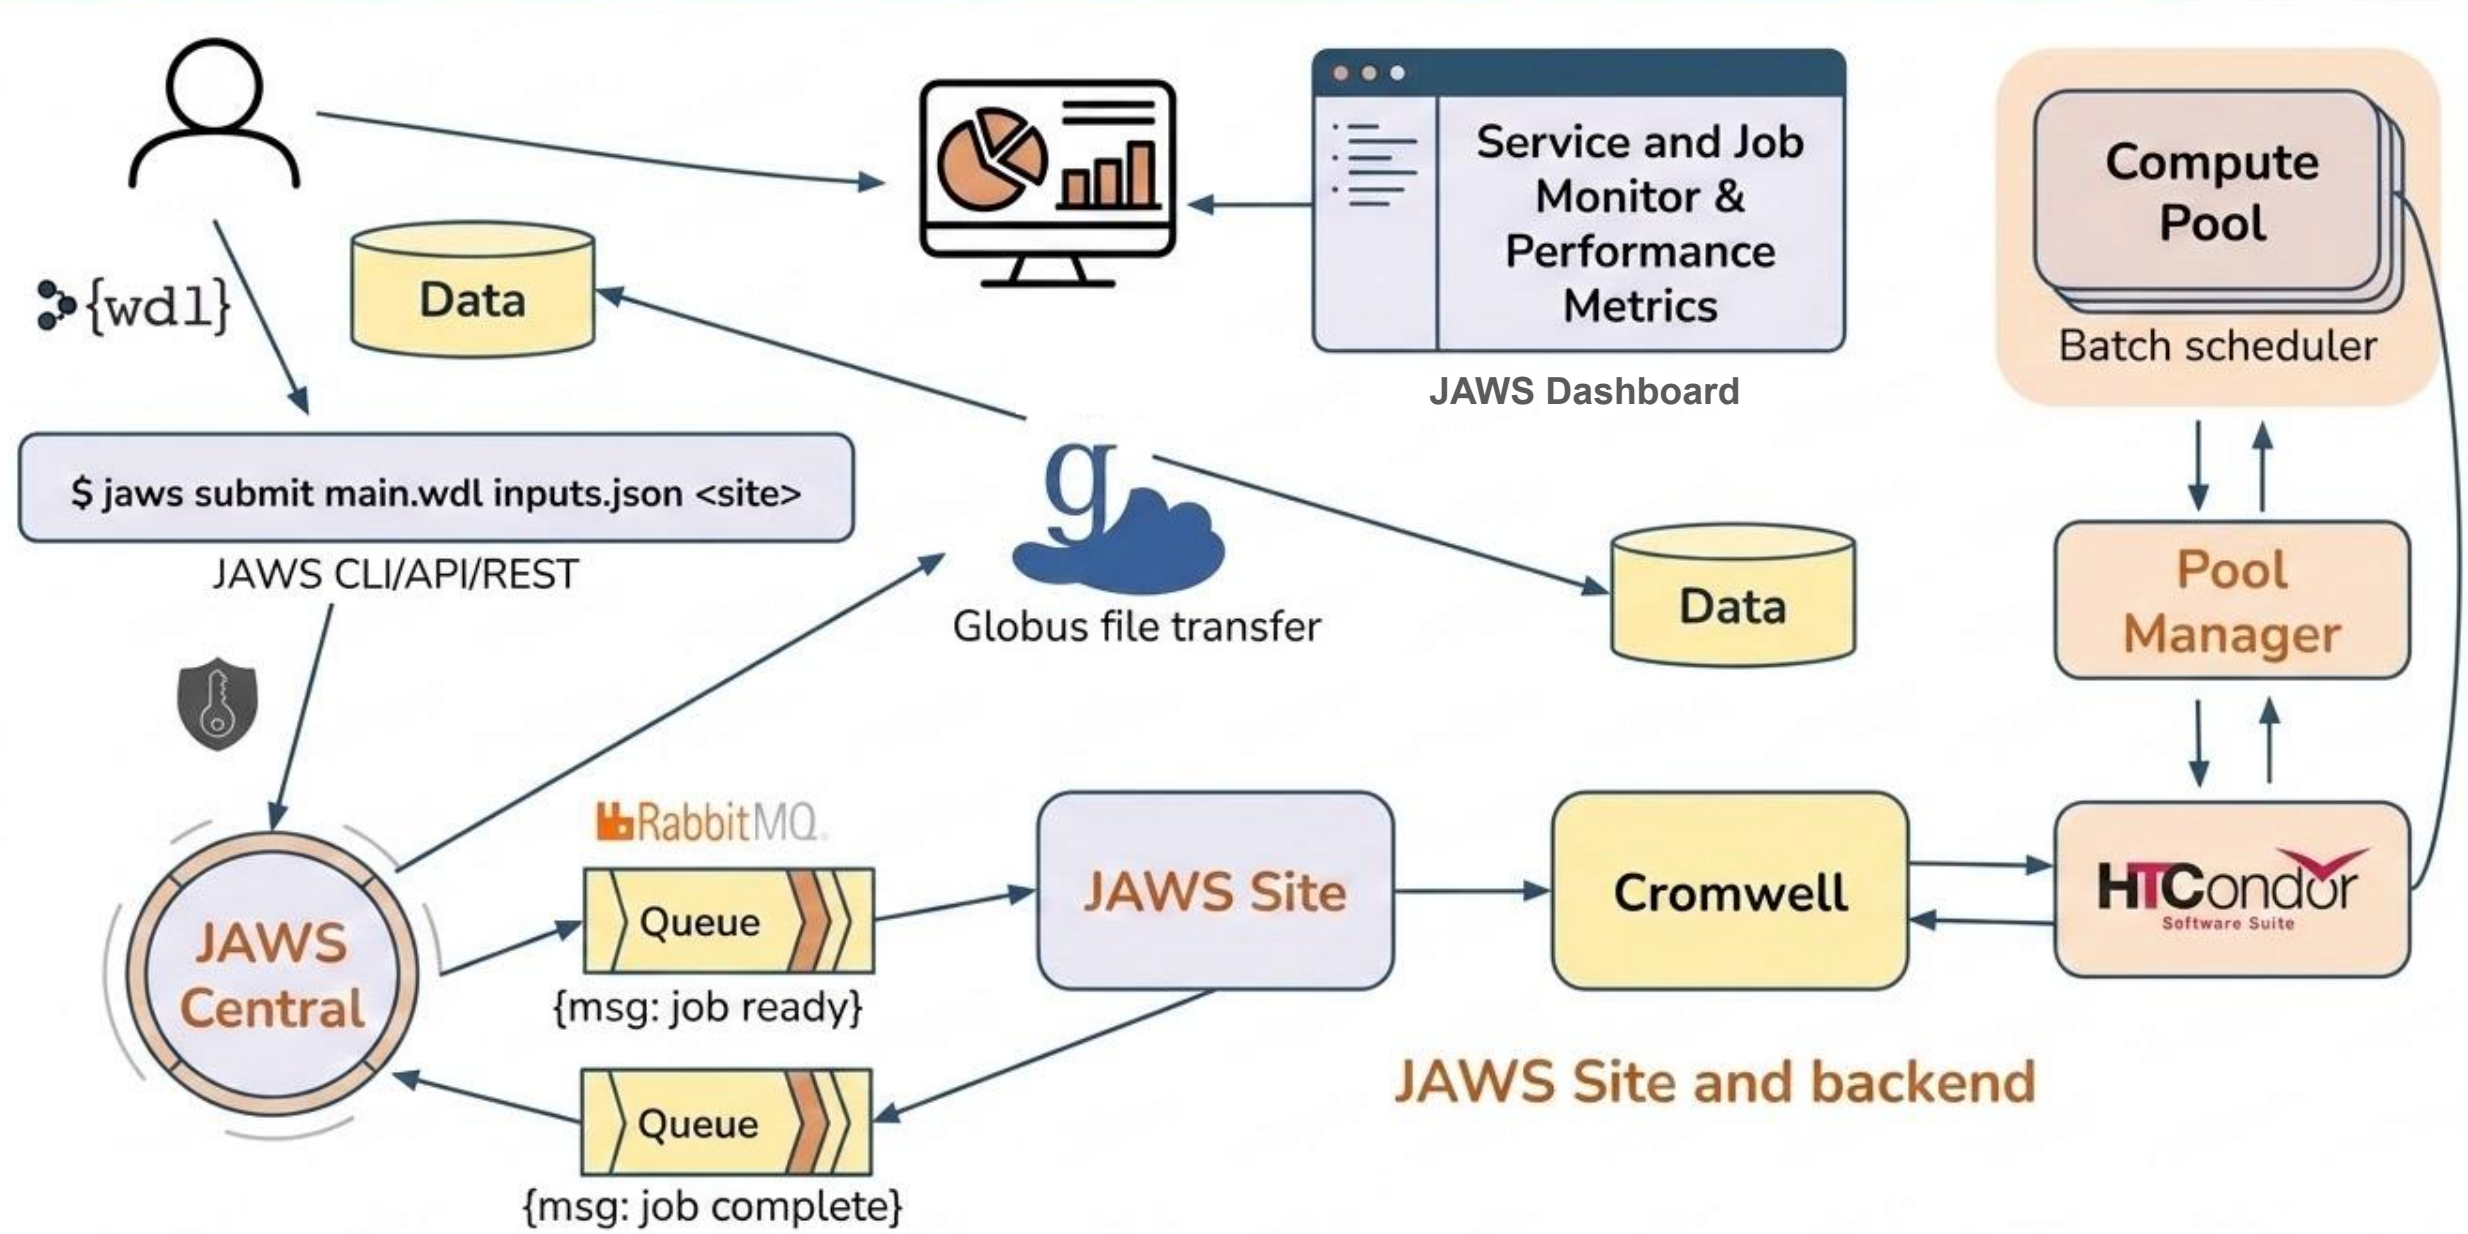

JAWS Pool Manager and HTCondor Glidein

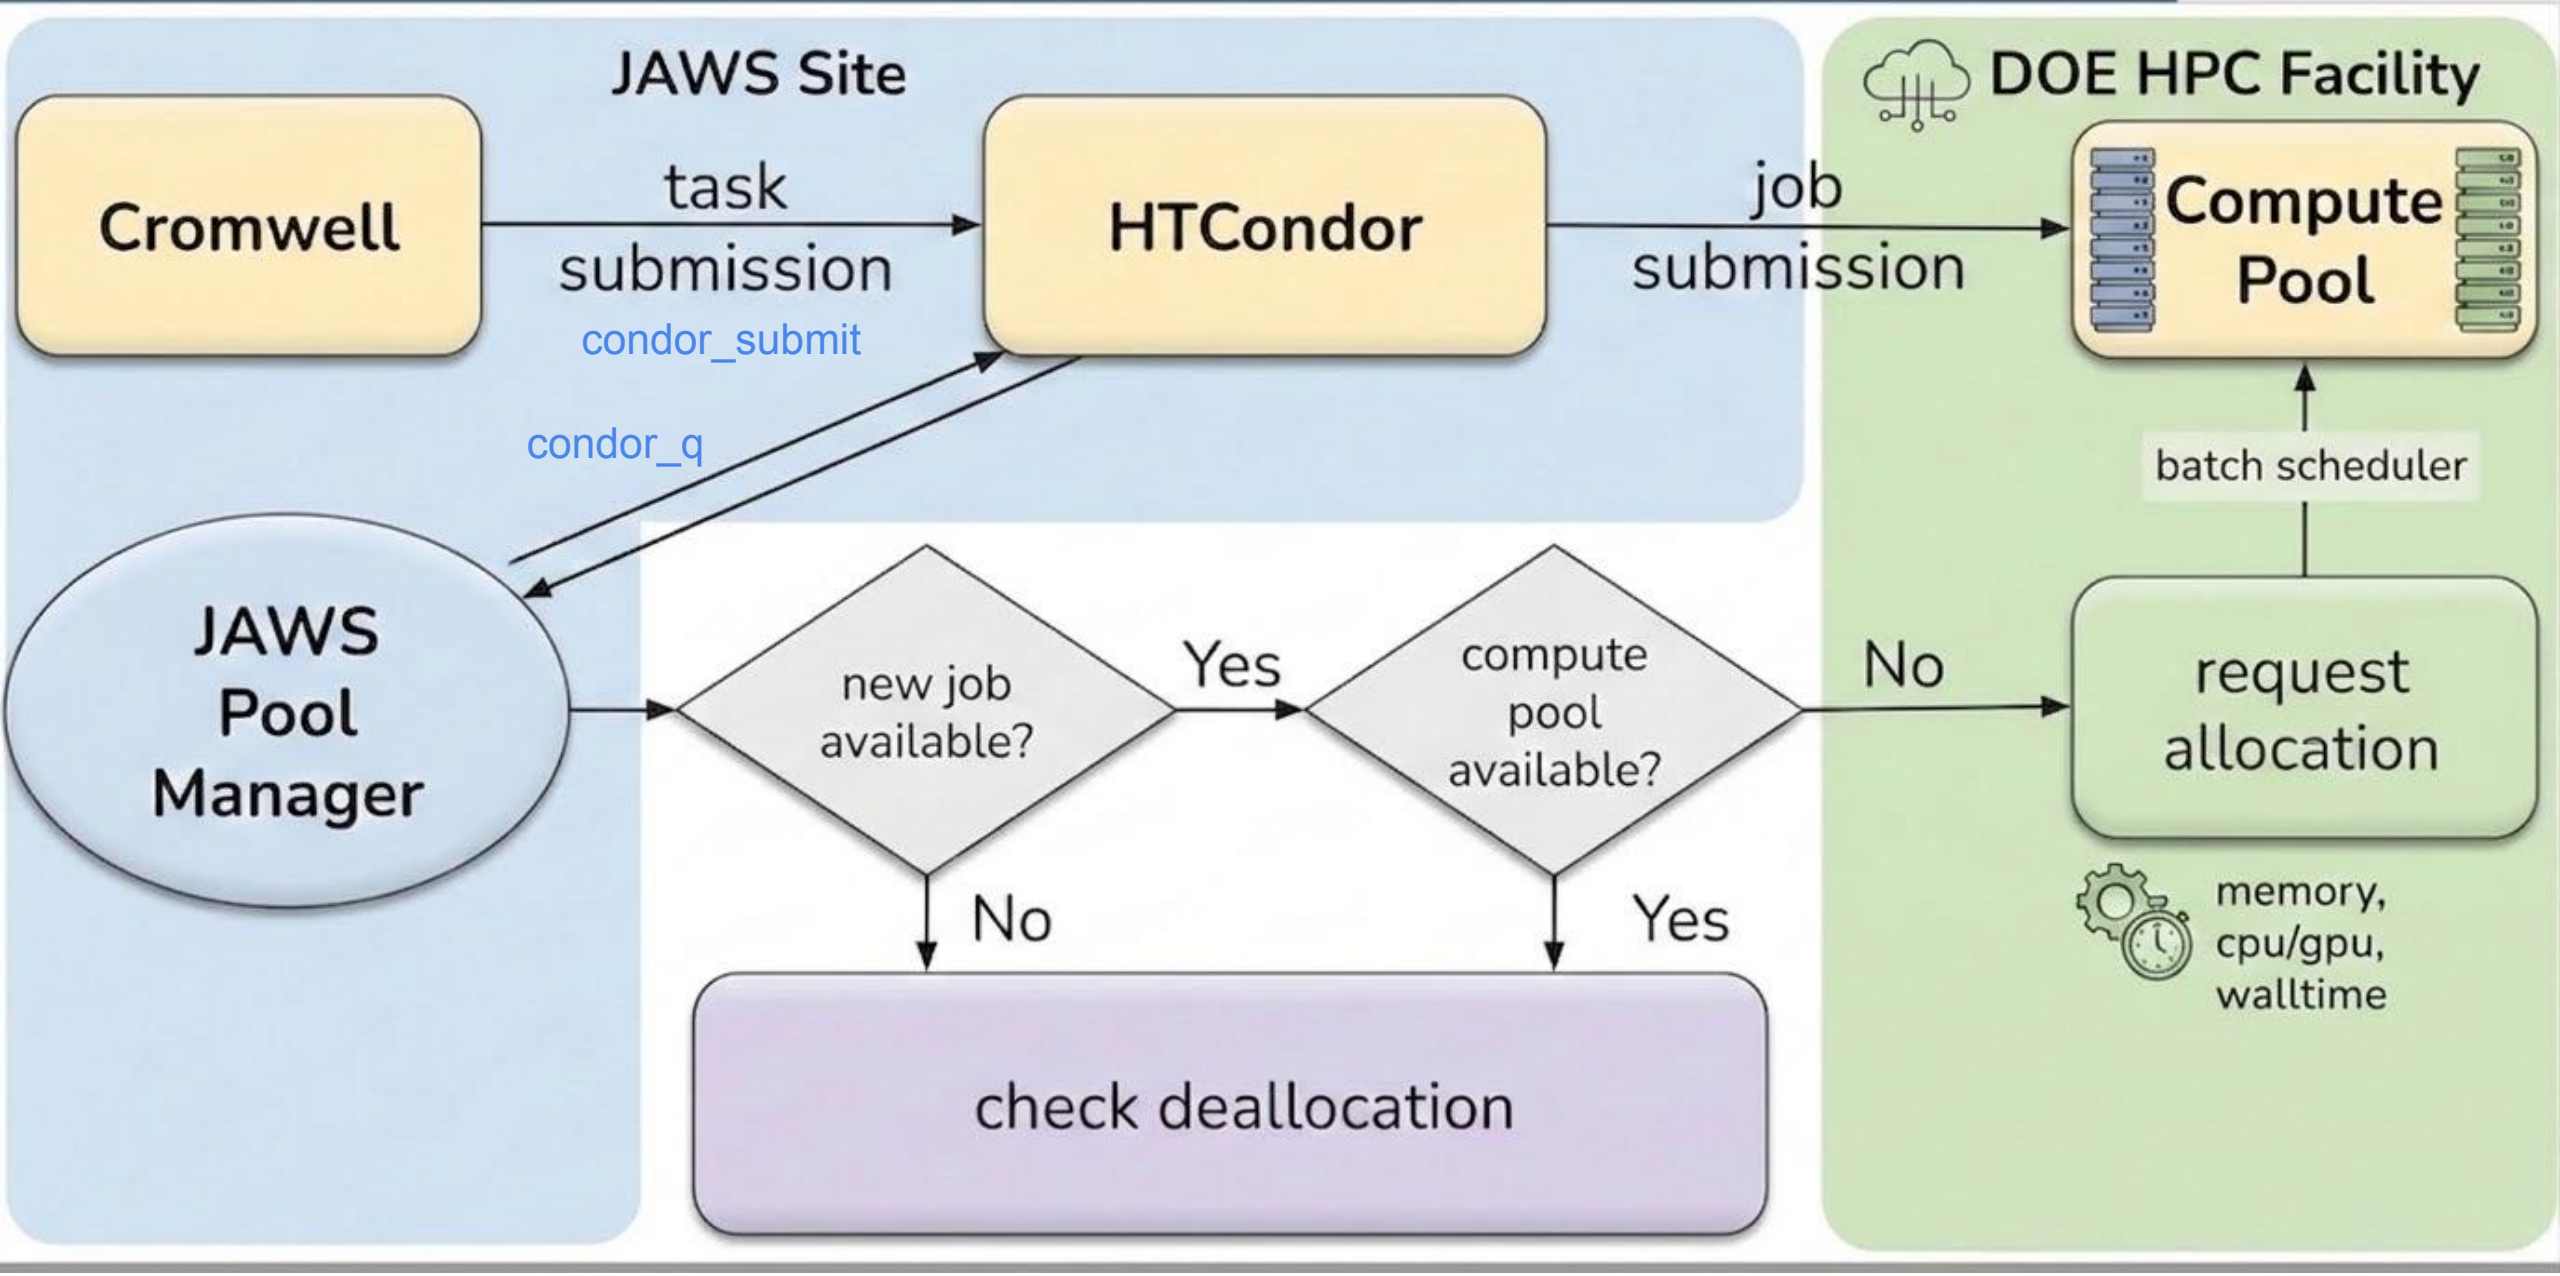

HTCondor Worker Pools

- **Pool types**
  - Five pool types based on the cpu/gpu & memory size
  - Each site has different set of node types
    - ex) Perlmutter: gpu and large (~500GB)

      Dori: small (debug), large and xlarge (~1.5TB)
  - Min/max pool sizes per pool type and site
    - ex) dori: 10-node large pool

- **Scheduling Policies**
  - Walltime limit for job scheduling
  - Policies for housekeeping
  - Slot defragmentation

JAWS Dashboard

Operational visibility helps users, support teams, and workflow developers keep large runs moving

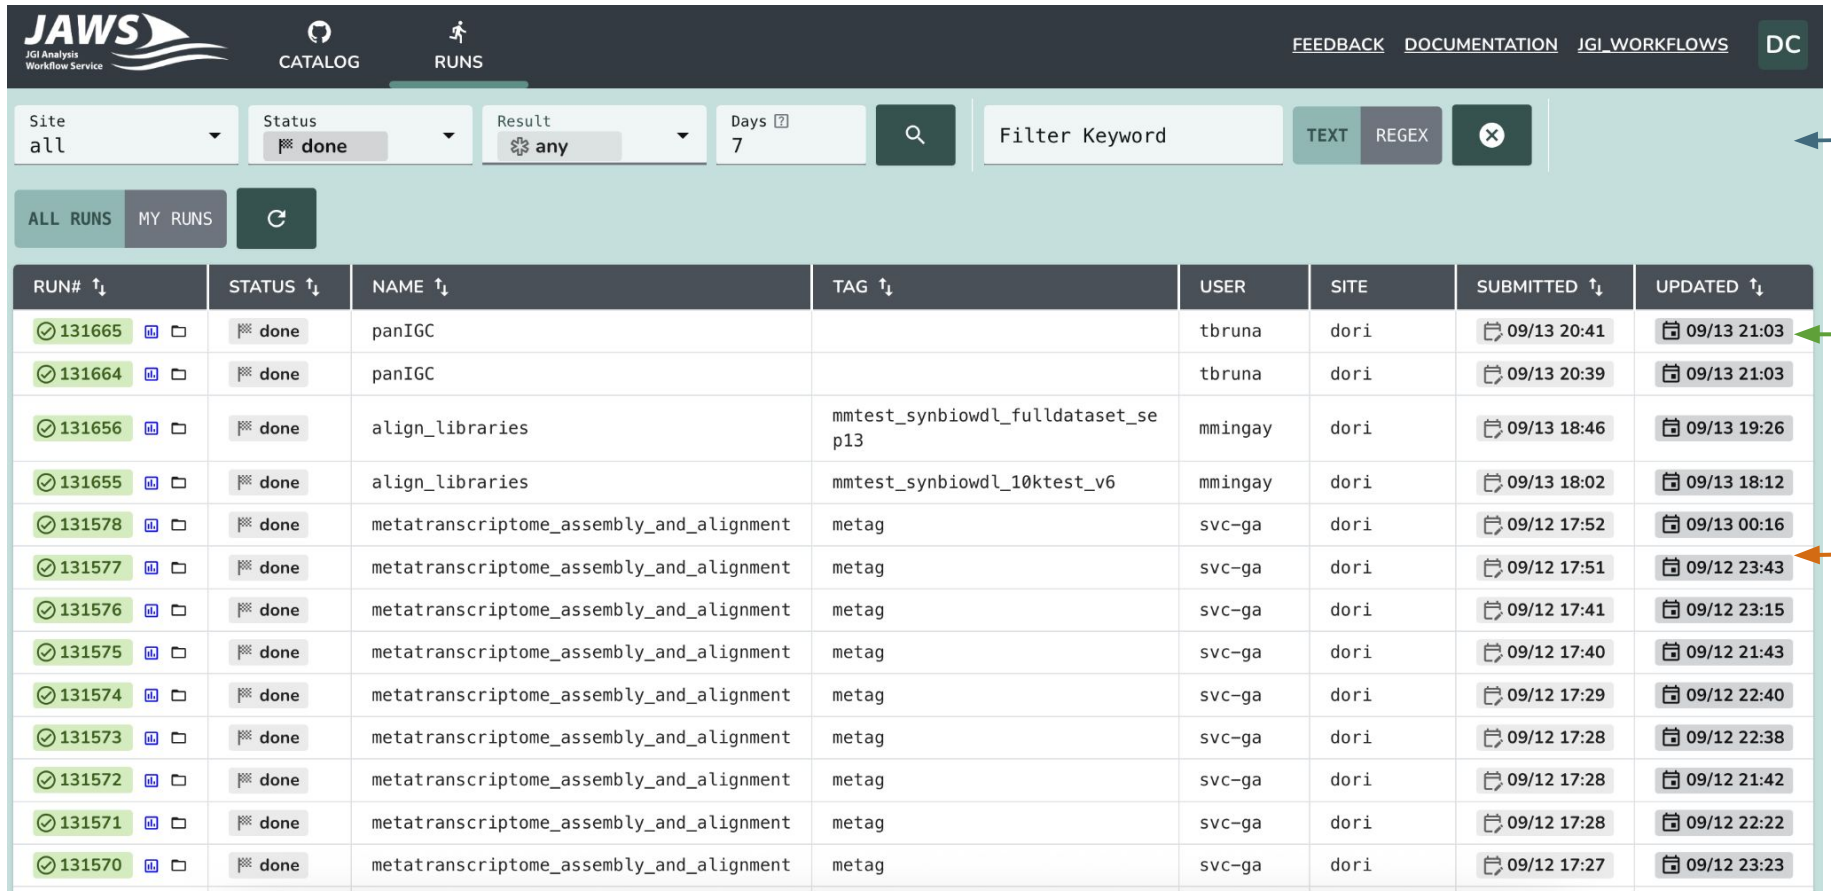

The dashboard makes workflow status visible at scale

JAWS Usage Is Doubling Every Year!

Training sparked adoption. Automation now drives scale.

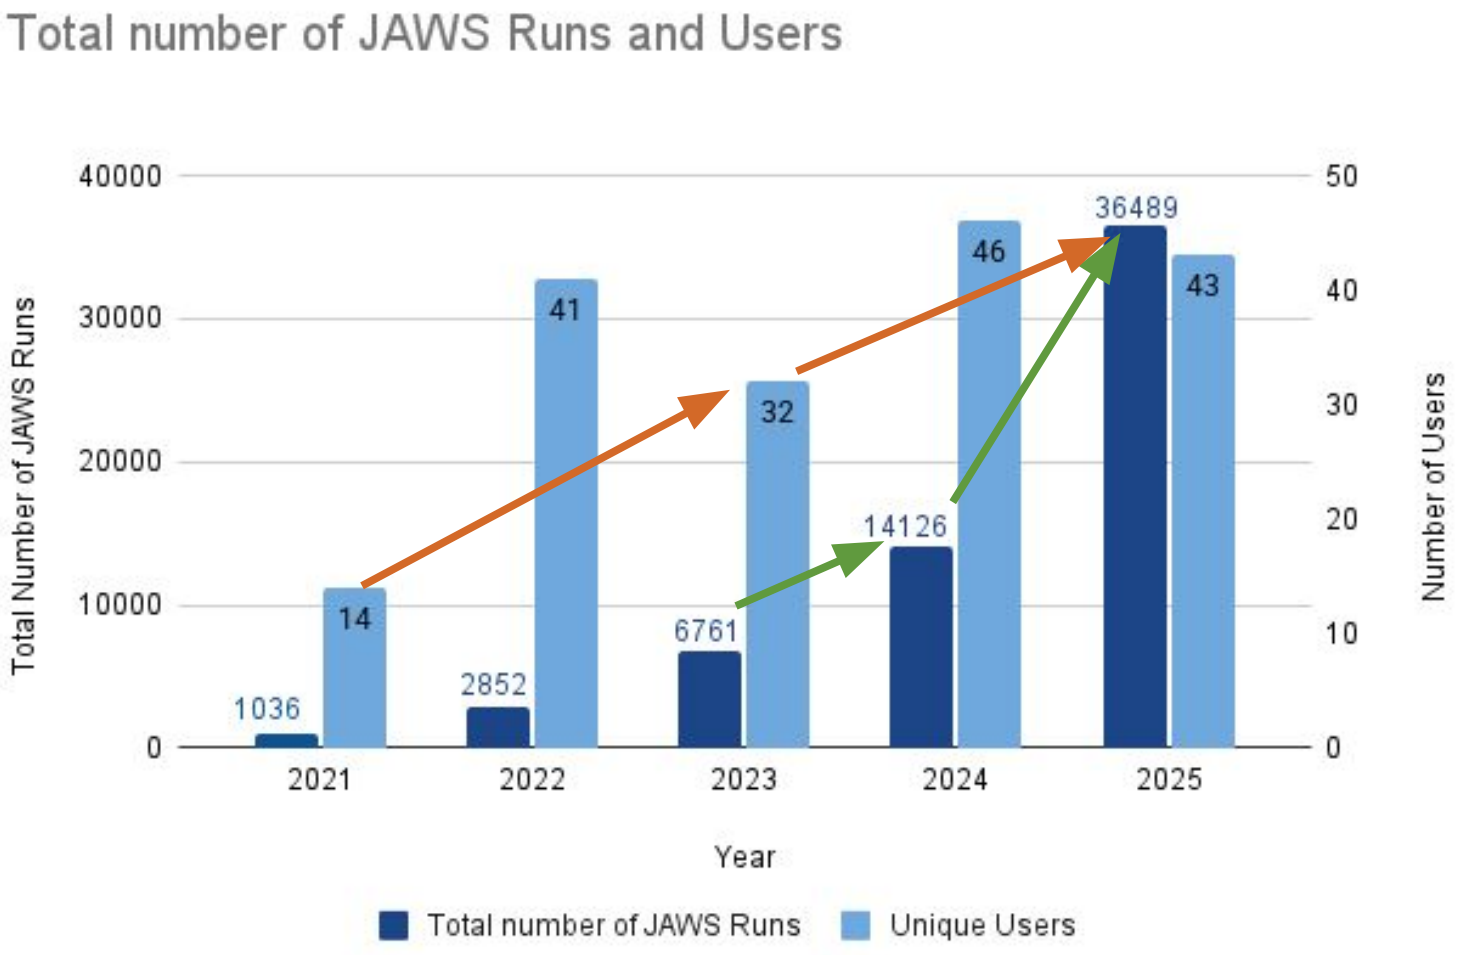

Early Growth (2021–2023):
➔ Internal training workshops

Scaling Up (2023–2025):
➔ Usage more than doubled

Today:
➔ The larger share of runs now comes from automation accounts.

"The JAWS team is a model of what good teams can be! Super responsive, transparent, always seeking and acting on feedback."
(JAWS User Survey 2025)

Powered by Partnership. Ready for Collaboration

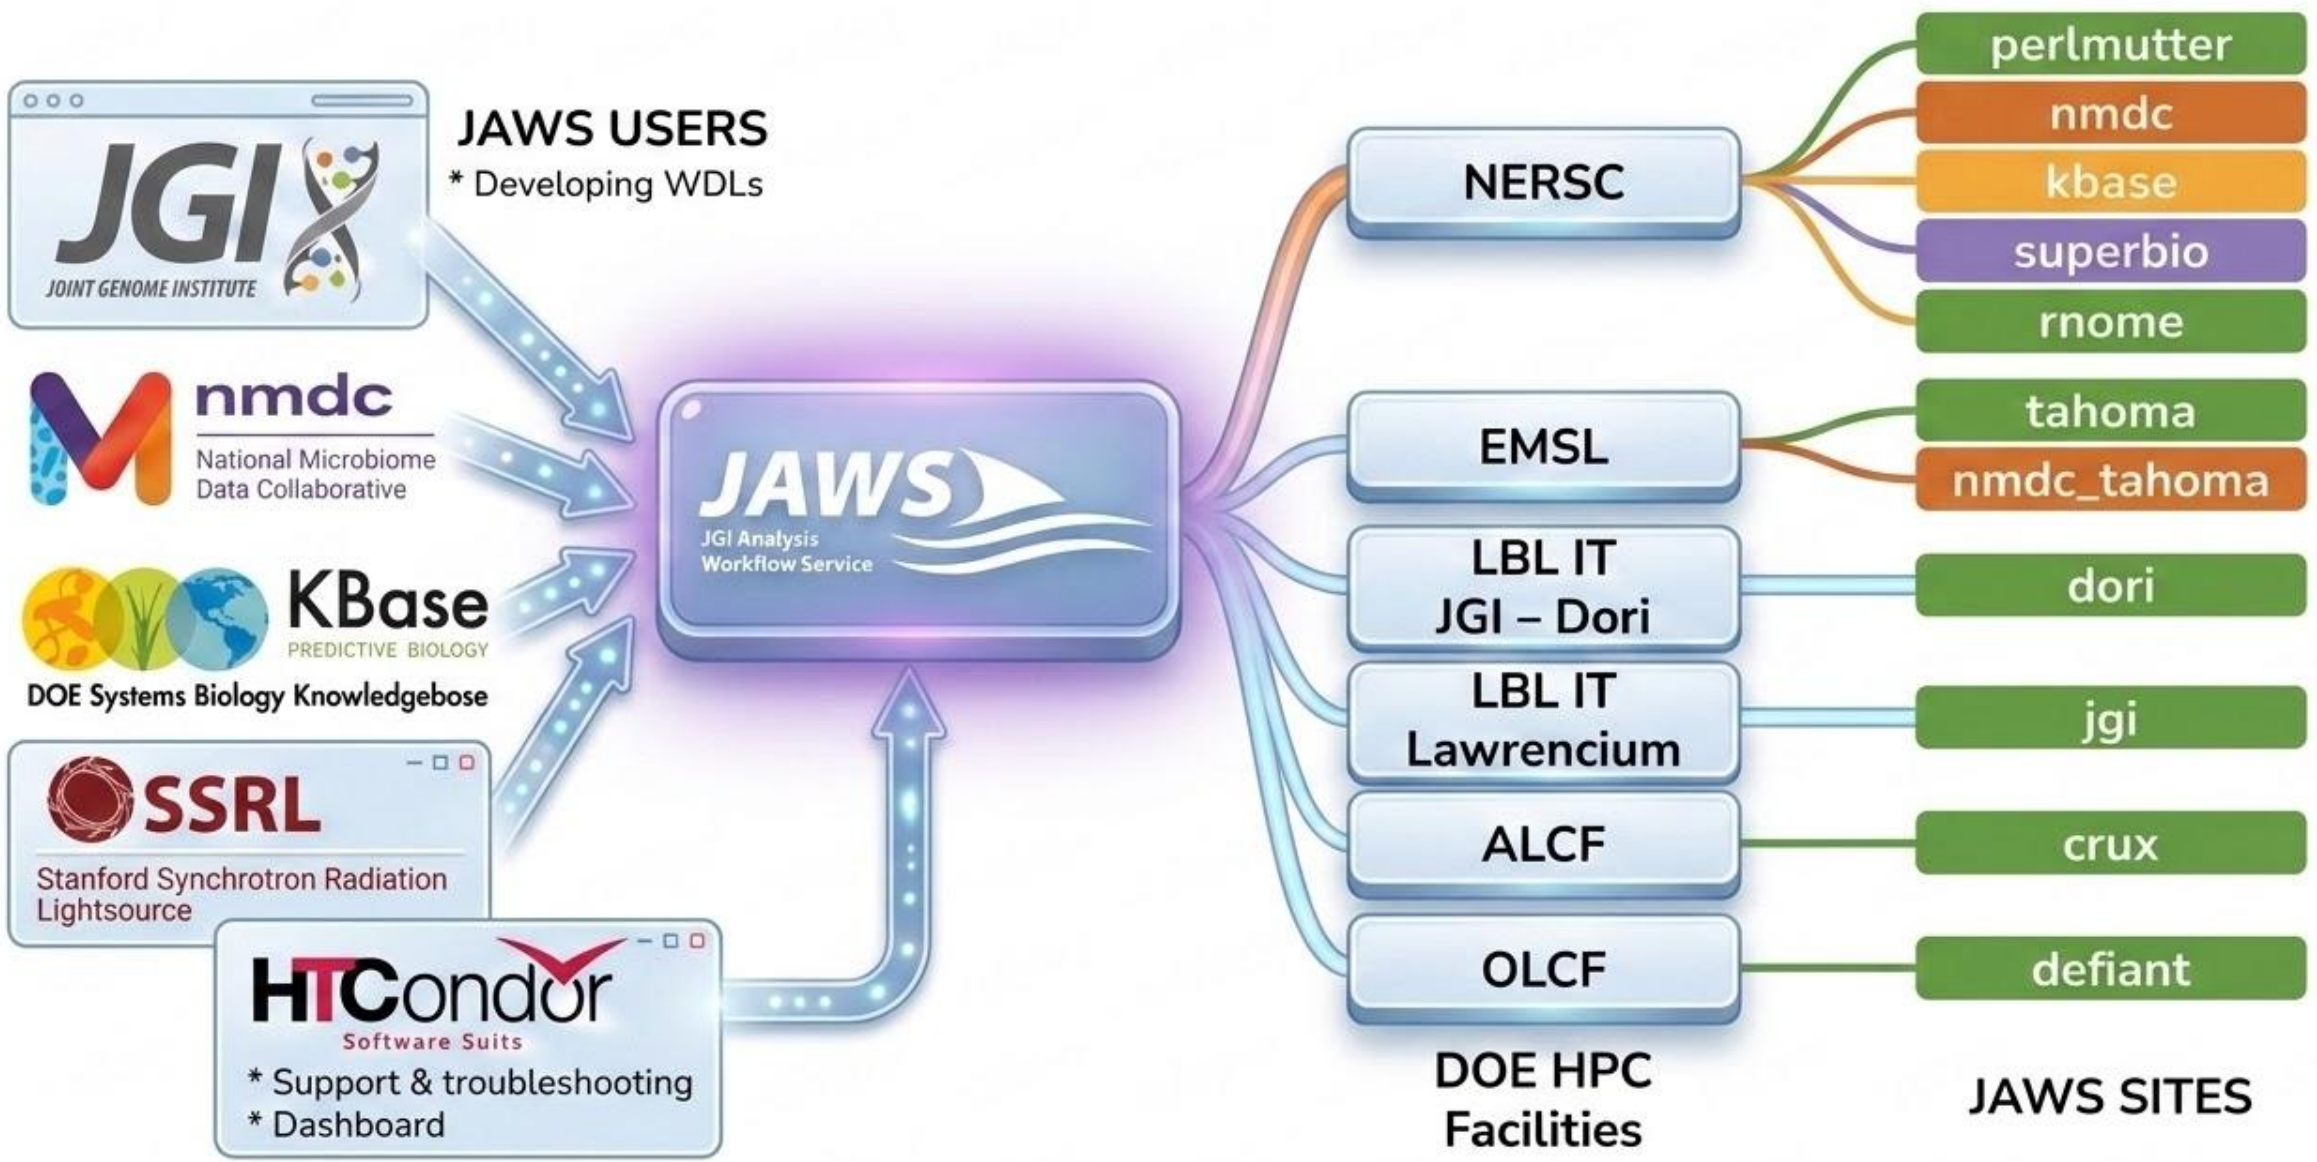

AI for Bio: Next-Gen Data Infrastructure

JOINT GENOME INSTITUTE

JGI is promptly moving forward to become an AI-centric facility that's Data Lakehouse enabled!

- Expanding sites
- Service resiliency
- Monitor and telemetrics

“Hey, JAWS!”

Powered by jaws.ai

DOE Computing Facilities

BERKELEY LAB

NERSC

Pacific Northwest NATIONAL LABORATORY

EMSL

Argonne NATIONAL LABORATORY

OAK RIDGE National Laboratory

Acknowledgements

HTCondor Team

Miron Livny, Todd Tannenbaum, Todd Miller

JAWS Team

Seung-Jin Sul, Mario Melara, Ramani Kothadia, Ludovico Bianchi, Joshua Boverhof, Daniela Cassol, Mike Sneddon, Set Sarrafan, Kjiersten Fagnan

NERSC

Nick Tyler

Q&A

Thank you!

JAWS Docs

JAWS
JGI Analysis
Workflow Service

JAWS Code

Acknowledgements

The work conducted by the U.S. Department of Energy Joint Genome Institute (https://jgi.doe.gov/), a DOE Office of Science User Facility, is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-C02-05CH11231.

Backup Slides

Cromwell and HTCondor

- **Cromwell**
  - Open-source workflow management software by the Broad Institute
  - Executes WDL workflows with various compute backends
  - Creates HTCondor job submit files
  - Submits/manages jobs using HTCondor

	Cromwell backend interface	HTCondor command
Task submission	“submit”	condor_submit
Task cancellation	“kill-docker”	condor_rm
Status check	“check-alive”	condor_q

AI for Bio: Next-Gen Data Infrastructure

JOINT GENOME INSTITUTE

JGI is promptly moving forward to become an AI-centric facility that's Data Lakehouse enabled!

BRIDGE - **B**iological and envi**R**onmental **I**nfrastructure for **D**ata mana**G**ement and **E**xploration

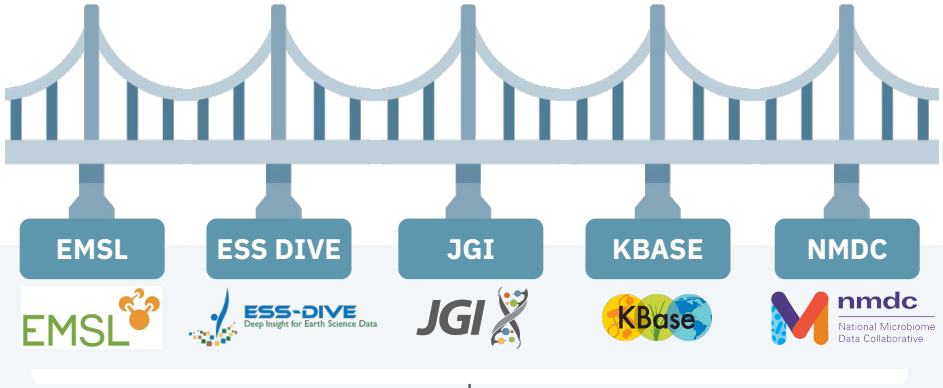

DOE COMPUTING ASSETS

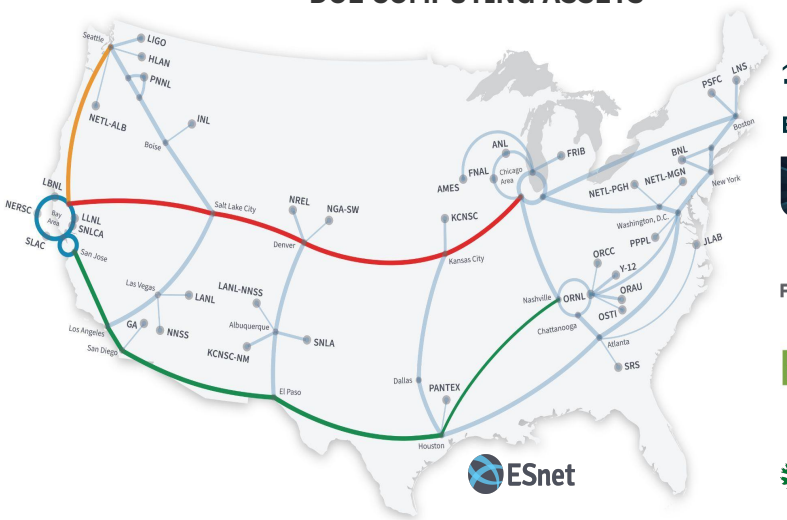

BERKELEY LAB
NERSC
Pacific Northwest NATIONAL LABORATORY
EMSL
Argonne NATIONAL LABORATORY
OAK RIDGE National Laboratory

LARGE SCALE ANALYSES

"Hey, JAWS!"

Powered by jaws.ai

JAWS
JGI Analysis Workflow Service

JGI BRIDGE Initiative

The **BRIDGE** initiative is a biological data infrastructure program developed by the U.S. Department of Energy (DOE) Joint Genome Institute (JGI).

Its name stands for Biological and enviRonmental Infrastructure for Data manaGement and Exploration.

The initiative aims to create an AI-ready data architecture and unified ecosystem to make genomic and environmental research data highly accessible and interoperable.

Key Features of JGI's BRIDGE Initiative

- Data Lakehouse: Integrates raw data repositories and data warehouses for seamless exploration.
- JAMO Integration: Uses JGI's Archive and Metadata Organizer to securely store and transfer data files.
- AI Readiness: Structures vast biological datasets to fuel modern Artificial Intelligence models and tools.
- Cross-Facility Access: Connects with other DOE hubs—such as the Environmental Molecular Sciences Laboratory (EMSL) and KBase—to allow predictive ecosystem biology.

JGI Data Lakehouse

The Department of Energy (DOE) Joint Genome Institute (JGI) is building a **data lakehouse** to transition into an AI-centric user facility.

This architecture combines the flexibility of a data lake (for raw, unstructured data) with the structure of a data warehouse (for indexed, highly-queryable data).

Key Features of the JGI Data Lakehouse

- AI-Ready Data: Standardizes and organizes massive genomic datasets to train foundational AI models across all kingdoms of life.
- Unified Ecosystem: Integrates data from the JGI, KBase, Environmental Molecular Science Laboratory (EMSL), and National Microbiome Data Collaborative (NMDC).
- Better Accessibility: Replaces legacy, disjointed search systems to provide researchers with a single, highly responsive entry point for over 13 petabytes of data.
- Community Driven: Actively collects user feedback to ensure the new architecture addresses past data frustration and fills current use cases.

Sites and Compute Pools

Site	Type	#Nodes	Mem (GB)*	Minutes	#Threads	#GPUs
Perlmutter (NERSC)	Large	3072	492	2865	256	0
	GPU (4x NVIDIA A100 (40GB))	1536	256	2865	128	4
	GPU (4x NVIDIA A100 (80GB))	256	256	2865	128	4
JGI (Lab-IT)	Large	8	492	4305	32	0
Dori (Lab-IT)	Large	100	492	4305	64	0
	Xlarge	16	1980	20160	128	0
Tahoma (EMSL)	Medium	184	364	2865	36	0
	Xlarge	24	1480	2865	36	0
	GPU (2x NVIDIA Tesla V100 32GB)	24	1480	2865	36	2
Crux (ALCF)	Medium	256	256	1425	256	0
Defiant (OLCF)	Medium	20	492	4305	128	0

The DOE Joint Genome Institute (JGI)

JGI JOINT GENOME INSTITUTE

About us

- High-throughput AI-centric genomic science user facility located at Lawrence Berkeley National Laboratory
- Provides the genomic capabilities, data, and expertise that supports the global research community in studying complex biological and environmental systems

Advanced genomic capabilities

Large-scale sequencing and synthesis

AI-ready scientific data

Curated data products for reuse

Professional expertise

Support from domain and workflow experts

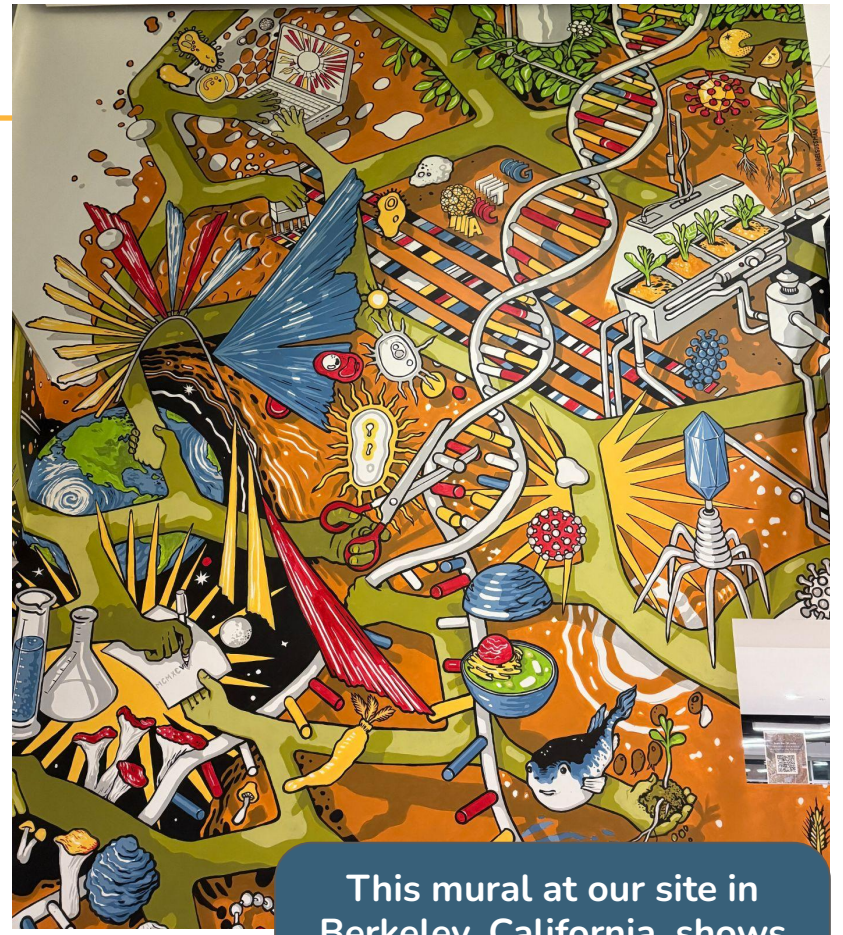

This mural at our site in Berkeley, California, shows some of the species that JGI users have studied.