Global Services for Global Science

Ian Foster

THE UNIVERSITY OF CHICAGO

globus labs

Argonne NATIONAL LABORATORY

Compare climate models to **understand the earth system**

Coupled Model Intercomparison Project (CMIP): Standard protocol for studying general circulation model output

CMIP6	11,307,181 total datasets 20,823.68 TB	5,400,359 distinct datasets 11,236.58 TB	5,906,822 replica datasets 9,587.1 TB
CORDEX	183,980 total datasets 1,391.12 TB	183,708 distinct datasets 1,390.56 TB	272 replica datasets 0.56 TB
CMIP5	201,129 total datasets 5,295.44 TB	52,163 distinct datasets 1,527.12 TB	148,966 replica datasets 3,768.32 TB
INPUT4MIPS	11,492 total datasets 19.91 TB	5,660 distinct datasets 9.97 TB	5,832 replica datasets 9.94 TB
OBS4MIPS	210 total datasets 0.2 TB	210 distinct datasets 0.2 TB	0 replica datasets 0 TB

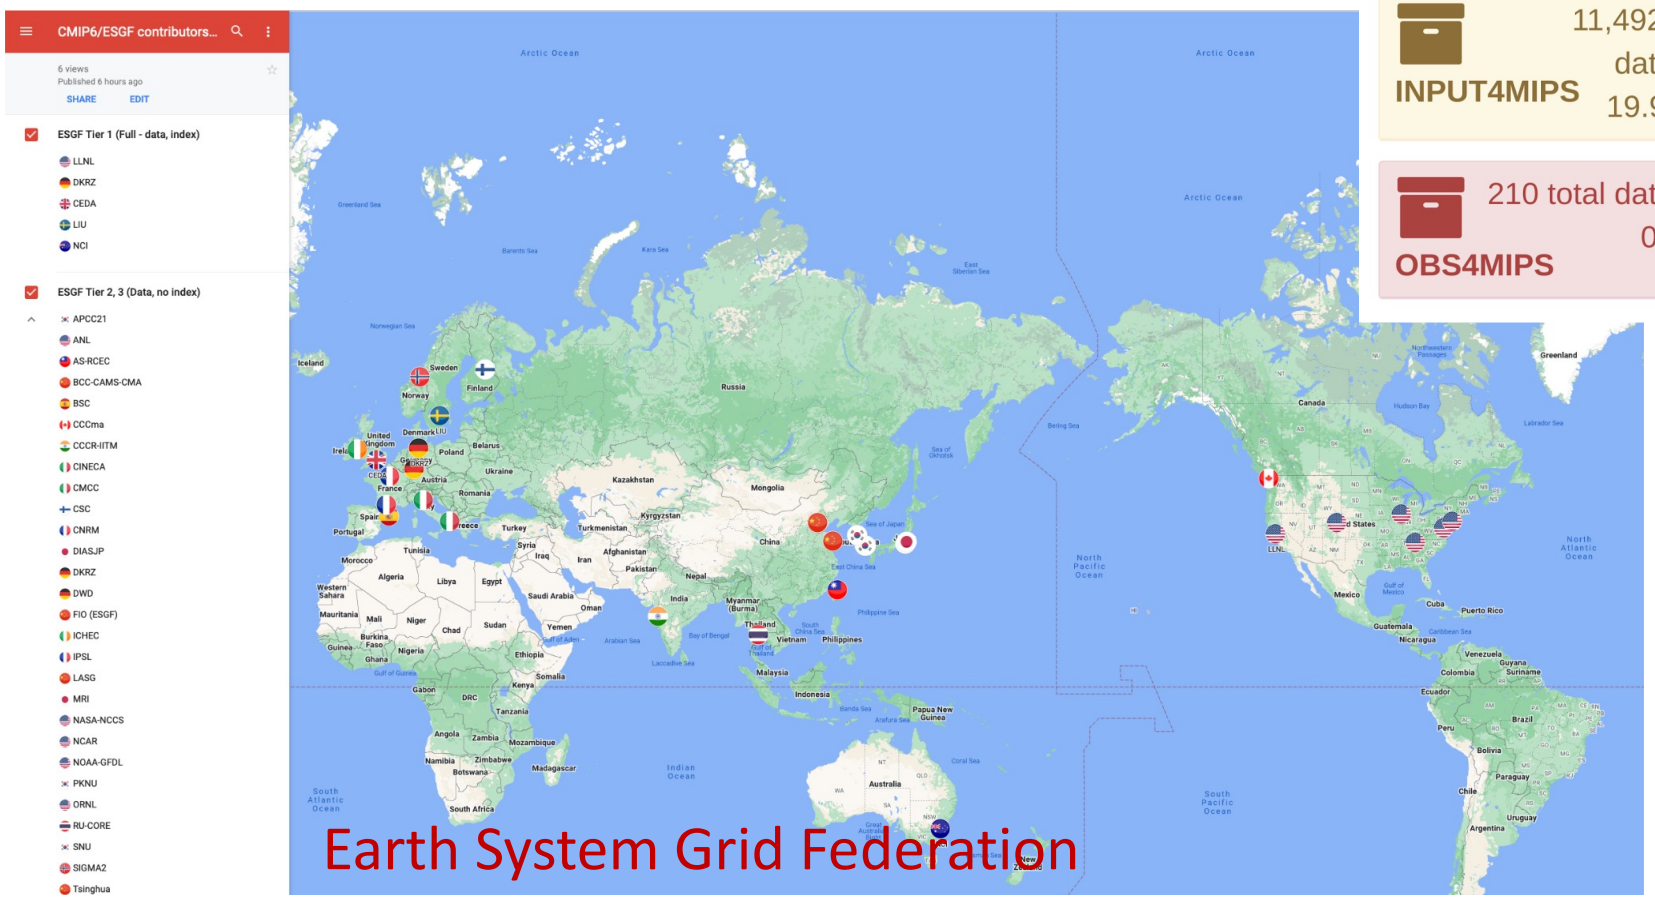

Earth System Grid Federation

>100 models, >20 countries

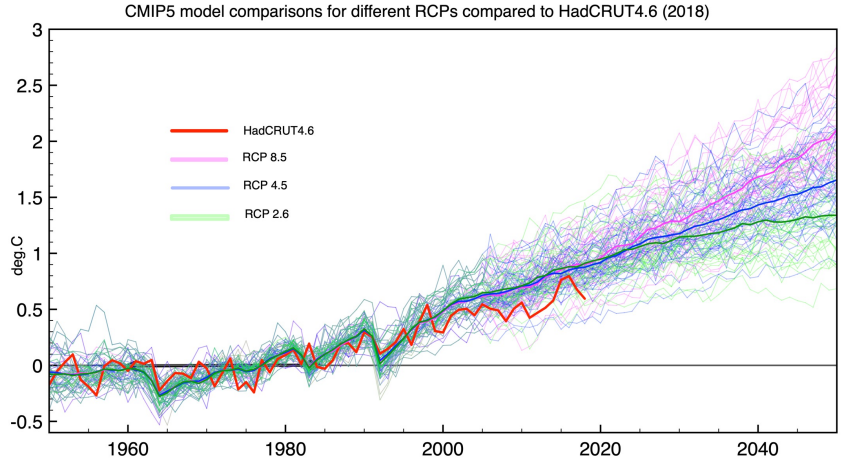

Link field data on pathogen variants with AI for **pandemic surveillance**

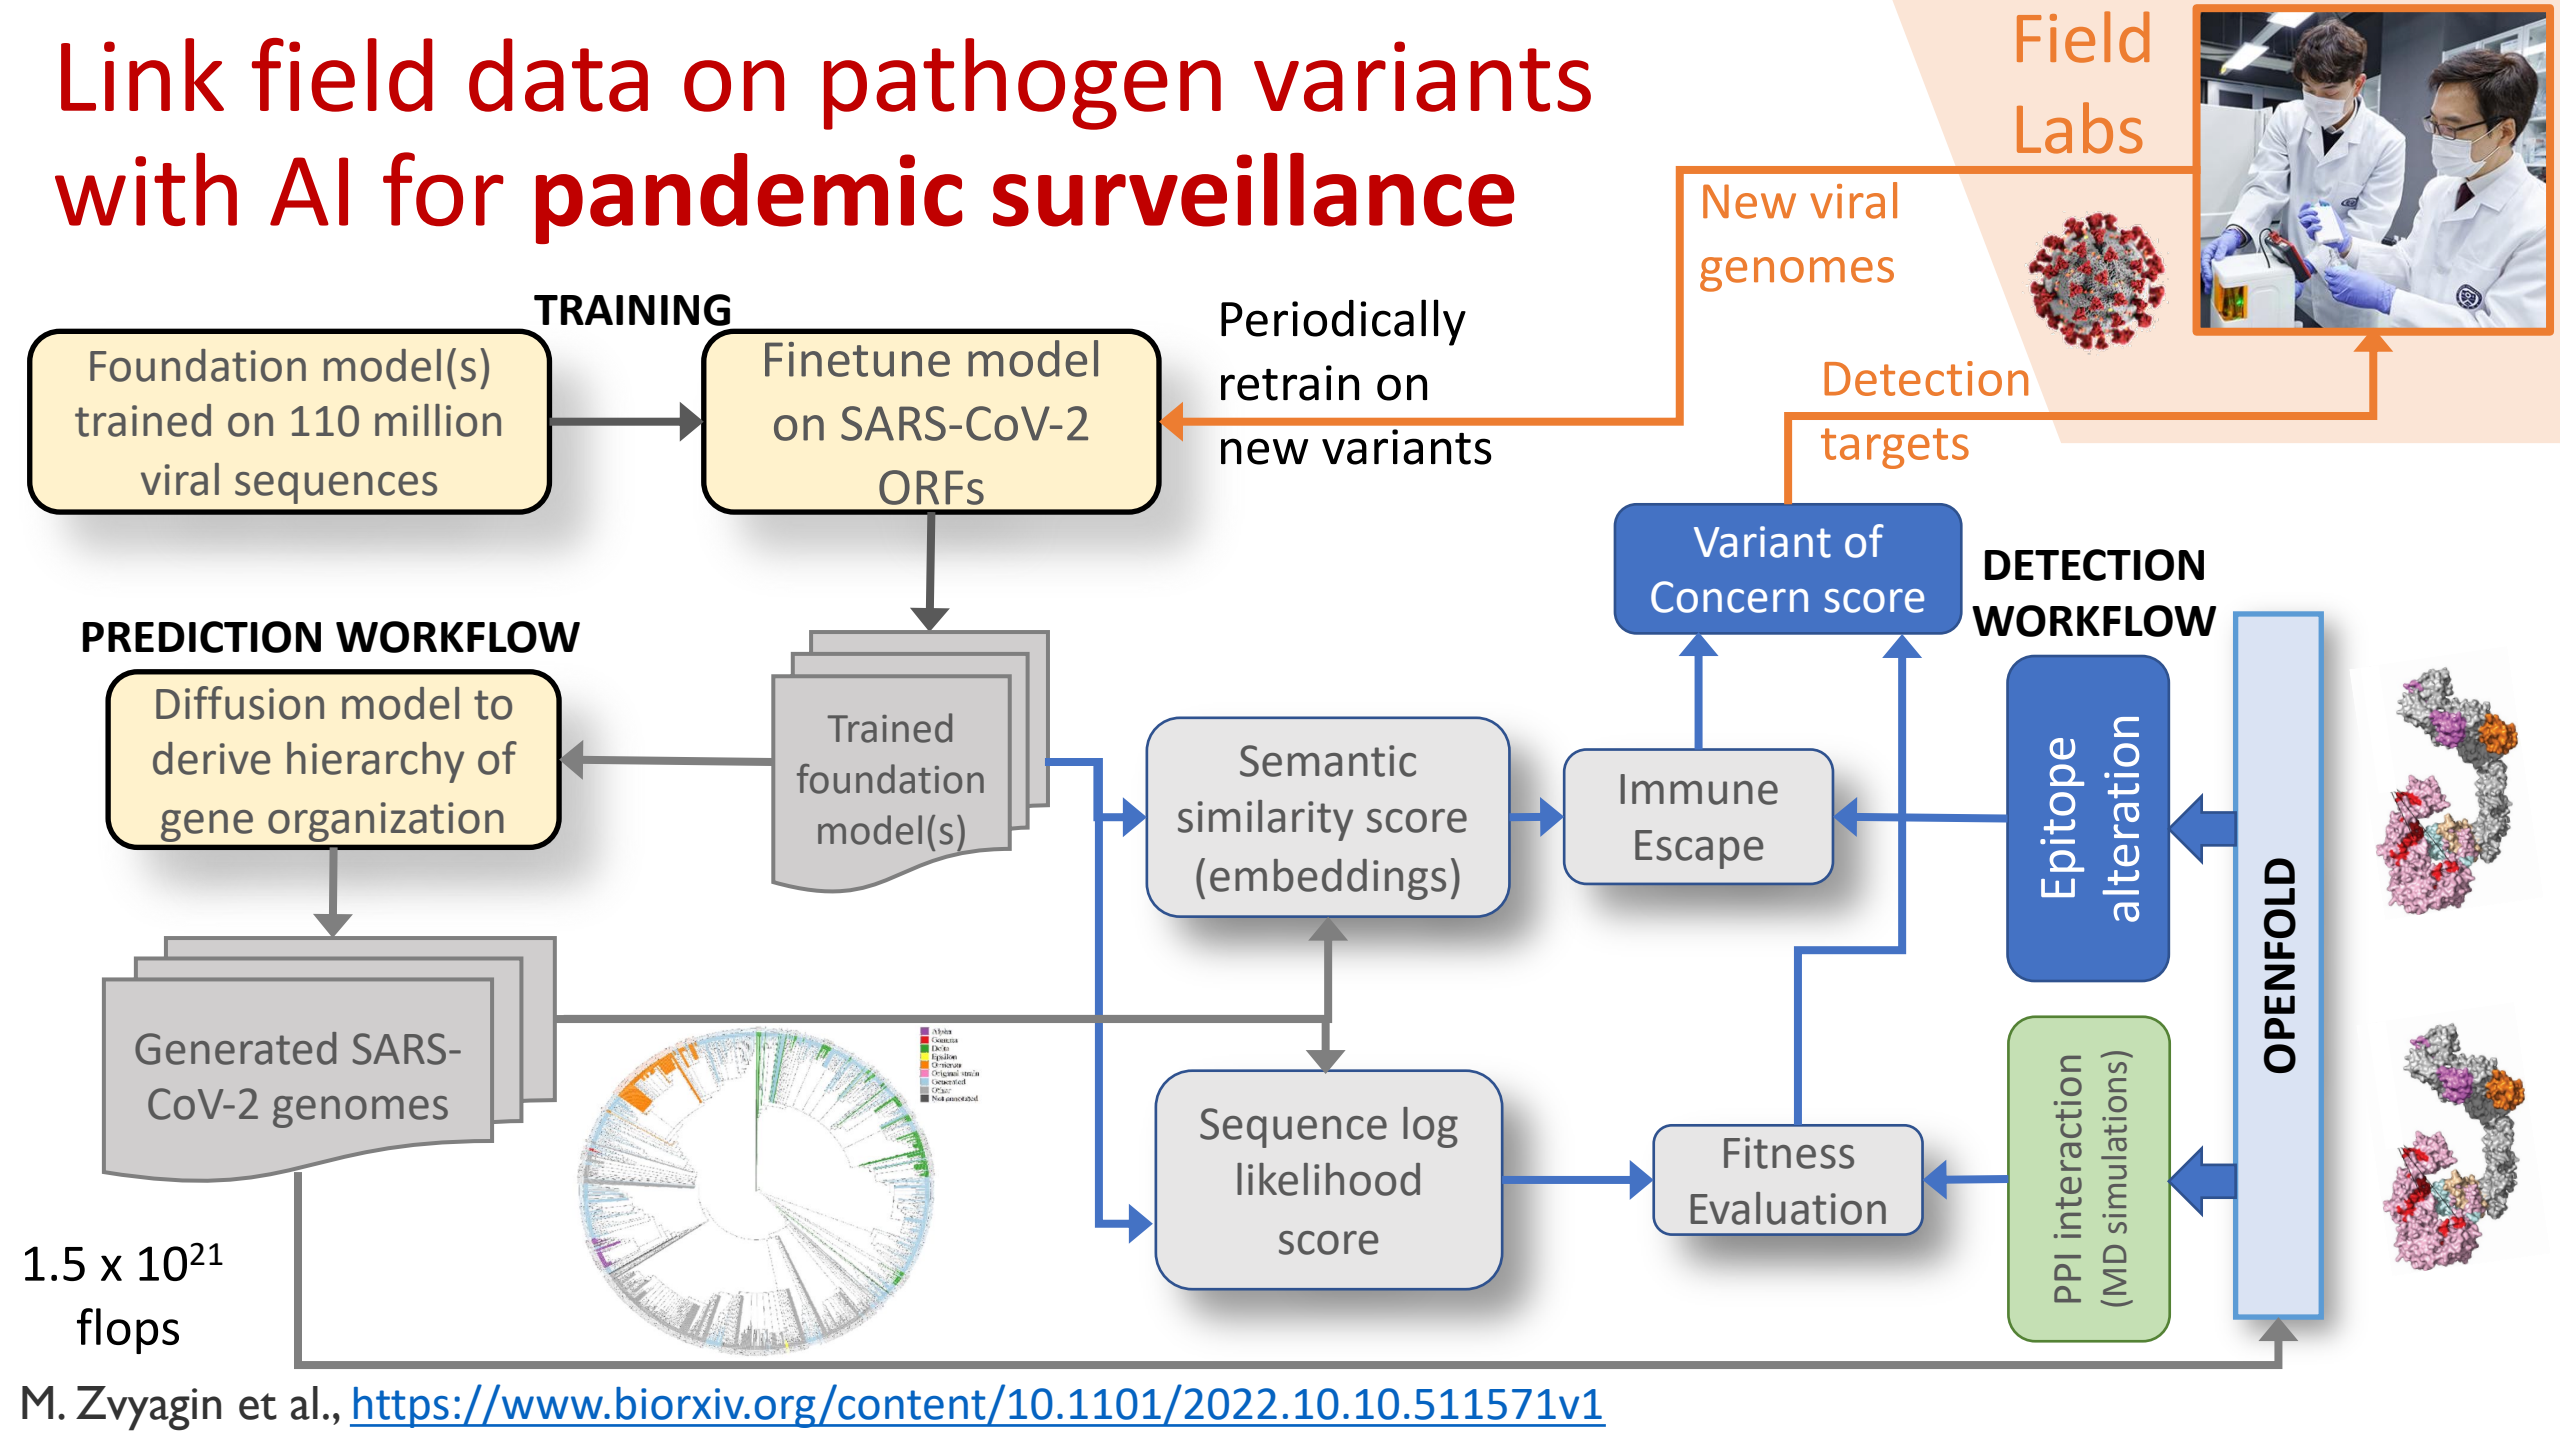

M. Zvyagin et al., https://www.biorxiv.org/content/10.1101/2022.10.10.511571v1

Connect scientific instruments to remote computers to create **smart instruments**

Model (re)training on local GPU:
1102 seconds

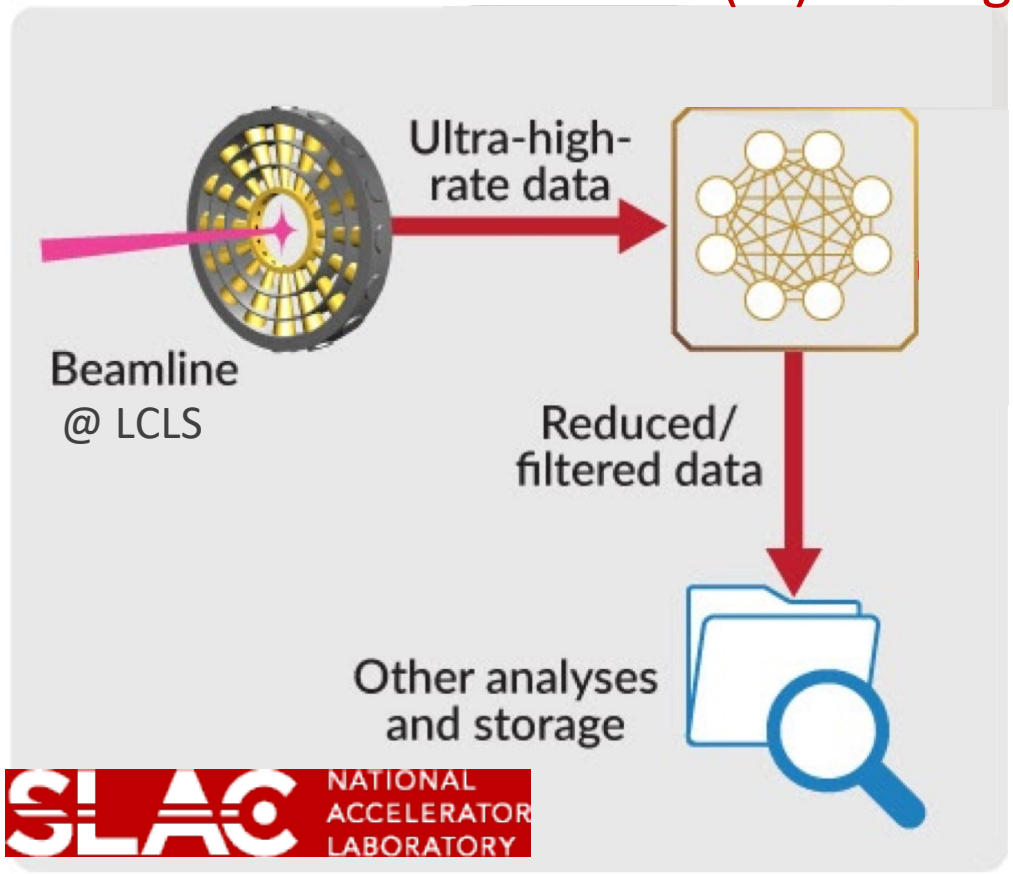

Z. Liu et al., https://doi.org/10.48550/arXiv.2105.13967

Connect scientific instruments to remote computers to create **smart instruments**

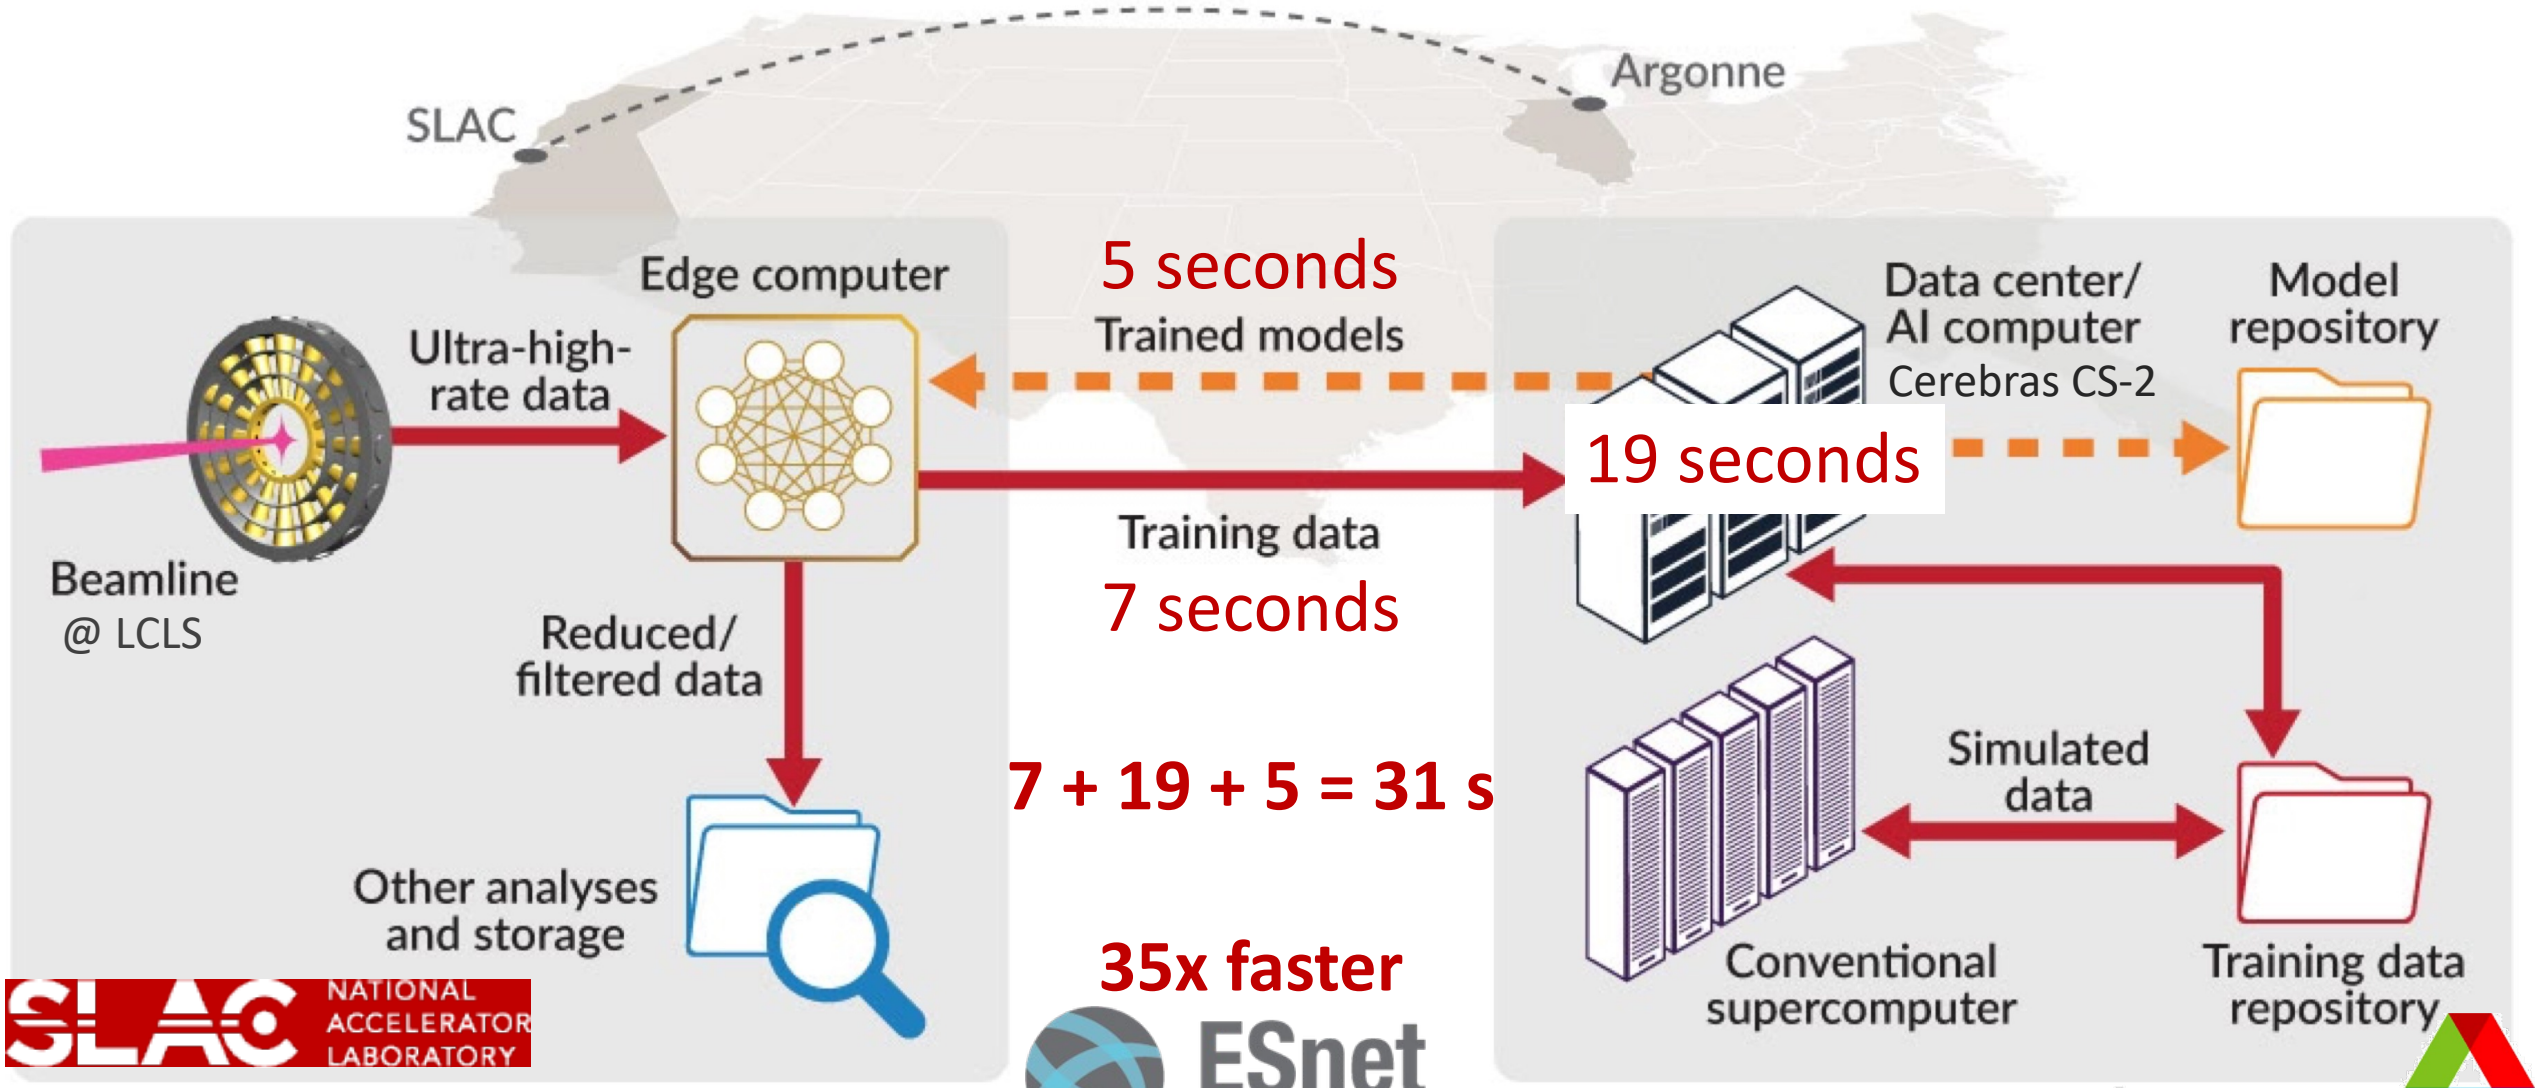

Z. Liu et al., https://doi.org/10.48550/arXiv.2105.13967

Integrate data & models to advance **urban science & resilience**

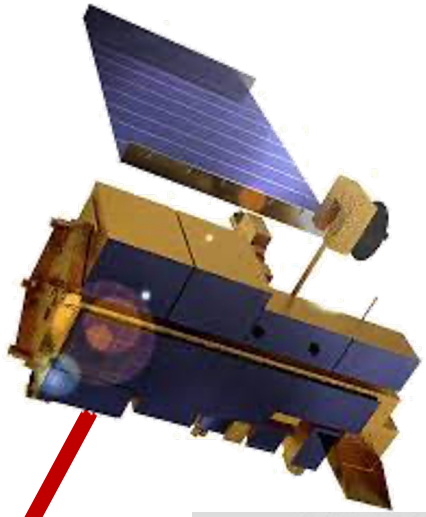

CROCUS
Community Research on Climate & Urban Science

‘A new paradigm for doing science’

Argonne to deploy sensors to track climate change on a neighborhood level

By Maddie Ellis
Chicago Tribune

Next to an almost-200-foot red-and-white-striped tower at Argonne National Laboratory sits a building filled with newly opened boxes of instruments. These tools measure climate conditions like air quality and precipitation, and compared with the lab’s historic tools, like the tower outside, they’re small — really small.

That’s because instead of measuring the region’s atmospheric conditions from Argonne’s sprawling site in DuPage County, researchers will use these tools in a different kind of lab — the city of Chicago. Data collected will be used in modeling to show the effects of climate at scales as small as individual neighborhoods, said Cristina Negri, the director for the project.

Research scientist Adam Theisen, left, speaks with atmospheric scient … onal Laboratory in Lemont. **ANTONIO PEREZ/CHICAGO TRIBUNE PHOTOS**

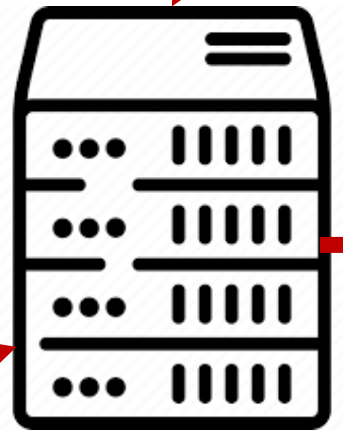

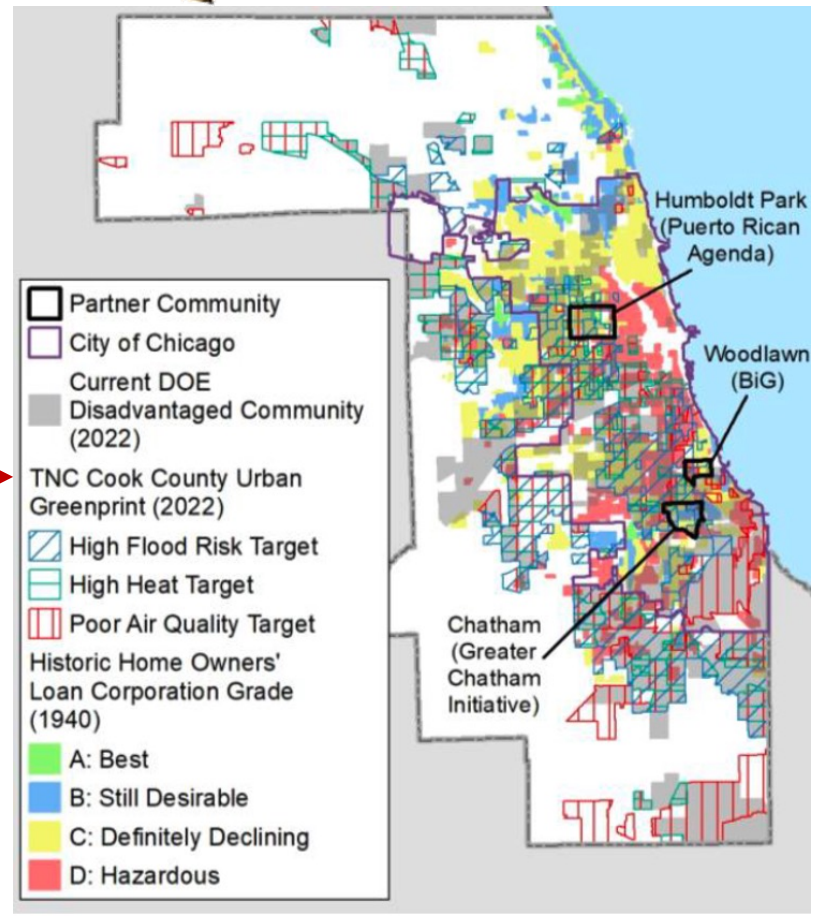

Integration of distributed resources, a sine qua non for these applications, is impeded by many sources of **friction**

"Whereas computational friction expresses the struggle involved in transforming data information and knowledge … **data friction** expresses a more primitive form of resistance -- **the costs in time, energy, and attention required simply to collect, check, store, move, receive, and access data**. Whenever data travel … data friction impedes their movement" (Edwards, 2010, p. 84).

Friction first characterized by Leonardo di Vinci in 1493

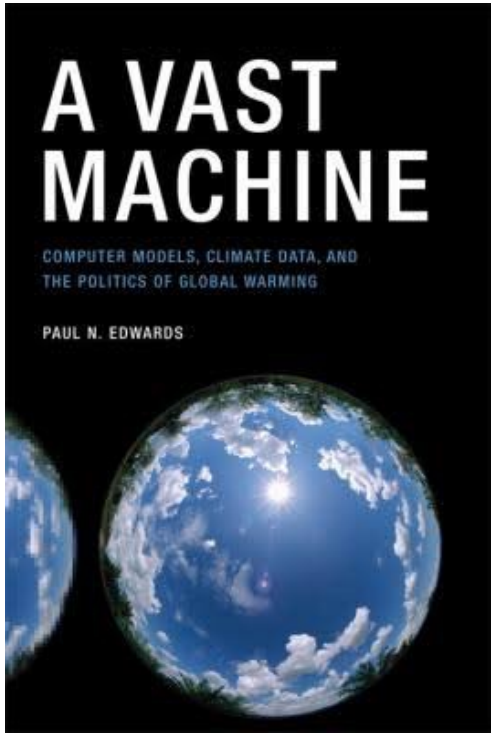

Increasingly, telecommunications is not the problem

Increasingly, telecommunications is not the problem

Aalyria's "Spacetime" [originally "Minkowski"] platform

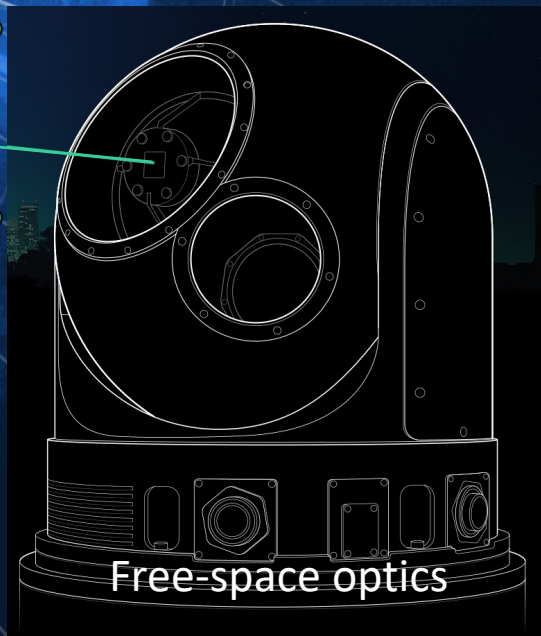

Free-space optics

Increasingly, telecommunications is not the problem

”Henceforth, **space** for itself, and **time** for itself, are doomed to fade away into mere shadows, and only a kind of union of the two will preserve an independent reality.” – Hermann Minkowski

A **computing continuum**

Aalyria’s “Spacetime” [originally “Minkowski”] platform

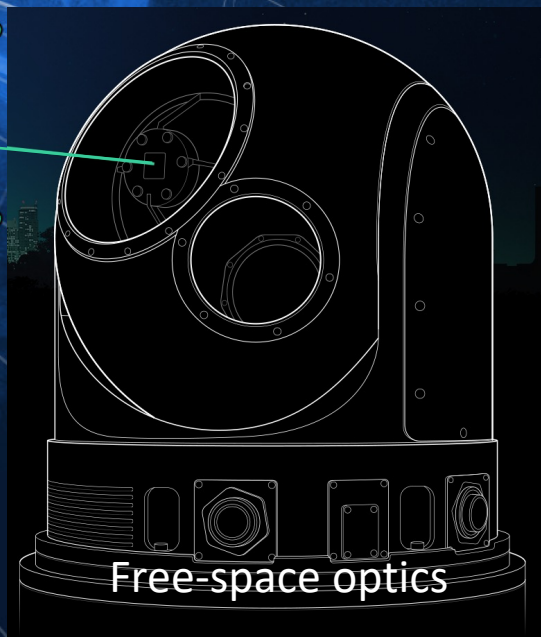

Free-space optics

But it remains unduly difficult to:

1) **Act on resources** regardless of location and interface

Friction: Varying interfaces, behaviors; reliability; security

But it remains unduly difficult to:

1) **Act on resources** regardless of location and interface

Friction: Varying interfaces, behaviors; reliability; security

2) Execute remote actions **reliably**

Friction: Failures, scalability, usability

But it remains unduly difficult to:

1) **Act on resources** regardless of location and interface

Friction: Varying interfaces, behaviors; reliability; security

2) Execute remote actions **reliably**

Friction: Failures, scalability, usability

3) Manage <u>who</u> is **trusted** to perform <u>what</u> actions, <u>where</u> and <u>when</u>

Friction: Varying credentials, authentication protocols, authorization policies; need to act on behalf of others

Need: 1) **Act anywhere**

Compare climate models to **understand the earth system**

Coupled Model Intercomparison Project (CMIP): Standard protocol for studying general circulation model output

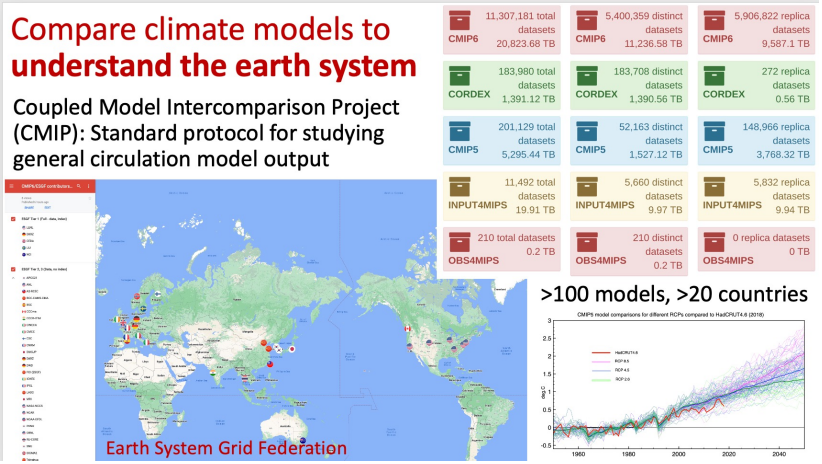

Link field data on pathogen variants with AI for **pandemic surveillance**

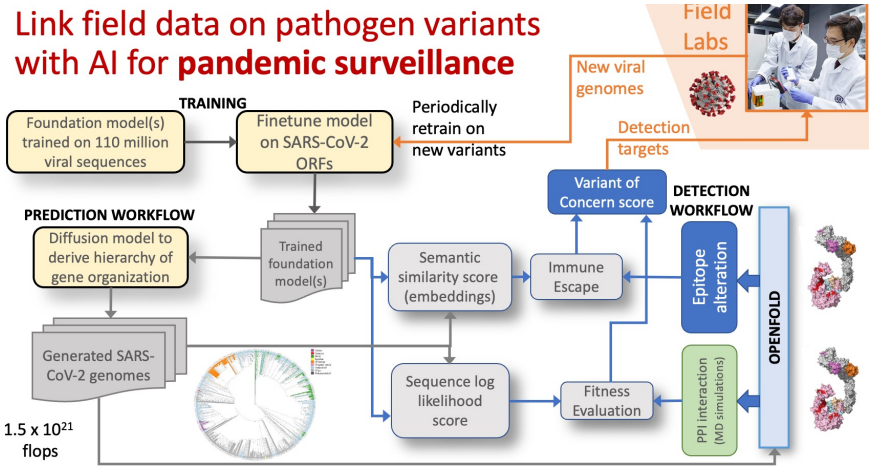

Integrate data & models to advance **urban science & resilience**

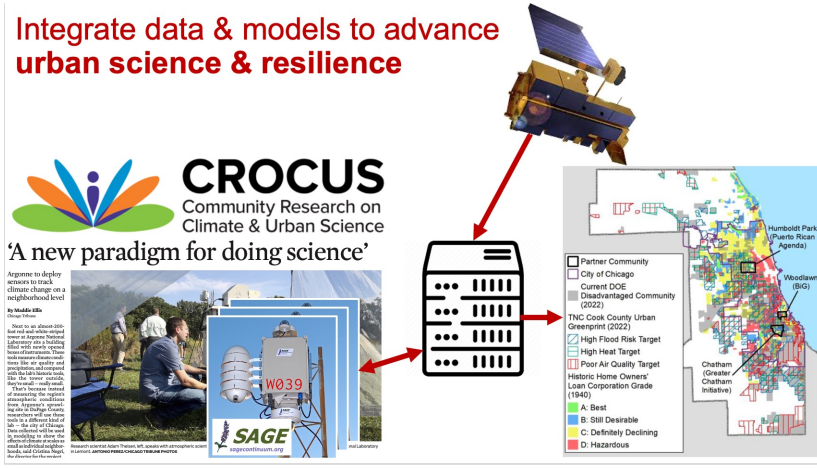

Past approaches (data actions):

- Gopher, FTP, Web, OPenDAP, …
- Distributed file systems

Past approaches (compute actions):

- SSH, grid protocols, cloud APIs
- Java, virtual machines, containers

Challenges: Performance, scalability, reliability, portability, usability

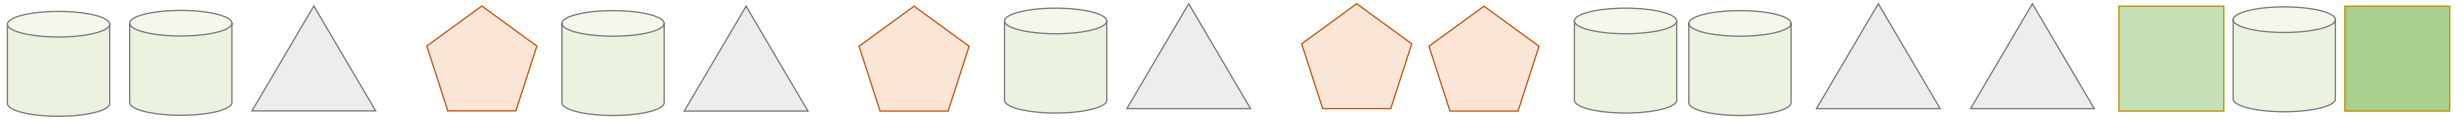

Need: 1) **Act anywhere**

Compare climate models to **understand the earth system**

Coupled Model Intercomparison Project (CMIP): Standard protocol for studying general circulation model output

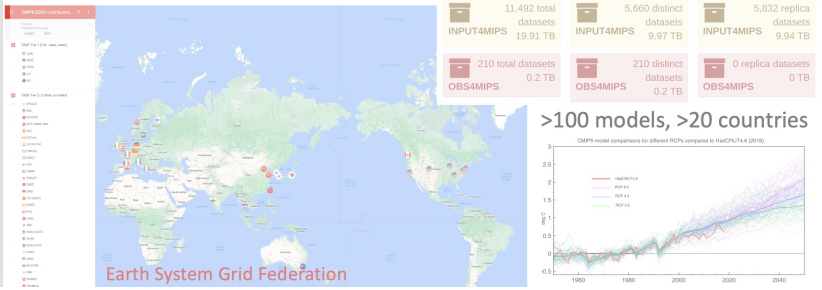

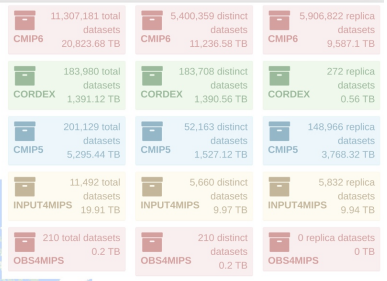

>100 models, >20 countries

Link field data on pathogen variants with AI for **pandemic surveillance**

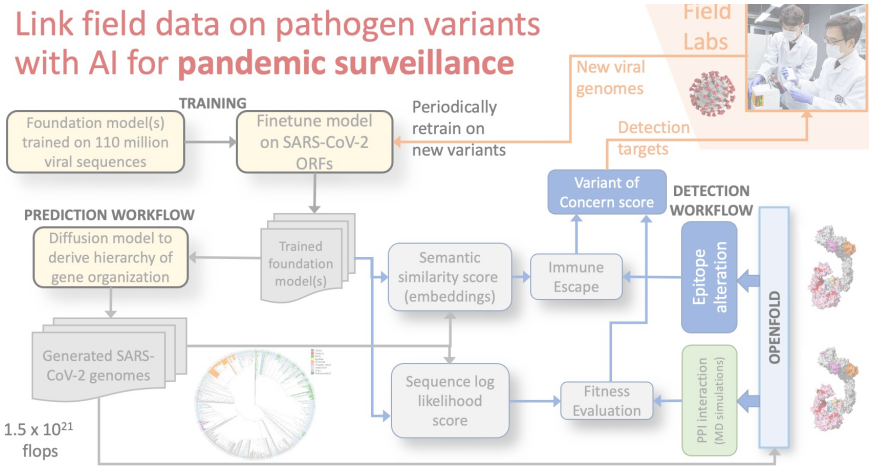

Integrate data & models to advance **urban science & resilience**

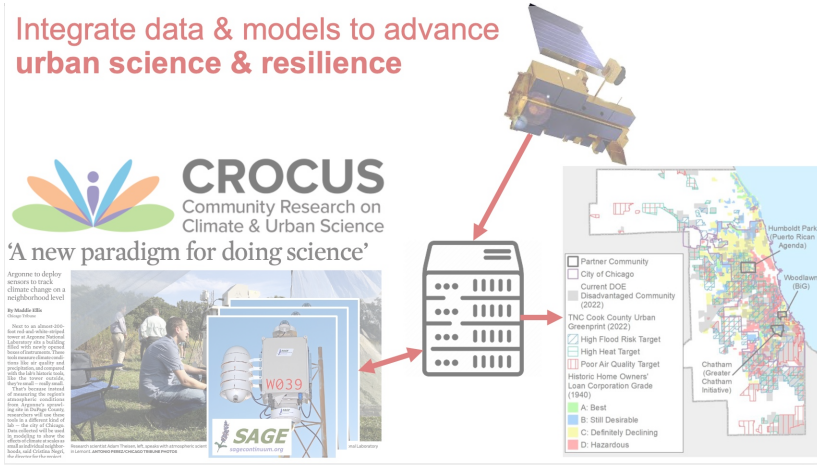

Our approach:

- HTTPS, GridFTP for universal, fast access
- Local agents for broad deployment
- Modularity to target many systems
- Integration with secure delegation
- Integration with hosted supervision

Lightweight agents abstract local details

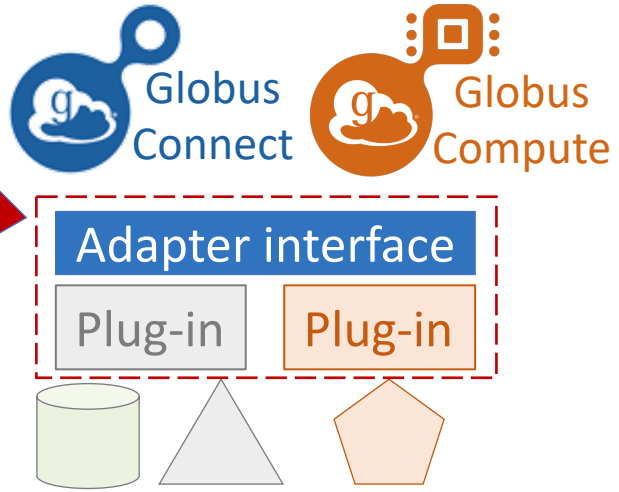

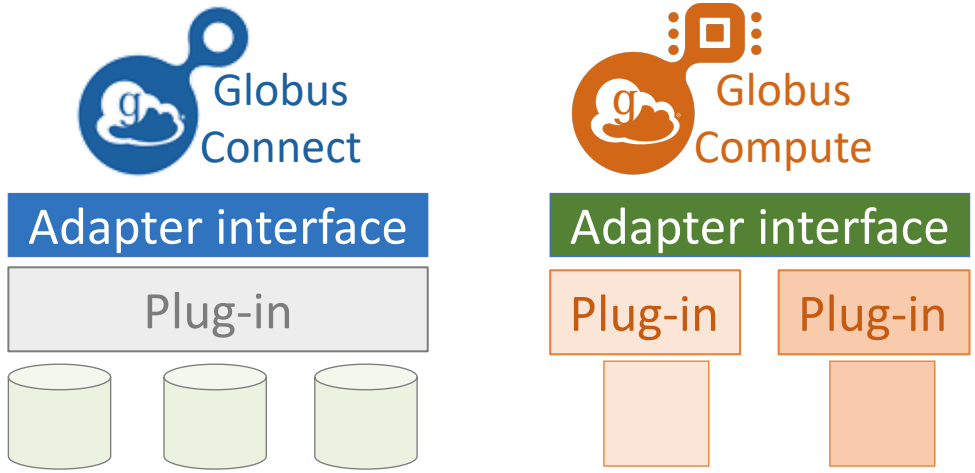

Need: 2) **Reliable execution** of (sets of) actions

Compare climate models to **understand the earth system**

Coupled Model Intercomparison Project (CMIP): Standard protocol for studying general circulation model output

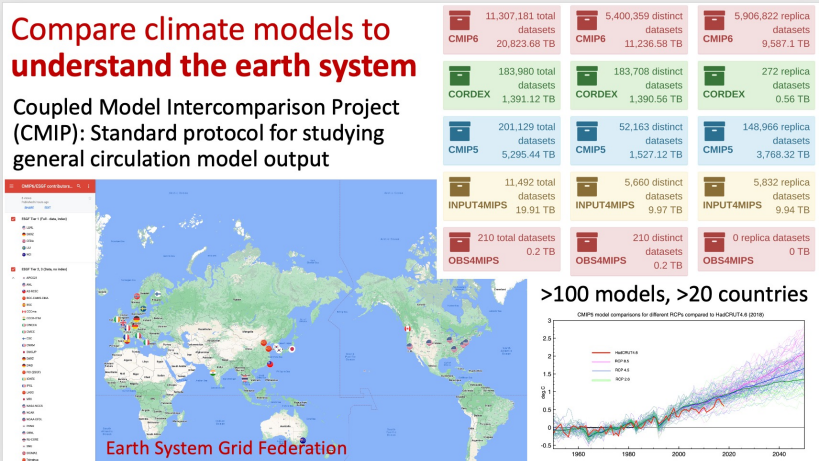

>100 models, >20 countries

Link field data on pathogen variants with AI for **pandemic surveillance**

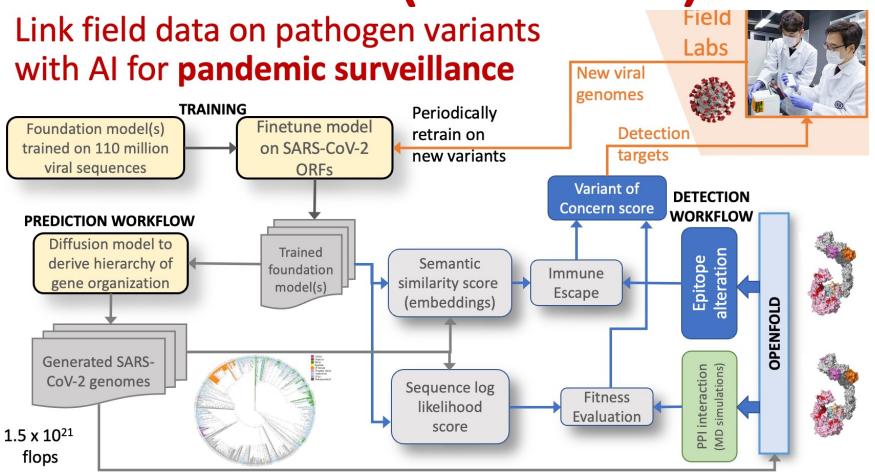

Integrate data & models to advance **urban science & resilience**

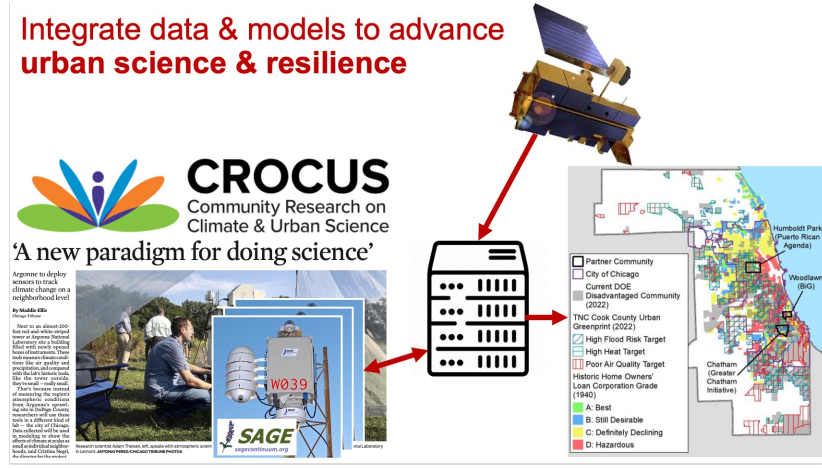

Past approaches:

- Workflow systems
- Distributed file systems
- Eventing, consistency protocols
- Reliable RPC, replication
- Cloud

Challenges: Complexity, fragility, scalability, reach

Need: 2) **Reliable execution** of (sets of) actions

Our approach:

- Cloud-hosted, replicated supervision
- Simple retry-based protocols
- Reduce endpoint complexity
- High assurance for sensitive data
- Integration with secure delegation

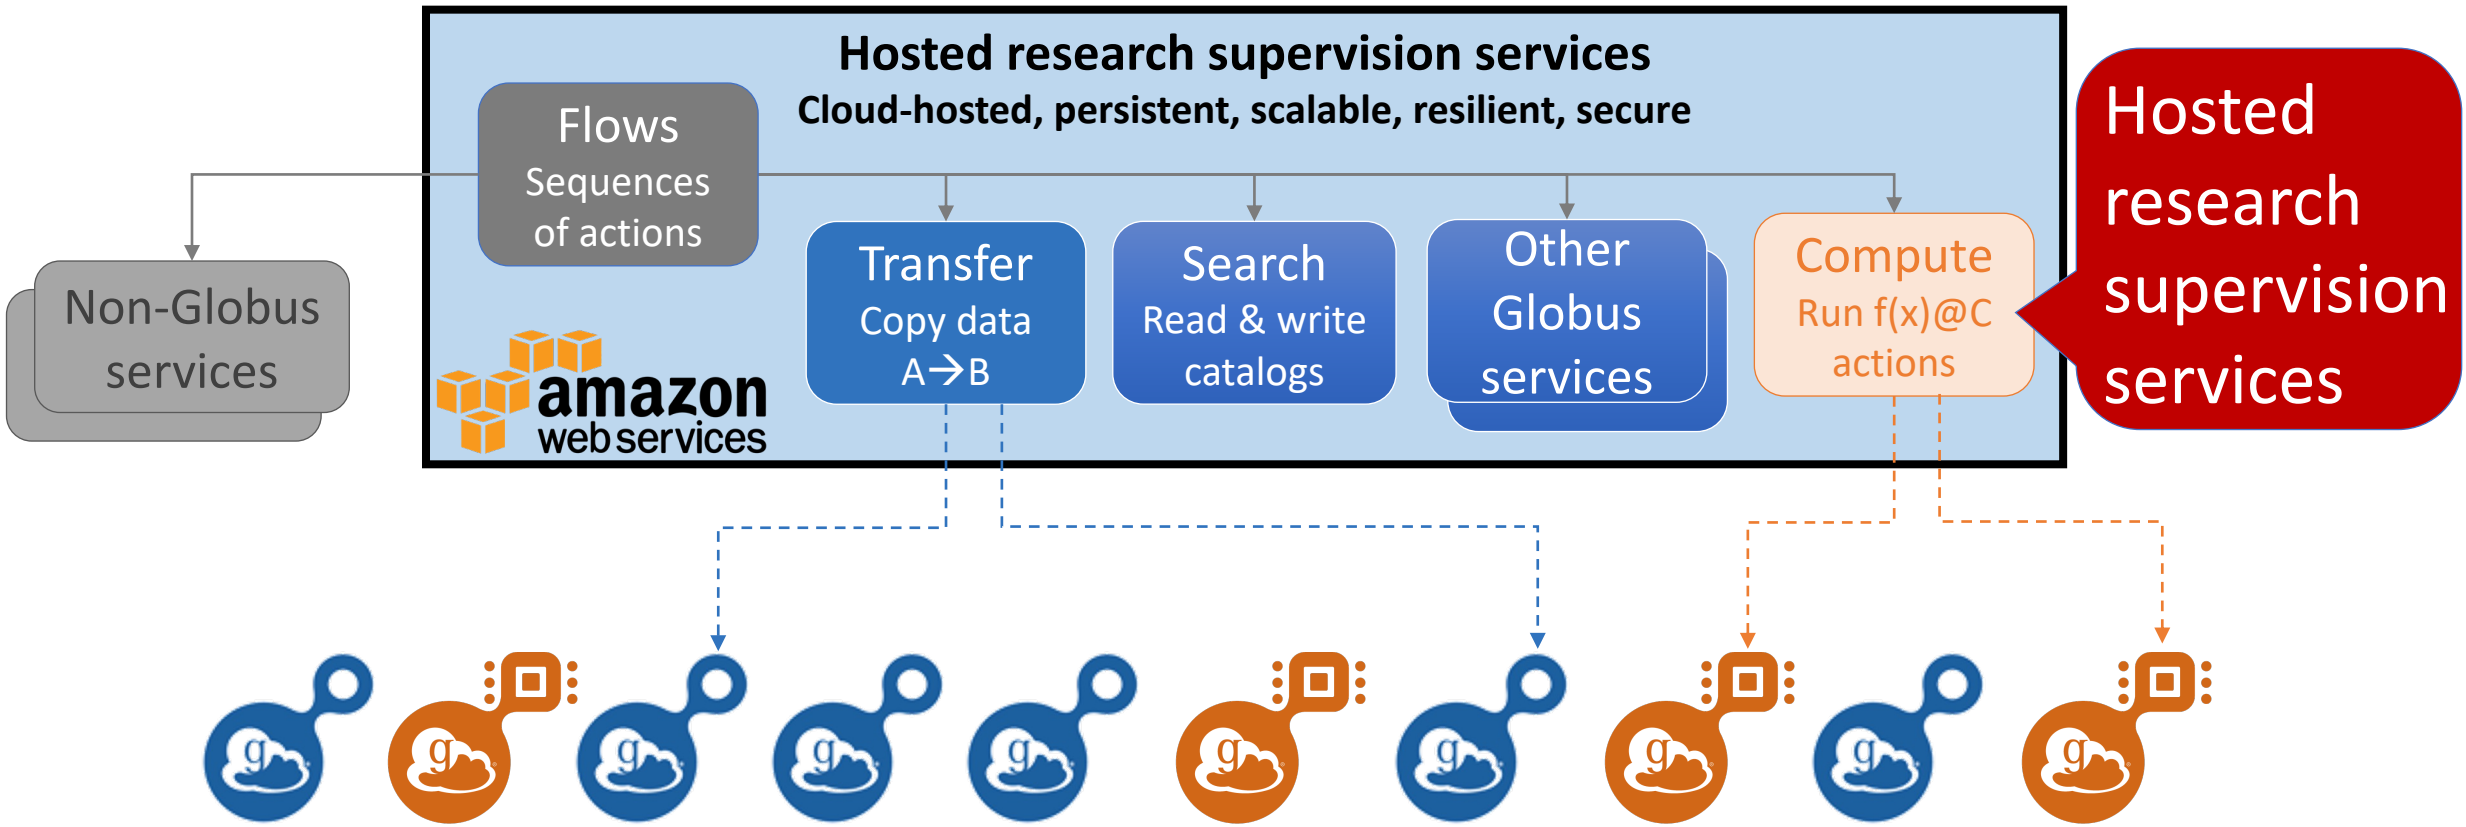

Hybrid model for distributed systems at scale

Global management & orchestration

Local agents with plug-in

Institutional resources

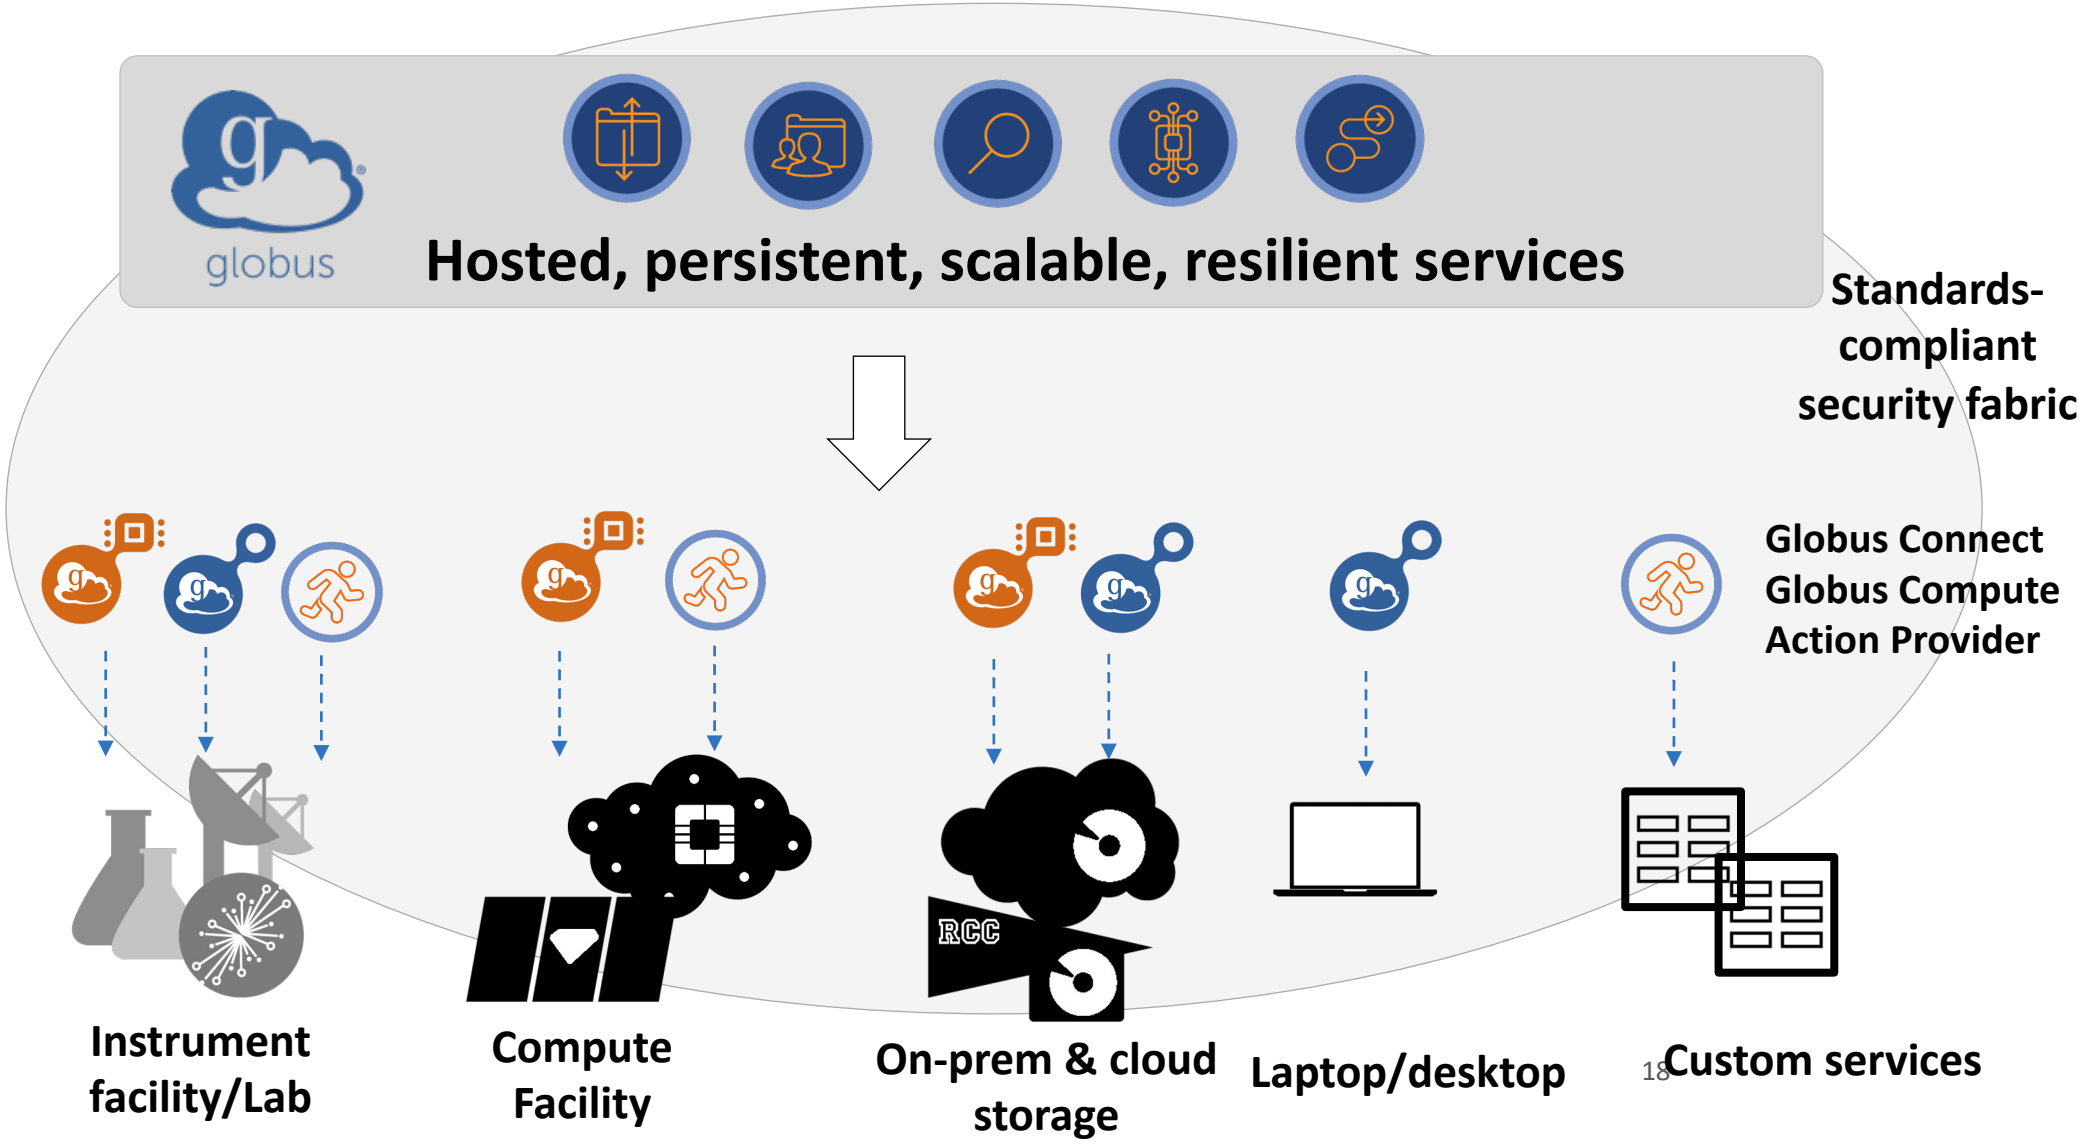

Reliable and managed execution of sets of actions

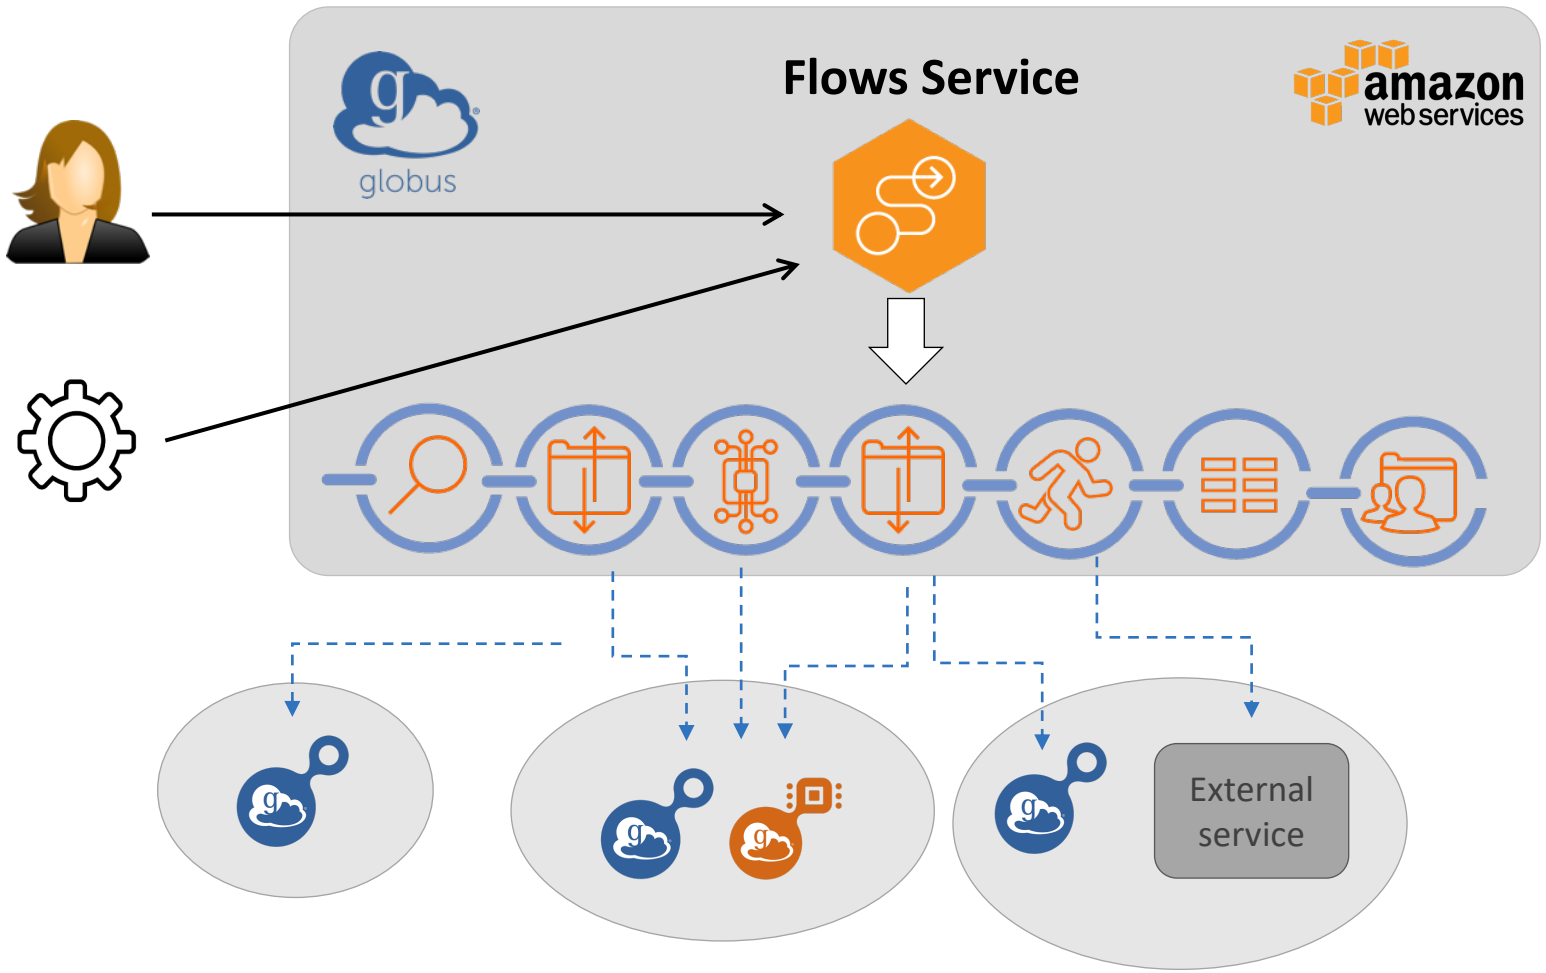

- Declarative language for flow definition
- Interface to integrate external actions
- Simple retry based protocol
- User interface for run, monitor and manage
- Fine-grained delegation for secure access

R. Chard et al., https://doi.org/10.48550/arXiv.2208.09513

Need: 3) Control **who** can perform **what** actions, **when** & **where**

Past approaches:

- Passwords, PKI, Kerberos
- Grid Security Infrastructure
- OAuth,
- Specialized delegation protocols

Challenges: Multi-site, dynamic computing; complexity, usability

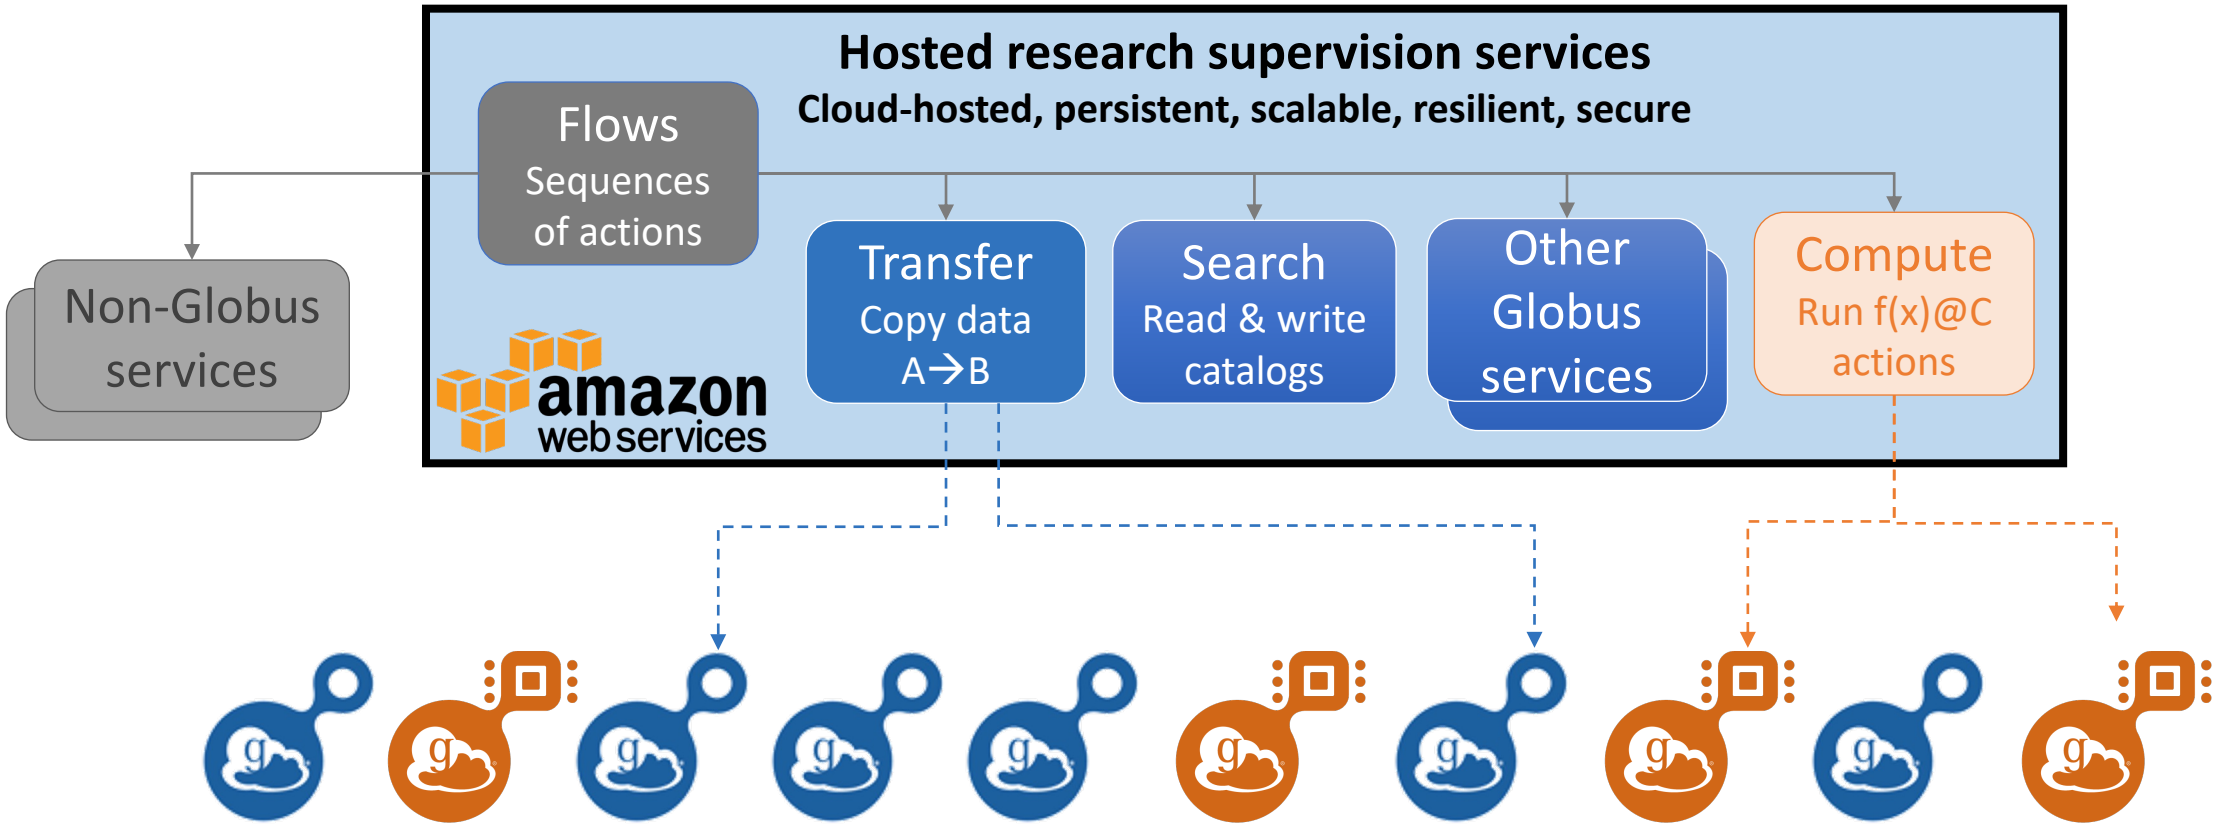

Need: 3) Control **who** can perform **what** actions, **when** & **where**

Federated auth & secure managed delegation

Our approach:

- Secure delegation
- Scoped credentials
- Leverage OAuth2
- Brower compatibility
- IdP federation

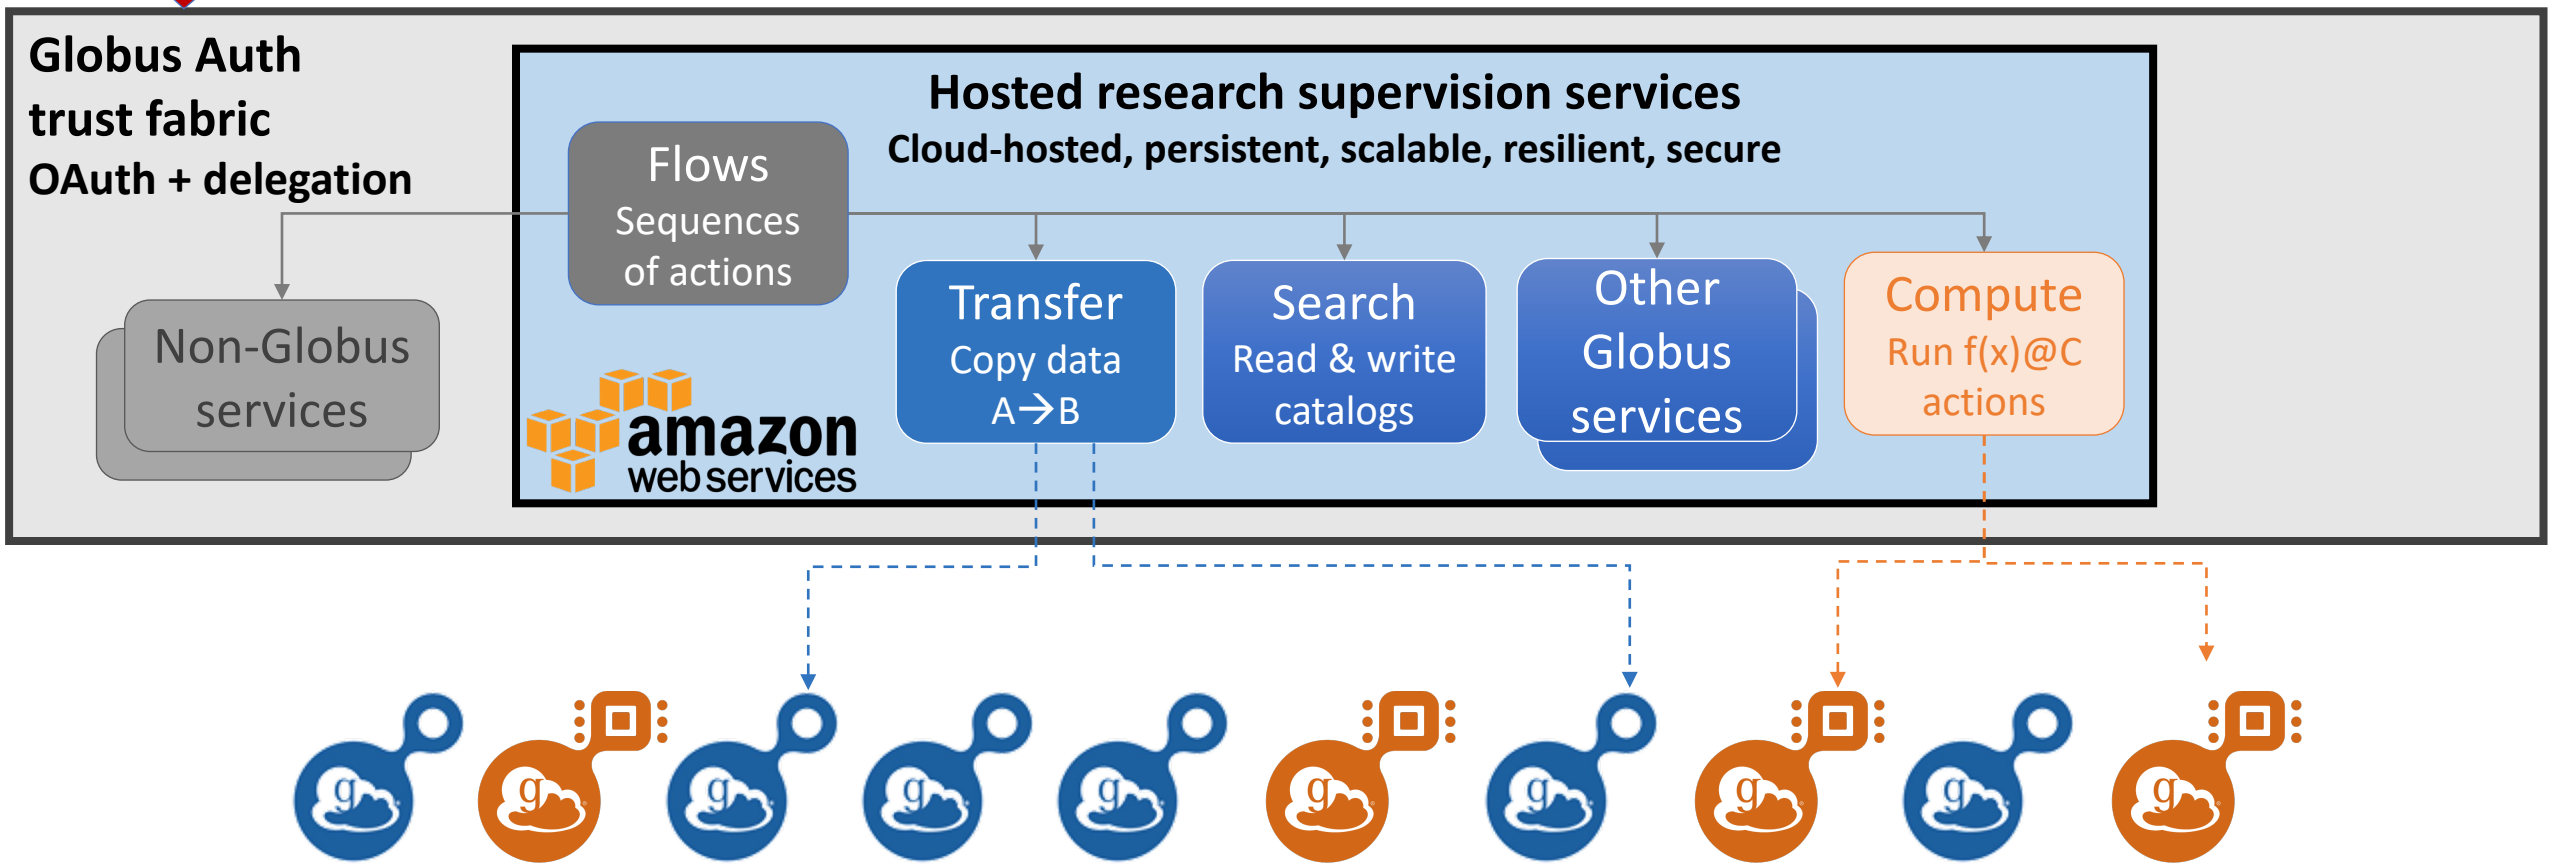

Managed identity & access management (IAM) with fine-grained delegation

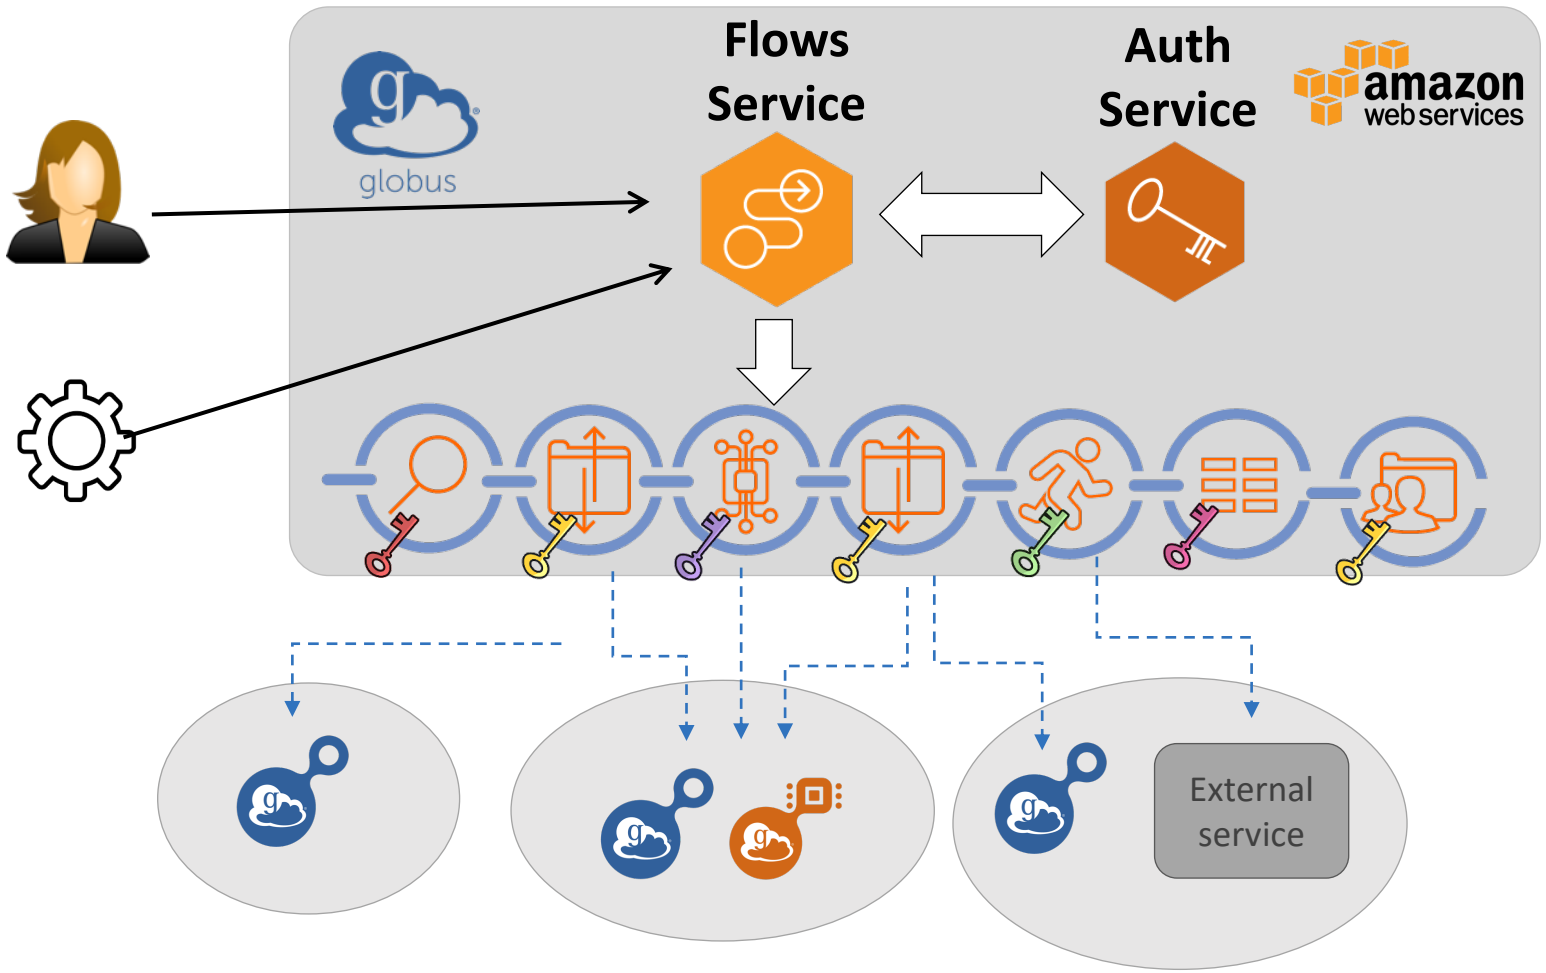

- **Federated authentication**
- **Standards based** (OAuth/OIDC)
- Compatible for use with browsers, command line, automation
- **Fine-grained delegation** via scopes for dependent call chains

1700 identity providers
1.3 B access tokens
2.7 M consents

In total: **Global services** enable low-friction **global science**

1) **Act on resources** regardless of location and interface

→ Widely deployed **local agents**
provide a global footprint for actions

2) Execute remote actions **reliably**

→ Cloud-hosted **management & orchestration services** buffer against inevitable failures

3) Manage <u>who</u> is **trusted** to perform <u>what</u> actions, <u>where</u> and <u>when</u>

→ **Distributed authentication with delegation**
enables secure management of privileges

Leveraging distributed resources

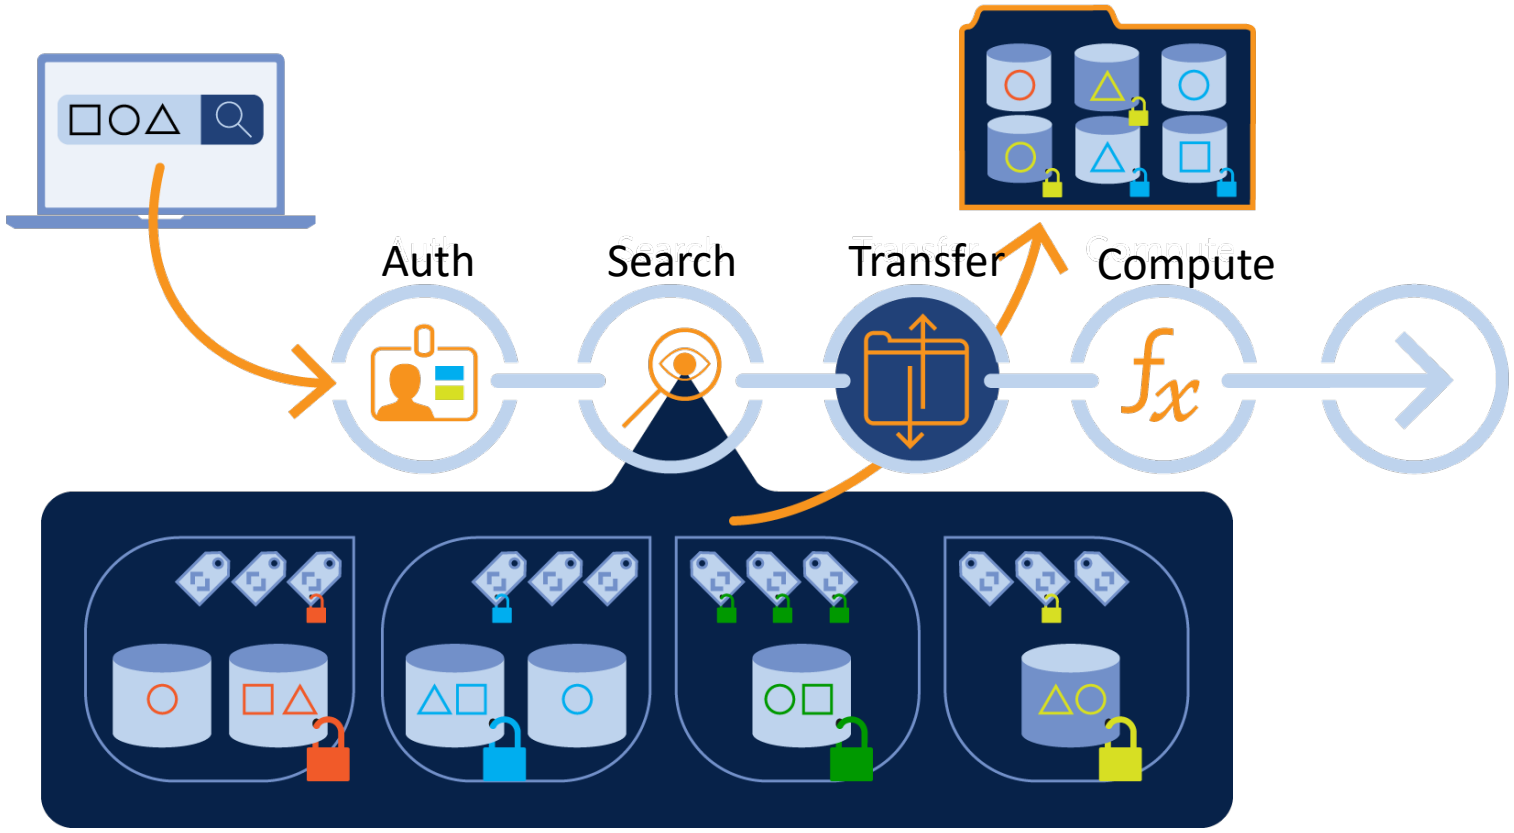

- **Access resources** regardless of location and interface
- Execute remote actions **reliably**
- Execute remote actions **securely**
- **Coordinate** distributed resource use
- **Usability** for end users, and administrators
- …

Globus platform

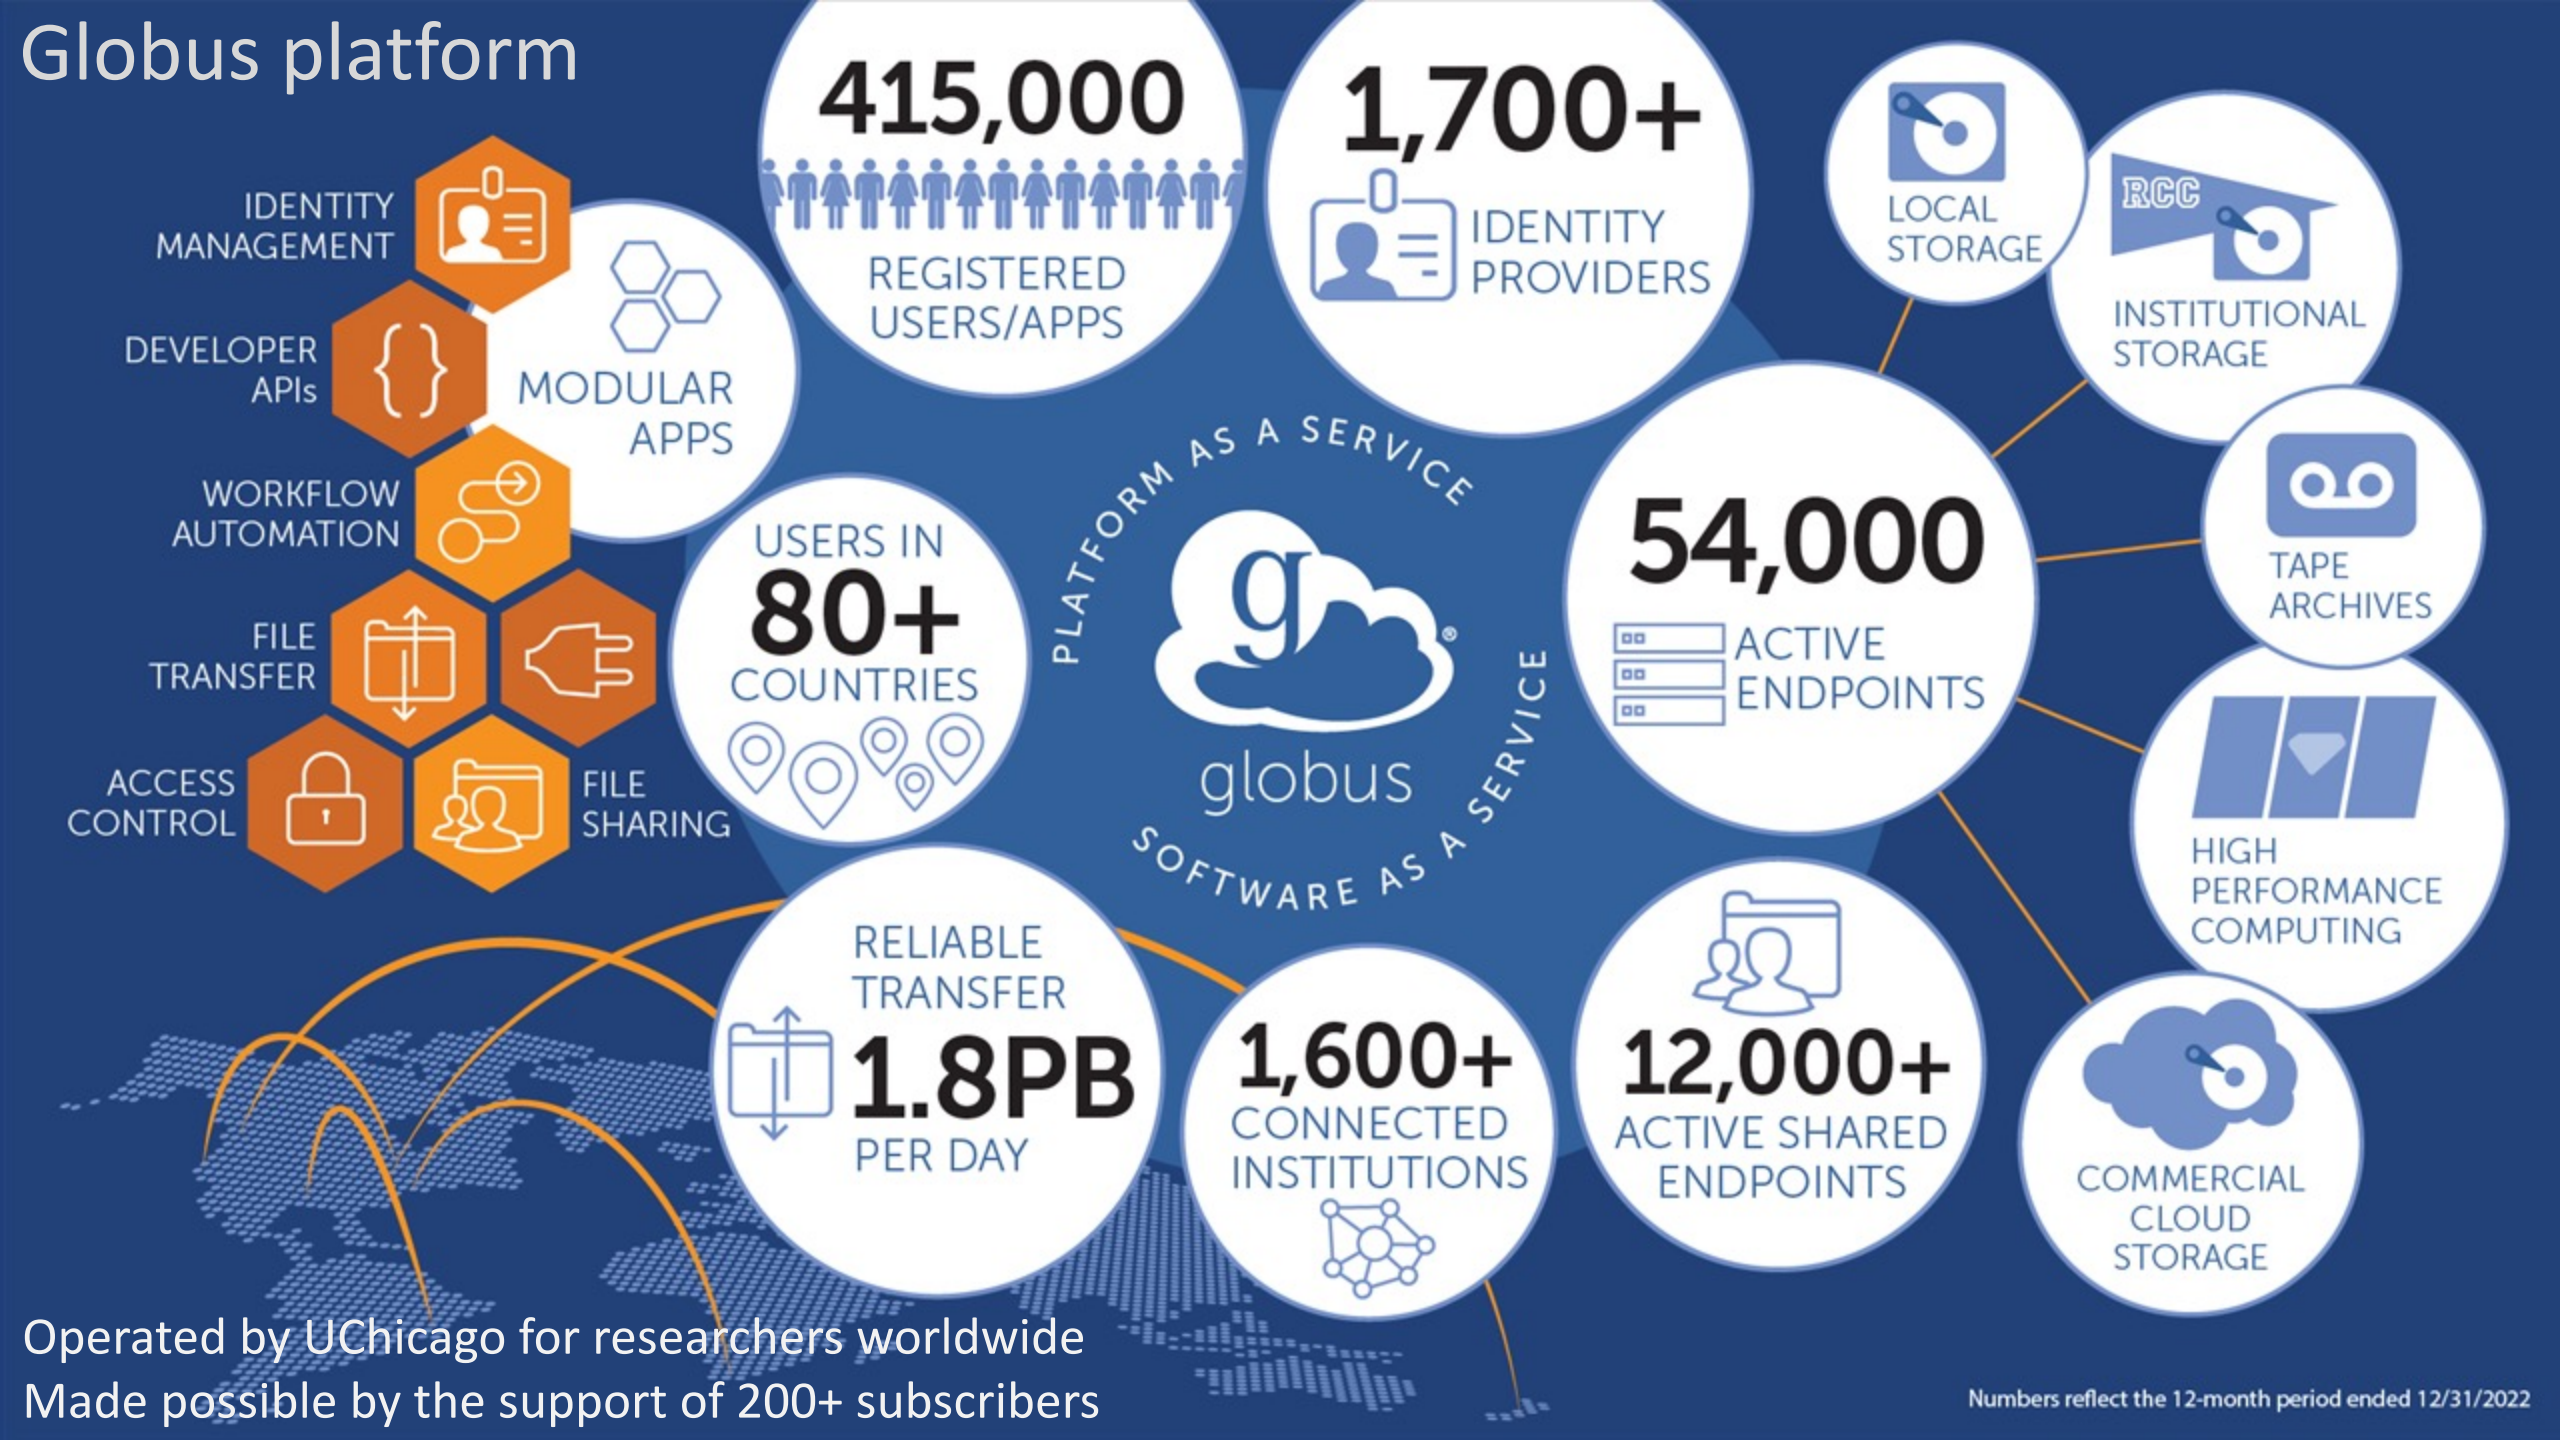

Operated by UChicago for researchers worldwide

Made possible by the support of 200+ subscribers

Numbers reflect the 12-month period ended 12/31/2022

Endpoints at an experimental facility

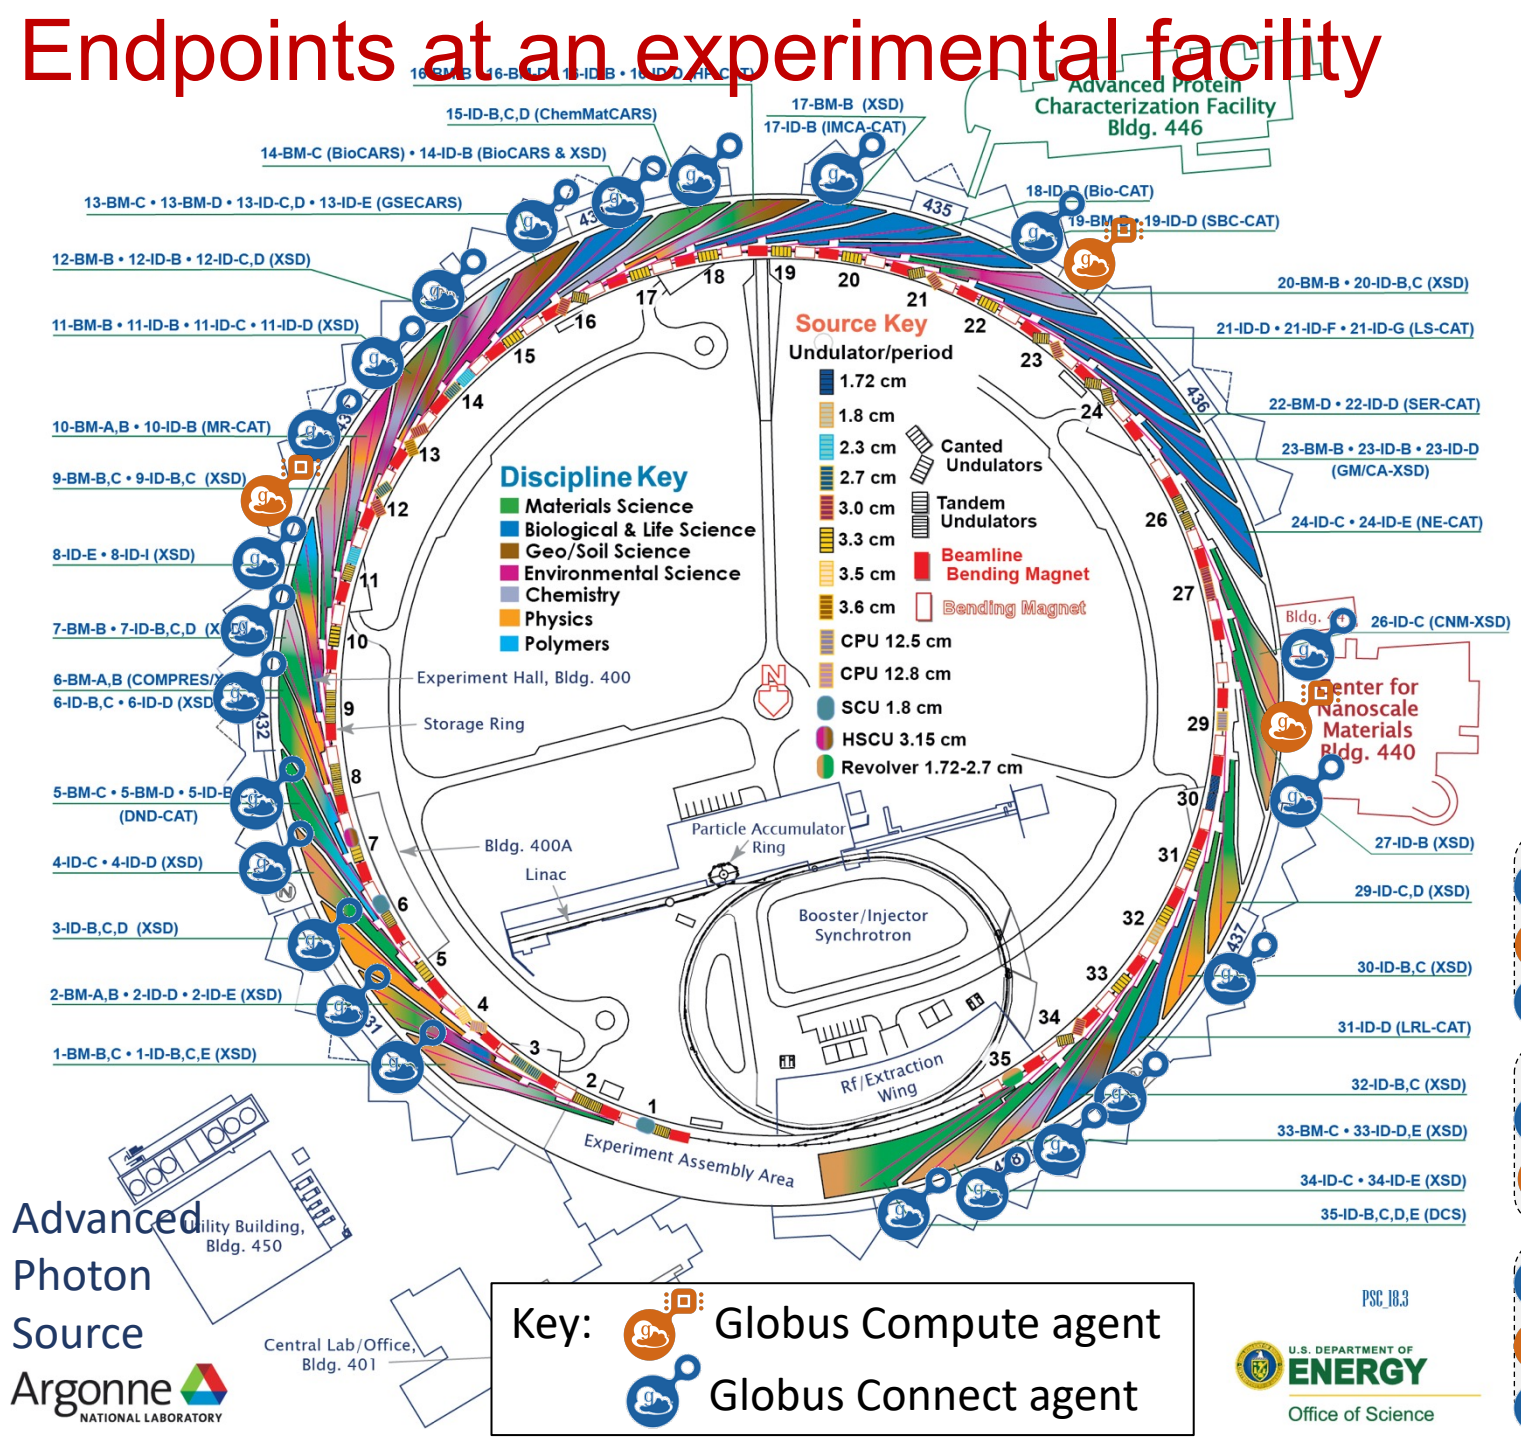

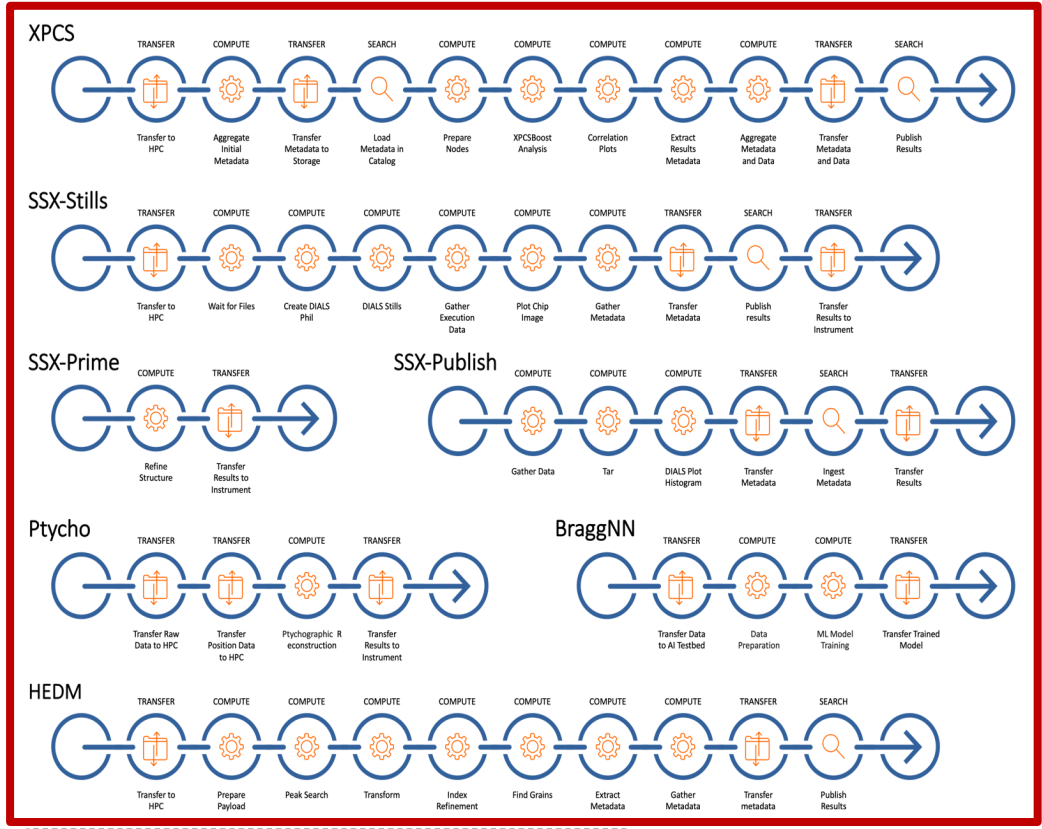

Polaris

Theta

Eagle store

Argonne Leadership Computing Facility

Bebop Cluster

Laboratory Computing Research Center

Orthros Cluster

APS DM system

APS Computing

Globus-accessible storage and computing (10,000s of systems)

Portal server

amazon web services

Transfers performed by the Globus service involving host(s) in:

any location

After this brief title sequence, an animation of all transfers involving any location through 06/2023, as explained below

Monthly plots

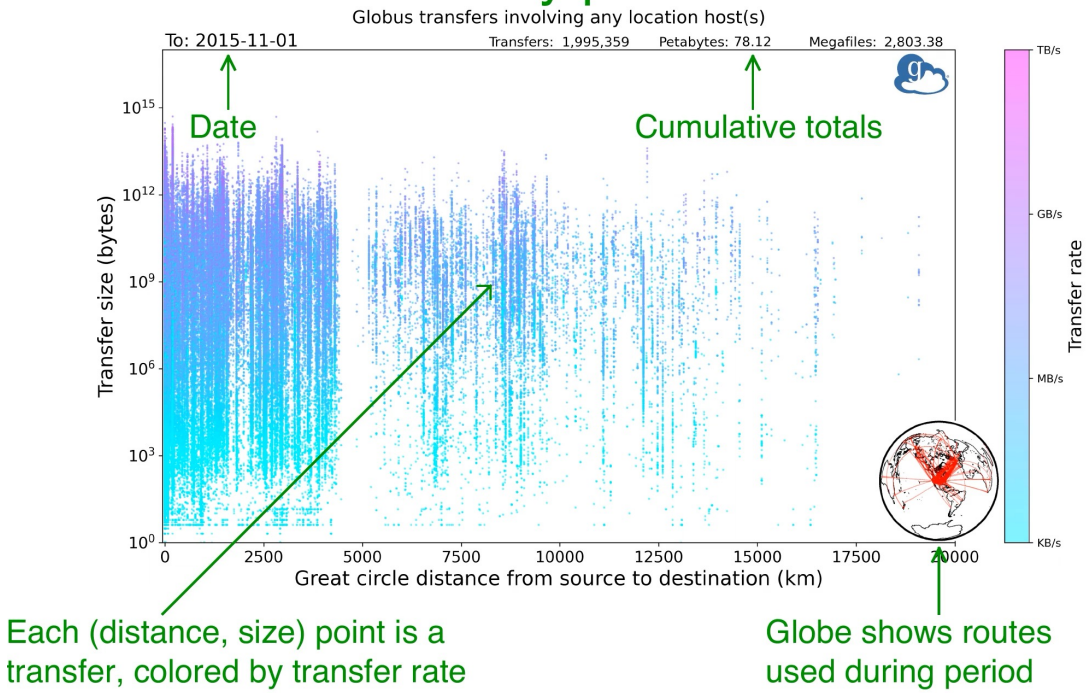

Final aggregate plot

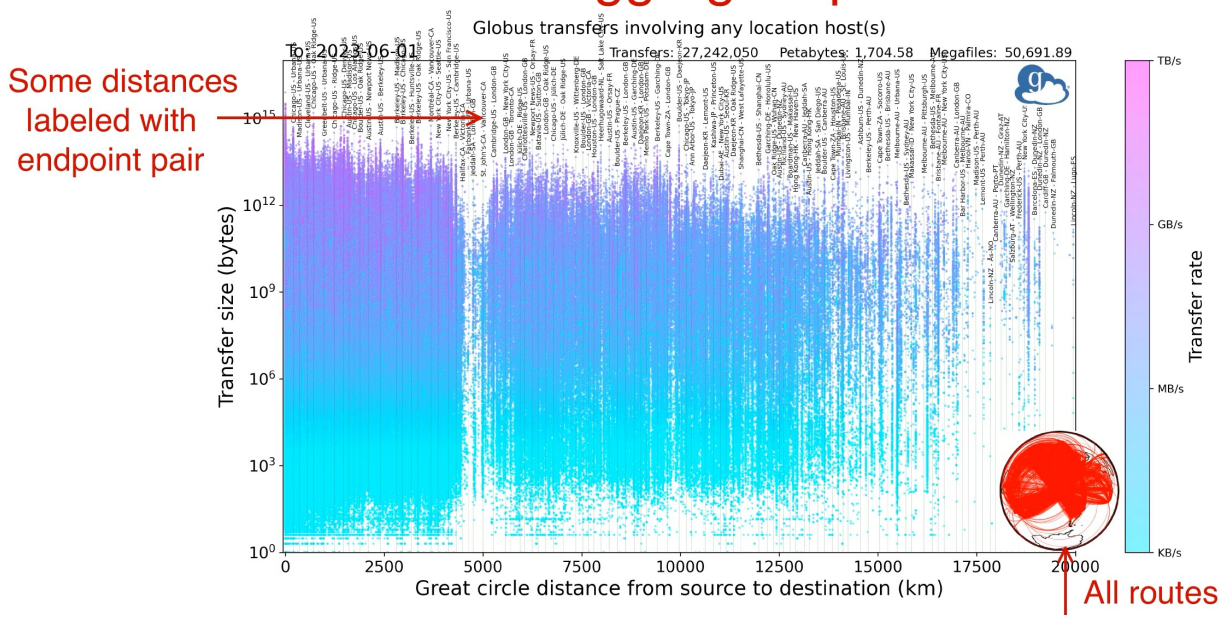

https://globus.org

Globus research data movement service

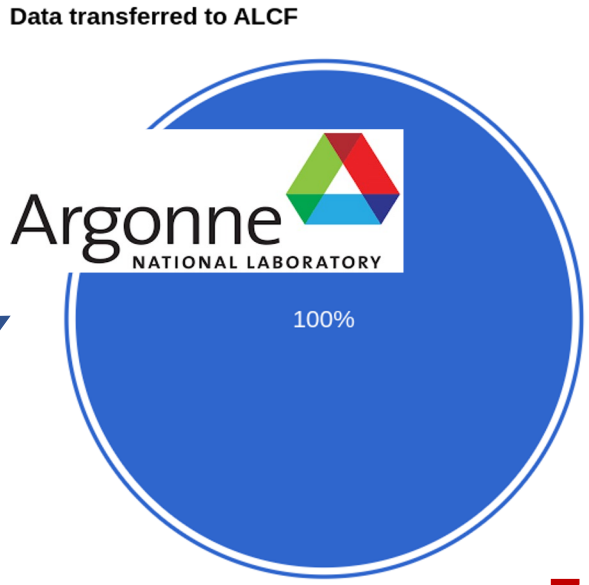

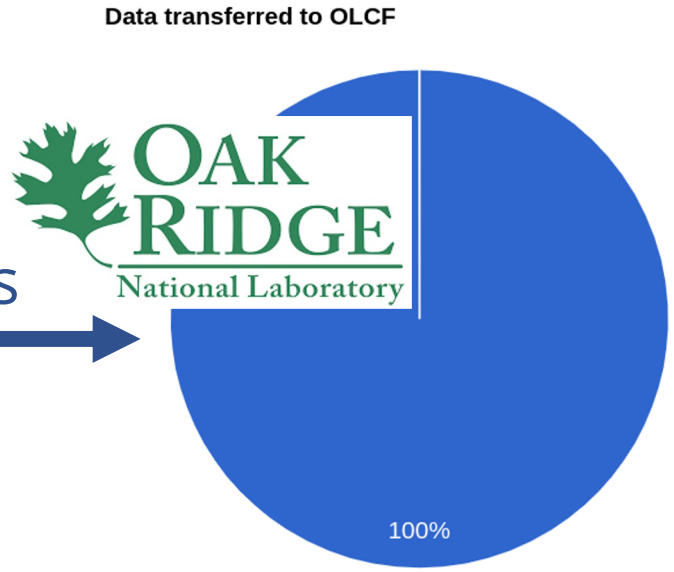

1.5 GB/s

4 to 6 GB/s

7.5 PB transferred between mid-Feb and May 4, ‘22

17,347,671 directories and 28,907,532 files

Replication to ALCF

ACTIVE, PAUSED and the latest SUCCEEDED transfers

No	Datasets	From	Requested	Completed	Status	Directories	Files	Bytes Transferred	Faults	Rate
1	/cmip5_css01_data/cmip5/output1/NSF-DOE-NCAR/CESM1-CAM5	LLNL	2022-05-03 08:46:03	2022-05-04 11:37:43	SUCCEEDED	7208	13540	29913341340	16	309 kB/s
2	/cmip5_css02_data/cmip5/output1/NCC/NorESM1-M	LLNL	2022-05-02 09:52:03	2022-05-02 11:31:27	SUCCEEDED	4017	7548	5367692747060	0	900 MB/s
3	/cmip5_css02_data/cmip5/output1/NCAR/CCSM4	LLNL	2022-05-02 01:53:03	2022-05-03 00:50:23	SUCCEEDED	52571	48925	33455438769668	11	405 MB/s
4	/cmip5_css02_data/cmip5/output1/NASA-GISS/GISS-E2-R-CC	LLNL	2022-05-02 01:28:03	2022-05-02 01:52:31	SUCCEEDED	2098	9576	1087745609416	0	741 MB/s
5	/cmip5_css02_data/cmip5/output1/NASA-GISS/GISS-E2-R	LLNL	2022-05-02 00:42:03	2022-05-02 09:51:16	SUCCEEDED	30164	132059	24482369232188	5	743 MB/s

Replication to OLCF

ACTIVE, PAUSED and the latest SUCCEEDED transfers

No	Datasets	From	Requested	Completed	Status	Directories	Files	Bytes Transferred	Faults	Rate
1	/cmip5_css01_data/cmip5/output1/NSF-DOE-NCAR/CESM1-CAM5	LLNL	2022-05-03 08:47:18	2022-05-04 11:41:11	SUCCEEDED	7208	13540	271068730	16	2.80 kB/s
2	/cmip5_css02_data/cmip5/output1/NCAR/CCSM4	LLNL	2022-05-02 13:58:03	2022-05-03 03:14:27	SUCCEEDED	52571	48925	33455438769668	1	700 MB/s
3	/cmip5_css02_data/cmip5/output1/NCC/NorESM1-M	ALCF	2022-05-02 11:32:03	2022-05-02 12:15:48	SUCCEEDED	4017	7548	5367692747060	0	2.04 GB/s
4	/cmip5_css02_data/cmip5/output1/NASA-GISS/GISS-E2-R	ALCF	2022-05-02 09:52:03	2022-05-02 12:30:08	SUCCEEDED	30164	132059	24482369232188	3	2.58 GB/s
5	/cmip5_css02_data/cmip5/output1/NASA-GISS/GISS-E2-R-CC	ALCF	2022-05-02 05:34:04	2022-05-02 05:44:32	SUCCEEDED	2098	9576	1087745609416	0	1.73 GB/s

https://dashboard.globus.org/esgf

As of May 4, 2022

Cumulative data transferred over time

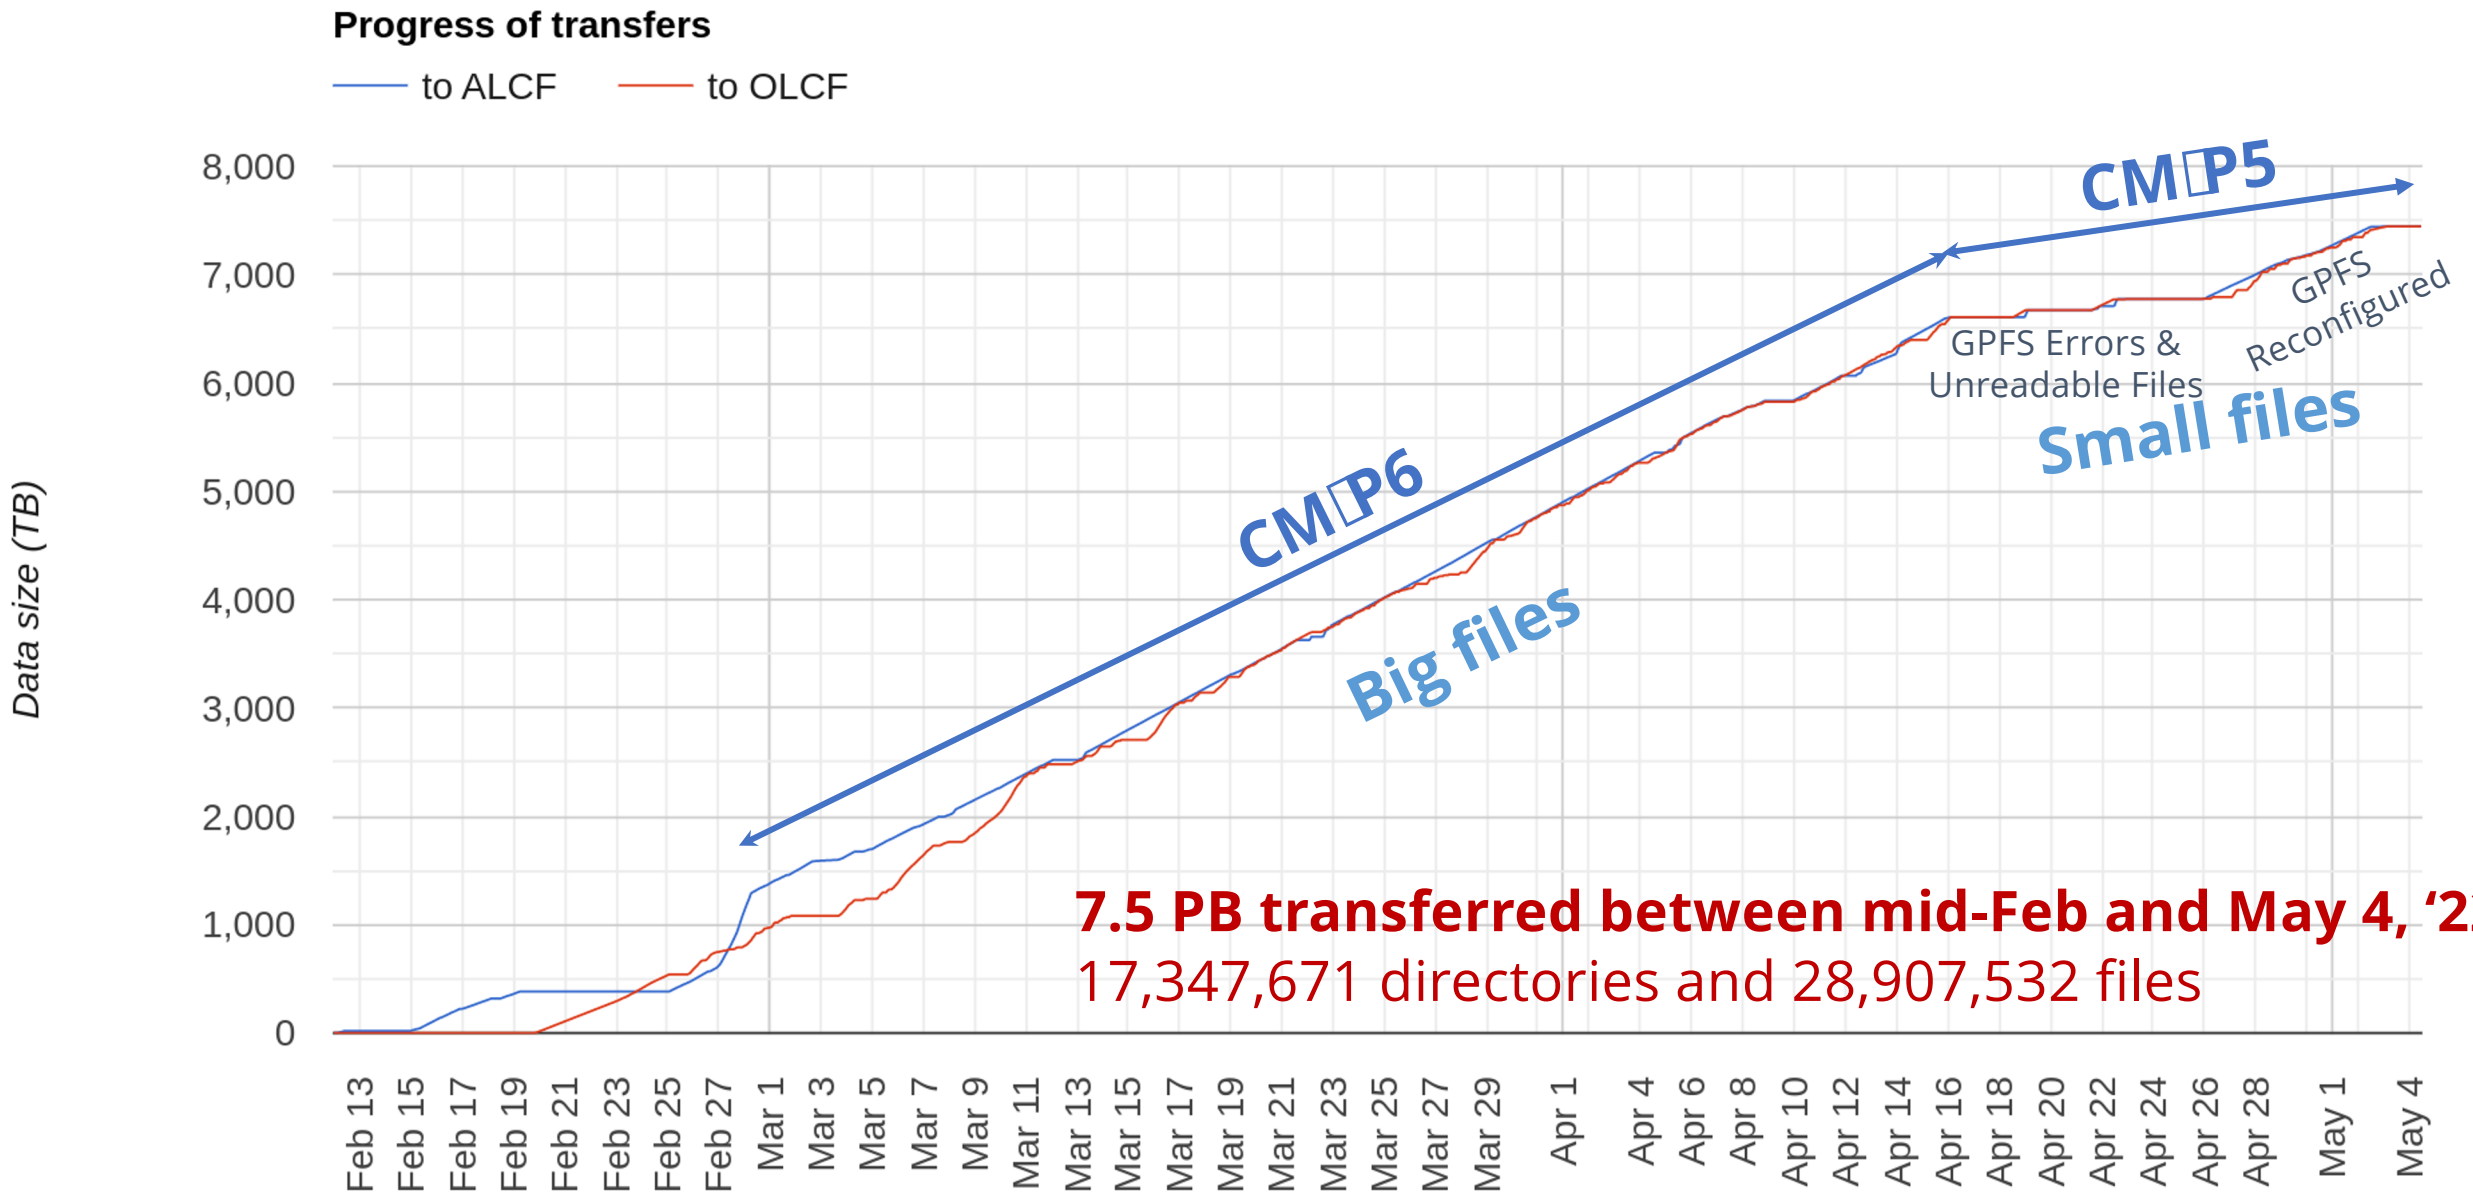

Globus Compute: A hosted research supervision service that implements a **universal computing fabric**

```
def F(in_args):
  # do something
  return results

gcc.register_function(F)
```

Register functions

Globus Compute

Run functions

```
F(ep, "A")
```

```
f = gcc.run("A",
        endpoint_id=ep,
        function_id=F)

R = gcc.get_result(f)
```

Deploy Globus Compute agents


```
$ pip install globus-compute-endpoint
$ globus-compute-endpoint configure myep
$ globus-compute-endpoint start myep
```

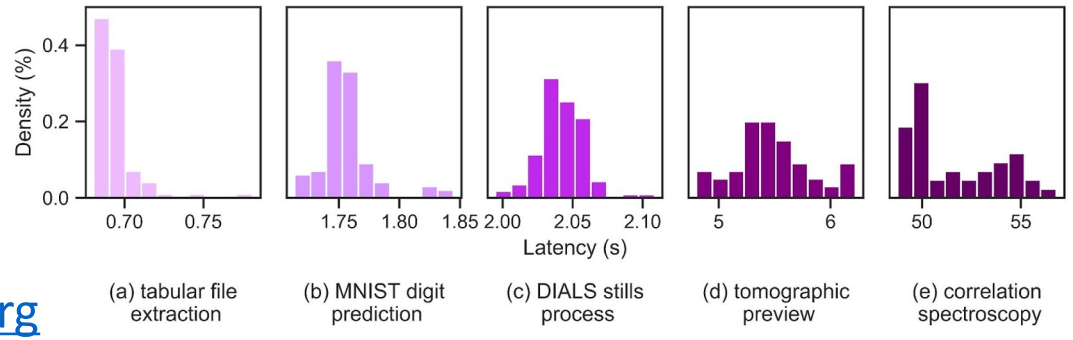

Z. Li et al., https://doi.org/10.48550/arXiv.2209.11631 https://funcx.org

Globus Compute application

Privacy preserving federated learning

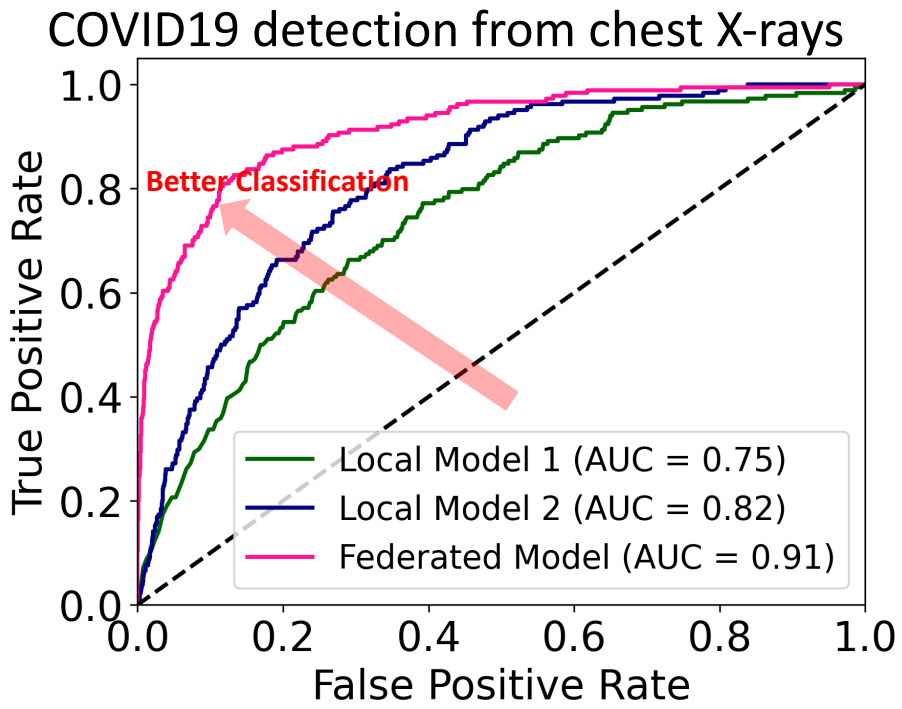


Ravi Madduri


Kibaek Kim


Ryu Minseok


Eliu Huerta

Pradeep Natarajan


Puneet Batra


Maryellen Giger


Sam Armato

Privacy Preserving Federated Learning as a Service DASHBOARD | PROFILE | IAN@GLOBUSID.ORG | LOGOUT

Federation Information

Endpoint Information

Client	Organization	Email	Endpoint Status
Jan F Nygård	Cancer Registry of Norway	jfn@kreftregisteret.no	
Severin Langberg	Cancer Registry of Norway	Langberg91@gmail.com	
Zilinghan Li (You)	University of Illinois	1250976113@qq.com	
Zilinghan Li - NCSA	National Center for Supercomputing Applications	zl52@illinois.edu	
Ravi Madduri	Argonne	madduri@anl.gov	
Marcus Klarqvist	broad institute of mit and harvard	mklarqvi@broadinstitute.org	
Jordan Fuhrman	The University of Chicago	jdfuhrman@uchicago.edu	

Experiment Information

Experiment Name	Experiment ID	Status	Config	Log	Report	Tensorboard
MNIST1	5a525a61353a4a5a82b3ee895773eedf	DONE				
MNIST2	4cf1ee4409b04b89bbc1b2f0f76969b1	DONE				
MNIST3	75474c0d2bbe4c2481e766b1166b6672	DONE				
MNIST4	be5eb91f8e9e4e8ca647f061b52ccb93	DONE				
Ravi_Demo	23e0bc6faf234130a4a99917e759b928	DONE				
MNIST5	de7ff6bb6d2a42bbaf205158e22bdbfa	DONE				
Demo_Polaris	57d2605794d744f6b7dd08147cafb3c6	DONE				
Demo_Polaris_New	922ddcfe9ecf4ad2b912a5eb14cf720f	DONE				
Ravi_Demo_Latest	7151875c342747169a6707af62ebf21d	DONE				
Final_Demo	4e4432e25b2d4eb6ab4cf3f5c1c86d87	DONE				
Ravi_Demo1	06f501225b694a459b3591fec6b69e23	DONE				
MNIST-Report-Demo	fec4ff7c793e4027bb223d1fe5ab7e97	DONE				
MNIST-Report-Demo2	dfd328dc940346ea87cd4f68a2600773	DONE				
MNIST-Demo	27e6ad17a07d4d3f83385e7660078895	DONE				

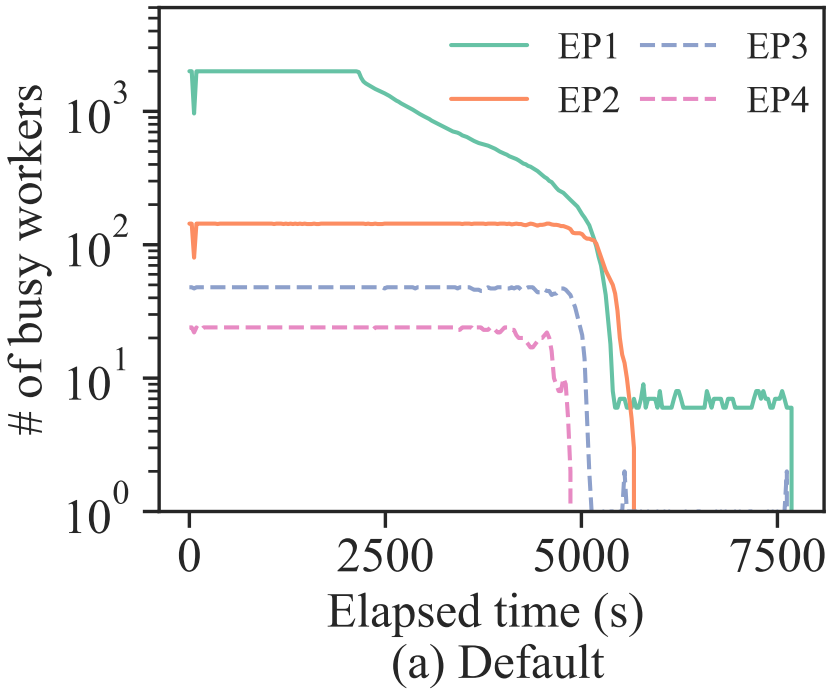

(a) Default

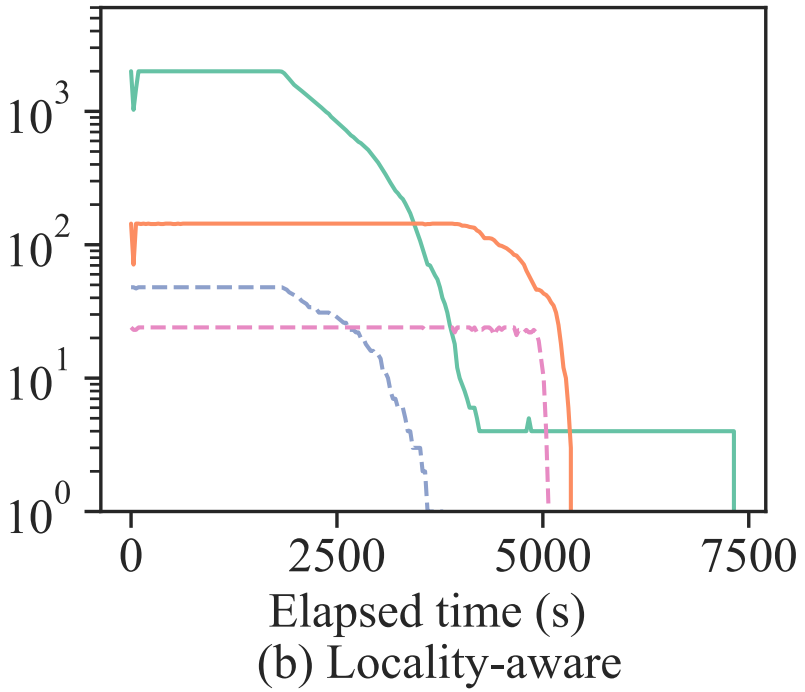

(b) Locality-aware

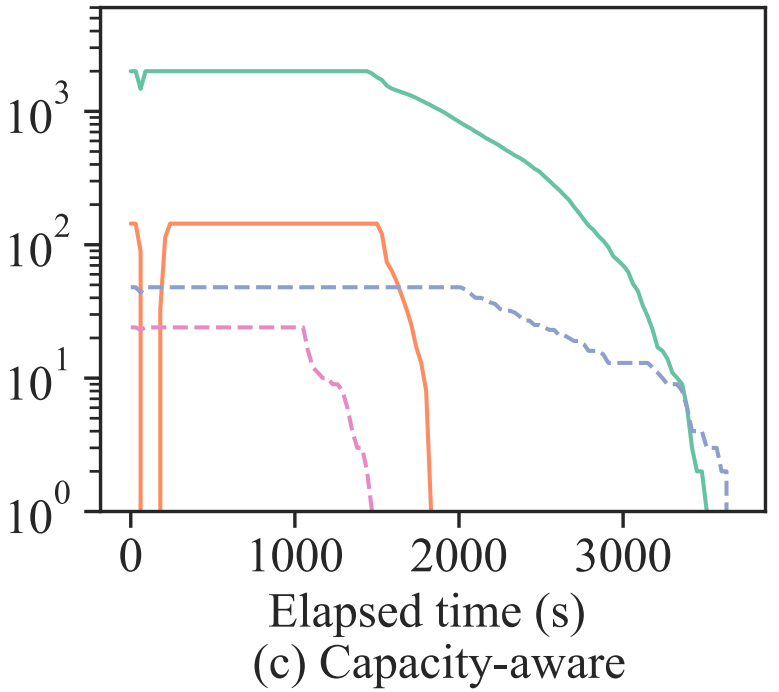

(c) Capacity-aware

Globus Compute application

30,000-task drug-screening runs over multiple HPC systems, with different scheduling heuristics

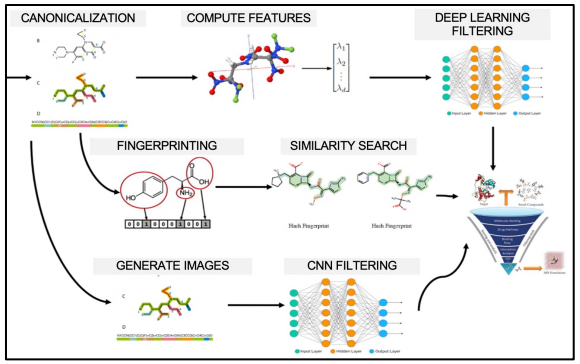

Z. Li et al., 2023

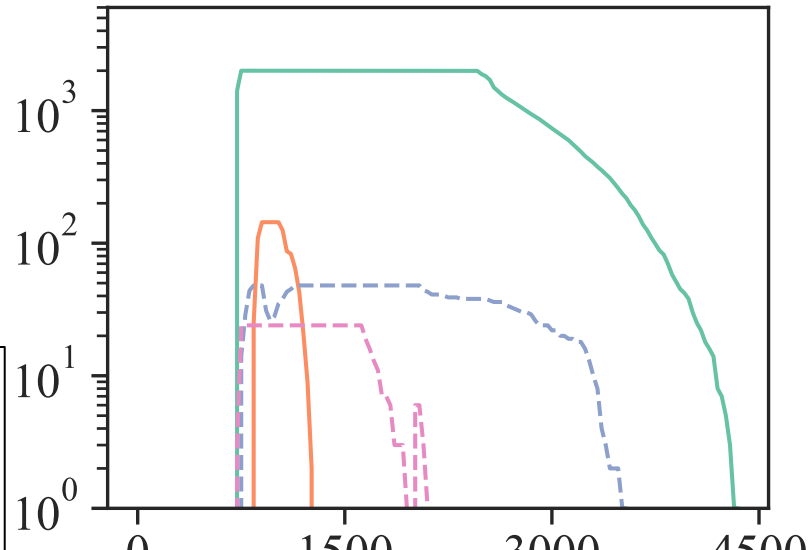

(d) Heterogeneity-aware

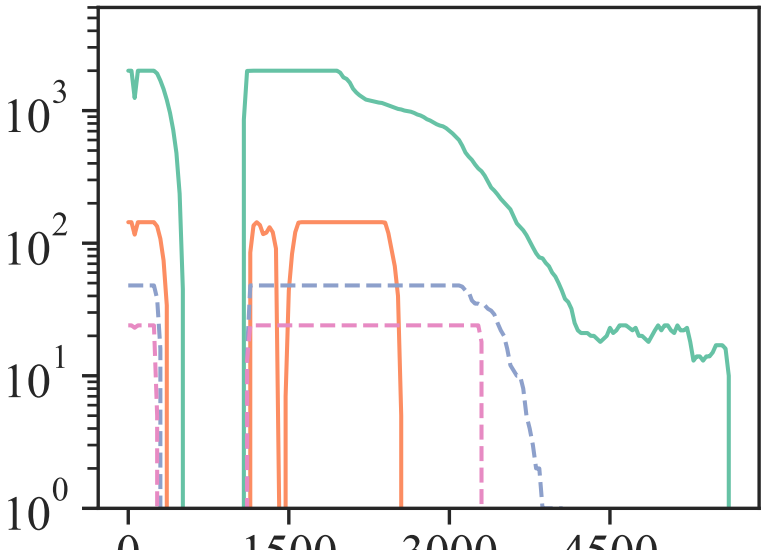

(e) Dynamic-adjusting

Example: Multi-facility climate data analysis

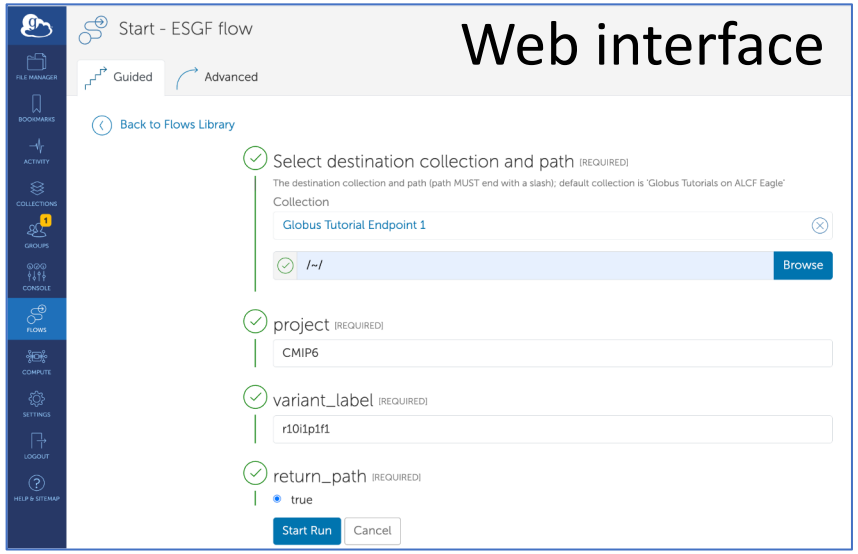

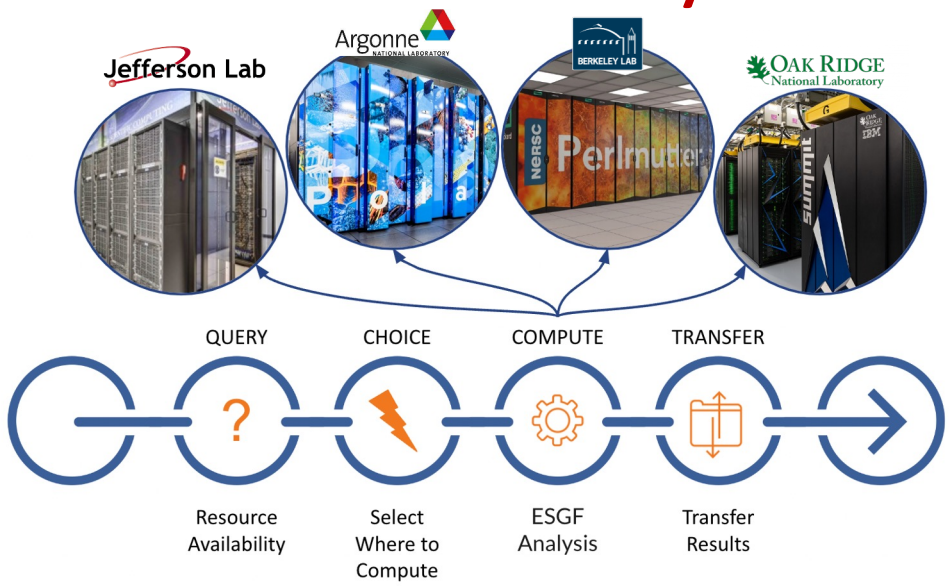

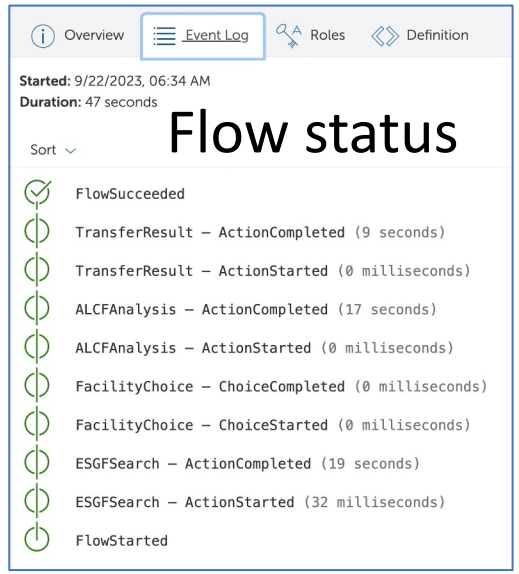

- **Analyze large climate datasets where data and compute are available, and return results**
- **Globus Flows** provides managed automation of the analysis pipeline, **Globus Auth** for A&A
- Web interface for user to specify inputs, start analysis, and monitor progress
- Flow selects compute location based on both resource availability and data locality
- **Globus Compute** federates and standardizes the interface to submit tasks to any resource with endpoint deployed, including JLab iFarm, ALCF Polaris, NERSC Perlmutter, OLCF Summit
- Results deposited at specified **Globus collection**

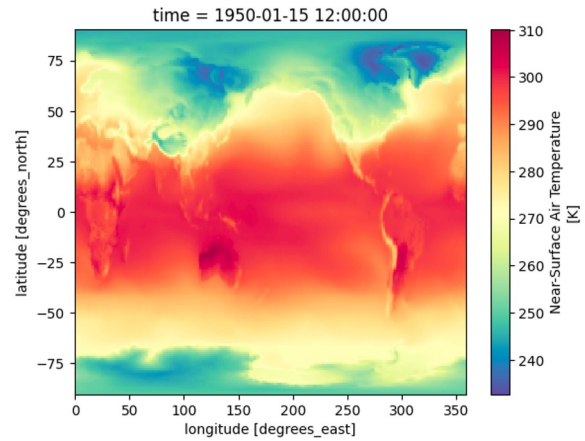

Argonne: Ian Foster, Mike Papka, Max Grover, Scott Collis, Tom Uram, Christine Simpson, Bill Allcock, Benoit Cote, Ryan Chard
JLab: Amitoj Singh, Xinxin Mei, Chris Larrieu, Bryan Hess
ORNL: Forrest Hoffman, Nathan Collier
UChicago/Argonne: Rachana Ananthakrishnan, Kyle Chard, Nick Saint

Flows enable creation of smart instruments

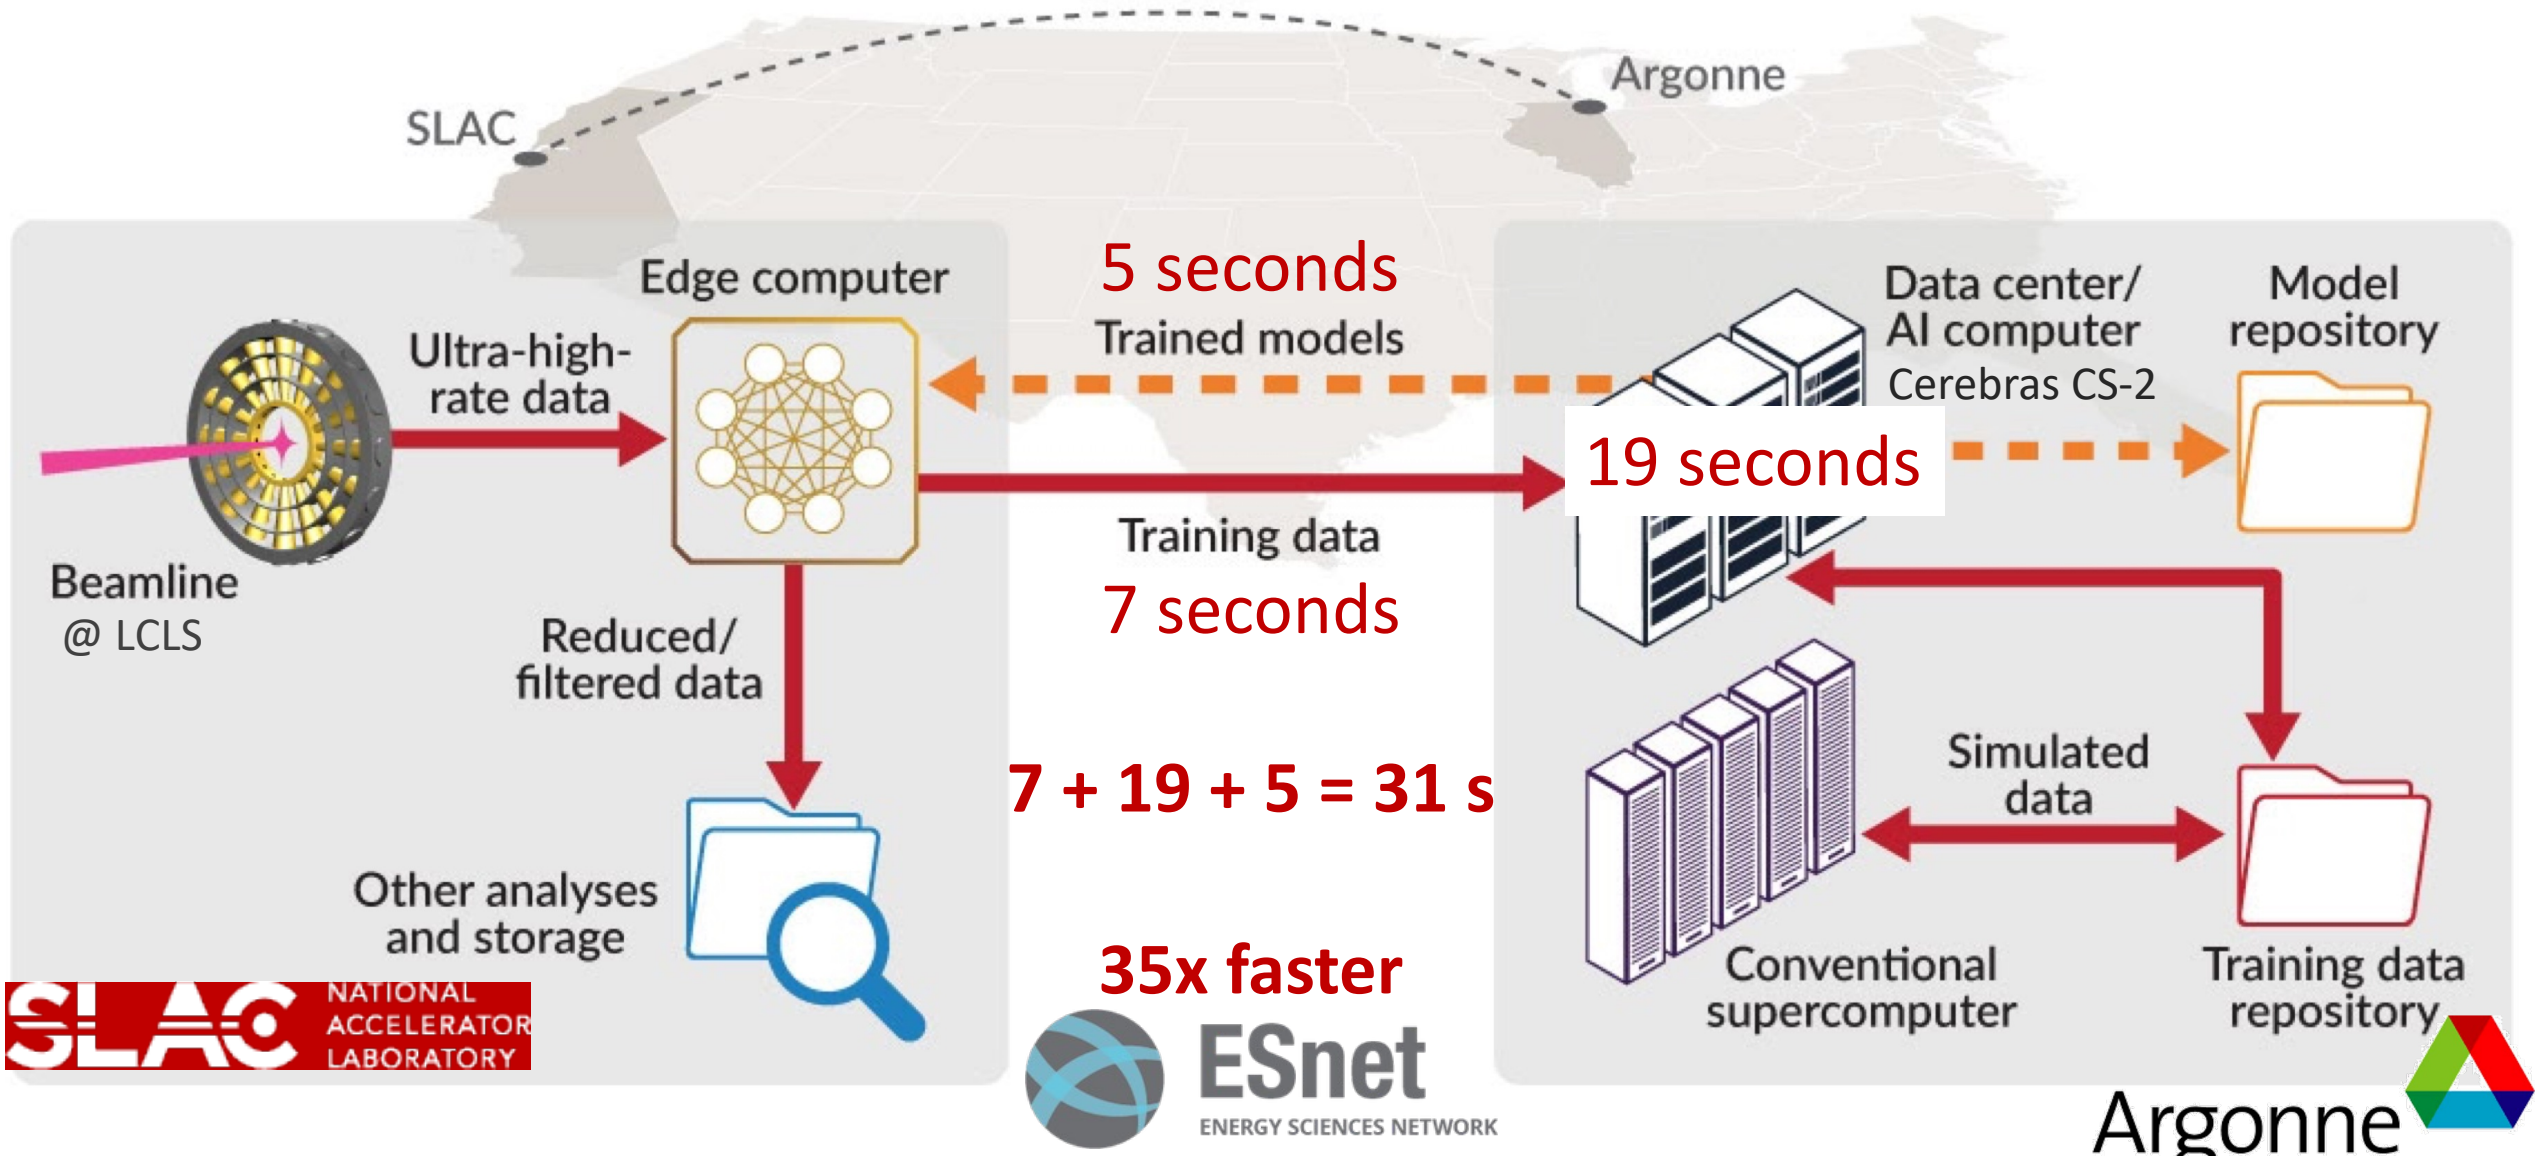

Z. Liu et al., https://doi.org/10.48550/arXiv.2105.13967

Flows enable creation of smart instruments

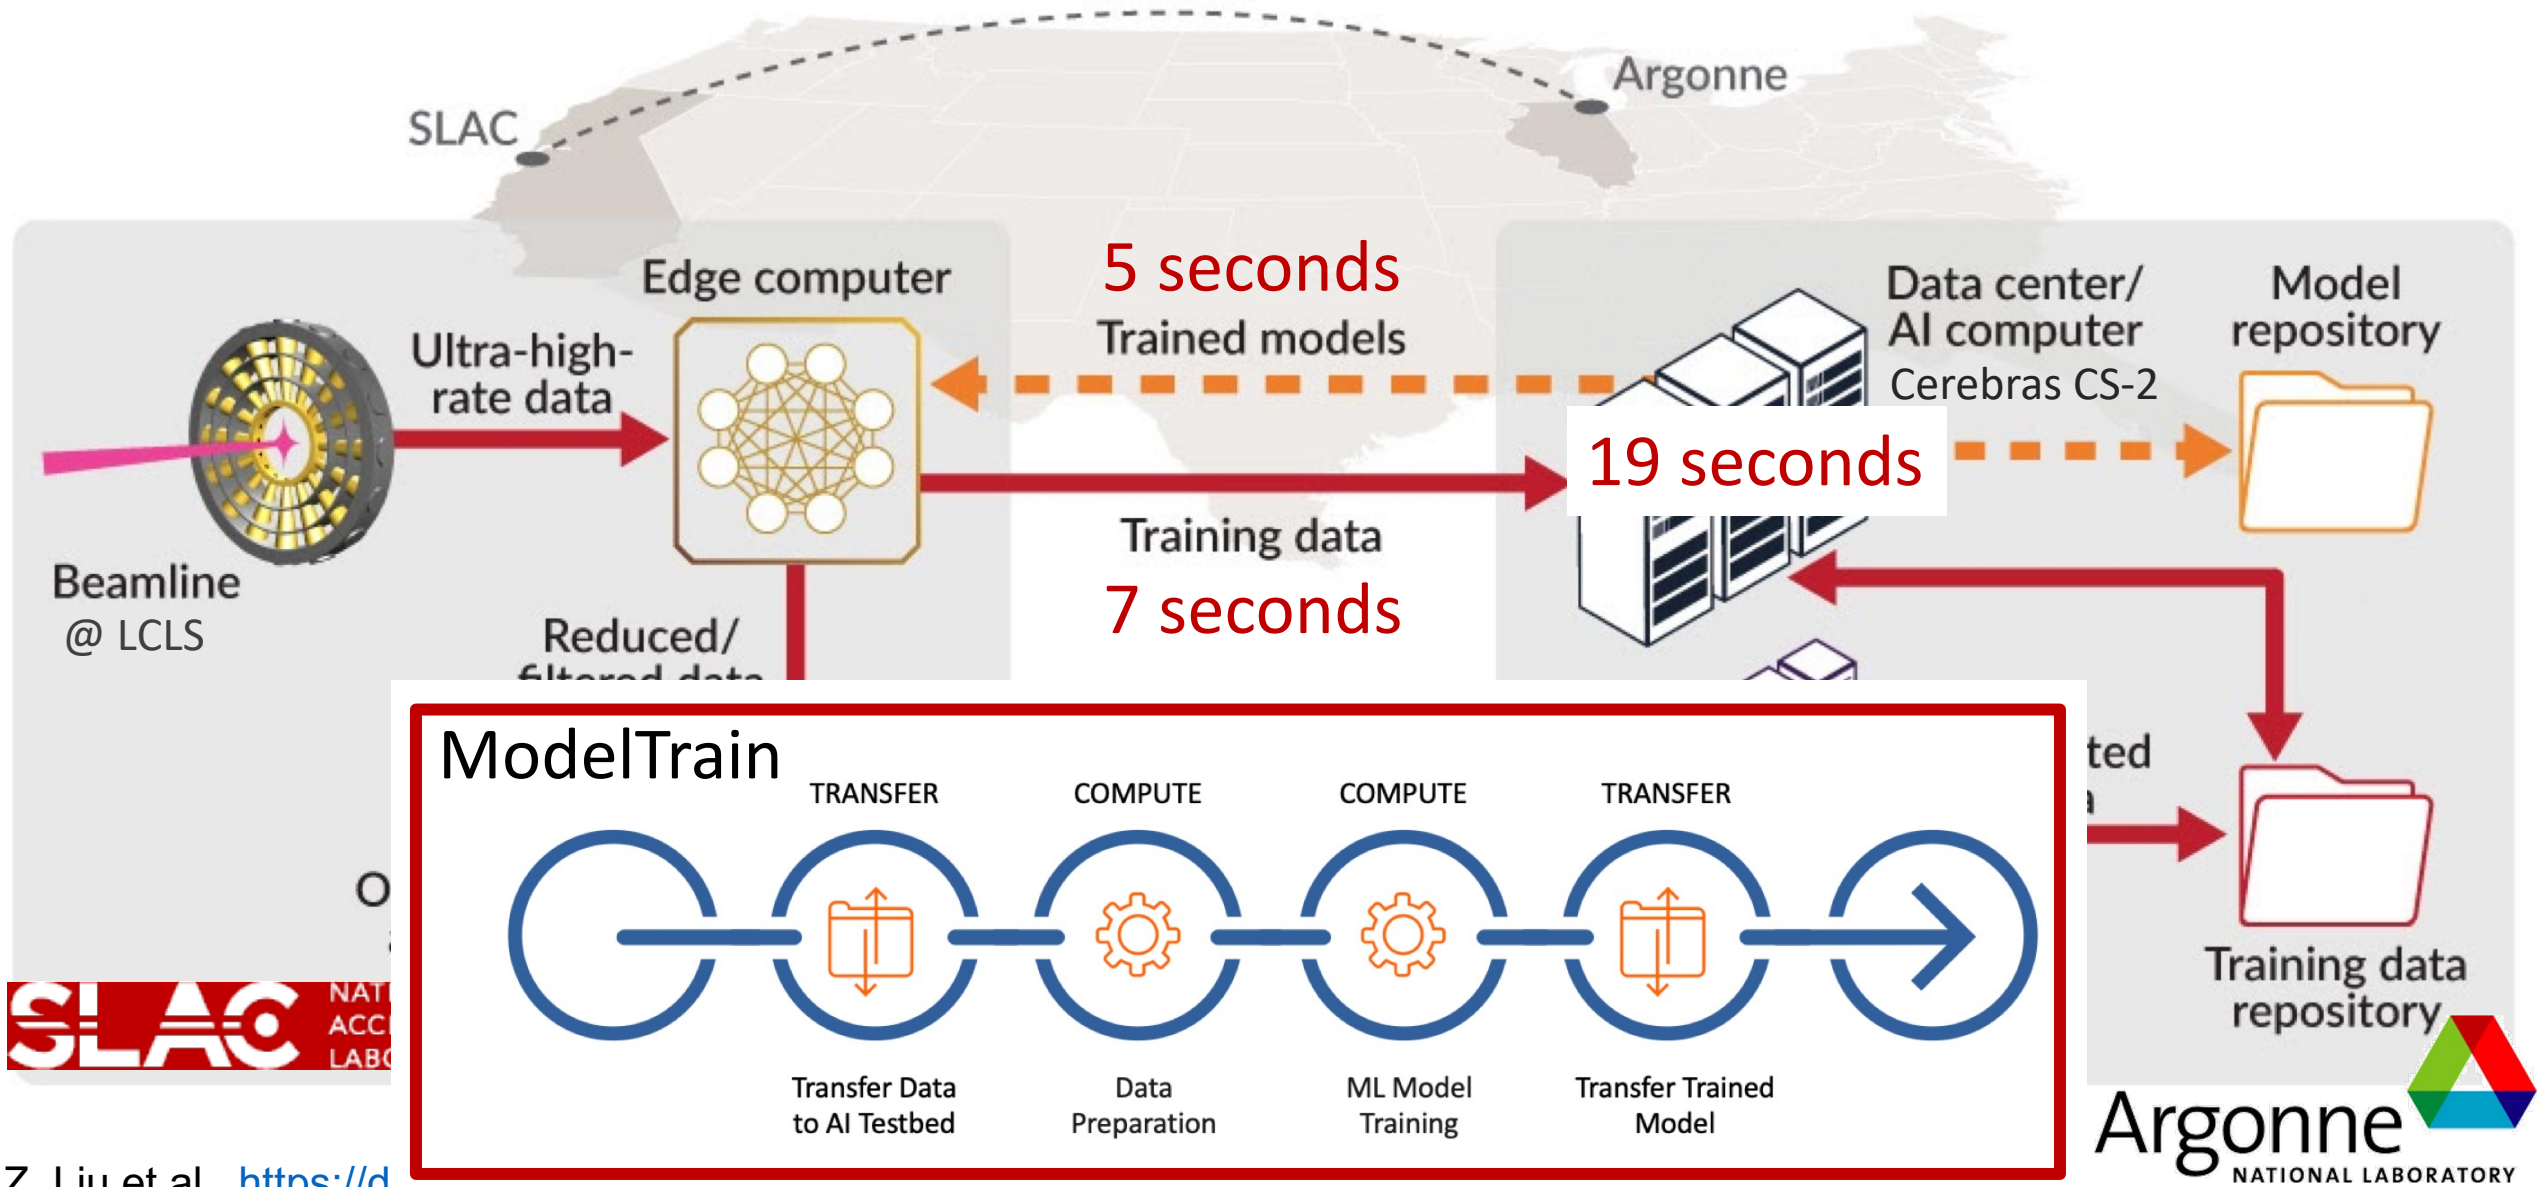

Z. Liu et al., https://d

New applications mean new computing workloads

Globus Flows can invoke arbitrary functions via **Globus Compute**

Functions may be executed in various locations: at a beamline, local server, cluster, cloud

Seven flows in use at the Advanced Photon Source

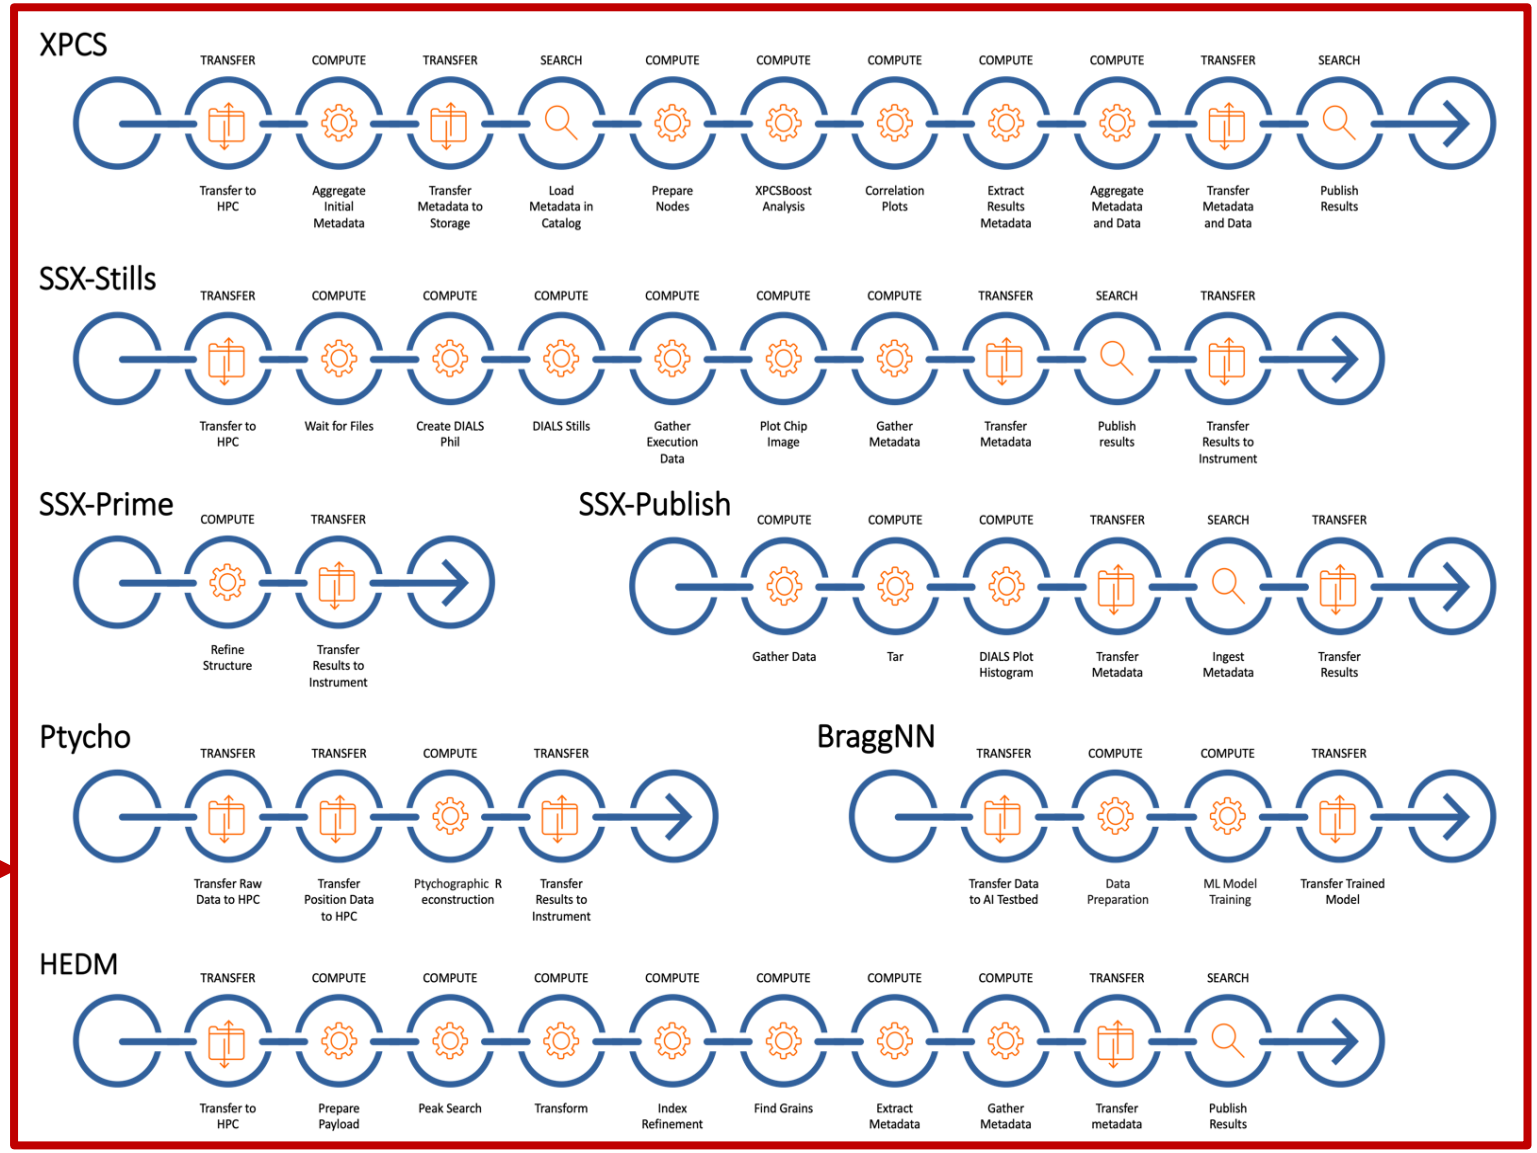

R. Vescovi et al., https://doi.org/10.1016/j.patter.2022.100606

New applications mean new computing workloads

Globus Flows can invoke arbitrary functions via **Globus Compute**

Functions may be executed in various locations: at a beamline, local server, cluster, cloud

Execution times at the Argonne Leadership Computing Facility

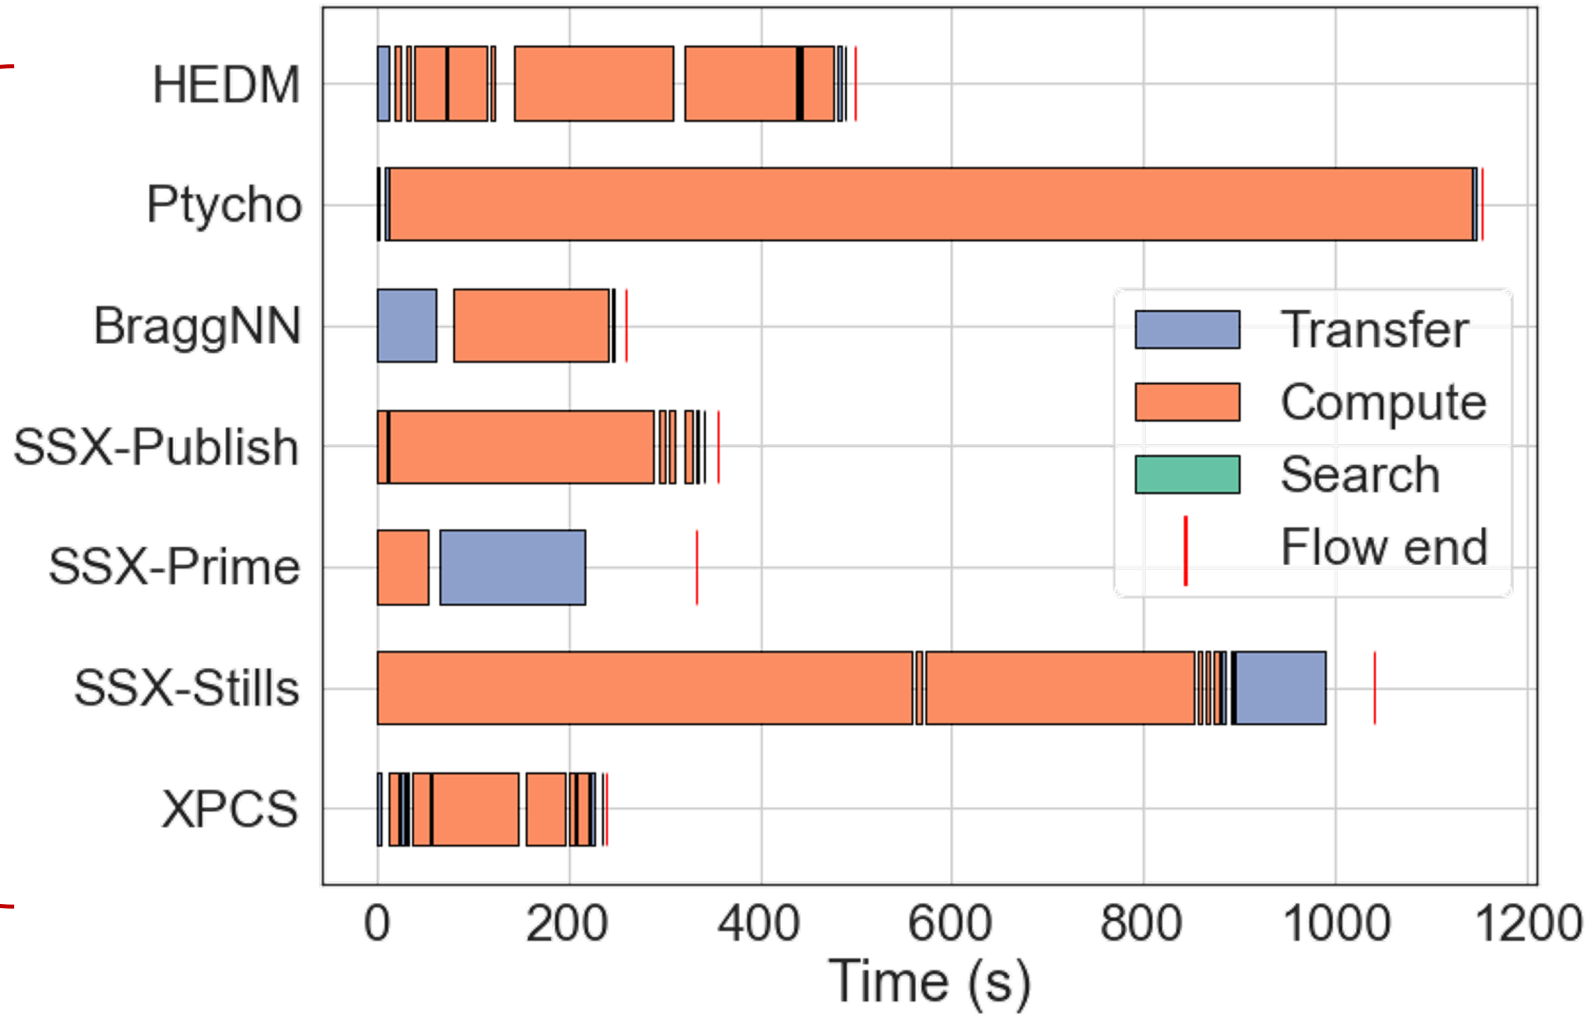

R. Vescovi et al., https://doi.org/10.1016/j.patter.2022.100606

We are applying these methods to construct science factories

Instrument-specific action providers abstract instrument details; domain protocols compiled to robotic commands

Workflow Execution Interface (WEI) architecture

arxiv.org/abs/2308.09793

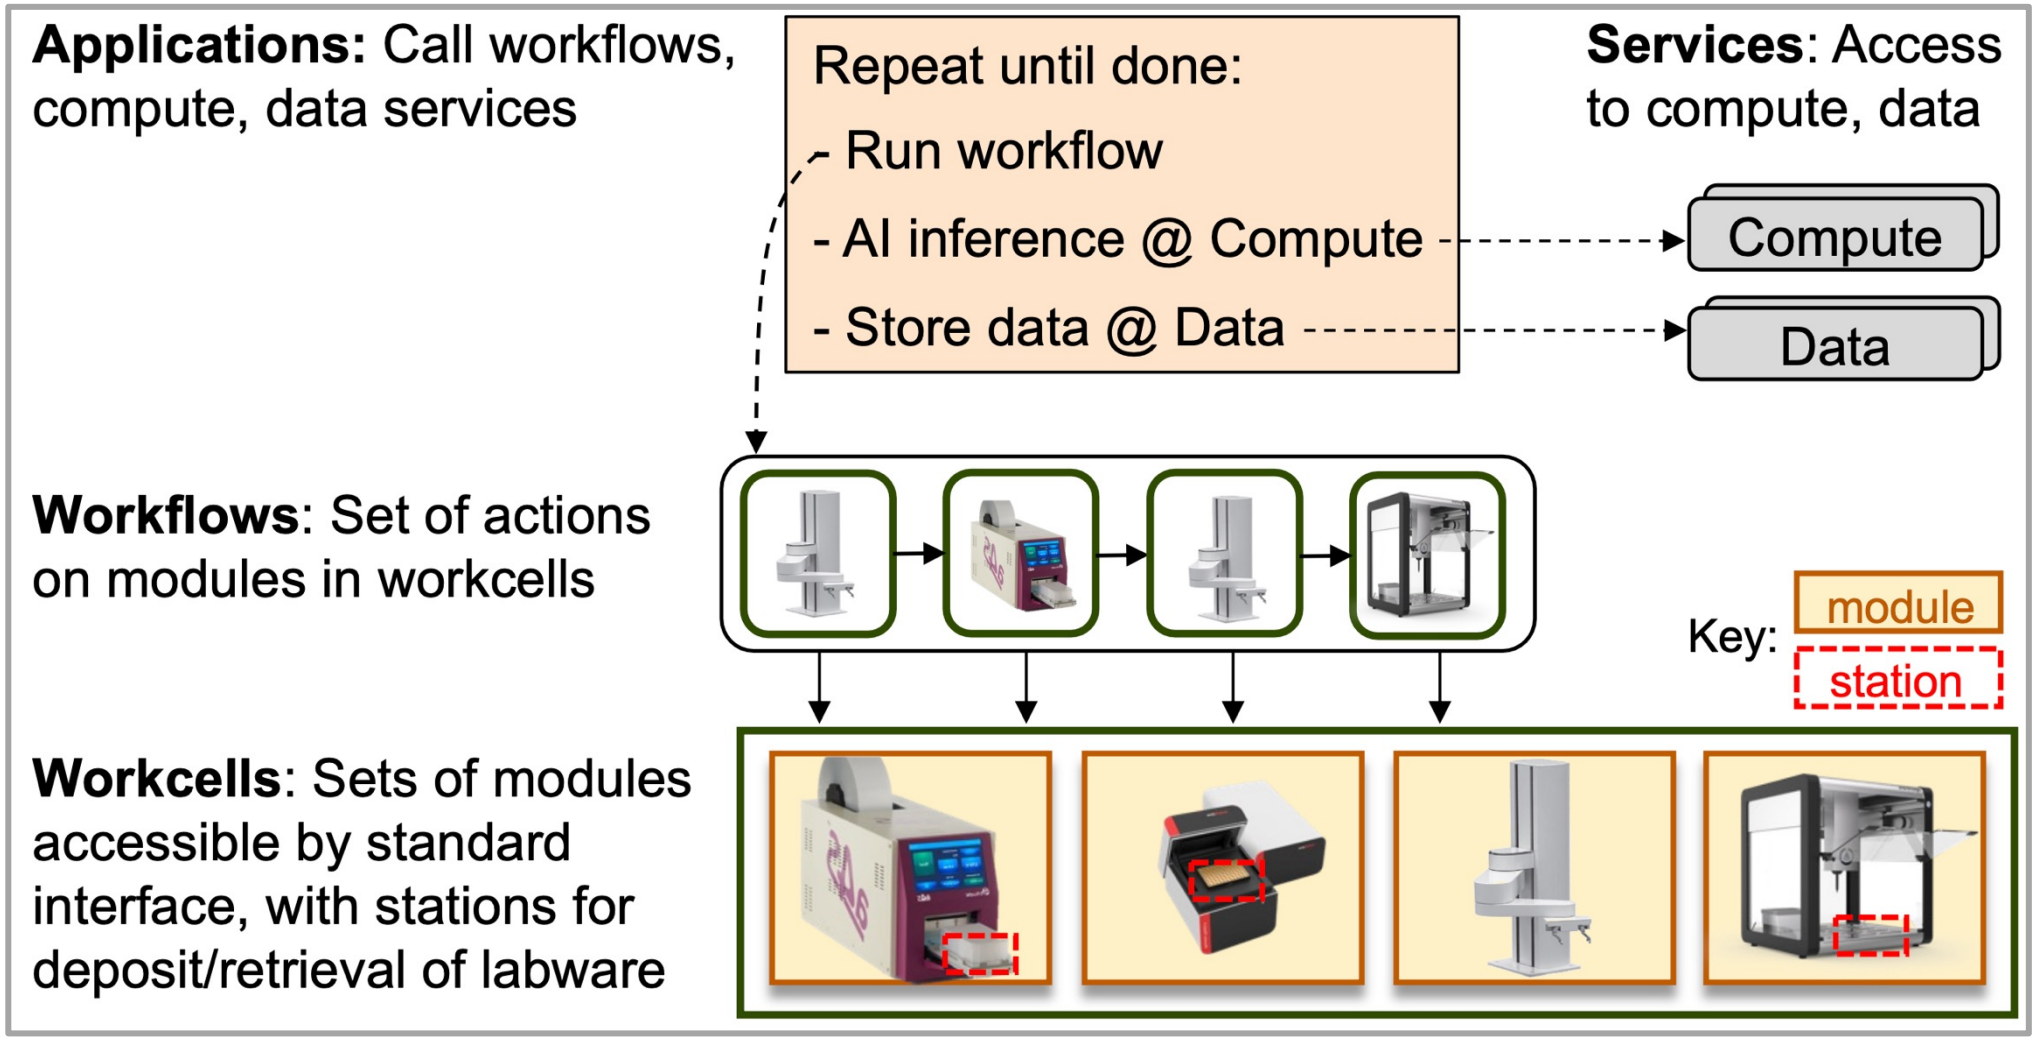

We are applying these methods to construct science factories

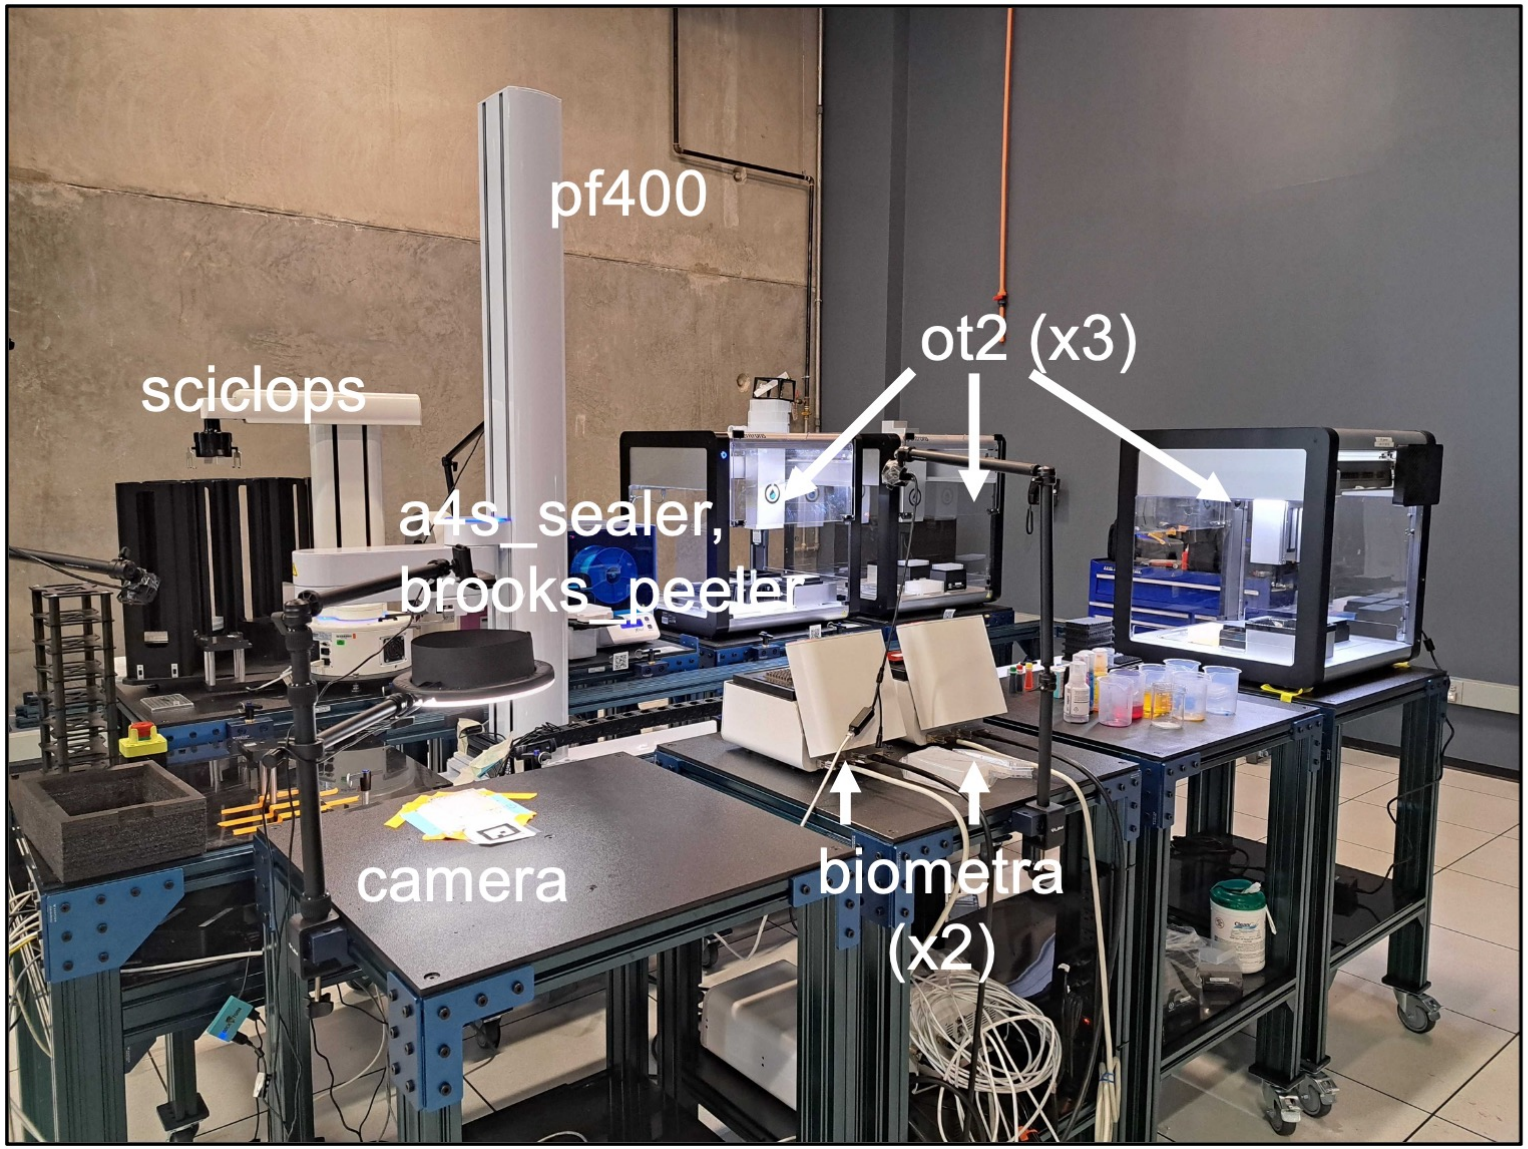

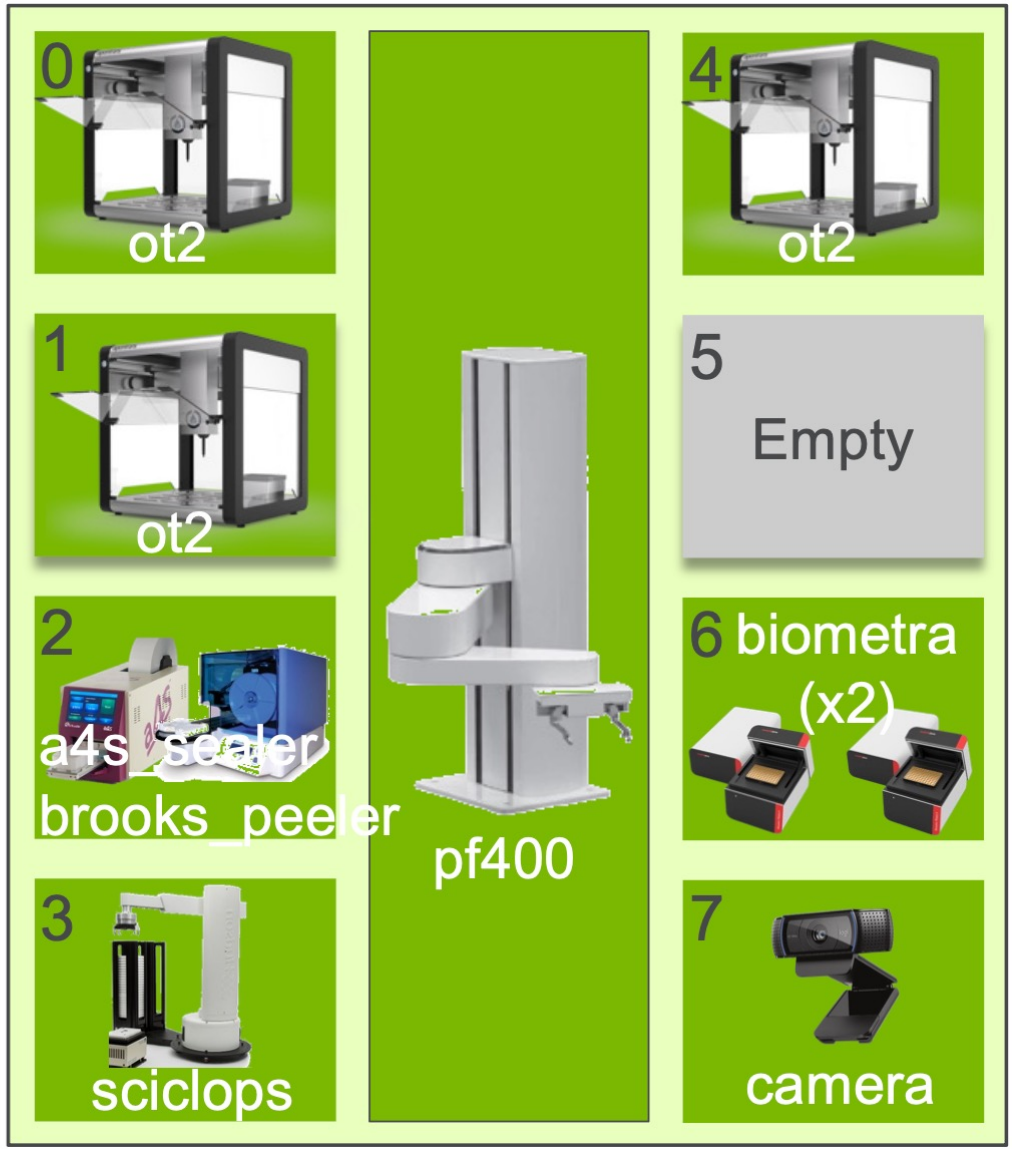

https://arxiv.org/abs/2308.09793

Modularity simplifies creation of digital twins

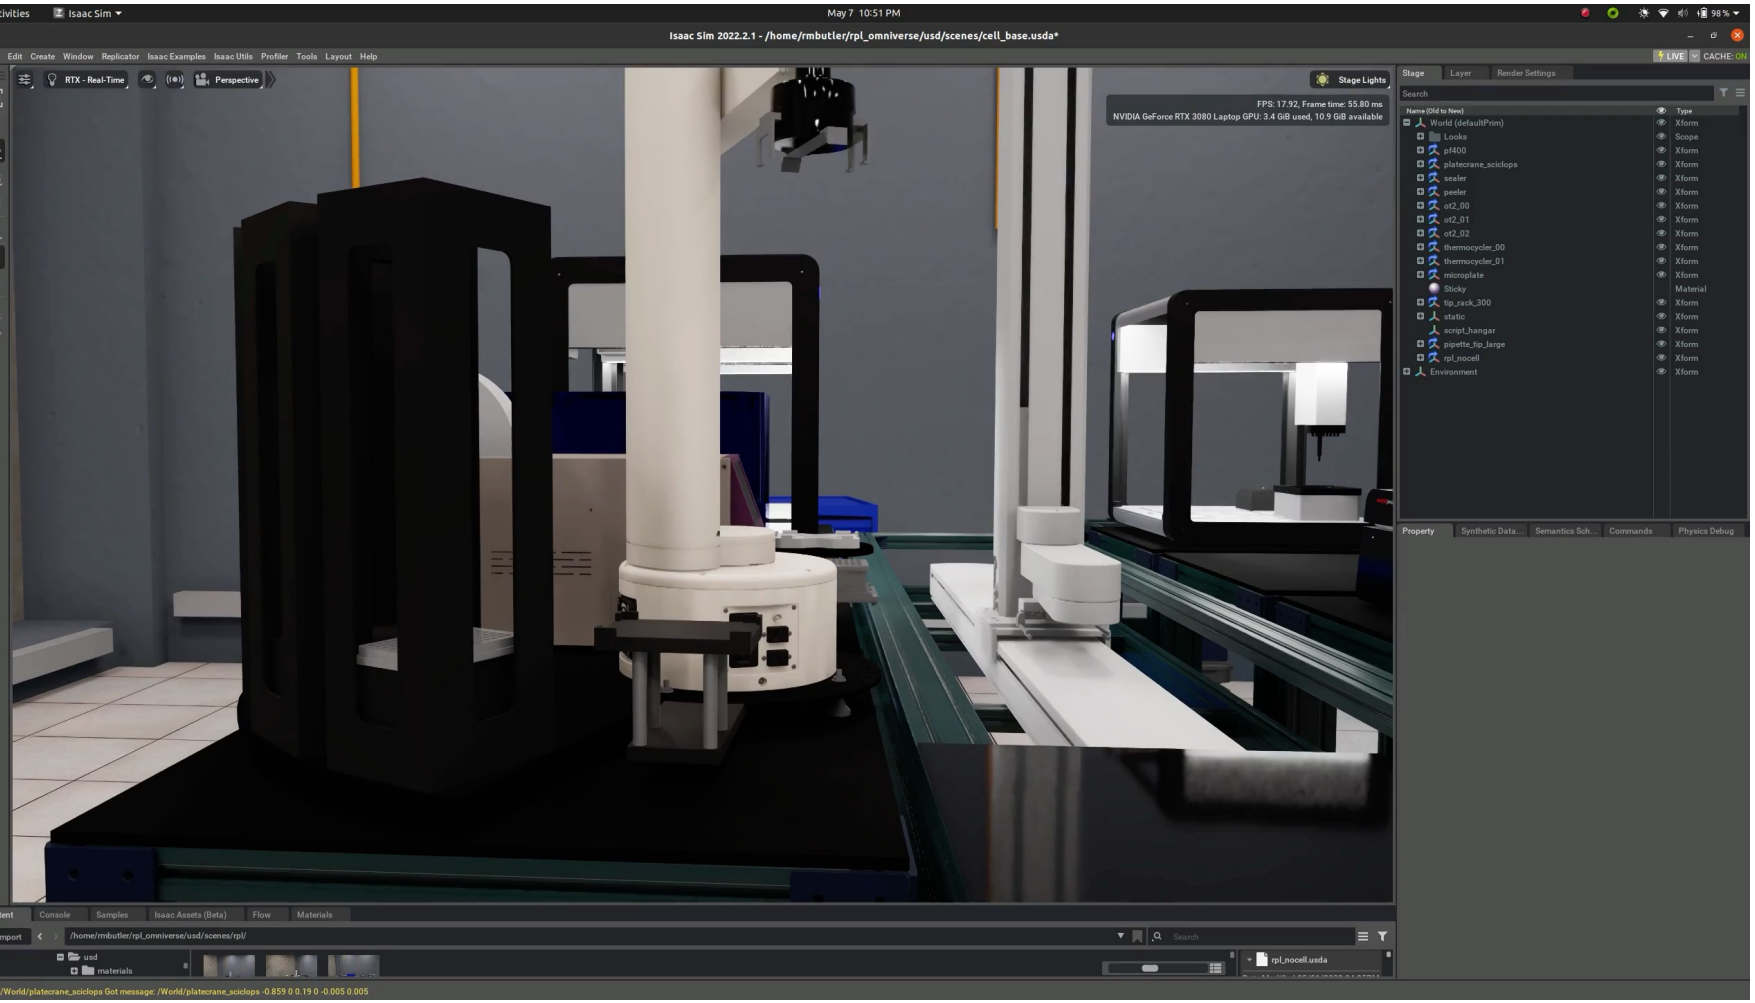

What else can we do with these methods and tools?

Integrate with large language model (LLM) technologies

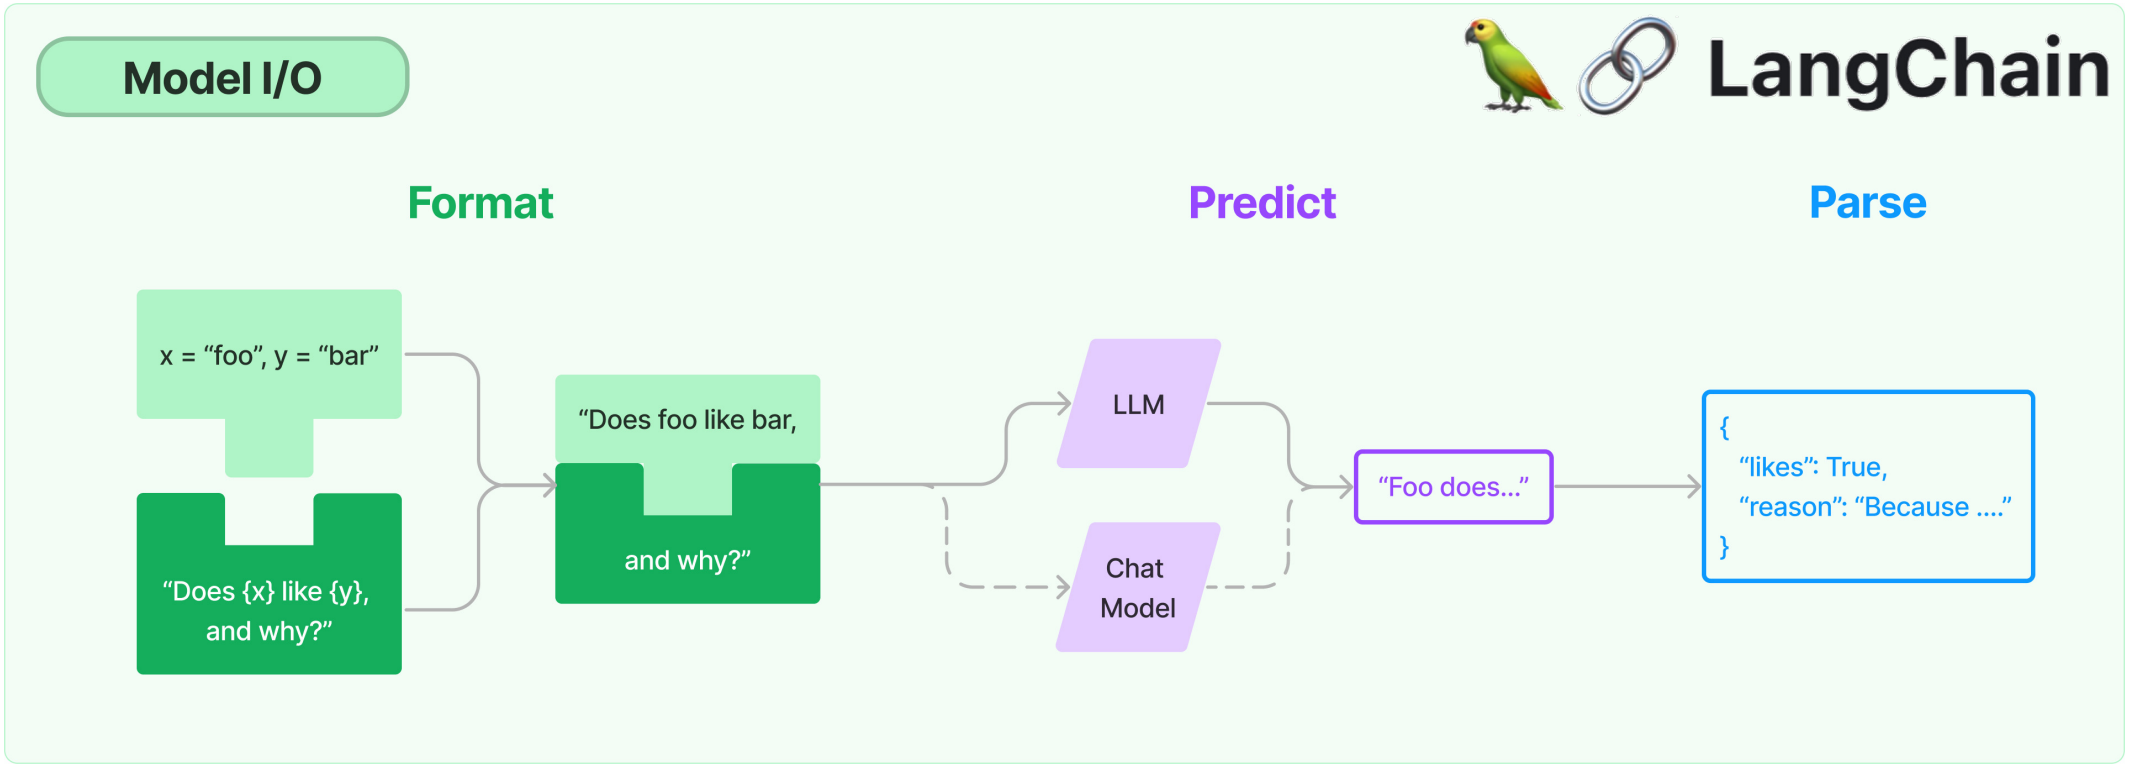

https://python.langchain.com/docs/modules/model_io/

What else can we do with these methods and tools?

Integrate with large language model (LLM) technologies

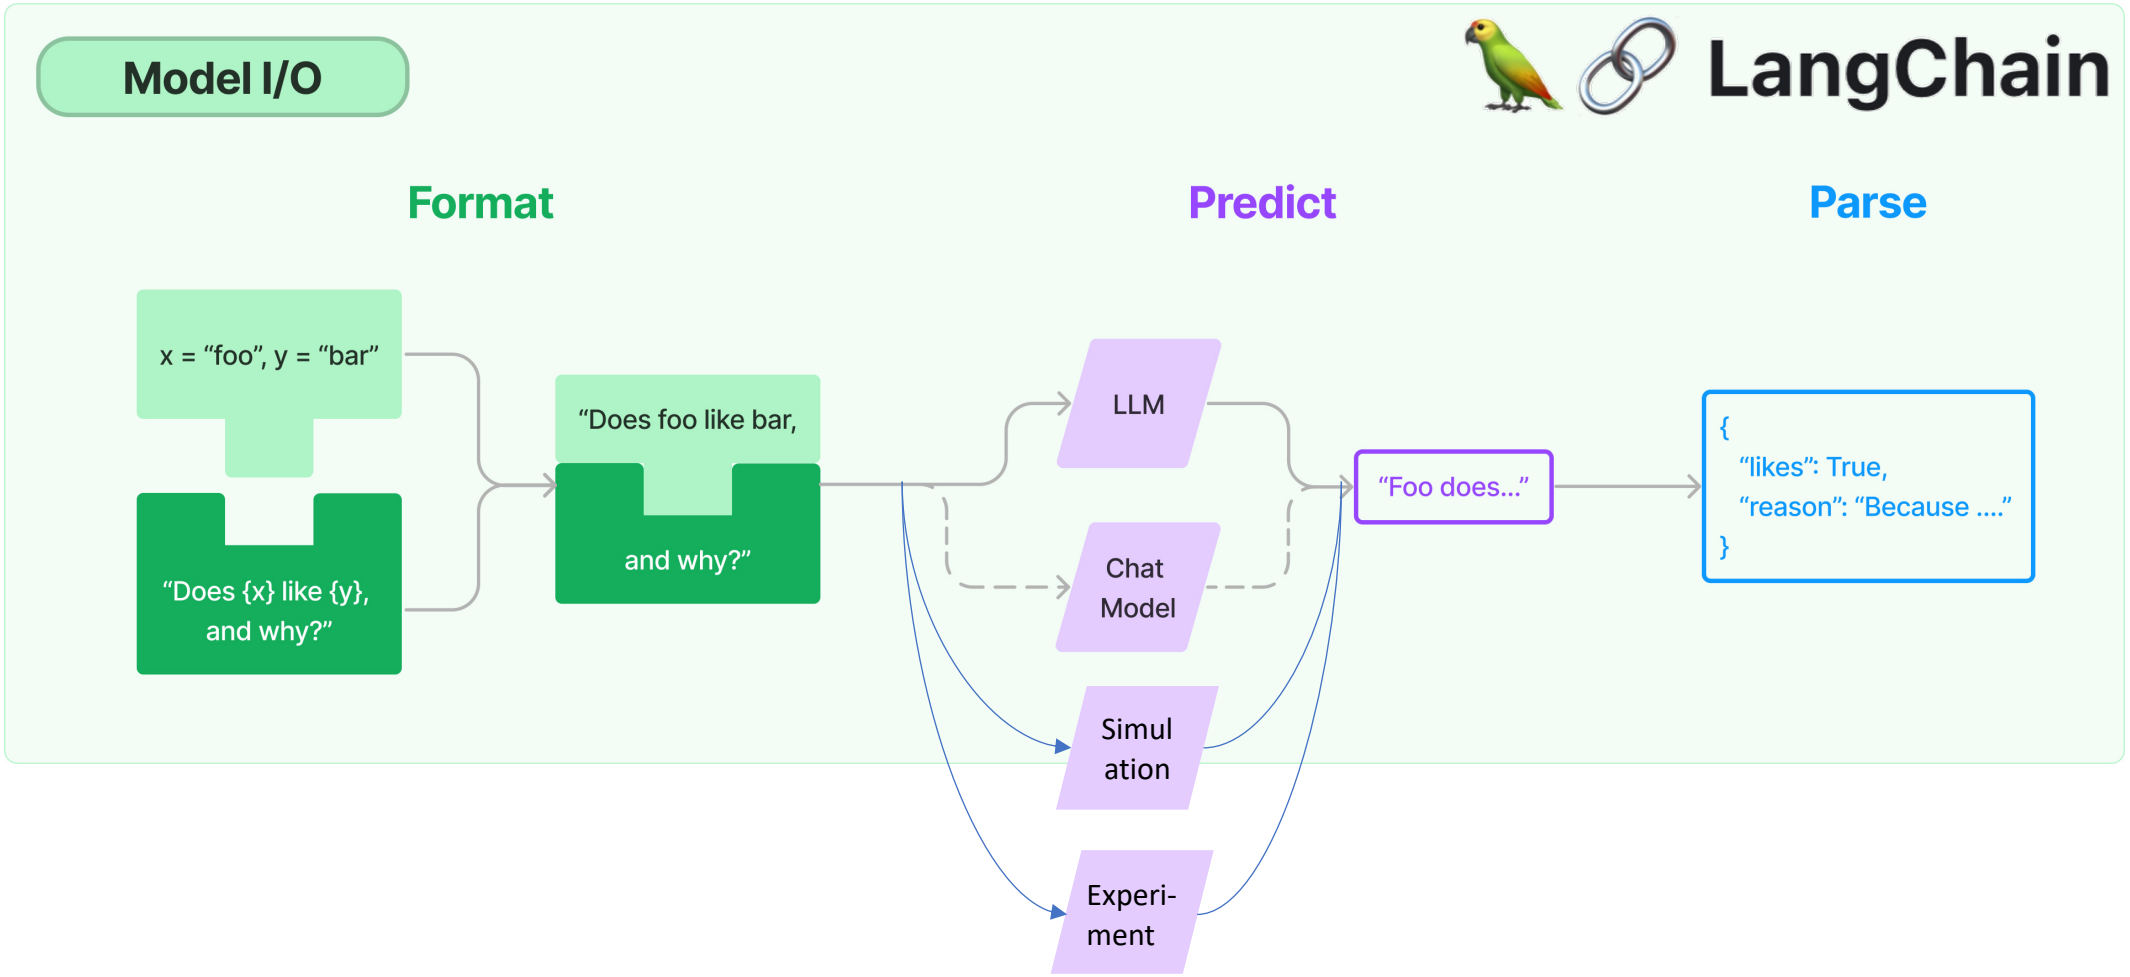

https://python.langchain.com/docs/modules/model_io/

What else can we do with these methods and tools?

Integrate with large language model (LLM) technologies

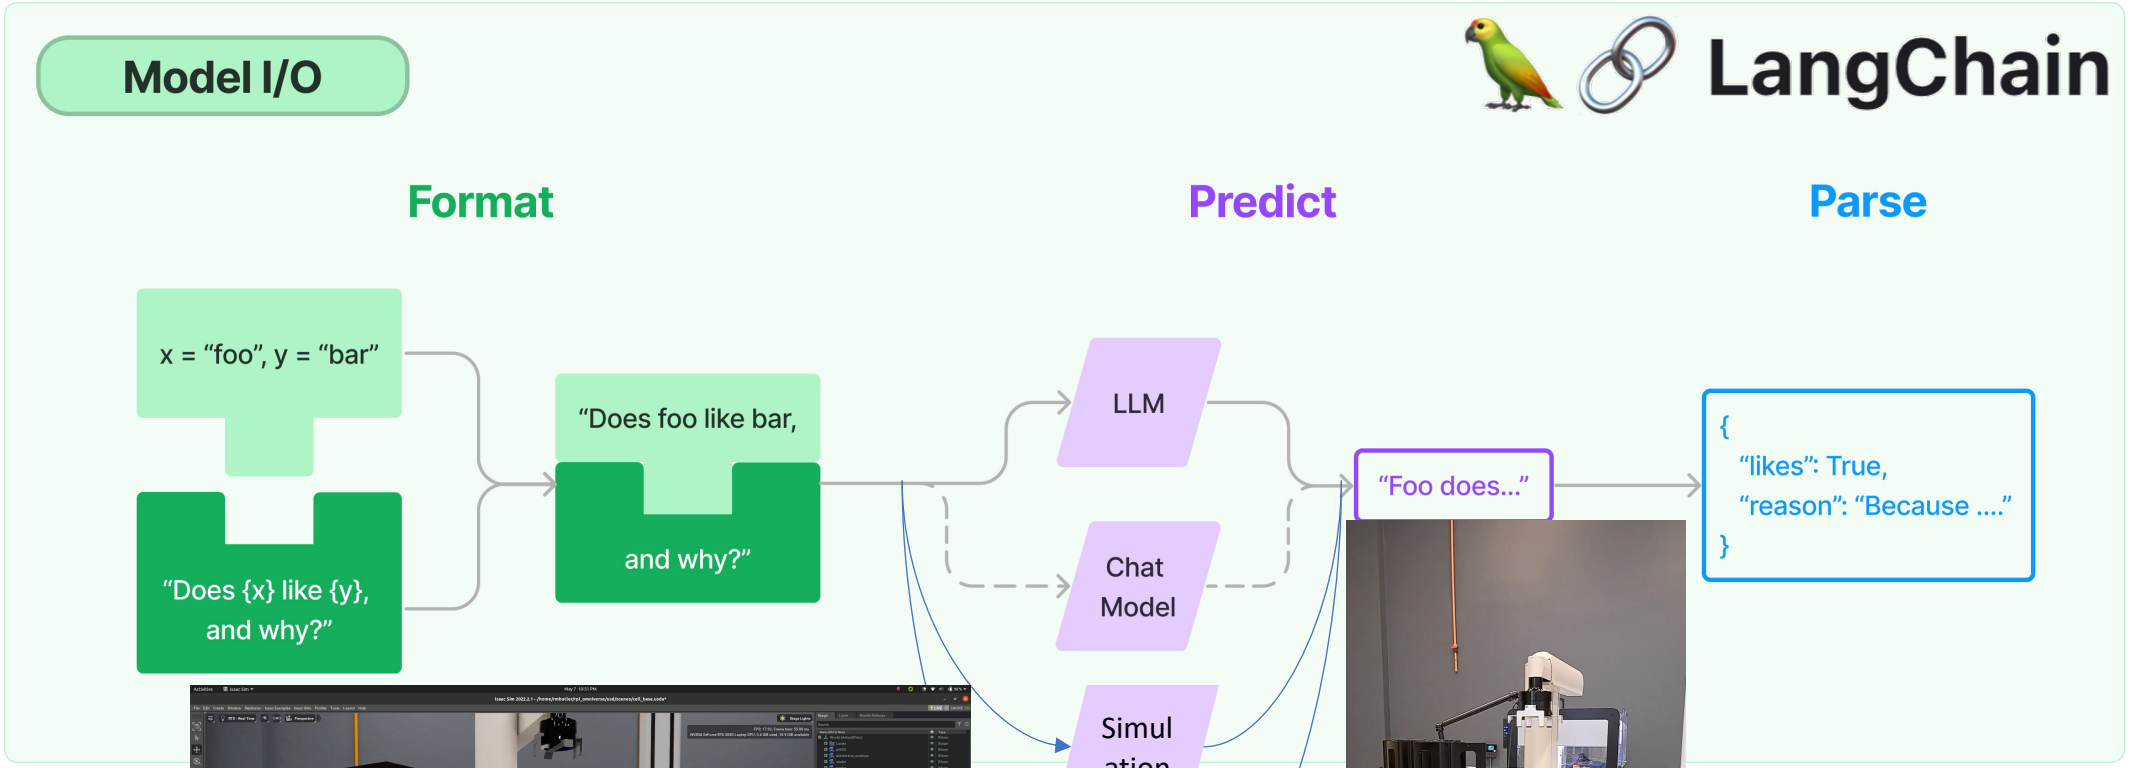

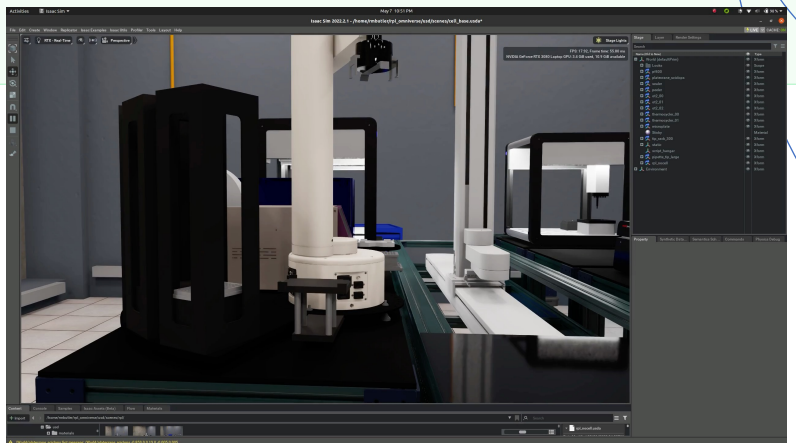

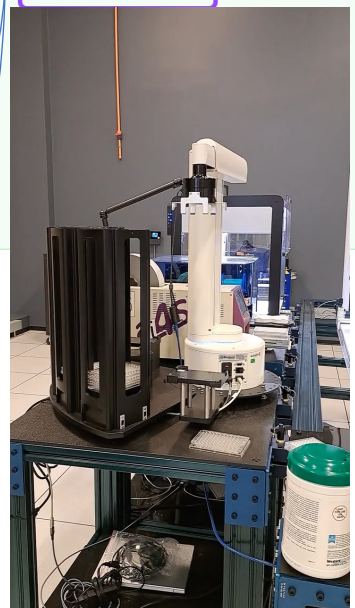

Rory Butler

Tobias Ginsburg, Vas Vescovi, et al.

There remain important questions to answer

- How to extend today's network into a truly universal continuum?
- Energy vs. accuracy vs. cost vs. … tradeoffs
- What computing and storage do we need, where?
- What will be economic foundation of this new computing fabric?
- What new scientific instruments will be created?
- What new applications and new science will be enabled?
- What new abstractions, services, and tools do we need?
- Do we need new continuum-aware algorithm design methods?
- (How) Will we integrate quantum sensors, networks, computers?

Thank you for your attention!

To learn more about our work: https://labs.globus.org https://globus.org

Questions or thoughts: foster@anl.gov

Experiment with tools:
https://braid-project.org

Thanks to:

Argonne National Laboratory and the **University of Chicago** – and students and staff

Federal agencies for continued support: **DOE**, **NSF**, **NIH**, **NIST**

Wonderful colleagues: Rachana Ananthakrishnan, Ben Blaiszik, Kyle Chard, Ryan Chard, Carl Kesselman, Arvind Ramanathan, Rick Stevens, Vas Vasiliadis, Logan Ward, & **many more**