Global Services for Global Science

Ian Foster

THE UNIVERSITY OF CHICAGO

globus labs

Argonne NATIONAL LABORATORY

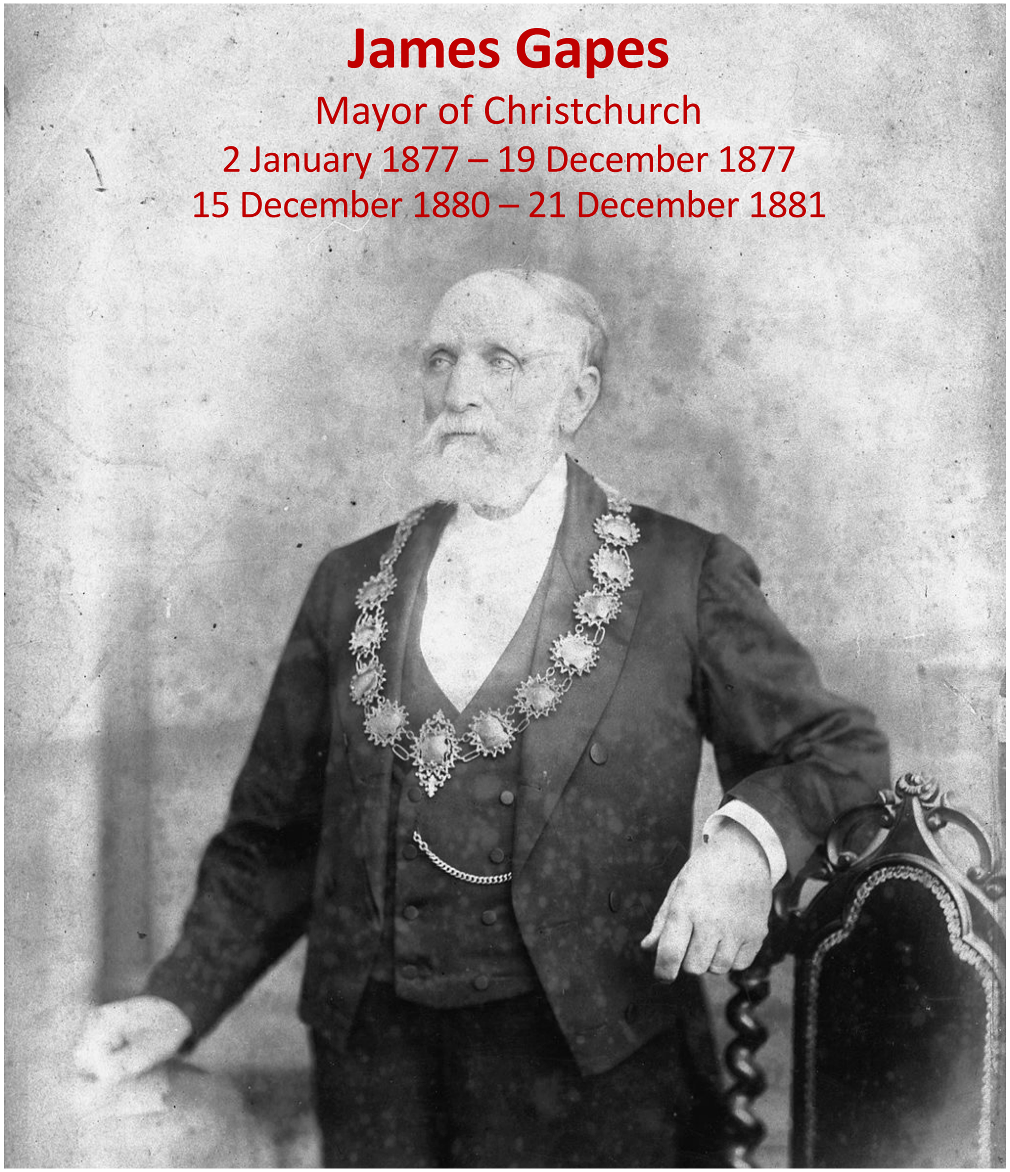

7 August 1876

[ADVERTISEMENT.]

CORRESPONDENCE.

A DISCLAIMER.

TO THE EDITOR OF THE STAR.

IN consequence of my name being uncommon in this Province, I beg to state that I am not the James Gapes that was fined for drunkenness at the Resident Magistrate's Court.

JAMES GAPES,
Painter, &c., Whately road.

9671

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

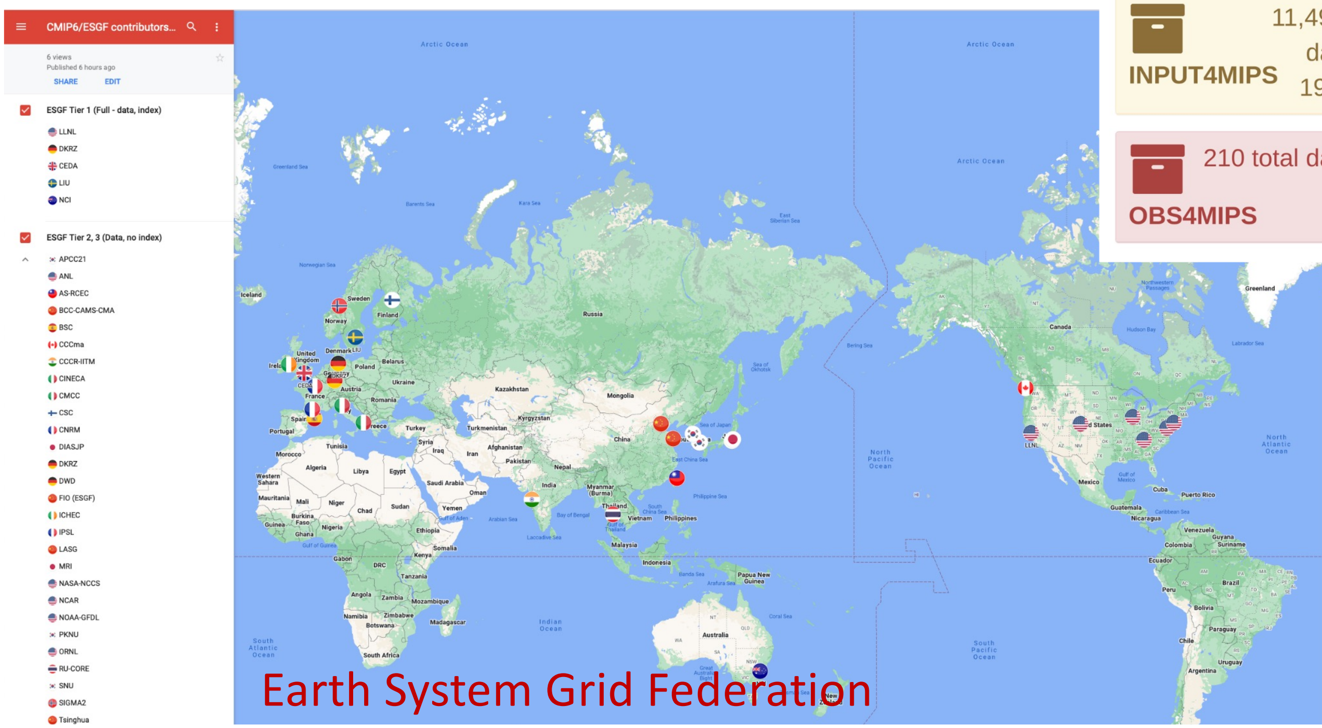

>100 models, >20 countries

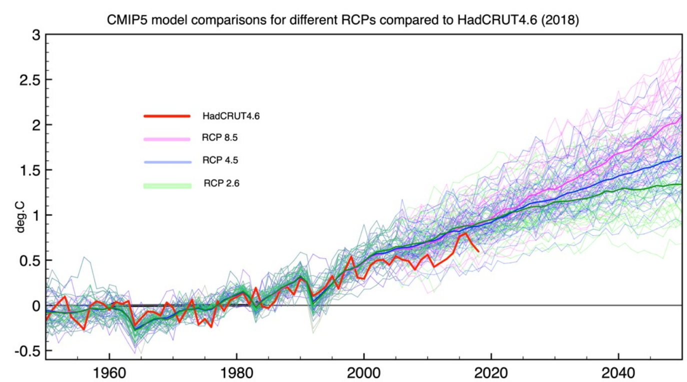

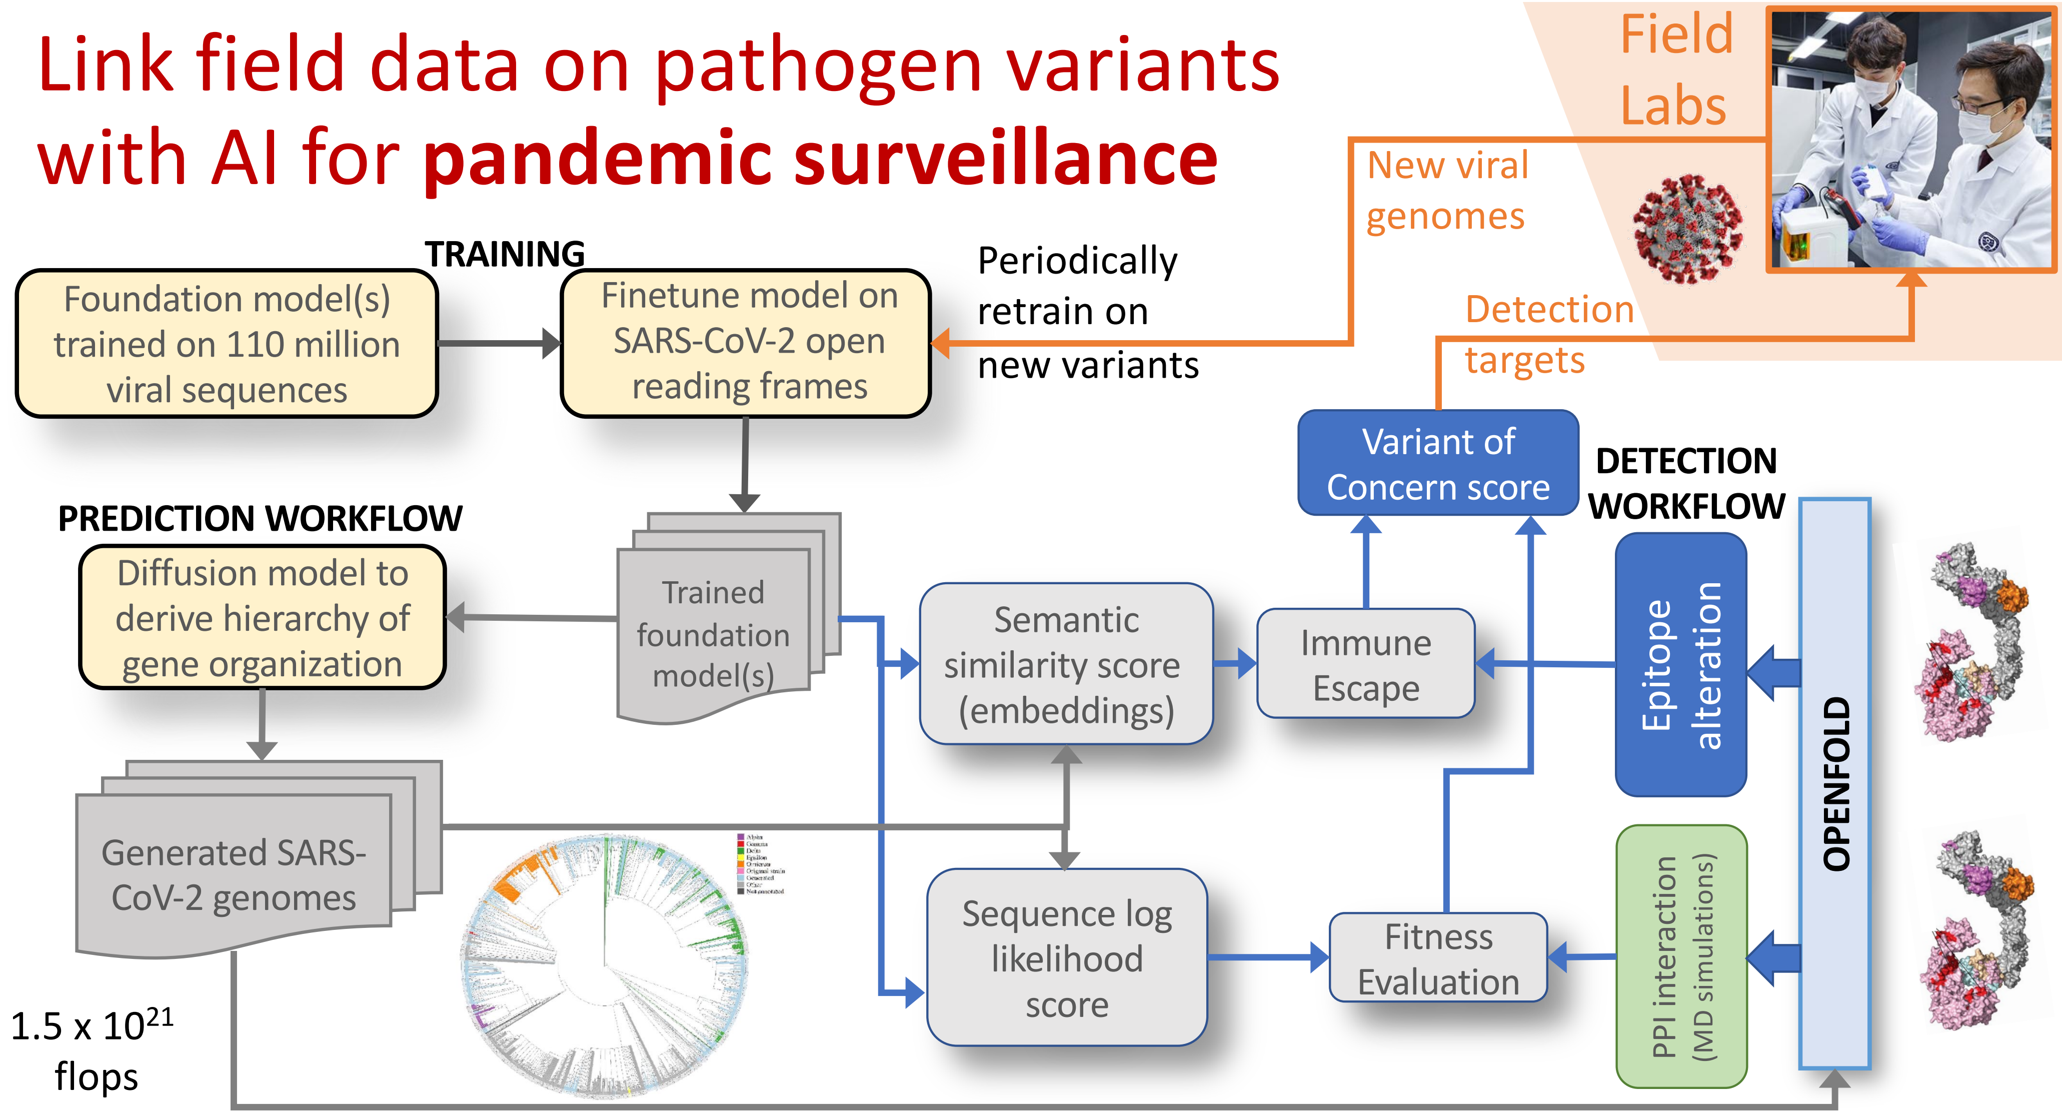

M. Zvyagin et al., https://www.biorxiv.org/content/10.1101/2022.10.10.511571v1

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

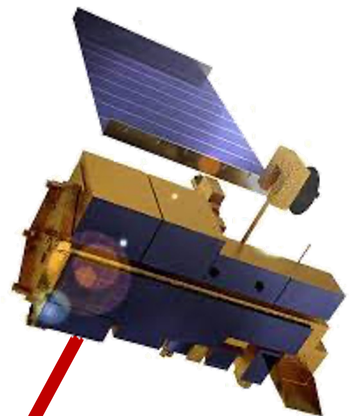

CROCUS
Community Research on
Climate & Urban Science

‘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… in Lemont. **ANTONIO PEREZ/CHICAGO TRIBUNE PHOTOS** …onal Laboratory

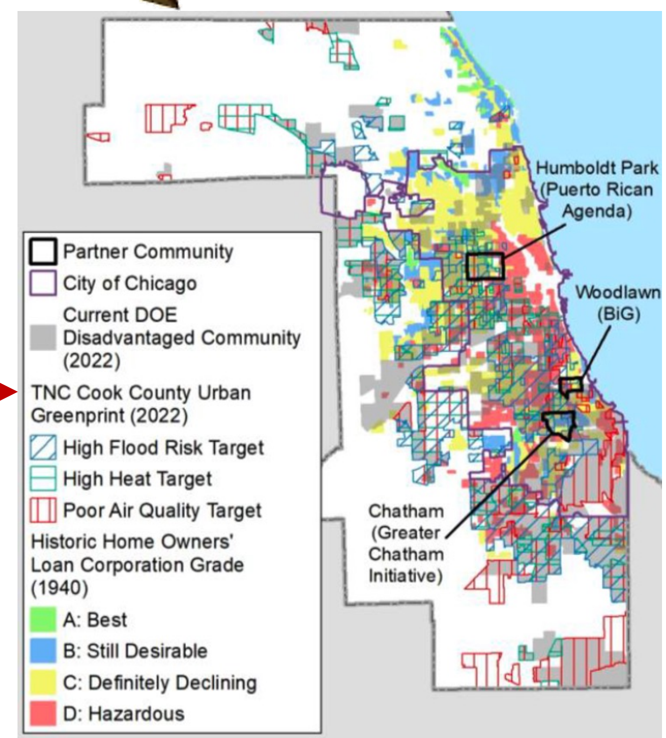

Fallacies of distributed computing

1. The network is reliable
2. Latency is zero
3. Bandwidth is infinite
4. The network is secure
5. There is a single administrator
6. The topology won't change
7. Transport cost is zero
8. The network is homogeneous

Joy, Lyon, Gosling, 1991? (items 1–4)

Deutsch, 1994 (items 5–7)

Gosling, 1997 (item 8)

We are entering a new era of **universal** and **reliable** communication & computation

- Widely deployed optical fiber, cellular radio, [and free-space optics will] provide ultra-fast, reliable, predictable **communication**
- GPS and atomic clocks provide highly accurate **time signals**
- Widely deployed, highly replicated, highly available **cloud services + accurate time** permit highly reliable **global state**

ESnet as exemplar of science networks

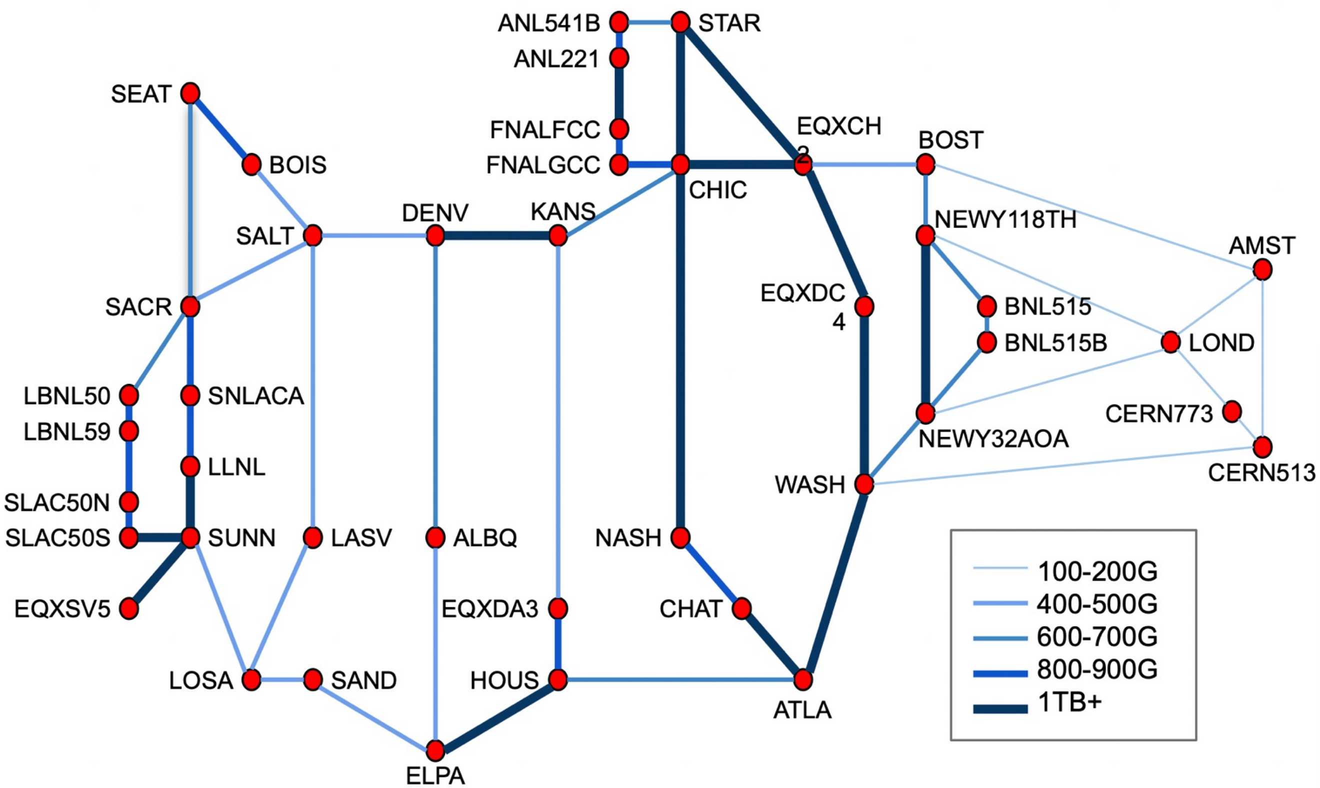

ESnet6 network capacity, as of May 2022

ESnet
ENERGY SCIENCES NETWORK

- **15,000 miles** of dark fiber lit up
- **300 leased spaces** installed across the US with ESnet owned equipment
- New fiber spans to **increase reliability and reduce latency**
- **46.1 Tbps** aggregate capacity deployed
- **400Gbps - 1 Tbps** services available

ESnet improvements in bisection bandwidth over 30 years

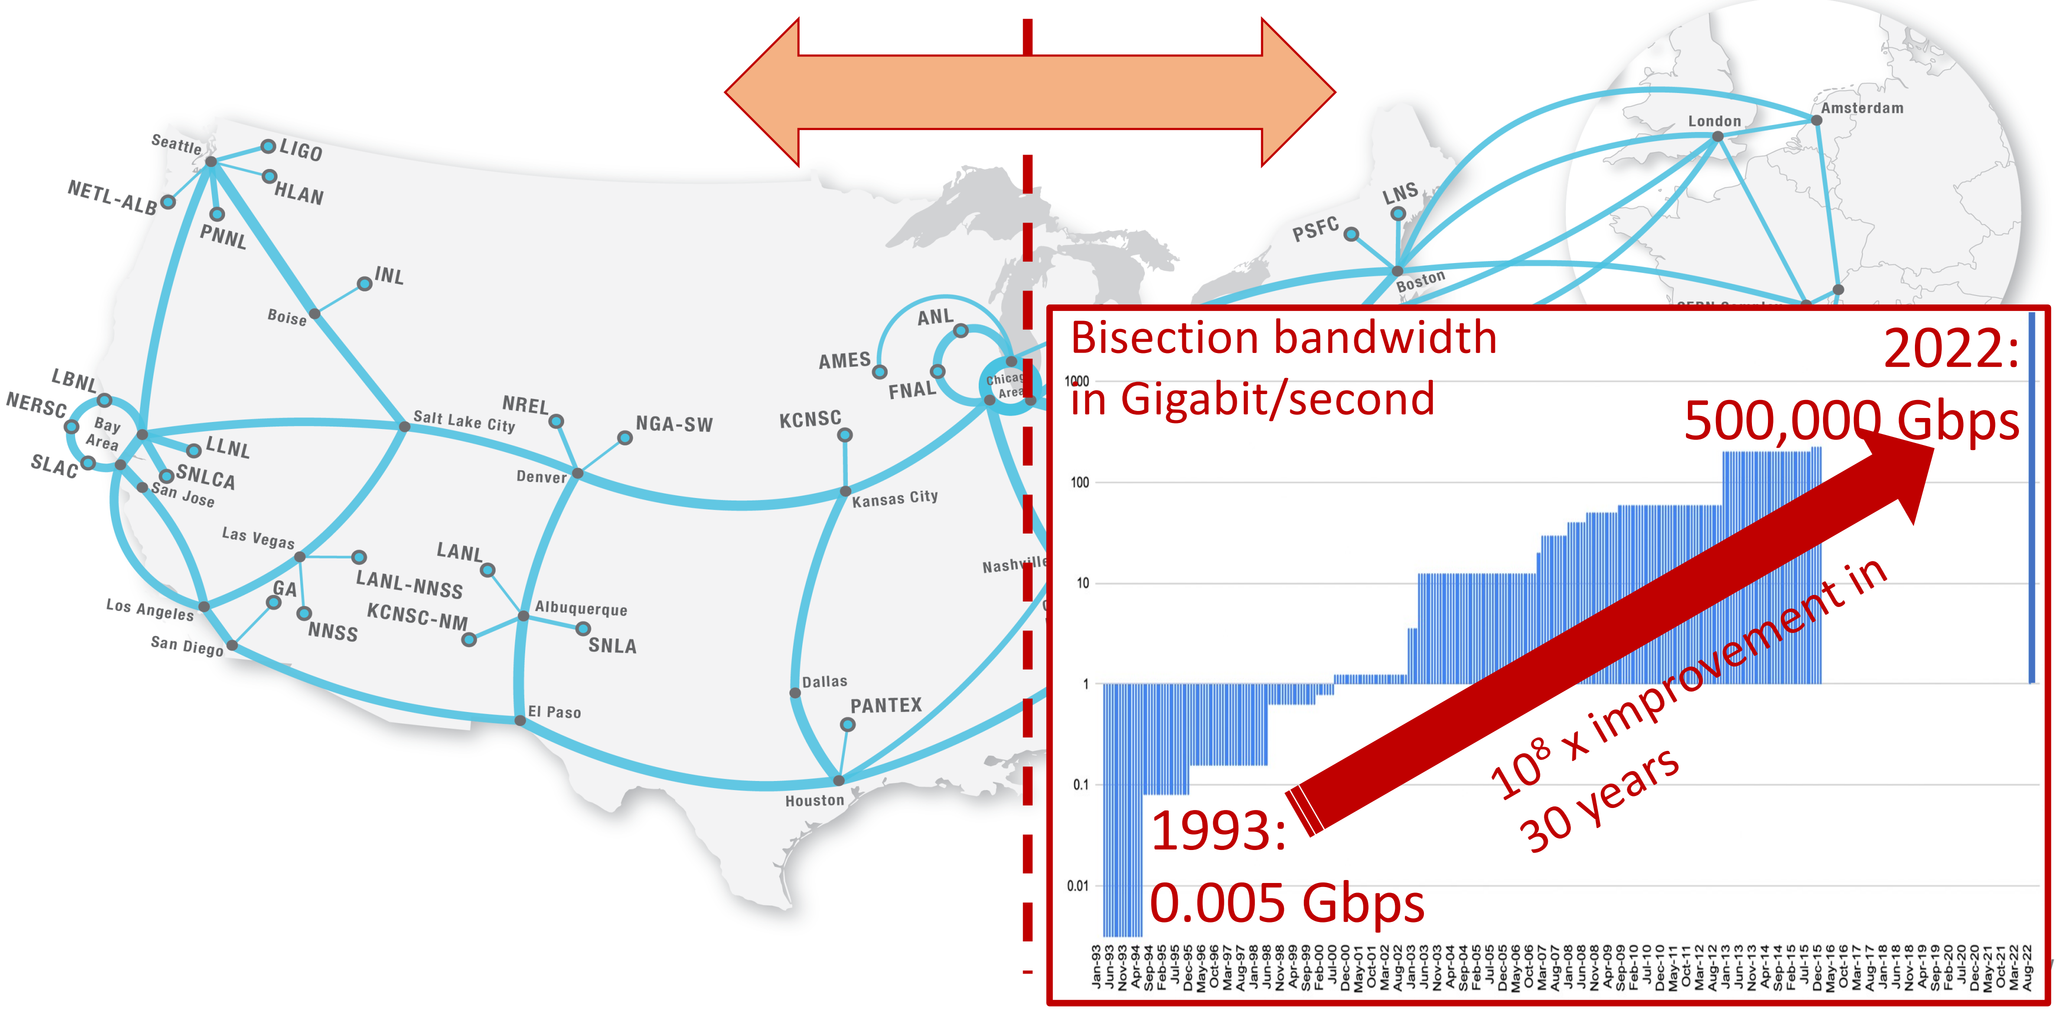


Aalyria's "Spacetime" [originally "Minkowski"] platform
Free-space optics


”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, 109
Aalyria’s “Spacetime” [originally “Minkowski”] platform
Free-space optics

The space-time continuum in converged systems

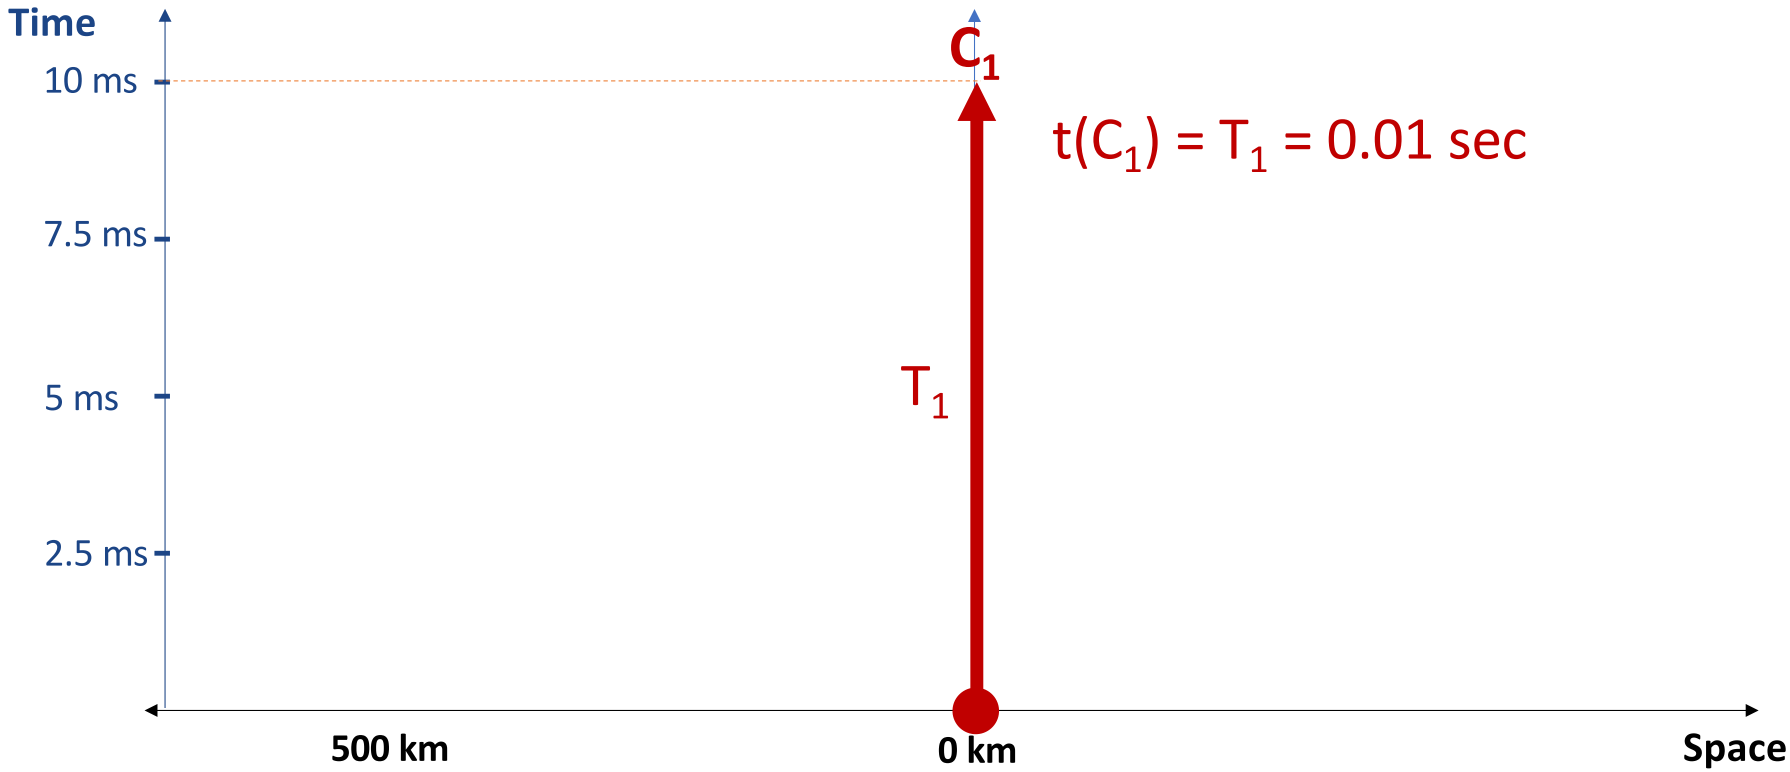

Misquoted [2022]: "*Henceforth, **location** for itself, and **speed** for itself, are doomed to fade away into mere shadows, and only a kind of union of the two will preserve an independent reality*."

The space-time continuum in converged systems

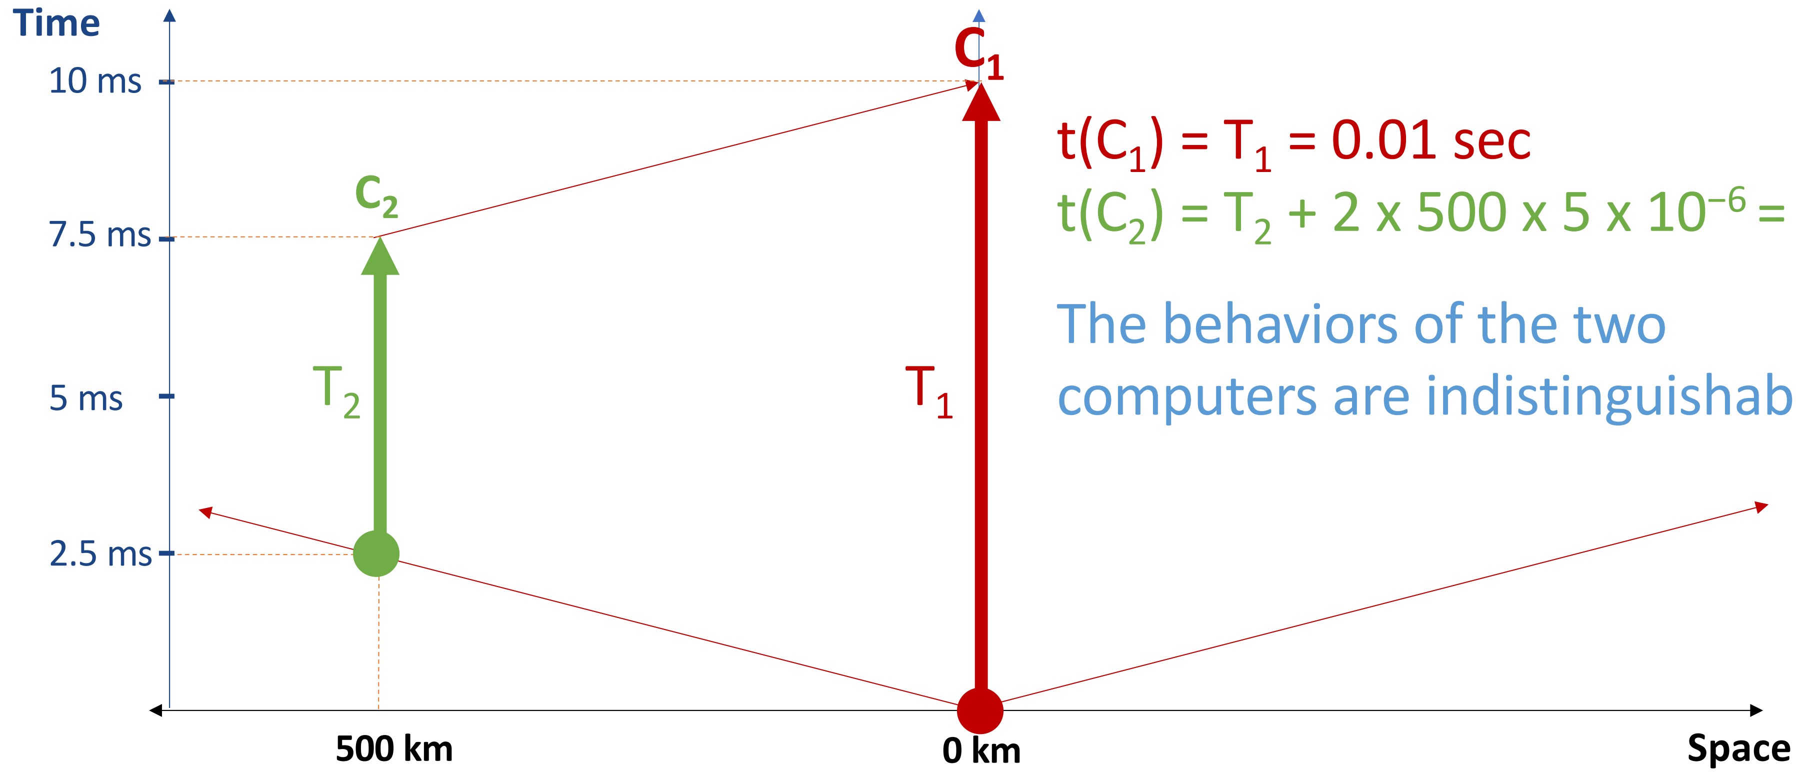

$t(C_1) = T_1 = 0.01$ sec

$t(C_2) = T_2 + 2 \times 500 \times 5 \times 10^{-6} = 0.01$ sec

The behaviors of the two computers are indistinguishable

Misquoted [2022]: "*Henceforth, **location** for itself, and **speed** for itself, are doomed to fade away into mere shadows, and only a kind of union of the two will preserve an independent reality*."

Example: **High energy physics trigger analysis**

Local: 2000 msec

Remote: 30 + 10 + 10 = 50 msec

40x acceleration

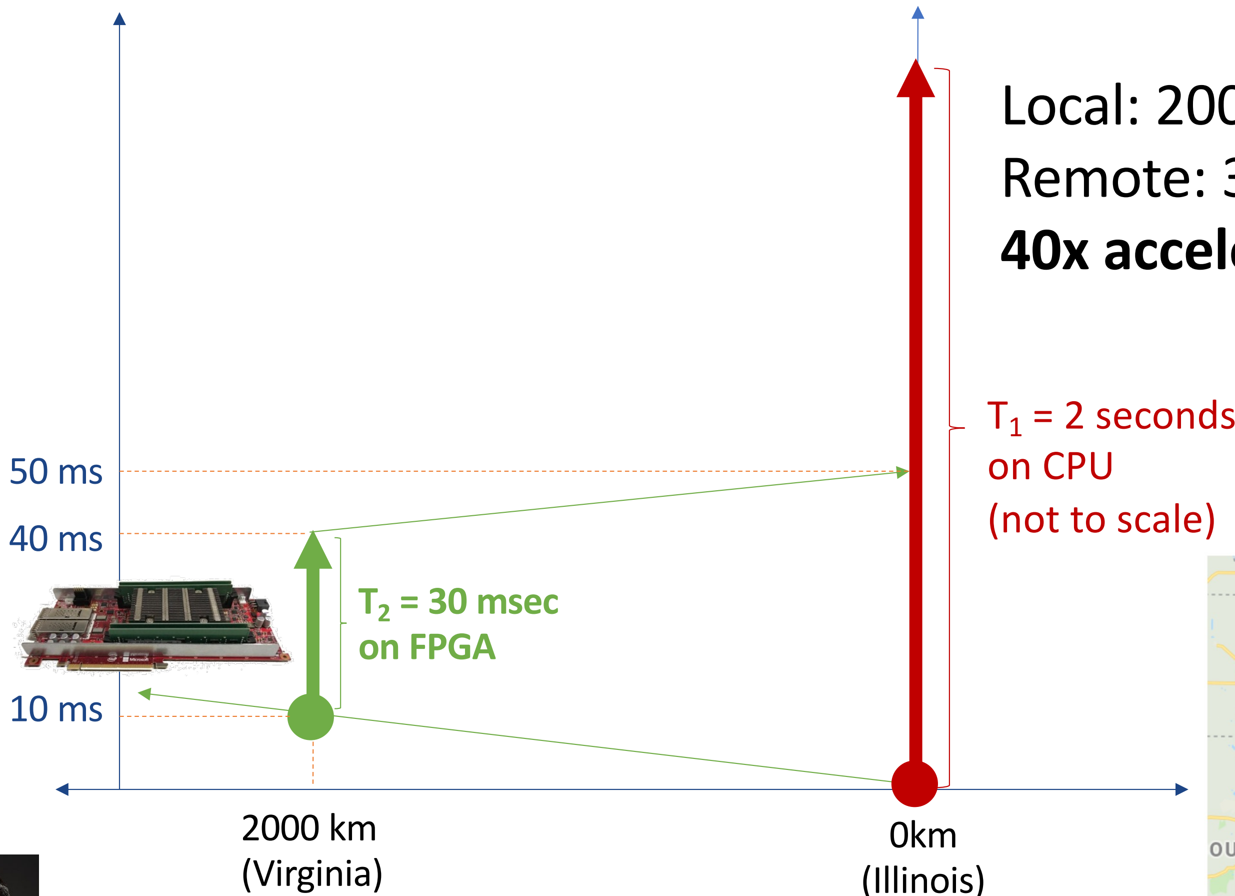

Nhan Tran, FermiLab, et al. arXiv:1904.08986

Example: **High energy diffraction microscopy**

Model (re)training on local GPU:
1102 seconds

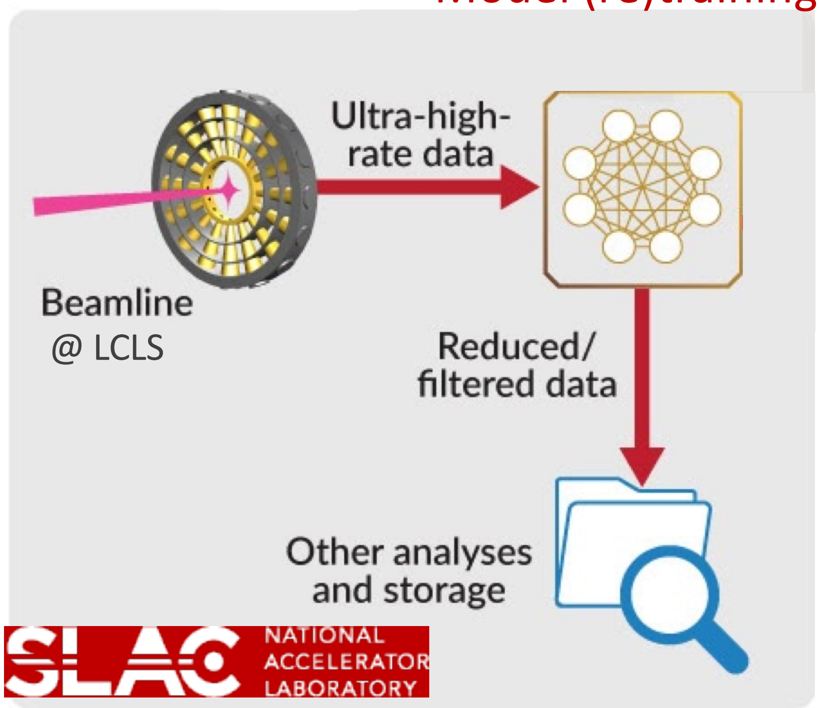

Z. Liu et al., https://doi.org/10.48550/arXiv.2105.1396

Example: **High energy diffraction microscopy**

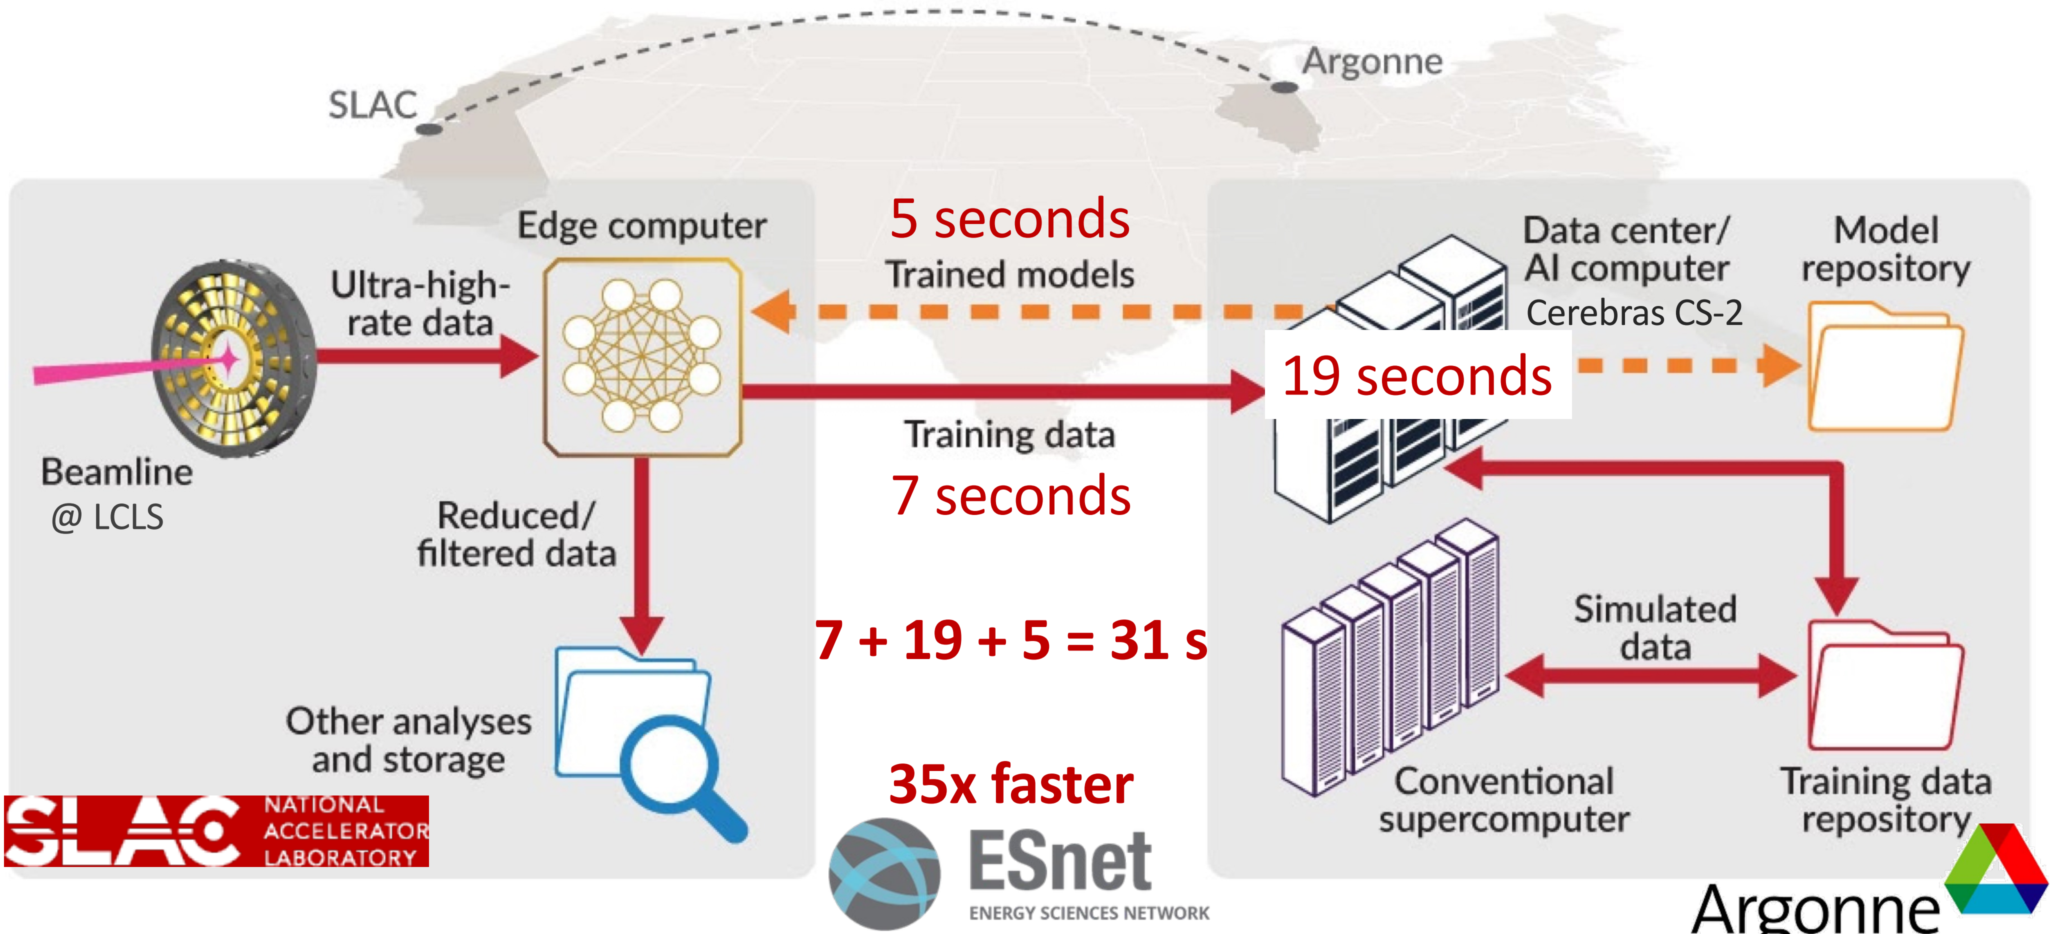

Z. Liu et al., https://doi.org/10.48550/arXiv.2105.13967

Building such distributed applications requires practical solutions to challenges 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).

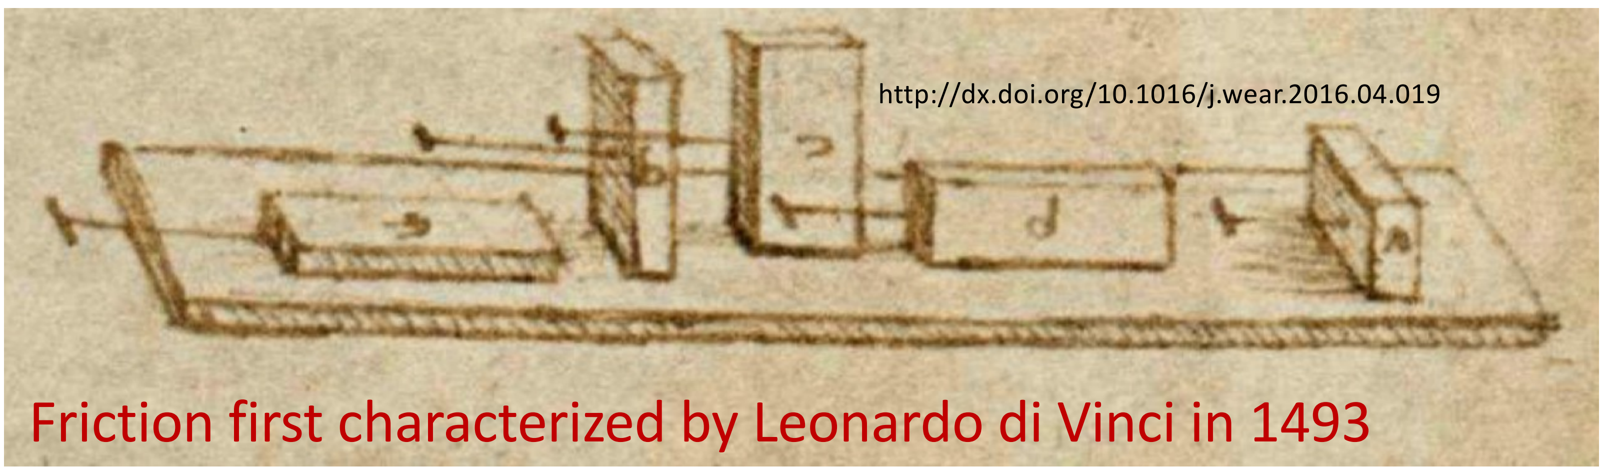

Friction first characterized by Leonardo di Vinci in 1493

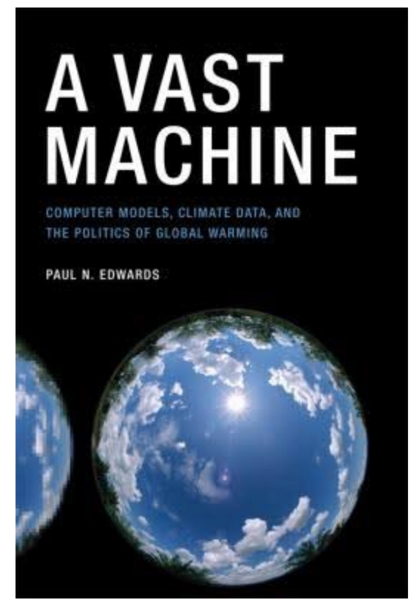

Three obstacles to tackling such codesign problems

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

Friction: Varying interfaces, behaviors; reliability; security

Three obstacles to tackling such codesign problems

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

Friction: Varying interfaces, behaviors; reliability; security

2) Execute remote actions **reliably**

Friction: Failures, scalability, usability

Three obstacles to tackling such codesign problems

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**

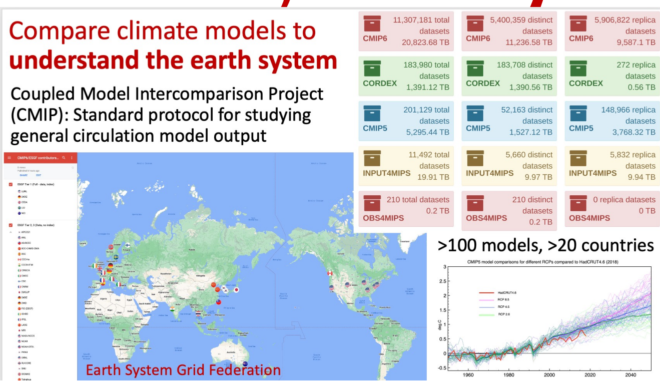

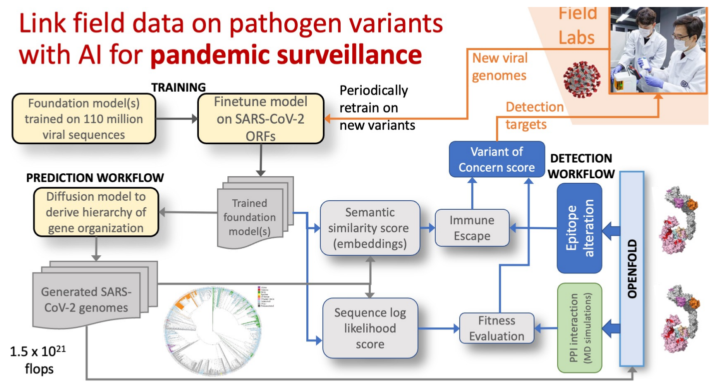

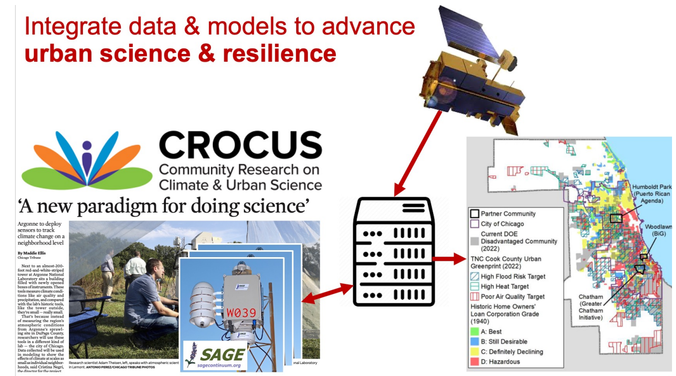

Past approaches (data actions):

- 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

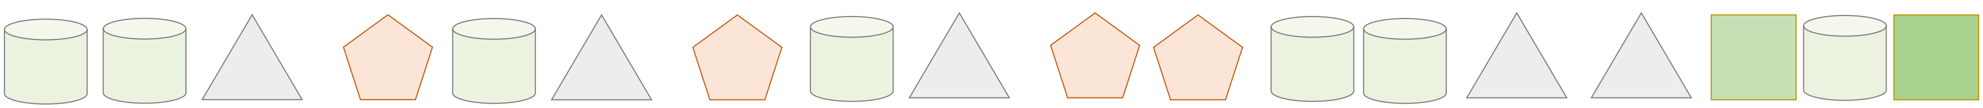

Need: 1) **Act anywhere**

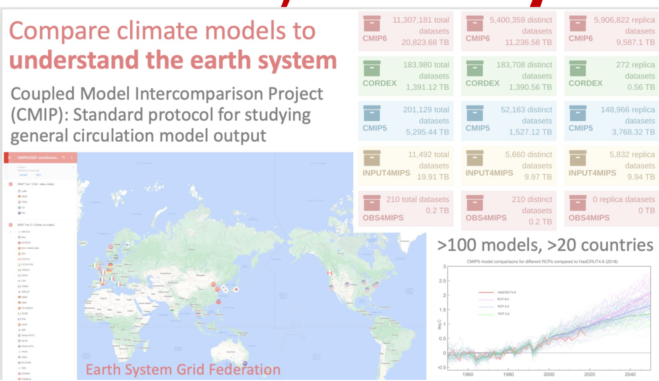

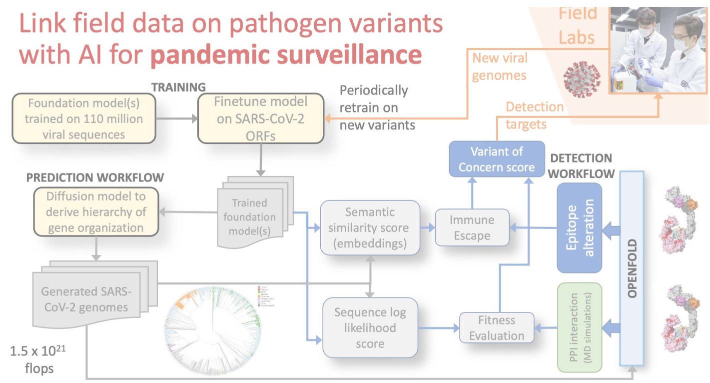

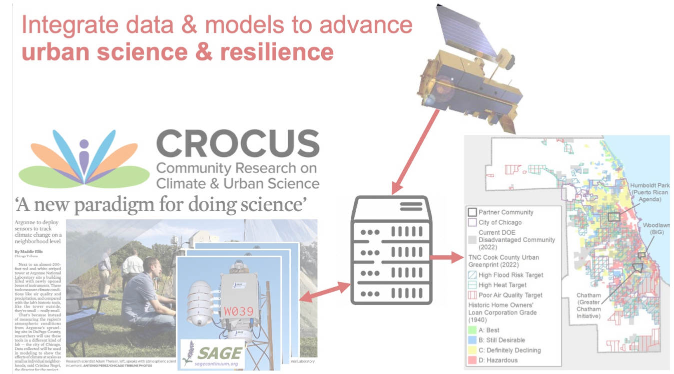

Our approach:

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

Lightweight agents abstract local details

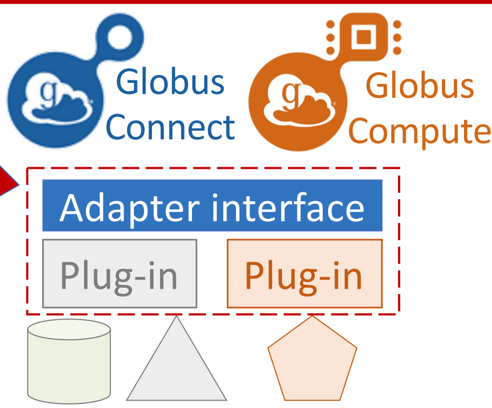

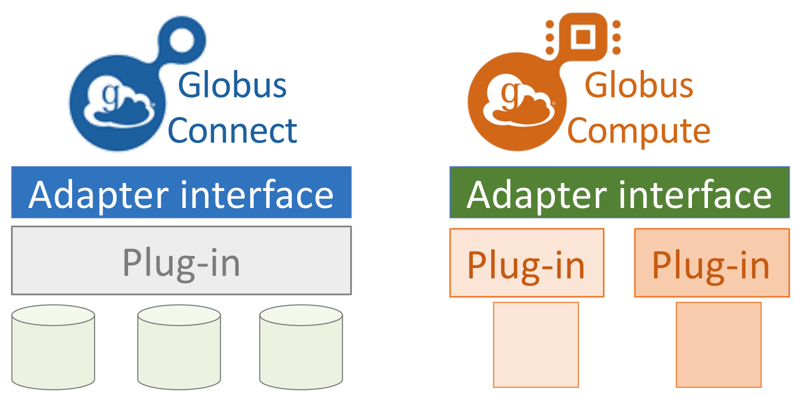

Need: 2) **Reliable execution** of (sets of) actions

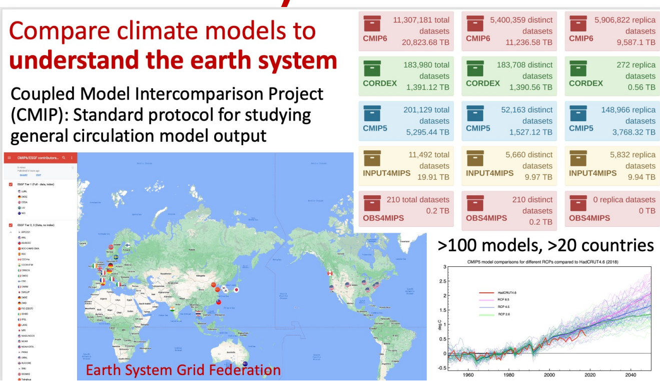

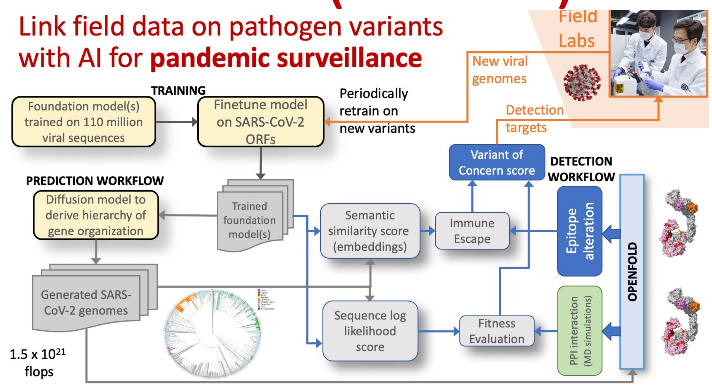

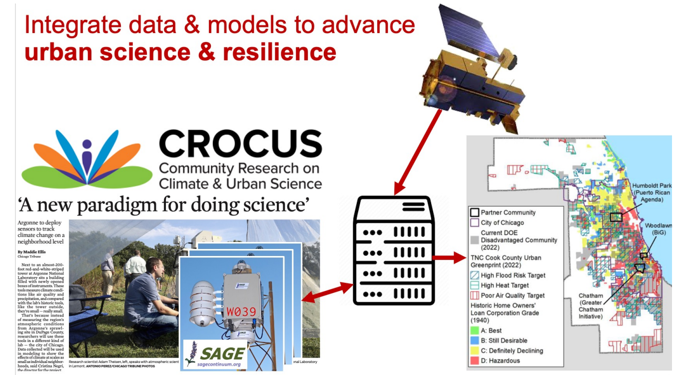

Past approaches:

- 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

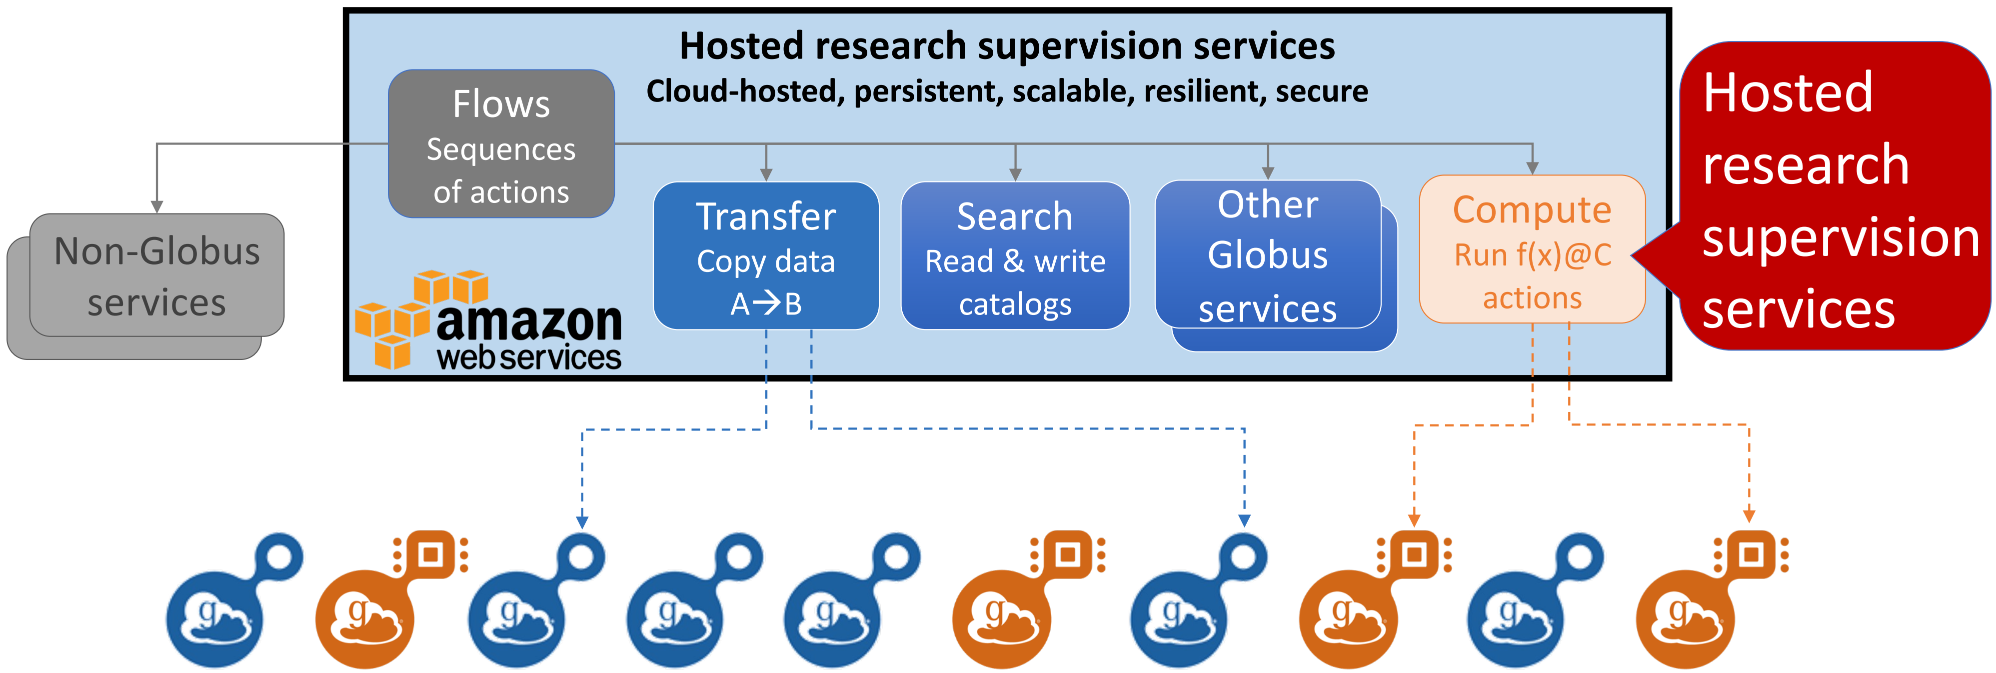

Need: 3) Control **who** can perform **what** actions, **when** & **where**

Past approaches:

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

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

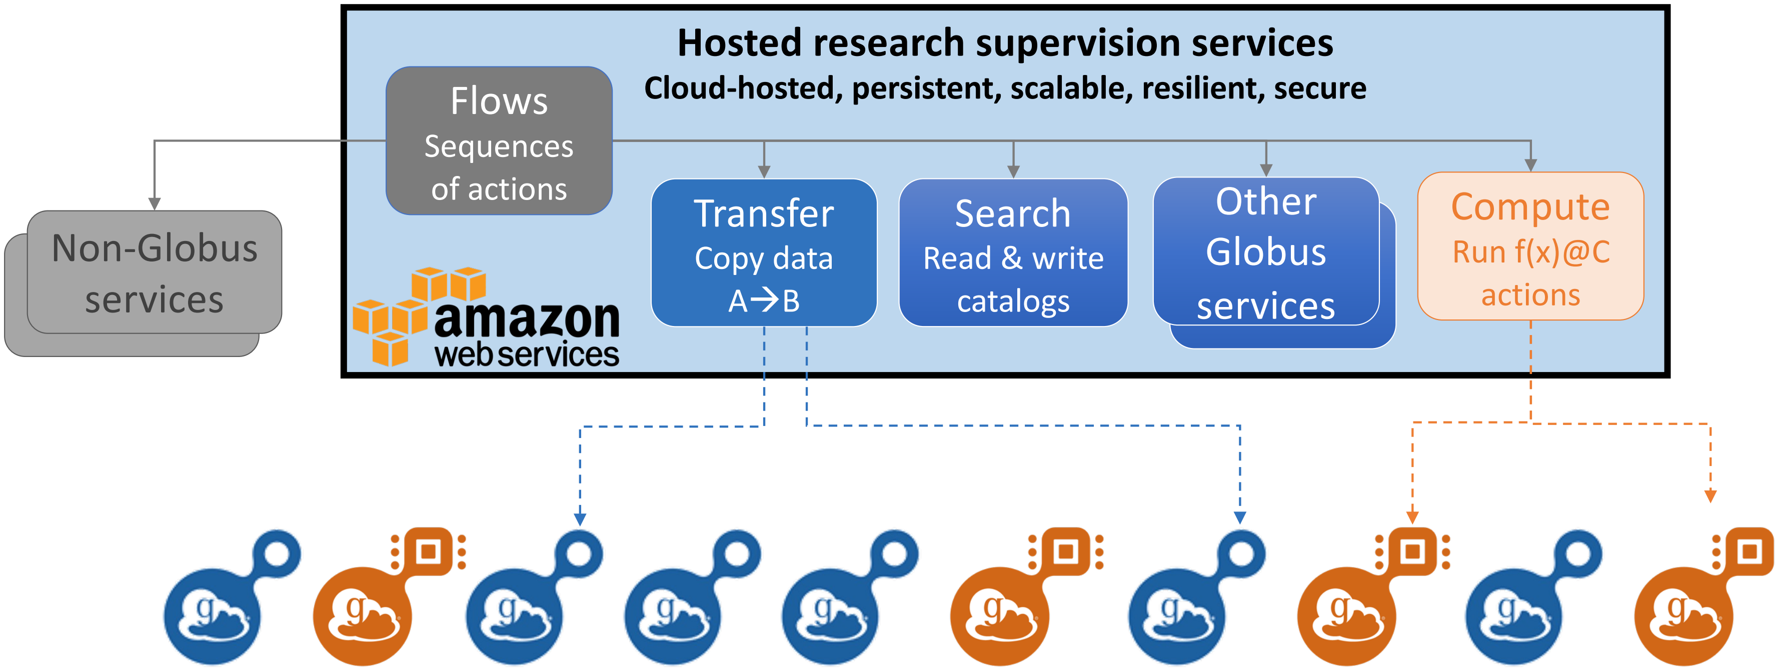

Need: 3) Control **who** can perform **what** actions, **when** & **where**

Federated auth & secure managed delegation

Our approach:

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

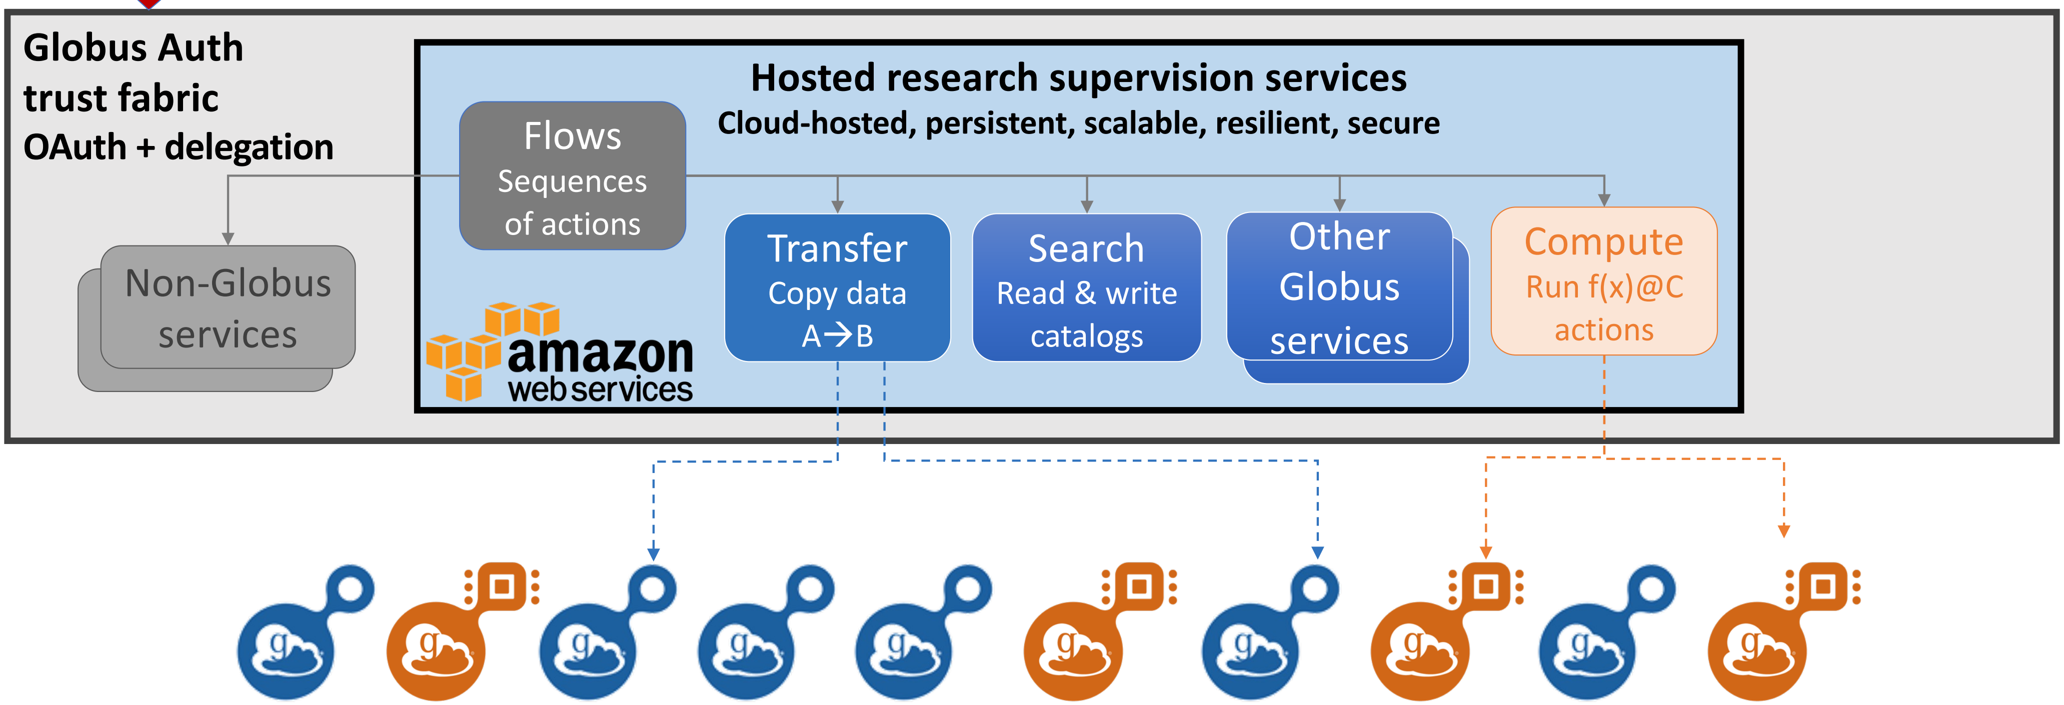

Globus Auth: A managed research acceleration service providing **distributed authorization with delegation**

Who do I trust to act on my behalf, **when**, and for **what** purpose?

Leverage OAuth for:
- **Security** via scoped credentials
- **Usability** via browser compatibility

1700 identity providers
1.3 B access tokens
2.7 M consents

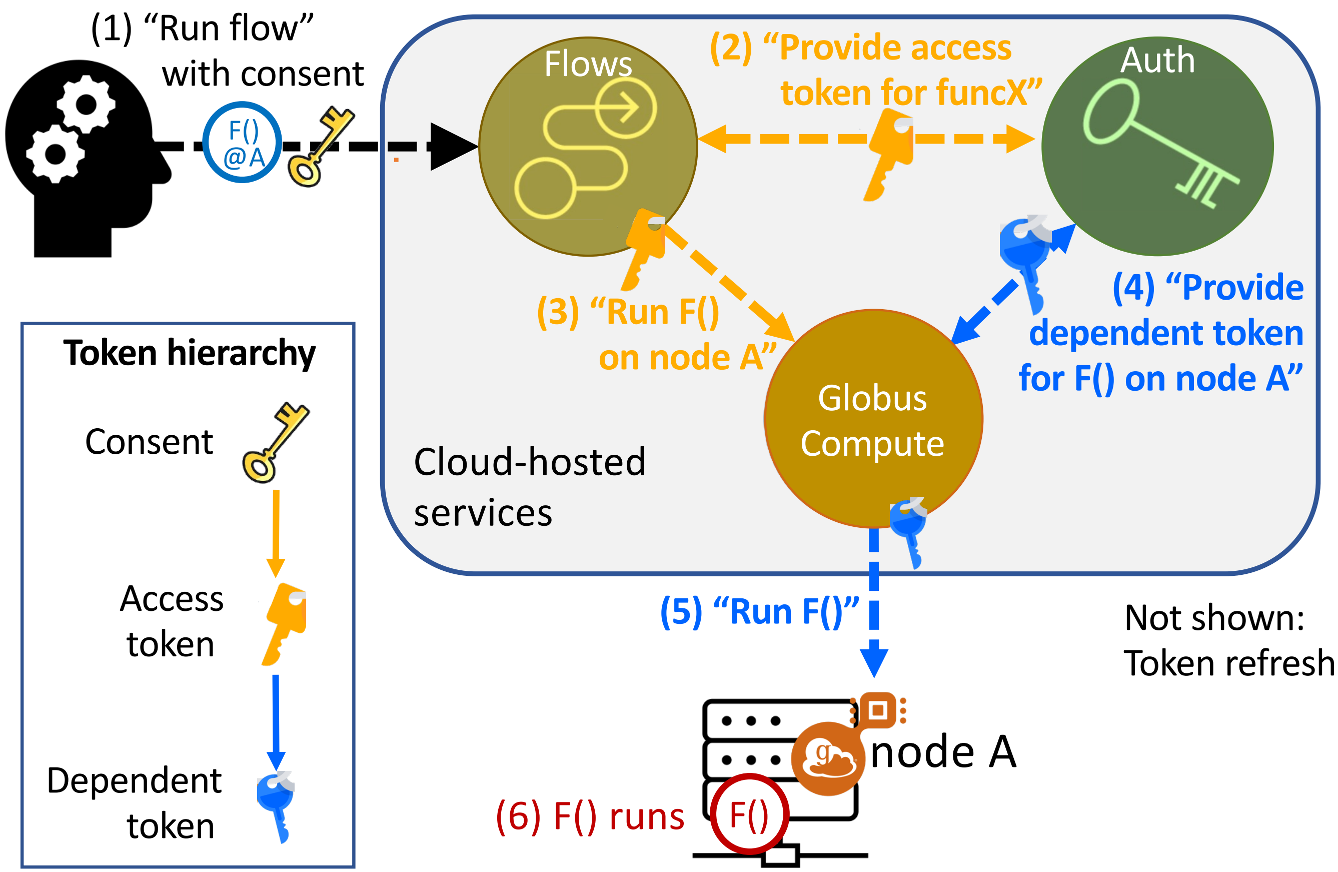

S. Tuecke et al., https://doi.org/10.1109/eScience.2016.7870901, R. Chard et al., https://doi.org/10.48550/arXiv.2208.09513

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 **managed research acceleration services** buffer against inevitable failures

3) Manage who is **trusted** to perform what actions, where and when

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

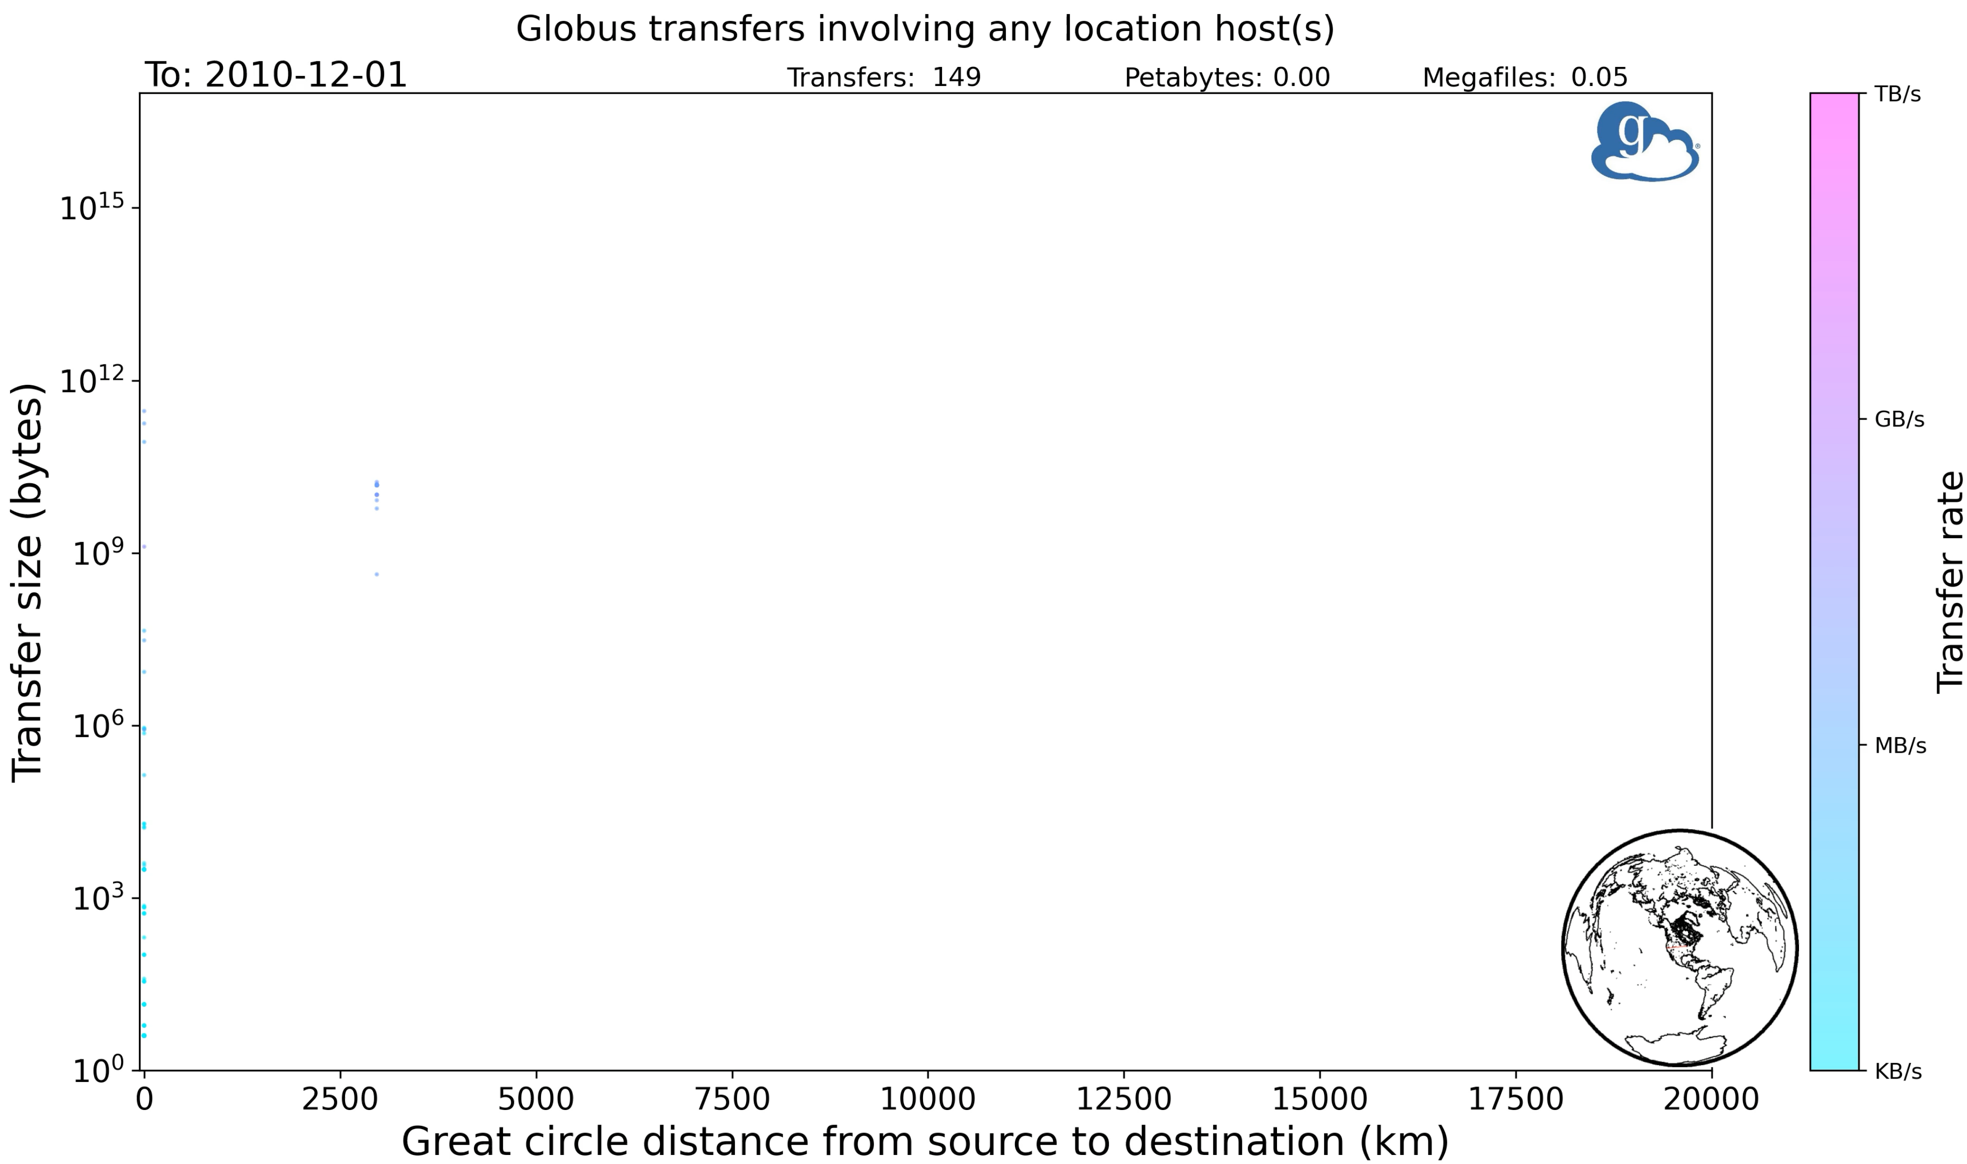

Globus transfers involving any location host(s)
To: 2010-12-01
Transfers: 149
Petabytes: 0.00
Megafiles: 0.05
Transfer size (bytes)
10^{15}
10^{12}
10^{9}
10^{6}
10^{3}
10^{0}
0
2500
5000
7500
10000
12500
15000
17500
20000
Great circle distance from source to destination (km)
TB/s
GB/s
MB/s
KB/s
Transfer rate

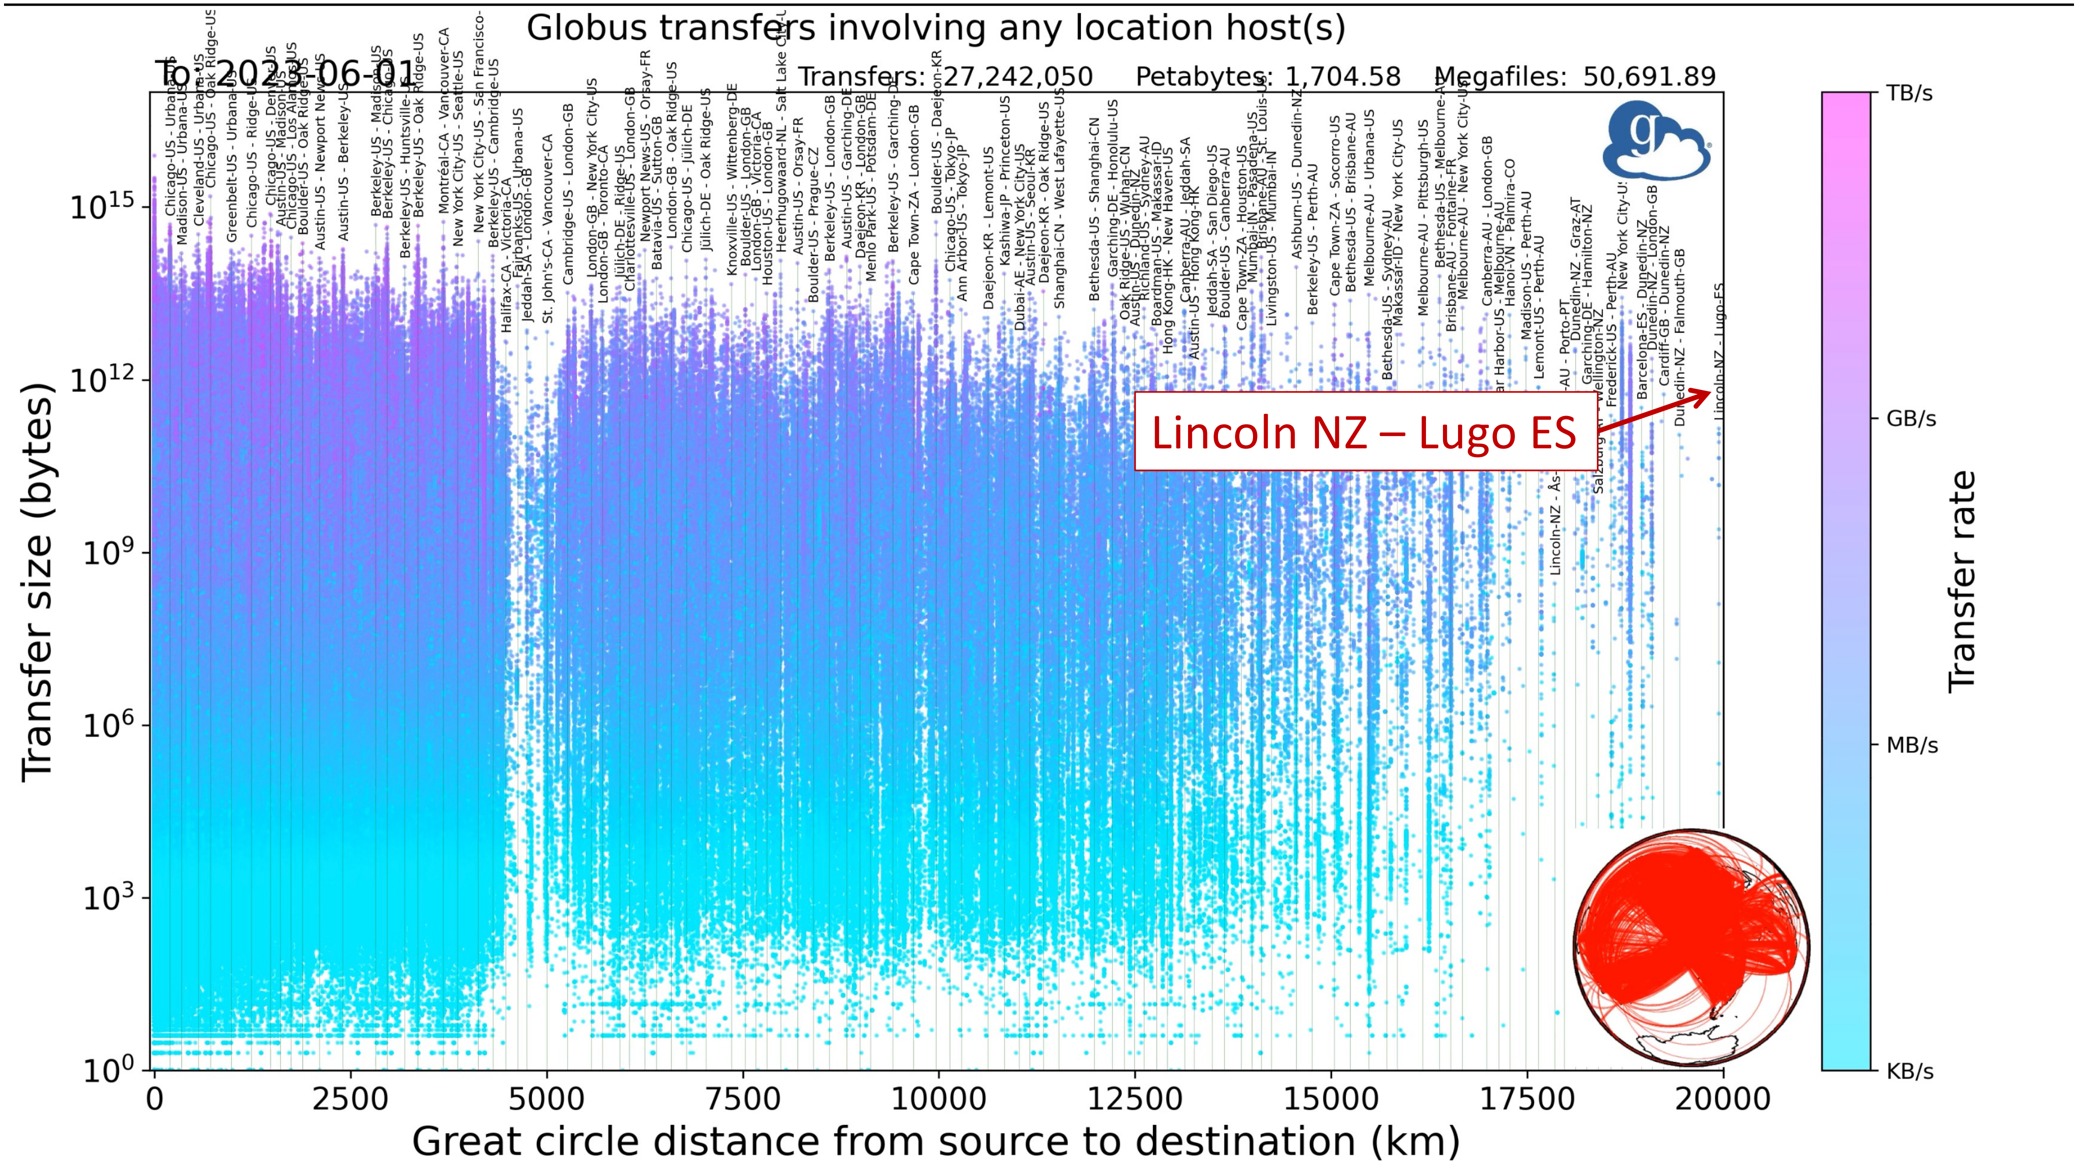
Globus transfers involving any location host(s)
To: 2023-06-01
Transfers: 27,242,050
Petabytes: 1,704.58
Megafiles: 50,691.89
Lincoln NZ – Lugo ES
Transfer size (bytes)
Great circle distance from source to destination (km)
Transfer rate
TB/s
GB/s
MB/s
KB/s

Lincoln, New Zealand
to Lugo, Spain:
19,932 km

Circumference of the earth:
40,075 km
Semi-circumference:
20,037 km

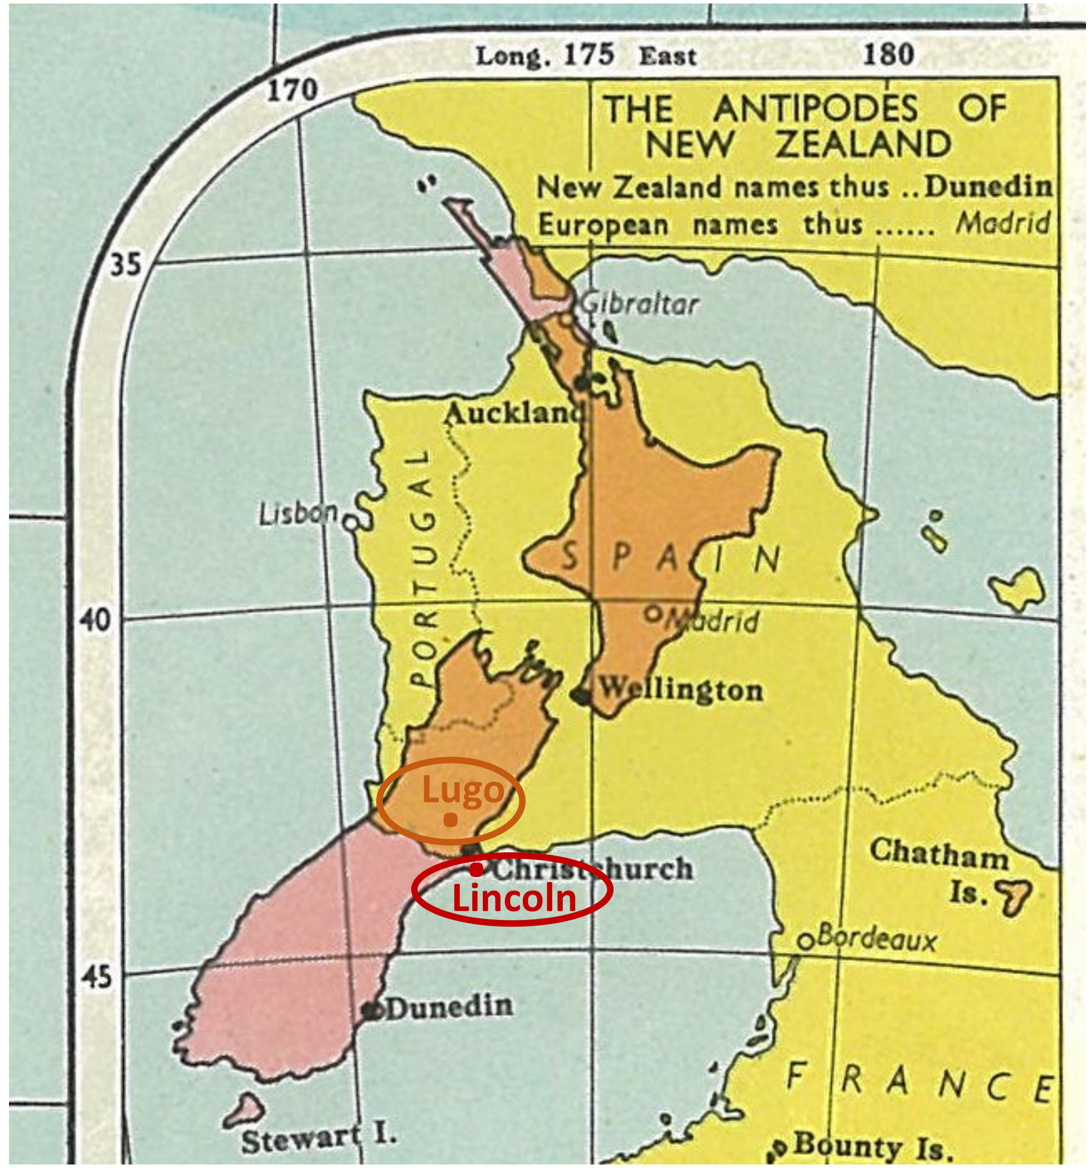

Collins Australian Clear School Atlas, 1964

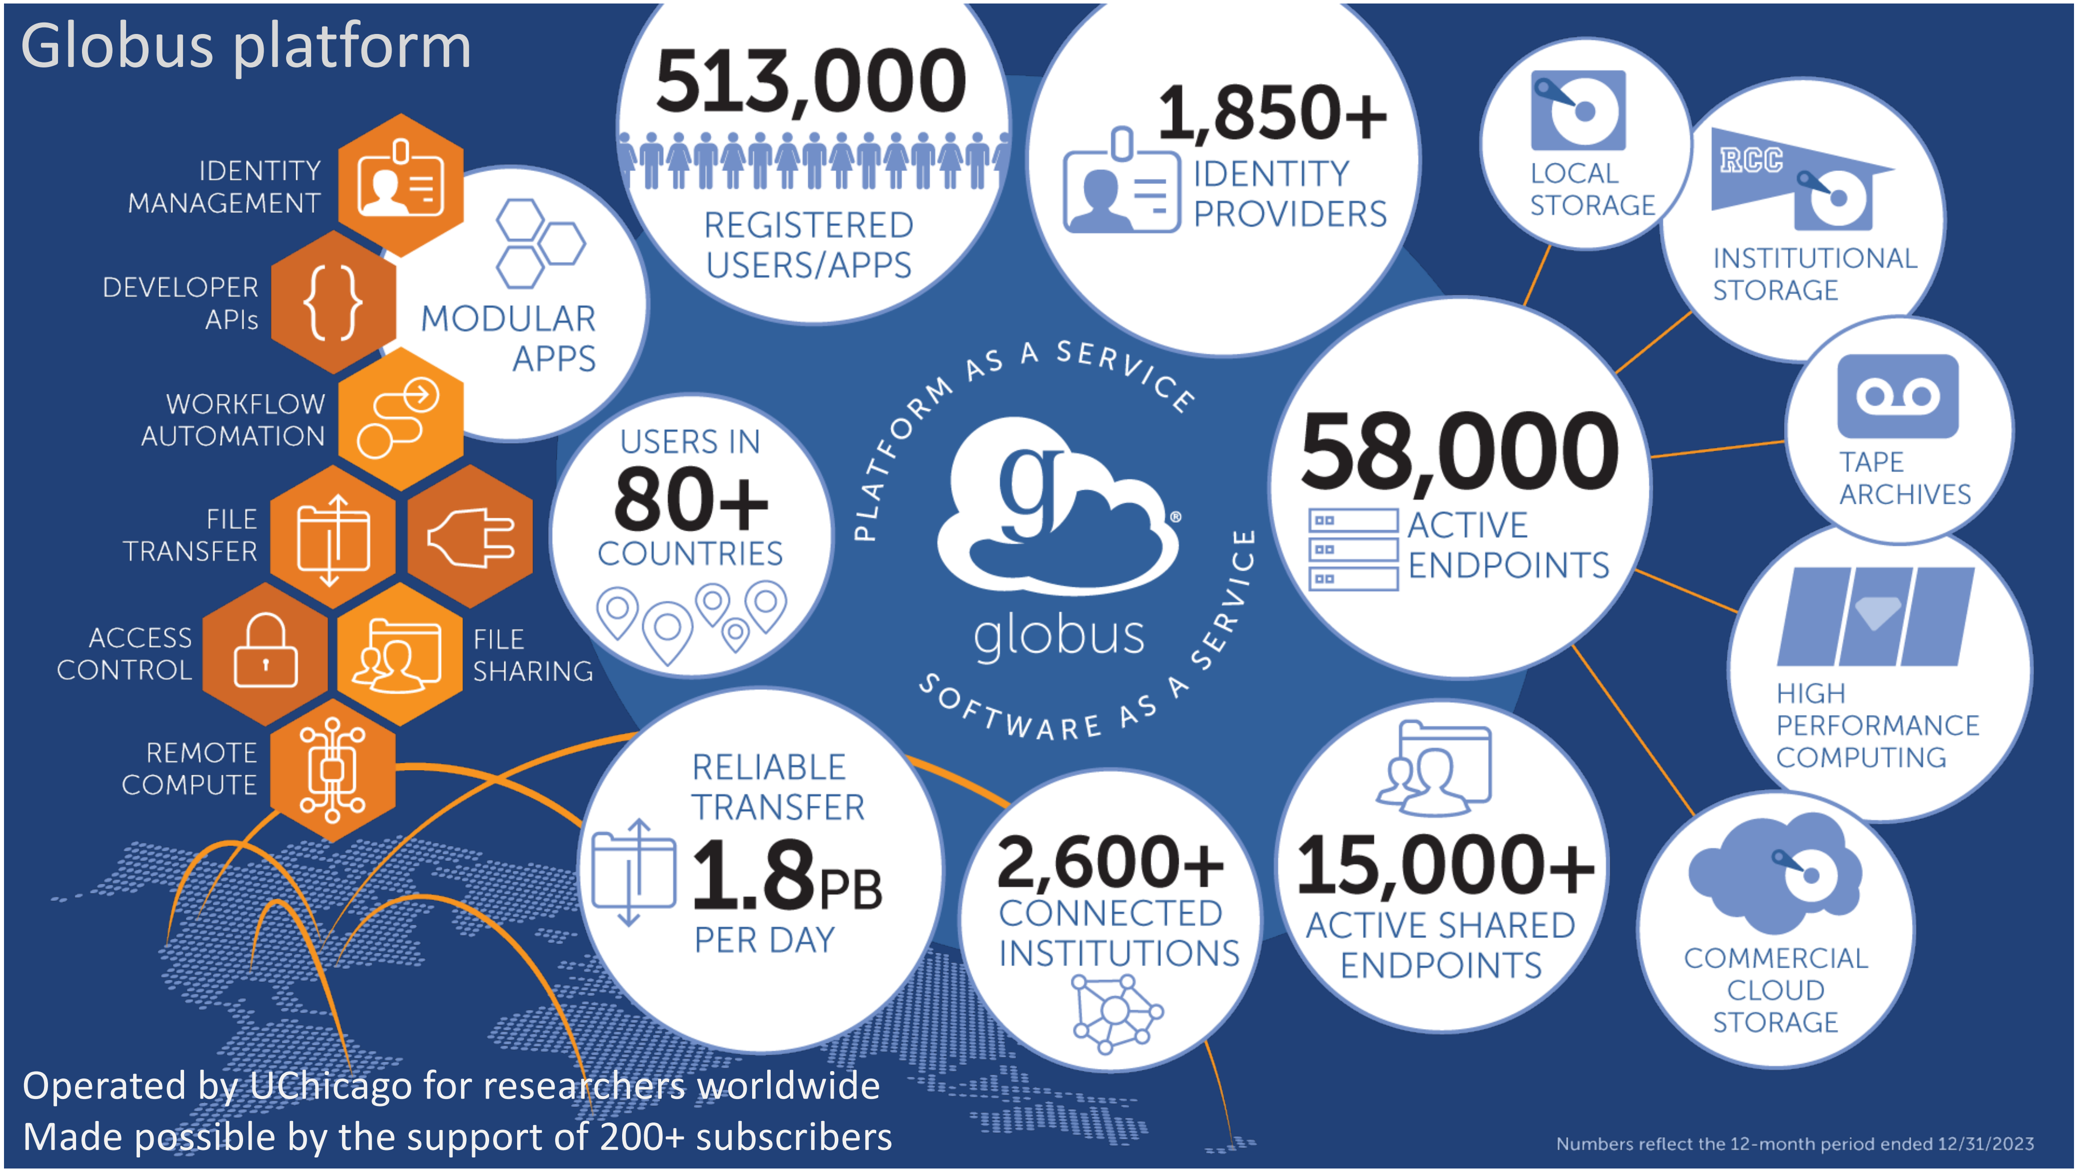
Globus platform
IDENTITY MANAGEMENT
DEVELOPER APIs
WORKFLOW AUTOMATION
FILE TRANSFER
ACCESS CONTROL
FILE SHARING
REMOTE COMPUTE
MODULAR APPS
513,000 REGISTERED USERS/APPS
1,850+ IDENTITY PROVIDERS
USERS IN 80+ COUNTRIES
PLATFORM AS A SERVICE
SOFTWARE AS A SERVICE
globus
58,000 ACTIVE ENDPOINTS
LOCAL STORAGE
RCC
INSTITUTIONAL STORAGE
TAPE ARCHIVES
HIGH PERFORMANCE COMPUTING
COMMERCIAL CLOUD STORAGE
RELIABLE TRANSFER 1.8PB PER DAY
2,600+ CONNECTED INSTITUTIONS
15,000+ ACTIVE SHARED ENDPOINTS
Operated by UChicago for researchers worldwide
Made possible by the support of 200+ subscribers
Numbers reflect the 12-month period ended 12/31/2023

Endpoints at an experimental facility

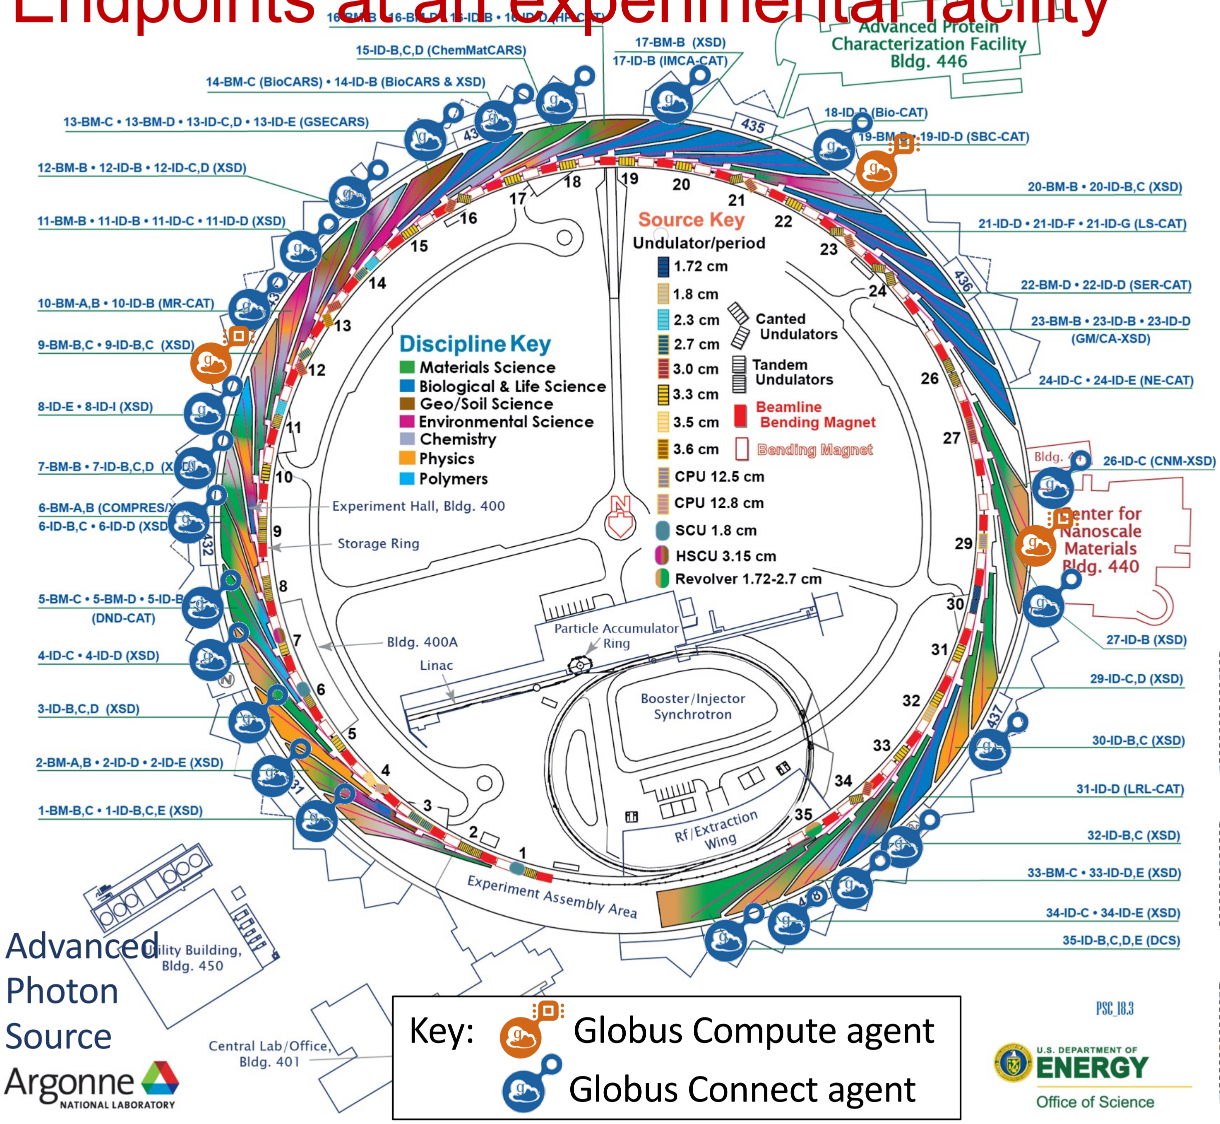

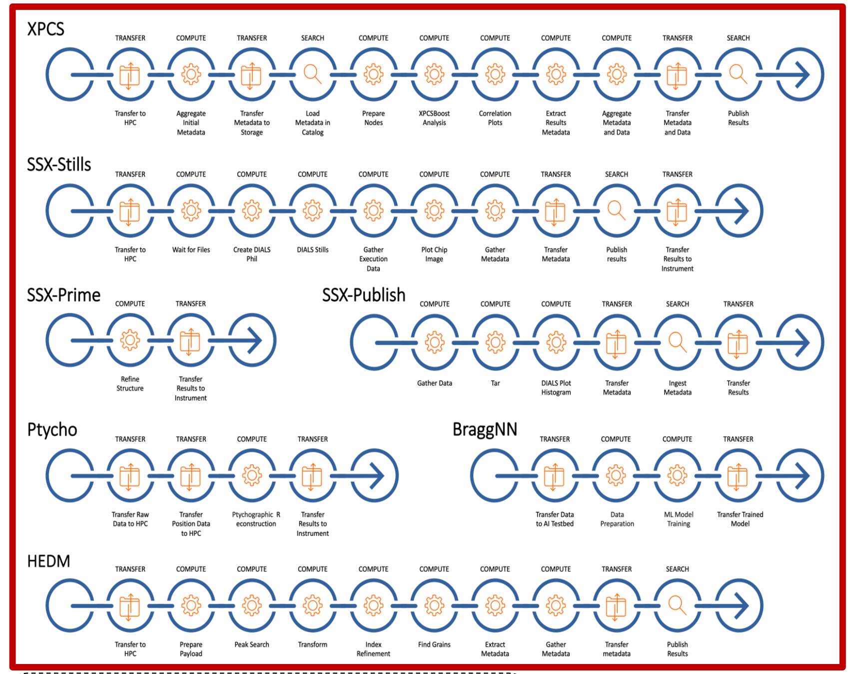

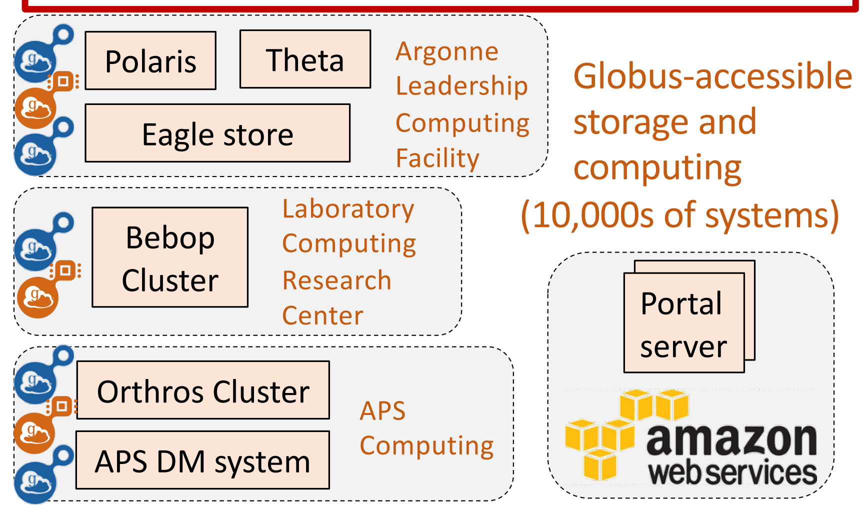

ESGF2

Lawrence Livermore National Laboratory

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

ESnet ENERGY SCIENCES NETWORK

globus

https://dashboard.globus.org/esgf

As of May 4, 2022

Cumulative data transferred over time

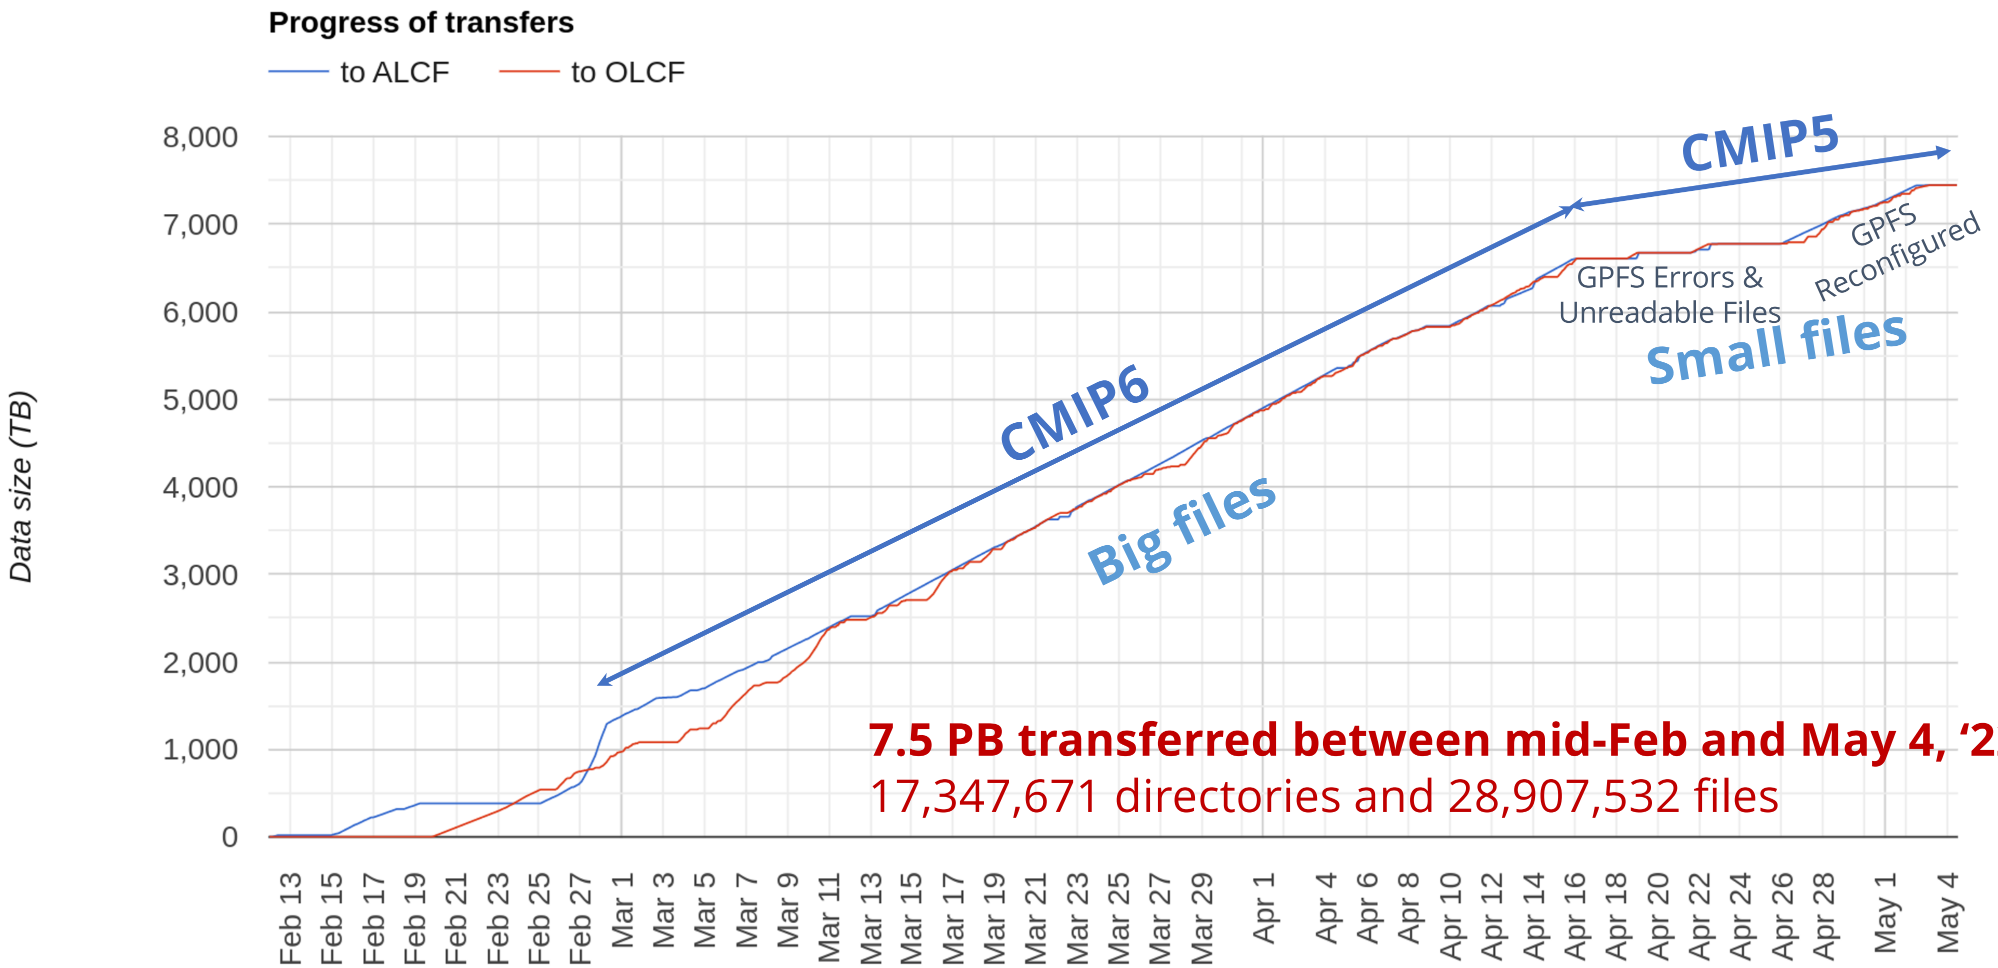

7.5 PB transferred between mid-Feb and May 4, '22

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

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
```

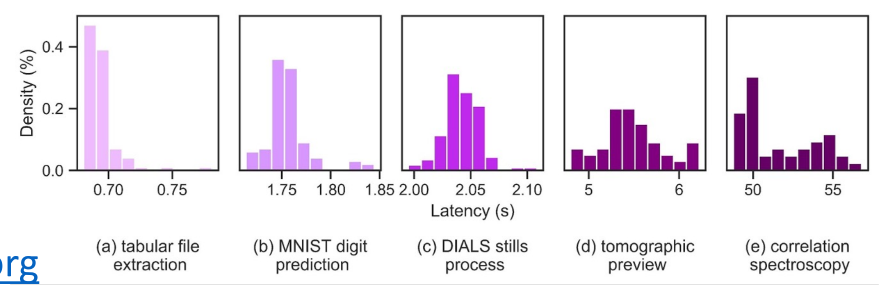

Z. Li et al., https://doi.org/10.48550/arXiv.2209.11631 https://funcx.org

Globus Compute application: Privacy Preserving Learning

COVID19 detection from chest X-rays

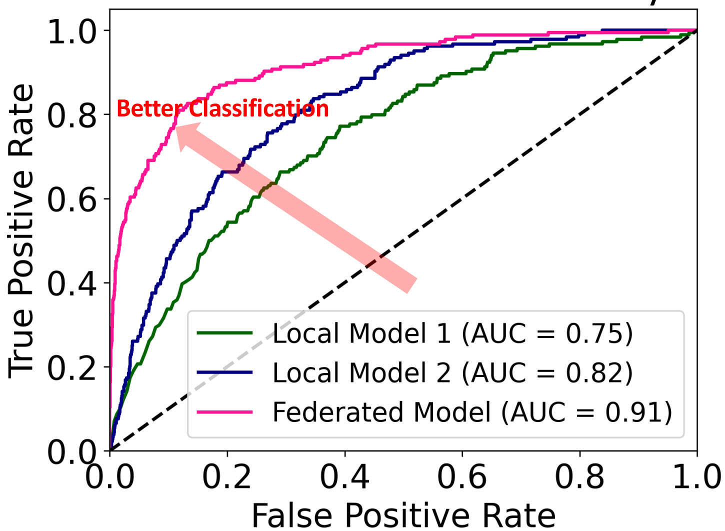


Ravi Madduri


Kibaek Kim


Ryu Minseok


Eliu Huerta

Pradeep Natarajan

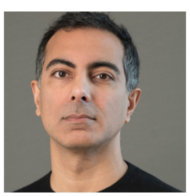
Puneet Batra


Maryellen Giger

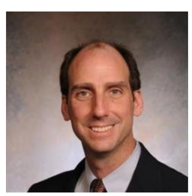
Sam Armato

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				

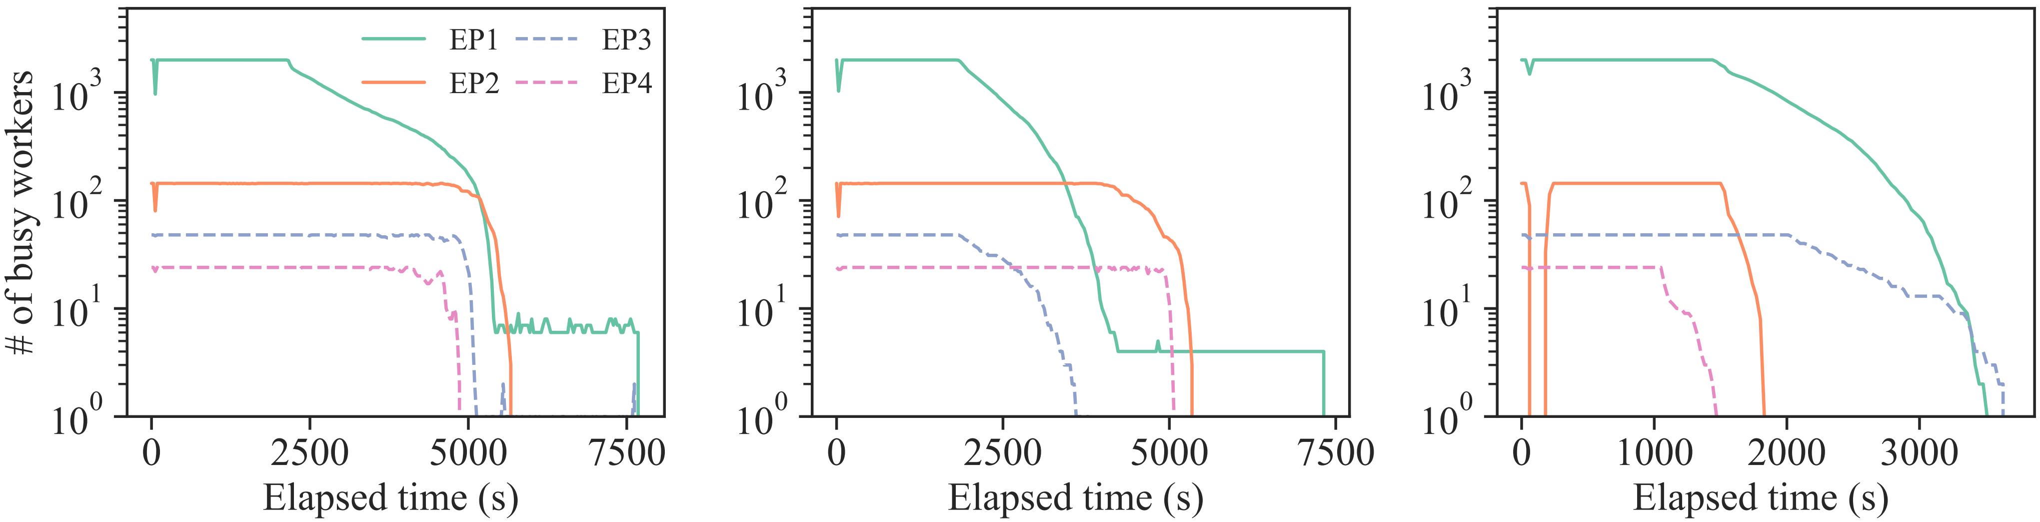

(a) Default (b) Locality-aware (c) Capacity-aware

Distributed execution of 30,000-task drug-screening protocol over multiple HPC systems, with different scheduling heuristics

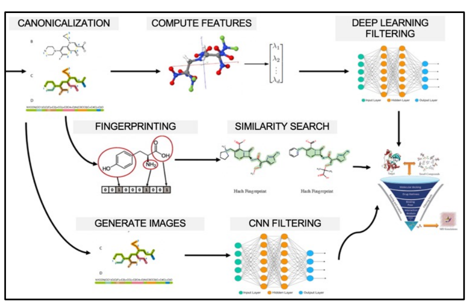

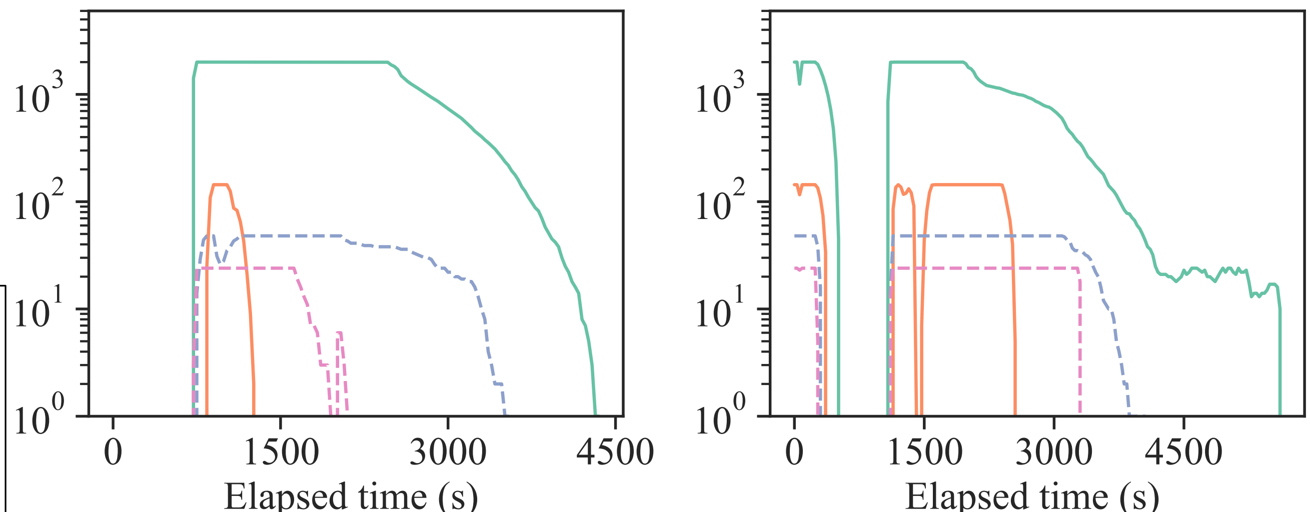

(d) Heterogeneity-aware (e) Dynamic-adjusting

Z. Li et al., 2023

"Globus Flows": Hosted research supervision services for **flow specification, execution, and management**

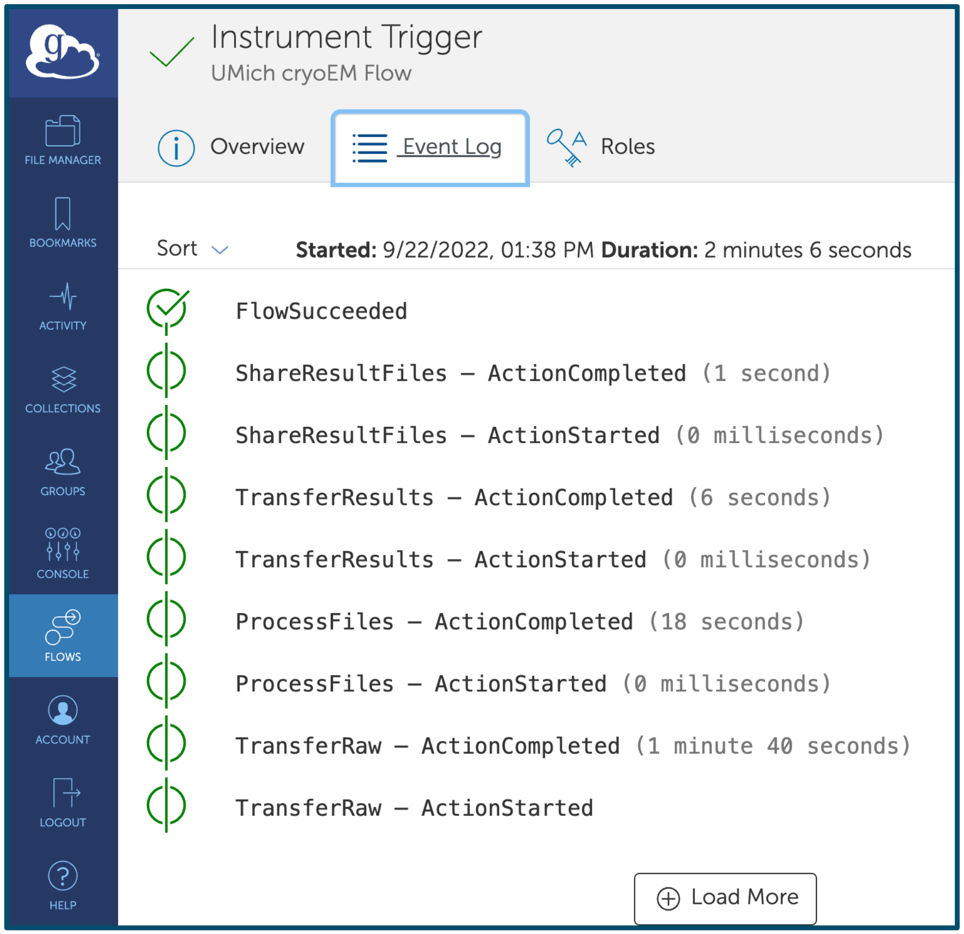

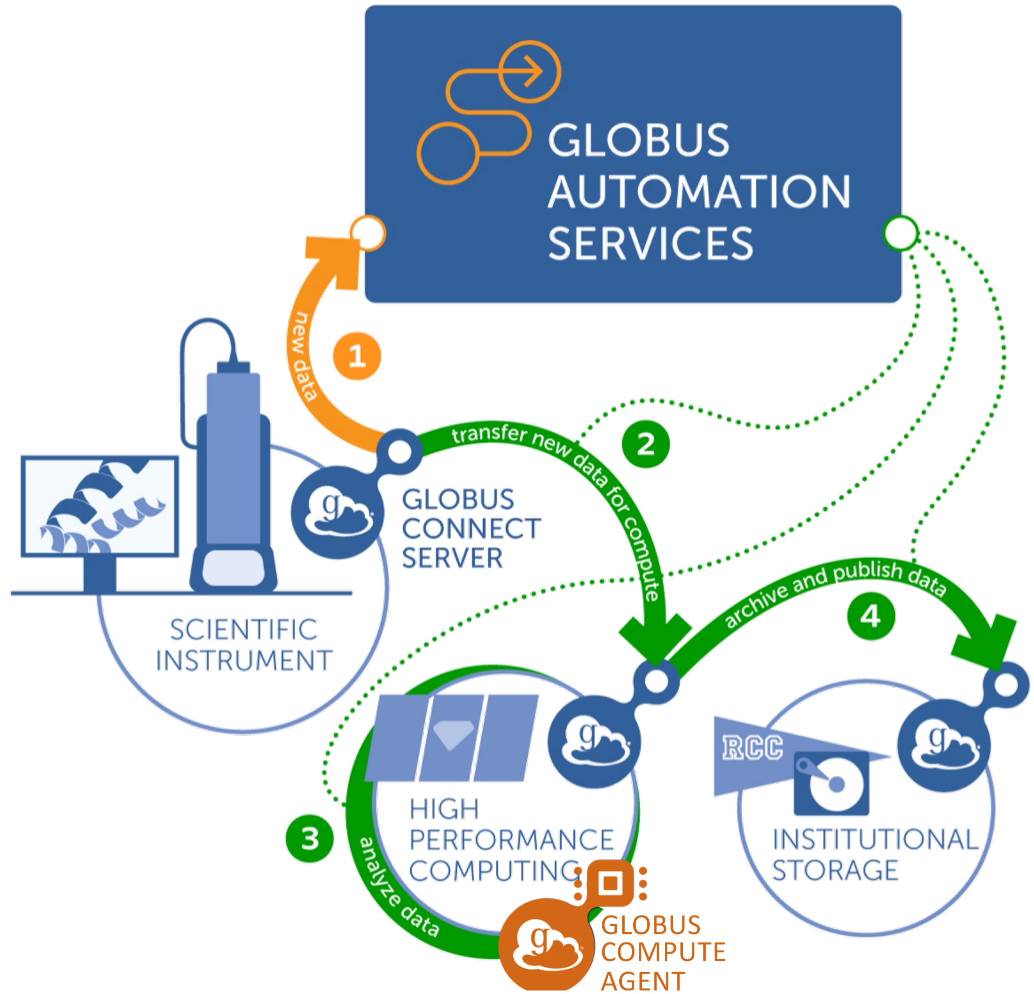

R. Chard et al., https://doi.org/10.48550/arXiv.2208.09513

Flows enable creation of smart instruments

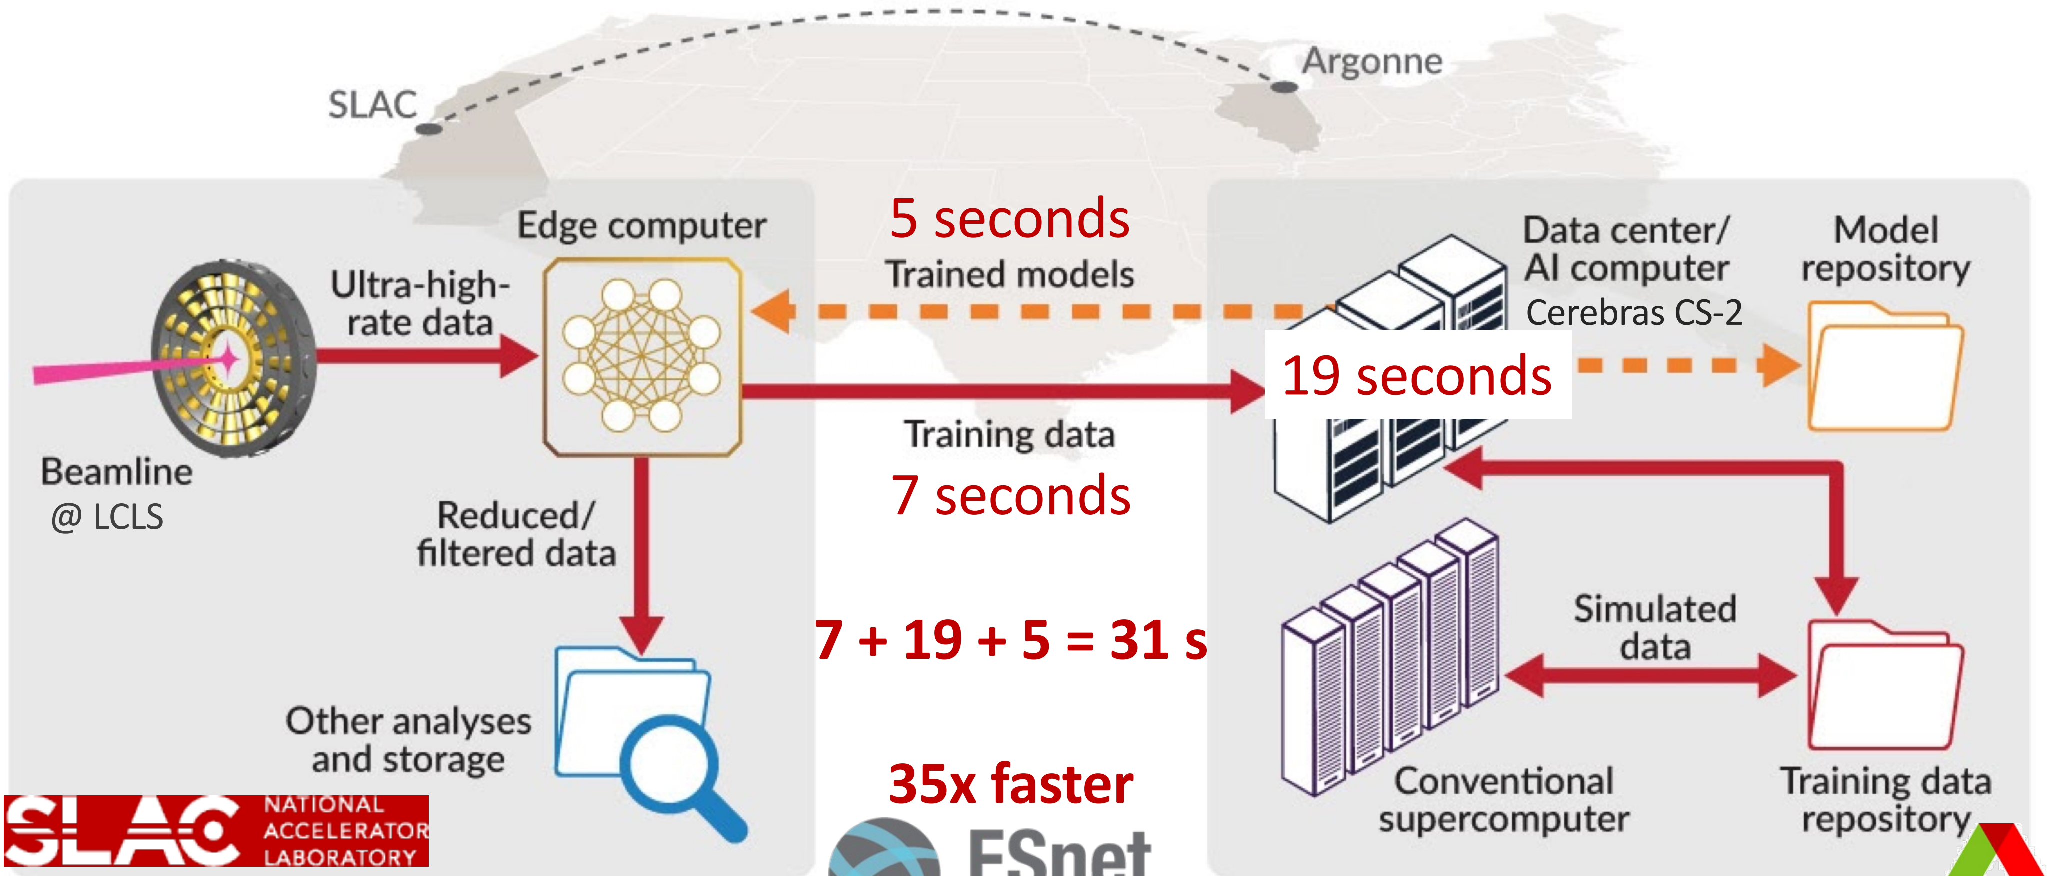

SLAC NATIONAL ACCELERATOR LABORATORY

ESnet ENERGY SCIENCES NETWORK

Argonne NATIONAL LABORATORY

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

Flows enable creation of smart instruments

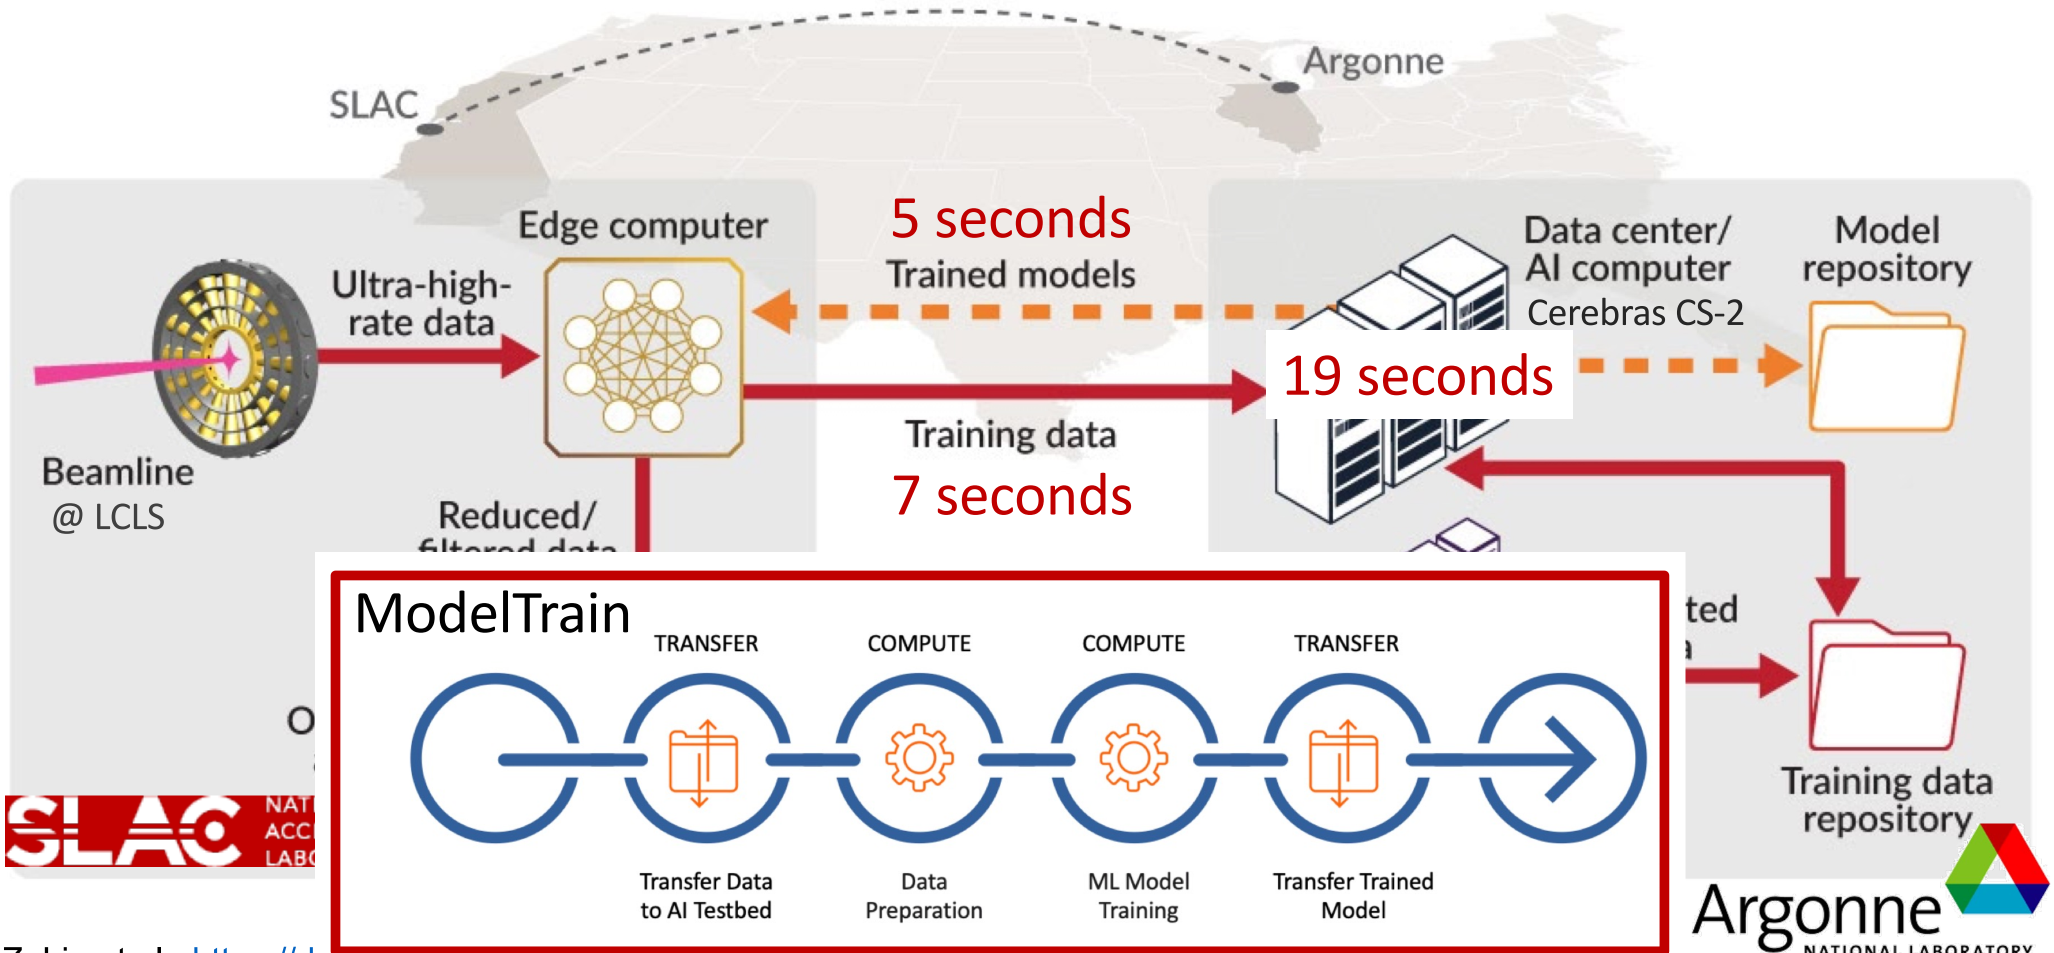

Z. Liu et al., https://d

What else can we do with a global computing fabric?

Realize embodied agents

Persistent, stateful

Sense, act

A **computational system** that can **interact** with its **environment** and **learn** from those interactions

Physical or virtual

Adapt responses

What else can we do with a global computing fabric?

Integrate with large language model (LLM) technologies

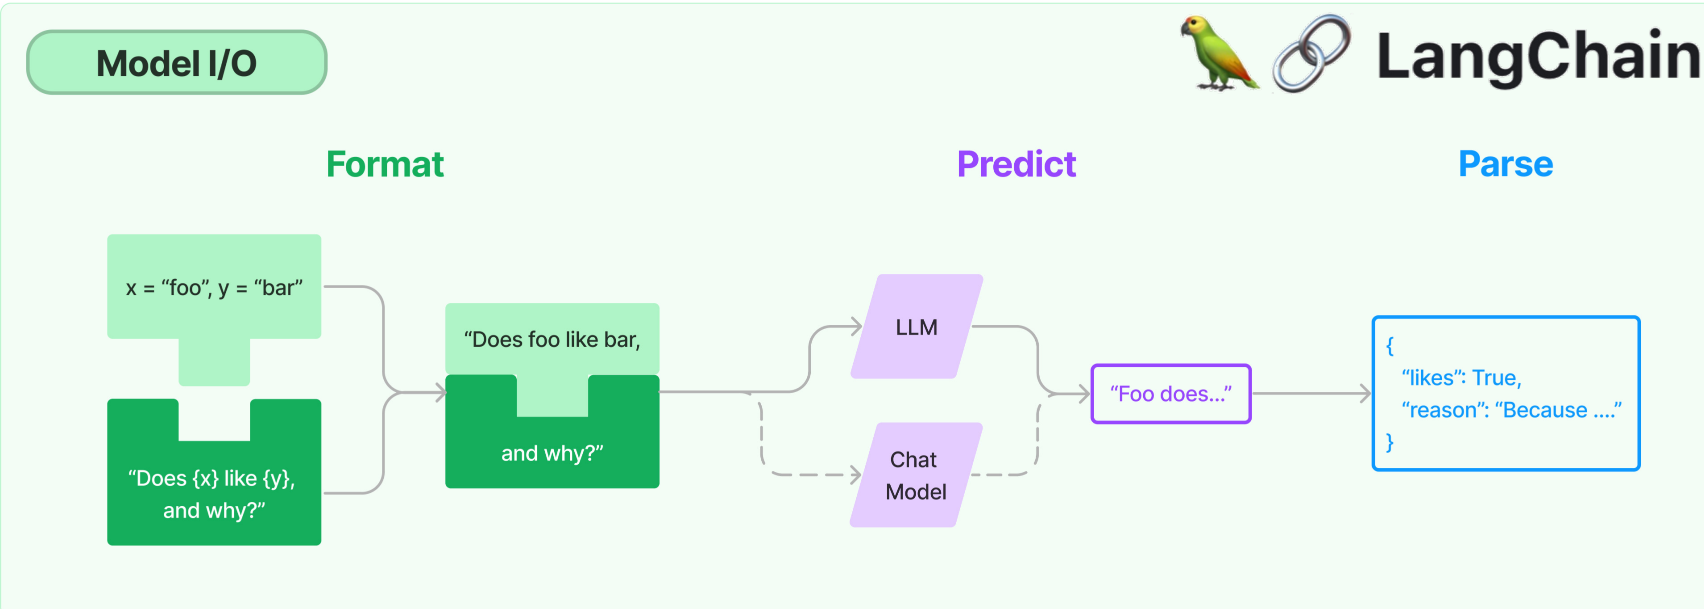

What else can we do with a global computing fabric?

Integrate with large language model (LLM) technologies

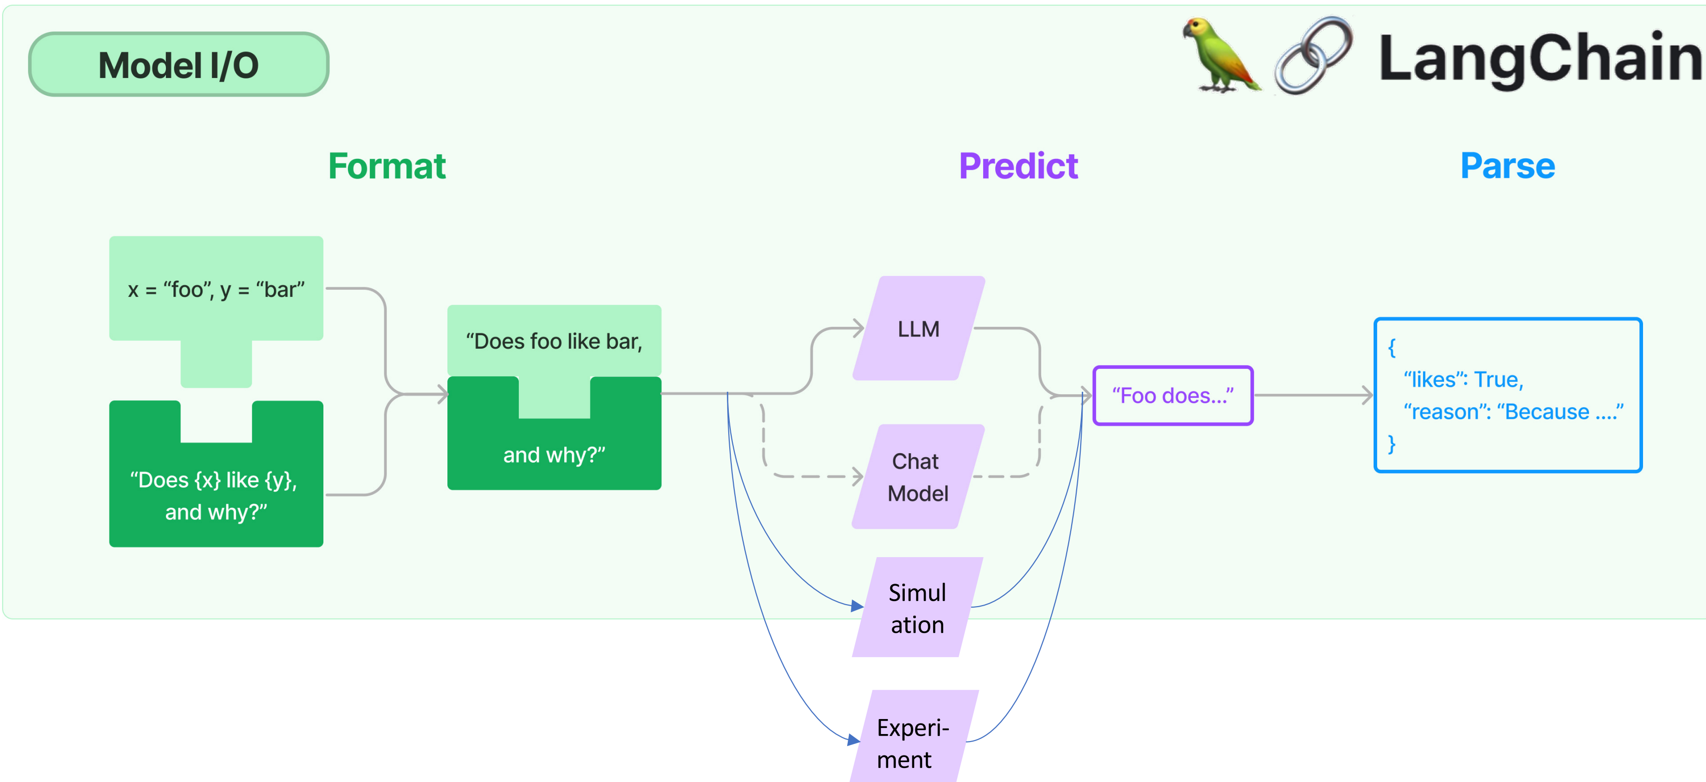

Example: Design of antimicrobial peptides

Set of peptides as input

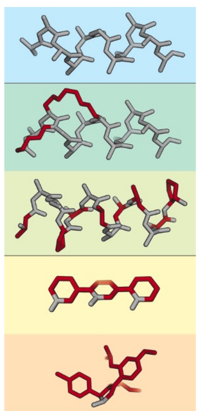

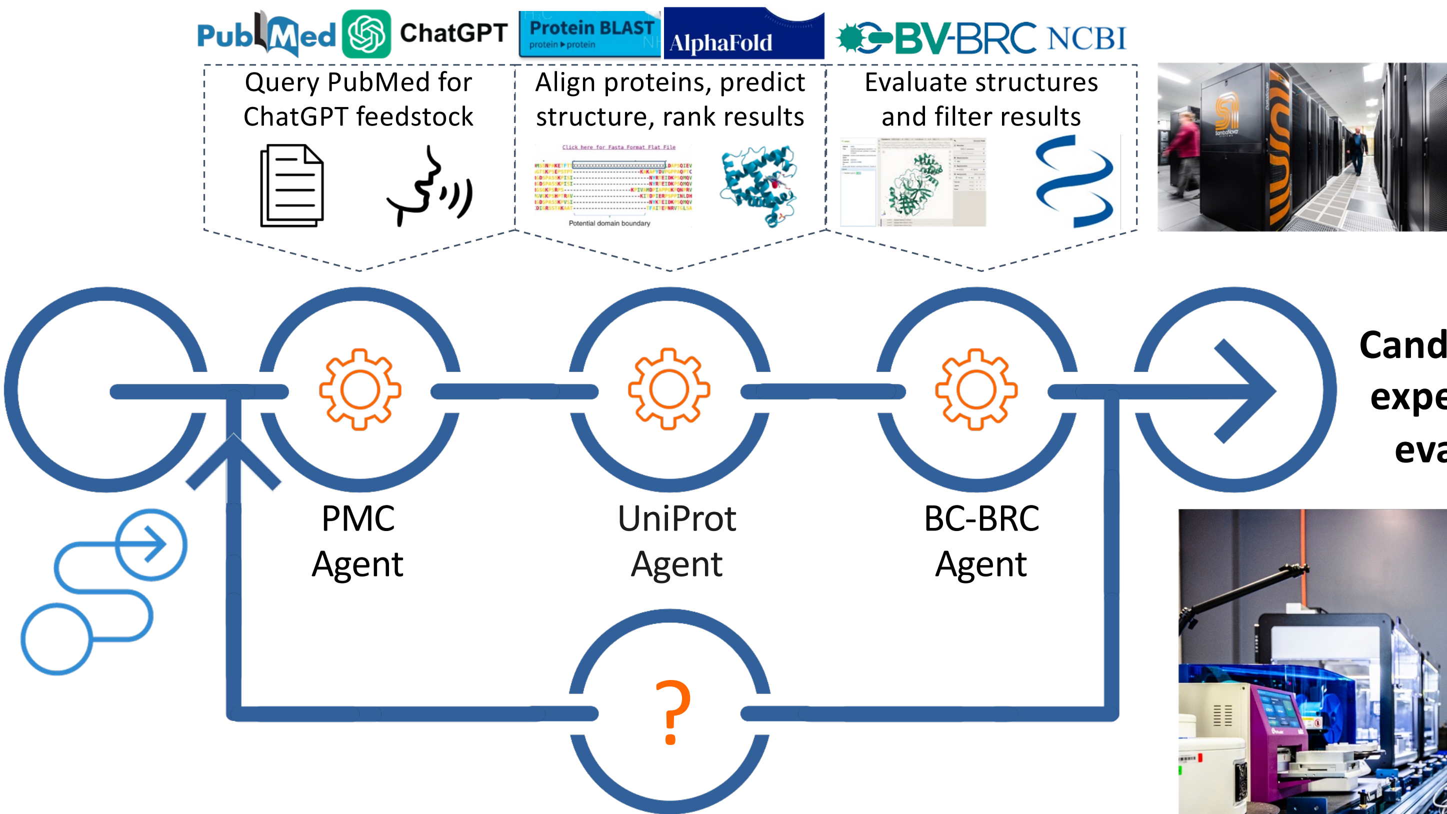

Agents run on HPC/AI resources

Candidates for experimental evaluation

Self-driving lab performs experiments

Arvind Ramanathan, Priyanaka Setty, et al.

Questions

- It used to be that computers were reliable and networks were unreliable. Now computers are unreliable and “networks” are reliable. How does that change what we do, and how?

- Global science services make it trivial to write programs that integrate resources at global scale. What new programs do we want to write? What new instruments can we create?

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

Patterns
CellPress OPEN ACCESS

https://doi.org/10.1016/j.patter.2022.100606

Article

Linking scientific instruments and computation: Patterns, technologies, and experiences

Rafael Vescovi,[1] Ryan Chard,[1] Nickolaus D. Saint,[6] Ben Blaiszik,[1,6] Jim Pruyne,[1,6] Tekin Bicer,[1,3] Alex Lavens,[4] Zhengchun Liu,[1] Michael E. Papka,[2,7] Suresh Narayanan,[3] Nicholas Schwarz,[3] Kyle Chard,[1,5] and Ian T. Foster[1,5,*]

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, Logan Ward, & **many more**