





# Live Streaming of Large Electron Microscope Data to NERSC

Sam Welborn, Bjoern Enders, Peter Ercius, Chris Harris, Deborah Bard



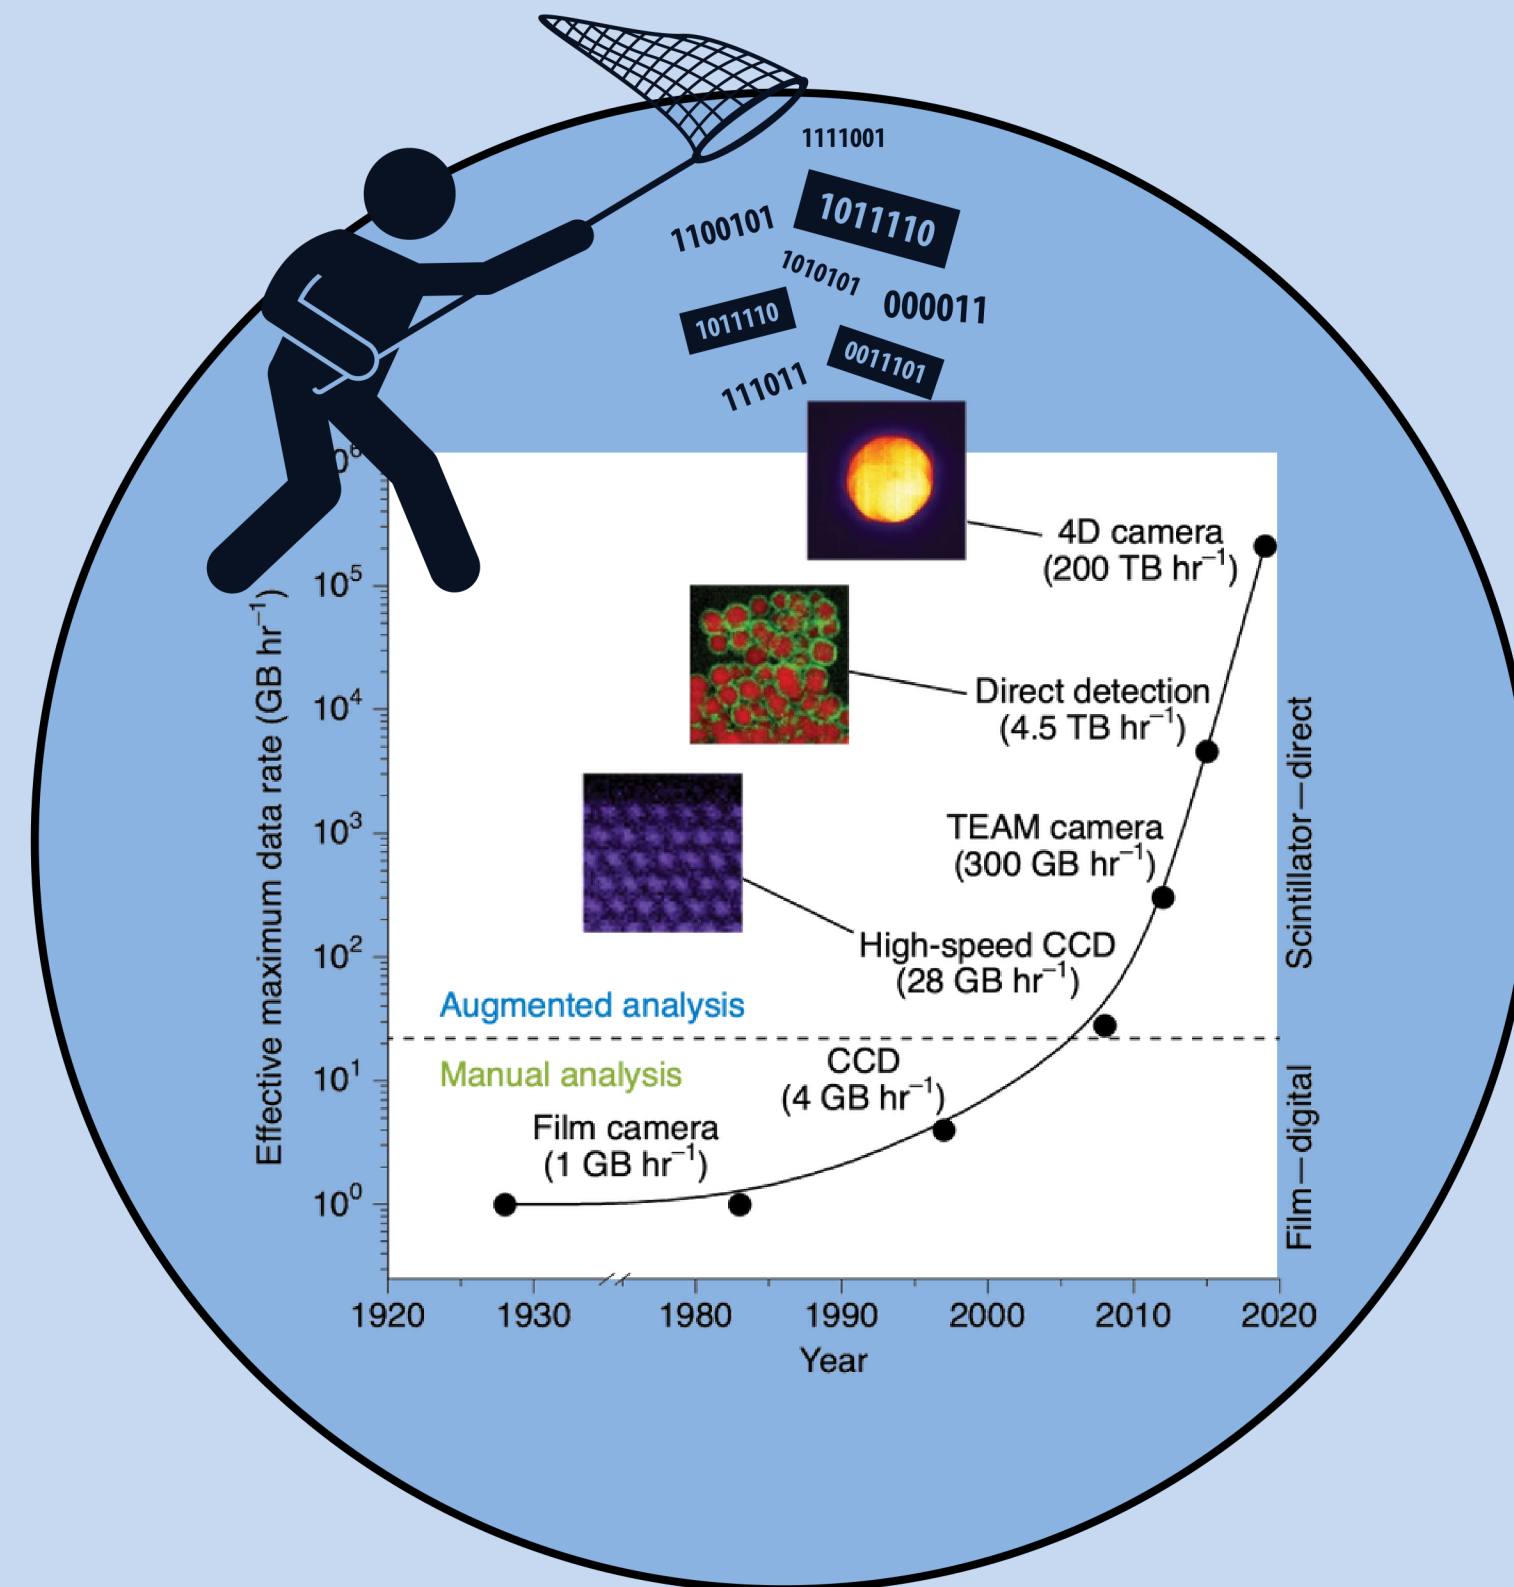


# The data deluge

## Faster detectors have led to unprecedented data generation rates

- Lightsources (e.g., ALS, NSLS-II, APS) will cumulatively generate an **exabyte** of data per year by 2028
- 10-1000 PFLOPs peak on-demand compute resources


Schwarz, N. et al. doi.org/10.1007/978-3-030-63393-6\_10



Read more

Spurgeon et al. Nature Materials volume 20, pages 274–279 (2021).







Berkeley Lab


**NERSC**

**MOLECULAR  
FOUNDRY**




100 ft






# The Molecular Foundry/NCEM



**National Center for  
Electron Microscopy**  
one of seven facilities in TMF

## Instrumentation


- State-of-the-art methods and instrumentation in nanoscale science free of charge



## Expertise


- Leaders in:

**4DSTEM,** high resolution,  
image simulation, tomography,  
**electron detector** in situ,  
**technology** soft materials



# STEM

Scanning Transmission  
Electron Microscopy



converged electron probe

sample

annular dark field (ADF)  
detector


diffraction pattern

pixelated detector


2D images recorded over a 2D grid of probe  
positions:

Four dimensional  
scanning transmission electron microscopy  
**(4D-STEM)**






# The 4D Camera



**TEAM 0.5**



**4D Camera**

Generates data at 480 Gb/s  
700 GB dataset collected in 15 seconds




What is that?



Something your models didn't predict...








# File transfer workflow

User






# File transfer workflow




TEAM 0.5



4D Camera







# File transfer workflow




TEAM 0.5



4D Camera




→ 120 Gb/s link → 20 Gb/s link




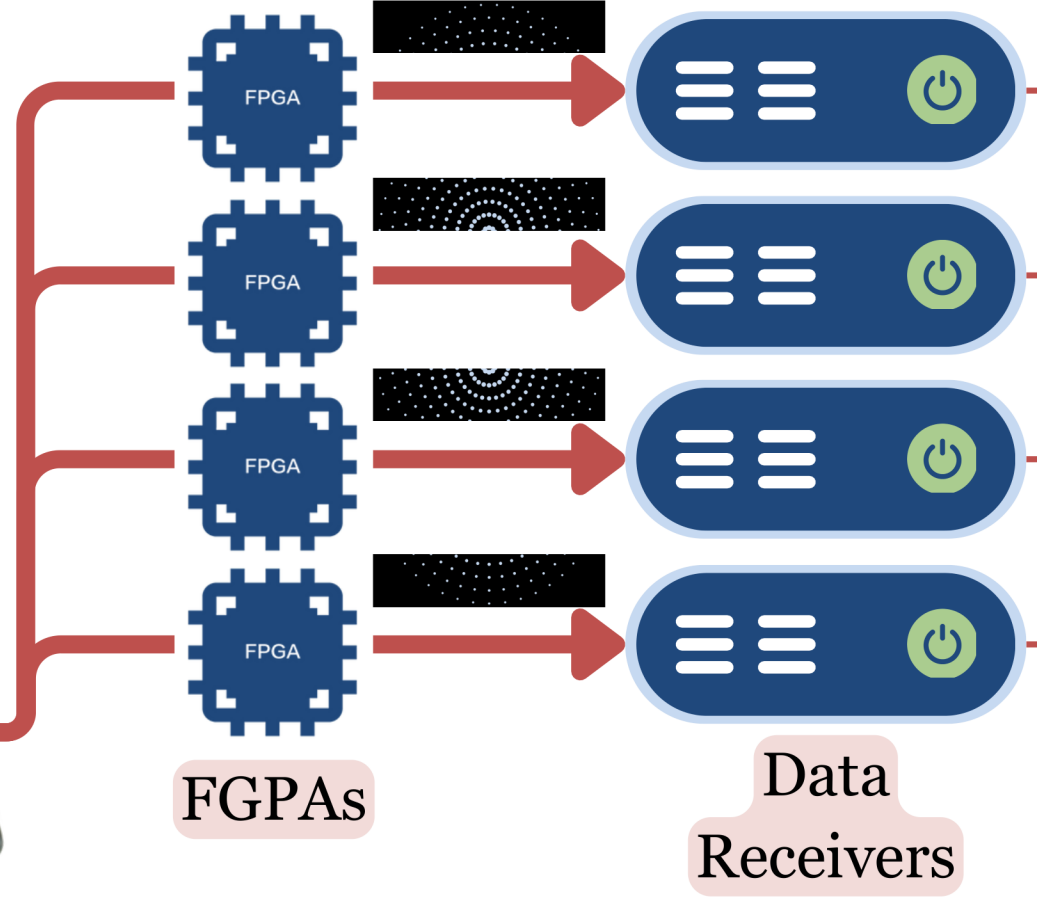

# File transfer workflow



TEAM 0.5




4D Camera



Writing data at NCEM







# File transfer workflow




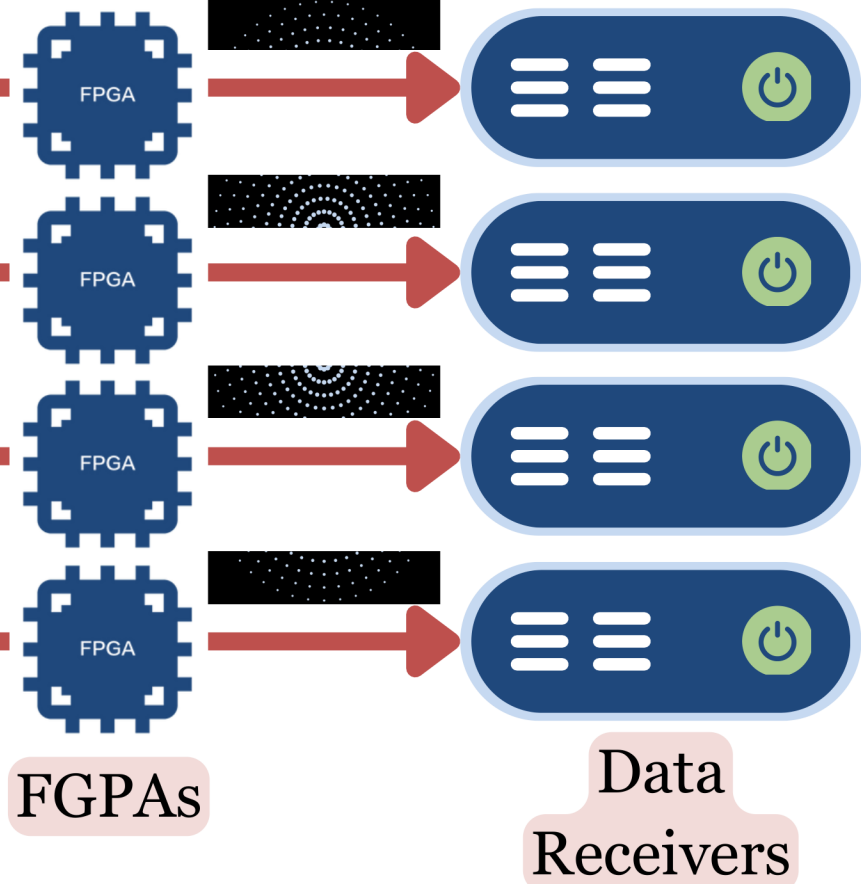

Distiller



TEAM 0.5




4D Camera



Writing data  
at NCEM

SSD

120 Gb/s link → 20 Gb/s link



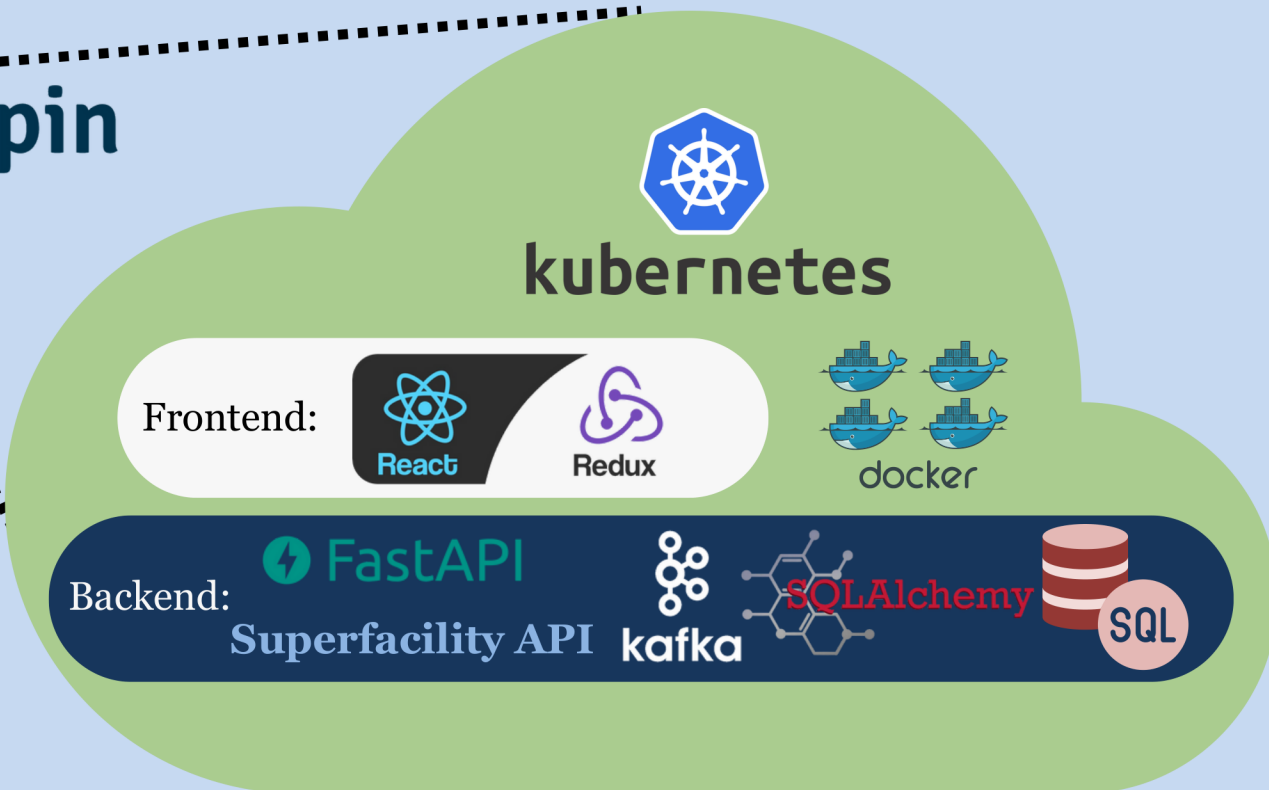
# File transfer workflow




Distiller




Spin



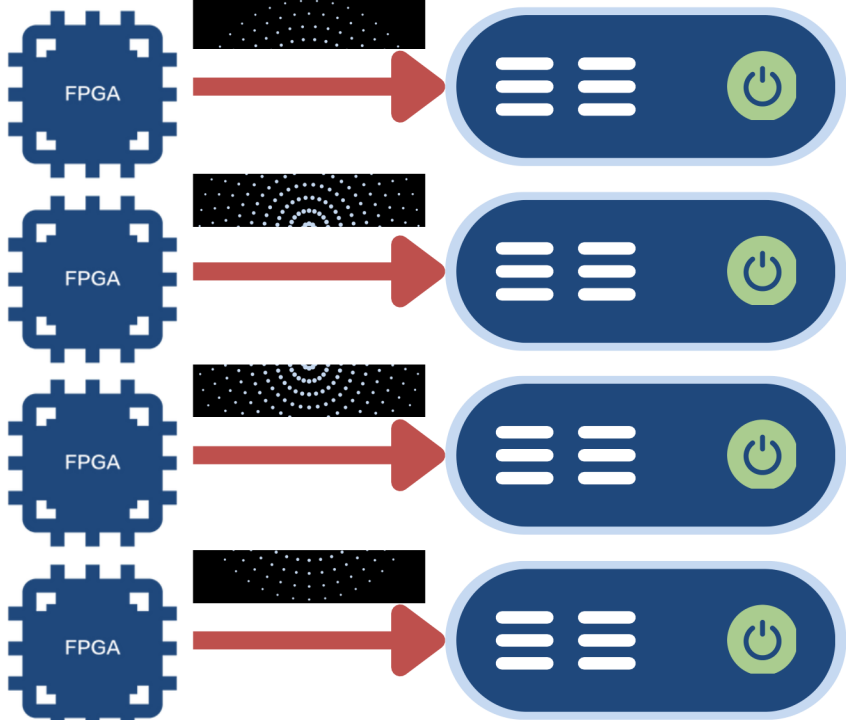
TEAM 0.5




4D Camera



FGPAs




Data Receivers



Writing data at NCEM



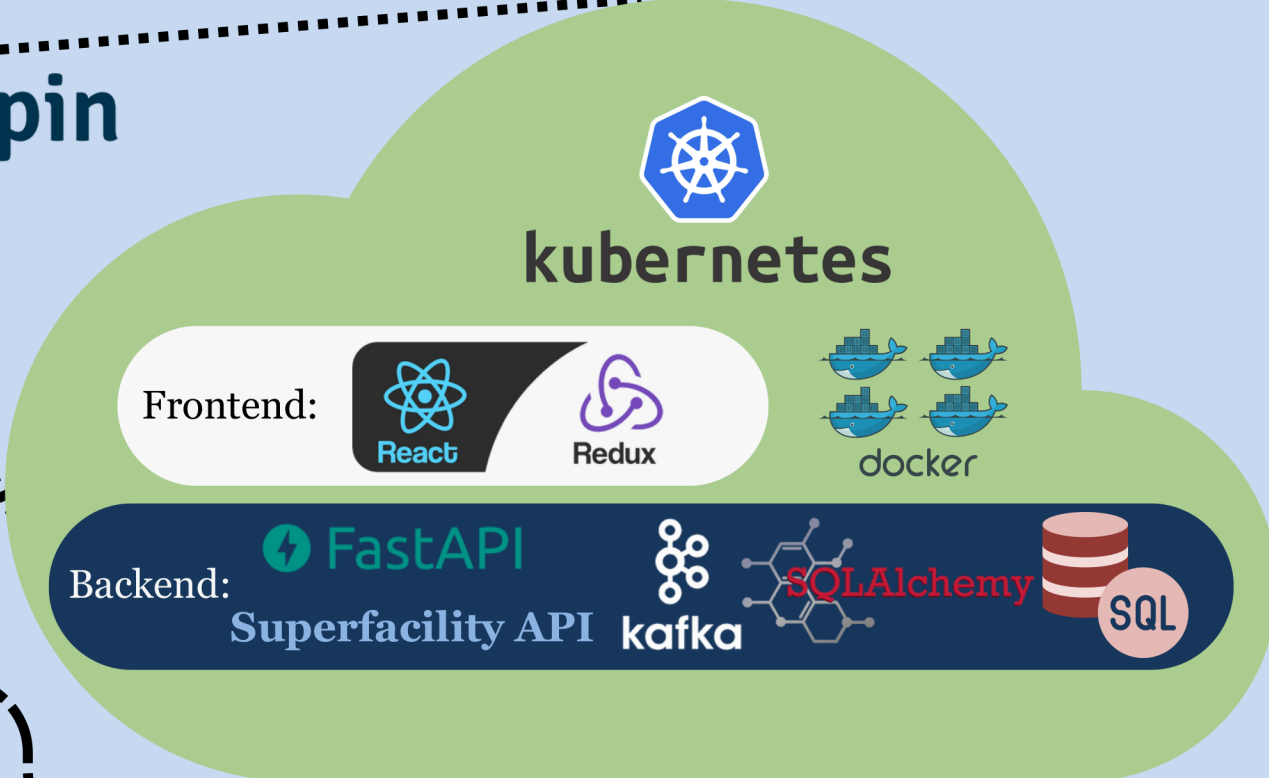


# File transfer workflow




Distiller


Spin




TEAM 0.5




4D Camera




FGPAs



Data Receivers

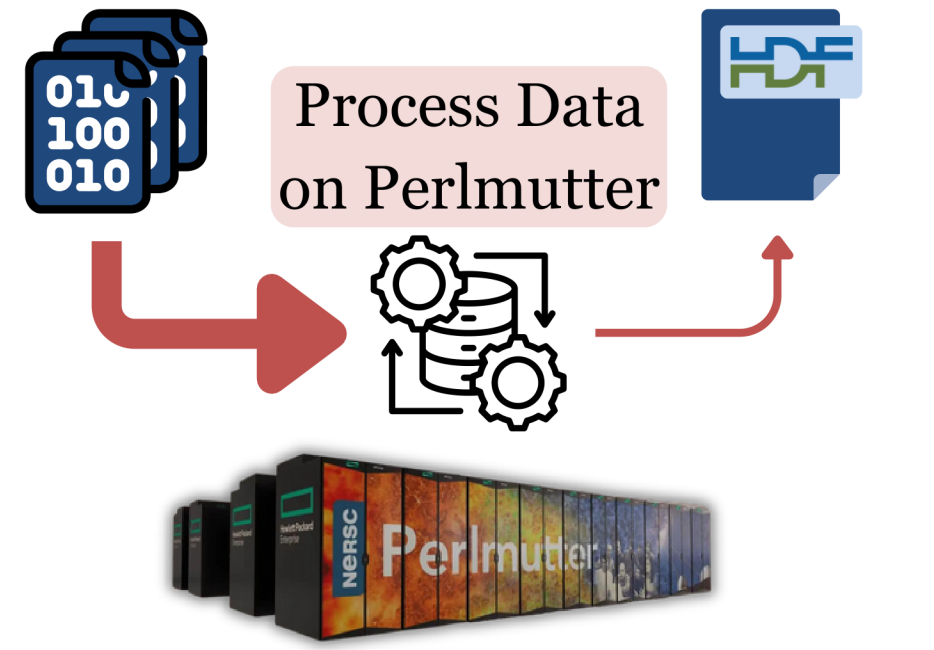


Writing data at NCEM




Raw data

Reduced Data



Process Data on Perlmutter




**NERSC**

**MOLECULAR FOUNDRY**


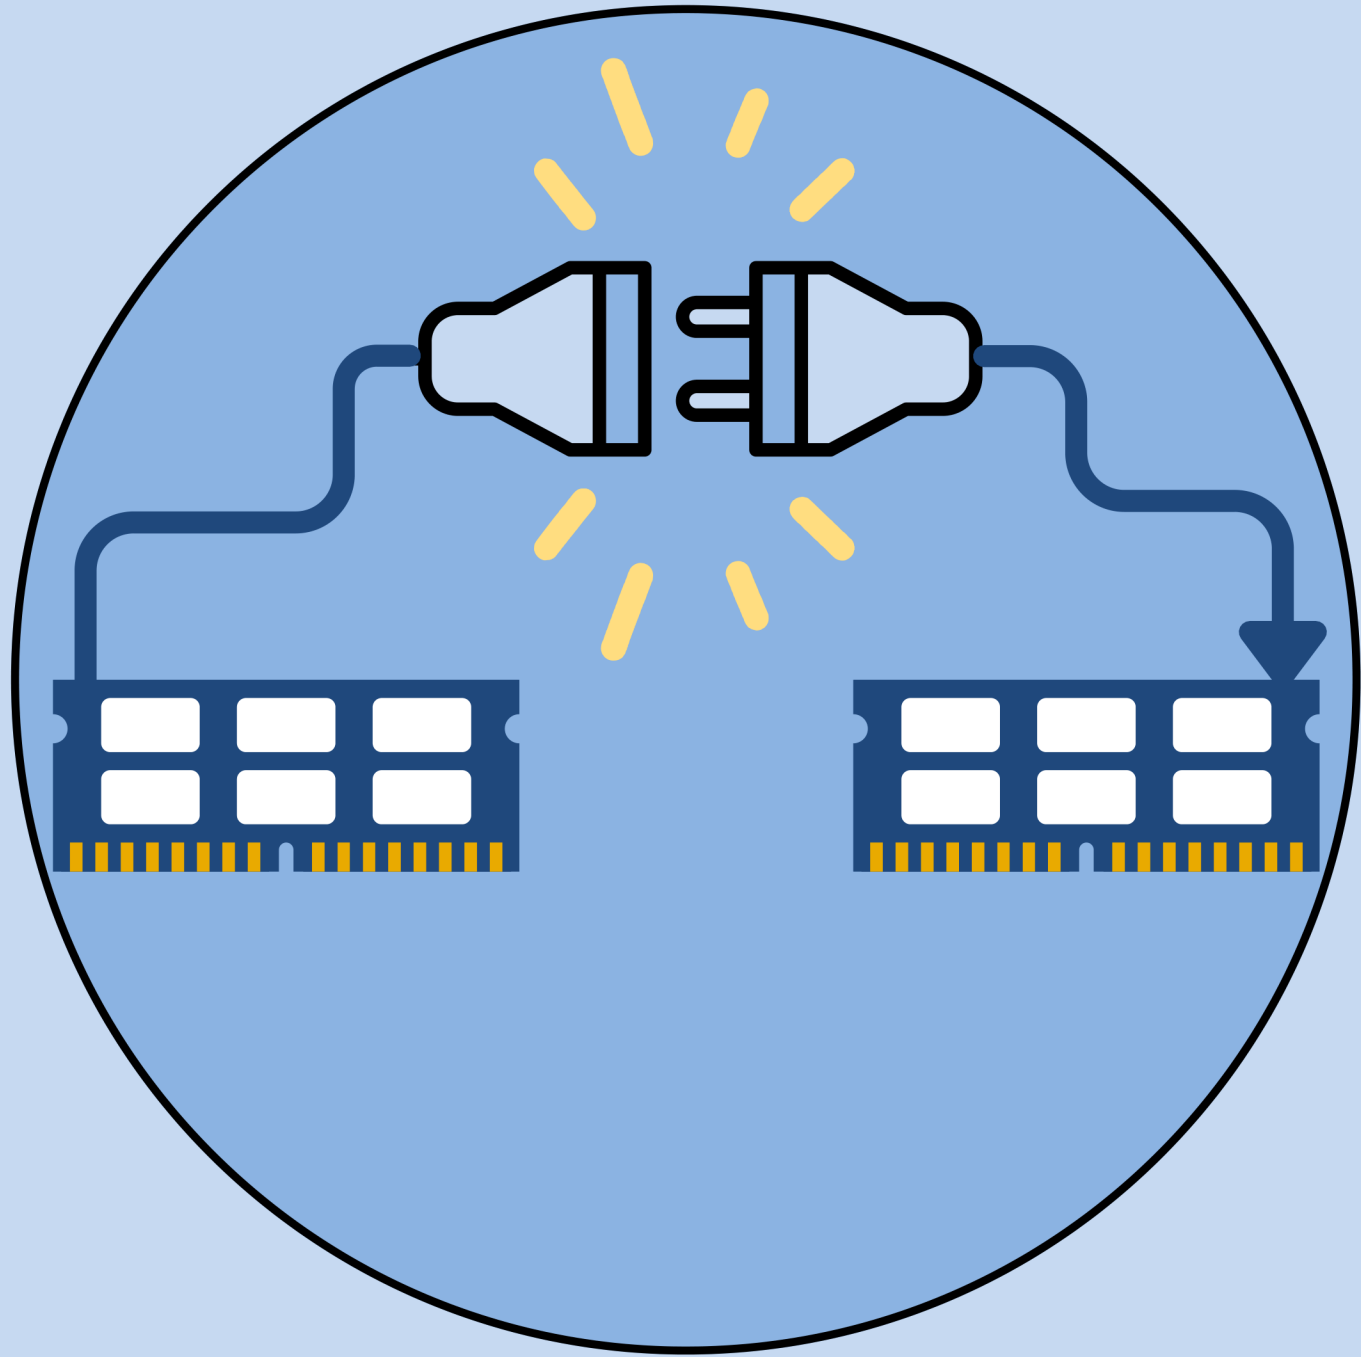



File I/O is slow






RAM-to-RAM transfer is fast







# Streaming workflow

## Data Receivers

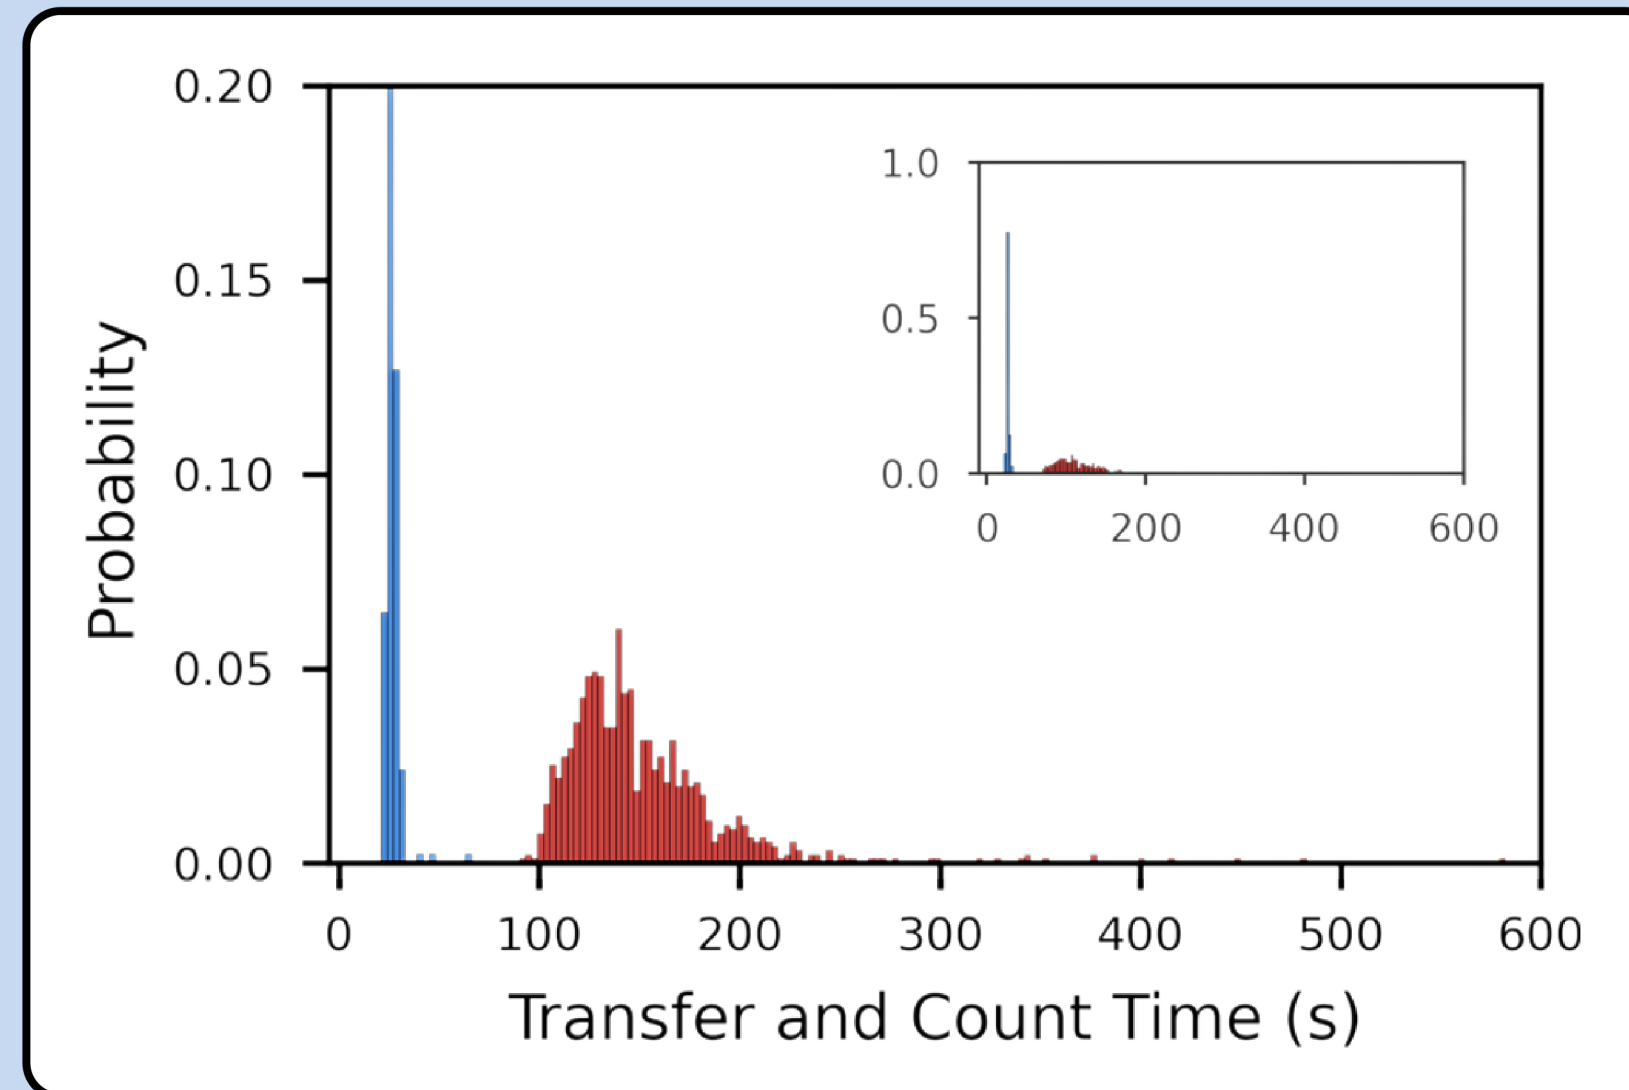


ØMQ sockets  
connect our network








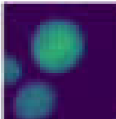




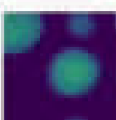




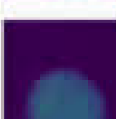




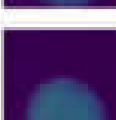




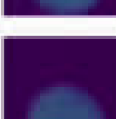
































# Streaming workflow: results



Storage Usage  5.9 TB Free

[EXPORT](#) [FILTER](#)

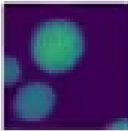




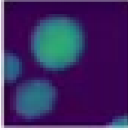




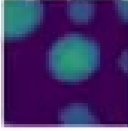




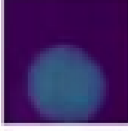









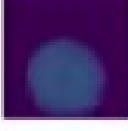




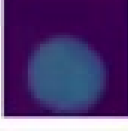









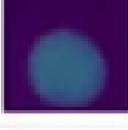









<input type="checkbox"/>	ID	Scan ID	Notes	Location	Created	Progress		
<input type="checkbox"/>		4680	5439		128.55.132.192 	11/8/2023		
<input type="checkbox"/>		4677	5436		128.55.132.192  perlmutter	11/7/2023		
<input type="checkbox"/>		4676	5435		128.55.132.192  perlmutter	11/7/2023		
<input type="checkbox"/>		4675	5434		128.55.132.192  perlmutter	11/7/2023		
<input type="checkbox"/>		4674	5433		128.55.132.192 	11/7/2023		
<input type="checkbox"/>		4673	5432		128.55.132.192  perlmutter	11/7/2023		
<input type="checkbox"/>		4672	5431		128.55.132.192  perlmutter	11/7/2023		
<input type="checkbox"/>		4671	5430		128.55.132.192  perlmutter	11/7/2023		
<input type="checkbox"/>		4670	5429		128.55.132.192  perlmutter	11/7/2023		
<input type="checkbox"/>		4669	5428		128.55.132.192  perlmutter	11/7/2023		
<input type="checkbox"/>		4668	5427		128.55.132.192  perlmutter	11/7/2023		

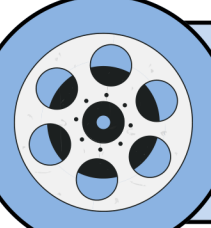




Storage Usage  5.7 TB Free

[EXPORT](#) [FILTER](#)

<input type="checkbox"/>	ID	Scan ID	Notes		Location	Created	Progress		
<input type="checkbox"/>		4680	5439	not streaming		128.55.132.192 	11/8/2023		
<input type="checkbox"/>		4677	5436			128.55.132.192  perlmutter	11/7/2023		
<input type="checkbox"/>		4676	5435			128.55.132.192  perlmutter	11/7/2023		
<input type="checkbox"/>		4675	5434			128.55.132.192  perlmutter	11/7/2023		
<input type="checkbox"/>		4674	5433			128.55.132.192 	11/7/2023		
<input type="checkbox"/>		4673	5432			128.55.132.192  perlmutter	11/7/2023		
<input type="checkbox"/>		4672	5431			128.55.132.192  perlmutter	11/7/2023		
<input type="checkbox"/>		4671	5430			128.55.132.192  perlmutter	11/7/2023		
<input type="checkbox"/>		4670	5429			128.55.132.192  perlmutter	11/7/2023		
<input type="checkbox"/>						128.55.132.192  perlmutter	11/7/2023		

 Starting session

4D Camera

Storage Usage 5.9 TB Free


[EXPORT](#) [FILTER](#)

ID	Scan ID	Notes	Location	Created	Progress
4660	5419		128.55.132.192	11/7/2023	✓
4659	5418		128.55.132.192	11/7/2023	✓
4658	5417		128.55.132.192	11/7/2023	✓
4657	5416		128.55.132.192	11/7/2023	✓
4656	5415		128.55.132.192	11/7/2023	✓
4655	5414		128.55.132.192	11/7/2023	✓
4654	5413		128.55.132.192	11/7/2023	✓
4653	5412		128.55.132.192	11/7/2023	✓
4652	5411		128.55.132.192	11/7/2023	✓
				11/7/2023	✓

```

1 // Acquire a set of 4D Camera scans.
2 // Set the sleep_time to account for offload time.
3
4 string command, ipaddressPlusPort, reply
5 TagGroup Out, OutTimes
6
7 // Olipscan user variables
8 number rotation = 0 // degree, 0 matches FEI software
9 number width = 256 // pixel, final 4D scan image is width + 1
10 number height = 256 // pixel
11
12 // 4D Camera user variables
13 number nread = 1 // frames per scan position
14 number nskip = 0 // number to skip between probe positions
15 number nflyback = 100 // typically this is set to 100 (8 frames for flyback time)
16
17 // Multi-scan variables
18 number num_scans = 15
19 number sleep_time = 30 // seconds (50 seconds is good for S12d12)
  
```

 Streaming



# What does starting a session mean?



User (client)

**Streaming Session**

Start a new streaming session

Machine  
perlmutter active ▾

---

threshold  
4


---

Session time (HH:MM)  
02:00

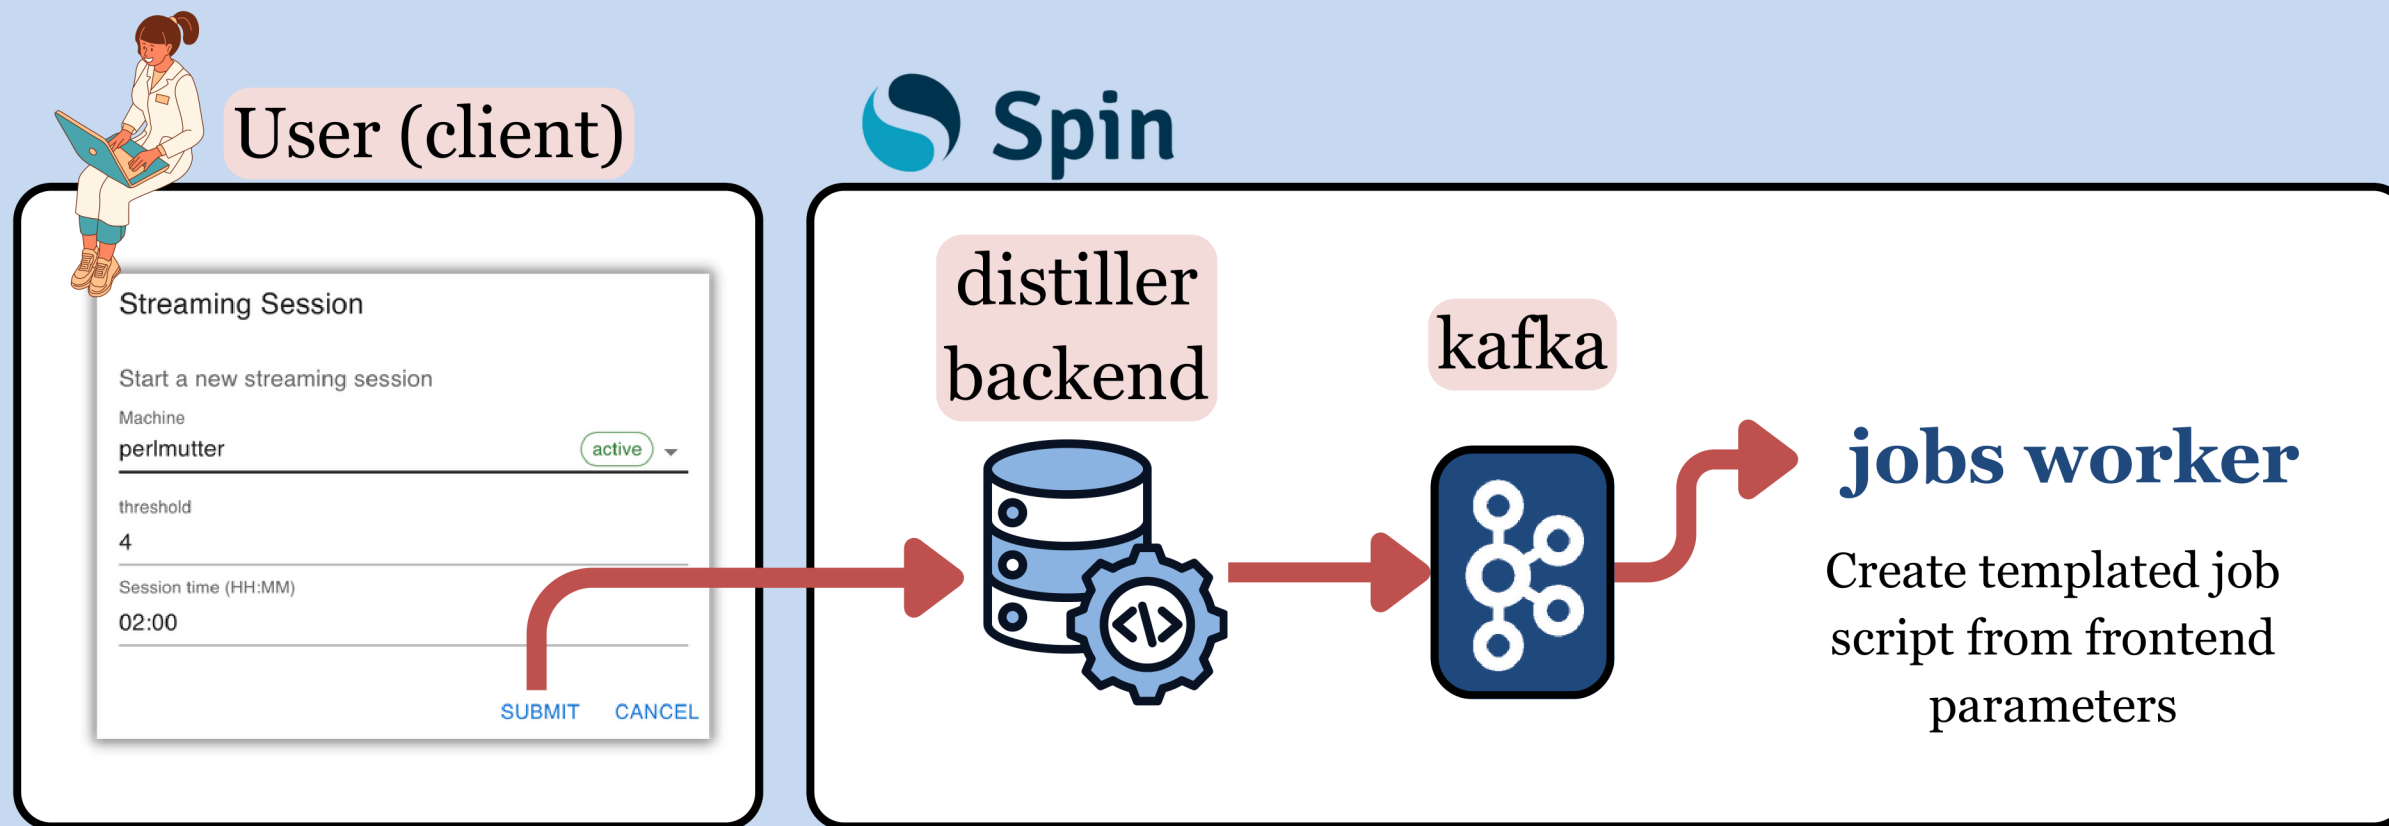
---


[SUBMIT](#) [CANCEL](#)



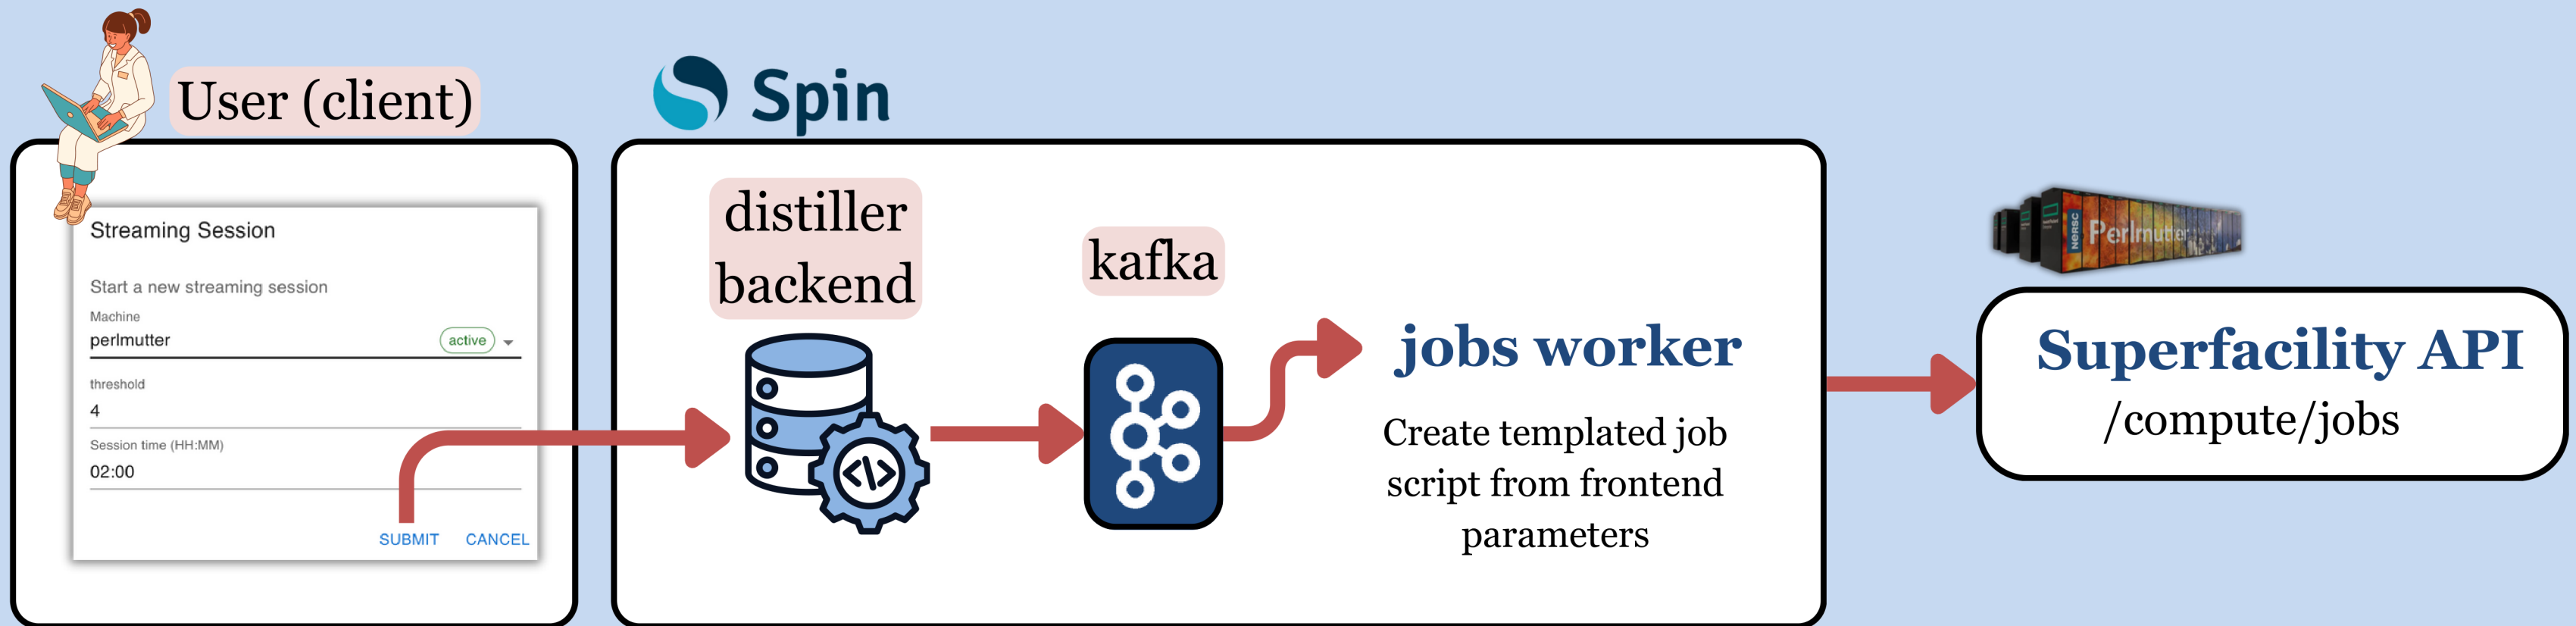



# What does starting a session mean?



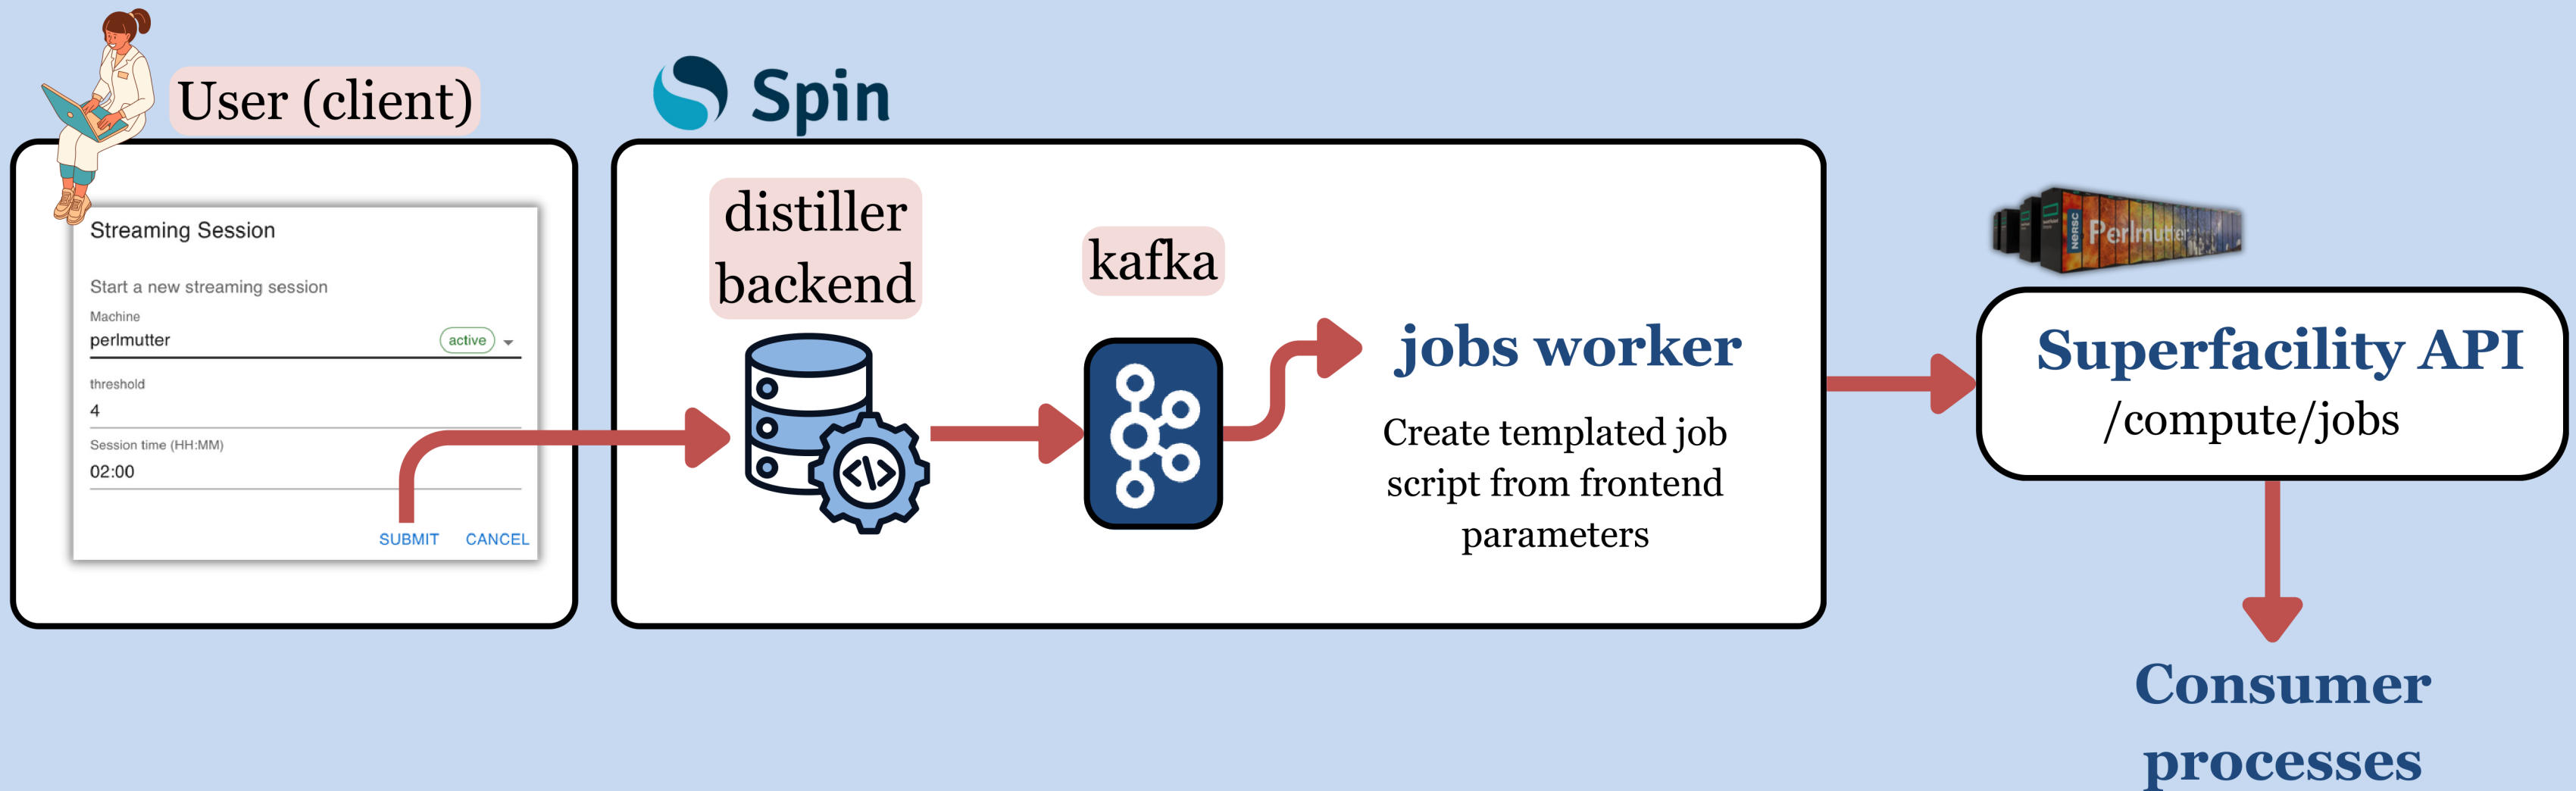


# What does starting a session mean?






# What does starting a session mean?








# How do we connect the network of processes?




Every process is a client in a distributed key-value store

## Clients

- Unique ID
- No. messages
- Current scan no.
- Status
- Ports/IP addresses



# How do we connect the network of processes?




Every process is a client in a distributed key-value store


- Clients**
- Unique ID
  - No. messages
  - Current scan no.
  - Status
  - Ports/IP addresses

- Server**
- Check sequence
  - Publish update
  - Log full state

\*different colors represent different clients  
(producers, consumers, aggregators)

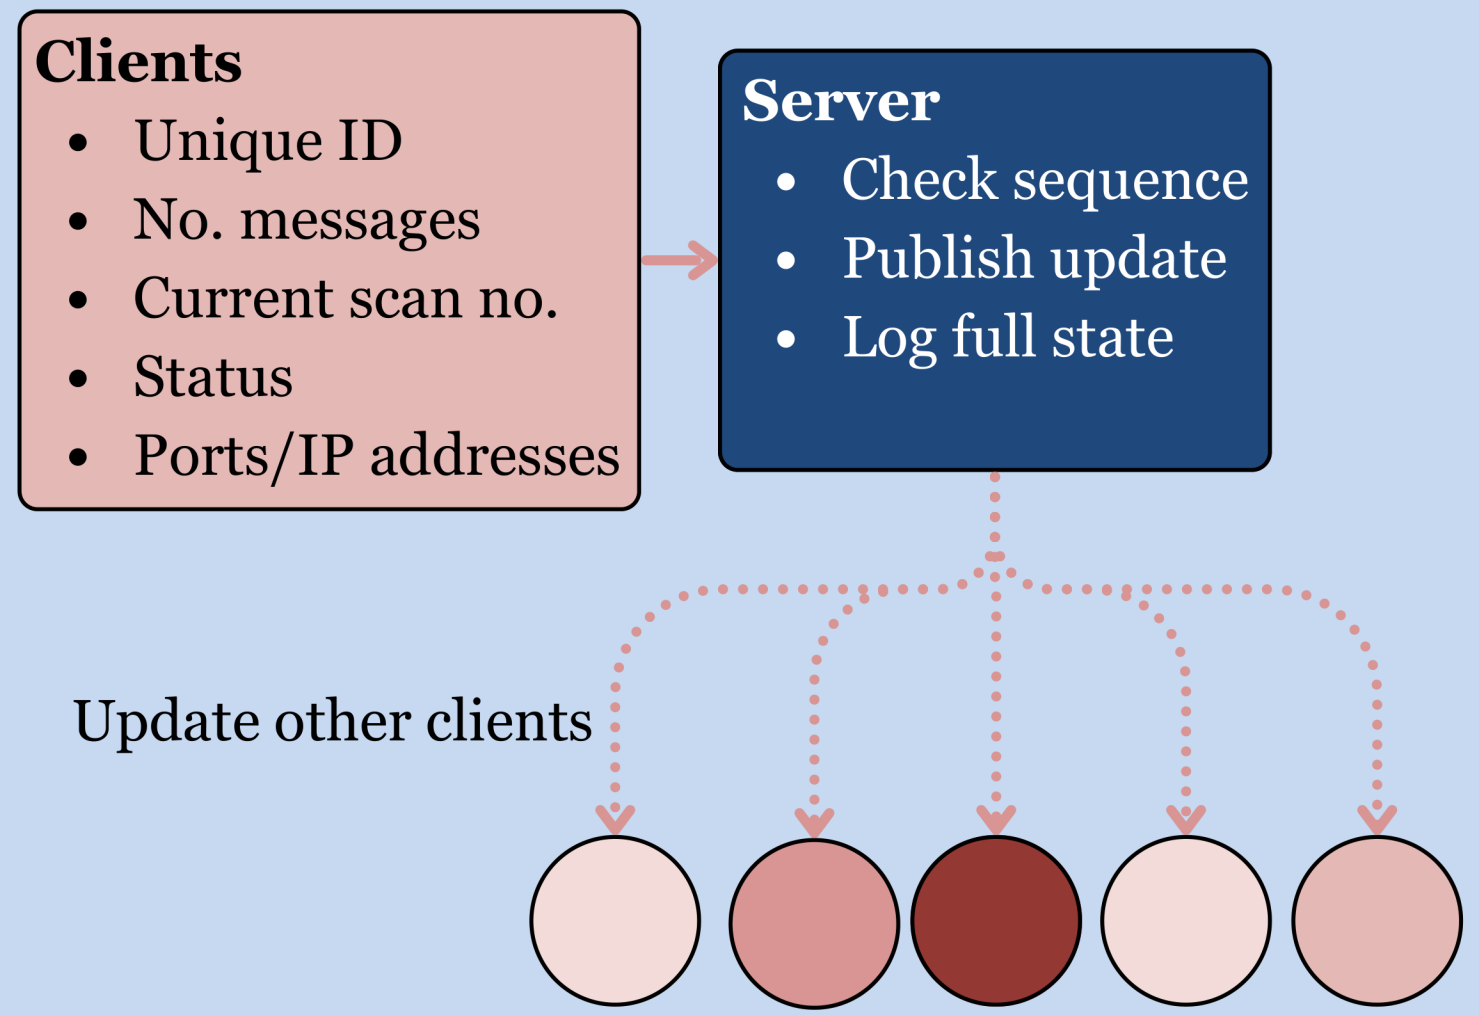


# How do we connect the network of processes?



\*different colors represent different clients  
(producers, consumers, aggregators)


Every process is a client in a distributed key-value store









# Implications




5x faster  
turnaround time.



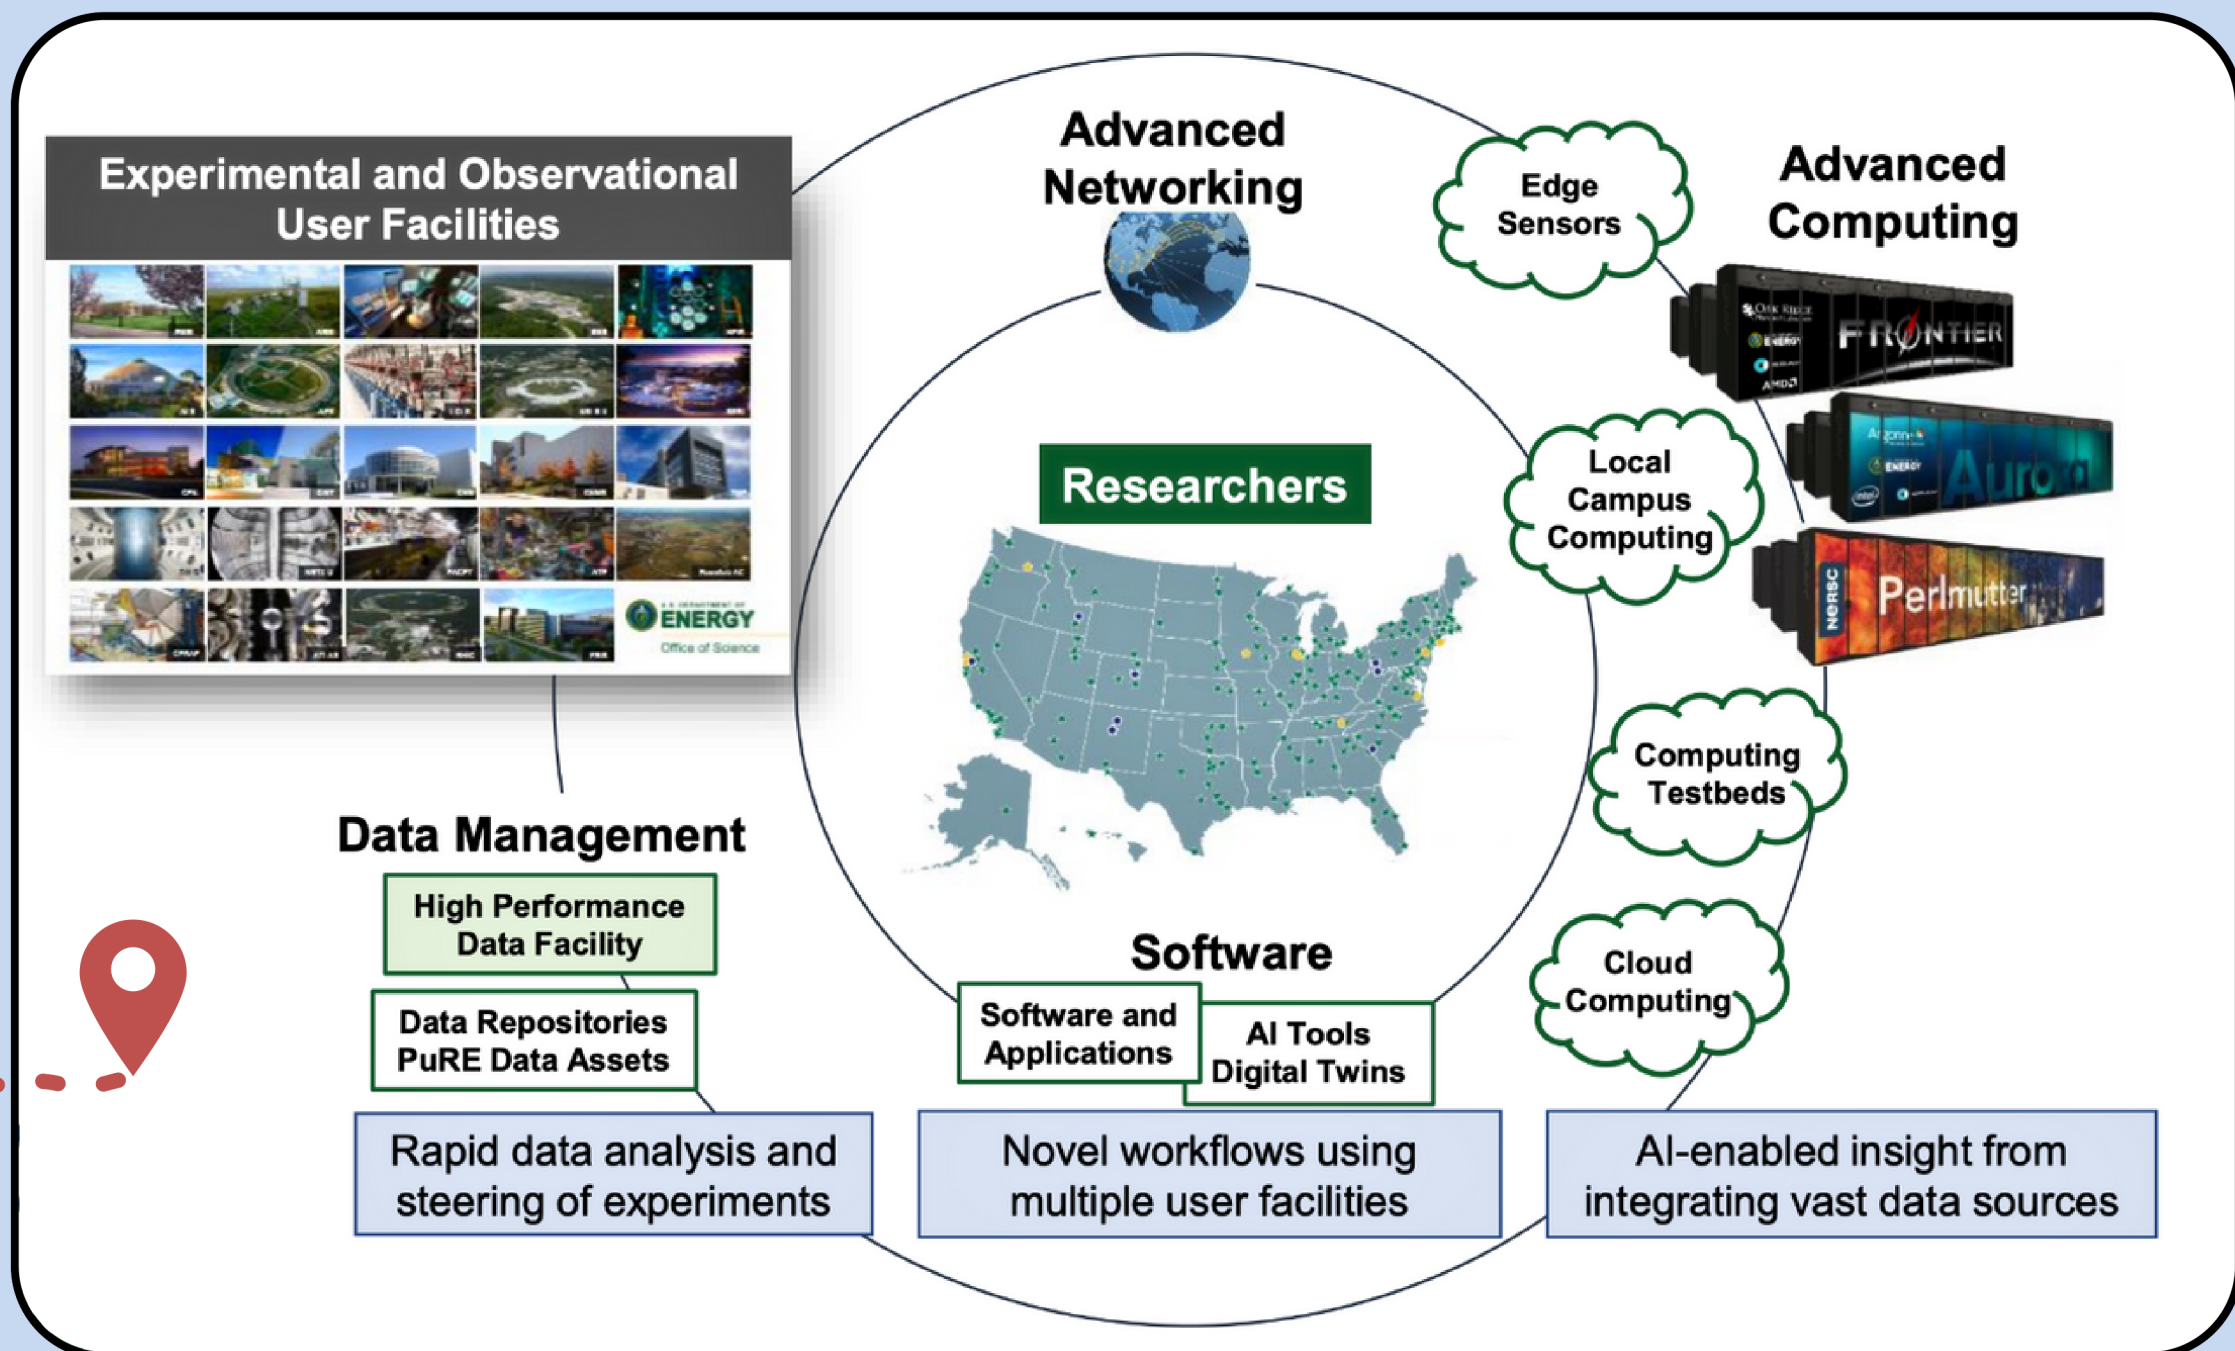
automated processing  
removes human in the  
loop.




happier  
experimentalists.



new pathway for  
integrated research  
infrastructure.





# Acknowledgements



Debbie Bard  
(DSEG)





Bjoern Enders  
(DSEG)



Peter Ercius  
(NCEM)



Chris Harris  
(DAS)



Tavia Stone Gibbons,  
Ashwin Selvarajan  
(SNG)



U.S. DEPARTMENT OF  
**ENERGY**

Office of  
Science

DE-AC02-05CH11231  
BES-ERCAP0024753  
BES-ERCAP0024754