12th Annual International Workshop on Innovating the Network for Data-Intensive Science (INDIS 2025)

Argonne NATIONAL LABORATORY

To Stream or Not to Stream: Towards A Quantitative Model for Remote HPC Processing Decisions

Flavio Castro, Weijian Zheng, Joaquin Chung, Ian Foster and Rajkumar Kettimuthu
Argonne National Laboratory

November 16th, 2025

Argonne NATIONAL LABORATORY

Computation-Experiment Coupling

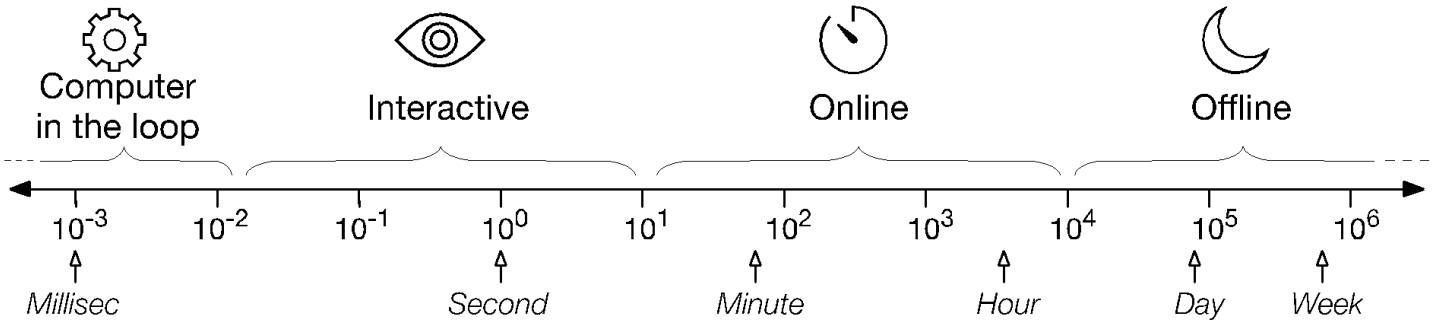

- Computer-in-the-loop: automated control during acquisition
- Interactive: rapid human-driven actions
- Online: real-time feedback during acquisition
- Offline: post-processing; hours to days; not experiment-coupled

Modern approach: Streaming + hybrid processing

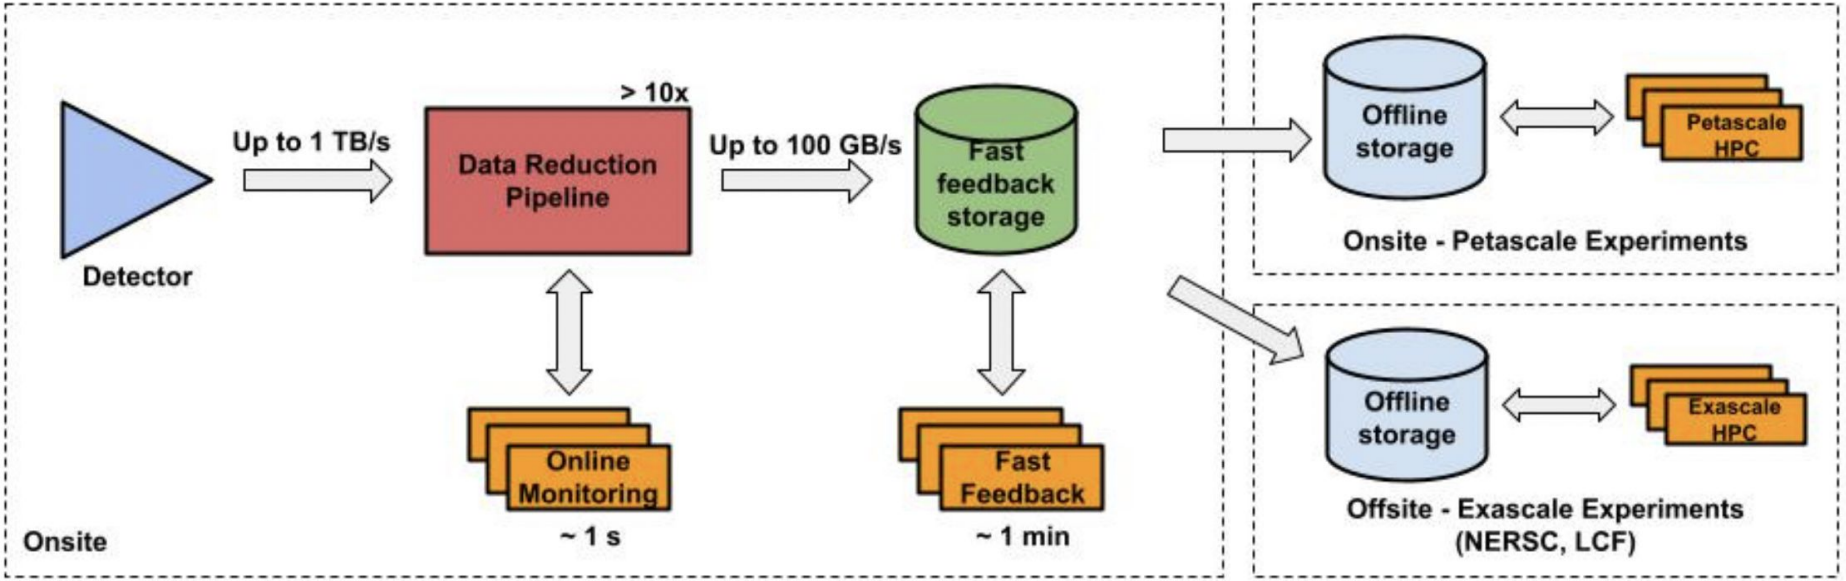

Adapted from Thayer et. al.
"Massive Scale Data Analytics at LCLS-II"

File-based Transfers are Inefficient for Online Analysis

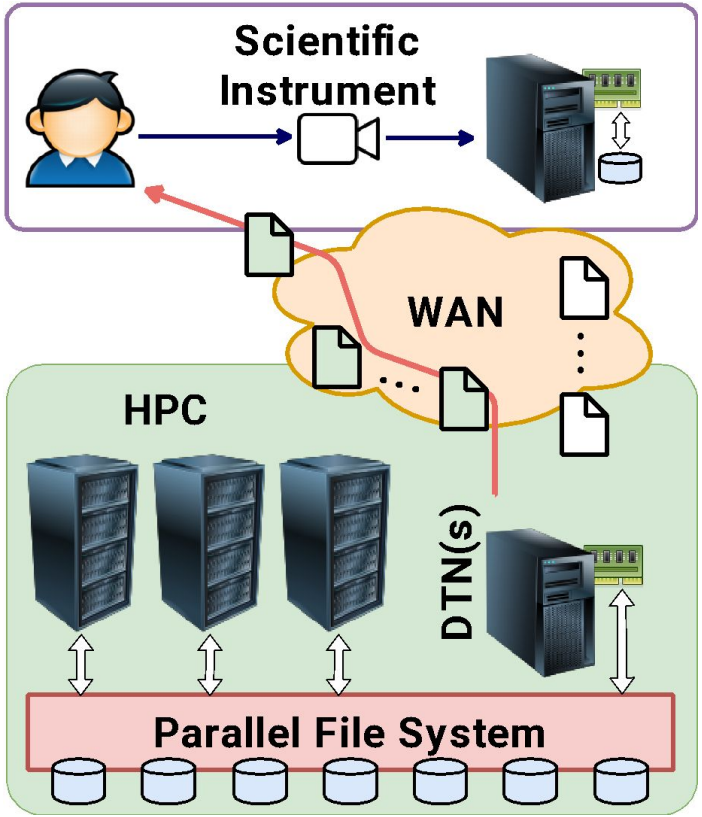

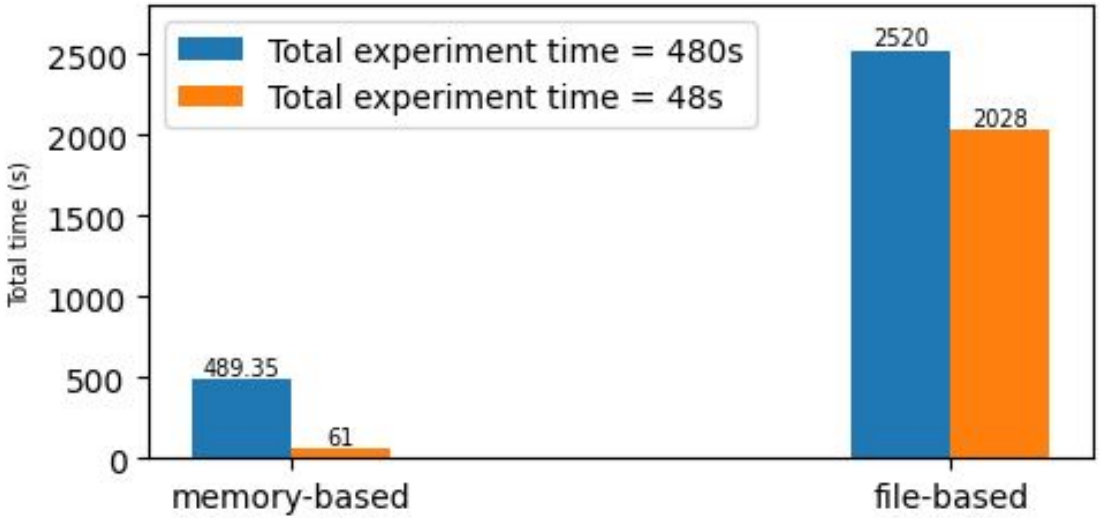

SciStream Globus Streaming

SciStream has evolved into Globus Streaming

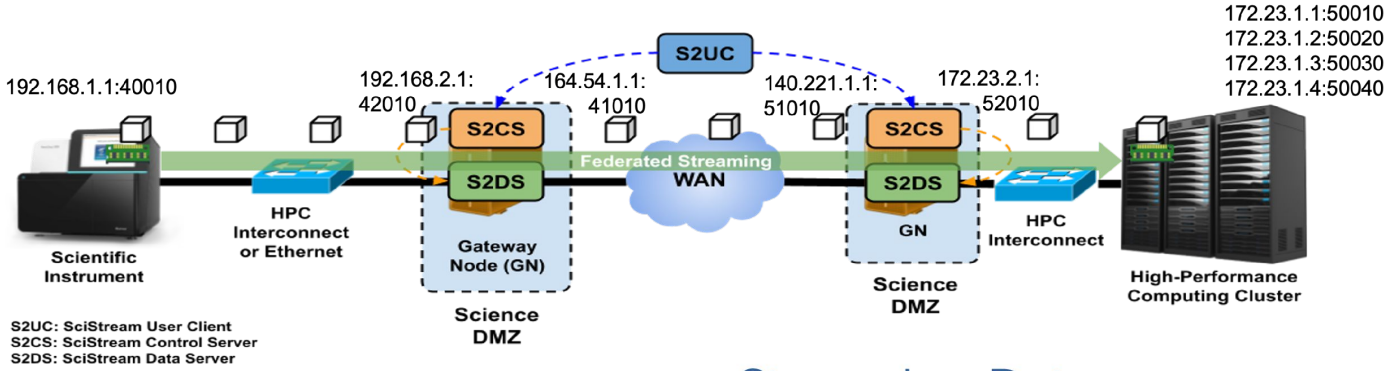

- Instrument and HPC are in Different Security Domains
  - On private networks and not directly visible from outside
- Automates the process of creating tunnels between security domain

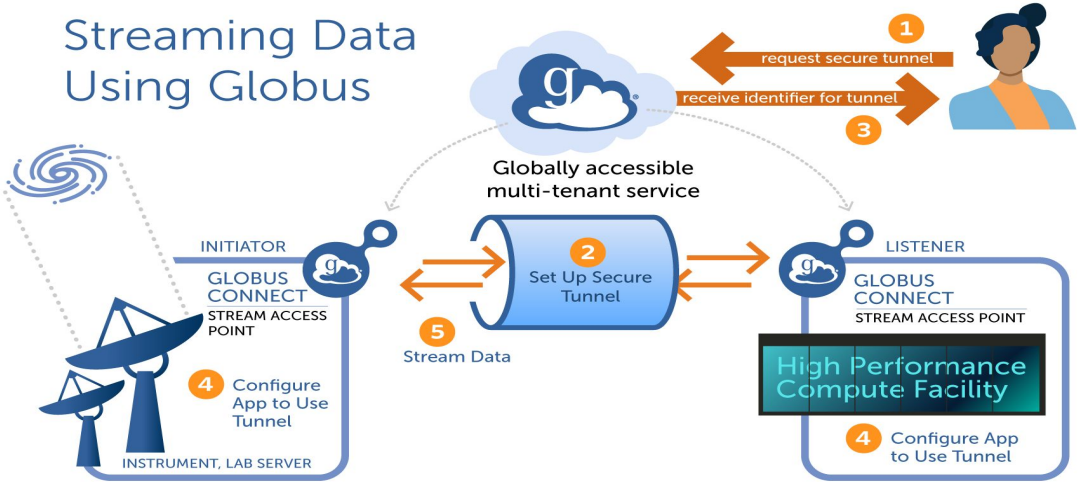

The Problem with Averages

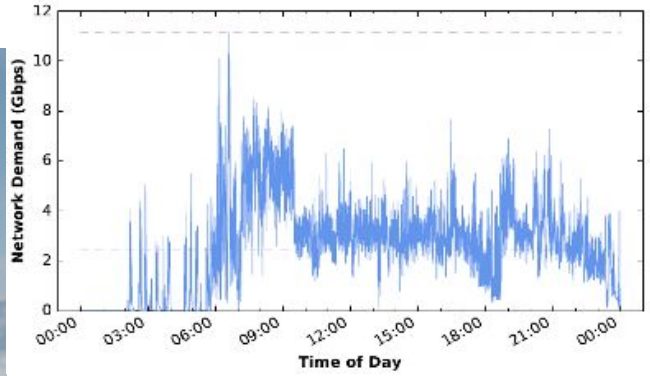

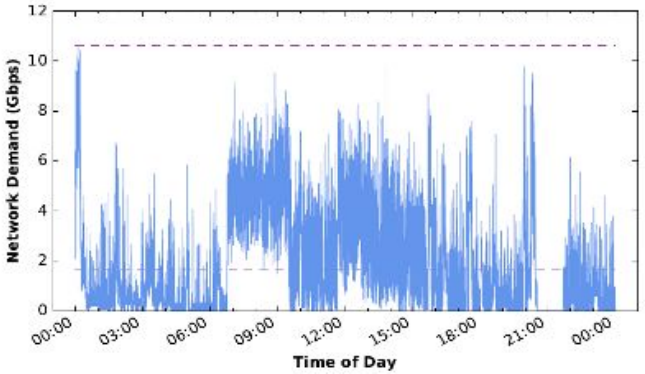

Average throughput hides large multi-hour variability.

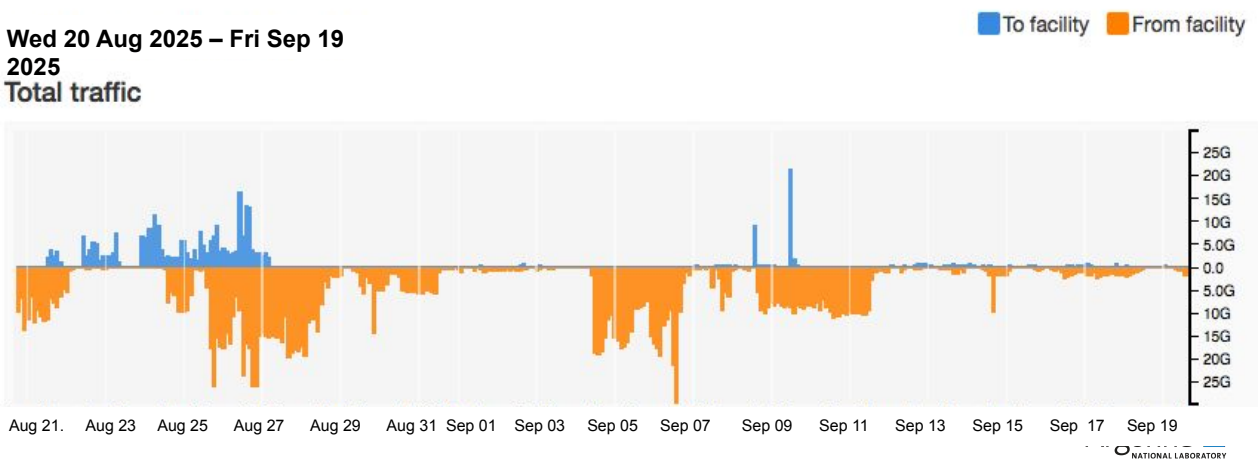

Streaming Speed Score and Completion Time Model

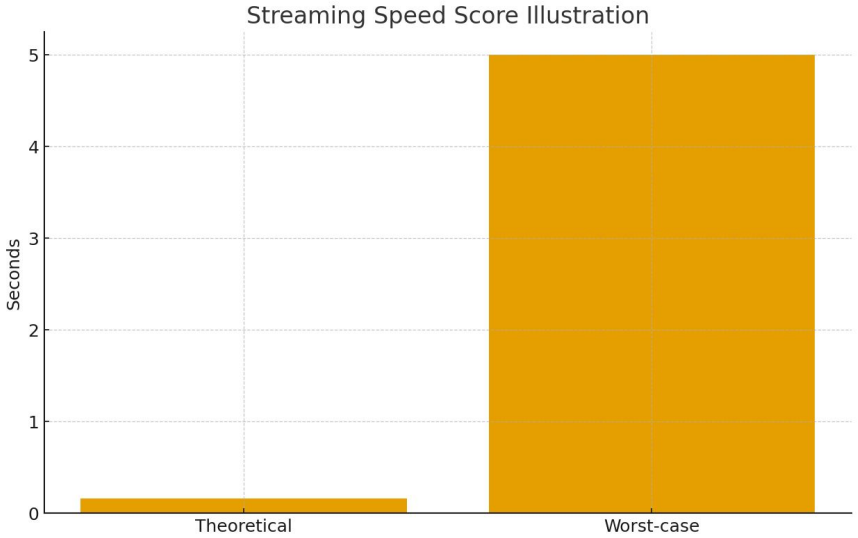

Remote Completion Time (T_{pct}):

$T_{pct} = (\theta \cdot S)/(\alpha \cdot Bw) + (C \cdot S)/(r \cdot R_{local})$

S = data size

C = computation required per GB

Bw = available network bandwidth

R_{local} = local compute throughput

α = transfer efficiency (↓ as SSS ↑)

θ = file I/O overhead factor

$r = R_{remote} / R_{local}$ (compute speed ratio)

$SSS = T_{worst\text{-}case} / T_{theoretical}$

High SSS → high tail latency → lower α

$T_{pct} < T_{local}$ → Stream to remote HPC; $T_{pct} > T_{local}$ → Process locally

Congestion Causes Non-Linear Growth in Transfer Time

- Worst-case transfer times grow non-linearly near high utilization
- Predictable performance requires scheduling
- Real-time workflows must account for worst-case behavior, not mean bandwidth

Case Study – Analytical Results

Scenario	Worst-case Transfer Time	Remote Compute (4 TF/s)	Local Compute (2 TF/s)	Outcome
2 GB/s, 34 TF	1.2 s	6.8 s	17 s	✓ Tier 2 feasible
3 GB/s, 20 TF	6 s	4 s	10 s	Borderline
4 GB/s, 20 TF	No feasible transfer	4 s	10 s	✗ Not feasible

- This table summarizes three representative streaming scenarios, showing how worst-case transfer time and compute capabilities determine whether remote HPC can meet real-time analysis deadlines.

Thanks!

Acknowledgments

Argonne Leadership Computing Facility