Shaping the Future of Distributed Computing: Bottom up, or Top down

Debbie Bard
Group Lead,
Data Science Engagement
NERSC

The Superfacility Model: an ecosystem of connected facilities, software and expertise to enable new modes of discovery

Superfacility@ LBNL: NERSC, ESnet and CS research divisions working together to support experimental science

- A model to integrate experimental, computational and networking facilities for reproducible science
- Enabling new discoveries by coupling experimental science with large scale data analysis and simulations

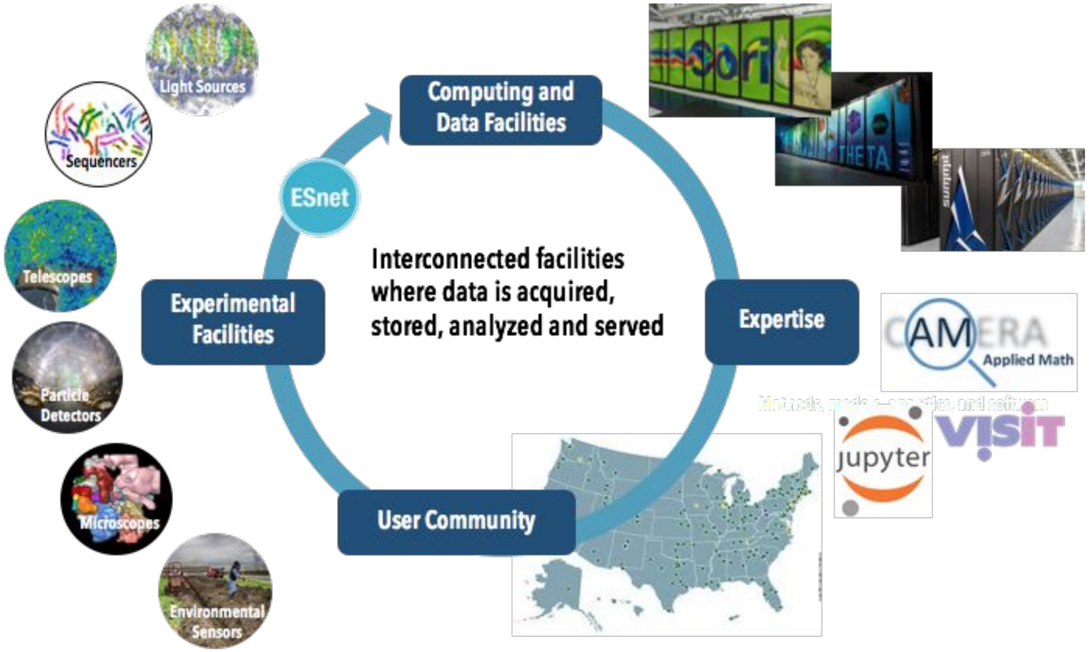

Which approach, bottom-up or top-down, is most likely to shape the future of distributed computing?

- Bottom-up is already happening organically - many science teams are juggling on-site computing with allocations at DOE HPC centers, coordinating data and compute tasks across sites
    - Superfacility work is designed to support/enable this model.
- The pressing need is to make this model easier to use
    - It's much easier to build the hardware than it is to solve the software, interoperability, sociology and policy issues.
        - Eg federated ID, API interfaces, unified allocation model, data movement and management, time-sensitive access to resources…
        - This has to be done top-down.

ExaFEL: a cross-facility workflow analysing LCLS data at SLAC and NERSC for real-time feedback to guide a running experiment

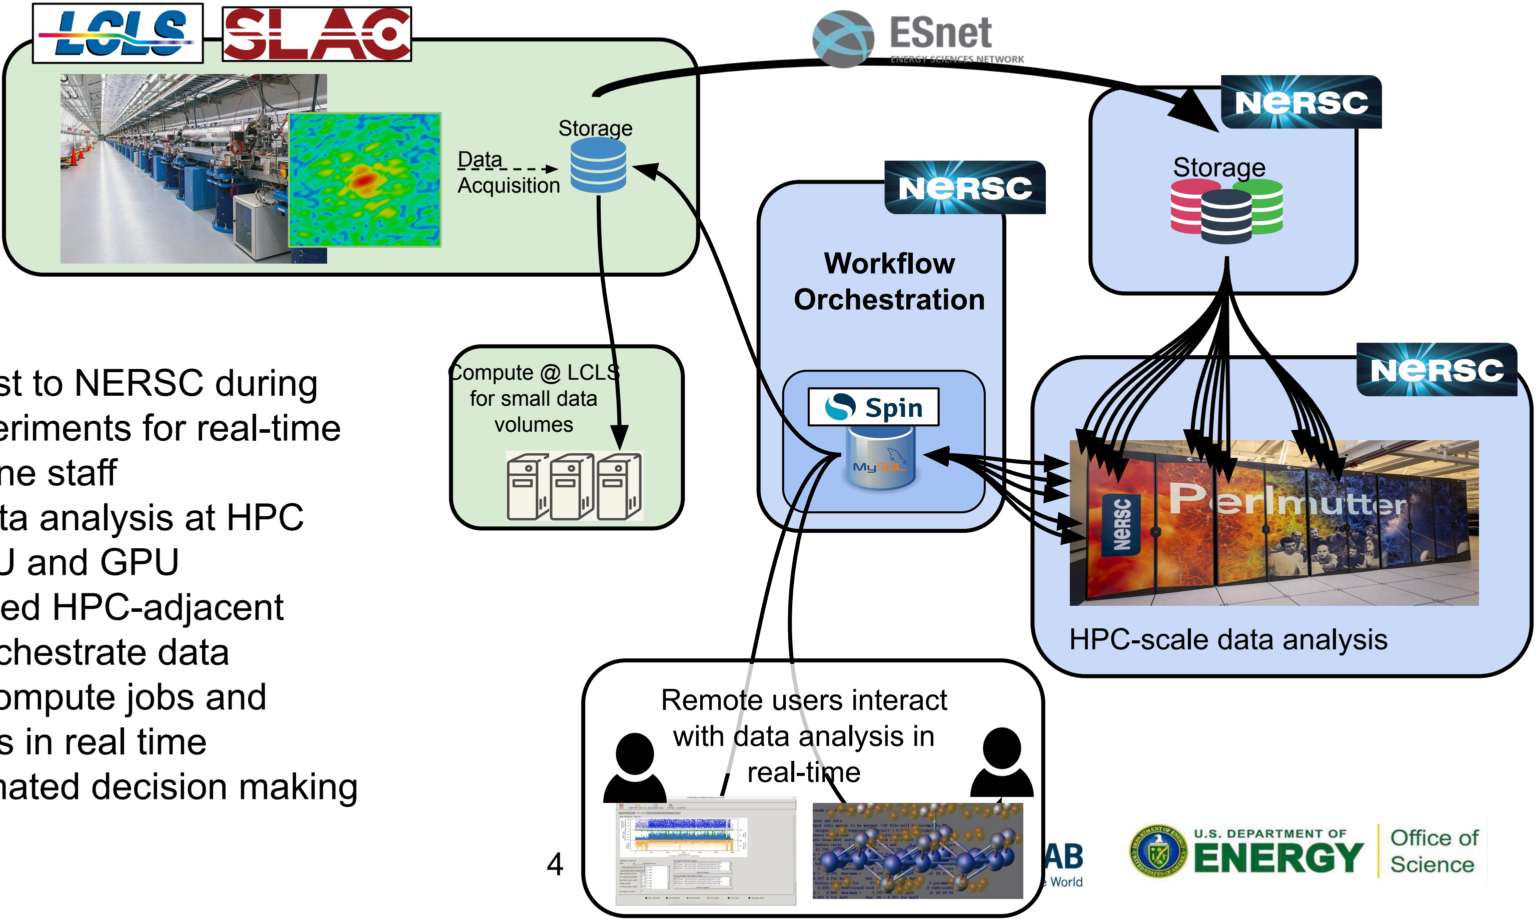

LCLS needs to burst to NERSC during some big-data experiments for real-time feedback to beamline staff

- **Real-time** data analysis at HPC scale, on CPU and GPU
- **Spin** (K8-based HPC-adjacent service) to orchestrate data movement, compute jobs and display results in real time
- **API** for automated decision making

To what extent will (should) economic considerations drive the future of distributed computing?

- A distributed system is required to meet the computing needs of DOE investments
  - An integrated research infrastructure is necessary for experiment facilities to be able to use the appropriate compute resources, when they are needed.
  - Multi-million dollar investments in science facilities are often made without sufficient attention/funding given to the computing required to run the experiment and process the data.
- People costs are harder to solve than equipment/hardware costs.
  - Small science teams don’t always have the personnel to redevelop their workflows to port to a new architecture.
  - Hard to hire for skills in this area.

To what extent does the notion of technology refresh rate factor into how distributed systems should be built?

- Natural tension between needing new, more performant technology to support user needs, and ability of user base to transition to it.
  - New tech means new challenges for the end user.
  - Multi-year projects struggle with systems and interfaces changing beneath them mid-project
  - Need to focus on sustainable, reliable interfaces (eg API) to smooth the hardware transition.
- Scientists strongly prefer predictable, reliable performance above the fastest possible performance.

How much will developments in interconnect technologies, possibly including wireless technologies, be important in shaping this future?

- The bottleneck for cross-facility workflows is more likely to be data ingest at facility edge.
- This poses a different set of problems than for a simulation workload that exchanges a lot of data *within* the system border.

2022 Internet Traffic: World Bank's Estimate 1.2 Pb/s, Cisco's Estimate 717 Tb/s

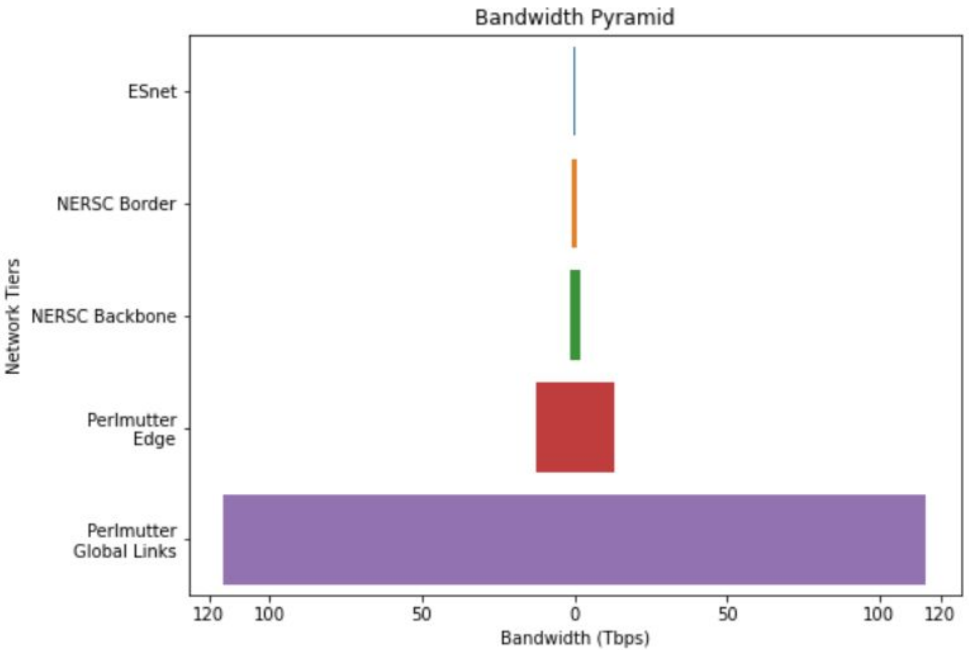

Figure courtesy of Taylor Groves and Ron Kumal @ NERSC

Emerging Use Case: Distributed Sensor Networks, e.g., Watershed SFA, self-guided field observatories

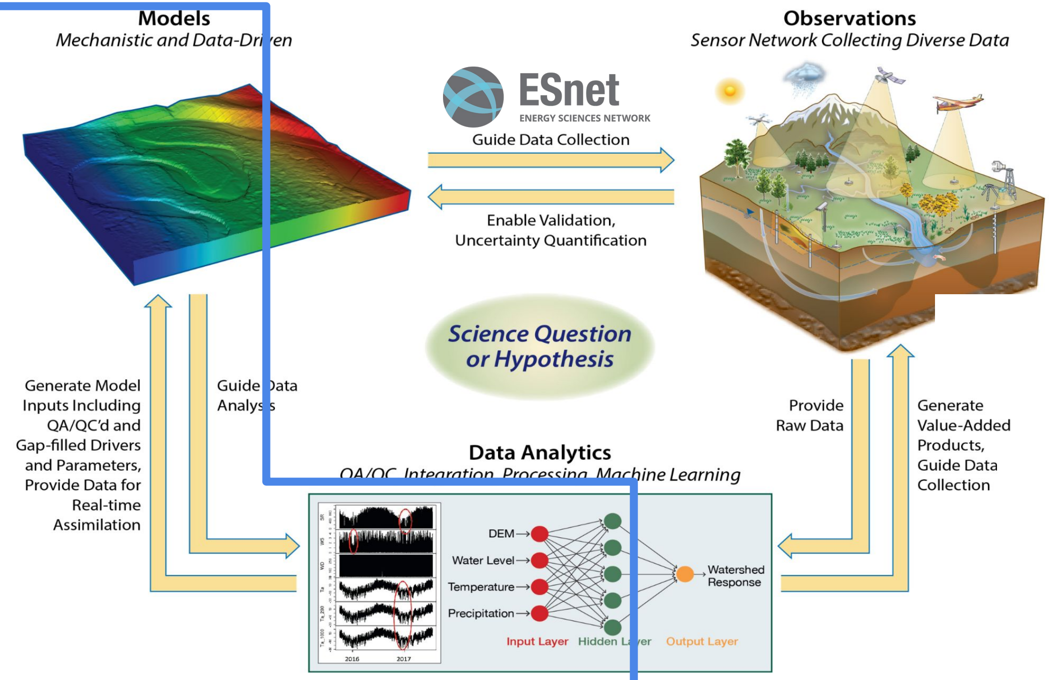

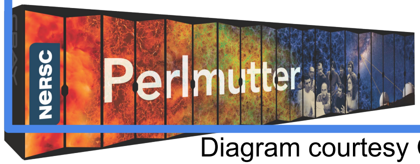

Diagram courtesy C. Varadharajan, Y. Wu et al

Wireless technologies are enabling new science use cases.

Real-time HPC could enable:

- Rapid updates to AI models based on streaming data and real-time simulations on HPC
- More accurate automated decision making
- Optimised sensor placement and calibration

How much will the required software infrastructure be critical in shaping this future?

- Software capabilities (like prioritisation of incoming urgent network traffic) are important for these use cases
  - NERSC has such a large workload, enabling network QOS to manage/balance the competing demands within the interconnect is important to enable consistent performance for users
- For a distributed computing system, it’s much easier to build the hardware than it is to solve the software, interoperability, sociology and policy issues
  - Eg federated ID, API interfaces, unified allocation model, easier data movement and management, time-sensitive access to resources…