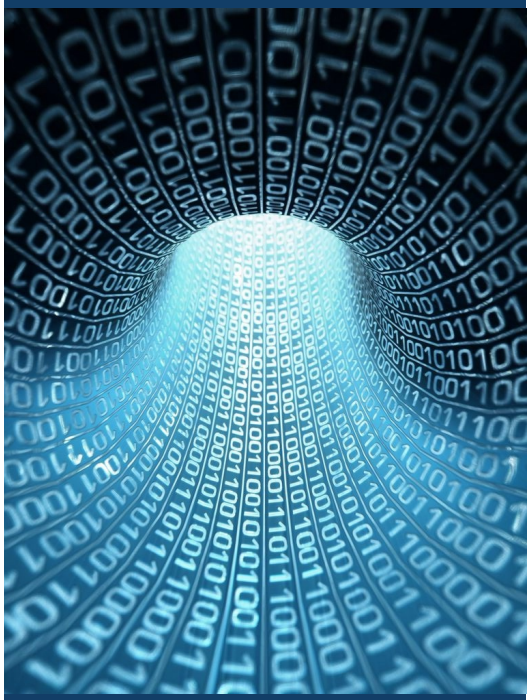

Sandia National Laboratories

Exceptional service in the national interest

Challenges turning Science Data into Knowledge

SOS17, Jekyll Island, GA

March 27, 2013

Janine Bennett, Robert Clay

Sandia National Laboratories

U.S. DEPARTMENT OF ENERGY

NNSA National Nuclear Security Administration

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND NO. 2011-XXXXP

Current simulation workflows pose a challenge to turning science data into knowledge

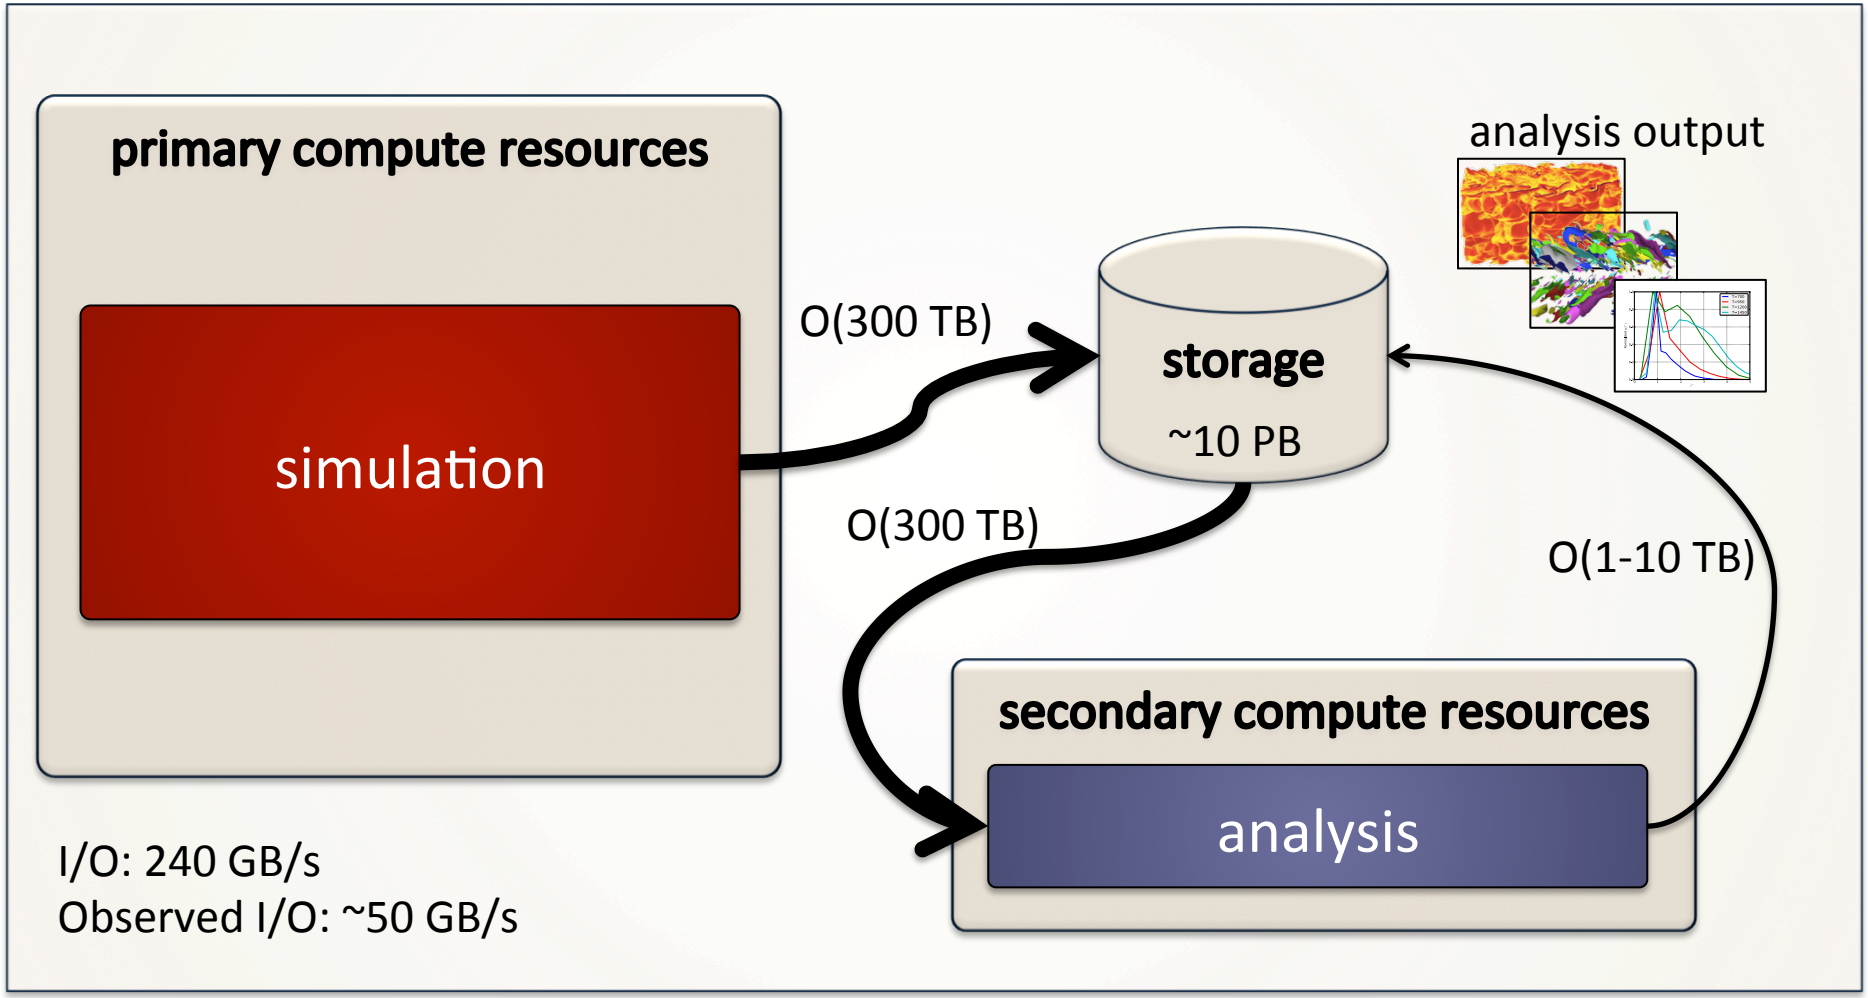

Current simulation workflows pose a challenge to turning science data into knowledge

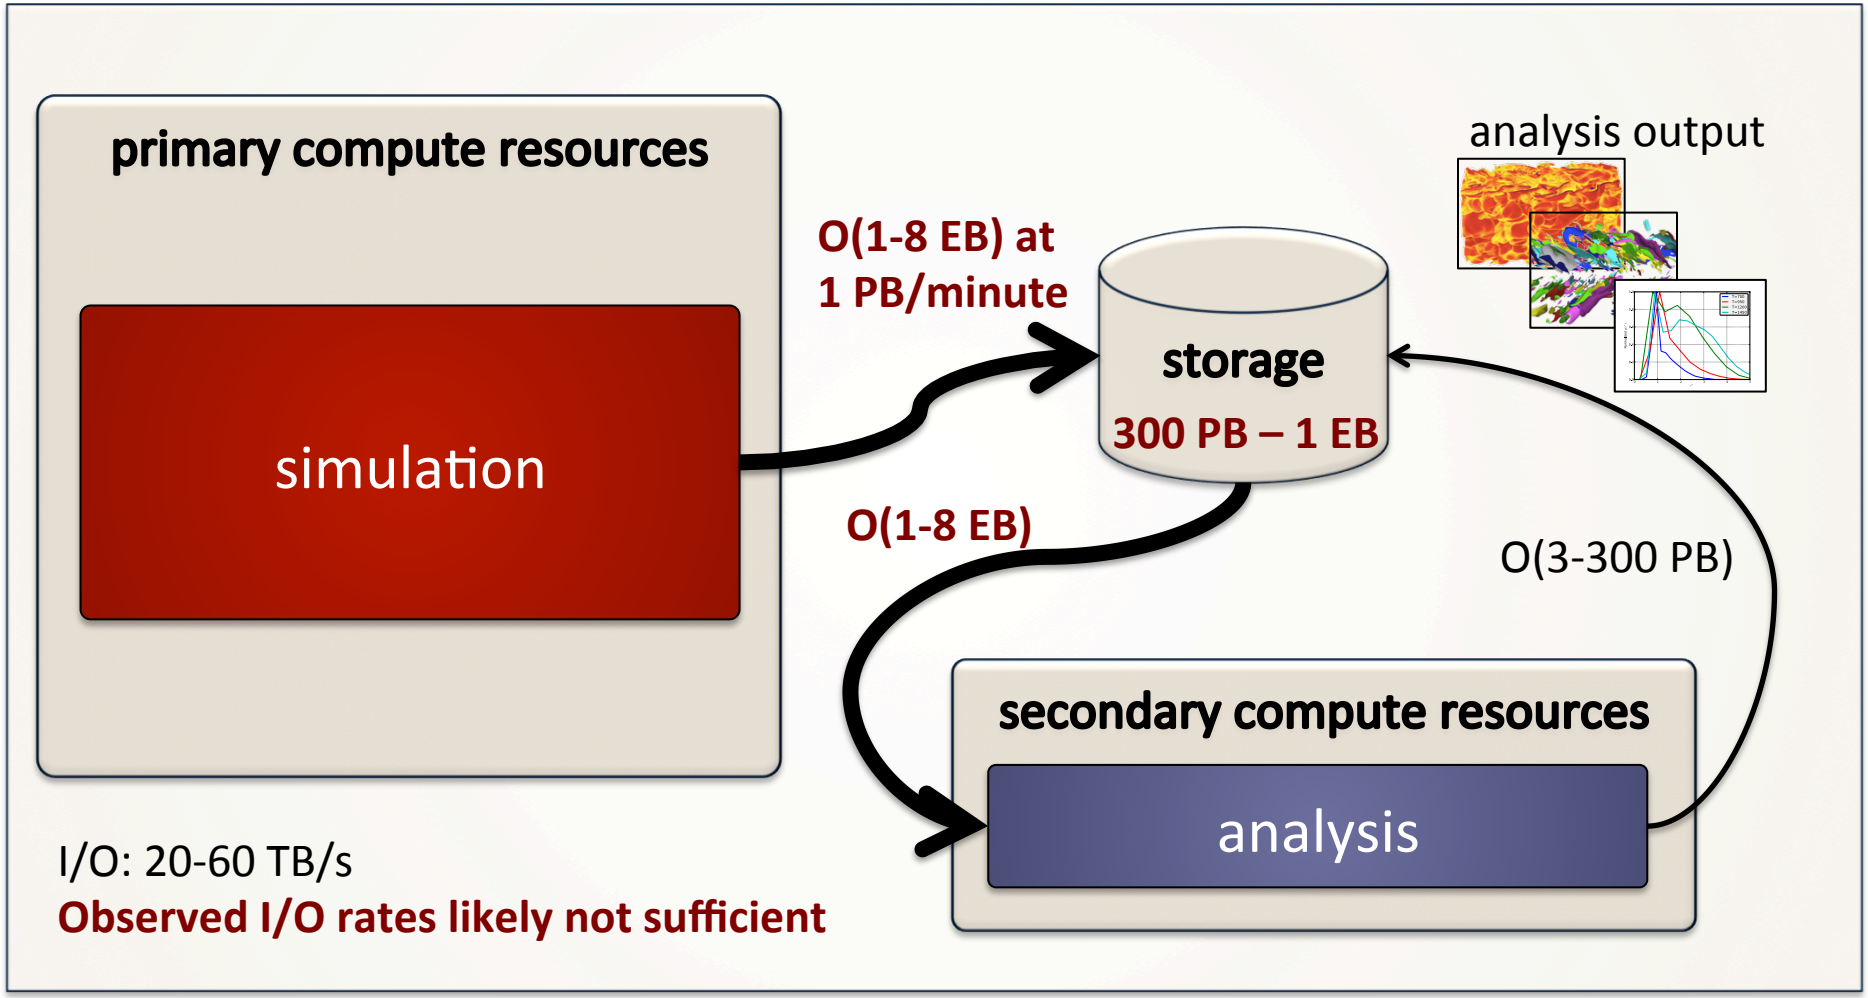

The existing workflow will not scale

New workflows should account for the entire scientific process: from simulation to insight

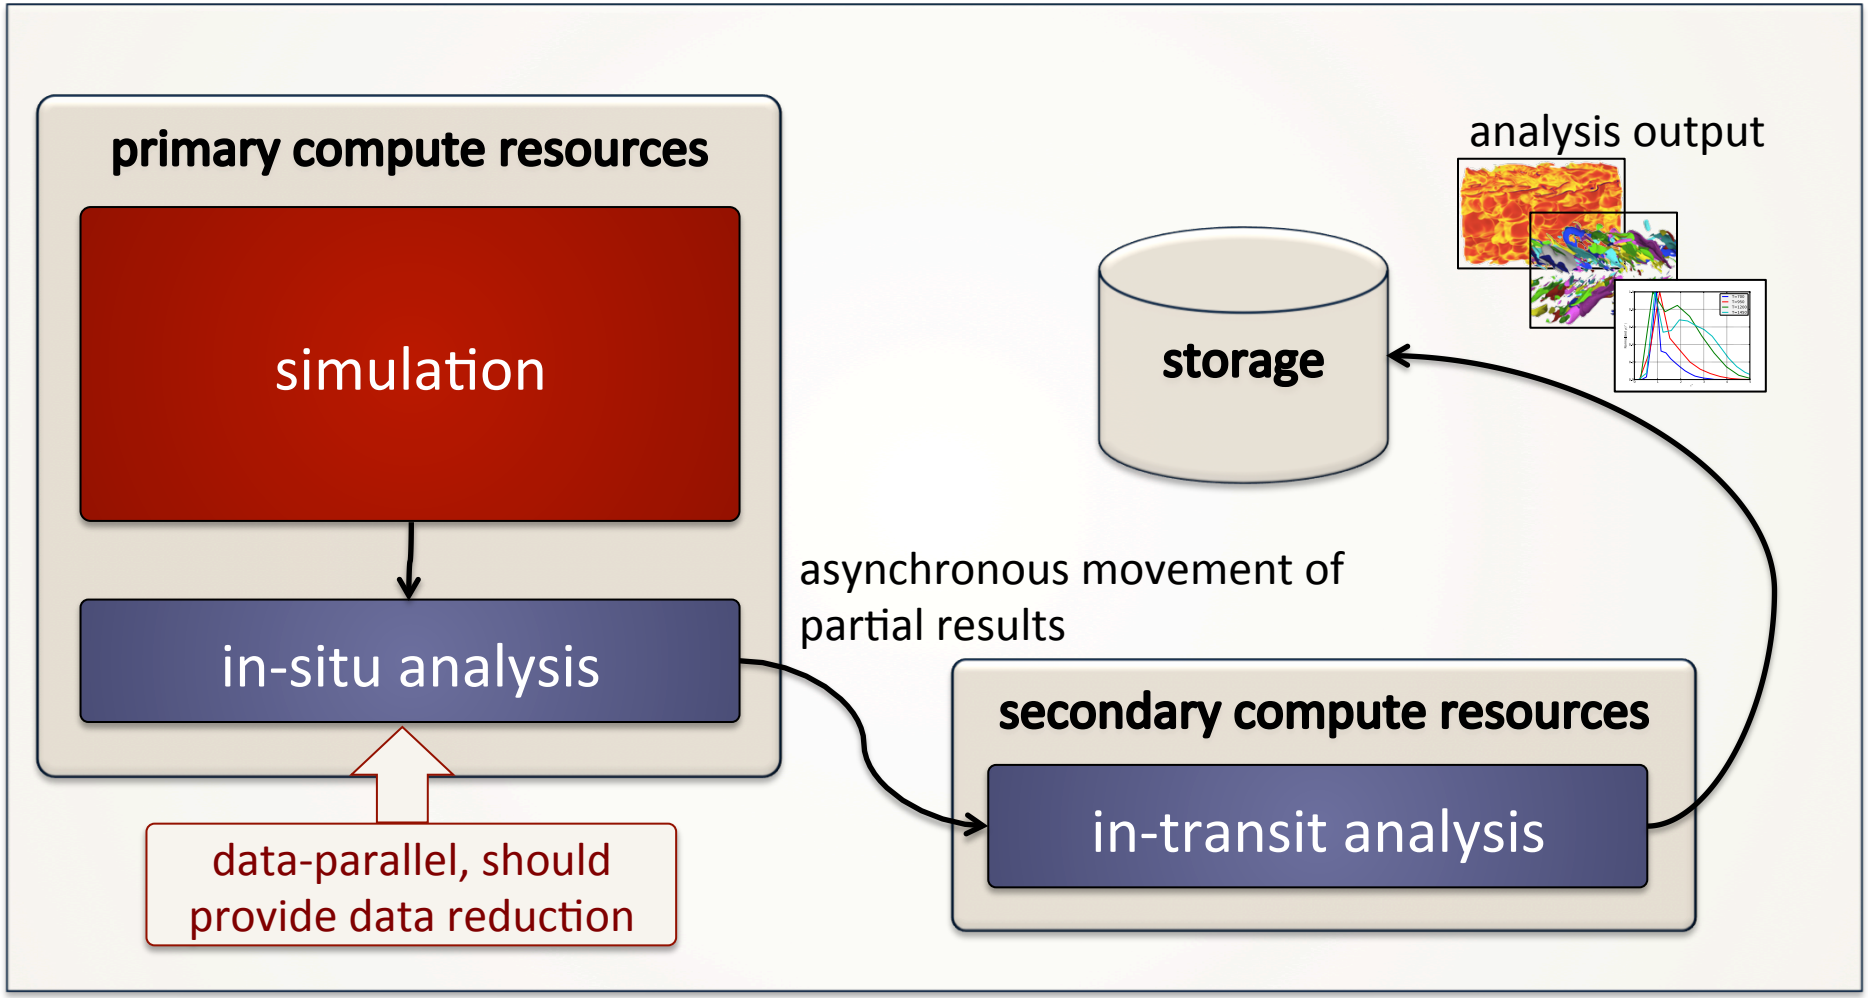

Concurrent in-situ/in-transit workflows show promise

Open challenges & questions for the community

- In-situ and in-transit workflows work well when analytic needs are known *a priori*
  - ***How do we best support exploratory analytic needs?***
- New workflows will no longer store everything to disk
  - ***How do you guarantee reproducibility and/or verification of results if you are not storing everything?***
- New workflows pose challenges to the data analysis community
  - ***How do we minimize impact to the simulation?***
  - ***How do we evolve existing analytics algorithms and/or infrastructures to support hybrid workflow paradigms?***

Acknowledgements

- Dr. Jacqueline Chen and the ExaCT co-design center.

- Funding provided by DOE Office of Advanced Scientific Computing Research.

- Sandia is a multi-program laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under Contract DE-AC04-94AL85000.

BACKUP SLIDES

HCCI Engine Combustion – Exascale Use Case 1

- Homogeneous Charge Compression Ignition (HCCI engine combustion) – tailor the charge stratification to control ignition and burn rate? Possible 25-50% increase in efficiency

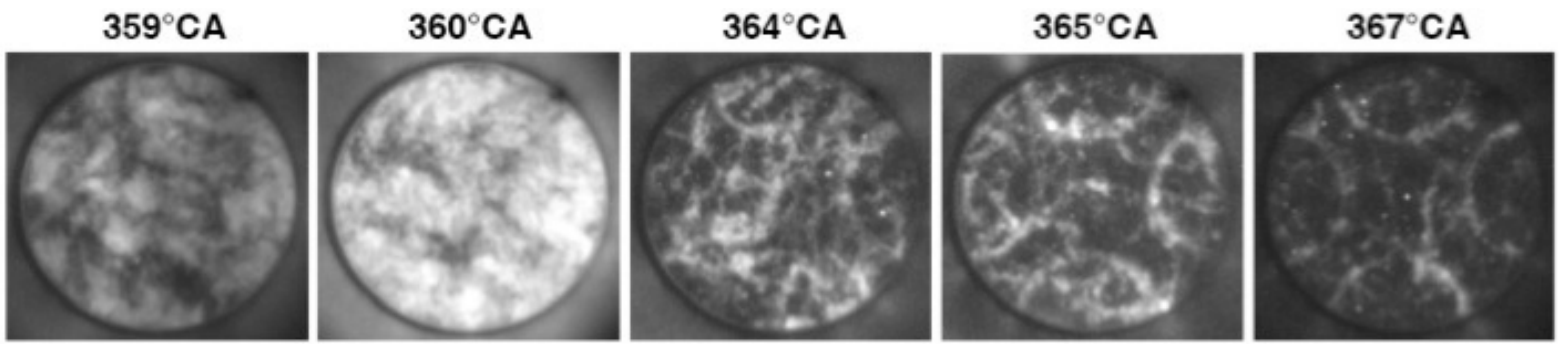

Chemiluminescence images of stratified sequential autoignition in an HCCI engine, courtesy J. Dec SNL

- Relevant turbulence, pressure, temperature (Re_t = 4300, 30-60 atm, 750-2000K)
- **Domain and grid size**: 5cm^3, 5 micron grid,
- **Size of state**: 10^{12} grids x 100 dof x 8 bytes/dof ~ **1 PB**
- **High water memory use**: ~ 3 x (size of state) ~ **3 PB**
- **Number of time steps**: 6 ms/5 ns timesteps = **1.2e6 steps**
- **Total run time**: 1380e-06 s/time step/grids x 1.2e6 time steps x 10^{12} grids / 3600s/hr = 0.46e^{12} cpu-hrs (**20 days at billion way concurrency**)
- **Total amount of data for analysis**: **1.0 exabyte**

Lifted Diesel Jet Flame – Exascale Use Case 2

- Relevant turbulence, pressure, temperature (Re_j = 500,000, 10 atm, 750-2000K)
- **Size of state**: $5.8*10^{12}$ grids x 27 dof x 8 bytes/dof ~ **1 PB**
- **High water memory use**: ~ 3 x (size of state) ~ **3 PB**
- **Number of time steps**: 0.7 ms/3.5 ns timesteps = **0.2e6 steps**
- **CPU-hours**: 370e-06 s/time step/grids x 0.2e6 time steps x $5.8*10^{12}$ grids / 3600s/hr = $0.12e^{12}$ cpu-hrs (~ **5 days at billion way concurrency**)
- Problem size i.e. N_{grids} (assuming 10^6 nodes and 10^3 cores/node) ~ 180^3 per node, 18^3 per core.
- **I/O Frequency: 1.3 PB every 33 timesteps or every minute** (dumping at Kolmogorov time scale of 0.115 μs)
- **Total data generated**: **8 exabytes**

Low Swirl Injector for Gas Turbine –Exascale Use Case 3

Sandia National Laboratories

- Simulate combustion of gaseous fuels at gas turbine conditions
  - Turbulence characteristics of a low swirl injector
  - 20 Atmospheres
  - Hydrogen and hydrogen-rich fuels
  - Inherently low Mach number flow

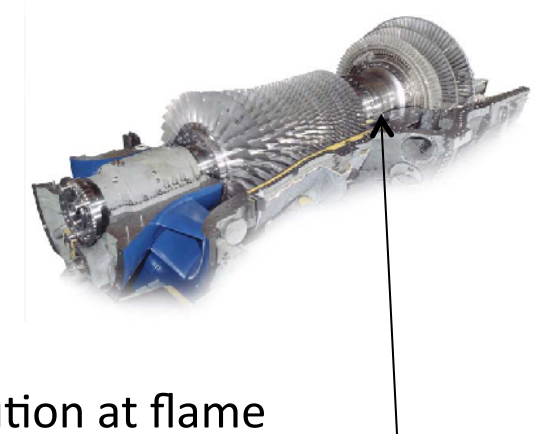

- Computational requirements
  - Domain size and grid: 25 cm domain, 15 microns resolution at flame
    - $(25e\text{-}2 / 1.5e\text{-}5)^3$ = 4.6e12 gridpoints
    - Significant benefit from AMR
  - **Size of State**: 4.6e12*20*8 ~ 7.4e14 α = (0.74 α) PB, **0.1 PB**
    - Where α -- AMR refinement efficiency -- ~10%
  - **High water memory usage** -- ~ **2.2 Pbytes**
  - **Number of timesteps**: (8.e-3 / 2.5e-5) = **3,200**
  - **Total run time** = **1 month**
  - **Total amount of data for analysis** – **75 Tbytes every 30 minutes**

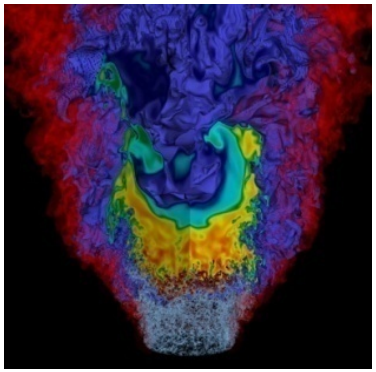

Current combustion simulation data is large and complex

- Data size
  - Billions of grid points per time step
  - Hundreds of time steps written to disk
- Data complexity
  - Multivariate
  - Turbulence is a complex phenomenon
  - Length scales: microns to centimeters
  - Temporal scales: nanoseconds to milliseconds
- Example: Lifted Ethylene Jet
  - 1.3 billion grid points
  - 22 chemical species, vector & particle data
  - 7.5 million cpu hours on 30,000 processors
  - 112,500 timesteps (data stored every 375th)
  - 240 TB of raw field data + 50 TB particle data

HPSS storage facility at NERSC

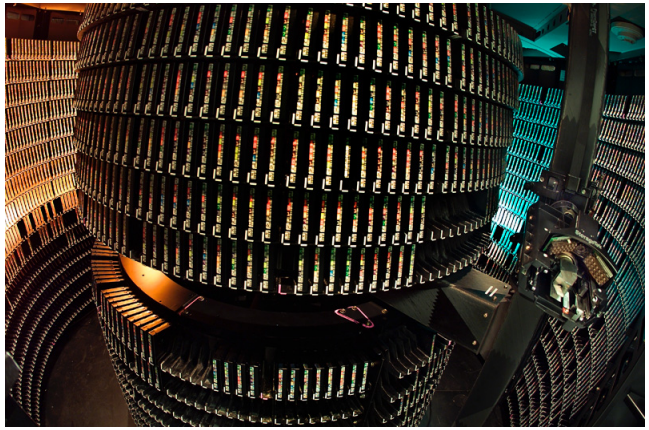

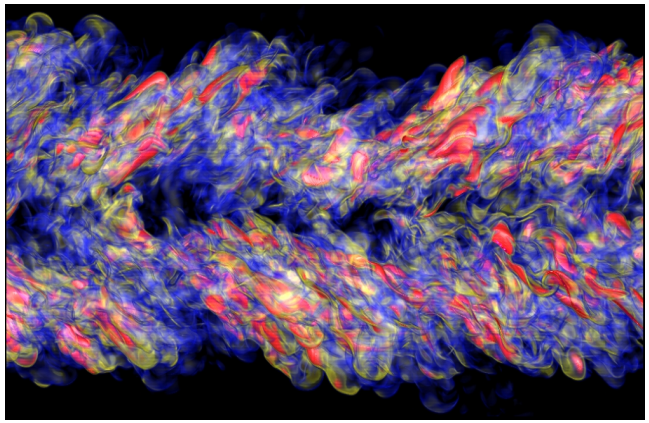

There is a rich design space of potential concurrent end-to-end workflows
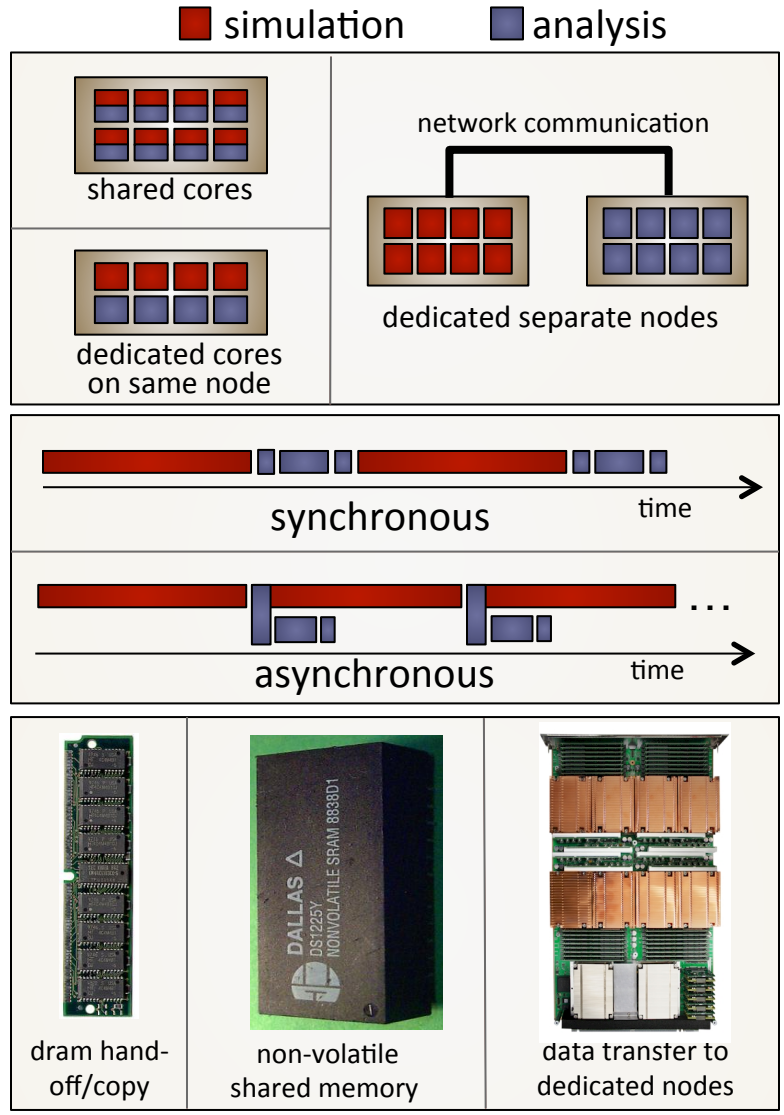
- Location of analysis compute resources
  - Same cores as the simulation (in-situ)
  - Dedicated cores on the same node (in-situ)
  - Dedicated nodes on the same machine (in-transit)
  - Dedicated nodes on external resource (in-transit)
- Synchronization and scheduling
  - Execute synchronously with simulation every n^{th} simulation time step
  - Execute asynchronously
- Data access, placement, and persistence
  - Shared memory access via hand-off / copy
  - Shared memory access via non-volatile near node storage (NVRAM)
  - Data transfer to dedicated nodes or external resources