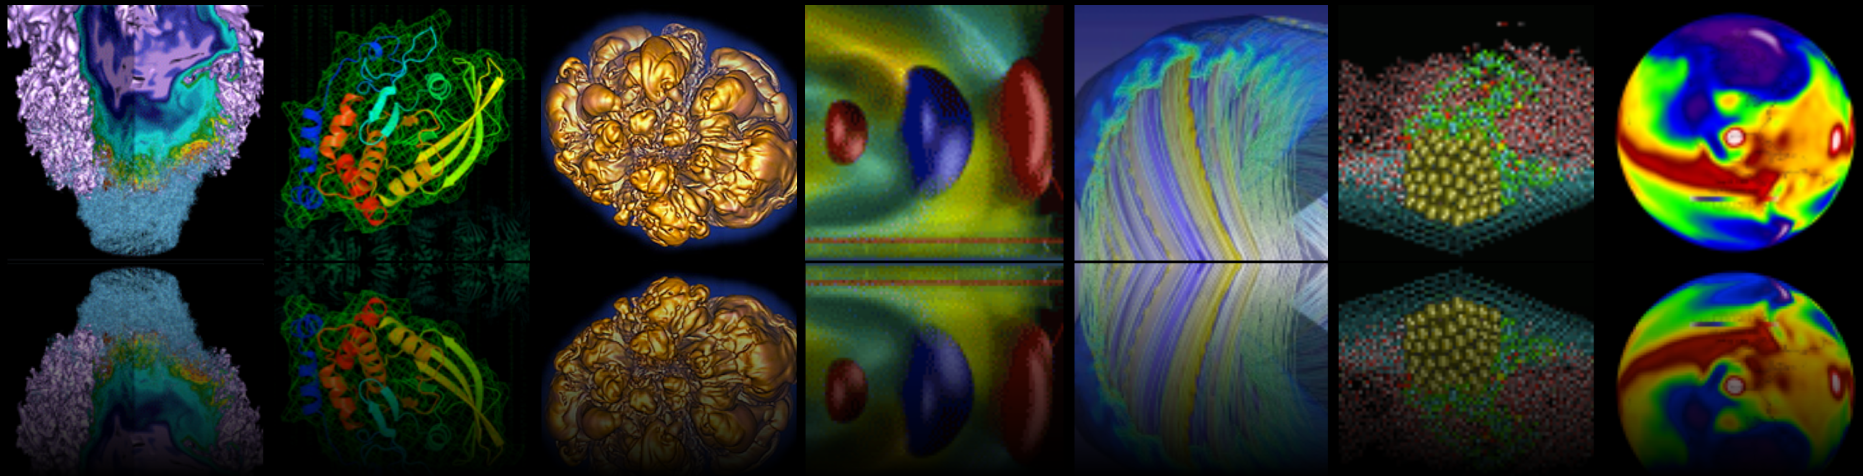

More Data, More Science and… Moore's Law

NERSC 40th

Kathy Yelick
Associate Laboratory Director for Computing Sciences

David Skinner
NERSC Strategic Partnerships Lead

Lawrence Berkeley National Laboratory

NERSC Strategy: *Science First*

- **Response to scientific needs**
  - Requirements setting activities
- **Support computational science:**
  - Provide effective machines that support fast algorithms
  - Deploy with flexible software
  - Help users with expert services
- **NERSC future priorities are driven by science:**
  - Increase application capability: “usable Exascale”
  - Simulation and data analysia

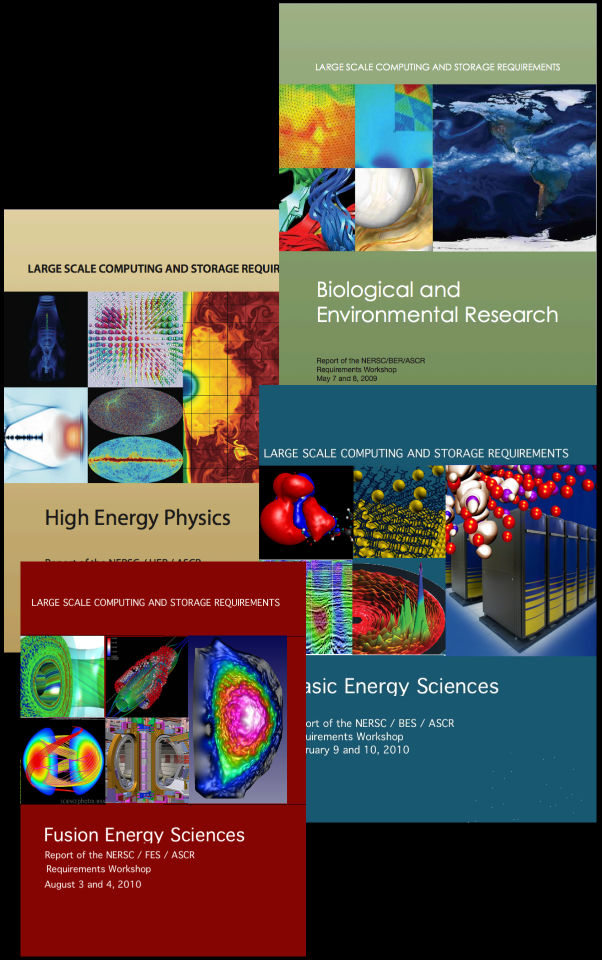

HPC Resources for User Science

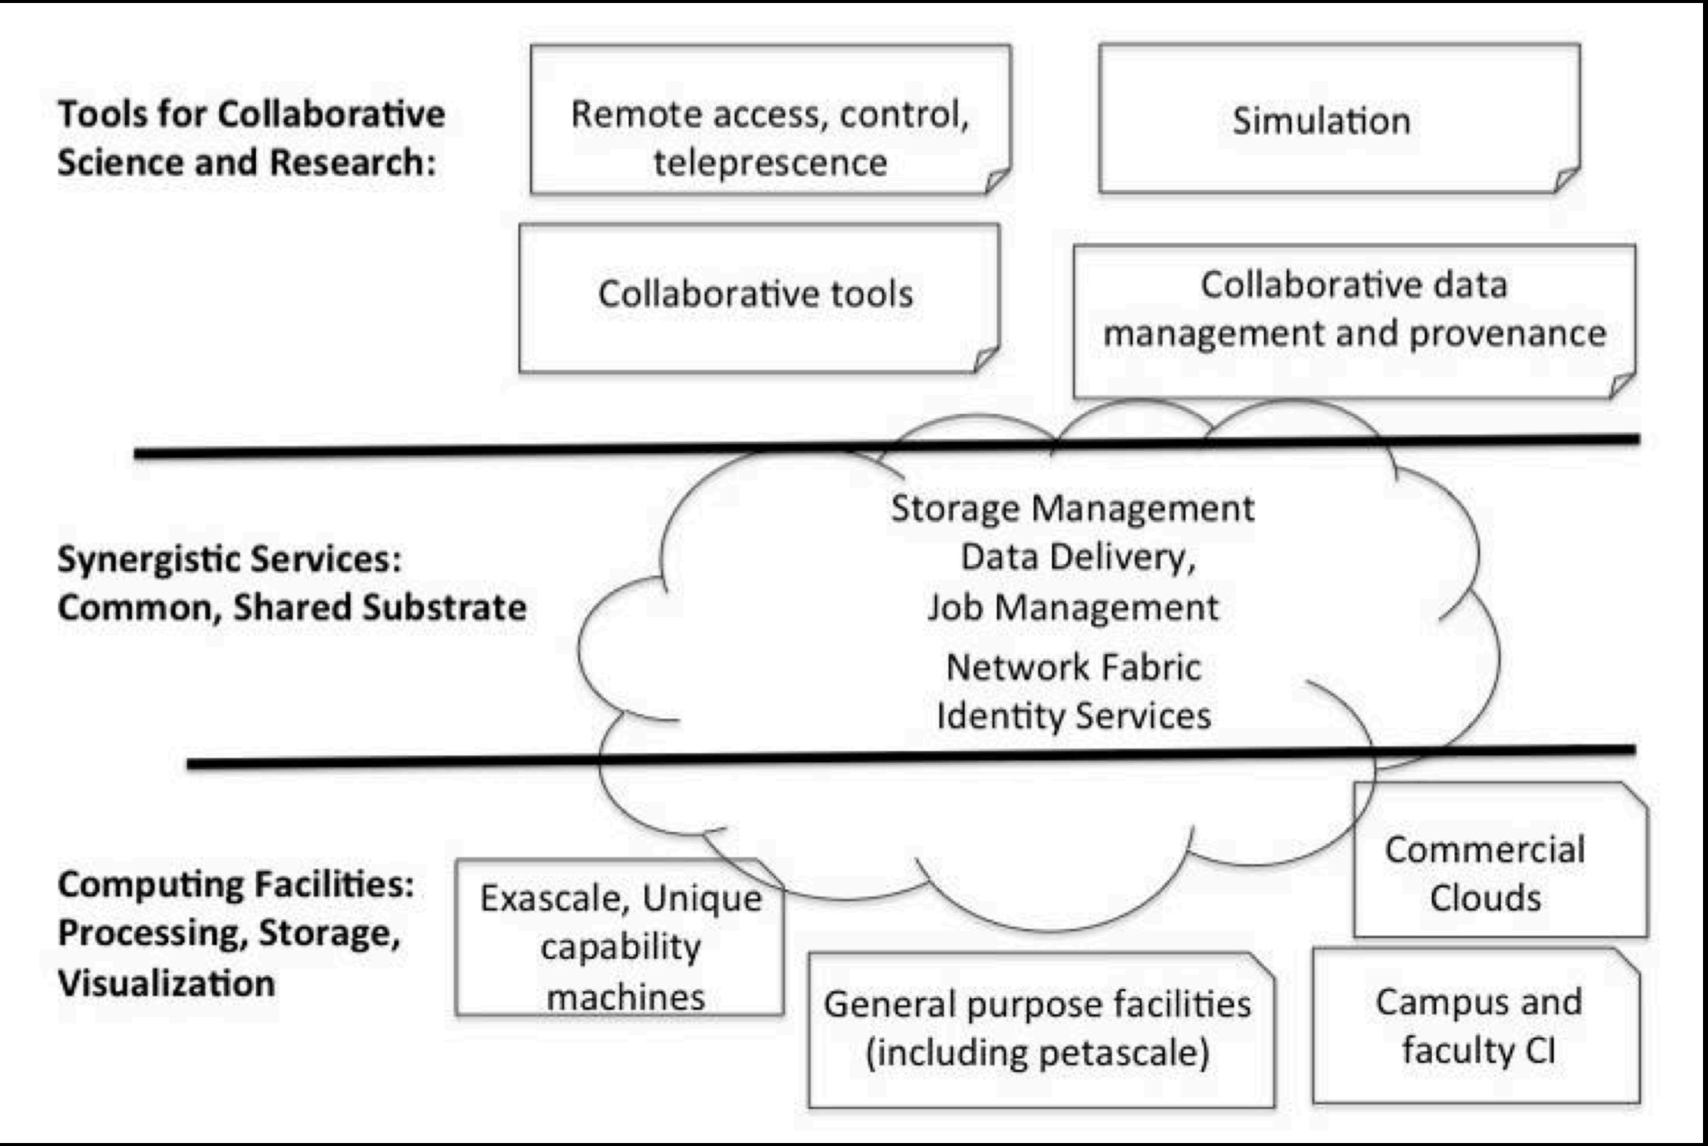

SCESS Report Mount, Skinner, et al.

DOE Big Data *Volume, velocity, variety, and veracity*

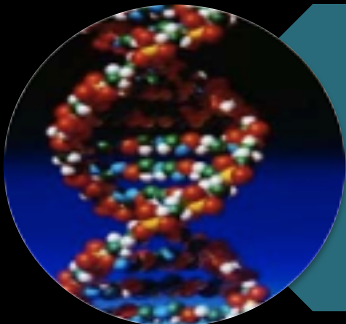

Biology

- *Volume:* Petabytes now; computation-limited
- *Variety*: multi-modal analysis on bioimages

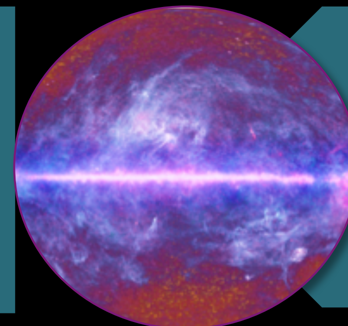

Cosmology & Astronomy:

- *Volume:* 1000x increase every 15 years
- *Variety:* combine data sources for accuracy

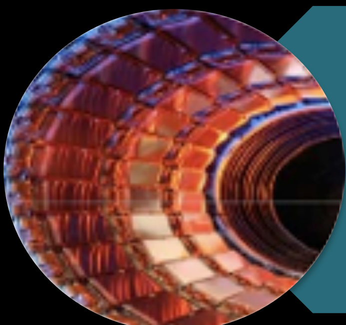

High Energy Physics

- *Volume*: 3-5x in 5 years
- *Velocity*: real-time filtering adapts to intended observation

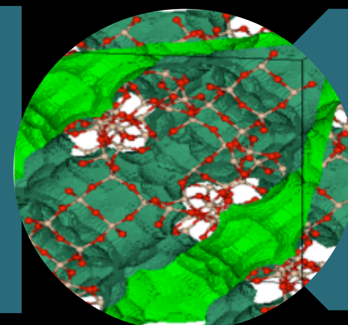

Materials:

- *Variety:* multiple models and experimental data
- *Veracity:* quality and resolution of simulations

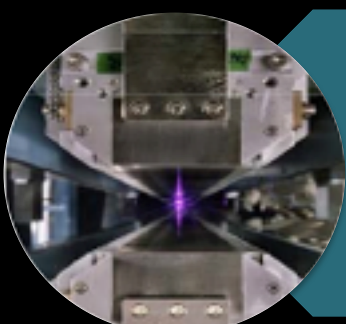

Light Sources

- *Velocity:* CCDs outpacing Moore's Law
- *Veracity:* noisy data for 3D reconstruction

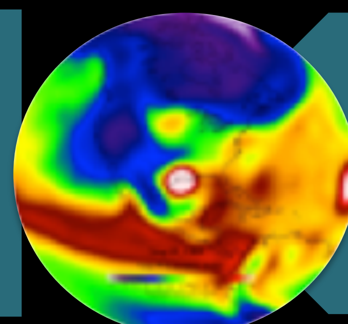

Climate

- *Volume:* Hundreds of exabytes by 2020
- *Veracity:* Reanalysis of 100-year-old sparse data

Top 15 Science Data Projects in NERSC Filesystem

Daya Bay
Urban Sensor + Sim
Supernova (PTF)
Cosmology Sim
Planck (CMB)
Climate 100
ALS (Light Source)
Climate Reanalysis
BAO
Alice (LHC)
SN Factory
STAR Detector
Extreme Weather
Materials Project
JGI (Genomes)

Biggest online data sets are from:

Experimental facilities
Observations
Simulations
Reconstructed observation
Sensors

Total for these projects:

1.5 Petabytes of Disk

4.5 Petabytes of Tape

Data Growth is Outpacing Computing Growth

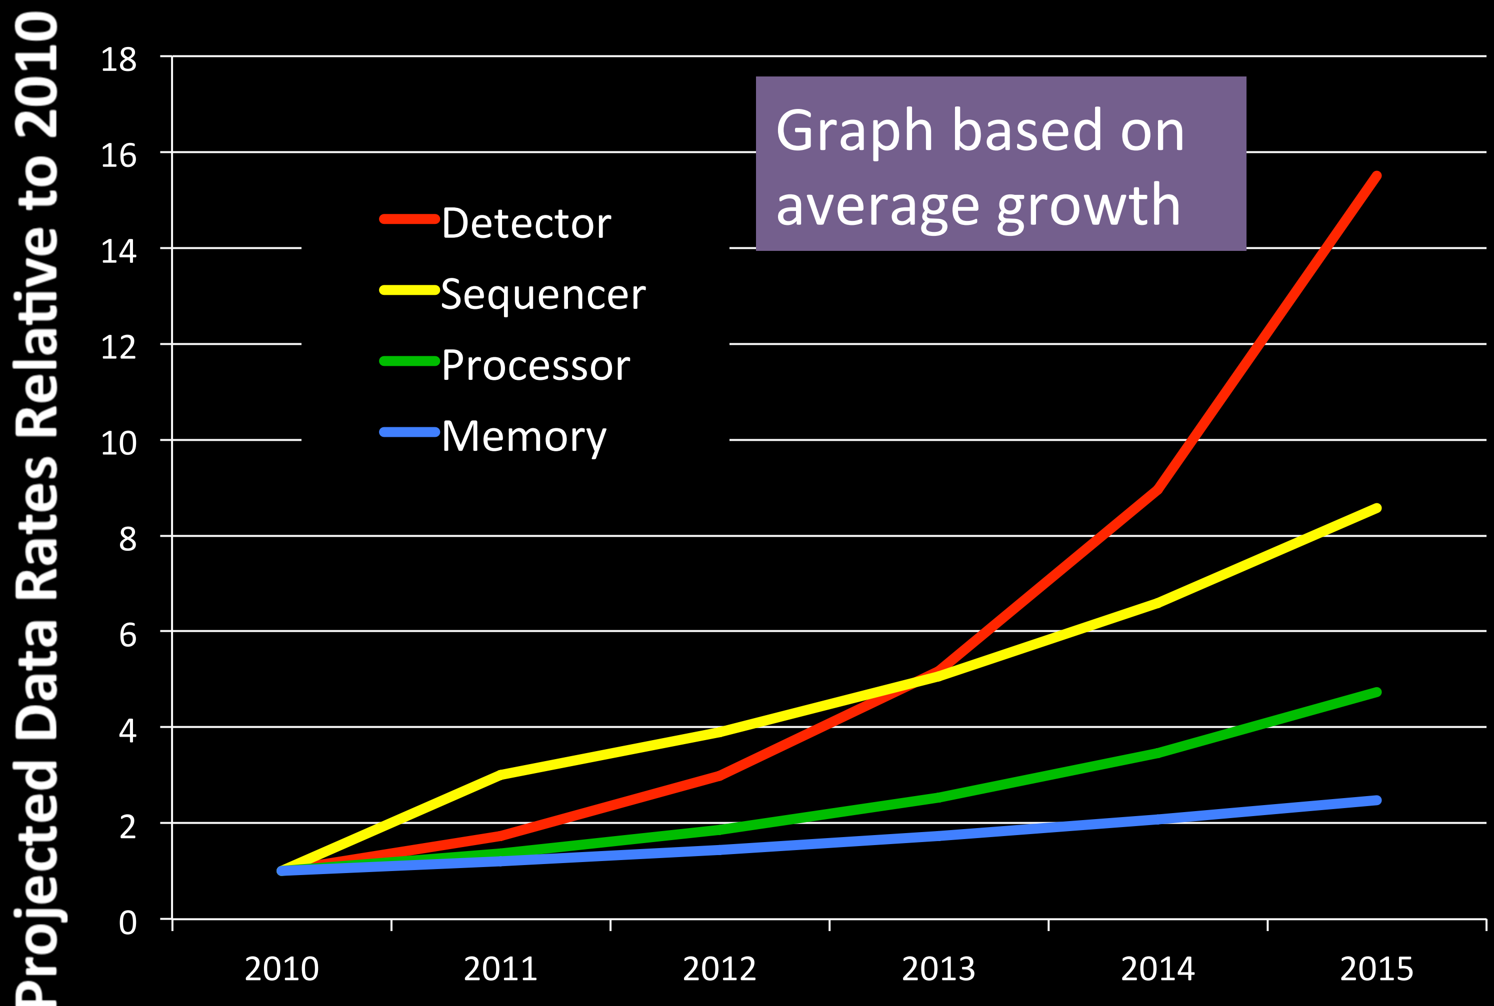

NERSC and Esnet: WAN data trends

Daily WAN traffic in/out of NERSC over the last decade

Roughly 10x 2011-2016

Automated data pipelines for large scale genomics , LHC, image processing

Community access to data and analysis, gateways

Data at NERSC is secure, reliable, fast, open, flexible

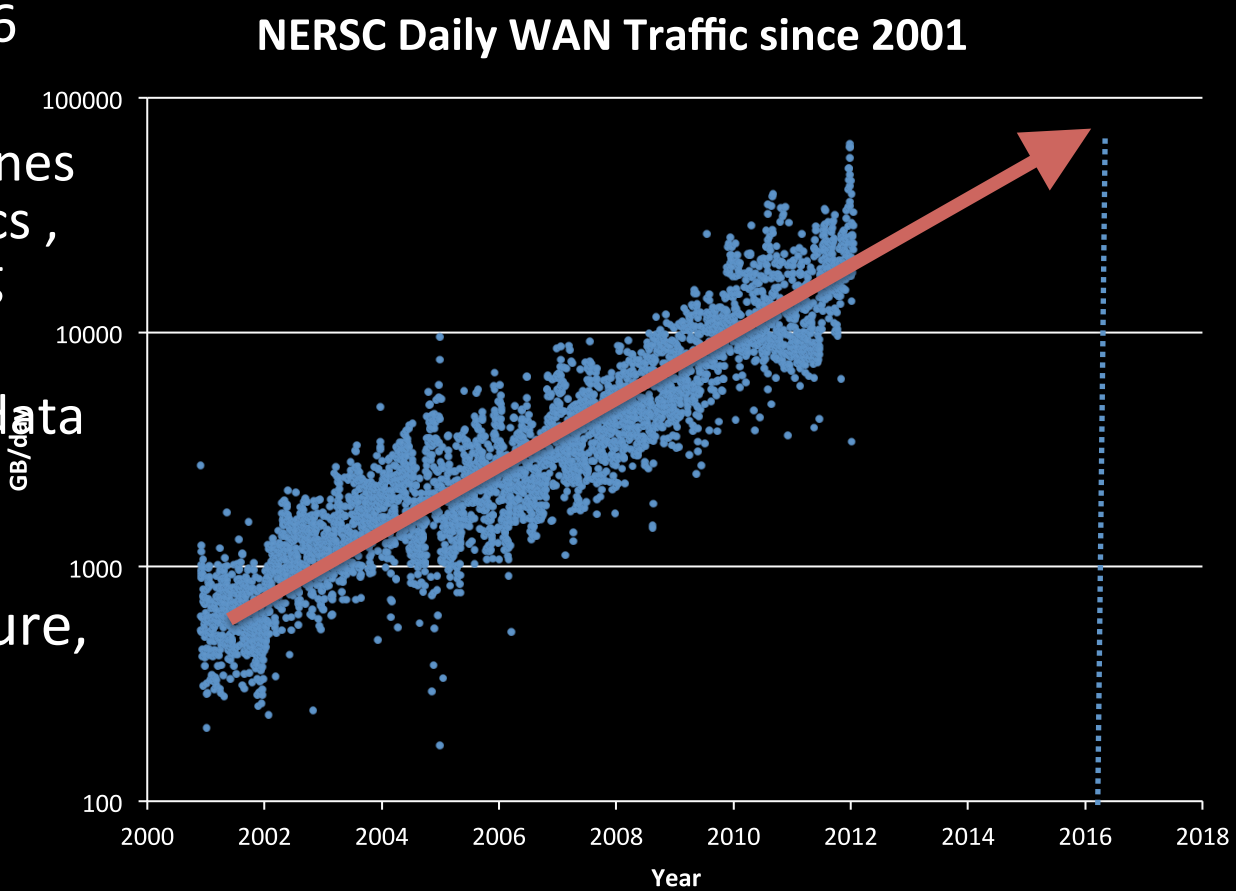

Data from DOE facilities: Tomorrow is already here

Experimental facilities will be transformed by high-resolution detectors, advanced data analysis techniques, robotics, software automation, and programmable networks.

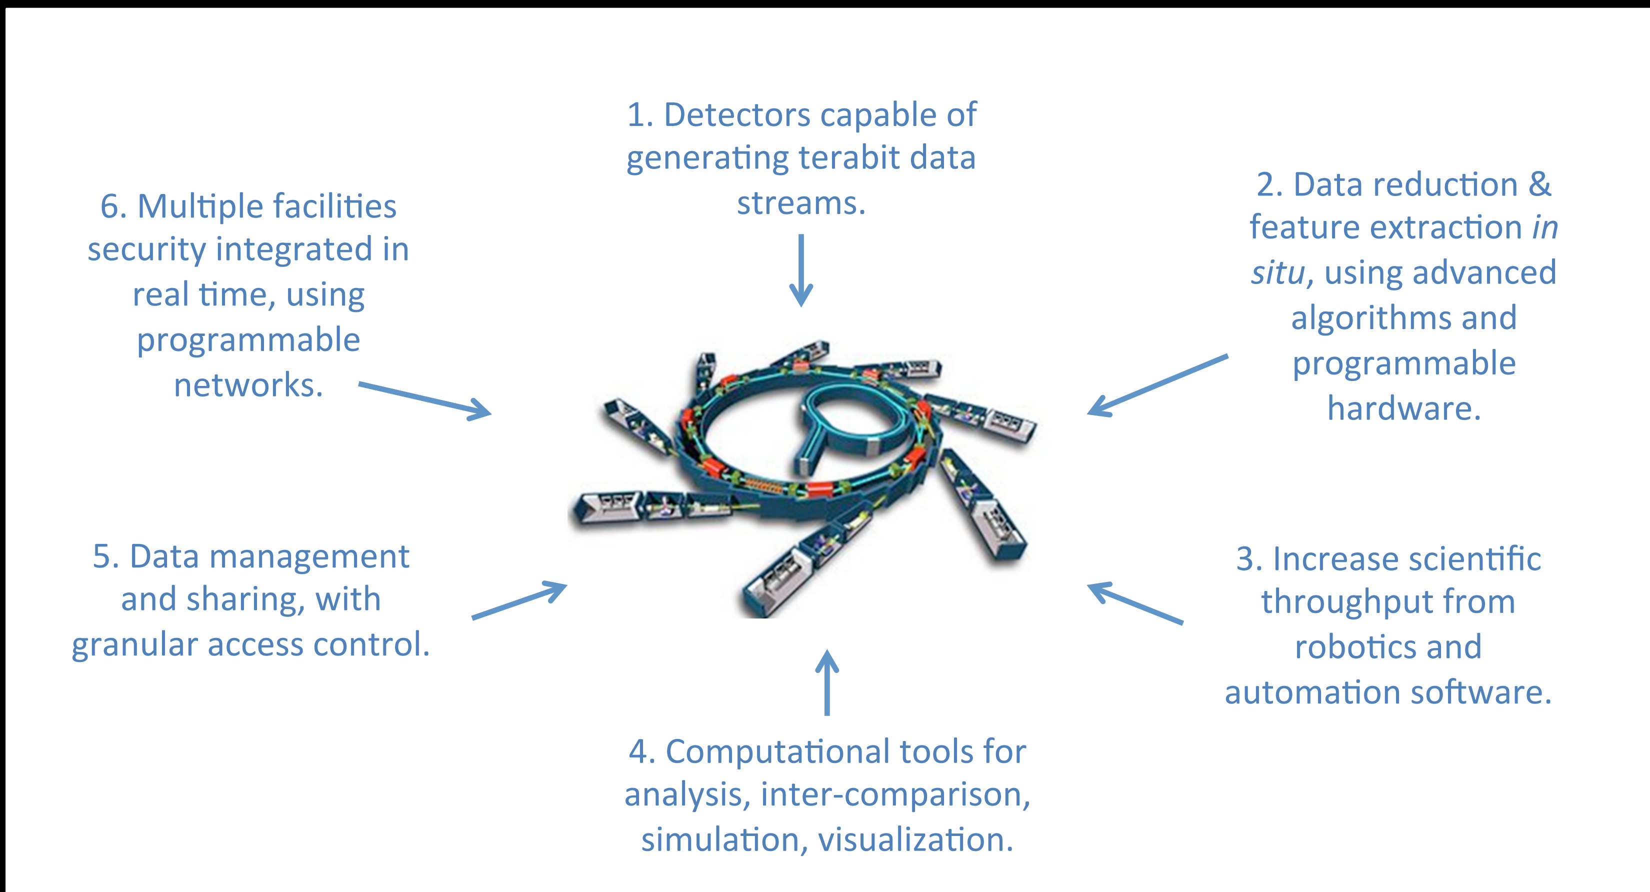

Transforming Science: Finding Data

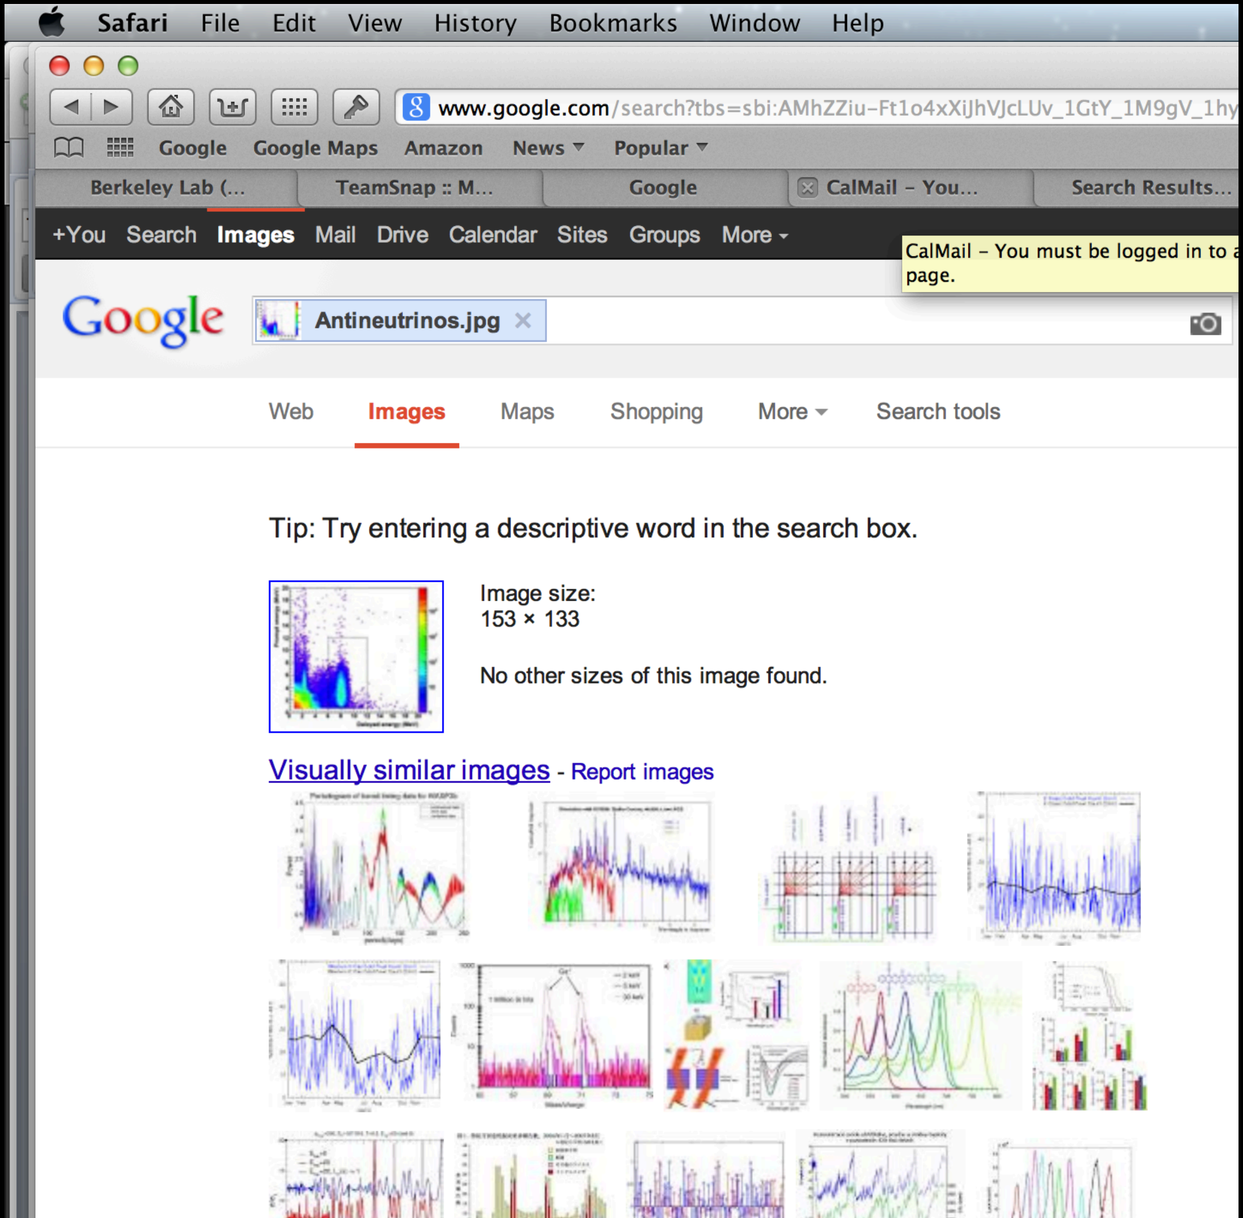

(Just a Few) New Data Methods : Tools vs. APIs

Simulation, data analysis, and visualization tools integrated in flexible portals.
Flexible execution frameworks on HPC (HTC, ensembles, VM images, etc.)
Advanced scalable databases (KVP w/ mapreduce) , ML at scale, in-situ analysis
Big Data thumbnails, synopsis generation, Metadata automation, inferred provenance,
De-noising, inter-dataset correlation, deep search, differential data sharing
Automated agents for opportunistic data QA/QC, data citation,
Social data, curation, community data management, attribution, Big Data reproducibility

RESTful Interface Circa 1955

190 REST APIs for Scientiific Data and Computing

http://www.w3.org/community/hpcweb/

programmableweb

[W3C Community and Business Groups](http://www.w3.org/community/hpcweb/)

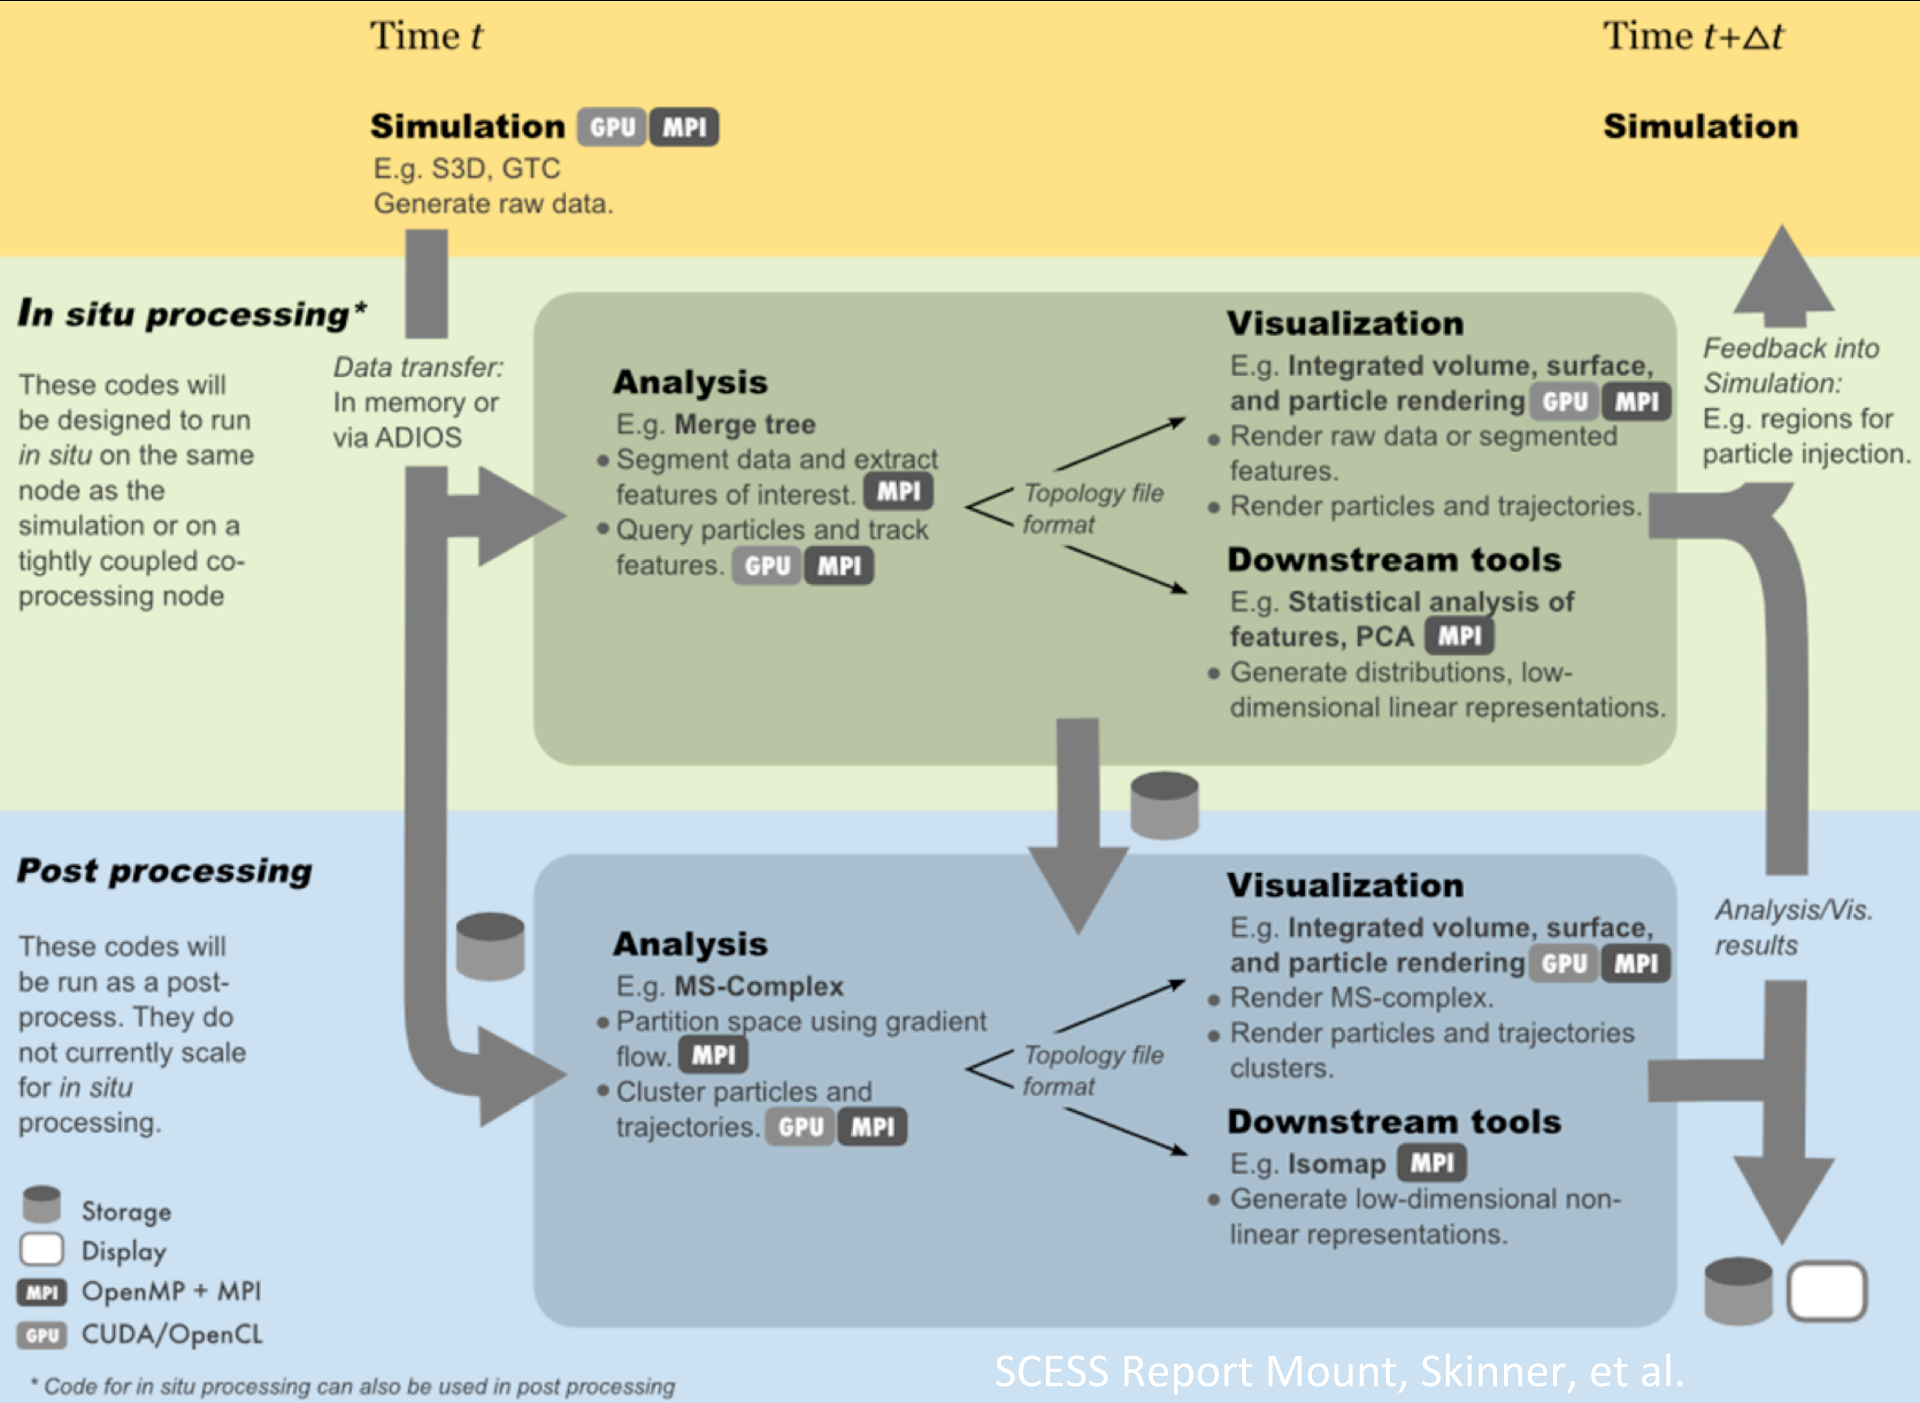

SCESS Report Mount, Skinner, et al.

Scientific Workflow envisioned

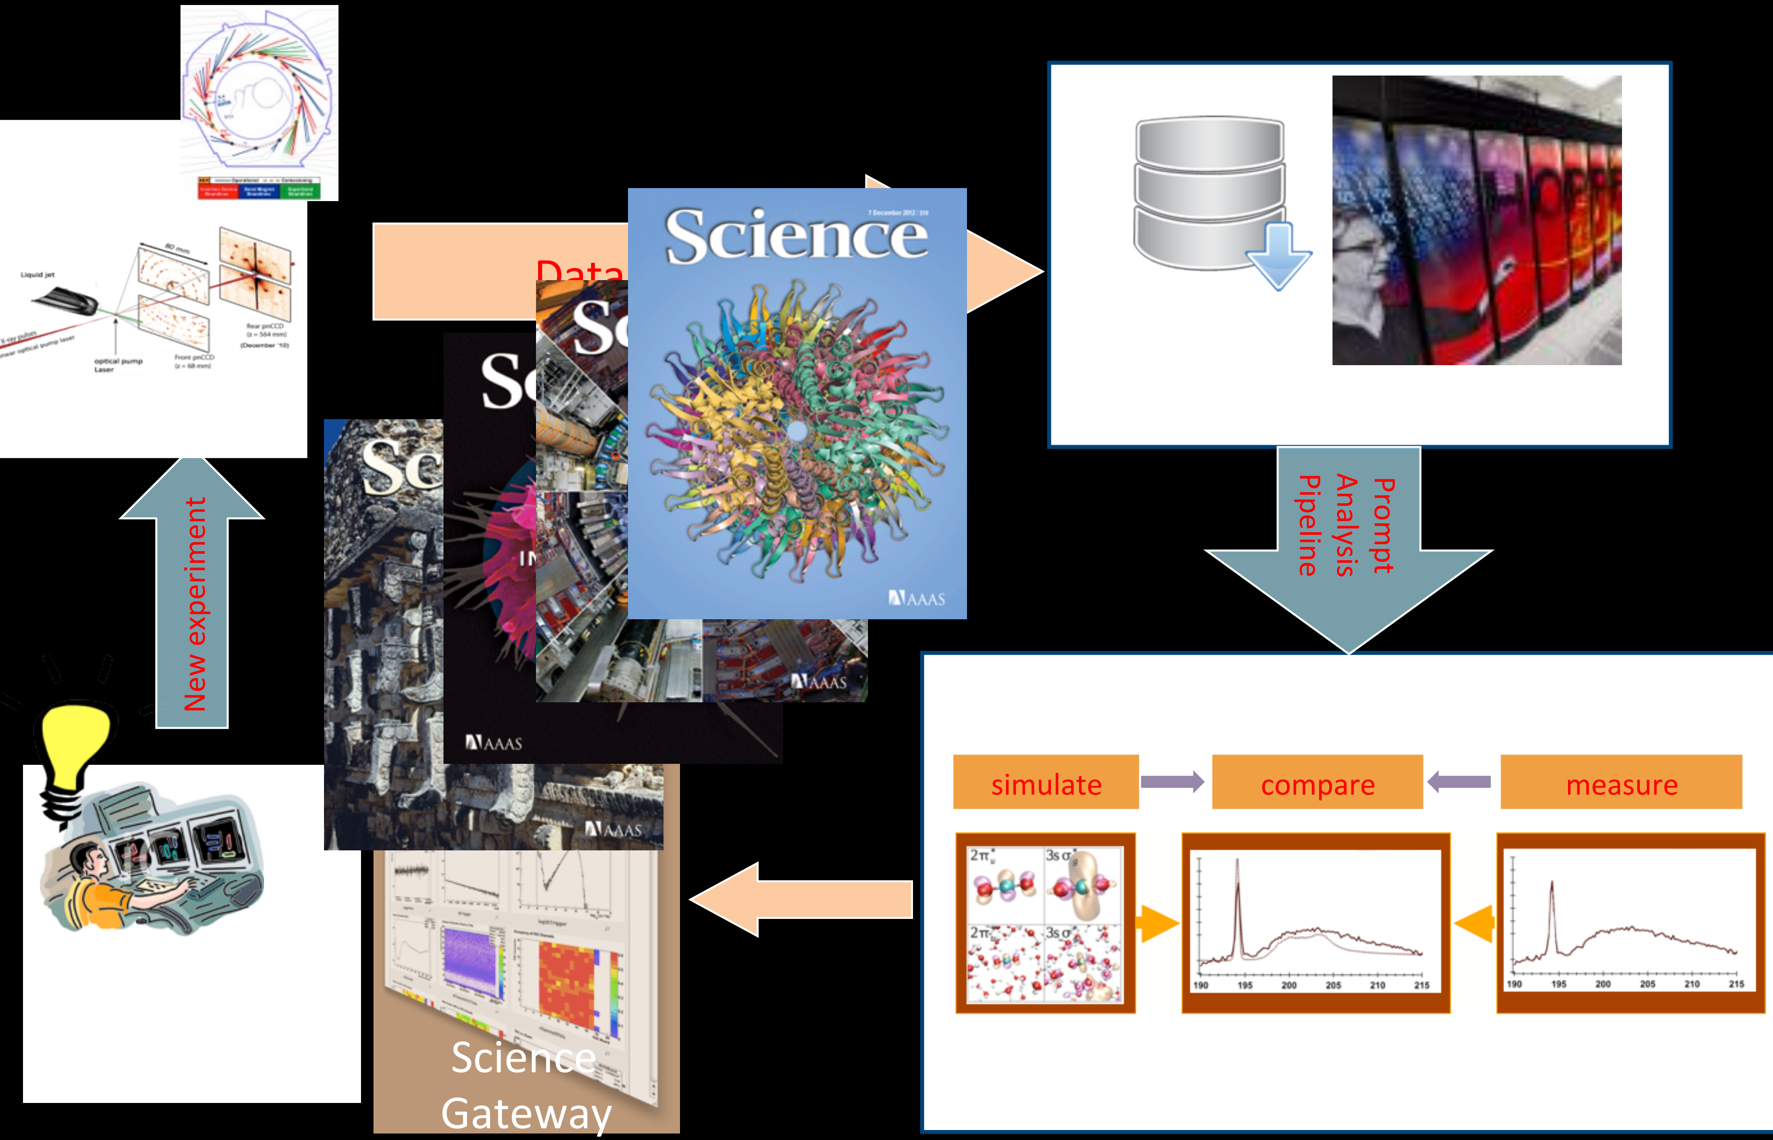

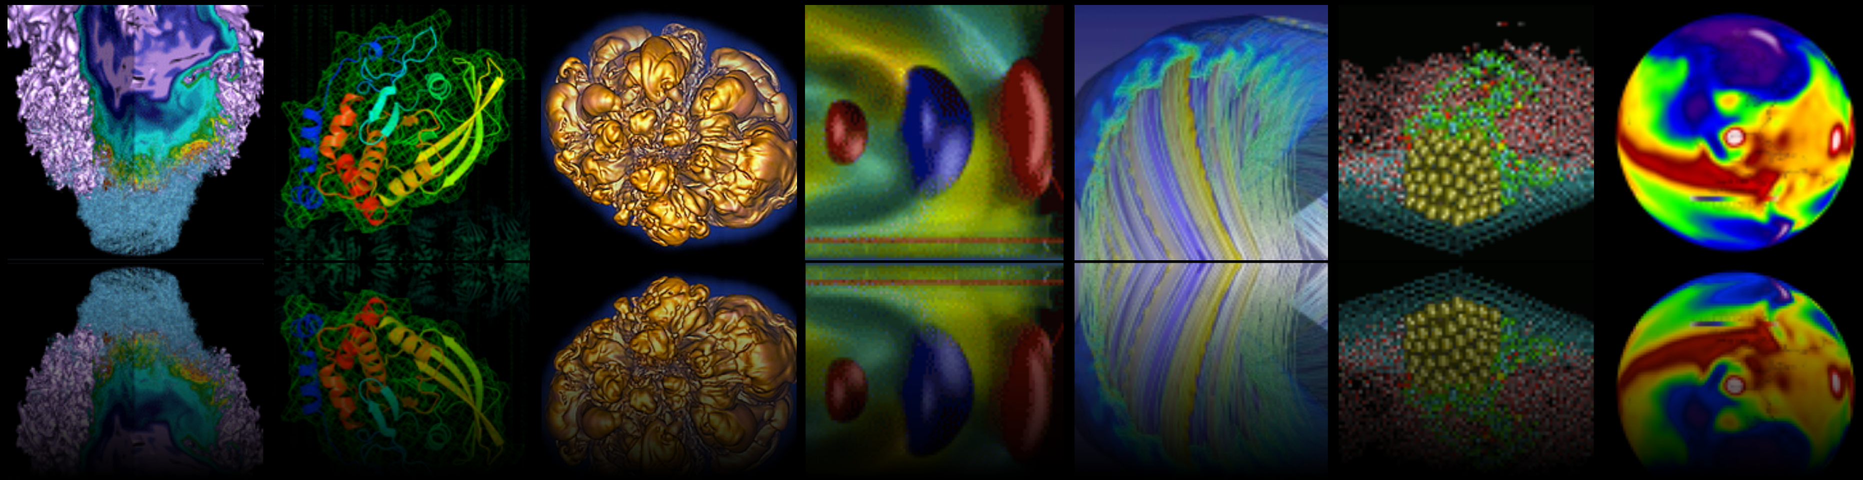

Extreme Data Science

The scientific process is poised to undergo a radical transformation based on the ability to access, analyze, simulate and combine large and complex data sets.

New Models of Discovery

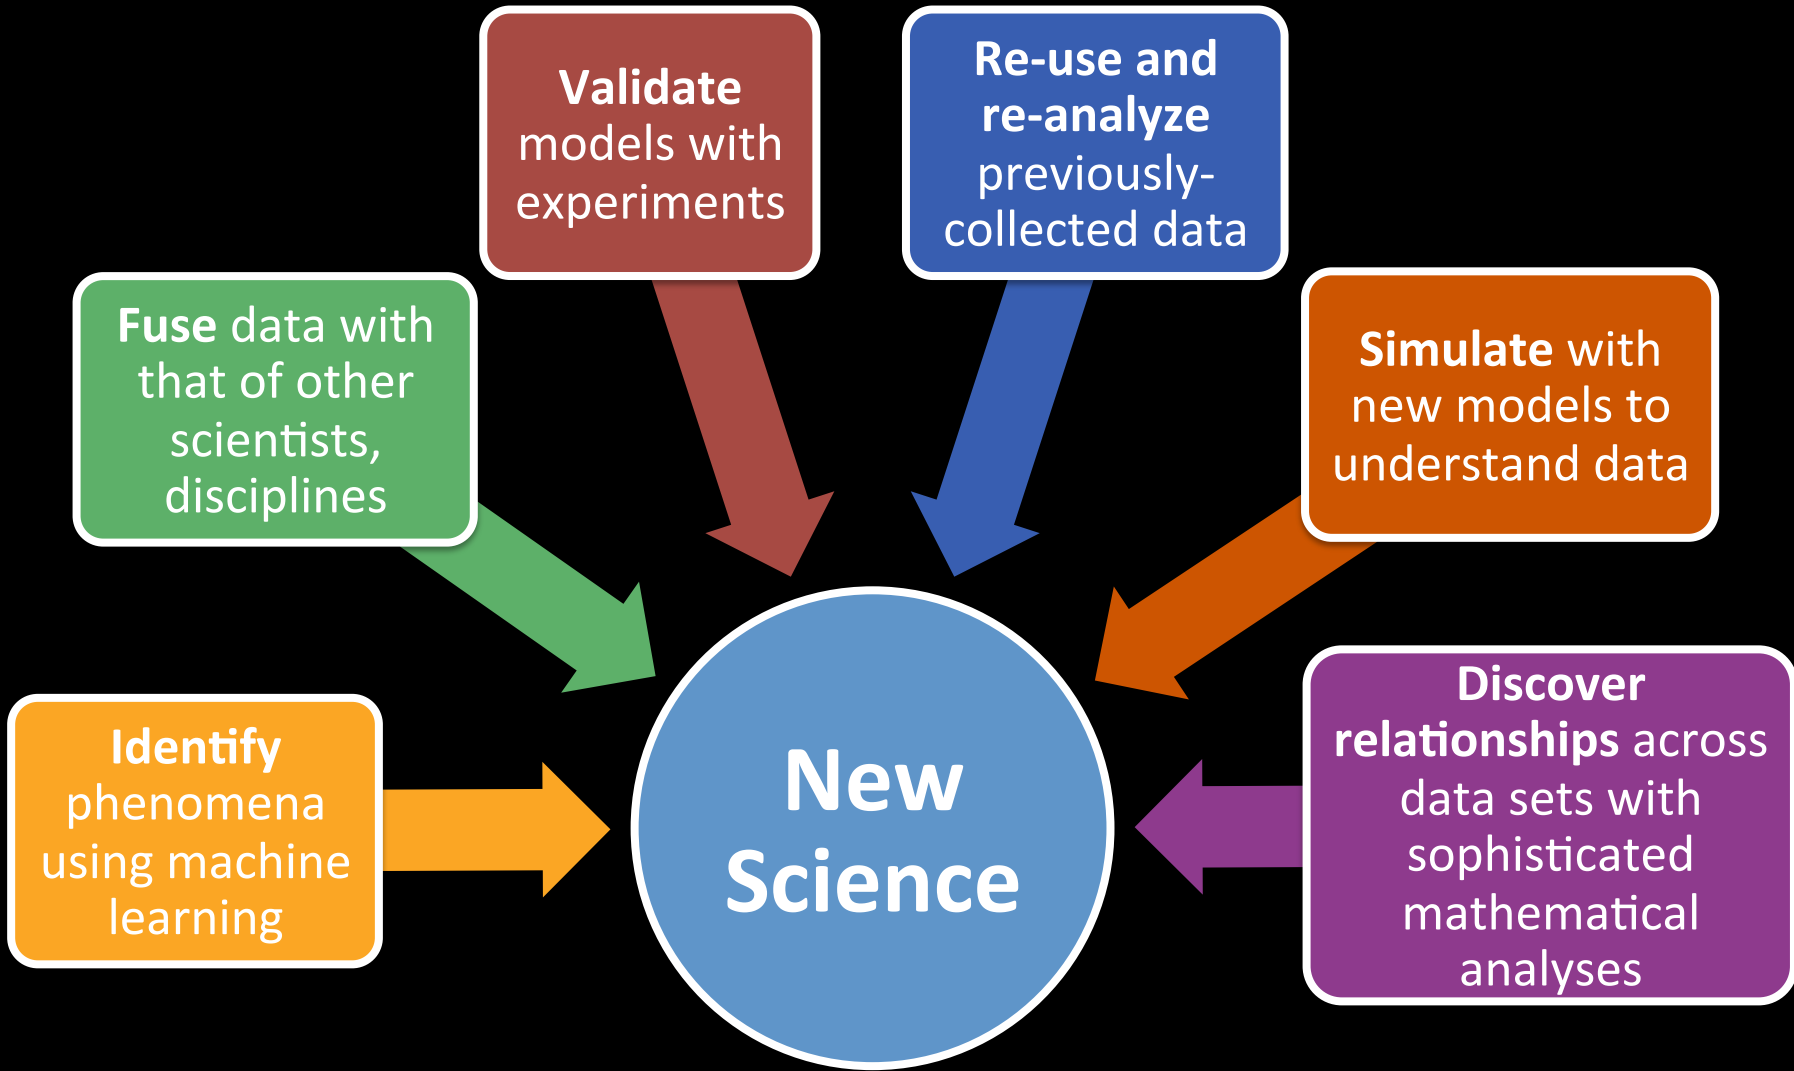

Identify Phenomenon using Machine Learning

- **TECA Toolkit today**
  - Automatic detection of cyclones, atmospheric rivers, and more
  - Analysis time years to minutes

- **Climate Analysis in 2031**
  - Machine learning for all events
  - Automatic metadata generation
  - Fusion of simulations, sensors, etc.
  - Real-time analysis and response

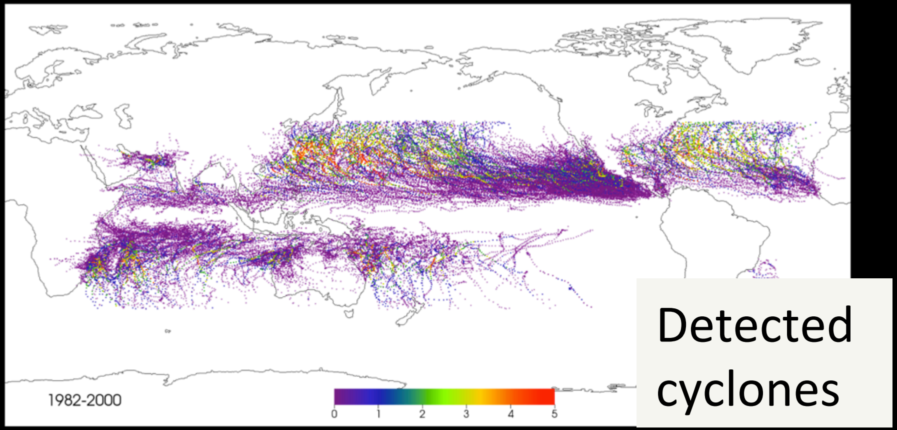

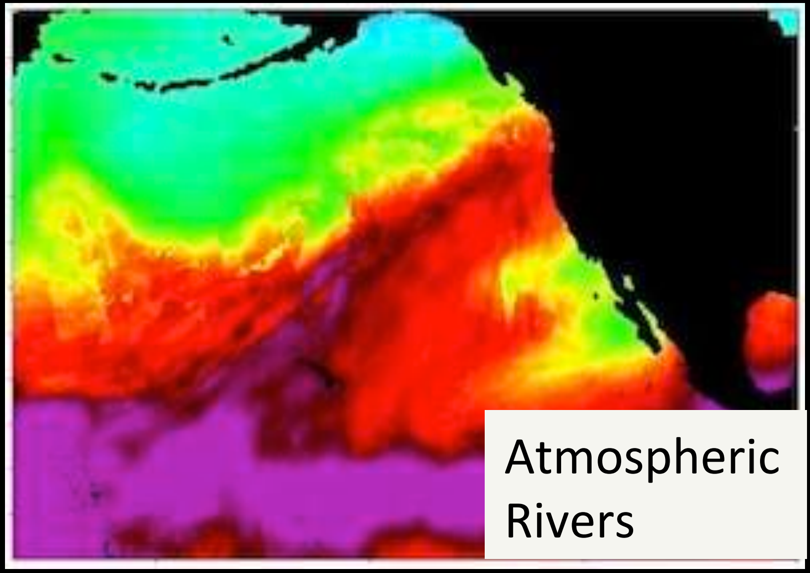

Connecting Data: Tools for Radical Scaling

- **Genomes to Life, KBASE**
  - Make genomics useful to biology
- **Microbes to Biomes**
  - Measuring and modeling the plant microbial biome
- **Quarks to Cosmos**
  - Frontiers bridge energy, intensity, and cosmos
- **Pixels to Knowledge**
  - Replace pixels with models, build kbases on models
  - Leverage repetition toward extreme structural resolution
- **Beamline to Browser**
  - Connect world's Biggest Data instruments to the internet
- **Climate to Weather**
  - Couple world class global models to regional problem solving
- **Materials to machines**
  - Materials Project (replace materials design with search)
  - JCESR (batteries), JCAP (engineered sunlight-to-fuels tech)
  - Defects, functional electronics, nano-to-mesoscale

Big & Fast Filesystems

Powerful Flexible Computing

Advanced Analysis

Machine Learning

Ontologies, Ksystems

Databases/KVP

MapReduce

Software Defined Networking

High-throughput

Automated workflows

Filtering, De-Noise and Curating Data

AmeriFlux & FLUXNET: 750 users access carbon sensor data from 960 carbon flux data years

Arno Penzias and Robert Wilson discover Cosmic Microwave Background in 1965

Re-Use and Re-Analyze Previously Collected Data

- **Materials Genome Initative**
  - Materials Project: 4500 users 18 months!
  - Scientific American "World Changing Idea" of 2013 – what about 2031?

Unbounded computing requirements for simulation and analysis

Multi-modal analysis of Brain Connectivity

Analyze brain connectivity at multiple scales:
From cells and regions to complex neural circuits.

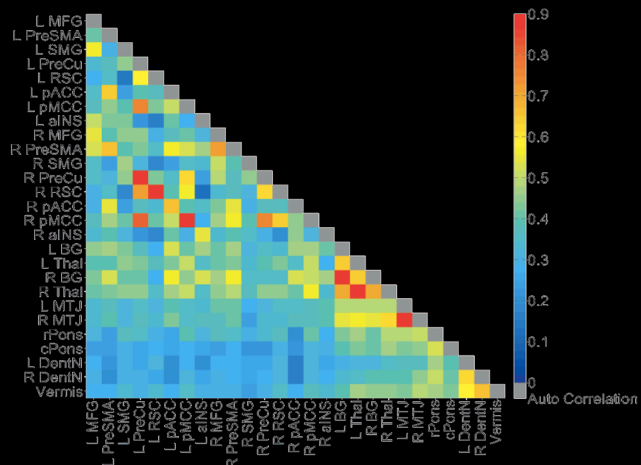

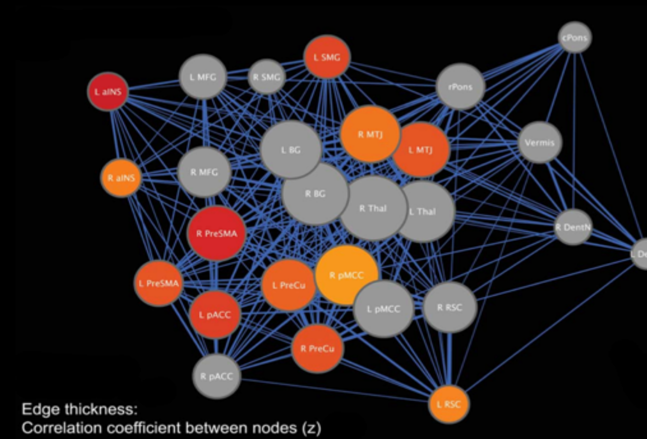

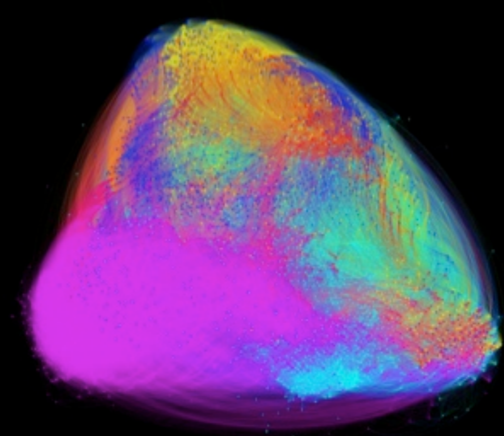

Brain Connectivity Graphs: Jesse Brown, Bill Seely (UCSF)

- **Improve understanding of brain pathology.**
- **Enable personalized treatment options.**

Big Picture: Advancing Scientific Knowledge Discovery

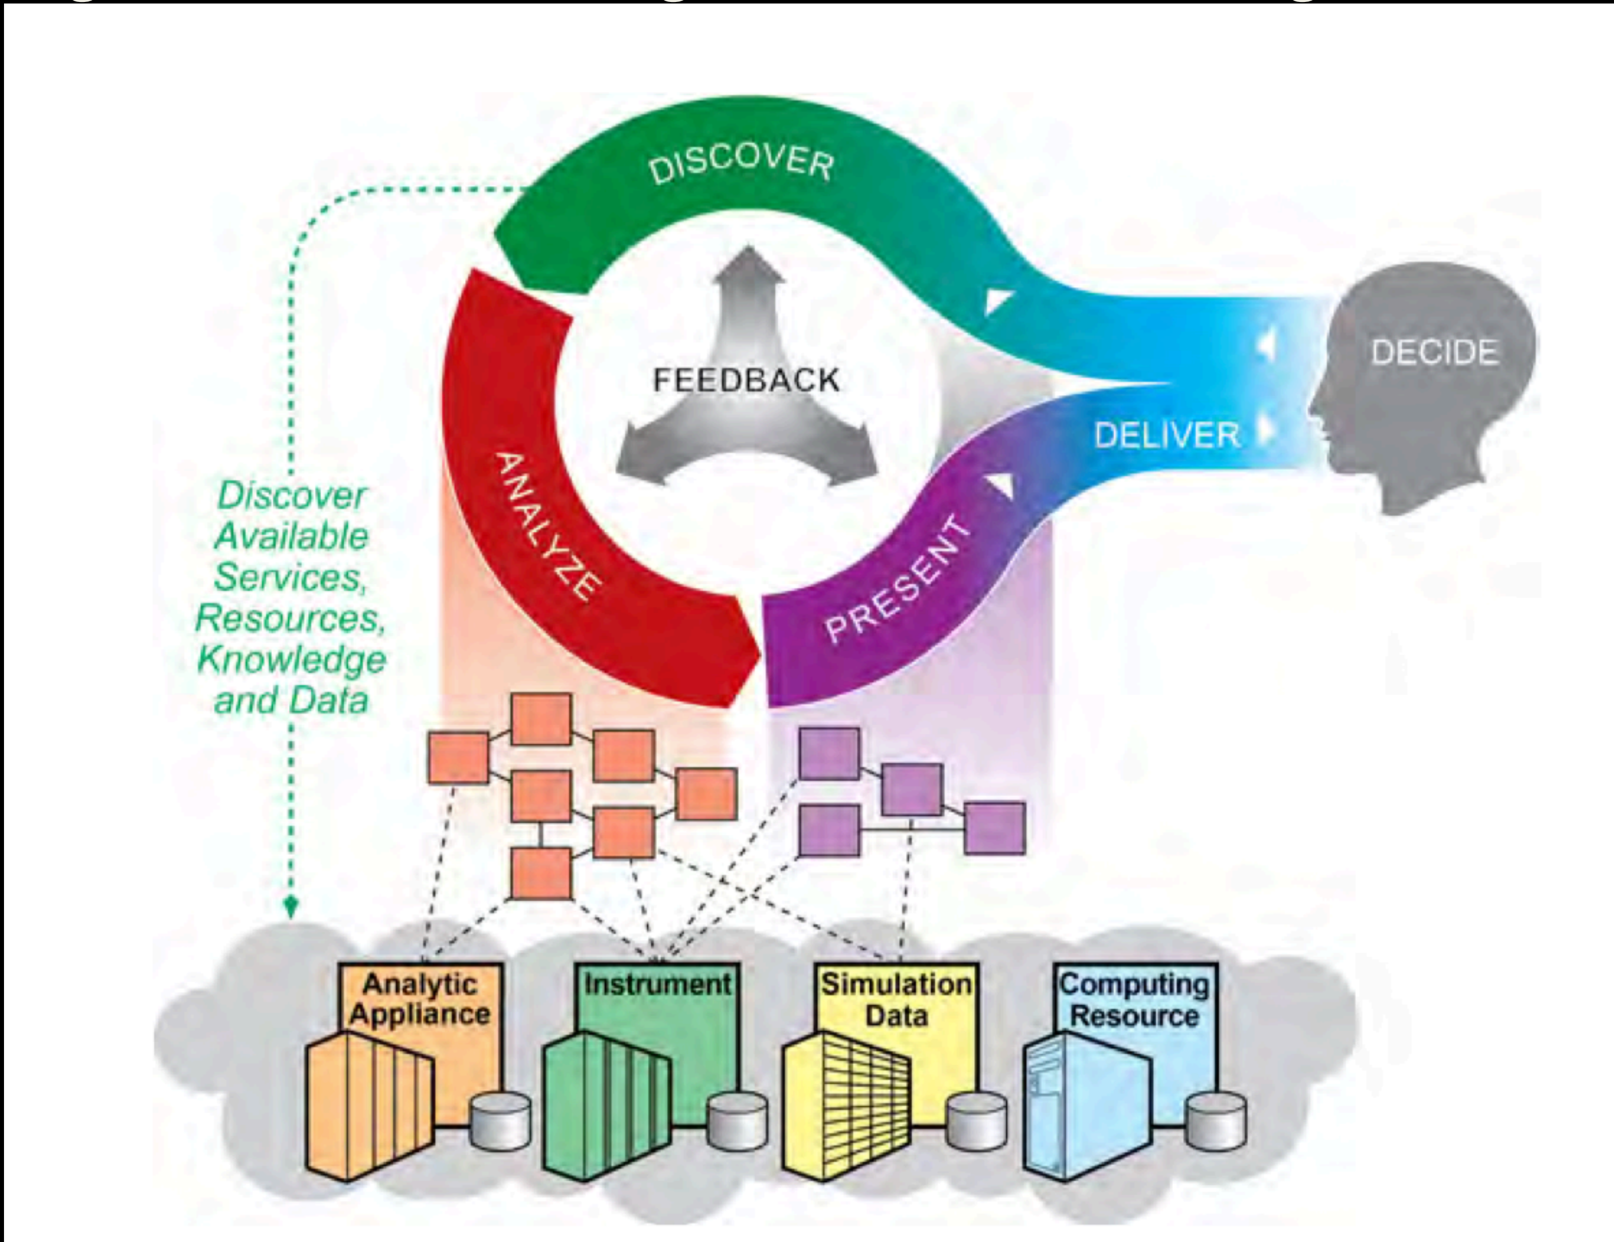

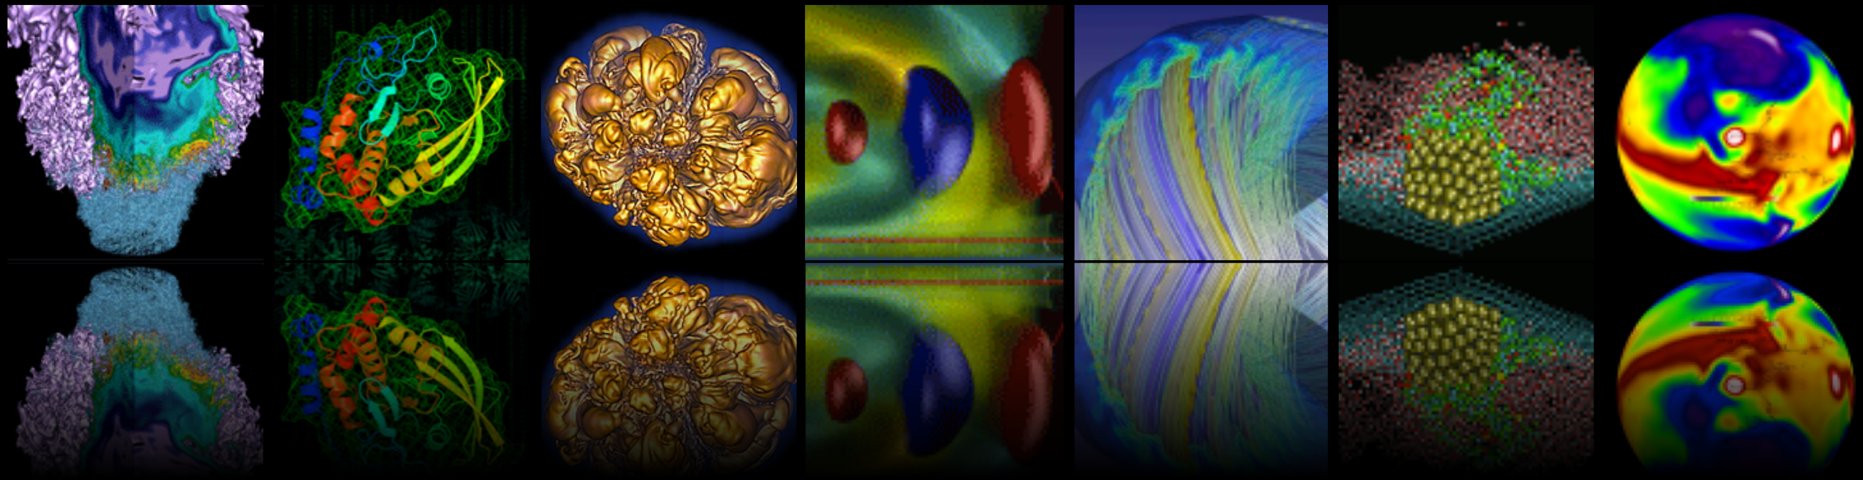

How do we get there?

Extreme Data Science Components

Extreme Data Scientific Facility (XDSF) Concept

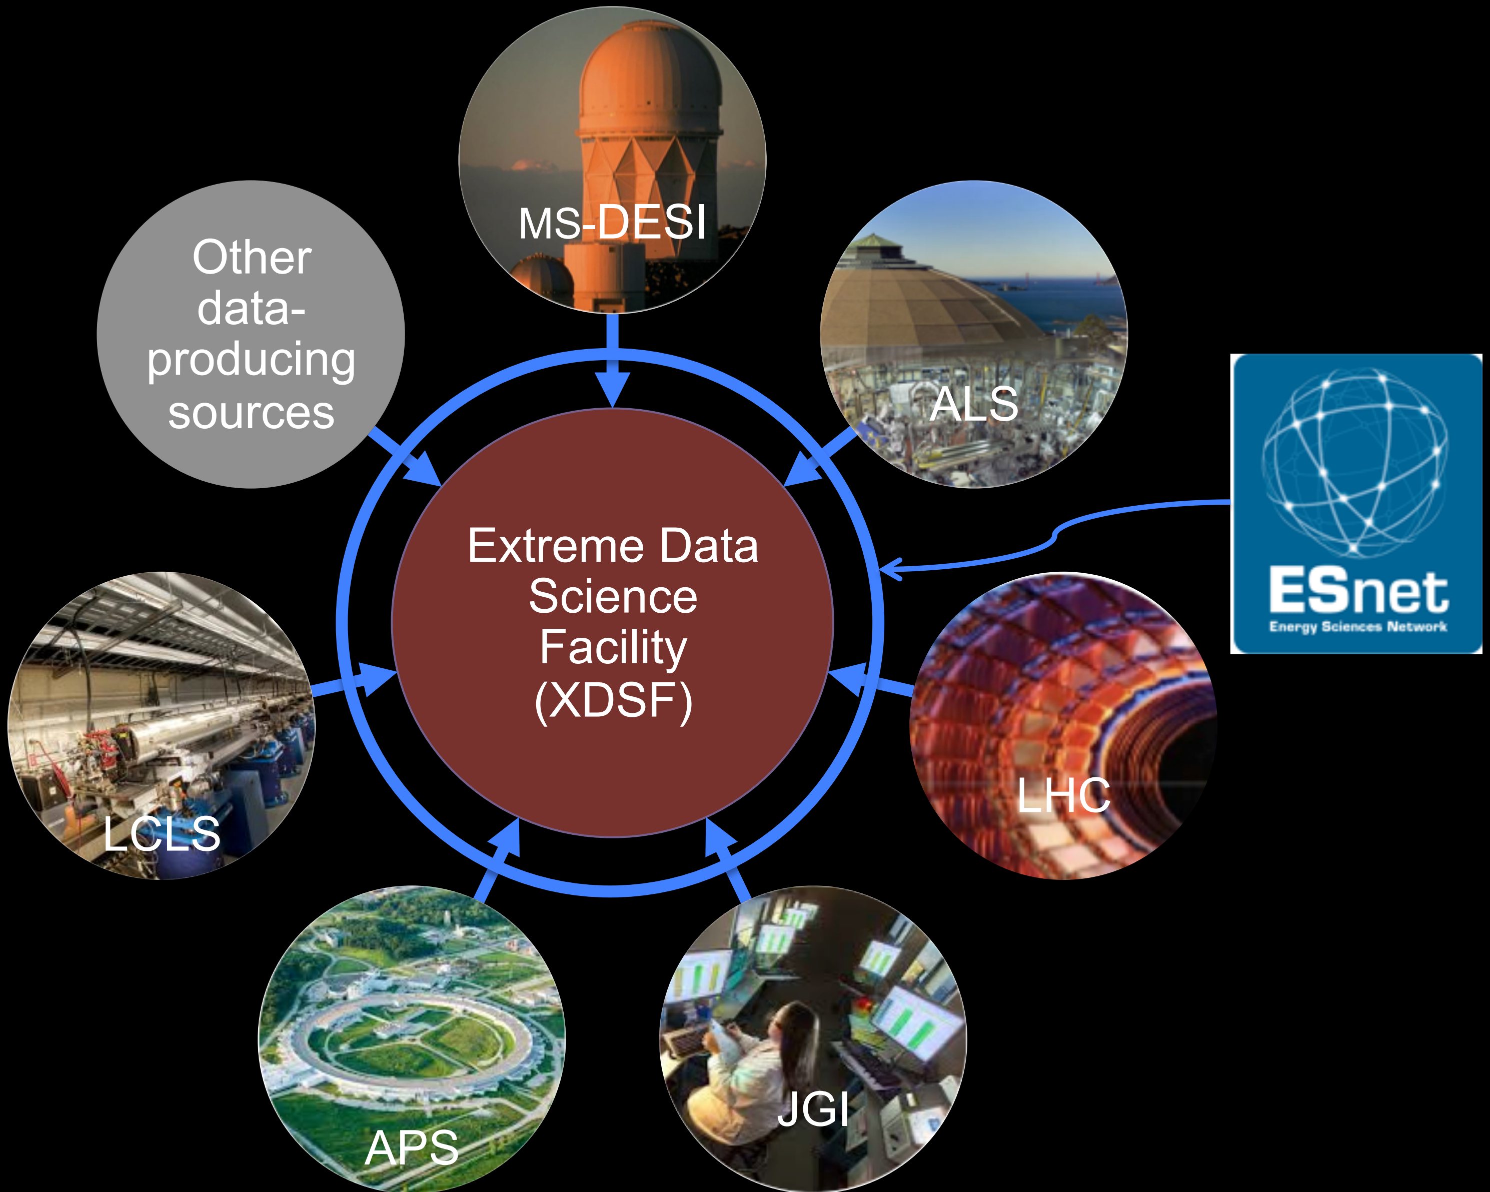

XDSF: Will bring scientists together with data researchers and software engineers

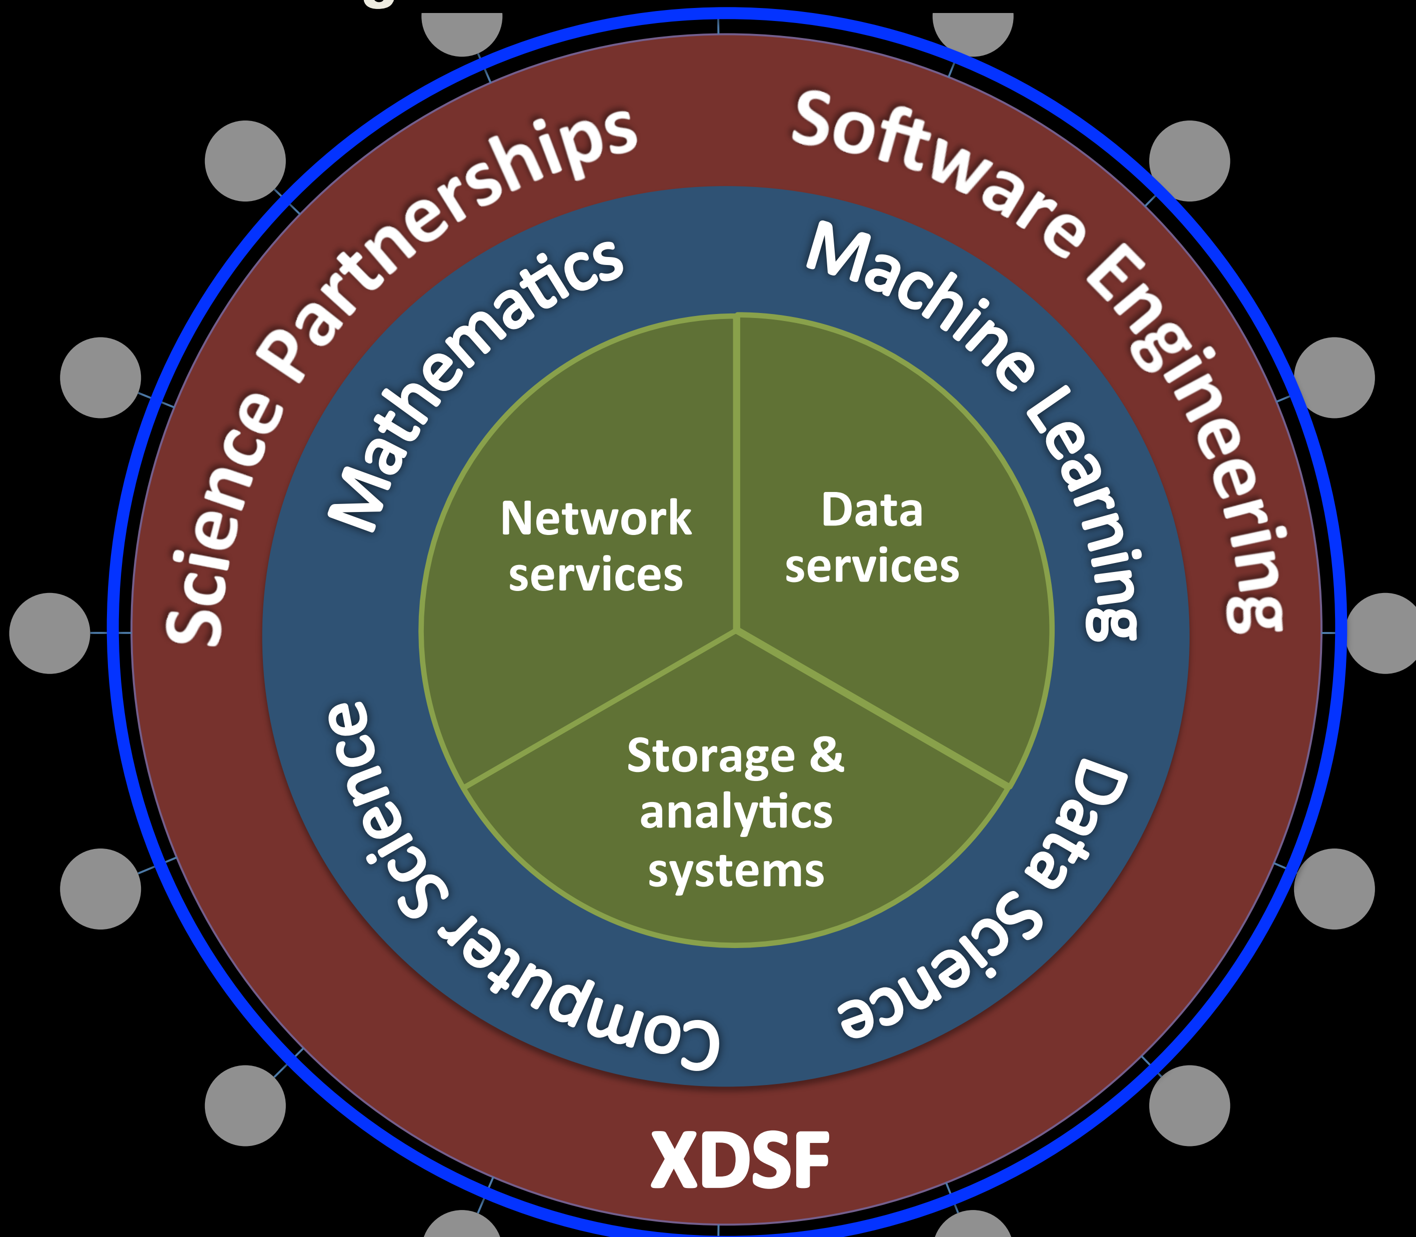

Software Defined Networking: A New Kind of Network (Especially for Science?)

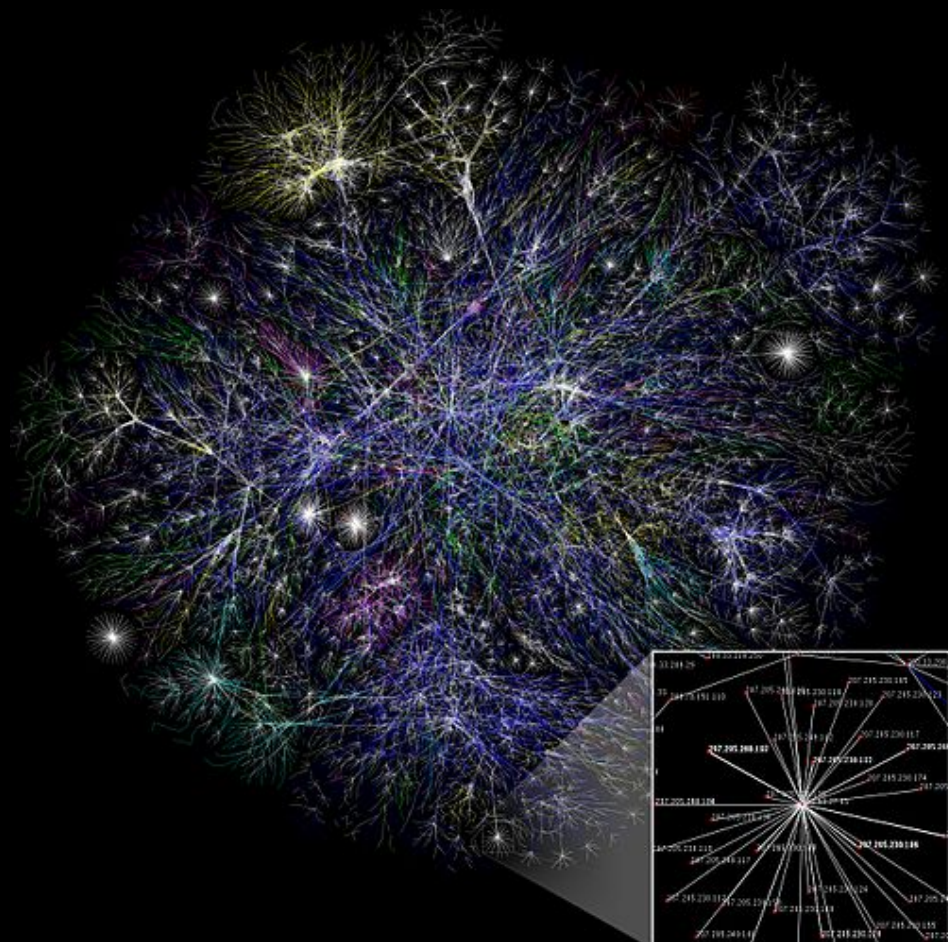

- **Internet is a black box—huge hidden complexity**
- **ESnet connects us to other labs and internet with novel "big science data" communications**
- **"Software Defined Network" demo'd in October:**
  - Automatically adapts to large data transfers
  - Faster, cheaper, more flexible
  - Demo with Infinera. Brocade

See Greg Bell talk here at 4pm.   Vern Paxson Plenary tomorrow

A view of the application space: Simulation and Data

7 Giants of Data	7 Dwarfs of Simulation
Basic statistics	Monte Carlo methods
Generalized N-Body	Particle methods
Graph-theory	Unstructured meshes
Linear algebra	Dense Linear Algebra
Optimizations	Sparse Linear Algebra
Integrations	Spectral methods
Alignment	Structured Meshes

NAP "Frontiers in Massive Data Analysis"

Data structures, Algorithms, Distributed Systems

- **Fastbit & Fastquery**
  - specialized compression and object-level search
  - bitmap indexing methods
  - Theoretically optimal and 10x-100x faster in practice

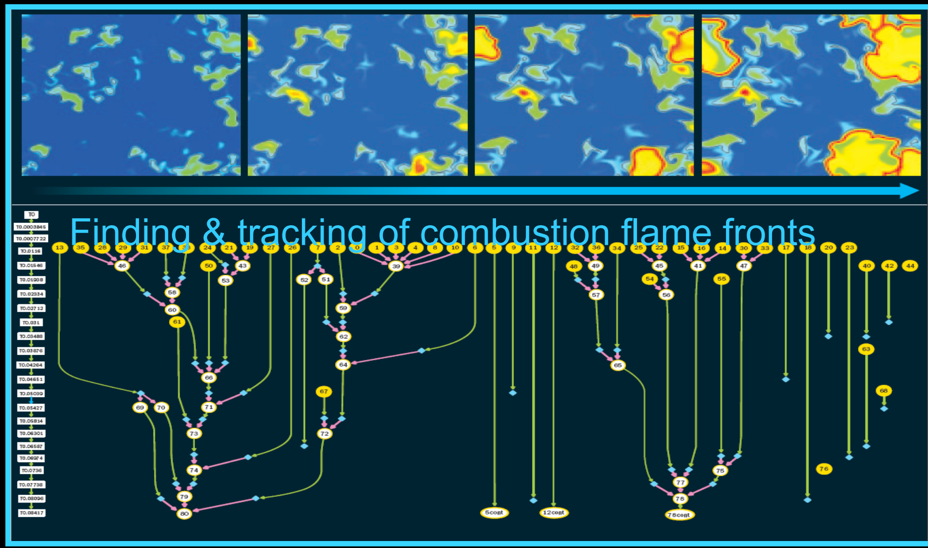

J. Wu, A. Shoshani, A. Sim, D. Rotum

- **Tigres: Design *templates* for scientific workflows**
  - Explicitly support Sequence, Parallel, Split, Merge

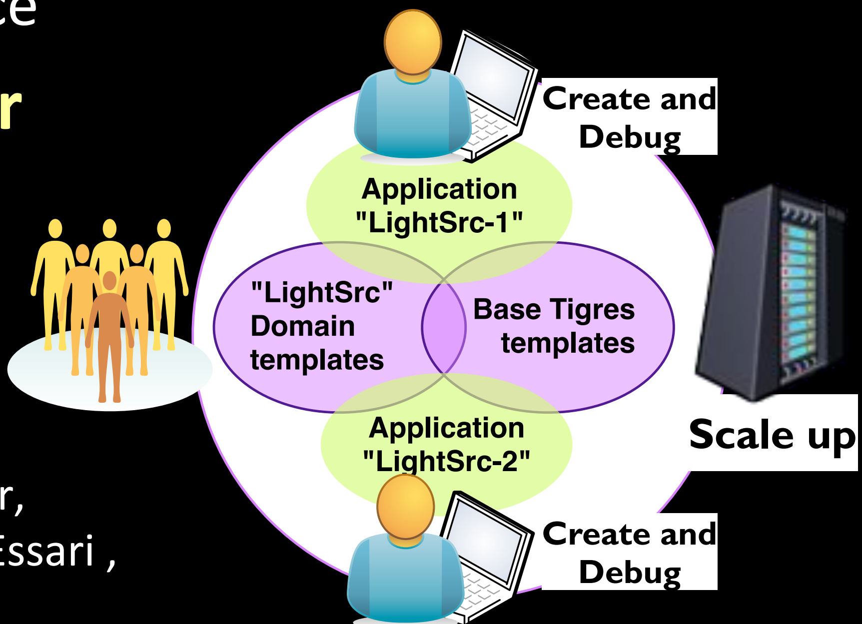

L. Ramakrishnan, Valerie Hendrix, Daniel Gunter, Gilberto Pastorello, Ryan Rodriguez, Abdellilah Essari , Deb Agarwal

Technology for Scientific Data

- Compute
- On-Package DRAM
- Capacity Memory
- On-node-Storage
- In-Rack Storage
- Interconnect
- Global Shared Disk
- Off-System Network

Compute Intensive Arch

Goal: *Maximum Computational Density and local bandwidth for given power/cost constraint.*

Maximizes bandwidth density near compute

Data Intensive Arch

Goal: *Maximum Data Capacity and global bandwidth for given power/cost constraint.*

Bring more storage capacity near compute (or conversely embed more compute into the storage).

Requires software and programming environment support for such a paradigm shift

Systems for Scientific Data

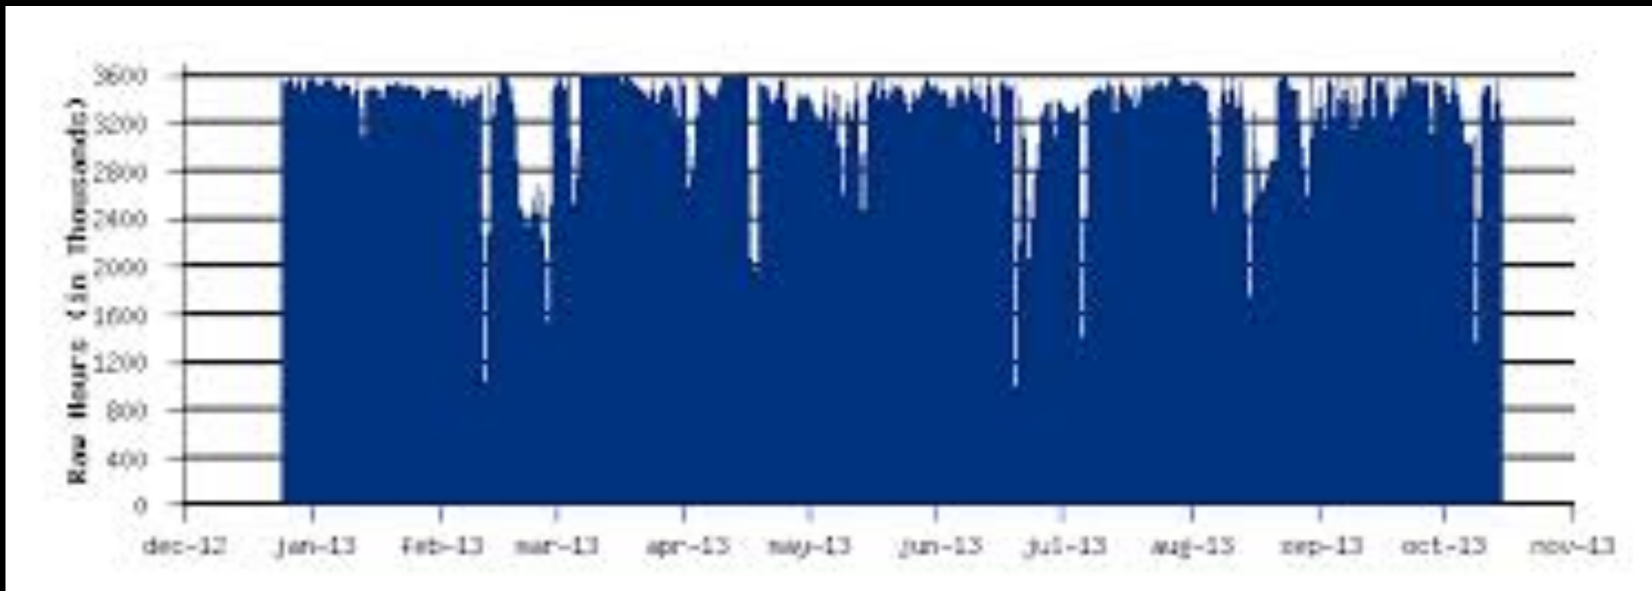

- 95% utilization, but the users wait
- Real-time analysis on streams
- Interactive access to data

Programming Challenge? Science Problems Fit Across the “Irregularity” Spectrum

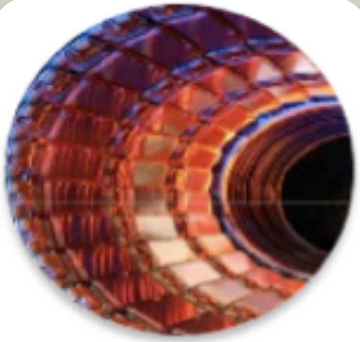

Massive Independent Jobs for Analysis and Simulations

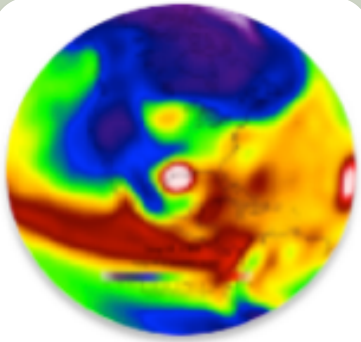

Nearest Neighbor Simulations

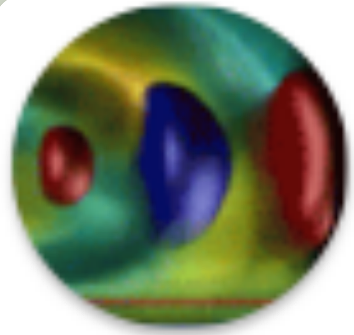

All-to-All Simulations

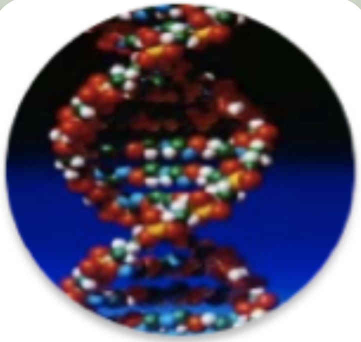

Random access, large data Analysis

… often they fit in multiple categories

The Programming Answer is Obvious…

More Regular

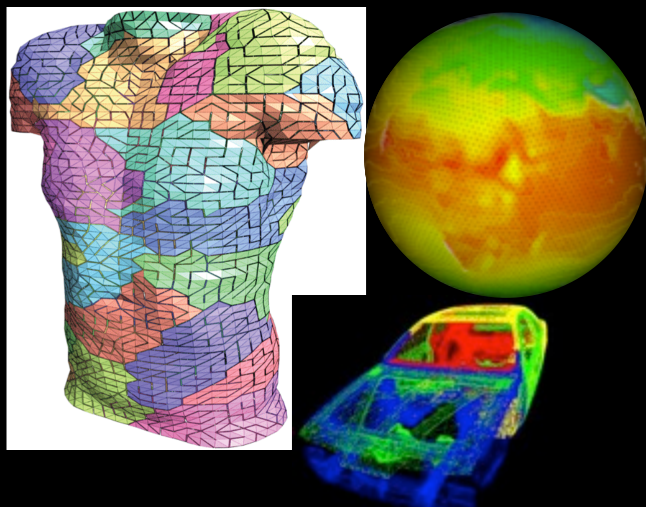

<u>Message Passing Programming</u>

Divide up domain in pieces

Compute one piece

Send/Receive data from others

MPI, and many libraries

More Irregular

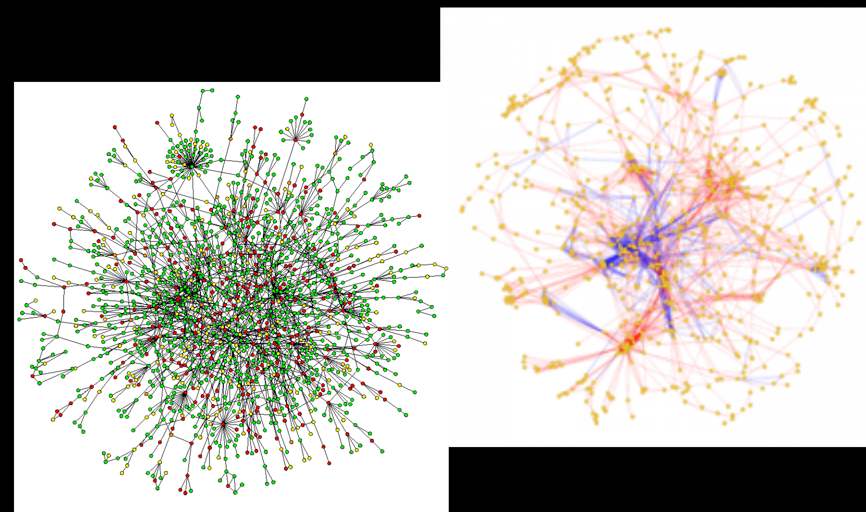

<u>Global Address Space Programming</u>

Each start computing

Grab whatever / whenever

UPC, CAF, X10, Chapel, GlobalArrays

Programming Models for Analytics

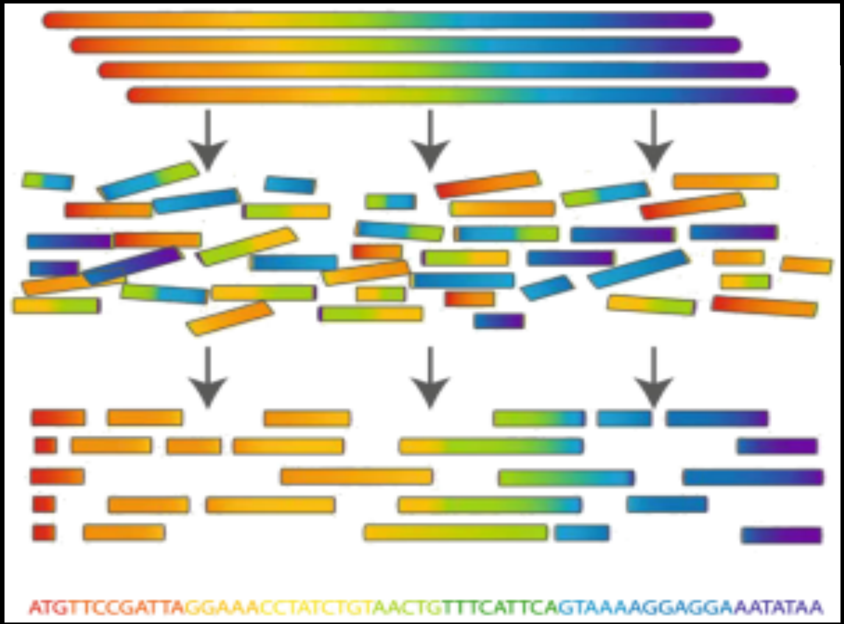

Strong Scaling of Meraculous Assember in UPC

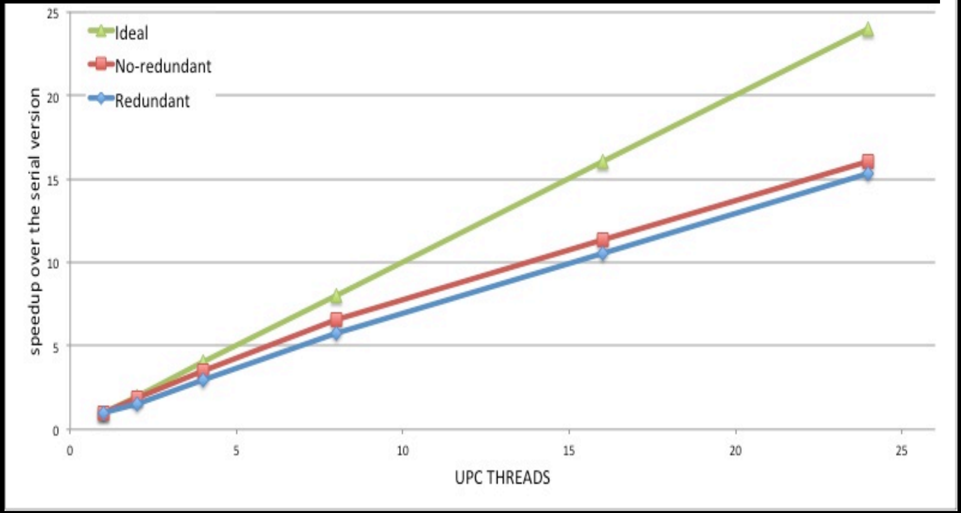

- **Computational Biologists buy large shared memory machines to assemble genomes**
- **For many problems (including metagenomics) these are not large enough**

Work by Evangelos Georganas, Jarrod Chapmanz, Khaled Ibrahim, Daniel Rokhsar, Leonid Oliker, and Katherine Yelick

End-to-end Computing for Science (Beam to bench)

Detector ⟷ **Network** ⟷ **Computer**

XDSF

Detector

- LCLS-II, ATLAS, Planck, K2, TEAM, PTF, FIB/SEM
- Synchrotrons & FELs
- Imaging Mass Spec
- Cryo EM
- Light and EM scopes

Network

- ESnet knows networks
- Built for science data
- Science DMZ, DTNs
- Making remote data local
- *"Insight before a TB"*
- *SDN 1Tb/s in 2016?*

Computer

- Petascale+ data analysis
- HTC & Realtime
- Science Gateways
- Community Databases
- *NERSC8+, Exascale*

Innovation from Data-driven Simulation Science

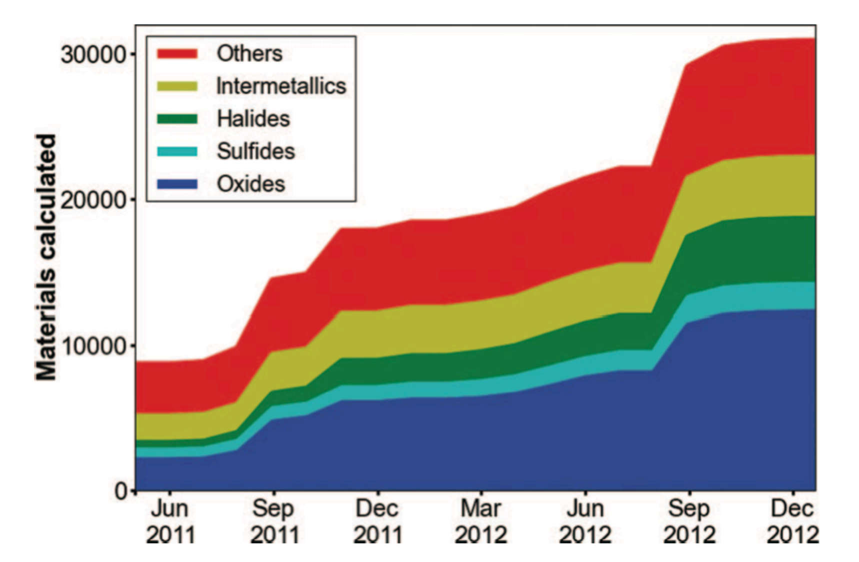

Reboot materials science as a collaborative HPC workflow, web based, durable data assets

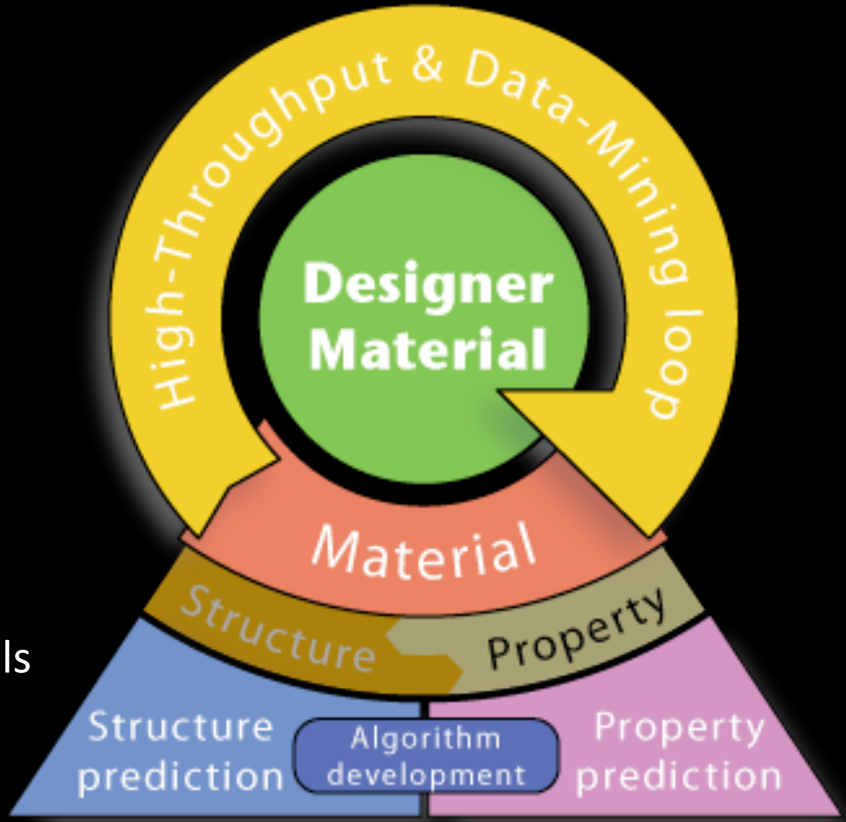

Ceder and Persson

- BES sponsored NERSC HPC+HTC resources
  - 15M hours in 2012, 40M in 2013
  - 10X needed for nano-synthesis, MOFs
- Topically-focused allocations for national MGI program discovery challenges / target materials
- Advanced I/O for data analytics
- Data science focused HPC R&D

NERSC path to more data, science, and computing…

1. **Meet the ever growing computing and data needs of our users by providing usable exascale computing and storage systems, transitioning SC codes to execute effectively on manycore architectures, and influencing the computer industry to ensure that future systems meet needs of SC**

Usable Exascale Initiative

2. **Increase the productivity, usability and impact of SC's data intensive science by providing comprehensive data systems and services to store, analyze, manage and share data.**

Extreme Data Analysis

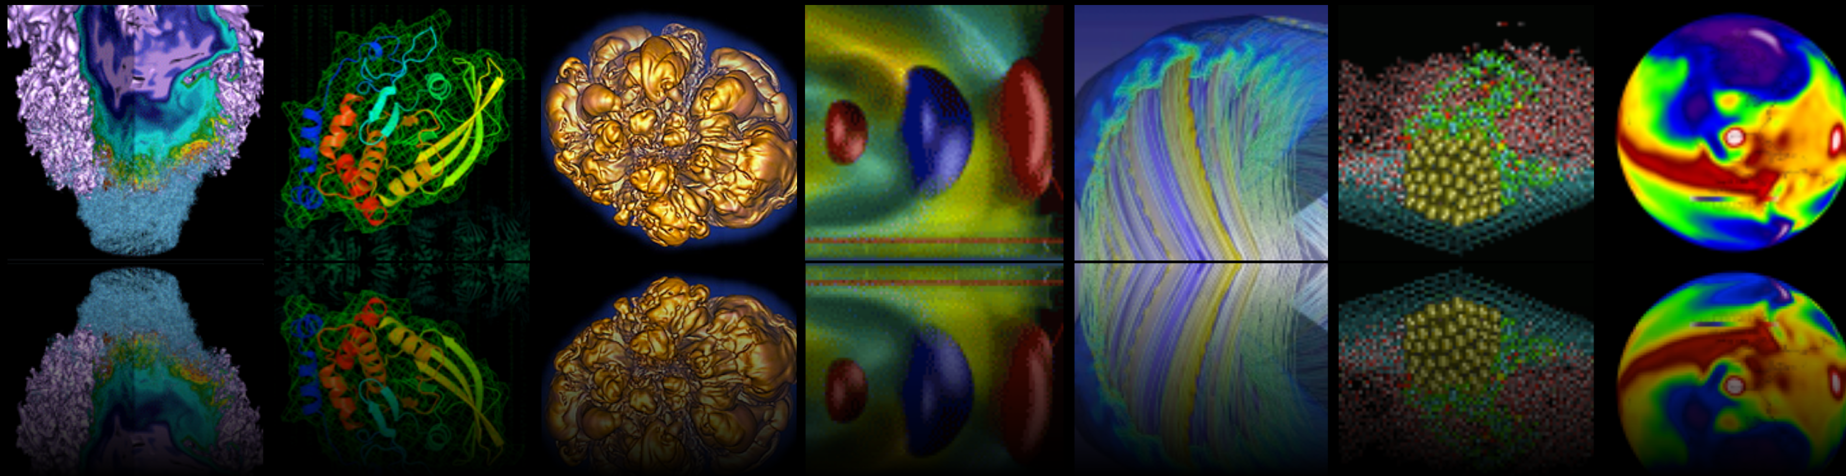

Happy 40th NERSC!

Kathy Yelick
Associate Laboratory Director for Computing Sciences

David Skinner
NERSC Strategic Partnerships Lead

Lawrence Berkeley National Laboratory