Convergence:
What it Means to Me

Converging View of Modeling, Simulation, Experiment, Data and Informatics

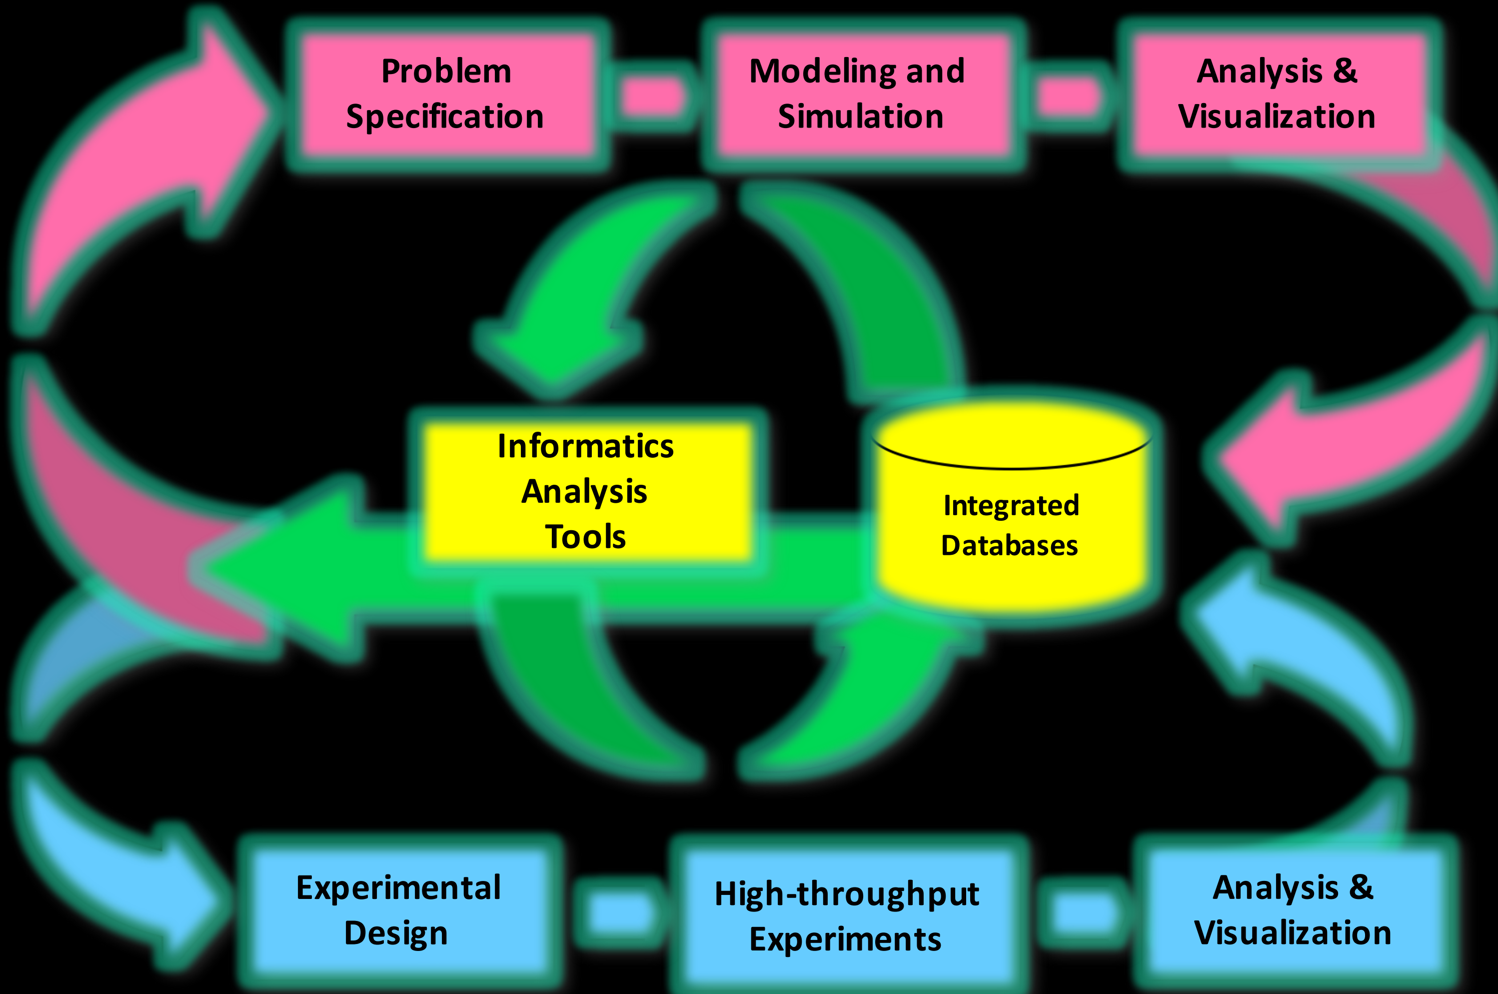

Integration of Simulation, Data Analytics and Machine Learning

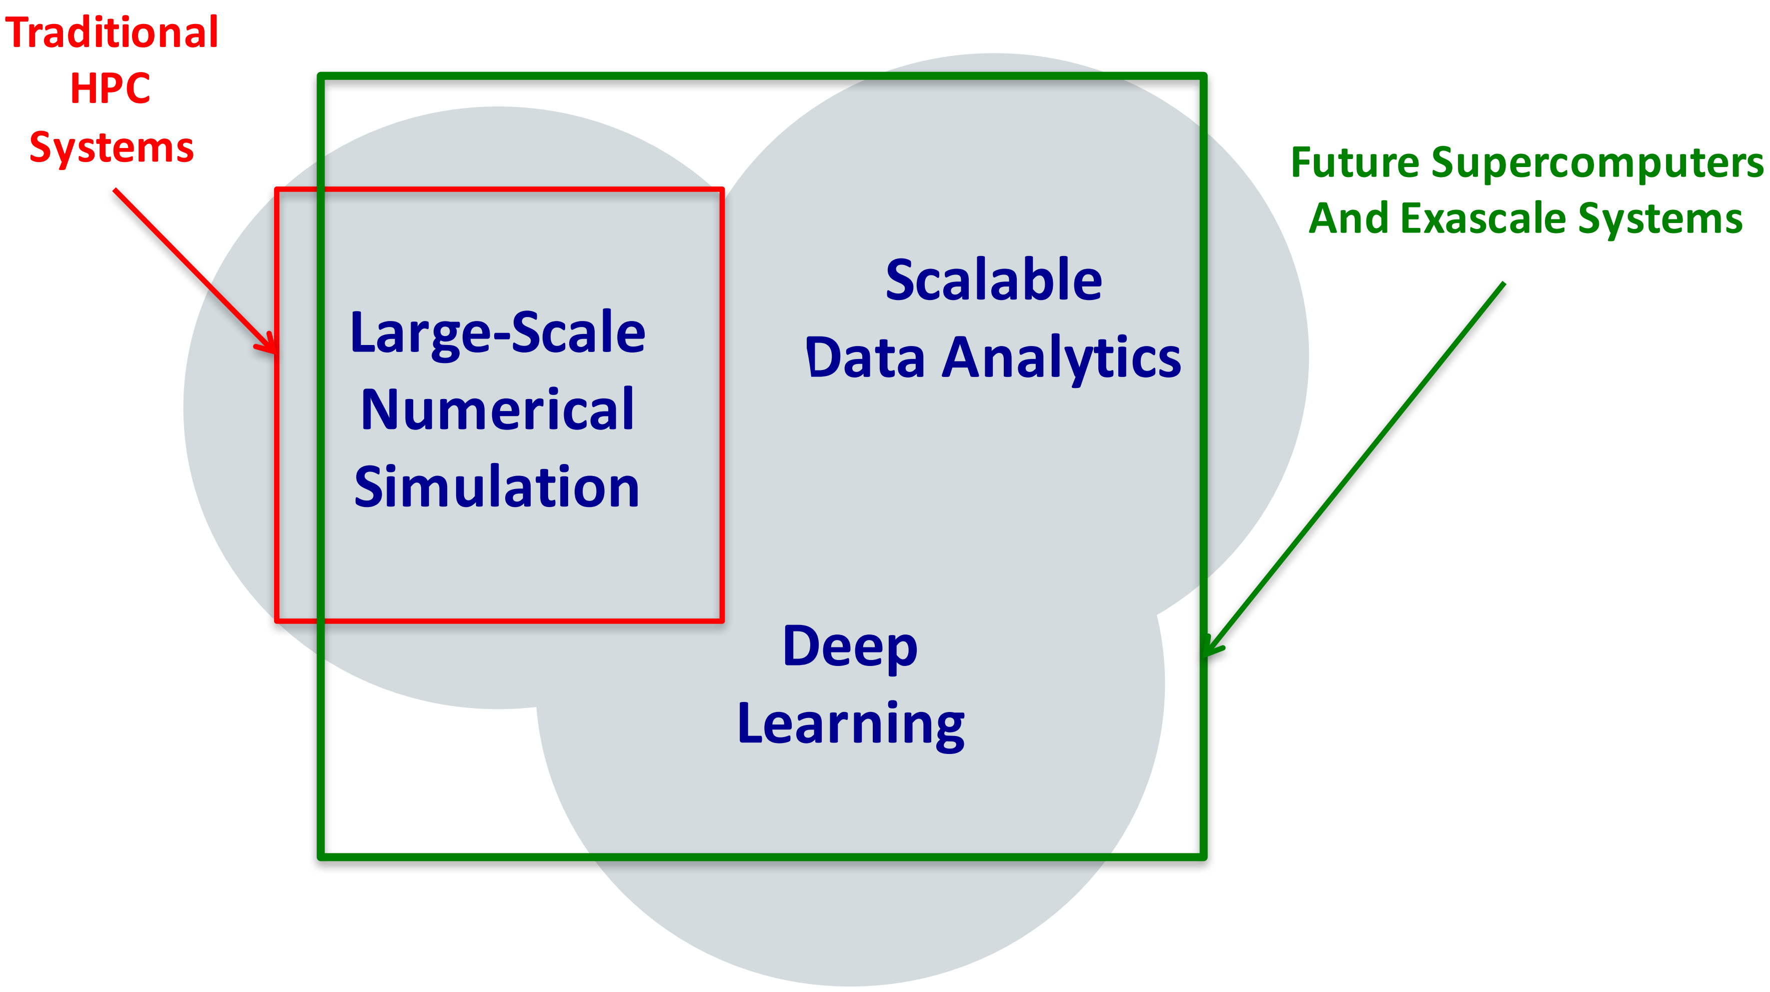

WE ESTIMATE BY 2022 ONE THIRD OF THE SUPERCOMPUTING JOBS ON OUR MACHINES WILL BE MACHINE LEARNING APPLICATIONS

SHOULD WE CONSIDER ARCHITECTURES THAT ARE OPTIMIZED FOR THIS TYPE OF WORK?

CAN WE LEVERAGE EXASCALE?

Machine Learning in Computational Science

Many fields are beginning to adopt machine learning to augment modeling and simulation methods

- Climate
- Biology
- Drug Design
- Epidemology
- Materials
- Cosmology
- High-Energy Physics

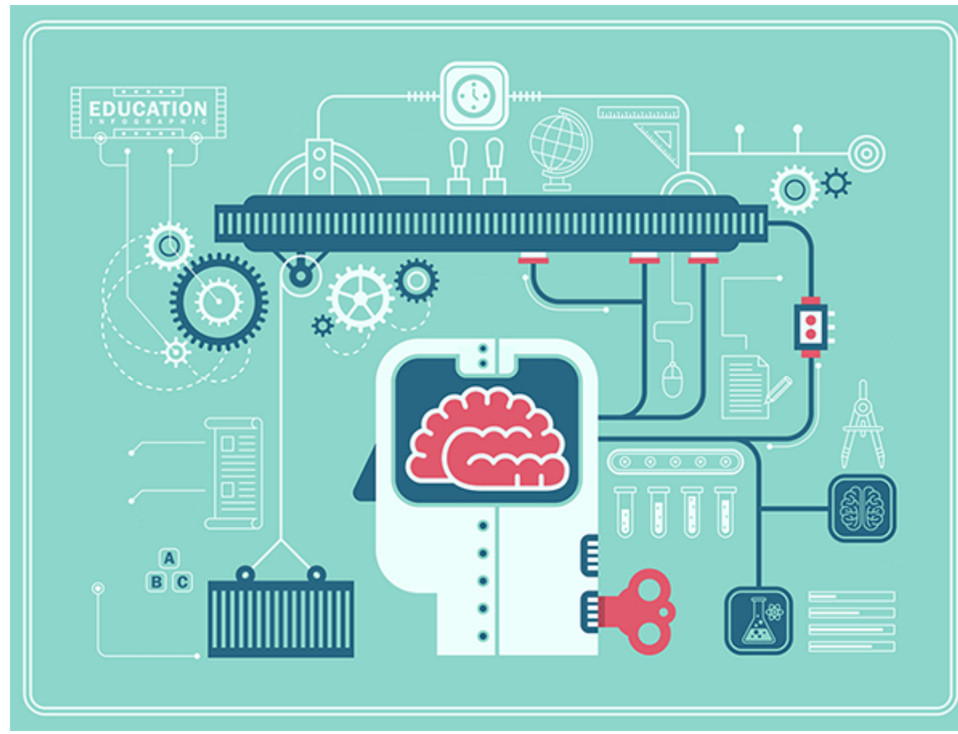

"Machine Learning: The New Infrastructure for Everything"

Applications Drivers

- We are seeing more applications that want to combine large-scale simulation with some form of large-scale data analysis but where the components are (mostly) separate codes and stacks
- A few examples we are tracking
  - Cosmology
  - Materials Science
  - High-Energy Physics
  - APS Imaging

Cosmology

"Cosmology is the study of the universe as a dynamical system"

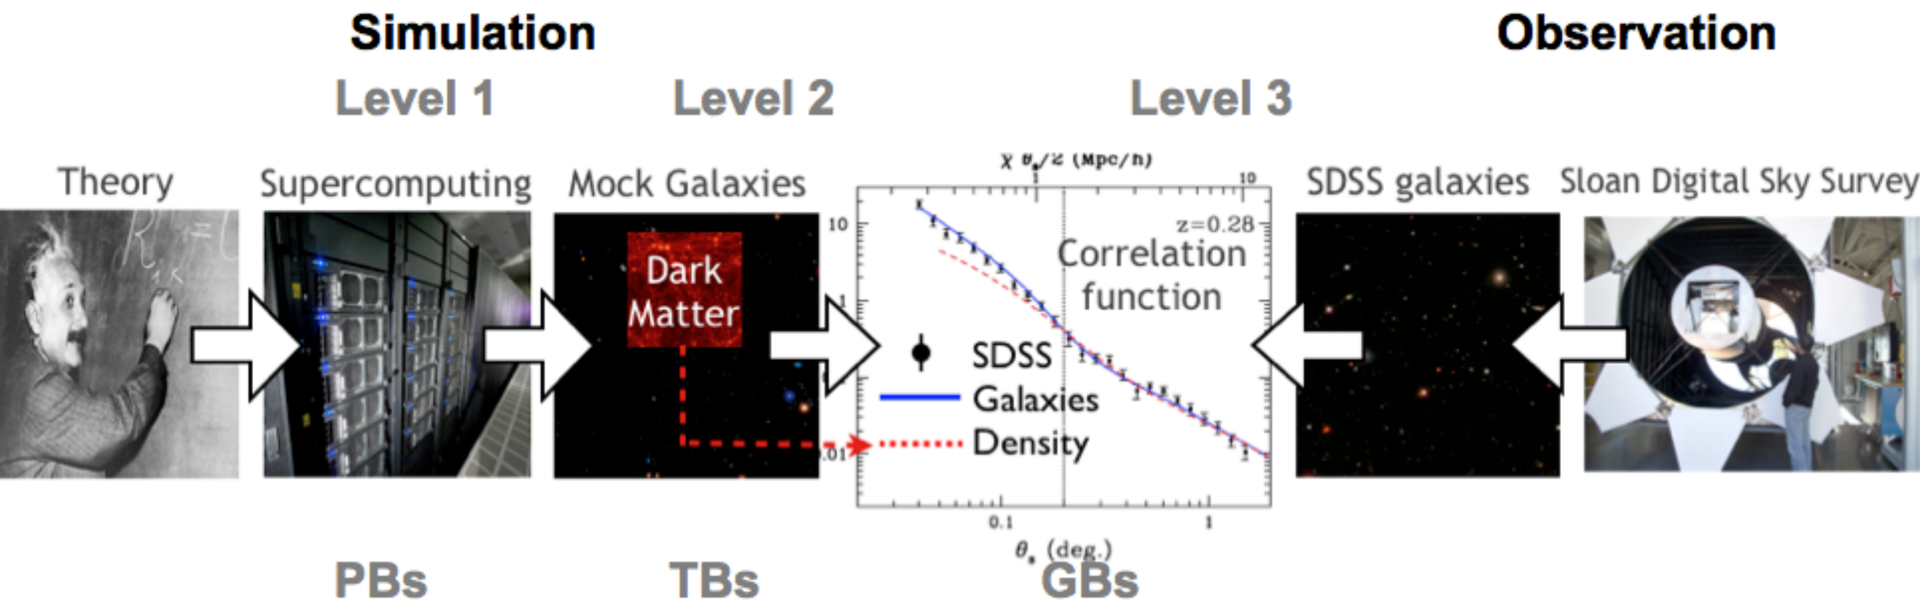

Three roles of cosmological simulations:

- Basic theory of cosmological probes
- Production of high-fidelity 'mock-skys' for end-to-end tests of the observation/analysis chain
- Essential component of analysis toolkits for scientific inference

Materials science example: Diffuse scattering

"Most of materials science is bottlenecked by disordered structures"

Simulation **Experiment**

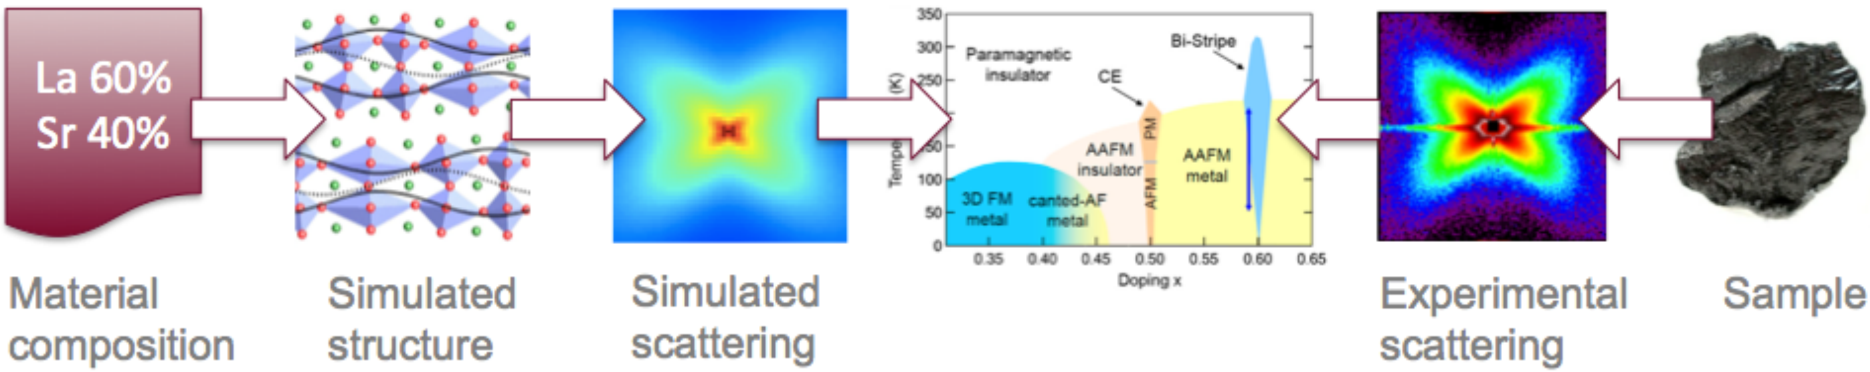

Use experiments to constrain models of material structure, and vice versa

- Experiments: Single crystal diffuse scattering of, e.g., bilayer manganites, yielding pair distribution functions
- Simulations: Molecular dynamics for candidate structures, yielding simulated scattering and simulated pair distribution functions

Applications Drivers

- We have applications that are persistent with deep web services that have a need for launching large-scale computing “backend” computations.. “Science Gateway” type problems (cloud hosted front end/shallow services)
- Some examples include
  - Systems Biology Knowledge Base
  - Materials “genome”
  - Engine design tools

Applications Drivers

- We have applications being developed where there is a coupling between machine learning and simulation
- Examples include
  - Cancer
  - Cosmology
  - Earth Science
- If we add in sensors as well then we have more examples such as urban science that need data analysis, sensor processing, simulation and machine learning

Pilot 1: Predictive Models for Pre-Clinical Screening

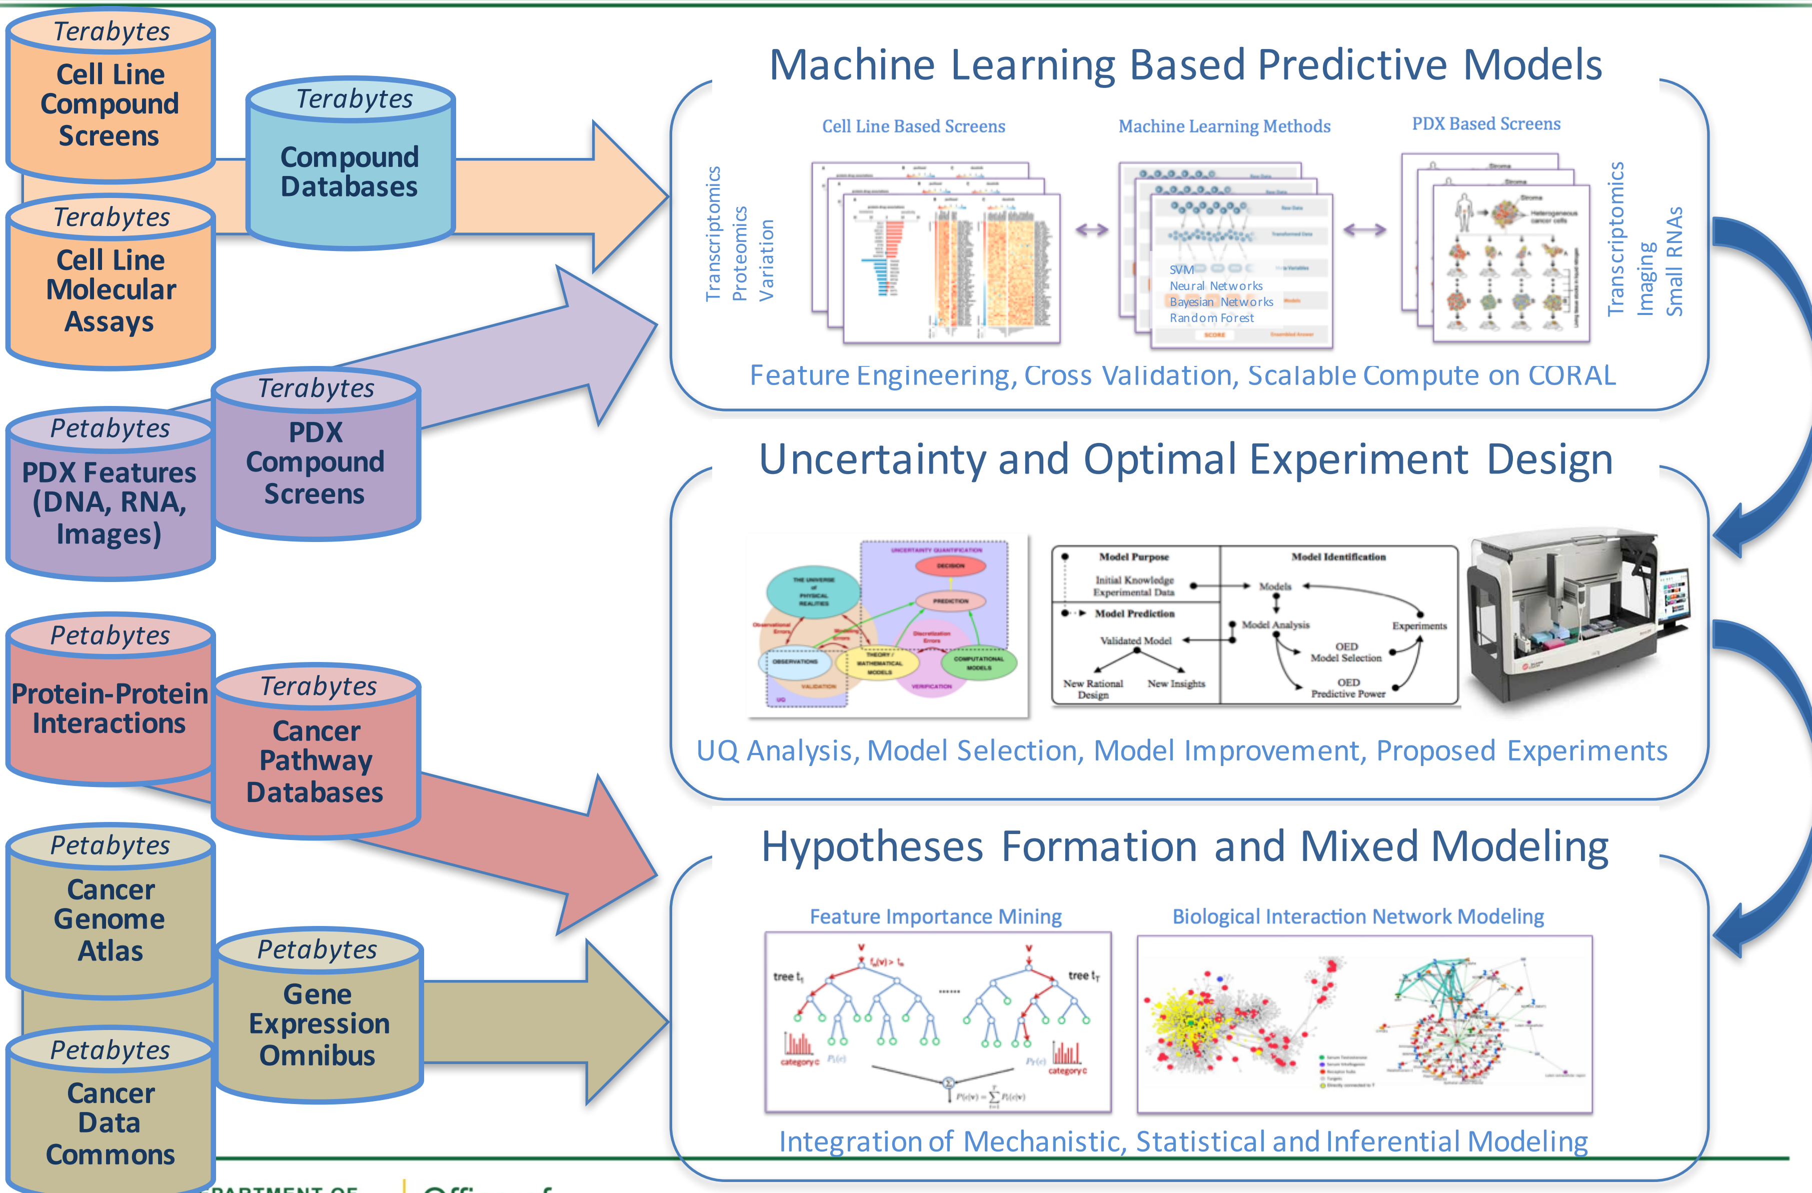

Pilot 2: RAS proteins in membranes

RAS activation experiments at NCI/FNL

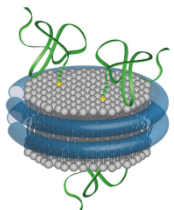

Experiments on nanodisc

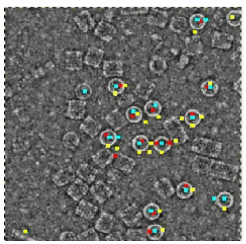

CryoEM imaging

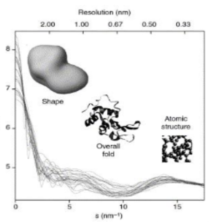

X-ray/neutron scattering

Multi-modal experimental data, image reconstruction, analytics

Protein structure databases

New adaptive sampling molecular dynamics simulation codes

Adaptive time stepping

Coarse-grain MD ⇄ Classical MD ⇄ Quantum MD

Adaptive spatial resolution

High-fidelity subgrid modeling

Predictive simulation and analysis of RAS activation

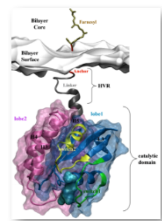

Granular RAS membrane interaction simulations

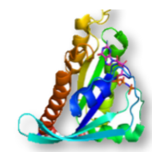

Atomic resolution sim of RAS-RAF interaction

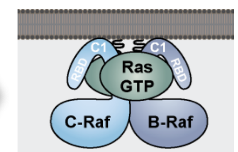

Inhibitor target discovery

Machine learning guided dynamic validation

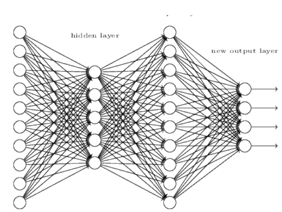

Unsupervised deep feature learning

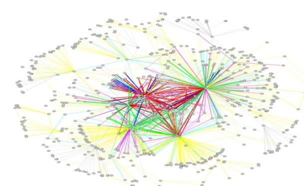

Mechanistic network models

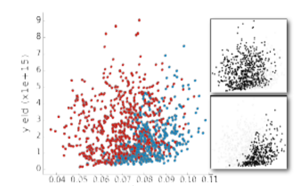

Uncertainty quantification

Pilot 3: Population Information Integration, Analysis and Modeling

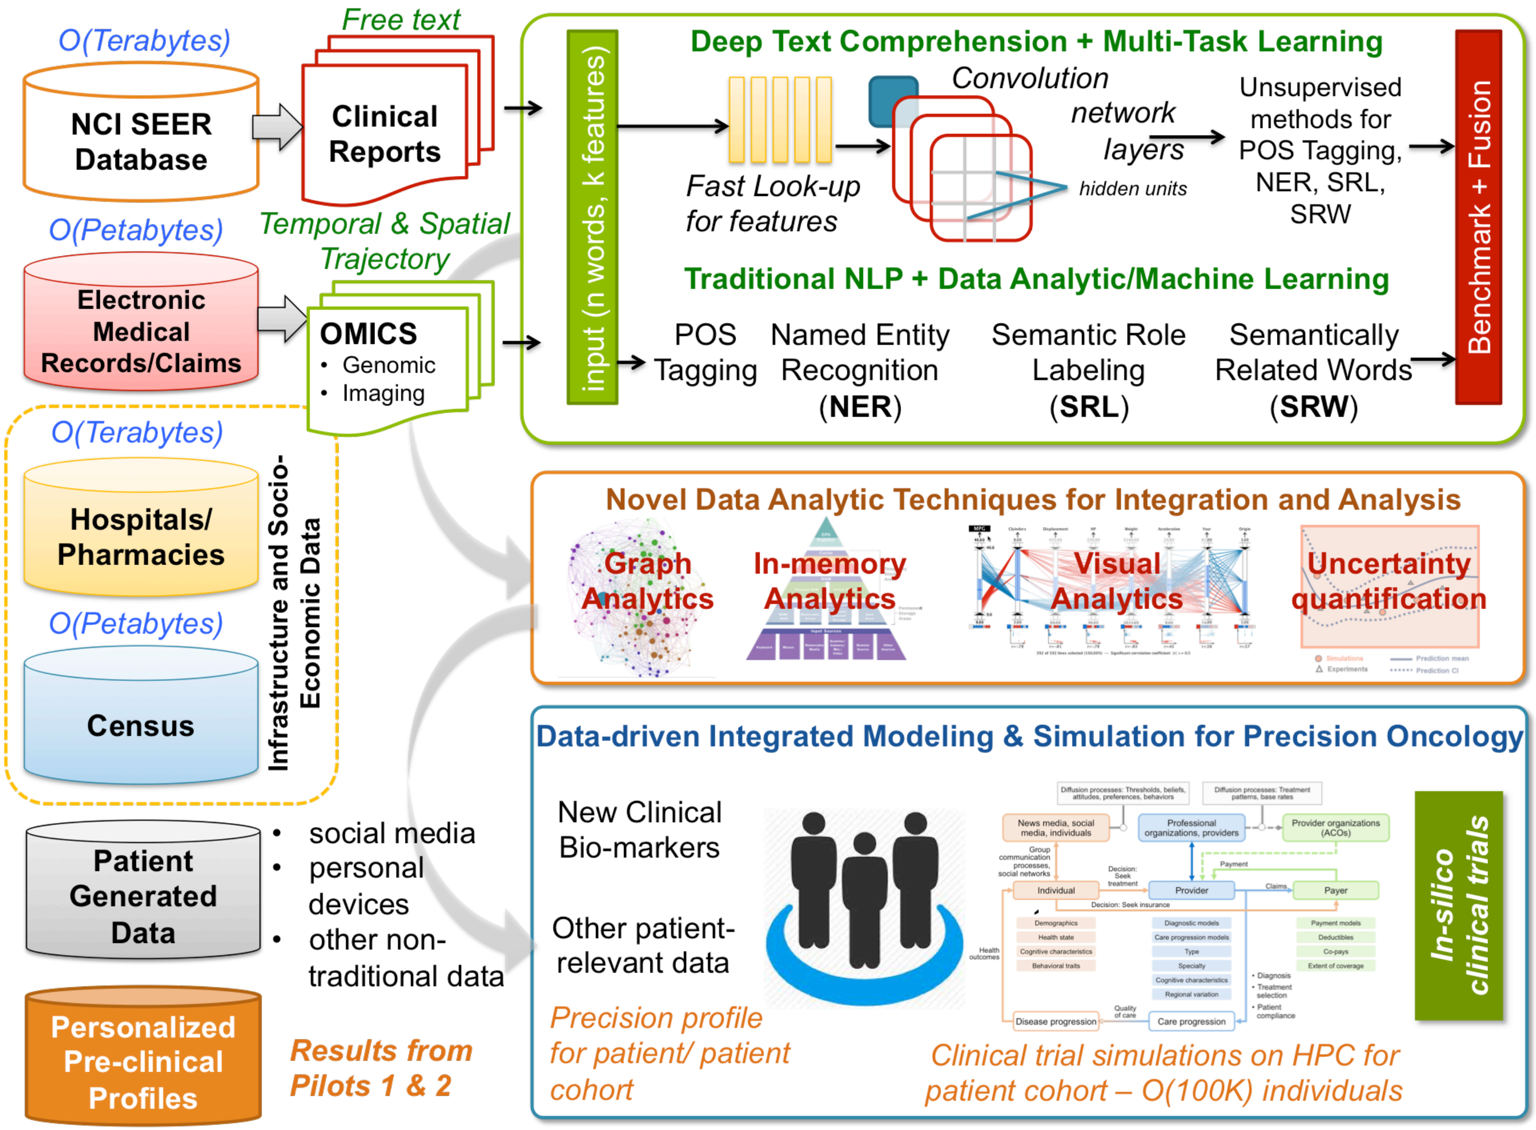

Convergence vs Co-Location

- This seems to be the fundamental question we are dealing with
- Convergence implies some need for tight funcational coupling, wereas co-location may be sufficient for loose coupling via workflows
- DOE labs have been exploring these needs a bit and had some facilities concepts for co-location (VDF concept)

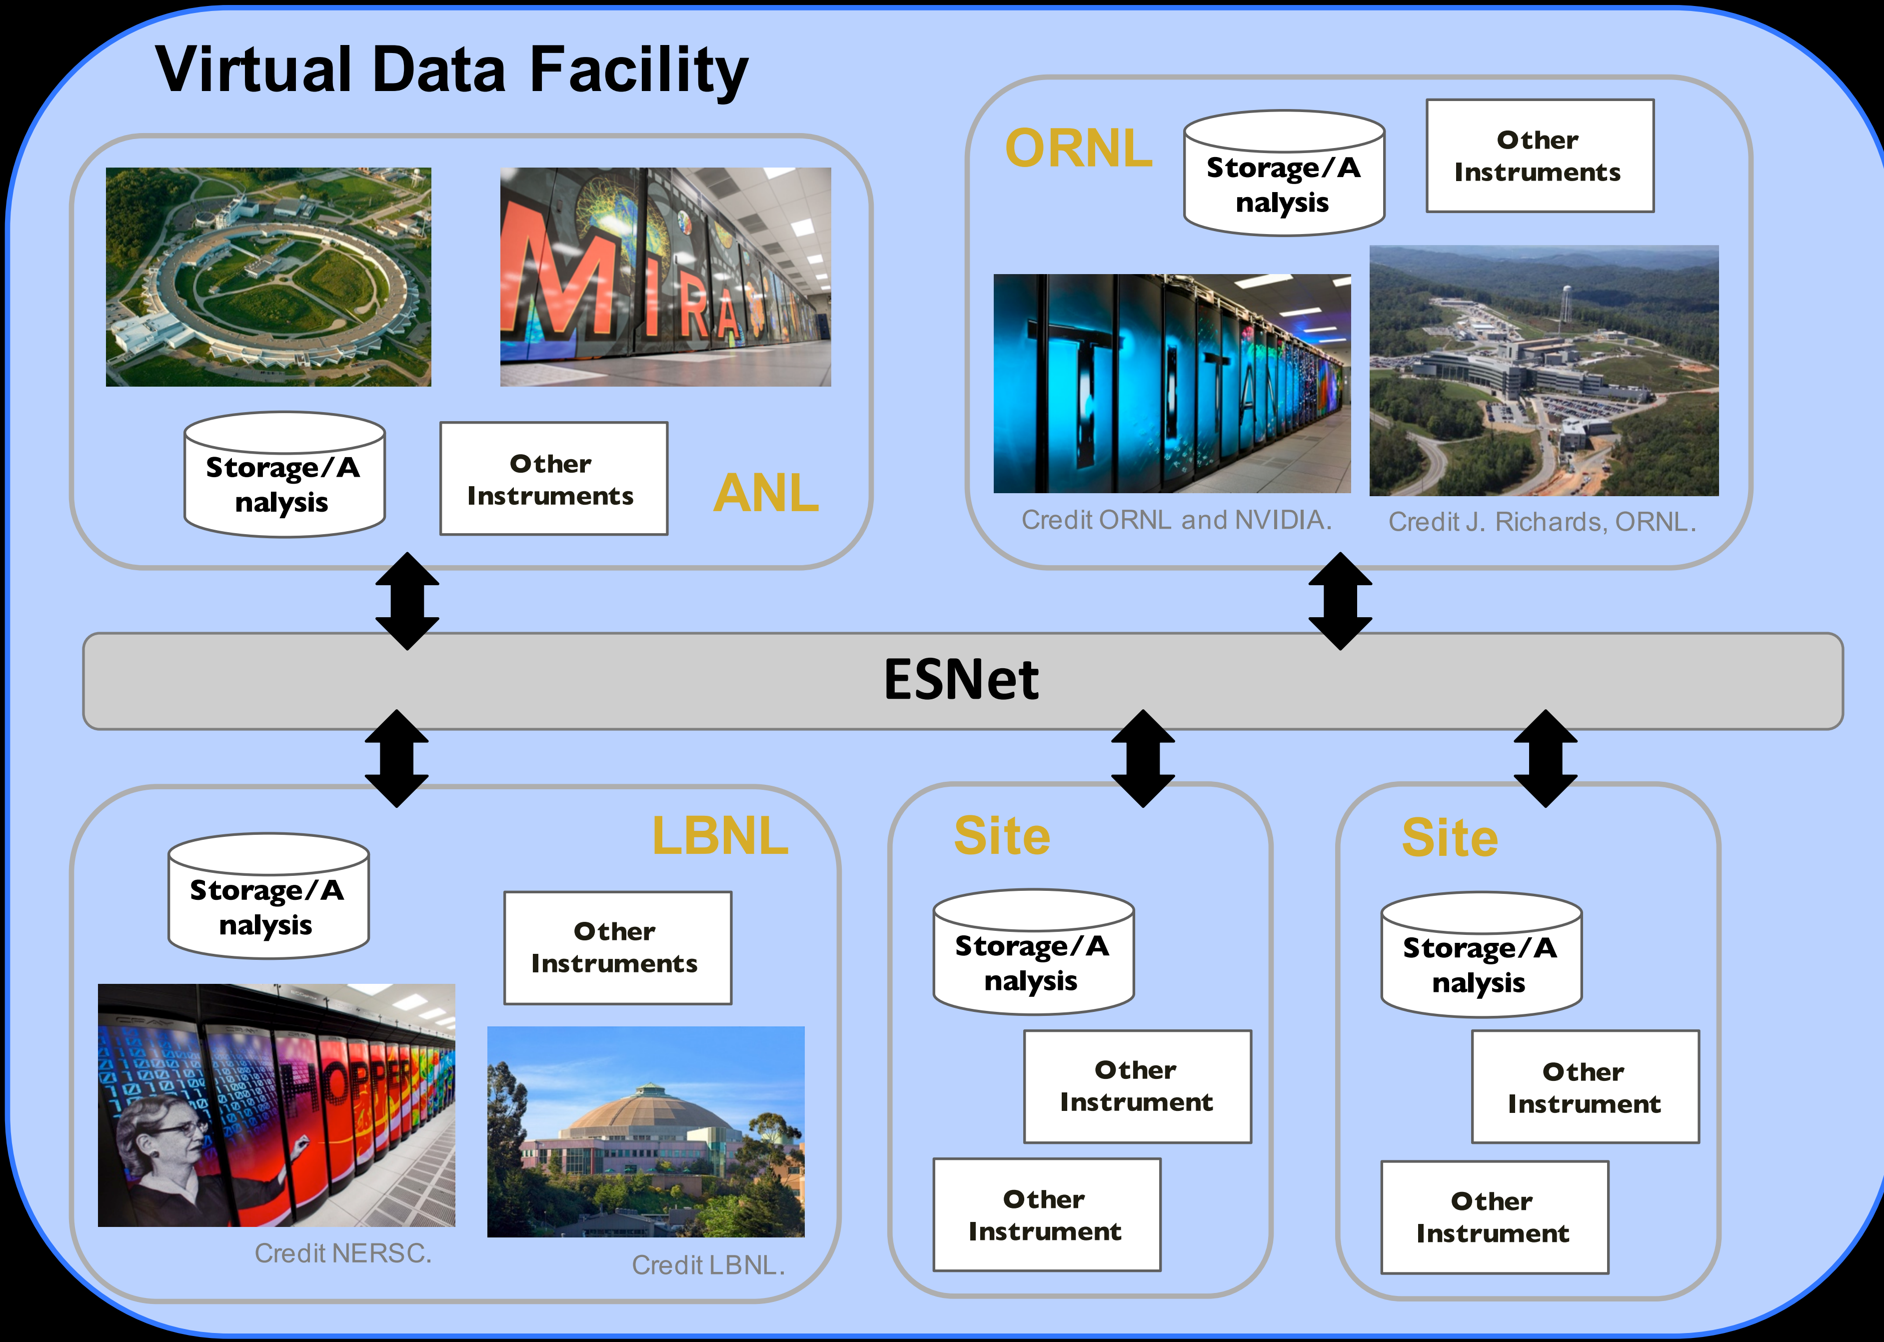

Virtual Data Facility
ANL
Storage/A nalysis
Other Instruments
ORNL
Storage/A nalysis
Other Instruments
Credit ORNL and NVIDIA.
Credit J. Richards, ORNL.
ESNet
LBNL
Storage/A nalysis
Other Instruments
Credit NERSC.
Credit LBNL.
Site
Storage/A nalysis
Other Instrument
Other Instrument
Site
Storage/A nalysis
Other Instrument
Other Instrument

Extrapolating: One Possible View of the Future

Individual University and Lab PIs

National and Int'l collabs

Research + Industry (e.g., Hubs)

Industry

Tools for Portals, Search, Assimilation, Provenance

Exascale Simulation Facilities

Massive Throughput Simulations

Data Serving and Archiving Facilities

Analytics & visualization systems

ESnet

On-Site Computing or data transfer nodes

Light sources, etc. (BES)

Colliders (HEP/NP)

Sequen-cers (BER)

Cosmology (HEP)

"Omics" data (All)

Model for Virtual Data Facility

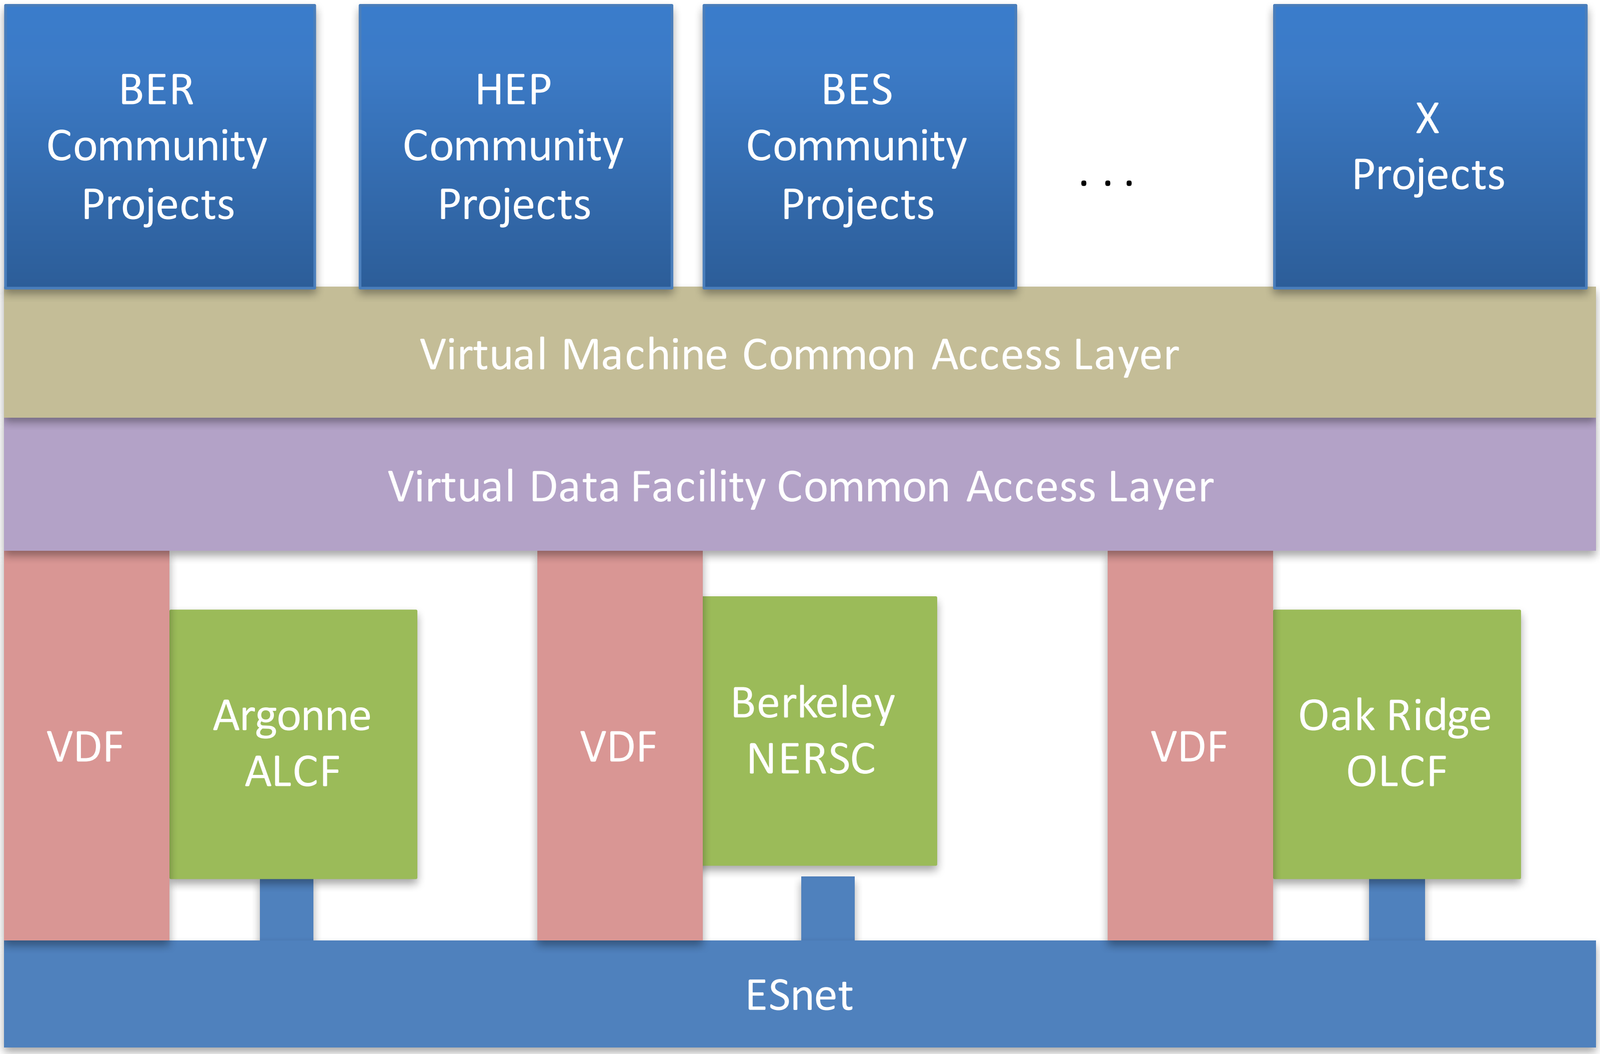

Co-location

- We are all doing that now to some degree
  - We have I/O farms, and visualization and analysis servers associated with our big machines
  - Often the associated servers are significant scale and have more relaxed access policies
- We also have “condominium” deals were community specific resources are virtually integrated into some infrastructure and shared (trading space for time)

Co-location

- Today co-location strategy might imply bringing together machine configured in different ways on a high-bandwidth machine room or virtual machine room network, each with diferent stacks
- Examples
  - Analytics Servers with apache SPARK
  - Visualization or GPU farms
  - Graph processing accelerators
  - Private clouds (hosting diverse VMs and applications)
  - Large-scale supercomputers
  - Large-memory servers
  - Database servers

BD Usage Models Differ from EC

Big Data

- Continuous access require based on data generation/submission rates
- CPU time, I/O and data volume all important
- Data products typically used in future computations via an integration or pipeline
- Data products made available for external users and curated over time

Extreme Compute

- Batch oriented access based on allocations for specific projects
- Mostly CPU time centric
- Output not necessarily used in future runs but often significant time used for visualization
- Output generally (but not always) used "privately" and rarely curated

Policies Need to be Different

- Long term (many years) access commitment at a continuous or increasing level of service
- Support for persistent services
- Storage allocation that grows over time
- Rich software environment with high-performance database support
- Mechanism to publish the data to a community
- Archival support for data, links and citations

Co-location to Convergence?

- From a pure operational efficiency stand point reducing the number of types of machines, OS stacks, vendors and diversity of configurations is desirable
- Many of these diverse servers could be hosted under a managed environment such as OpenStack and some could be converted to “self management” by the users

Co-Location to Convergence

- Our experience has been mixed.
- Private Clouds can be very inefficient without a feedback mechanism (people launch things and forget about them)
- Data replication is a big problem
- Management complexity is a problem
- User navigation of the resources is a problem
- Ad hoc collections of resources co-located can be fragile

What Have We Learned?

- Leveraging investments is probably good
  - Fabric, Data Storage, Power Infrastructure, Expertise, Documentation, Learning Curves
- Clouds (IaaS) is a mixed bag
  - Users can get going quickly, but also run into all the problems that need real teams and infrastructure to solve eventually
  - The speed at which things can be done however is attractive and does matter at some level
  - Sustainability is the balancing factor
- High-Utilization of a resource might not matter if the marginal cost of that resource is low
  - Just avoid it as a reporting metric

What have we learned?

- Convergence in other areas seems to work
  - SmartPhone is a convergence of many things
    - Phone, internet, camera, GPS, etc.
- There is a mutliplicative effect of possibilties when previously disperate resources are available in the same framework
  - This appears to be happening with coupling simulation with machine learning
  - Coupling graph engines with machine learning
- We fully agree with the examples from yesterday regarding things like AlphaGO
- Recent paper on learning gene regulatory network in planarium is annother good example

Observation

- We are expecting in our next generation machines to be able to do the following on the same hardware platform
  - Run a containerized data analytics stack
  - Run scalable deep learning tools with HW acceleration
  - Run user contribued containers
  - Run tightly coupled “Traditional” HPC applications
  - Run persistent data services
  - Run heterogeneous workflow as a single “job”

Convergence

- Ideal Environment
  - Interactive parallel prototyping environment
  - Seamless scale up to production (10^3x-10^6x)
  - Integrated platform for analysis and simulation
  - Same platform for publishing
  - Persistent data regions in memory
  - Programming language support for data analysis
  - Large-scale interactive computing
  - Seamless visualization and sharing

Recommendations

- Our community push forward on understanding and demostrating convergence of software frameworks and applications environments with the assumption that market forces will resolve the details of hardware convergence opportunities
- Choose a small number of next generation driver applications that might be able to leverage two or three “covergence” features and get those to work to learn how to do this

Automate and Accelerate

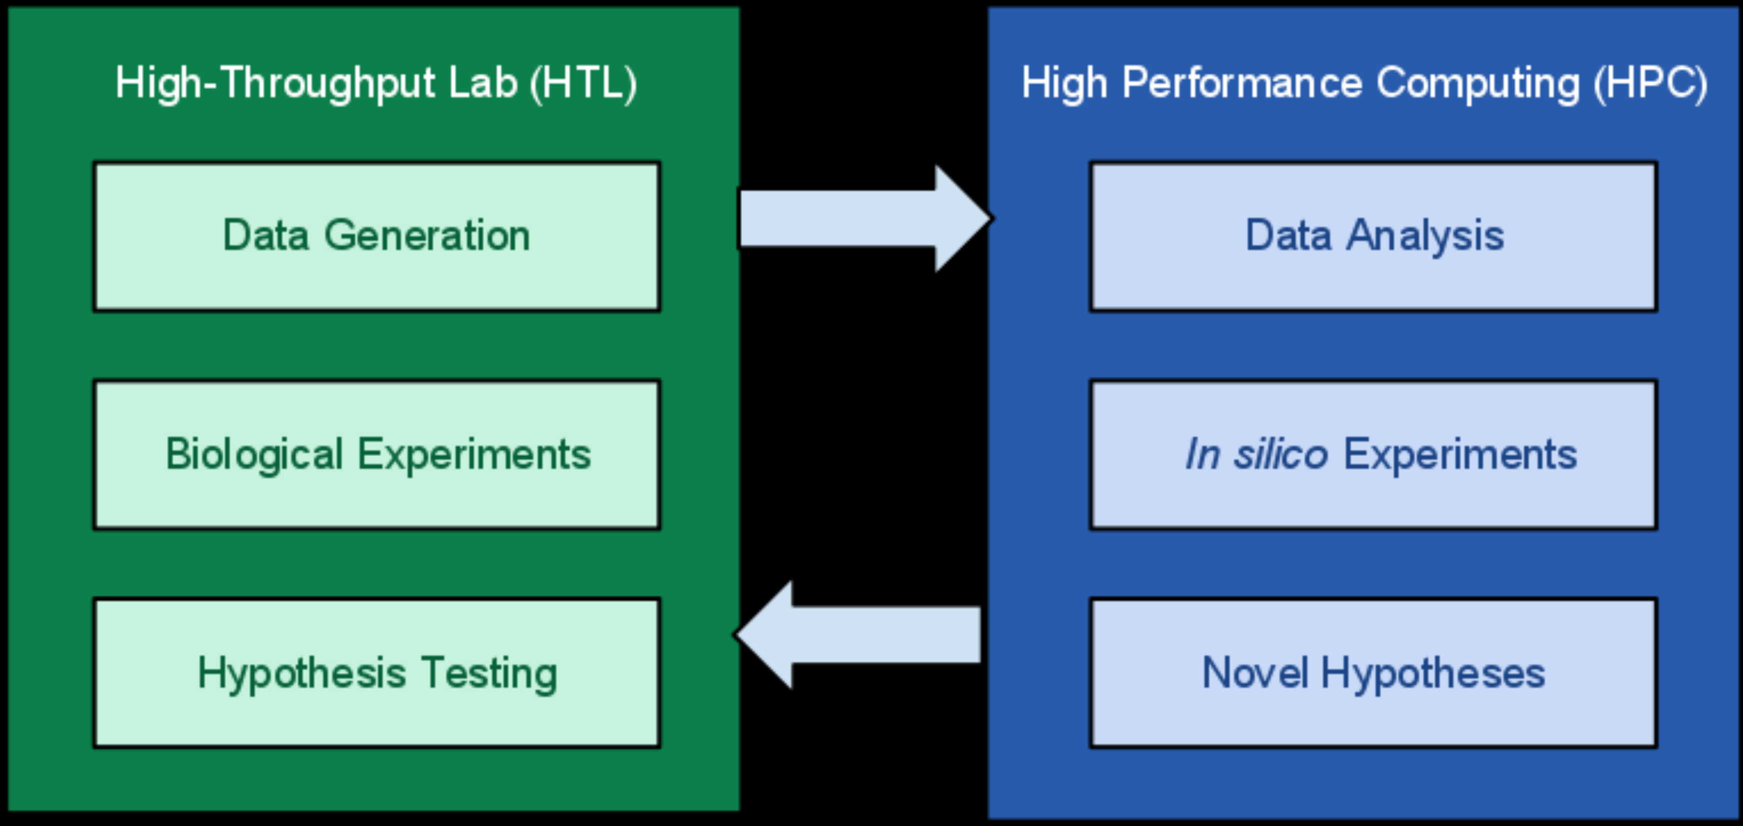

Ross King and his “science robot”