U.S. DEPARTMENT OF ENERGY | Office of Science

Future Lab Computing Working Group (FLC-WG)

Status Report

ASCAC Meeting 4/18-19/18

Co-chair: Mallikarjun Shakar (ORNL)

Co-chair: Eric Lancon (BNL)

ASCR PM: Richard Carlson

Executive Summary

- **DOE/SC Laboratories provide computing/storage resources to lab staff, researchers, and visiting scientists**
- **Demands on these resources are increasing**
- **Labs have the capability to leverage decades of research to create modern Distributed Computing and Data Ecosystems (DCDE) to meet the current and future demands of DOE scientists**

Spoilers

Emerging Trends in Lab Computing

- **DOE/SC Lab complex contains 10 labs each with a large collection of compute and storage resources scattered across their campus.**
- **The demands on each lab are the same**
  - Upgrades to SC facilities is rapidly increasing the amount of Experimental and Observational data that needs to be analyzed
  - SC facility user communities are becoming more diverse, may have less experience with HPDC environments, and are requesting more data analysis capabilities from the labs
  - Converged Computing and Data Services are in greater demand to support simulation and visualizations
  - Labs have difficulty meeting these growing demands while staying within existing budget constraints

Research Activities in Distributed Computing

- **ASCR has supported research projects to explore Grid and Cloud based computing for almost 2 decades**
  - Fusion Grid (FG), Earth Systems Grid Federation (ESGF), Open Science Grid (OSG)
- **These research projects have produced and contributed to broadly used tools and services**
  - Networking and data transfer: GridFTP, Globus, etc.
  - Access control and federated access: CI-Login, Science DMZ, etc.
  - Middleware for workflows: Pegasus, Swift, etc.
- **Several DOE science communities have extensive experience in operating distributed computing environments**
  - Climate modeling distribution (ESGF) LHC Atlas and CMS (OSG), Remote control room operations (FG)

Deployment for the Future

- **The Future Lab Computing Working Group (FLC-WG) was formed to:**
  - Encourage the Labs can take advantage of the decades of research and operational experience in Distributed Computing
  - Assist the Labs in developing a common set of policies and practices to ensure smooth operations across the DOE/SC complex
  - Assist the Labs in understanding the current state of the art in Distributed Computing and Data Ecosystems (DCDE)
  - Assist the Labs in developing a strategy for adopting a Federated DCDE model to meet the scientific computing and data needs of their current and future scientists, researchers, and lab staff

National Laboratory Research Computing Group (NLRCG)

Shortly after ASCR started the FLC-WG, the NLRCG started a similar activity. These activities were merged to better meet the needs of all interested parties. Four thrusts were defined that required dedicated working groups:

- Economies of Scale
- Blockers to Collaboration
- Scientific Data Management
- Future Laboratory Computing – DOE SC Interest

FLC Working Group Chartered

- **The National Laboratory Research Computing Group (NLRCG) and the Advanced Scientific Computing Research (ASCR) program office jointly establish the Future Lab Computing Working Group (FLC-WG) thrust to identify best practices and research challenges leading to the creation and operation of a DOE/SC wide federated Distributed Computing and Data Ecosystem (DCDE). The FLC-WG will accomplish this by:**
  - Reviewing past ASCR research activities
  - Reviewing current lab, scientific community, academic, and commercial solutions
  - Synthesizing past and current technical activities into a set of best practice documents
  - Identify current and future challenges that need to be addressed
  - Determine the activities needed to pilot an operational deployment of a DCDE

FLC-WG Goals

- **The outcome of the FLC-WG will be a report describing the current state of lab computing, a set of research challenges, and plans for achieving a federated DCDE.**

- **The FLC-WG will hold regular meetings and may bring in outside experts to obtain the information needed to complete its tasks.**

- **The NLRCG FLC-WG thrust will coordinate its activities with other NLRCG thrust working groups to obtain relevant background information and to ensure SC wide integration of the practices and recommendations.**

FLC-WG Members

Name	Lab	Specialty
Kerstin Kleese Van Dam	BNL	Lab/Science
Panagiotis Spentzouris	FNAL	Science
Arjun Shankar	ORNL	Lab
Amber Boehnlein	JLAB	Lab/Science
Rob Roser	FNAL	Lab/Science
Ray Bair	ANL	Lab
Richard Mount	SLAC	Lab/Science

Name	Lab	Specialty
Theresa Windus	AMES	Science
Craig Tull	LBNL	Lab/Science
David Cowley	PNNL	Lab/Science
Eli Dart	ESnet	Facilities
Lauren Rotman	ESnet	Facilities
Michael Hofmockel	PNNL	Lab
Eric Lancon	BNL	Lab
Adam Stone	LBNL	Lab

FLC-WG Observers and DOE/SC PMs

Observers		
Brian Kritzstein	PNNL	Lab
Frank Alexander	BNL	Lab
Adeyemi Adesanya	SLAC	Lab

Federal Program Managers		
Richard Carlson	DOE/SC	ASCR
Lali Chatterjee	DOE/SC	HEP
Eliane Lessner	DOE/SC	BES
Ben Brown	DOE/SC	ASCR
Jay Hnilo	DOE/SC	BER
Paul Bayer	DOE/SC	BER

FLC-WG Activities

- **Regular A/V meetings every other Thursday**
  - Obtain Information on the current and future needs of representative science communities
  - Obtain information on current and future plans for SC facilities
  - Obtain information on previous and current Research projects
  - Obtain information on Lab directions in computing and storage
- **Face-to-Face meeting at SC-17**
  - Develop report outline and assign writing assignments

Science use Cases

NWChemEx Science Challenges

- **DOE 2014-2018 Strategic Plan**
  - Development of computational models that demonstrate that biomass can be a viable, sustainable feedstock for biofuels, hydrogen and other products
- **Two Inter-related Science Challenges**
  - Efficient Production of Biomass

    Development of predictive molecular models of ion transport processes across cellular membranes that control stress responses to aid in the development of stress-resistant crops (Q. Liu, BNL)
  - Efficient Conversion of Biomass to Biofuels

    Development of predictive molecular models of catalytic conversion of biomass-derived alcohols to biofuels to aid in the discovery of energy efficient conversion processes (P. Sushko, PNNL)

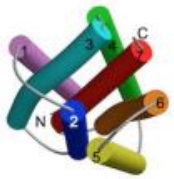
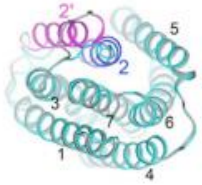
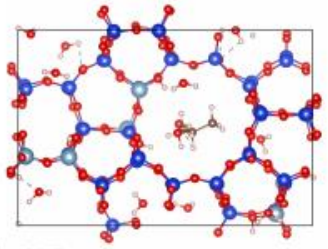

Example of LCLS Data Analytics: The Nanocrystallography Pipeline

SLAC

Serial Femtosecond Crystallography (SFX, or nanocrystallography): huge benefits to the study of **biological macromolecules**, including the availability of femtosecond time resolution and the avoidance of radiation damage under physiological conditions (**"diffraction-before-destruction"**)

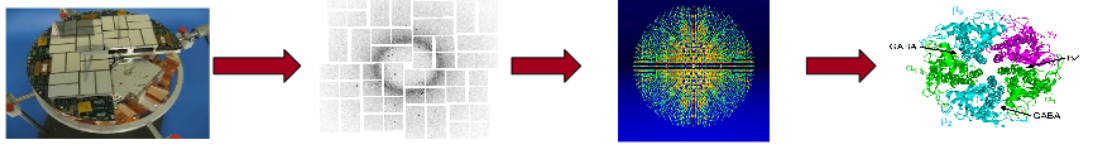

Megapixel detector → X-Ray Diffraction Image → Intensity map from multiple pulses → Electron density (3D) of the macromolecule

Well understood computing requirements

Significant fraction of LCLS experiments (~90%) use large area imaging detectors

Easy to scale: processing needs are linear with the number of frames

Must extrapolate from 120Hz (today) to 5-10 kHz (2022) to >50 kHz (2026)

But this won't be enough as the program evolves…

Better Accelerator and Detector performance increases discovery potential, but

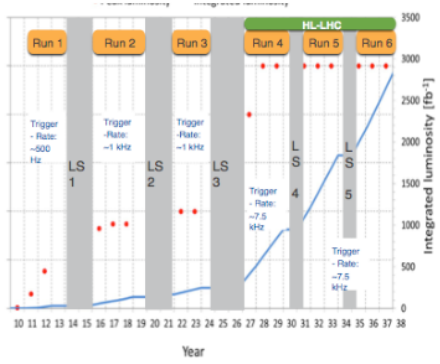

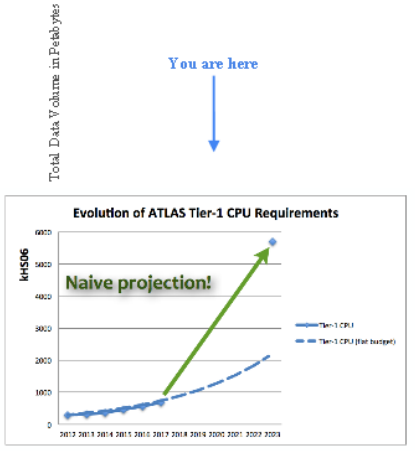

➡ computing will need to reach **10-100x current capacity!**

U.S. DEPARTMENT OF ENERGY | Office of Science

Fermilab

Science Use Case - Observations

- **Multiple domain science communities have growing data analytic needs**
- **Experimental and Observational science communities are diverse and widely distributed**
- **Scientific workflows systems are used to manage data collection, movement, analysis, and visualization tasks**
- **A new generation of DOE/SC domain scientists who are unwilling or unable to directly operate facility infrastructure are running experiments and using these facilities**

DOE/SC Facilities

It's a DOE user facility engineered and optimized for Big Data Science

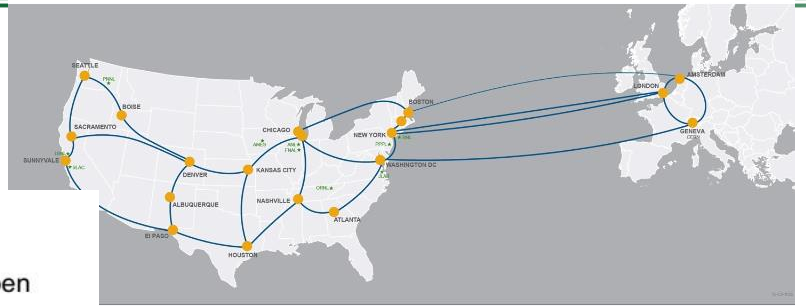

We do this by offering unique capabilities and optimizing the facility for data acquisition, data placement, data sharing, data mobility.

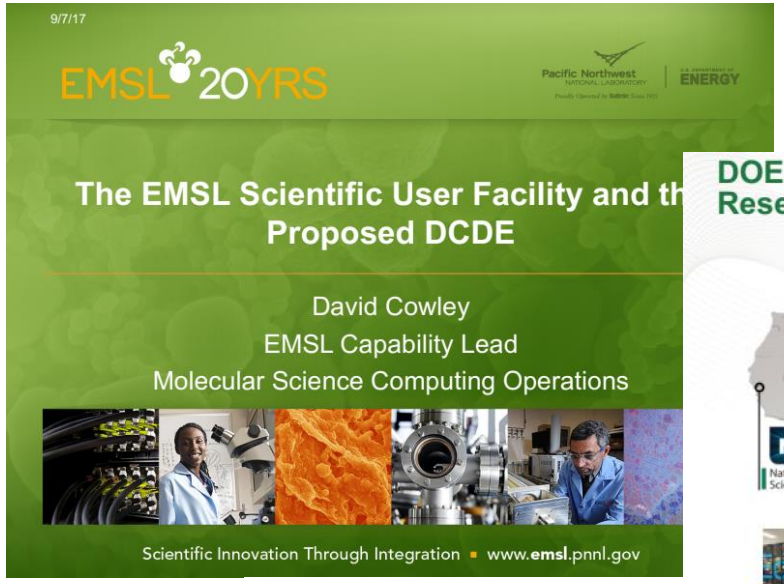

DOE's Office of Science, Advanced Scientific Computing Research (ASCR) Computation User Facilities

NERSC
Edison is 2.57 PF

ALCF
Mira is 10 PF

OLCF
Titan is 27 PF

- DOE is a leader in open High-Performance Computing
- Provide the world's most powerful computational tools for open science
- Access is free to researchers who publish
- Boost US competitiveness
- Attract the best and brightest researchers

CADES

Collaborates wi
research needs
a flexible enviro
and high perfor
data, expected
2020-2022)

ALICE

- CADES dedicated WLCG resources for ALICE:
- 60 Compute Nodes
- (1,500 cores/7TB ram 18 kSH06)

~2.5PB EOS Distributed Storage Environment

40Gb/s LHCOne HEP Data Network Uplink
- High performance, world wide connectivity to HEP centers

OAK RIDGE National Laboratory

UT-BATTELLE

Experimental Computing

Hall A – Individual experiments PODD framework Root analysis ~10-40 TB/week

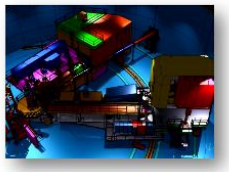

Hall C – Collaborates with Hall A 1.5TB/week

Hall B/CLAS12 JAVA CLARA In-house reconstruction & analysis GEMC (Geant4)

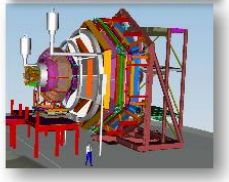

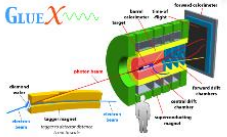

Hall D – C++ MT JANA framework OSG MT Geant4 ROOT

- Impressive Accomplishments in preparation for 12 GeV
  - GlueX – Digital Object Identifiers in publications
  - Hall B/CLAS-12: Working towards data challenges with full reconstruction

DOE/SC Facilities - Observations

- **DOE/SC facilities provide scientists with unique and highly specialized instruments, sensors, computers, and networks**
- **Facility staff provide scientists with expertise in system operation, data collection and analysis**
- **Lab computer and storage infrastructure is expected to play a significant role in meeting the scientists data analysis and management needs**

Research Activities

Large-scale applications using Swift

- A Simulation of super-cooled glass materials
- B Protein and biomolecule structure and interaction
- C Climate model analysis and decision making for global food production & supply
- D Materials science at the Advanced Photon Source
- E Multiscale subsurface flow modeling
- F Modeling of power grid for OE applications

All have published science results obtained using Swift

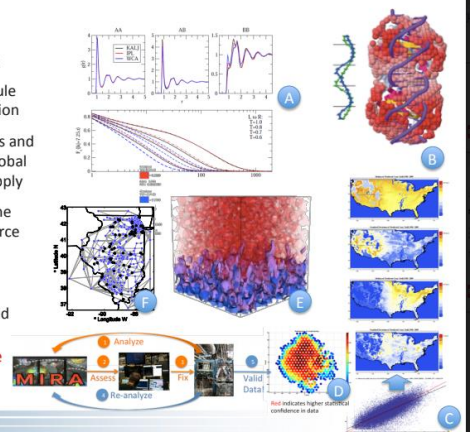

Example Pegasus Applications, varied domains, varied users expertise

DOE Panorama project: Molecular dynamics simulations for **drug delivery**

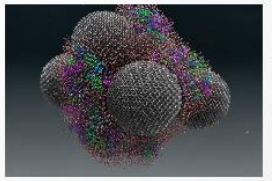

ORNL used Pegasus to confirmed that ***nanodiamonds*** can enhance the dynamics of

Processing of **MRI data for Alzheimer's** research

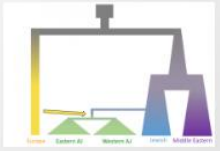

Inference of Human Demographic History
Infer human demographic history, such as global migrations, population size changes and mixing between populations through modeling.

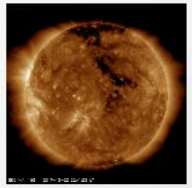

Helioseismology

The **Structural Protein-Ligand Interactome** (SPLINTER) project predicts the interaction of thousands of small molecules with thousands of proteins.

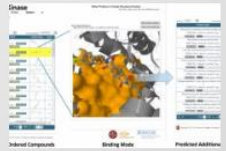

http://pegasus.isi.edu 16

Organization-to-organization Trust

Organization to organization trust delegates user management to intuitions.

E.g. Open Science Grid

13

Rucio beyond ATLAS

User community

- The AMS and Xenon1T experiments are already using Rucio in production:
  - Xenon1T Dark Matter Search
    - Thousands of files across 6 sites (Europe and US), using the MariaDB backend, operated by UChicago
  - AMS (Alpha Magnetic Spectrometer)
    - Millions of files across 10 sites, using the MySQL backend, operated by ASGC Taiwan
- CMS using the PostgreSQL backend operated by UChicago to evaluate Rucio
- + COMPASS, LSST and some others

→ Rucio Community Workshop: March 1-2, 2018

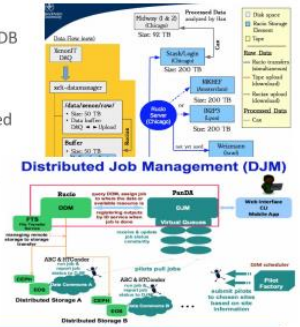

2017-12-14 Rucio - FLC WG 9

Federation - Shared Data and Services

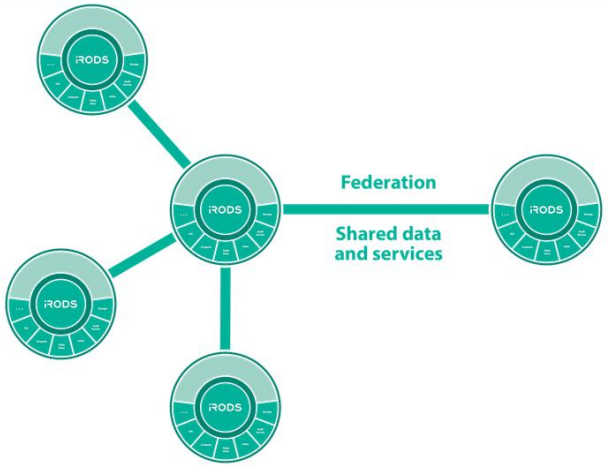

Research Activities - Observations

- **ASCR and other federal agencies have a long history of funding research in the area of High-Performance Distributed Computing (HPDC -> DCDE)**
- **Numerous funded research projects have been widely deployed and are in operational use**
  - GridFTP, Globus, HTCondor, BigPanDa, Pegasus, Swift, Kepler
- **Federated Identity Management services have been in routine use by science communities for decades**

DCDE Components

1. Seamless user access

2. Coordinated Resource allocations and cross-facility workflows

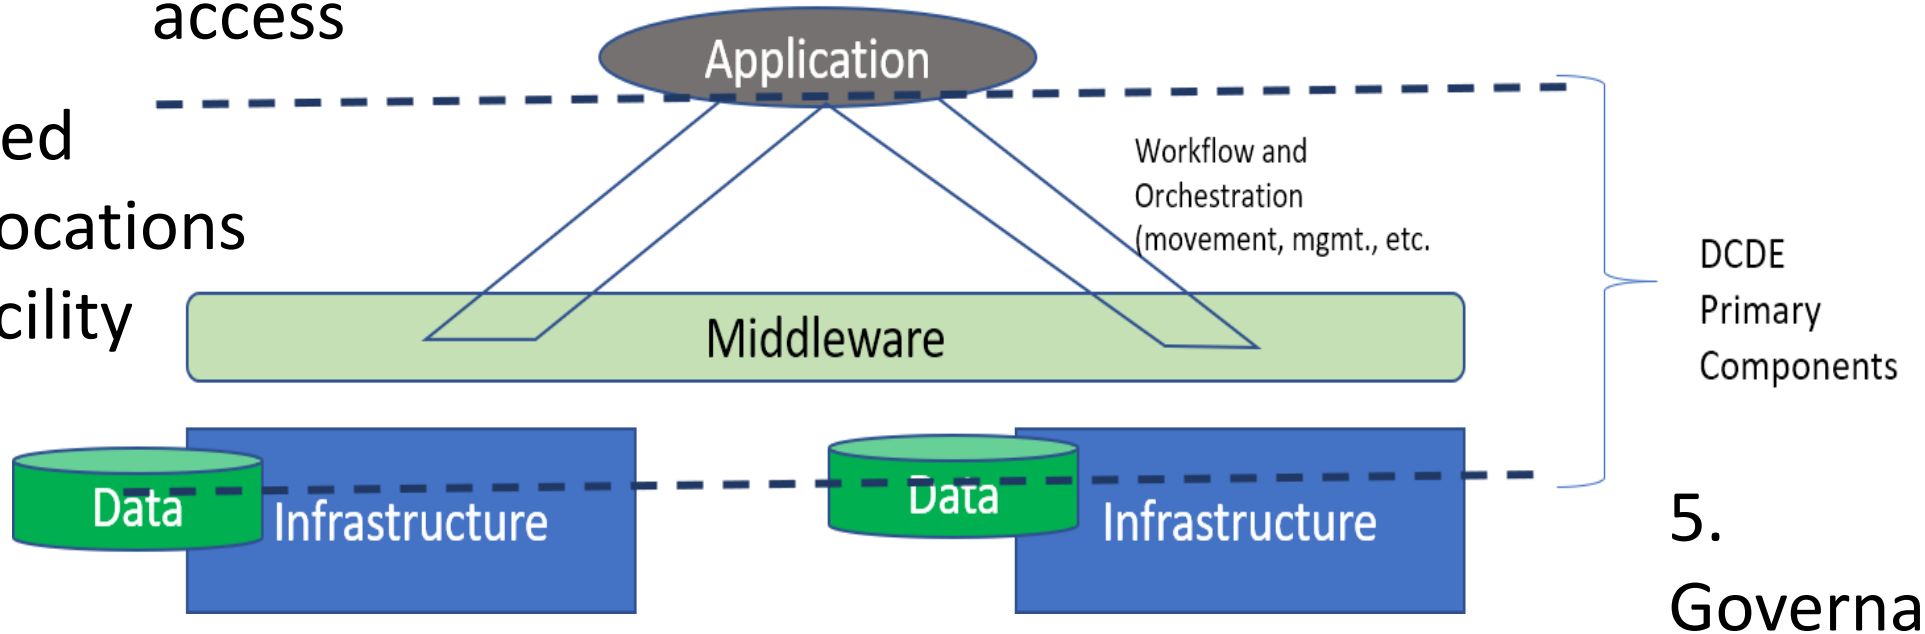

3. Data storage, movement, and dissemination for distributed operations

4. Variety, Portability: through virtualization, and containers, etc.

5. Governance and Policy Structures

FLC-WG Draft Report Findings

- **Seamless User Access required**
  - DOE science is increasingly carried out by teams of scientists with few connections to computer scientists and without expertise in DCDE based computing services
  - Existing Federated Identity Management services can, and should, be deployed widely across the DOE/SC lab complex
  - A mix of commercial (OneID) and research inspired (SAML-based) services can co-exist and will meet the needs of the DOE scientist and Lab staff

FLC-WG Draft Report Findings

- **Coordinated Resource Access and Cross-Facility Workflows**
  - Science specific Workload management systems (HEPCloud, PanDa, Condor) will continue to evolve and provide service to their science community
  - Domain agnostic Workflow Management Systems (Pegasus, Kepler, Swift, etc) will continue to gain adoption
  - Training for both scientists and facilities staff is required to effectively support the existing and emerging Workflow systems
  - A generalized collaboration environment would simplify usage across the DOE/SC lab complex

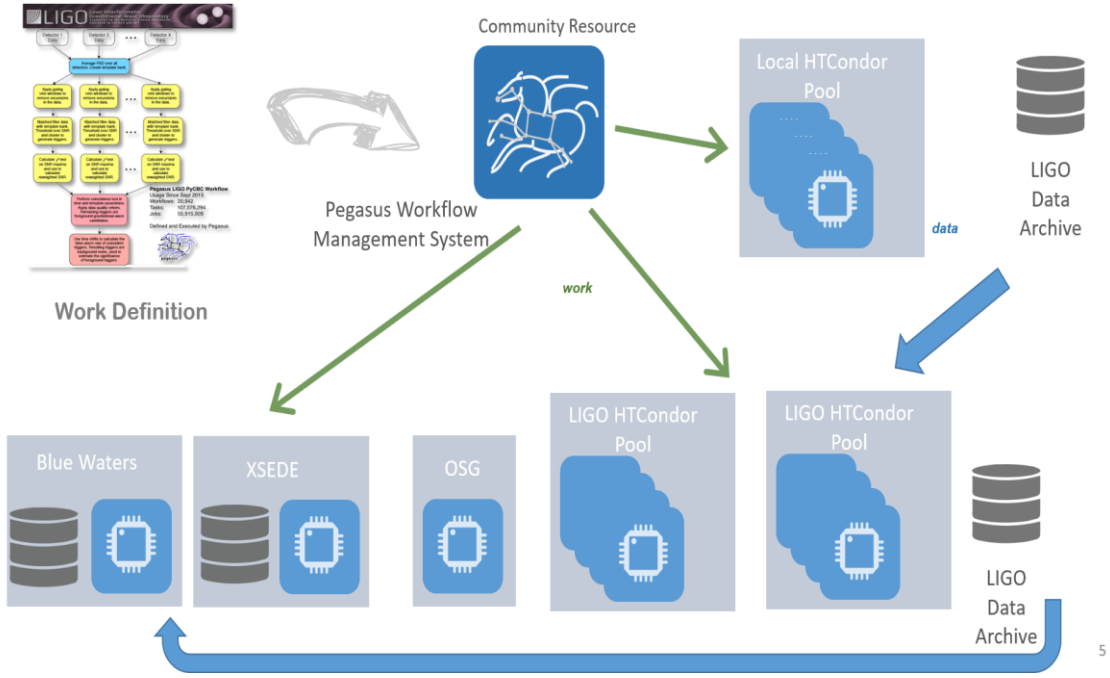

FLC-WG Draft Report Findings

- **Scientific Data Management: Movement, Dissemination and Archival Storage**
  - Large Data Management systems in use by Science communities (HEP, BER)
  - Basic network mechanism for separating bulk data (Elephants) from routine traffic (mice) exist and are in active use
  - Bulk Data Transfer services exist (Globus)
  - Streaming Data transfer issues still a research challenge
  - Science DMZ architecture well defined and widely implemented
  - Science Portals and web based services being improved

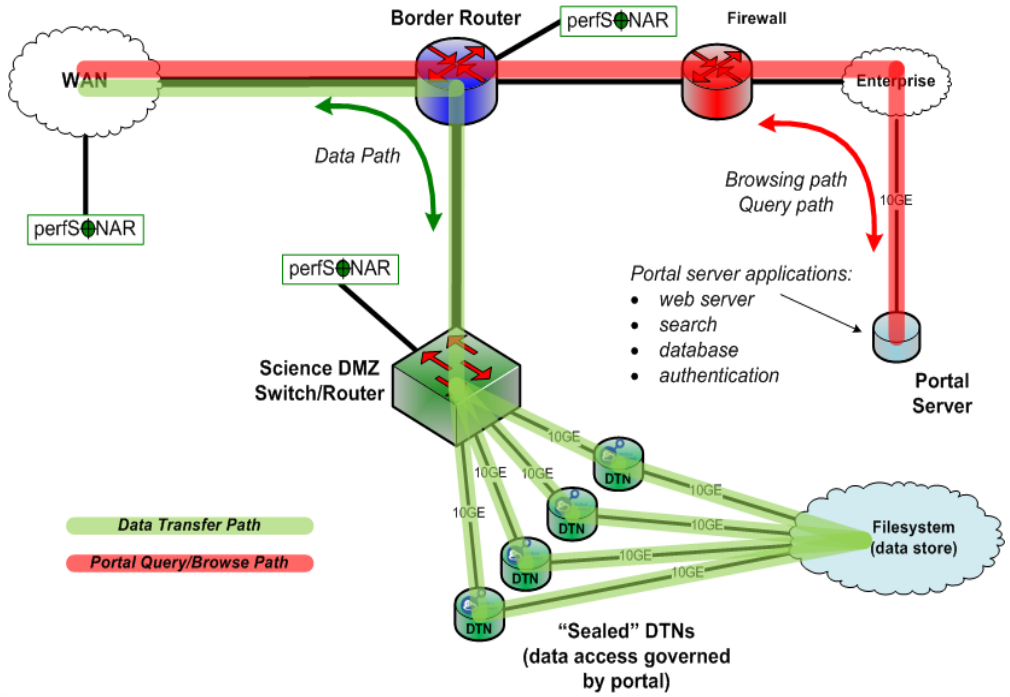

FLC-WG Draft Report Findings

- **Supporting Functional Variety and Portability**
  - Virtualization and Containerization technologies are improving and making run-time environments portable
  - Electronic notebooks (e.g., Juypter) are evolving to interoperate with workflow management systems allowing scientists to create new experiments and record results in an automated fashion
  - A Federated DCDE will span a wide variety of compute systems (i.e., desktops to supercomputers) and deeply interconnect them with other physical devices (i.e., storage systems, distributed sensors, and large/unique instruments) using a high performance network, accelerating the knowledge discovery process throughout the DOE/SC Lab complex

FLC-WG Draft Report Findings

- **Organizational Concerns and Governance Issues**
  - Operating a DCDE will require lab staff and management to develop new skills/training and work with their peers across the lab complex
  - New operating policies and procedures will need to be developed
    - Cost/allocation policies
    - Identity Federation procedures
    - Procurement and deployment of common tools/services/licenses
  - A lab based governance committee representing multiple stakeholders may be needed to coordinate activities across the Lab complex
    - ASCR's experience with ESnet can provide input into this activity

FLC-WG Path Forward

- **Establish a pilot that implements the following items/topics in a coherent and progressive manner**
  - Attribute based ID management (SAML based)
  - Preliminary Allocation policy (Unused cycles)
  - Preliminary Accounting policy (fair exchange)
  - Preliminary Common Access model
    - Science DMZ with DTNs
    - Common Workflow system support
    - Science Portal support

Conclusions

- **The FLC-WG has spent the past year looking deeply at how Lab computing and storage resources can be federated into a coherent DOE/SC DCDE**
- **The Lab NLRCG and ASCR both consider this a reasonable response to the growing demands of DOE scientists and limiting budgetary constraints**
- **The technical and operational expertise exists throughout the lab complex to implement a functional DCDE**
- **The components and services needed to implement a DCDE exist**
- **Policy and governance issues need to be resolved by experimentation and revisions as needed**

Questions?

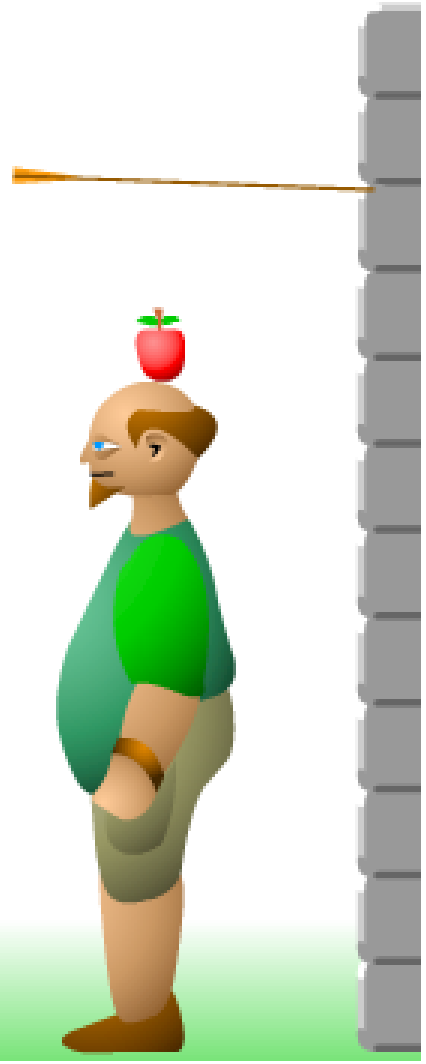

Aim High