HUGS2021: Grid Computing

Malachi Schram, Ph.D.
Department of Data Science
Thomas Jefferson National Accelerator Facility

June 2, 2021

Jefferson Lab

U.S. DEPARTMENT OF ENERGY | Office of Science

JSA

My research experience

Distributed Computing:

- Belle II US computing
- Project 8 computing

Design and Control:

- Accelerator control using reinforcement learning
- Exascale reinforcement learning
- Water cluster and material property optimization

Machine Learning:

- Anomaly detection and prognostication
- Segmentation, object identification, classification
- Domain aware ML

What is Distributed Computing?

- Coordinated collections of geographically distributed computing resources

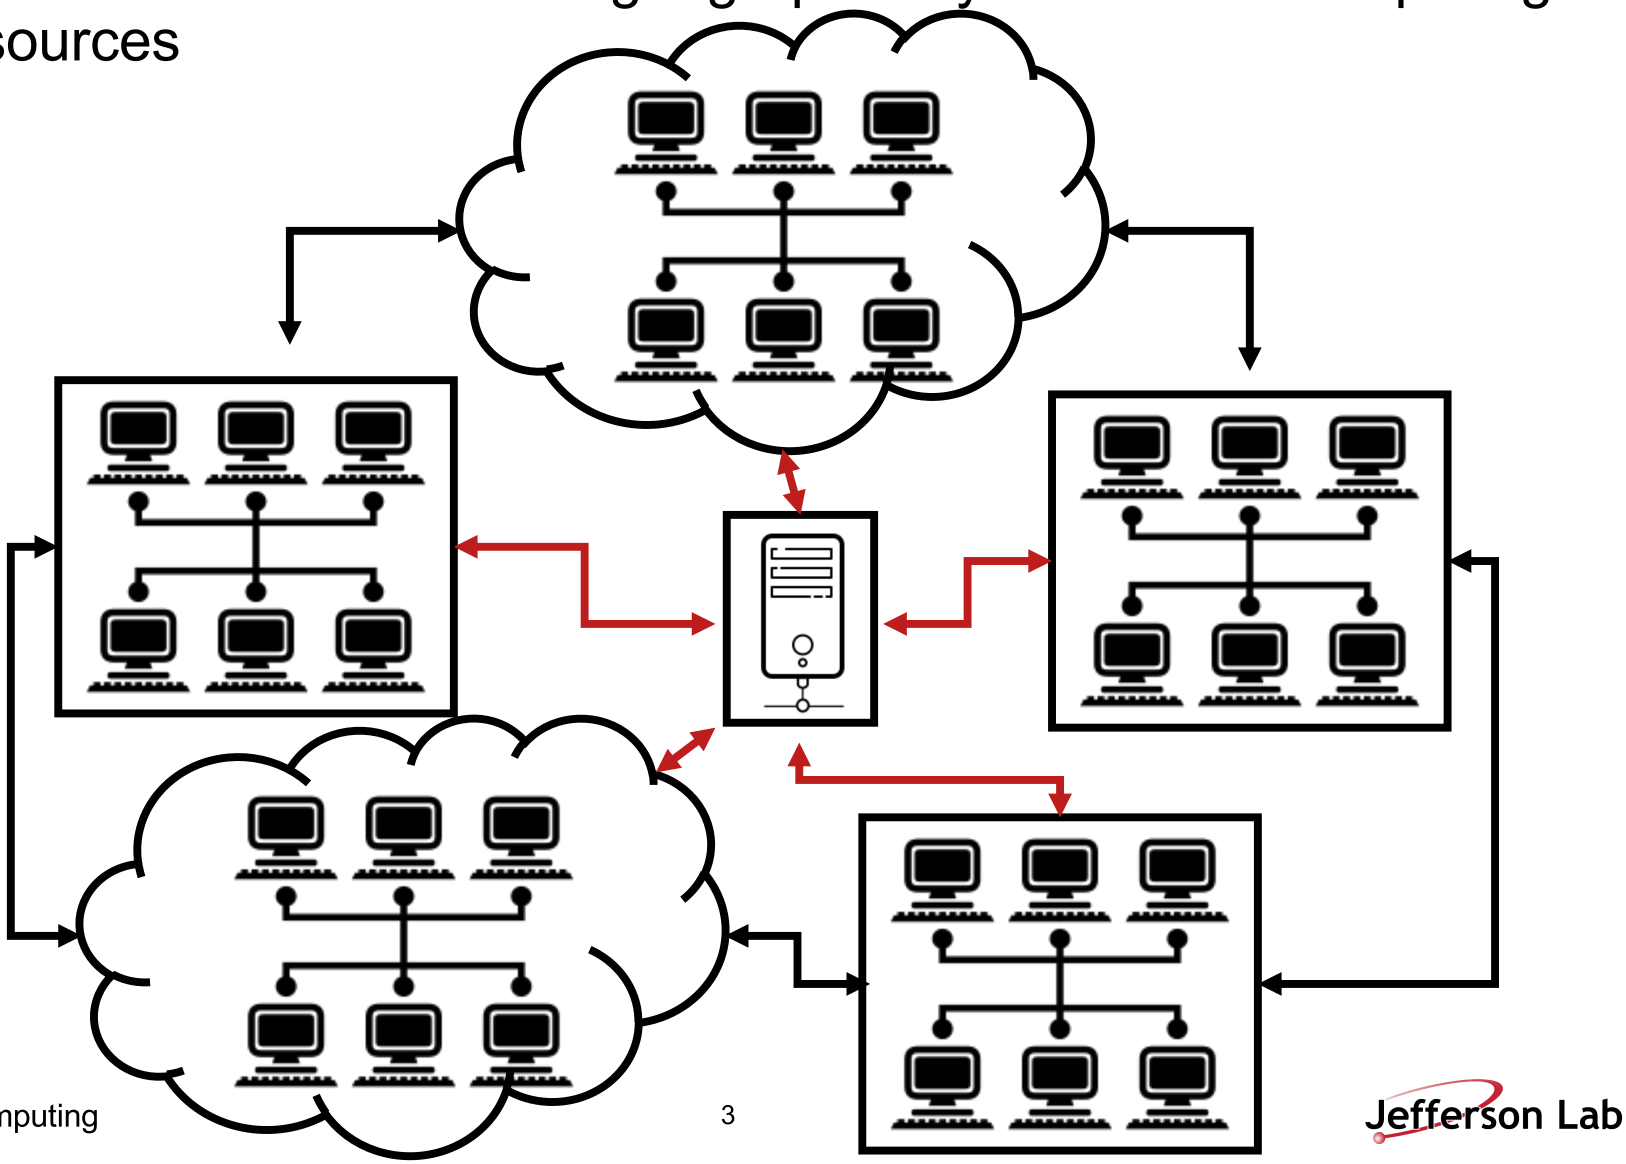

Real world example

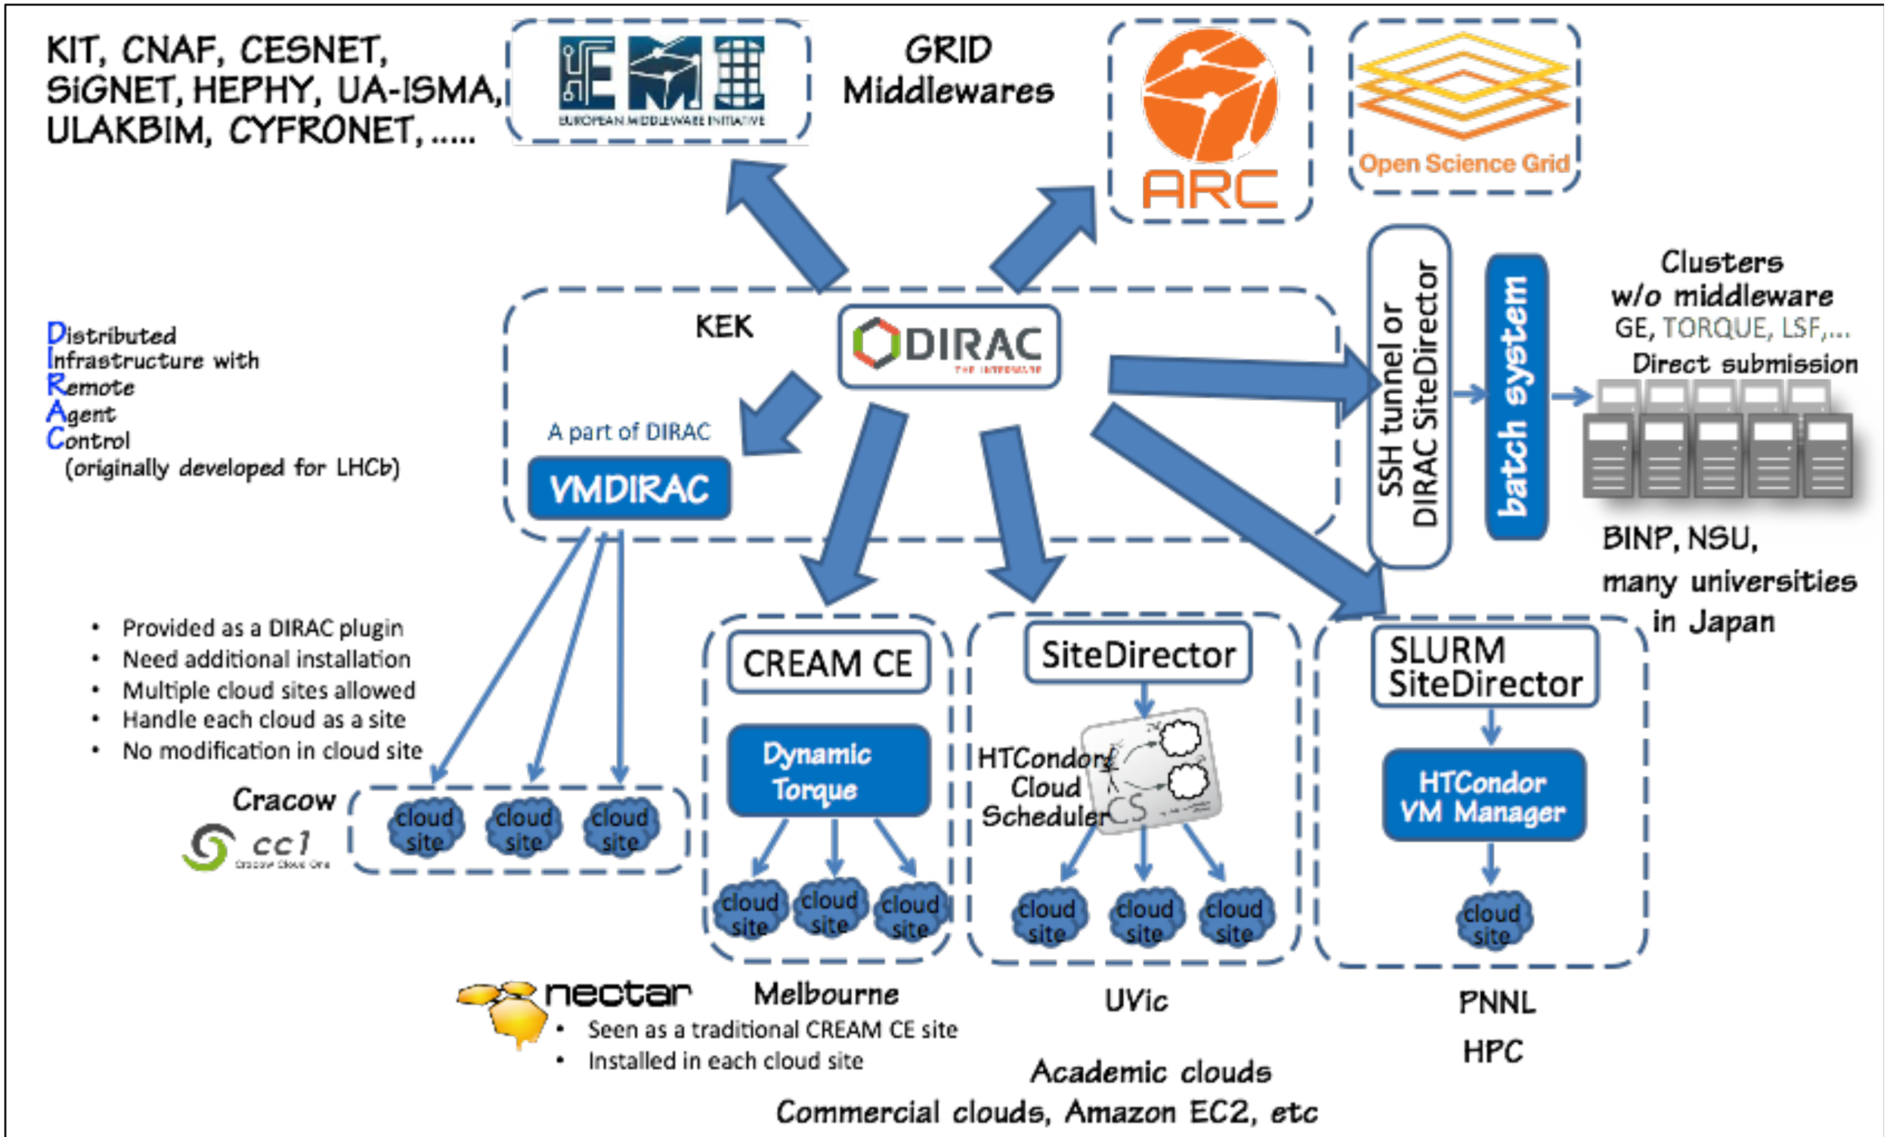

Starting with conventional scientific computing

- User node:
  - Terminal to access the local cluster using site defined authentication
- Login node and job scheduler (SLURM, HTCondor, Torque, etc.)
  - Requires permission to access
  - Head node(s) are used to submit jobs and run *very limited* test jobs
- Compute nodes:
  - Executes the computational work
  - OS, libraries, etc.
  - Software access (either on shared storage or on each individual node)
- Shared storage (Lustre, NFS, Ceph, etc.)
  - Typically visible throughout the cluster

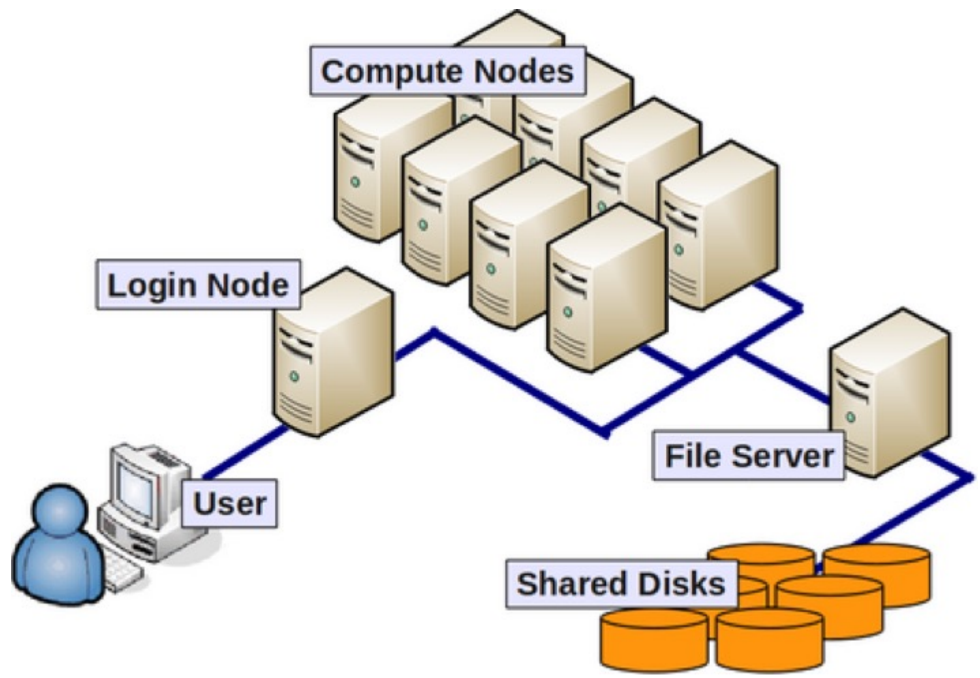

Additional considerations for grid over local computing?

We need to coordinate computing resources across various sites. Here are a few things to consider:

- Resource requirements and allocation
- Authentication services
- Computing resources and software distribution
- Data storage, transfers, and bookkeeping
- Grid Orchestration
- Workflow automation

Resource Estimations

- Large scientific projects such as the LHC experiments, Belle II, and soon EIC have incredible computing requirements.
- Example from Belle II:

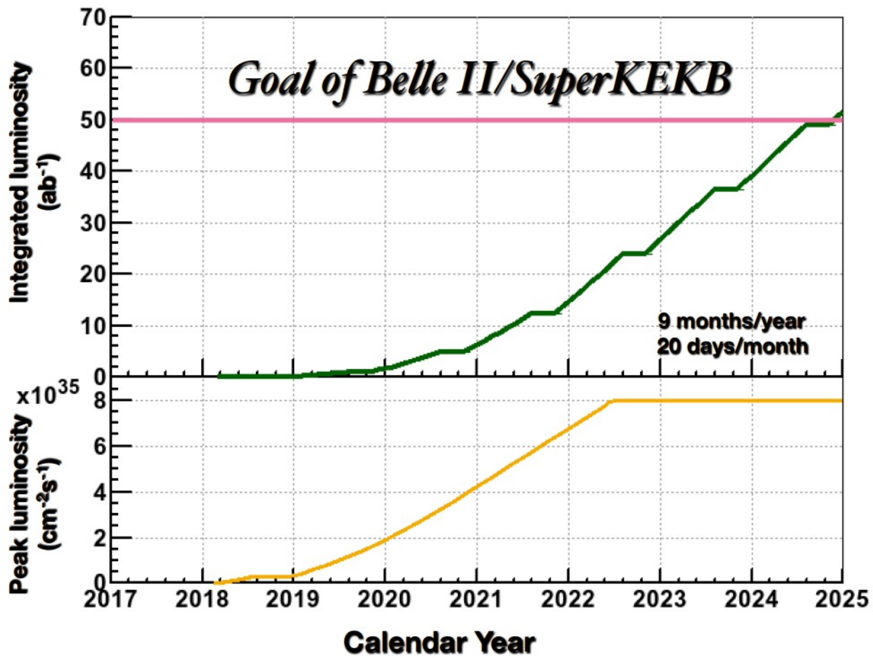

- Resource requirements depend upon luminosity profile & event size (data volume), (re)processing, MC generation, and user analysis
- MC production accounts for nearly 50% of US Belle II CPU
- Potentially use Leadership Class Facilities (LCFs) for MC production

	CY17	CY18	CY19	CY20	CY21
CPU [kHEPSpecs]	20.11	27.56	58.90	69.71	82.97
Storage [PB]	0.31	0.81	5.04	6.50	9.28
Networking In/Out [Gbps]	0.30/0.30	0.49/0.36	1.06/0.26	1.56/0.31	1.89/0.83

Coordinating network resources providers

- Workshops are organized between compute sites and networking providers
- The purpose of these workshops are preparation for addressing the wide-area networking requirements
- Example report can be found at:
  - http://www.es.net/assets/pubs_presos/Belle-II-Experiment-Network-Requirements-Workshop-v18-final.pdf

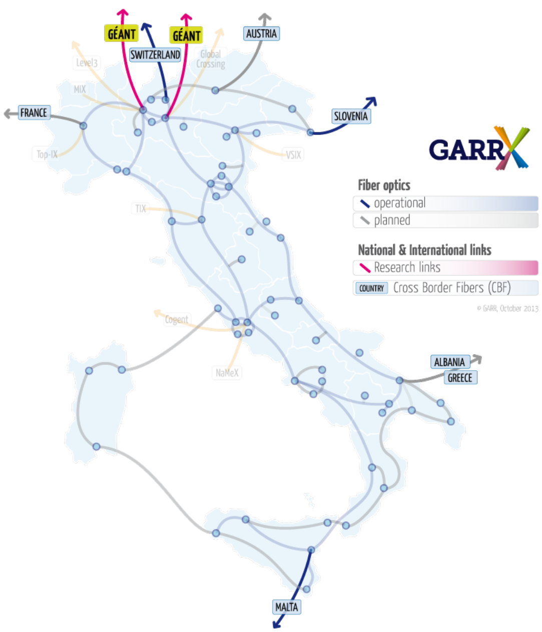

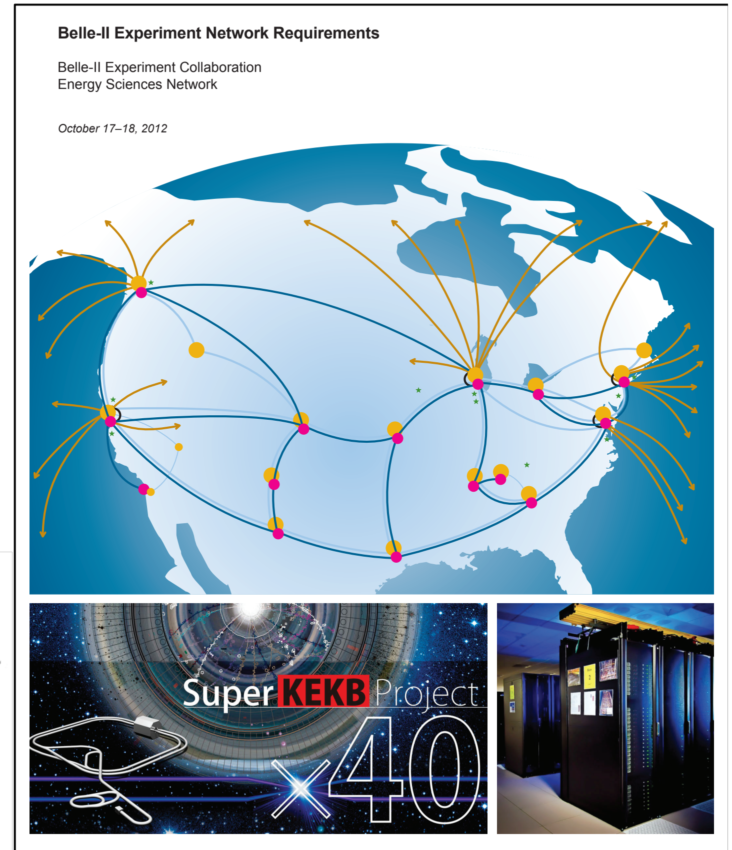
Belle-II Experiment Network Requirements

Belle-II Experiment Collaboration
Energy Sciences Network

October 17–18, 2012

Initial network traffic estimations

- Based on the capability at each site and local network provider
- The contributions are based on available resources and politics

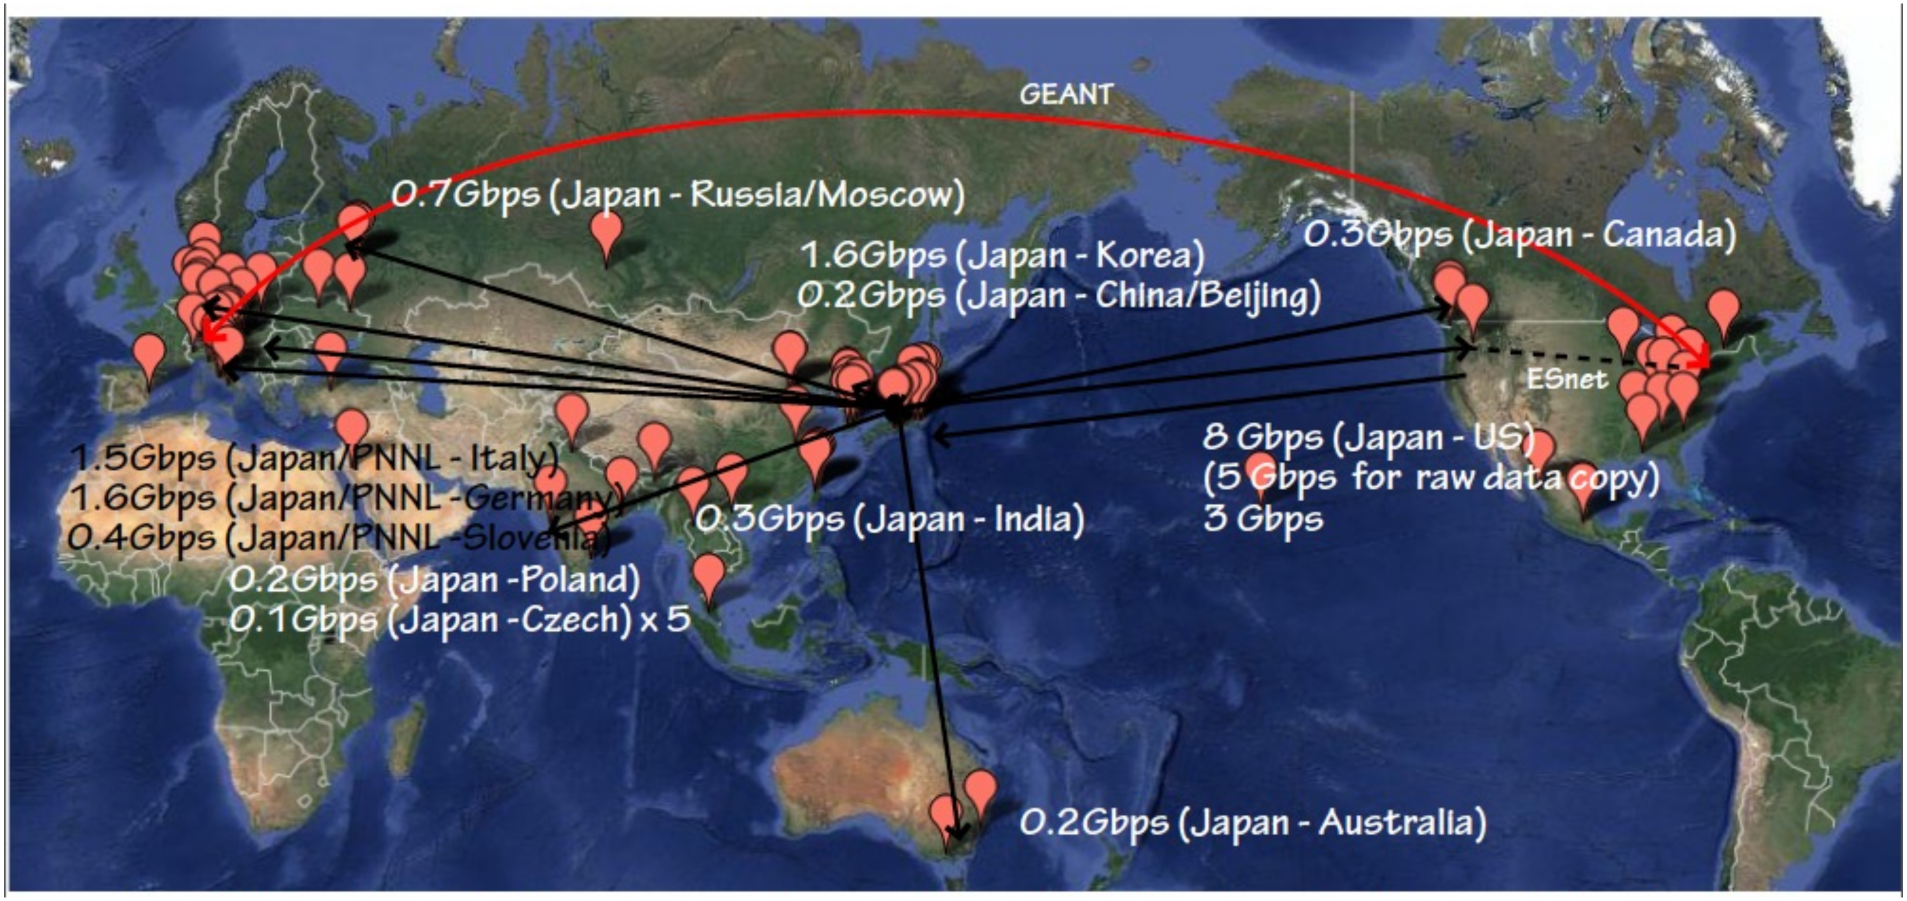

Virtual Organization & Authorization Services

- **Virtual Organization (VO)** is a dynamic set of individuals or institutions defined around a set of resource-sharing rules and conditions.
- **Virtual Organization Management System (VOMS)** provide a structured account database for users, group, and Vos.
- **Grid User Management System (GUMS)** is a site tool that captures the VOs and associated users and policies that will be allowed access to the site.

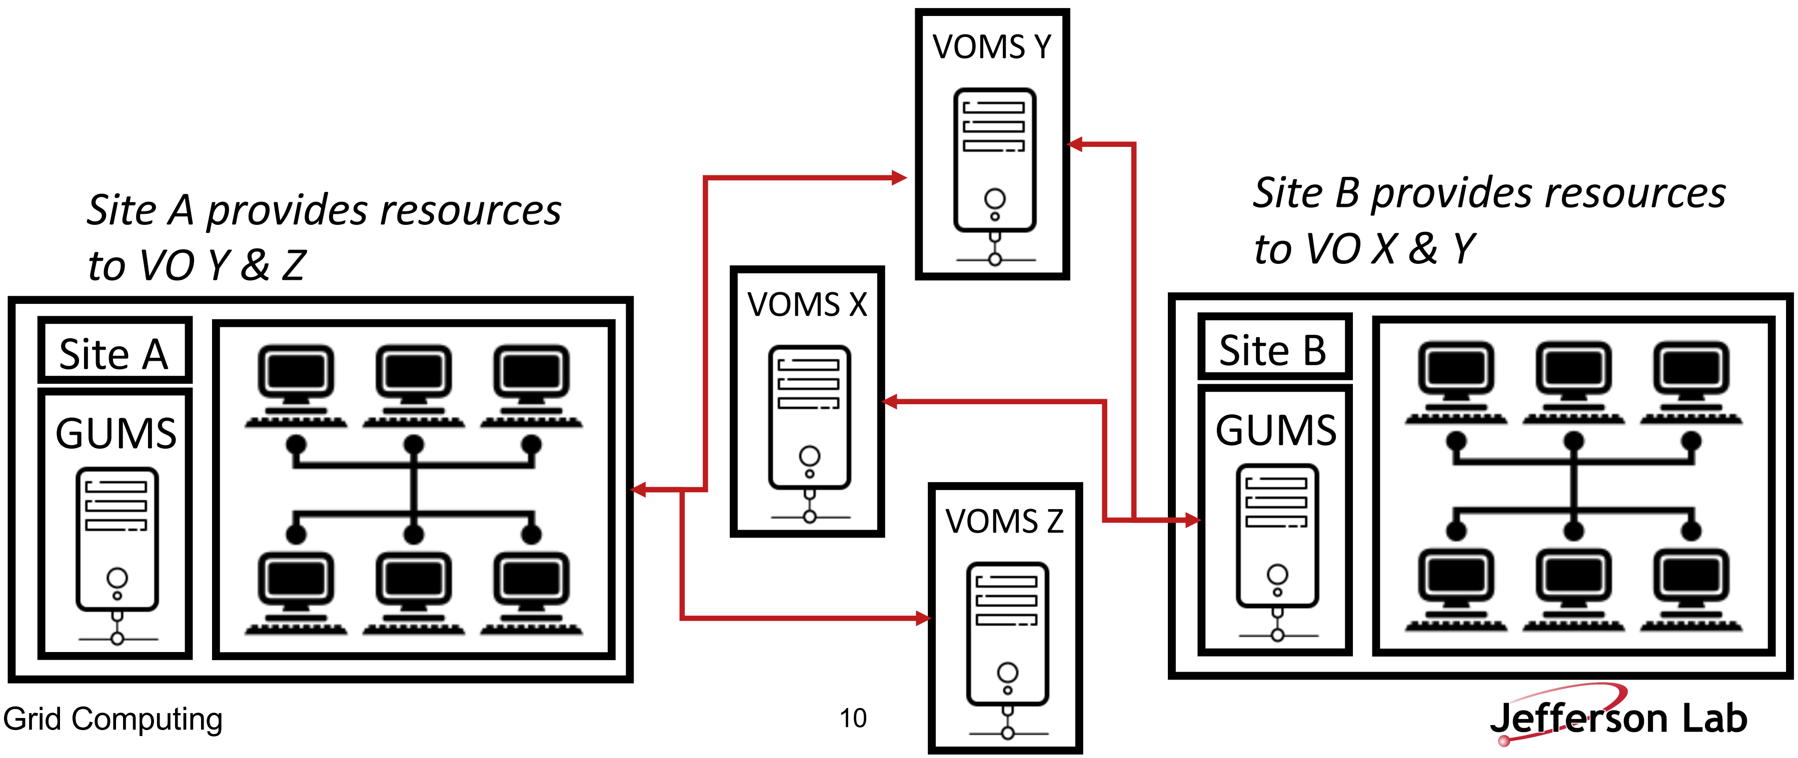

Computing resources and software distribution

- Computing Element (gateway node):
  - Access point for information to be exchanged between the central system and the local cluster
  - Can be used as a proxy for the worker nodes to access the outside world.
  - Push and pull techniques are available depending on the grid orchestration
- Worker nodes
  - Executes the job similar to traditional cluster with some network and software requirements
- Software coordination:
  - Define OS and required libraries
    - Compilers, domain software, etc.
    - This can cause a lot of headaches if you are supporting multiple VOs
  - CernVM File System (CVMFS): (https://cernvm.cern.ch/fs)
    - Provides a scalable software distribution platform
    - Multi-tiered caching system for scaling
    - Stratum 0, Stratum 1, Squid servers, and clients
  - Containers (ex: https://www.docker.com/):
    - Used to decouple the hardware level OS from the project requirements
  - Hybrid approach: containers in CVMFS

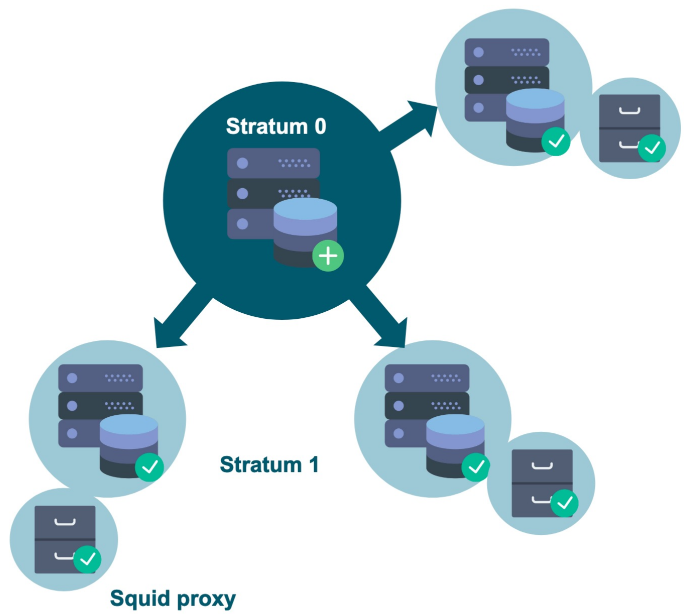

Data storage, transfers, and bookkeeping

- Storage Element (SE) is used to communicate with the local storage system and coordinate data transfers
- Data transfer nodes (DTN) leverage GridFTP for data transfers
- The data location is captured using a logical file name and associated information on physical location (site, path, etc.)
- Mass data transfer is done using tools like:
  - Globus online (https://www.globus.org/)
  - File Transfer Service 3 (https://fts.web.cern.ch/fts/)
- Network monitoring:
  - perfSONAR (https://www.perfsonar.net/)
  - Monitoring and Debugging Dashboard (MaDDash)

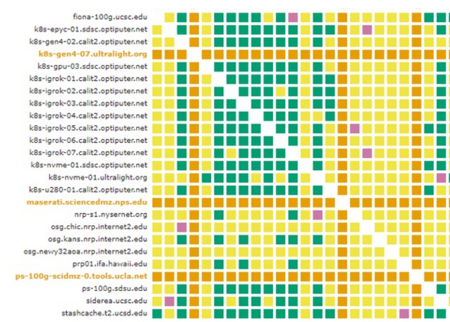

Network Data Challenges

- Run network stability tests to ensure there are no problems with the network and the DTN configurations and FTS setup.

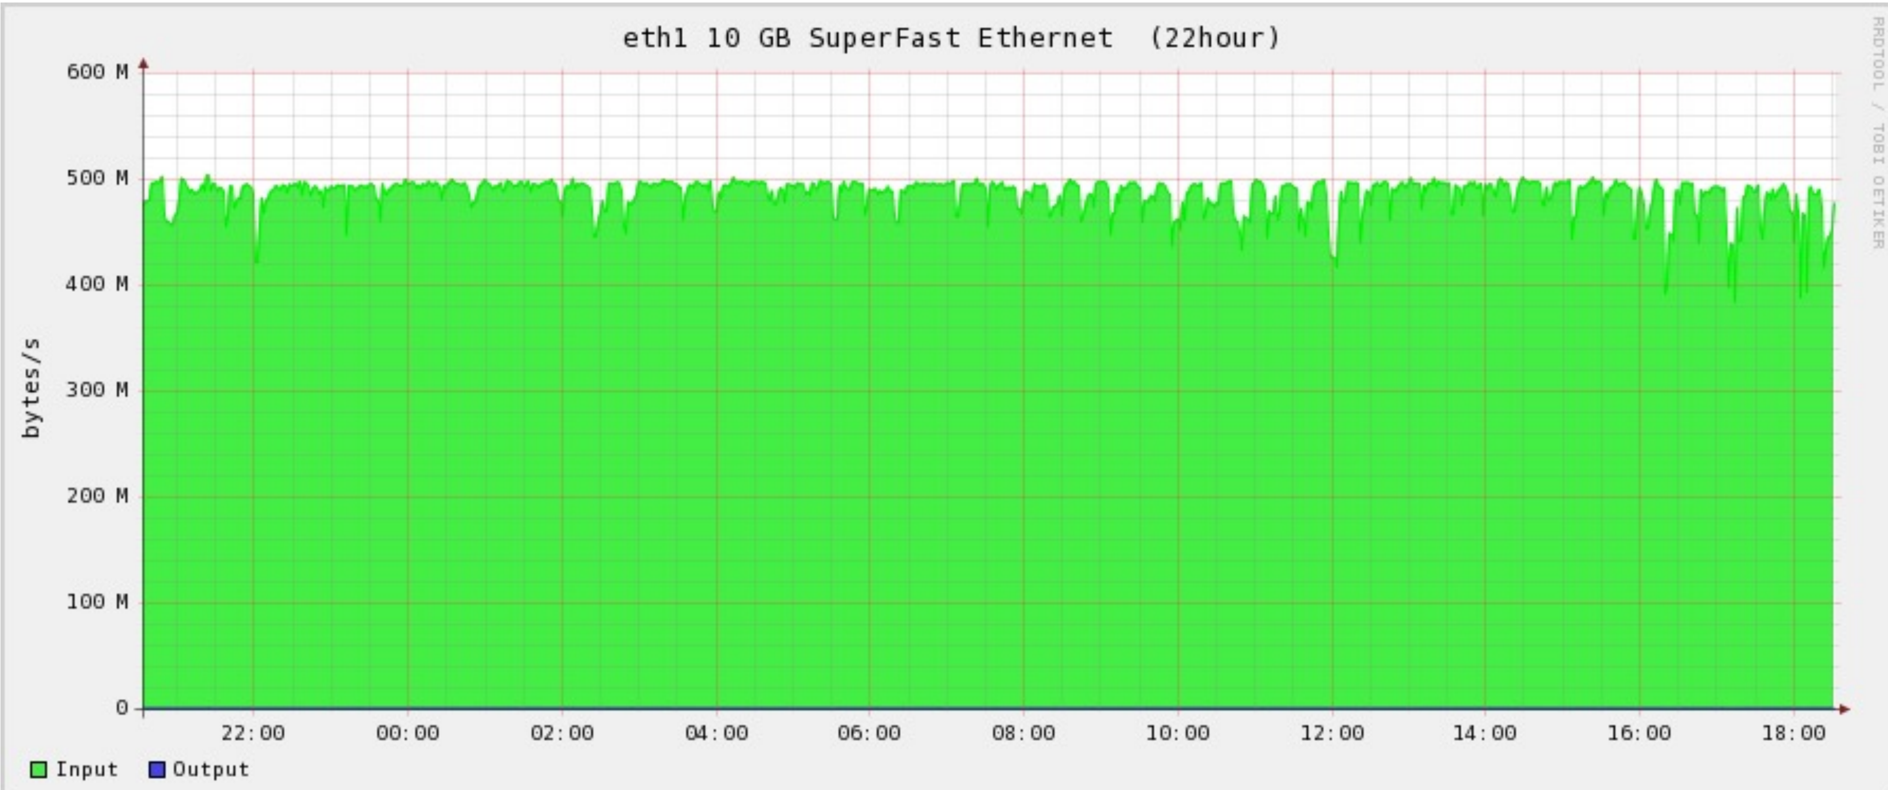

Jefferson Lab

Grid Orchestrations

- Once you have defined the computing requirements, established an authentication protocol, and setup the relevant computing hardware. You will need to identify a mechanism to orchestrate the workflows.
- There are a lot of elements required to provide a production ready grid system.
- There are several established LHC solutions, such as DIRAC (https://dirac.readthedocs.io), that provide great tools to orchestrate scientific workflows:

1. Workload Management System (WMS)
   - Provides users job submission functionality while hiding the heterogeneity of the computing resources. Compute jobs use Pilot Jobs model
2. Data Management System (DMS):
   - Provides basic functionality to upload and register a local file in a StorageElement (SE) and FileCatalog (FC). Additionally, it also provides bulk data replications using FTS and retrieves achieved data on Tape for later processing.
3. Transformation System (TS):
   - Provides automation functionality for redundant tasks commonly related to production activities

Grid Orchestration Example

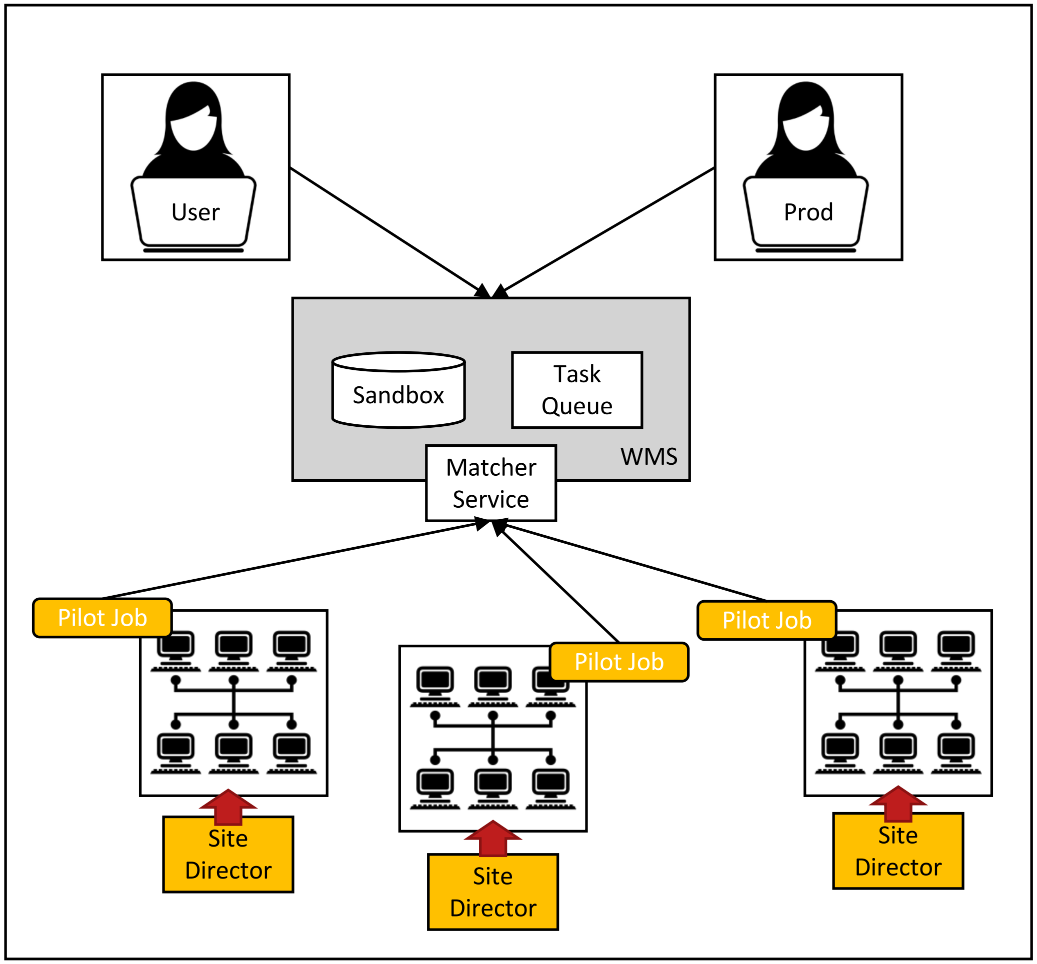

•**User or Prod** submits jobs to DIRAC WMS.

•**Sandbox** stores input file(s) required to run job. This is NOT the same as input files that are available on the SEs.

•**Task Queues** records the job requests and associated requirements.

•**Matcher Service** is used to find an appropriate job based on requirements.

•**Pilot job** is an "empty" job used for resource reservation. Performs sanity check and enables job matching.

•**Site Director** submits pilot jobs to the local queue based on configuration parameters.

Example of Automation Workflows

- For previous efforts, I adopted DIRAC (Distributed Infrastructure with Remote Agent Control) INTERWARE as the distributed computing workflow because it provides a complete solution and includes automation features.
- These automated features are <u>critical</u> for smaller scale experiments
- The files are registered to the DIRAC File Catalog with well defined <u>metadata</u> in order to trigger the raw data processing workflow. Once a data production run is finished, all raw processed files are automatically merged and analyzed.

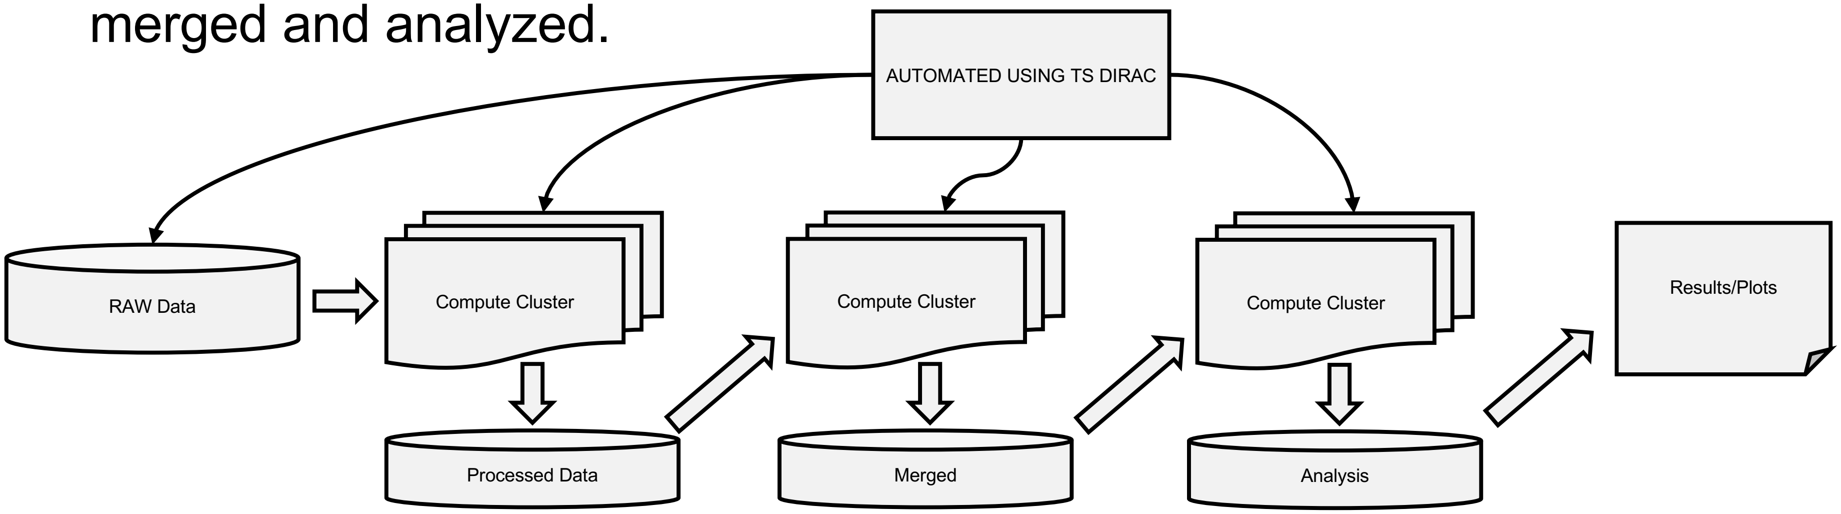

Side note: Designing a flexible infrastructure

- The core computing services are hosted and managed using Kubernetes (https://kubernetes.io/).
- Individual containers are used to instantiate services, agent, and databases and other core grid services.
- Docker containers are also used to build new instances of the development and production environment.
- This provides maximal flexibility and satisfies the collaborations specific OS and libraries requirements.
- Production containers are then converted to a Singularity (https://singularity.hpcng.org/) image.
- Computation jobs are performed on the HPC clusters using a dedicated DIRAC agent mapped to the desired singularity image.

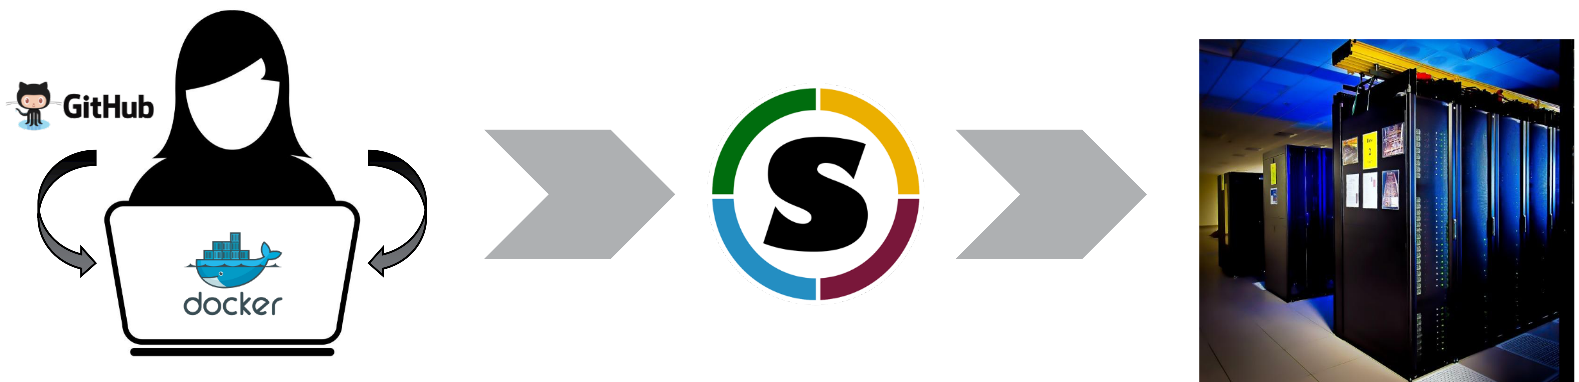

Malachi Schram, Ph.D.
Department of Data Science
Thomas Jefferson National Accelerator Facility
schram@jlab.org

- Upcoming Computing Trends in Nuclear Physics Talks:
  - The online world (Graham Heyes)
  - Data Science 1-3 (Malachi Schram)