

FabrIc for Frontier Experiments project at Fermilab

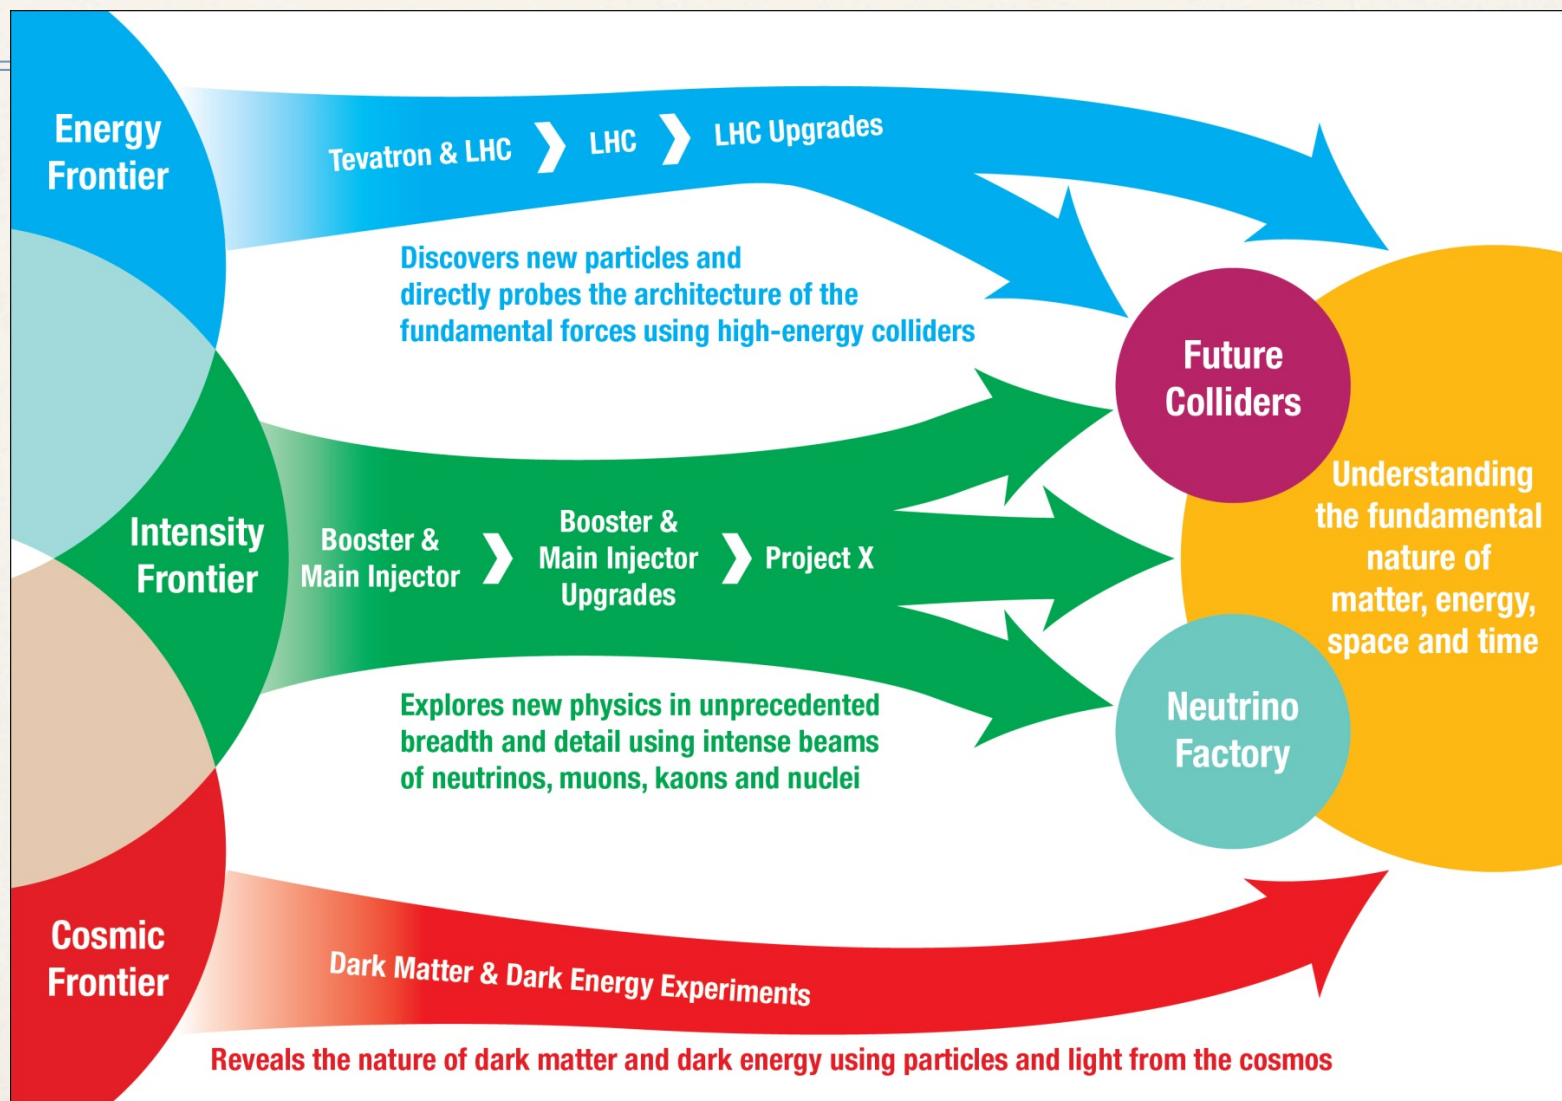
Mike Kirby, CHEP2013

Oct 14, 2013



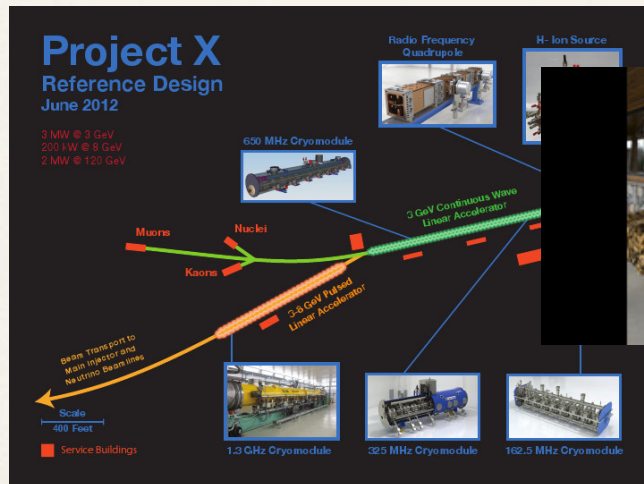

Outline

- ❖ Fermilab physics program in the coming decade
- ❖ Computing needs for frontier experiments at Fermilab
- ❖ FIFE strategy for providing computing resources
- ❖ Current services integrated in the FIFE model
- ❖ Recent successes and plans
- ❖ Summary

Fermilab Frontier Roadmap

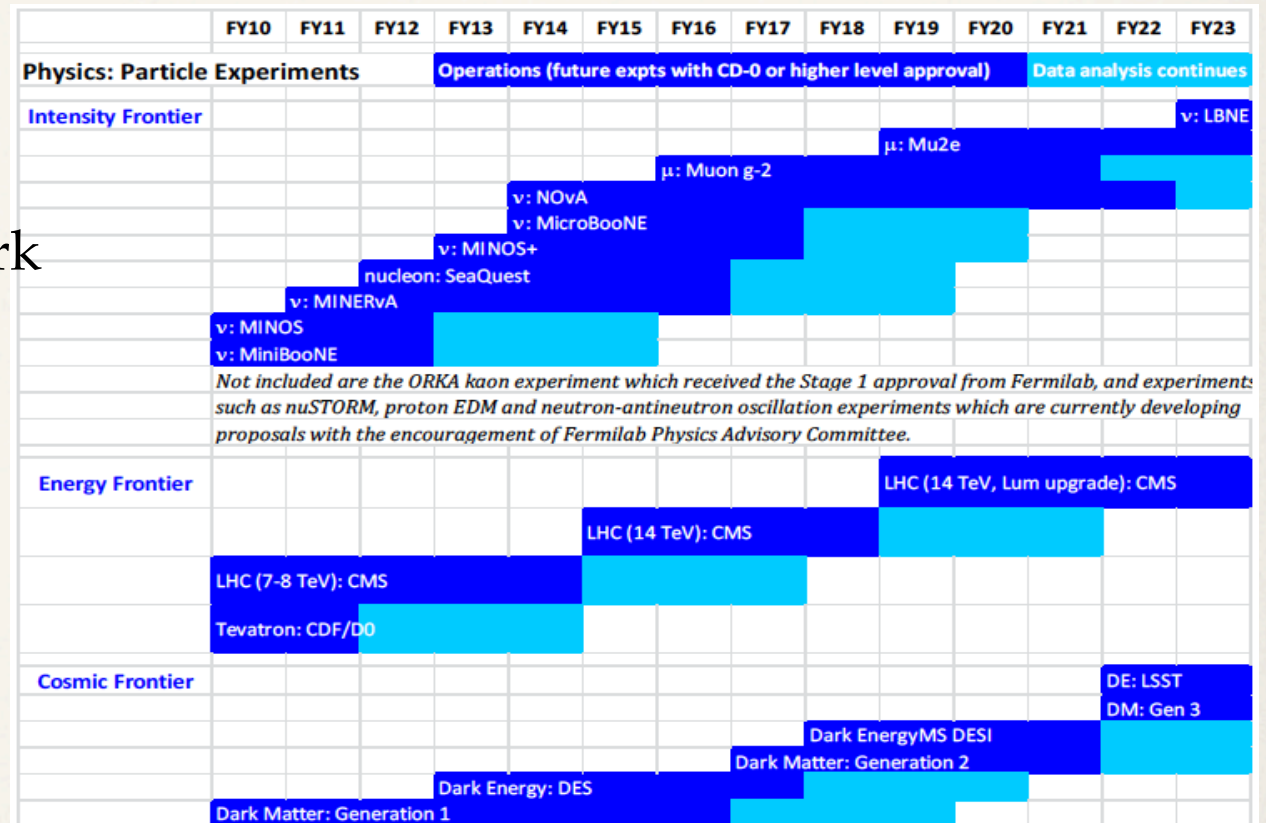


Diverse program of experiments



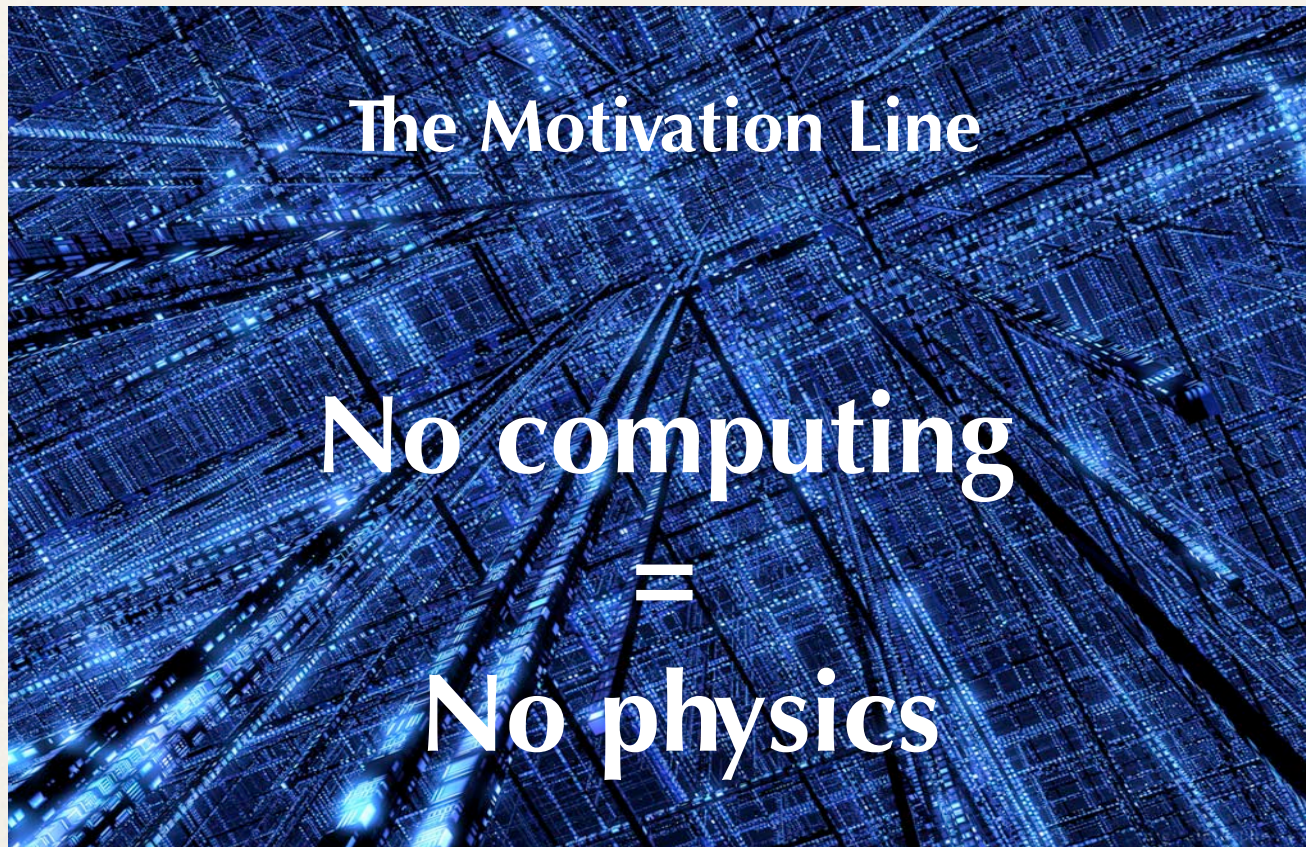
Fermilab Experiment Schedule

- measurements at all frontiers -Electroweak physics, neutrino oscillations, muon g-2, dark energy, dark matter
- 8 major experiments in 3 frontiers running simultaneously in 2016
- sharing both beam and computing resources
- impressive breadth of experiments at FNAL



Computing's role in the success of experiments

- ✦ put simply....



From Mayly Sanchez, Snowmass 2013

Computing requirements for experiments

Fermilab Scientific Computing Review

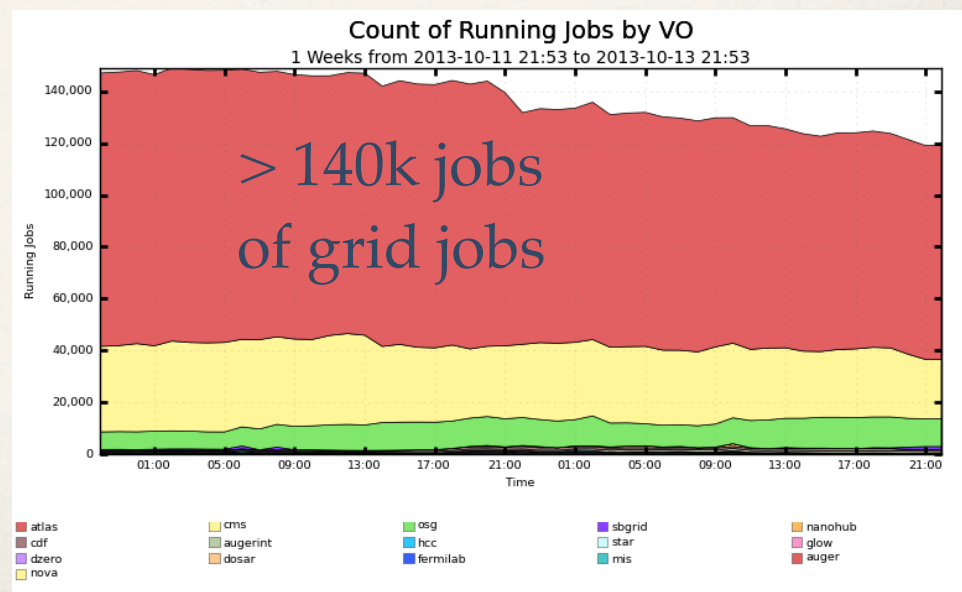
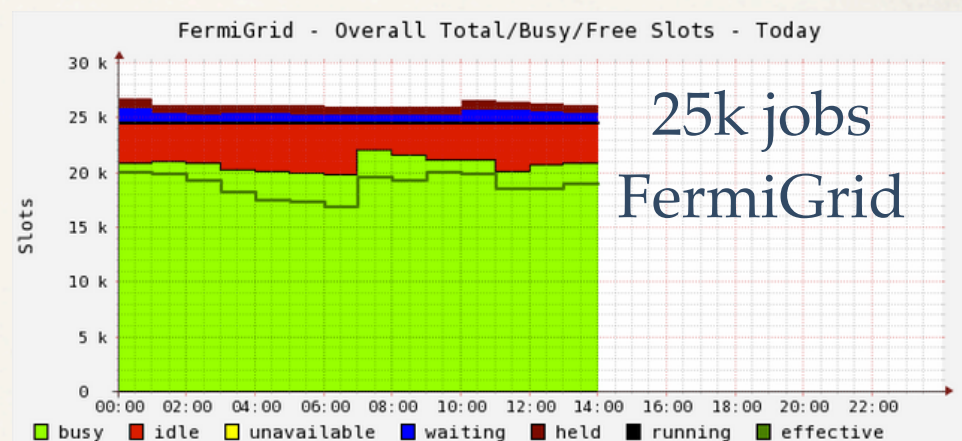
- ❖ higher intensity and higher precision measurements are driving request for more computing resources than previous “small” experiments
- ❖ beam simulations to optimize experiments - make every particle count
- ❖ detector design studies - cost effectiveness and sensitivity projections
- ❖ higher bandwidth DAQ and greater detector granularity
- ❖ event generation and detector response simulation
- ❖ reconstruction and analysis algorithms

| Experiment | Allocation (Slots) | Average Utilization (last quarter) | | FY14 Request (Slots) |
|------------------------|--------------------|------------------------------------|--------------|----------------------|
| | | Average Slots | Peak Slots | |
| ArgoNeuT | 200 | 75 | 1712 | 0 |
| Muon g-2 | 200 | 9 | 195 | 0 |
| LBNE | 500 | 30 | 1202 | 200 |
| MARS | 1225 | 212 | 1444 | 100 |
| MicroBooNE | 200 | 39 | 597 | 100 |
| Minerva | 800 | 667 | 3580 | 500 |
| Minos | 1200 | 294 | 2427 | 0 |
| Mu2e | 500 | 597 | 2491 | 500 |
| Nova | 800 | 410 | 1799 | 500 |
| Total Requested | 5625 | 2333 | 15447 | 1900 |

❖

available onsite resources

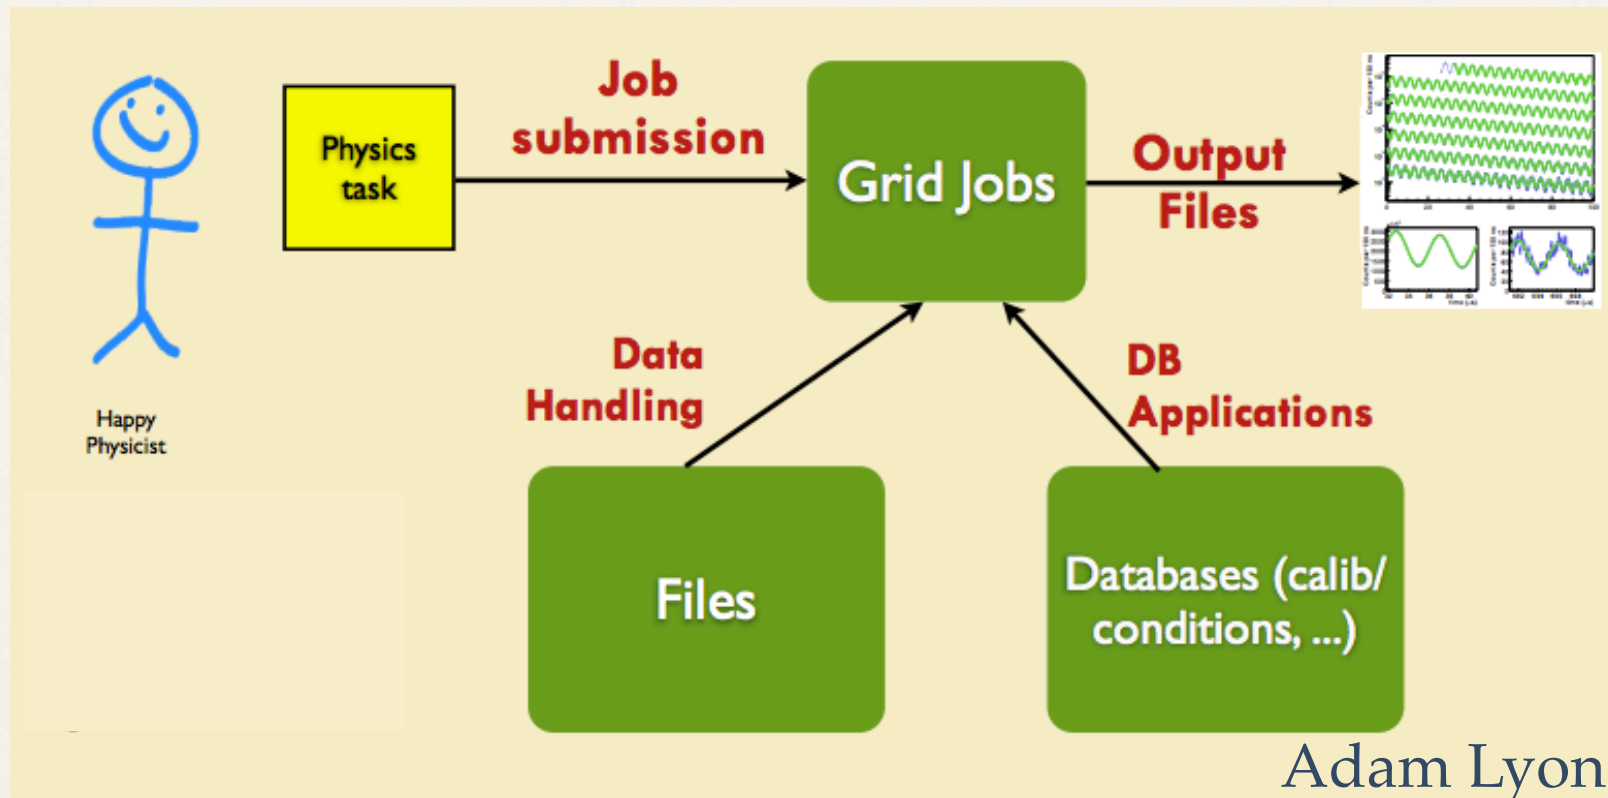
- ❖ victim of our own success
- ❖ allowed smaller experiments to develop computing architectures focused on local resources
- ❖ now a need to reintegrate old architectures into a distributed computing model
- ❖ > 25000 Open Science Grid slots on-site
 - ❖ ~5000 slots for new smaller experiments
 - ❖ off-site resources will be needed to meet needs for experiments
 - ❖ using OSG model for processing provides an opportunity for universities and laboratories to contribute resources to experiments remotely



transition to OSG important for smaller experiments

- ❖ Most experiments are relatively small (~20-300 researchers)
- ❖ unlike Tevatron or LHC experiments do not have the person-power resources to develop large scale computing infrastructure
- ❖ architecting and implemented unique computing infrastructure too great of a burden
- ❖ Fermilab SCD's goal is to provide tools that are easily integrated to meet the needs of smaller experiments, flexible enough to meet needs and OSG focused
- ❖ need to both inform the experiments and gather requirements from them
- ❖ what resources are available?
- ❖ how do you get a job onto the OSG?
- ❖ what storage elements are available?
- ❖ how does it all fit together?
- ❖ FIFE is presenting an integrated solution based on new architecture

FabIc for Frontier Experiments (FIFE)

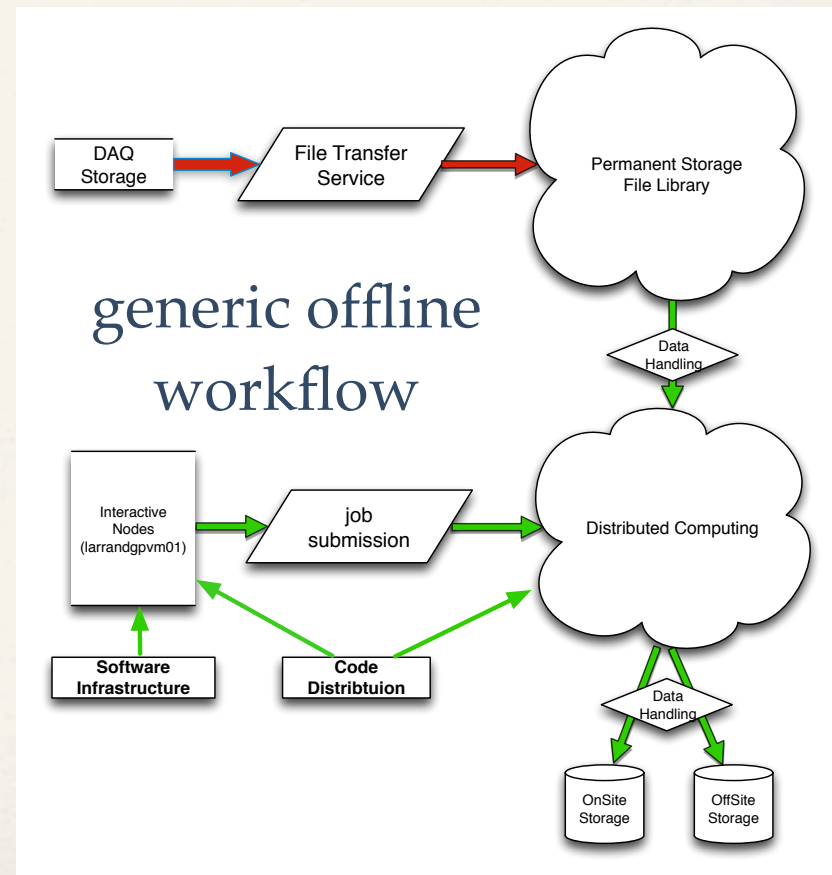


Adam Lyon

- ❖ provide infrastructure to experiments that get them to computing resources while focusing on their science
- ❖ utilize previous solutions and integrate together into seamless model

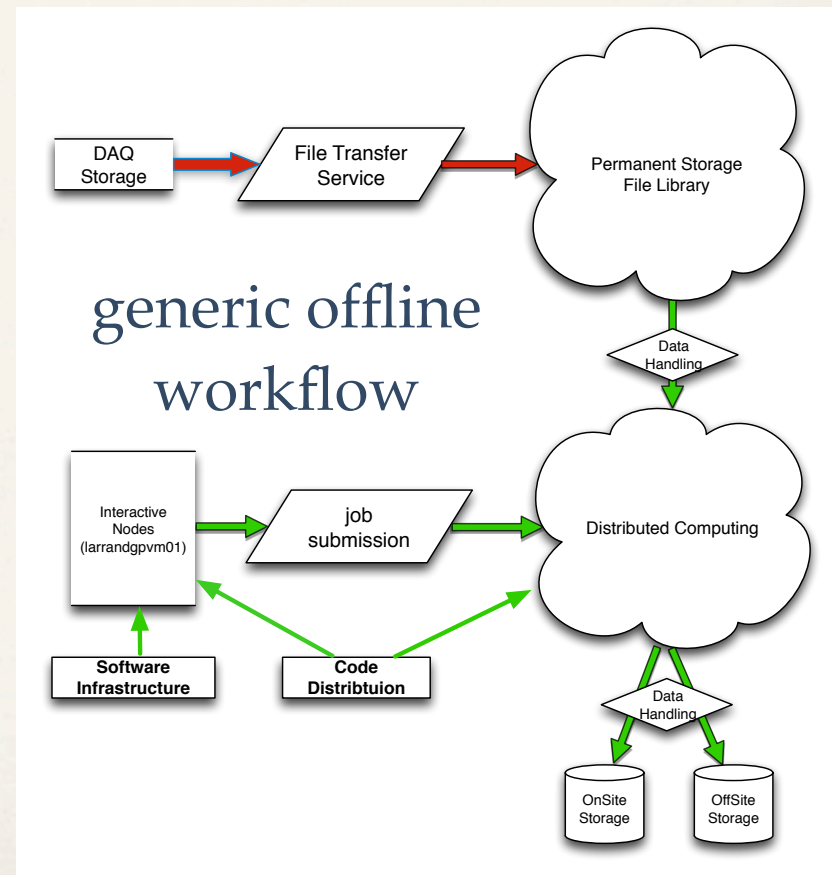
FIFE Strategy

- ❖ address all of the computing needs for experiments
- ❖ modular enough so that experiments can take what they need
- ❖ well enough designed so that while underlying solution may change, interface will be consistent
- ❖ provide mechanism for feedback from experiments to incorporate their tools and solutions
- ❖ help experiments utilize computing beyond the Fermilab campus
- ❖ integrate new tools and resources from outside Fermilab and other communities as they develop



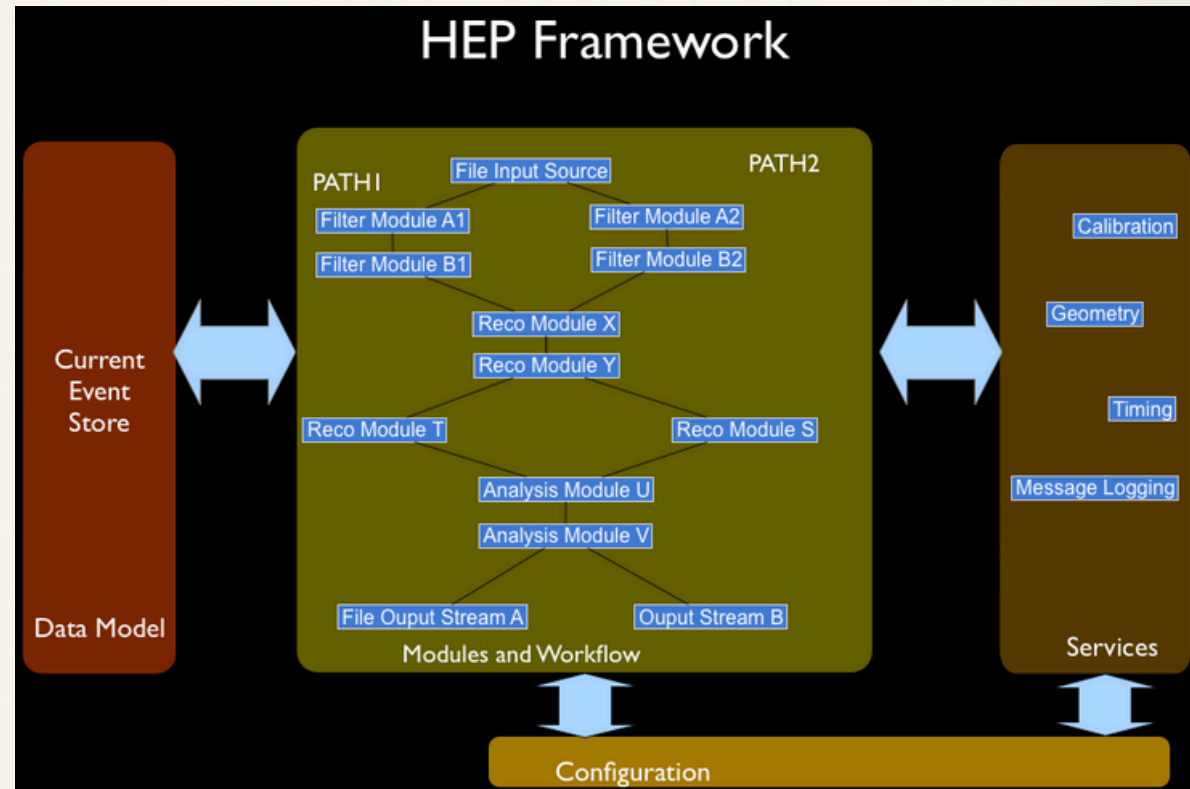
current FIFE implementations

- ❖ Software Framework - art has support of other services native
- ❖ build environment and distribution
 - ❖ git, svn, gmake, cmake and distribution with CVMFS - Andrew Norman's [talk](#)
 - ❖ build machine in development
- ❖ data handling - access to file catalog, tape storage, cache disk, and file transfer
- ❖ database infrastructure and access
- ❖ shared software (LArSoft for Liquid Argon TPC reco)
- ❖ additional infrastructure - authentication, electronic control log, new user accounts



art framework

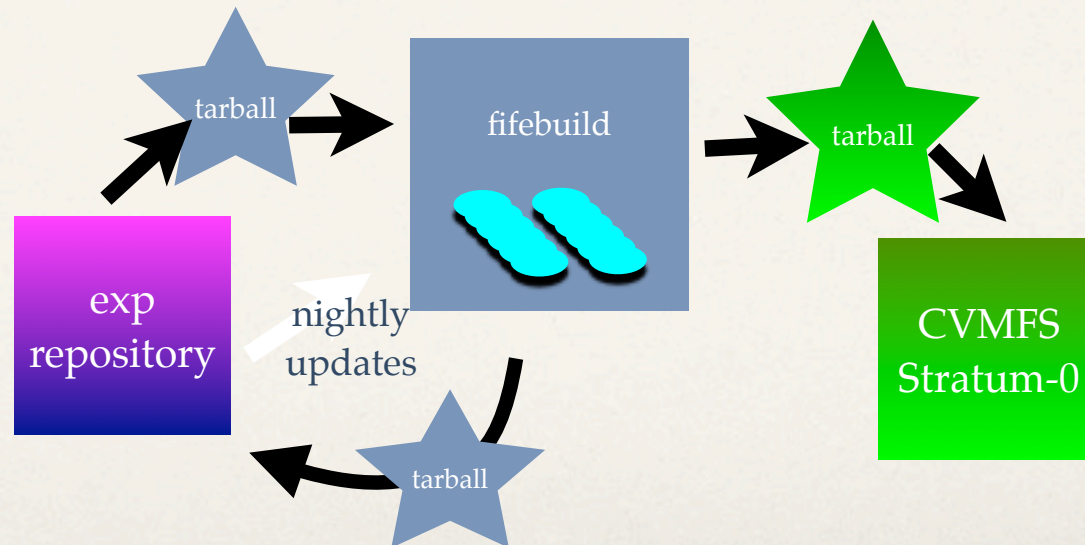
- ❖ well defined coding environment, designed to encourage best practices for algorithm development
- ❖ integrated data handling tools
- ❖ multiple path, provenance and metadata generation
- ❖ integration with generators, ROOT, Geant4
- ❖ active development with stakeholders
- ❖ integrated into the software environment of ArgoNeuT, g-2, $\mu 2e$, LBNE, MicroBooNE, NOvA, DarkSide-50



- ❖ talk by Adam Lyon

build and code distribution

- * nightly build of experiment software important - multiple flavors, debug/opt/prof, OSs
- * hardware and build configuration can have dramatic impact on build times - local disk and multithreaded
- * construct dedicated hardware with VM with scheduled build times for each experiment
- * considering integrating support for OS X
- * interface with CVMFS servers for code distribution
- * utilize relocatable UPS to distribute software and dependent packages
- * LArSoft transitioning to product distribution and CVMFS
- * OASIS server on the Open Science Grid
- * move away from the central disk servers using NFS



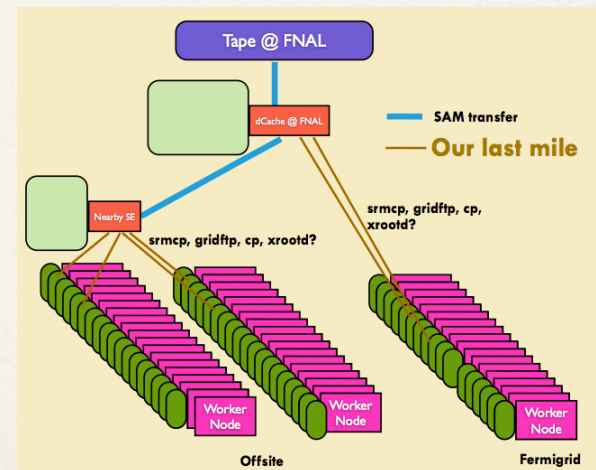
job submission tools

- ❖ where most of the heavy lifting for the experiments is being done
- ❖ transition to a new client-server submission system
- ❖ enable user code distribution through tarballs
- ❖ automatically interface with SAM, ifdhc, and OASIS
- ❖ generate DAG and workflows through dagNabbit.py
- ❖ focus on maintaining a consistent interface for the user, while potentially modifying the underlying solution
- ❖ utilize GlideinWMS (both FNAL and OSG factories) to access resource on-site and opportunistic off-site
- ❖ integrate new sites and resources quickly in an optimized manner
- ❖ currently used by 10 experiments
- ❖ poster by Dennis Box

Data Handling

- ❖ addressing the problem through three major services- file catalog, storage elements, and transfer clients
- ❖ SAMWeb project-Robert Illingworth talk
 - ❖ file catalog based on metadata and file location
 - ❖ web access to remove need for client
 - ❖ integrated with transfers to permanent storage
- ❖ storage element - transition away from NFS mounted storage to dCache implementation

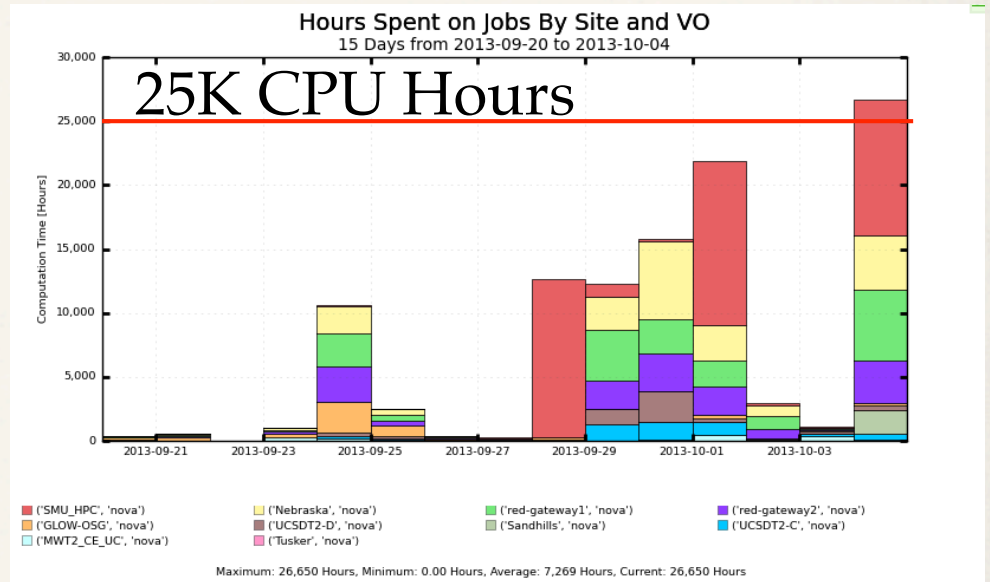
- ❖ IF Data Handling Client - unified package of file transfer tools for grid jobs
 - ❖ SRM, gridftp, xrootd interfaces
 - ❖ handle SAMGrid projects
 - ❖ focus on new dCache storage elements



- ❖ talk by Adam Lyon

integration successes so far

- ❖ NOvA experiment has successfully moved MC Generation to OSG
- ❖ integrated 6 sites into operations UNebraska, Southern Methodist University, Wisconsin, Harvard, UCSD.
- ❖ more than 10k jobs processed producing more than 1M simulated events
- ❖ Monitoring is an important part of assuring success of resource utilization



- ❖ using same submission command from user perspective, launched jobs “on-demand” at FermiCloud to provide proof-of-principle for Cloud bursting grid jobs

Plans for improvements

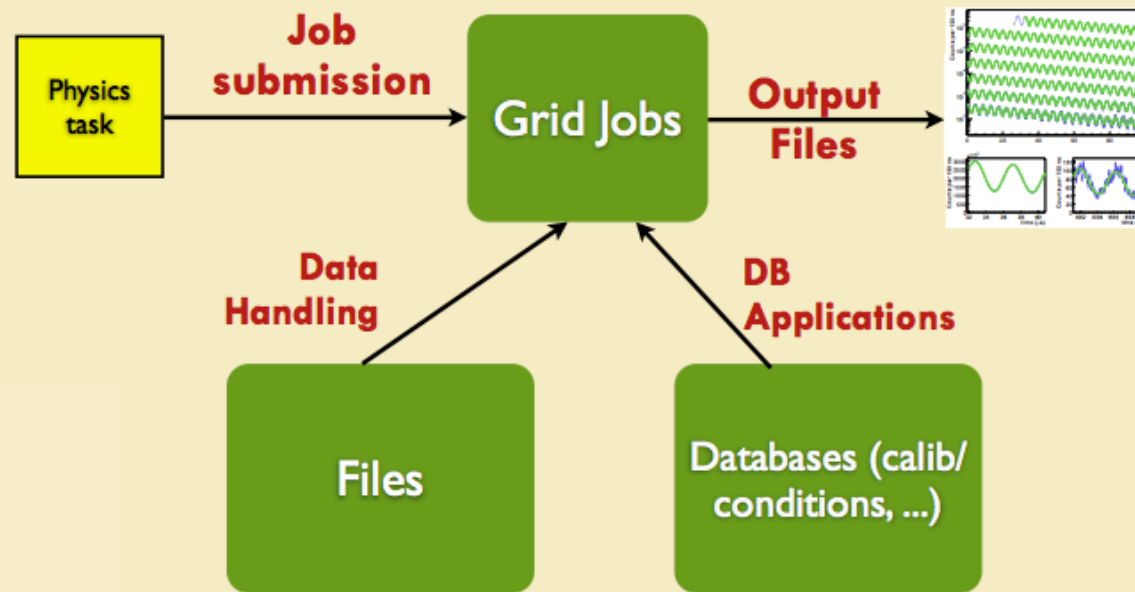

- ❖ FIFE architecture is currently undergoing re-evaluation and re-architecture process
- ❖ job submission infrastructure modify to client-server model
- ❖ local storage element making transition to shared dCache pools
- ❖ data handling project continues to integrate new resources without any change in user interface
- ❖ starting new integration push for several experiments in both the Intensity Frontier and Cosmic Frontier

Conclusions

- ❖ Fermilab physics program is extremely active at all three frontiers
- ❖ large scale computing needs for new “small” experiments encourages centralized services designed to ease integration to OSG and new resources
- ❖ FIFE project has developed a strategy and architecture to achieve this
- ❖ currently integrated OSG processing into the NOvA experiment
- ❖ from lessons learned from NOvA, working with other experiments
- ❖ looking forward to bring more resources and techniques to enable increased scientific output from frontier experiments

Conclusions

- ❖ Fermilab
- ❖ large scale centralized resources
- ❖ FIFE project
- ❖ currently
- ❖ from less
- ❖ looking for increased

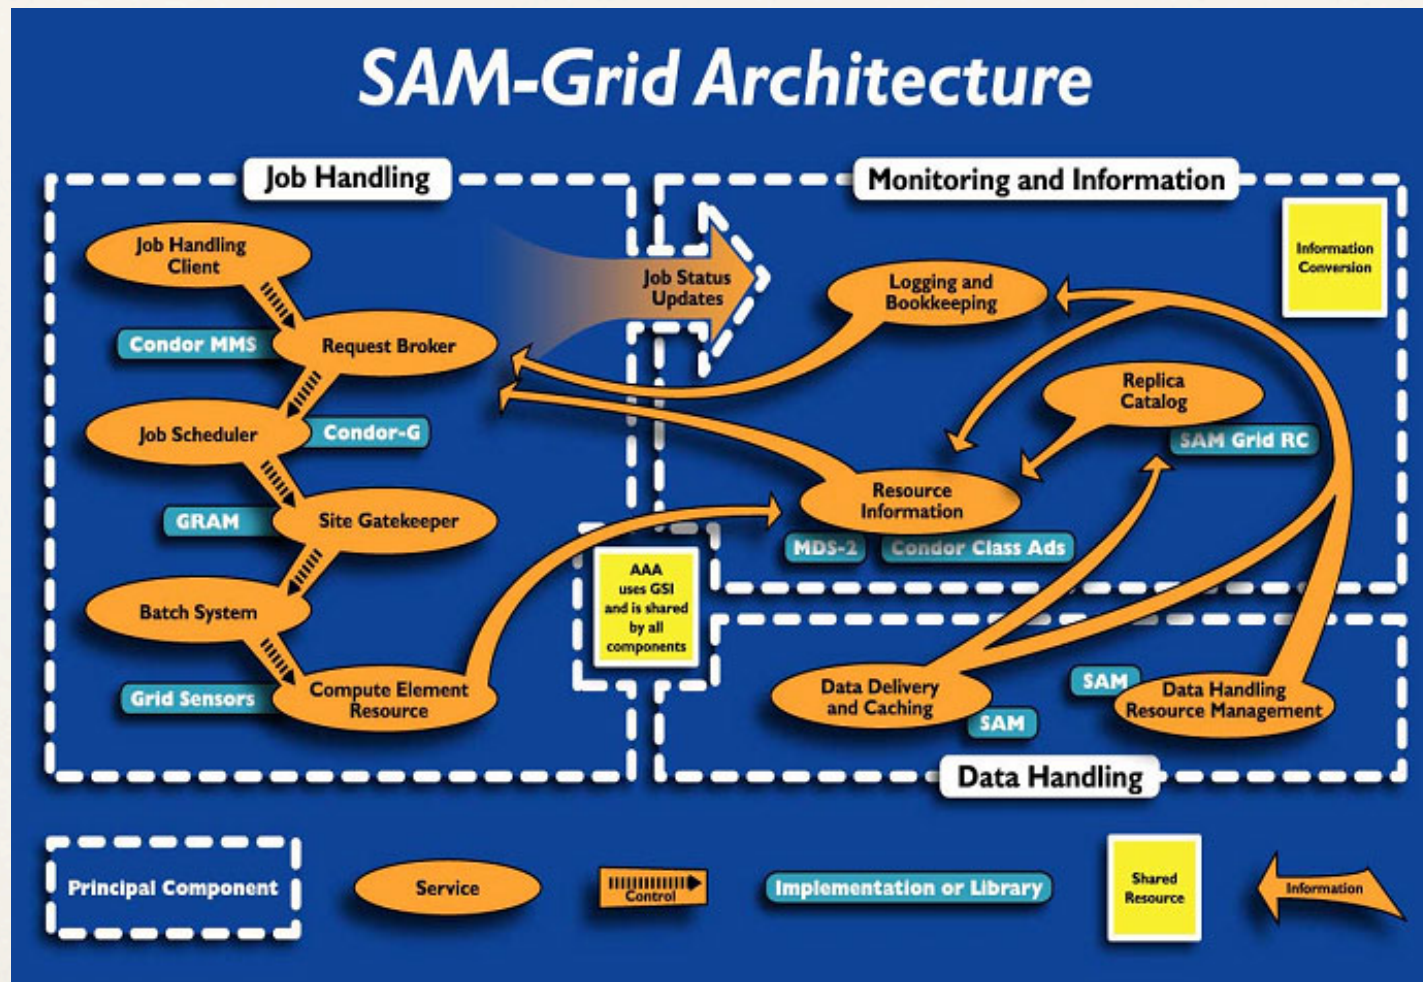


Adam Lyon

ers
urages
ew
ve this
at
ents
ble

backup slides

SAM Grid Diagram



GlideinWMS

- ❖ Decided to try and reduce the number of GlideinWMS factories being maintained on site
- ❖ reduce from the current 8 factories to one service (2 redundant factories)
- ❖ also transition to use OSG factory to submit Glideins to offsite computing elements
- ❖ configure the new factory service, then transition all of the frontends one at a time to new factory
- ❖ also configure the onsite factory to be able to submit offsite, taking configuration from OSG factory and importing