CHEP2015
OKINAWA, japan

Computing at the Intensity Frontier

Craig Group

University of Virginia
and
Fermilab

Fermilab

Introduction

- The Intensity Frontier (IF) refers to a category of particle physics experiments using high-intensity beams.
- Neutrino experiments require high neutrino fluxes.
- High-intensity beams are also required for experiments looking for rare processes or making precision measurements.
- IF experiments:
  - face some unique challenges
  - vary in the scale of data output and throughput
  - have overlapping computing requirements
- I'll discuss the Fermilab strategy to centrally provide for the computing needs of multiple IF experiments.

Outline

- Intensity Frontier Experiments
- Scale of IF computing efforts
- Challenges and strategies
- Case studies
- Conclusions

Computing is Critical

Software and computing are needed at all steps from data collection to physics results:

- Theory and modeling
- Experiment and beamline design
- Trigger and DAQ
- Online monitoring
- Event reconstruction and processing
- Physics analysis
- …

Intensity Frontier Experiments

Several categories of IF experiments:

- Quark flavor physics
- Neutrinos
- Charged lepton processes
- New light weakly coupled particles
- Nucleons, nuclei, and atoms

Intensity Frontier Experiments

Several categories of IF experiments:

- Quark flavor physics
- Neutrinos
- Charged lepton processes
- New light weakly coupled particles
- Nucleons, nuclei, and atoms

Requirements more similar to Energy Frontier

(Belle II computing discussed yesterday, T. Hara)

Intensity Frontier Experiments

Several categories of IF experiments:

- Quark flavor physics
- Neutrinos
- Charged lepton processes
- New light weakly coupled particles
- Nucleons, nuclei, and atoms

Requirements more similar to Cosmic Frontier

Intensity Frontier Experiments

Several categories of IF experiments:

- Quark flavor physics
- Neutrinos
- Charged lepton processes
- New light weakly coupled particles
- Nucleons, nuclei, and atoms

Will focus on neutrinos and charged-lepton experiments.

Why are neutrinos interesting?

- The wide range of quark masses is puzzling
- Neutrino mass does not fit in the standard model
  - We do not understand neutrinos' mass…..other than that it has mass
- The Higgs boson discovery has brought flavor and mass issues to the forefront
- There are many areas where we need to improve our understanding of neutrinos:
  - their mass ordering
  - the origin of their masses -- why they are so small?
  - how many types of neutrinos exist?
  - their interactions (CP violation?)
  - relationship to matter-antimatter asymmetry in universe (leptogenesis) and structure of the universe

Fermilab Neutrino Experiments

MINERvA, MiniBooNE, and MicroBooNE are located at Fermilab to measure neutrino cross sections and study neutrino oscillations over a short baseline.

Image from arXiv:1503.01520:
A Proposal for a Three Detector Short-Baseline Neutrino Oscillatio Program in the Fermilab Booster Neutrino Beam

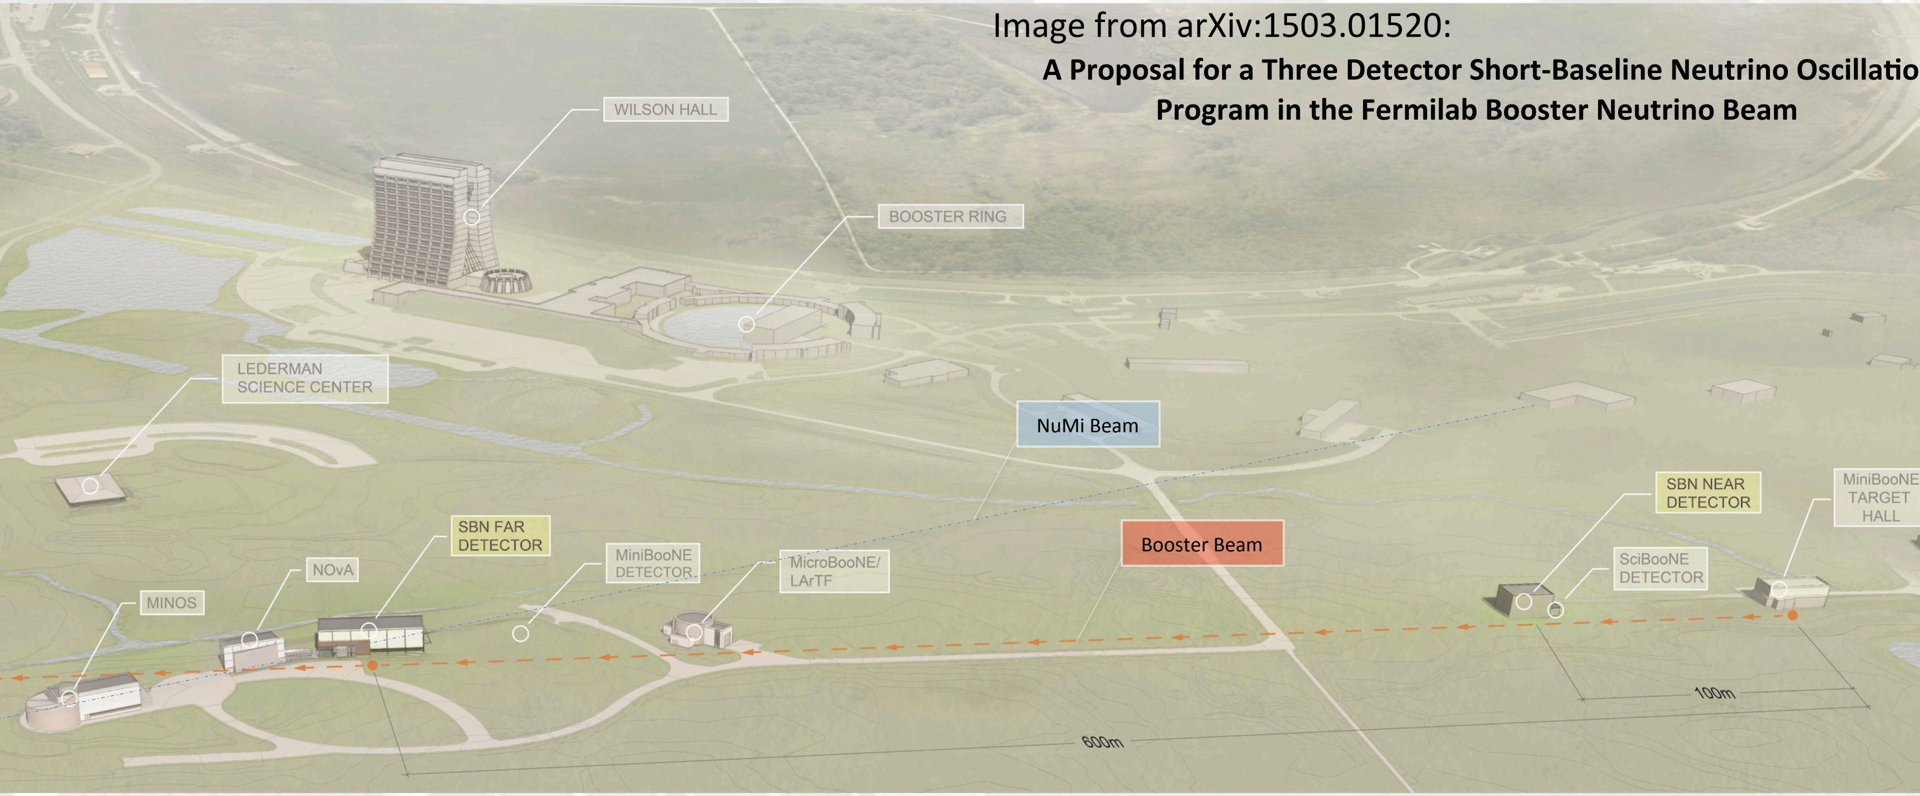

Long-baseline Neutrino Experiments

- MINOS and NOvA are long-baseline neutrino experiments
- Use Main Injector neutrino beam from Fermilab.
- Consist of "near" and "far" detectors.
- MINOS ongoing for 10+ years.
- NOvA finished commissioning last year. First results expected soon…
- These experiments primarily study neutrino oscillations.
- DUNE (formally LBNE) is a future long-baseline experiment which will use the LBNF beamline from Fermilab.

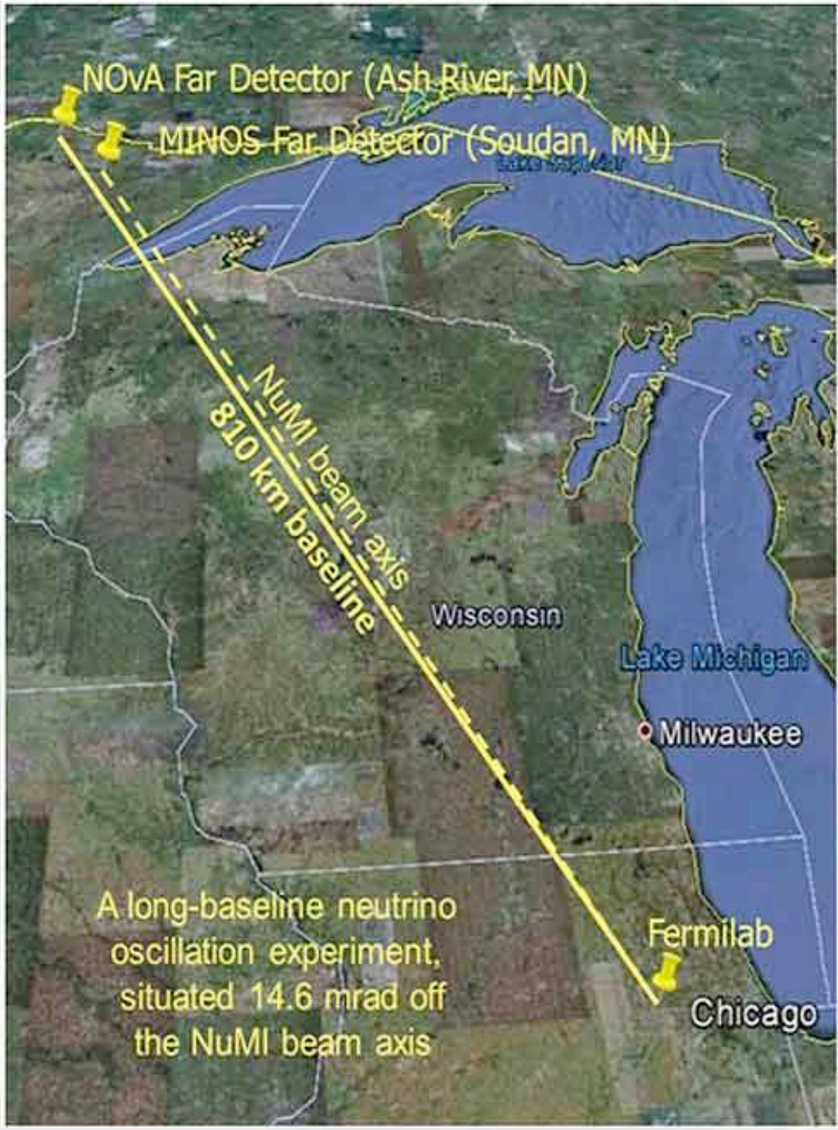

Why are muons interesting?

- One can probe the properties of the universe by making precise measurements or looking for extremely rare processes.
- Complementary alternative to using higher energies.
- Muons offer exciting prospects for discovery.
- For example: New physics can “run in the loops” and shift observables from standard model expectations – charge lepton flavor violation or the anomalous magnetic moment of the muon (g-2).

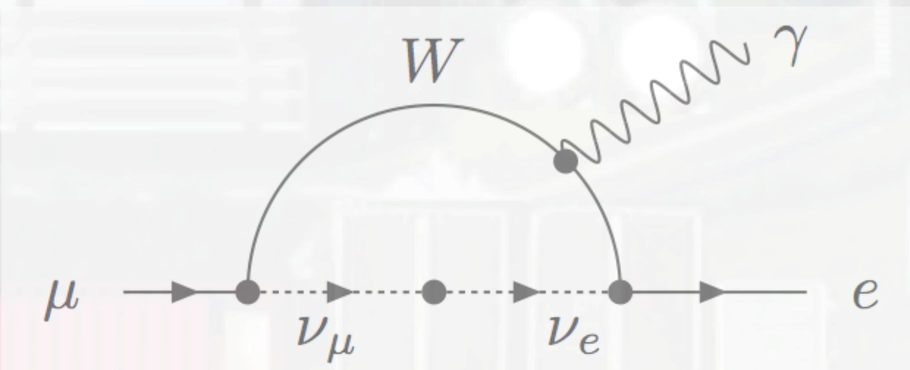

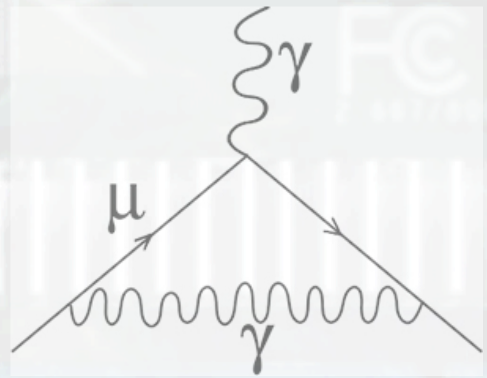

Muon Experiments

- Muon *g-2* probes TeV-scale (flavor-conserving) physics with the muon anomalous magnetic moment.
  → BNL E821 ~3 S.D. with theory. Goal to improve statistics by a factor of 20, requires ~10^{12} mu+.
- Mu2e probes effective mass scales up to 10^4 TeV range with charged lepton flavor violation.
  → Goal to improve sensitivity by 4 orders of magnitude. Requires ~10^{17} mu-.

While the g-2 and Mu2e experiments are quite different they share the muon beamline.

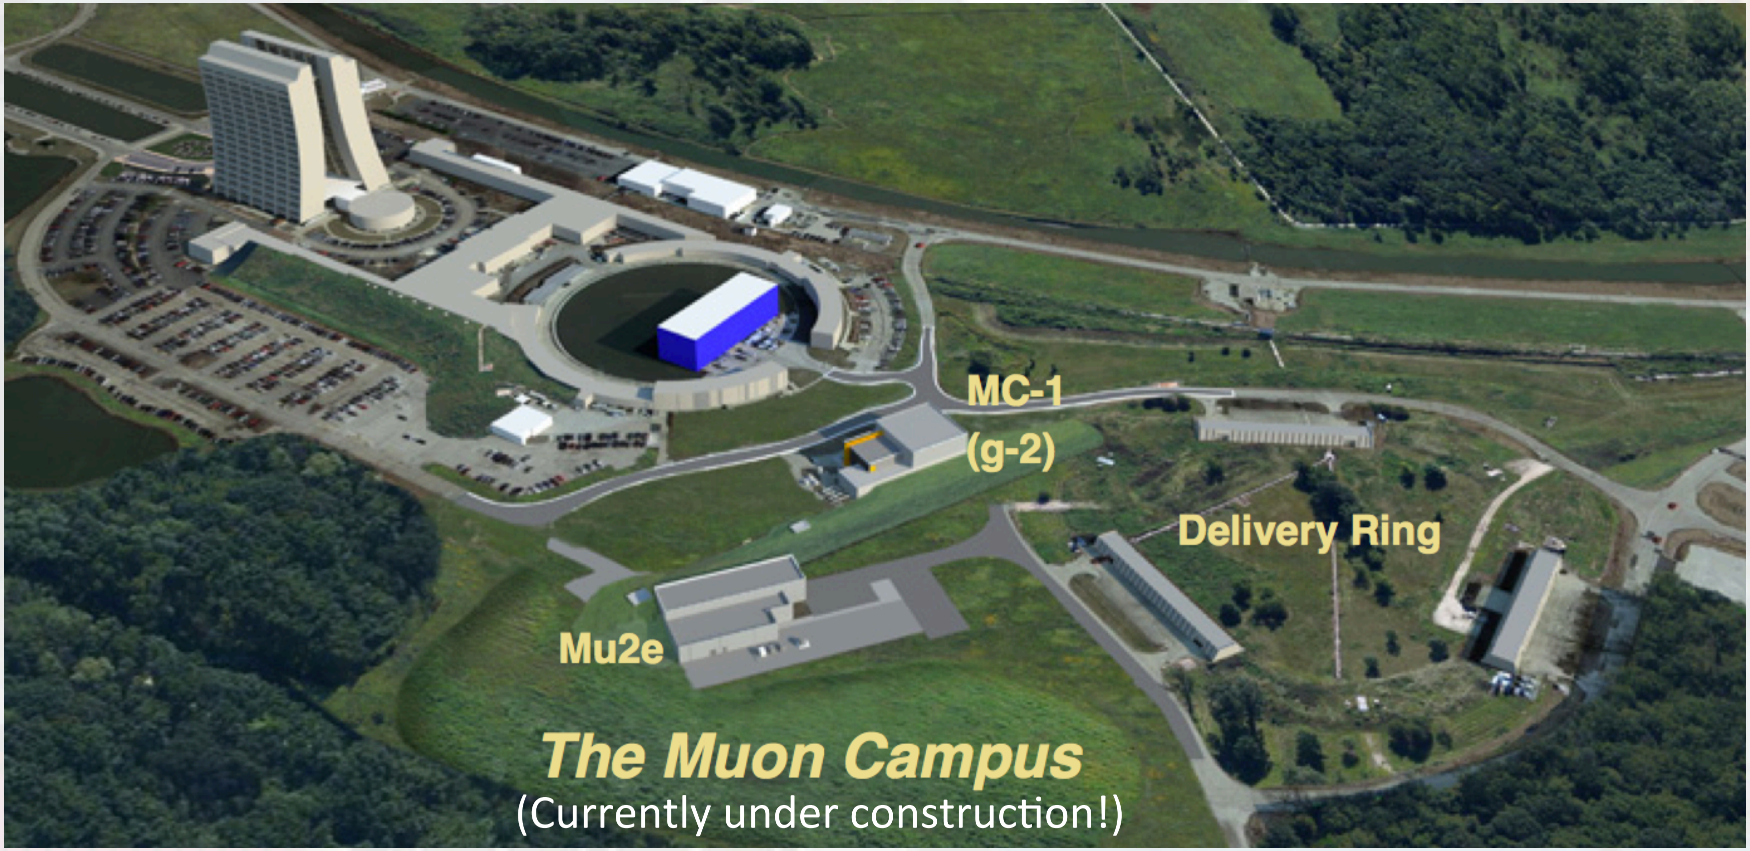

Small Experiments?

- Intensity Frontier experiments are not small!

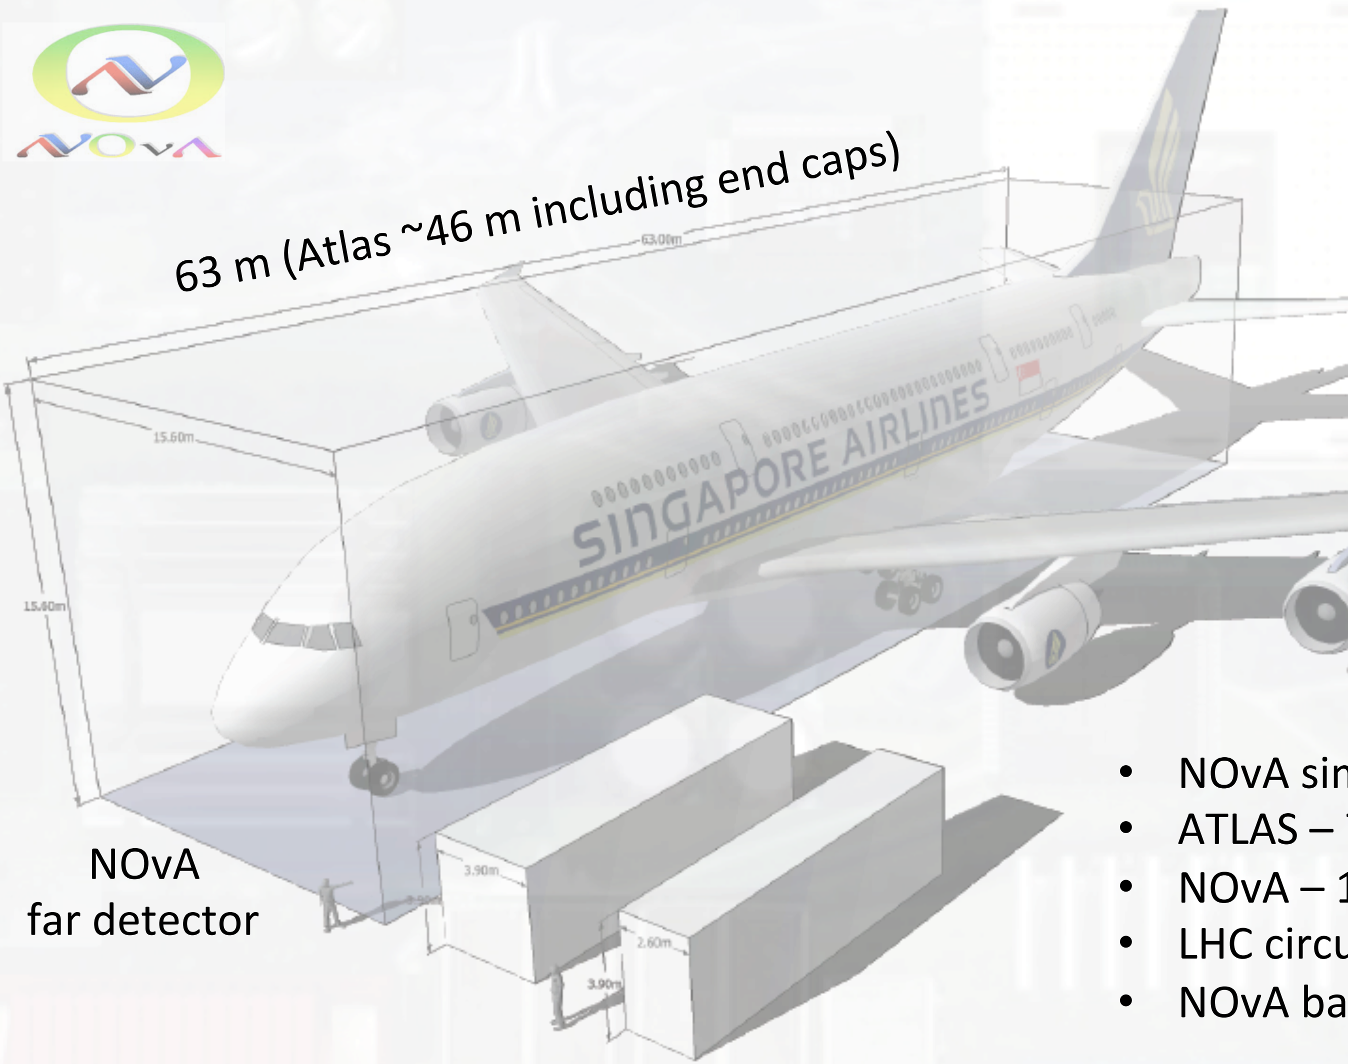

Airbus-380

- NOvA similar volume to Atlas
- ATLAS – 7,000 tonnes
- NOvA – 14,000 tonnes
- LHC circumference 27 km
- NOvA baseline 810 km

Intensity v/s Energy Frontier

CMS will use about 1 billion CPU hours this year.

→ FNAL IF experiments will use about 100 million

CMS used ~40 petabytes of tape in Run 1

→ So far, FNAL IF has used >5 petabytes

This is just the Fermilab IF program which represents just a fraction of the international program.

So, IF needs are smaller, but still large!

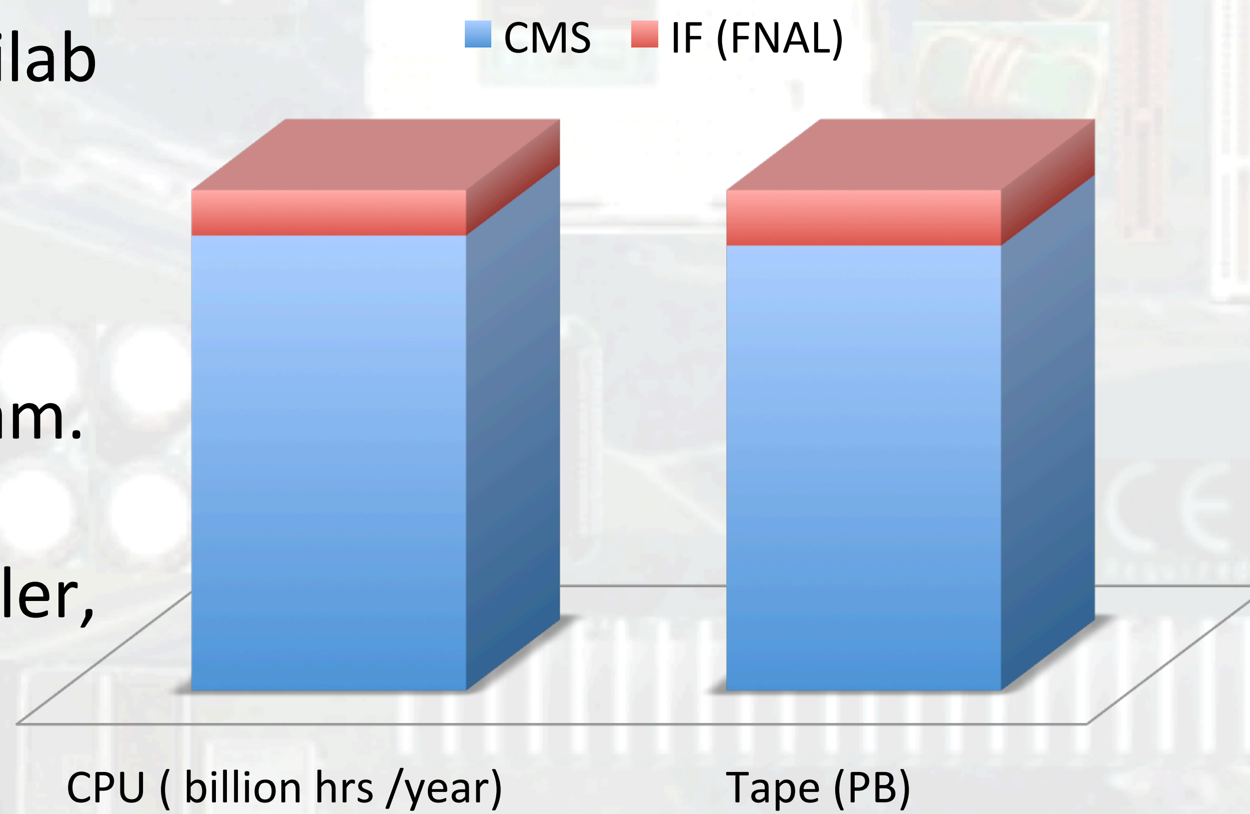

Limited Personnel Resources

~160 μBooNE

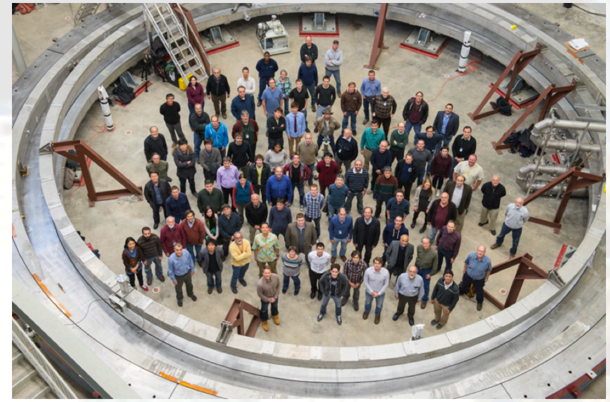

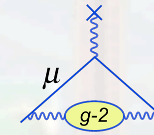

~155

>200 NOvA

Mu2e ~170

MINOS ~100

MiniBooNE ~60

MINERvA ~60

Limited Personnel Resources

- If you combine all of the IF experiments you get a scientific team that is a significant fraction of the size of an LHC team.
- Each experiment only has a few strong post docs and students with expert computing skills.
- Designing an independent analysis framework, data handling systems, and other computing tools for each experiment would not be an efficient use of personnel.
- Also, independent computing resources (CPU, storage) for each experiment at this scale is not efficient.
- Strategy: Use central tools and shared resources when possible.

IF Computing Survey

- For Snowmass 2013 there was a survey of computing for the Intensity Frontier.
- Found a high degree of commonality between computing models of IF experiments:

"Traditional event-driven analysis and Monte Carlo simulation using centralized data stores that are distributed to independent analysis jobs running in parallel on grid computing clusters"

- All experiments use ROOT and GEANT4

Snowmass Report: arXiv:1310.6964

Common Services and Projects Toolkit

FIFE (Fabric for Frontier Experiments):Manages/provides access to shared resources at Fermilab
Used by all neutrino and muon experiments at FNAL

FIFE provides access and support for a comprehensive set of services and tools:

- **DAQ and Controls**
- **Grid and Cloud**
- **Scientific Data Storage and Access**
- **Scientific Data Management**
- **Scientific Frameworks and Software**
- **Physics and detector simulation**
- **Databases**
- **Scientific Computing Systems**
- **Scientific Collaboration Tools**

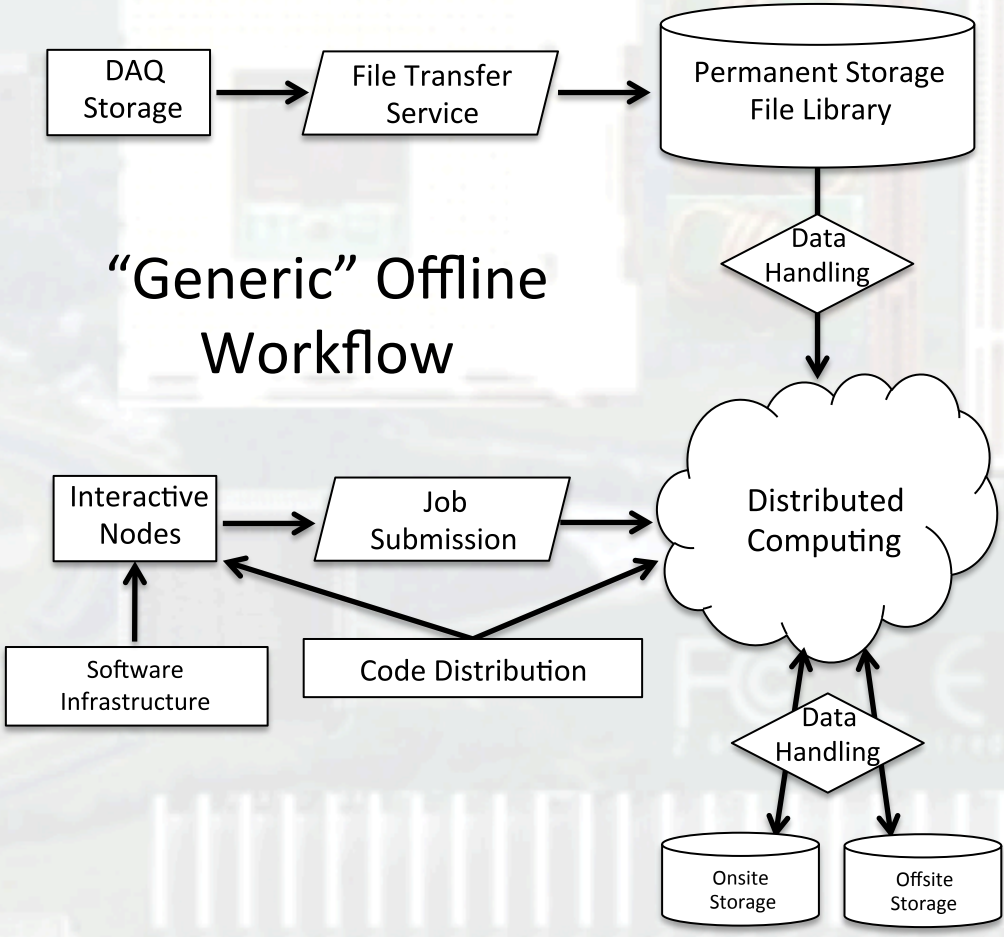

See talk on Thursday, Parag Mhashilkar, Track 4:
Advances in Distributed High Throughput Computing for the Fabric for Frontier Experiments Project at Fermilab

Slide content, O. Gutsche

Physics Software Requirements

- Science demands reproducibility.
  → We must have control over our software

- We want to work together.
  → Share ideas through code

- We want to do physics, not computing.
  → We just want to make plots! Somehow, that should be easy, sane, robust, and repeatable

Slide content, A. Lyon – ArtG4 Seminar 2013

Physics Software Requirements

- Science demands reproducibility.

  → Official results use version controlled software.

- We want to work together.

  → Code repositories; modular frameworks

- We want to do physics, not computing.

  → Infrastructure in a framework + an easy build system

See poster session A:
Software Management for the NOvA Experiment

Slide content, A. Lyon – ArtG4 Seminar 2013

What is a framework?

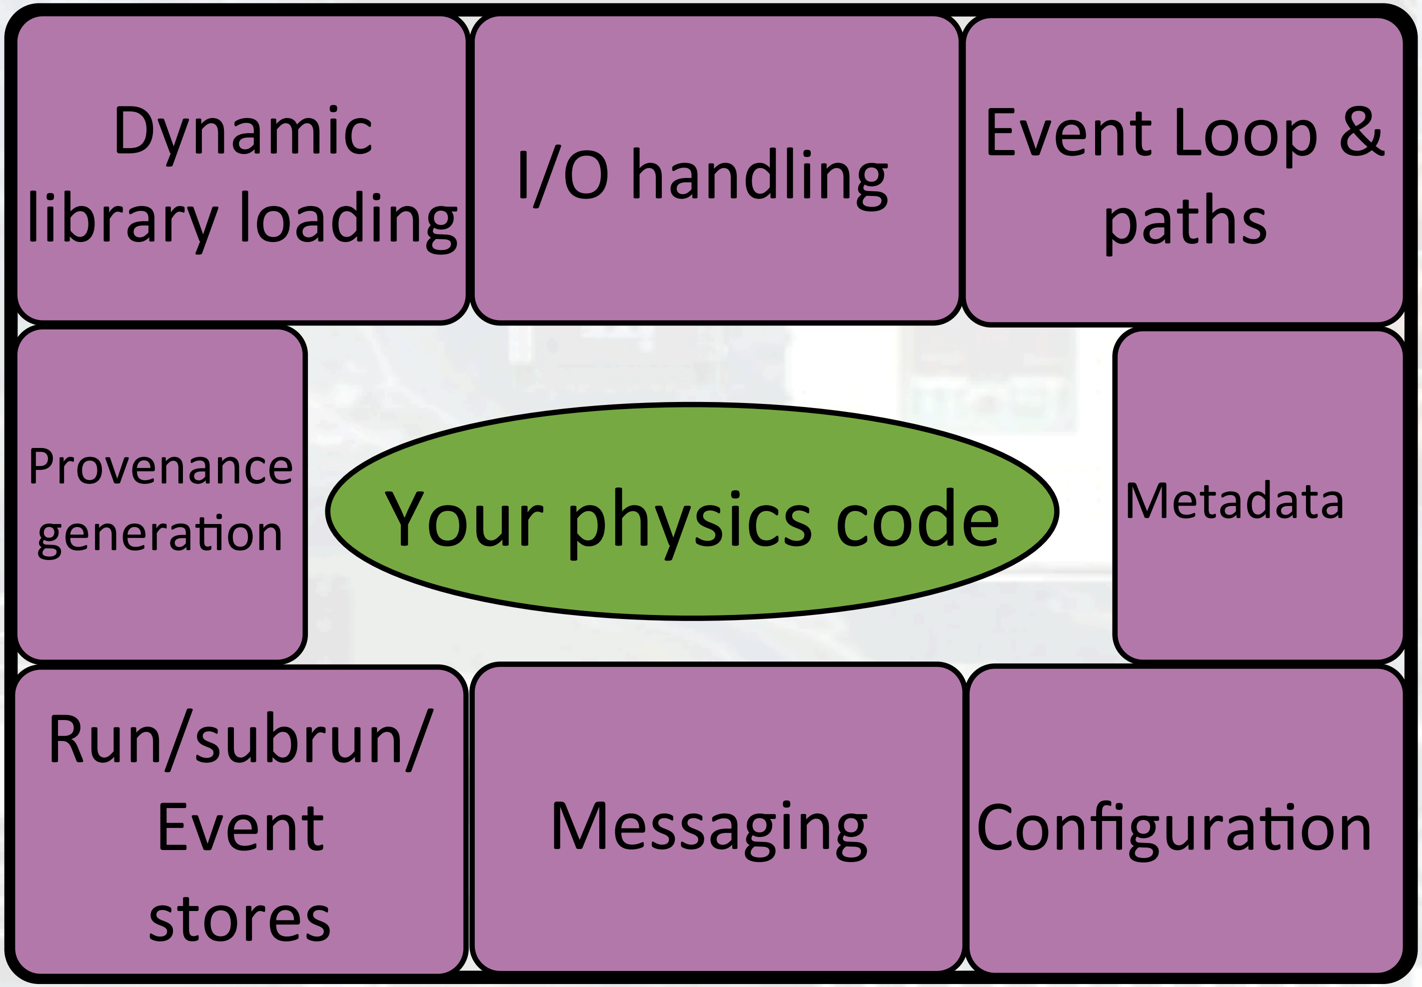

Code you write

Code you use from the framework

See slides from yesterday, Liz Sexton Kennedy, track 5:

A Comparative Analysis of Event Processing Frameworks used in HEP

Slide content, A. Lyon – ArtG4 Seminar 2013

Framework Benefits

Allows you to write your physics code without worrying about the infrastructure. Makes it easy to work with others.

→ But not for free – you have to learn how to use the framework!

Some people find such a system constraining:

- Infrastructure is hidden behind the scenes from you
- Your ideas may not be included
- You have to trust a system you didn't write
- You miss out on the fun of writing super-cool complicated C++ code

Some people find such a system liberating:

- You can concentrate on physics code
- Your C++ is pretty easy (you are *using* a complicated system, not *writing* it)
- You get to miss out having to maintain the complicated C++ code
- You can use code from others and share yours with others
- You can get services for free (e.g. data handling)

Slide content, A. Lyon – ArtG4 Seminar 2013

Why not write your own framework?

"Small" experiments may not have…

- The expertise
  - Writing large C++ systems is hard (need low dependences, efficient generic programming, follow software engineering best practices)
- The time
  - With lots of milestones and reviews, there's no time to devote to correctly writing such a large system
- The energy
  - We just wanna make plots! Not write infrastructure code

Slide content, A. Lyon – ArtG4 Seminar 2013

The art Framework

- art is a framework forked from CMSSW and tailored for neutrino/muon experiments.
- Experiments use art as-is as an external. art **is not** modified by the individual experiments. Used like ROOT, G4, boost...
- Used by Mu2e, Muon g-2, NOvA, MicroBooNE, LAriAT, Darkside-50, and DUNE prototype efforts.
  - Easier for individuals to work on multiple experiments -- this is a common practice for neutrino experiments
  - Experiments can share solutions to common problems
  - Common training (classes, workshops, workbook…)
- Seamless integration to data-handling tools (I/O) is an important aspect of art.
- New features and direction decided among stakeholders by consensus.
- Main support forum is centralized, but experts from experiments help answer questions.

Examples: building on art

- artdaq: art-based toolkit for creating DAQ systems
  - Provides common reusable components.
  - Based on event streaming architecture with software event filtering.
  - Integrated with art framework, so offline modules can run online
  - Used by Darkside-50, Mu2e, and LArIAT
  - NOvA and MicroBooNE also use some artdaq components.

See poster session A, track 1:
Recent developments in the infrastructure and use of artdaq.

- LarSoft: art-based framework for liquid argon experiments
  - There have been a series of liquid argon experiments and R&D efforts:

  Small tests → ArgoNeuT→ MicroBooNE → LBNF (35t) → … → DUNE
  - LarSoft provides simulation, data reconstruction, and analysis tools for LArTPC experiments
  - Community collectively contributes to software

Neutrino Simulations

1) Beamline (FLUKA/Geant)
+ Produces a flux prediction
+ Hadron production, focusing, etc.

A Three-Part Software Stack

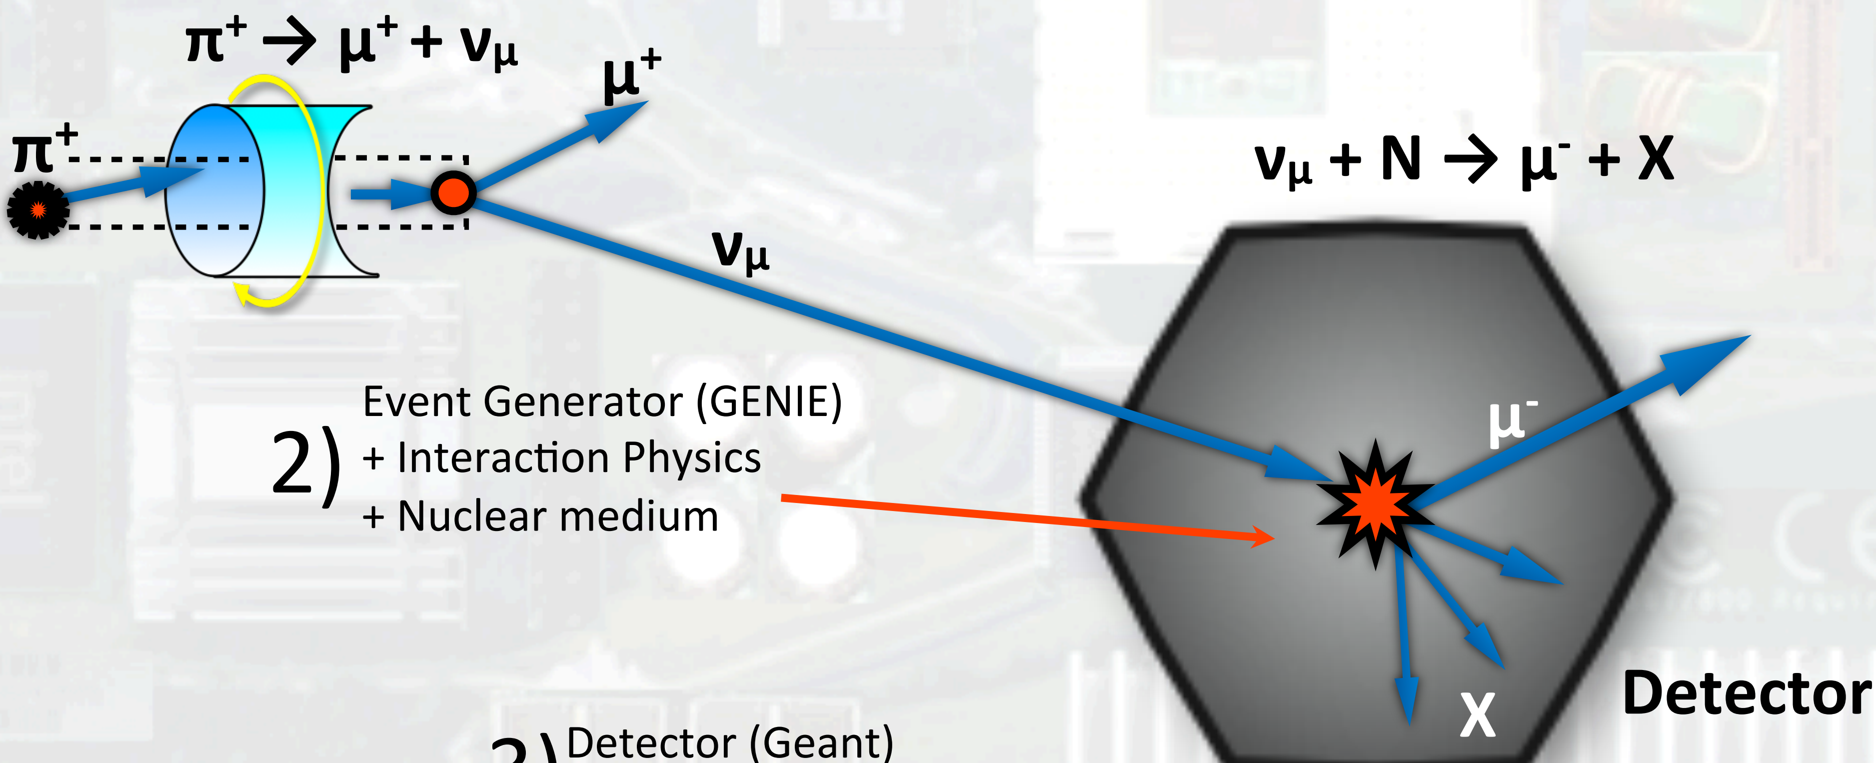

2) Event Generator (GENIE)
+ Interaction Physics
+ Nuclear medium

3) Detector (Geant)
+ Final state radiation traversing matter

Slide content, G. Perdue

CHEP2015
OKINAWA, Japan

GENIE

- **G**enerates **E**vents for **N**eutrino **I**nteraction **E**xperiments.
- http://genie.hepforge.org
- Well-engineered C++ software framework built on sound OO-principles and design patterns.
- Propagates a flux of neutrinos (specified by function, histogram, or ntuple) through a geometry (Geant4-compatible) and simulates the initial interaction and propagation of hard vertex products through the nuclear medium.
- Geant4 takes over when particles leave the nucleus.
- ROOT provides many core utilities. GENIE also heavily leverages other HEP software - LHAPDF, Pythia, etc.

Andreopoulos, C. and Bell, A. and Bhattacharya, D. and Cavanna, F. and Dobson, J. and others. "The GENIE Neutrino Monte Carlo Generator". Nucl.Instrum.Meth. A614. 87-104. 2010.

Slide content, G. Perdue

GENIE @ FNAL

- GENIE is primary event generator for:
  - ArgoNeut
  - LAr1-ND (SBND)
  - DUNE
  - MicroBooNE
  - MINERvA
  - NOvA
- GENIE is being considered for special studies by MINOS and MiniBooNE (they use previous generation software for their main generators).

Slide content, G. Perdue

Other Common IF Software

- FLUKA – used to simulate the production of hadrons in beam-line simulations. Important for neutrino production targets.
- CRY – used for simulating cosmic rays. Critical for both neutrino and muon experiments.
- GLoBES – used by long baseline neutrino experiments to study expected physics reach given beam spectrum and detector efficiencies
- …

Data handling

Your files reside here

Tape Robot

Disk Storage (cache)

Not here, where you need them

Compute Farm of worker nodes

Must identify the files to process

Must move those files to the right place

Must associate those files with jobs and process them

All this must be reliable, scalable, and efficient

Slide content, A. Lyon – ISGC 2014

SAM

- Sequential data Access via Metadata
  (note: not *Service Availability Monitoring!)*
- Forerunner of LHC data management (CDF and D0)
- Database with per-file metadata catalog (file name, size, events, run info, luminosity, MC details, ...)
- Dataset creation and queries
- Coordinates and manages data movement to jobs (gridftp, dccp, SRM, and soon XRootD)
- Cache management (now using dCache)
- File consumption and success tracking, recovery

See poster session B:
Data Handling with SAM and ART at the NOvA Experiment

See talk this afternoon, Andrew Norman, track 5:
Large Scale Management of Physicist's Personal Analysis Data

Slide content, A. Lyon – ISGC 2014

File Storage

- BlueArc
  - High performance, high cost (~$1000/TB)
  - POSIX compatible (all I/O operations) on local systems that mount via NFS.
  - Not directly accessible off-site; Can be overloaded
  - Easiest to use for development, local analysis
  - Allocations per experiment; nearly always full
- dCache
  - Highly distributed storage with central name space
  - Much lower cost (~$100/TB), ~4PB shared by IF experiments
  - Read / Write interfaces, but does not look like usual file systems
  - Accessible from off-site
  - A cache (optionally front-end to tape system) -- old files are flushed
- Tape
  - Current complex has four 10,000-slot tape libraries available for IF
  - Recent tape technology (> 8 TB per) provides ~320 PB of capacity
  - Cost is only an issue when speaking PB size volumes

Fermigrid

- A physical "unit" is one core with 2 GB physical memory and associated local disk space
  - Typical machines are 64 cores, 128 GB (64 units)
- ~15,000 general-purpose cores on Fermigrid
- More available opportunistically from CMS and Dzero farms.
- Several IF experiments now running on the OSG:
  - NOvA used ~3M CPU hours on OSG last year
  - Mu2e launching large campaign for this year - recently ~6000 cores with a goal of ~10M offsite CPU hours (following the NOvA experience)
- Recent success on Amazon Cloud (AWS)
- CernVM-FS used to distribute code.
- Introducing experiments to new resources is streamlined due to common tools.

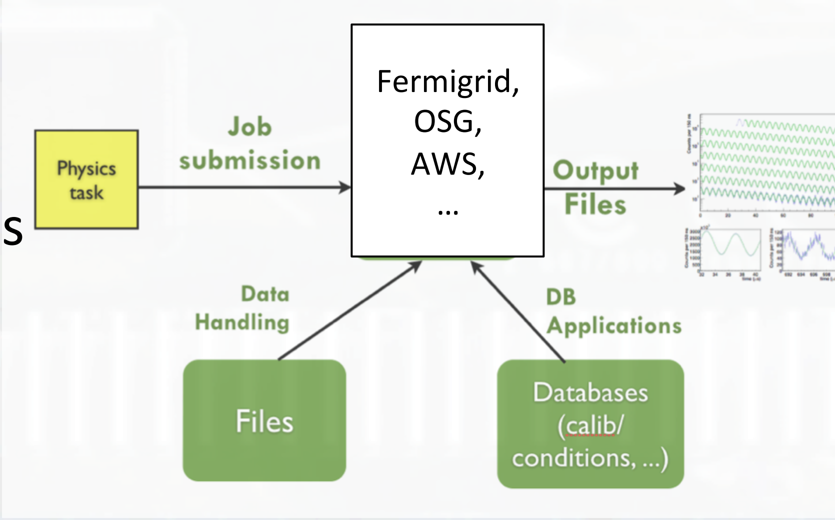

See poster session A:
Large Scale Monte Carlo Simulation of neutrino interactions using the Open Science Grid and Commercial Clouds

Example: NOvA

More than 5M files in SAM, >1.5 PB/year to tape, about 15 M CPU hours last year (20% offsite)
Experiment-agnostic "production team" heavily used

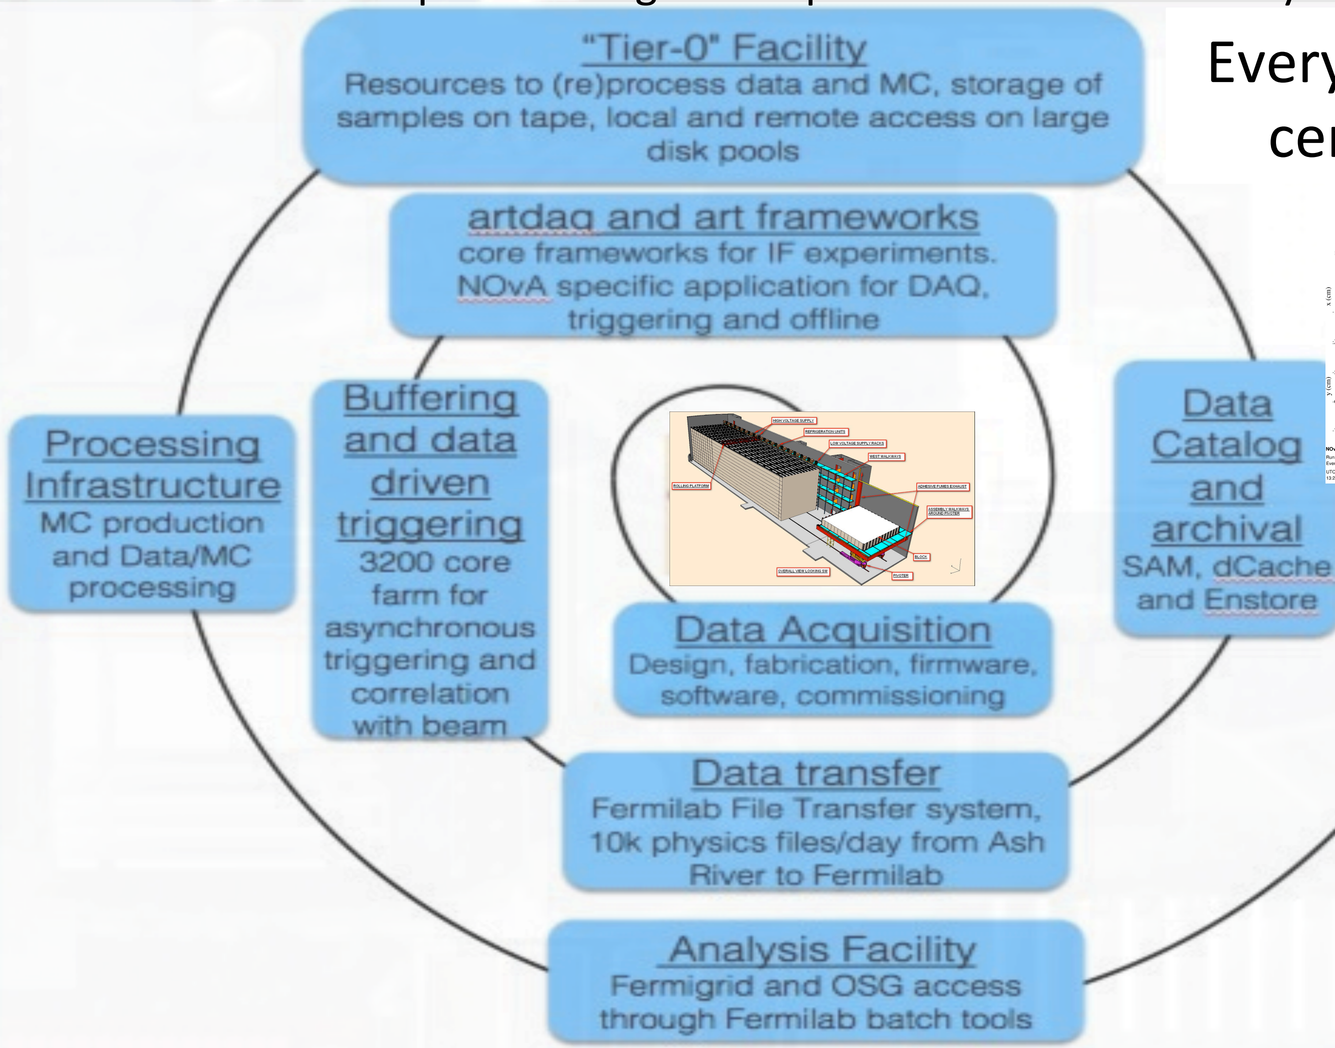

Every step utilizes central tools.

Neutrino in far detector!

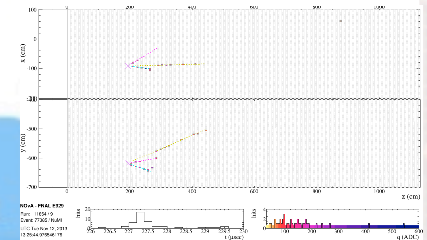

See slides from yesterday, Alec Habig, track 5:
Recent Evolution of the Offline Computing Model of the NOvA Experiment

Example: Mu2e

- Uses art and plans to use artdaq
- Uses a fork of the BaBar track-fitting code
- Uses FIFE tools for data handling and batch jobs

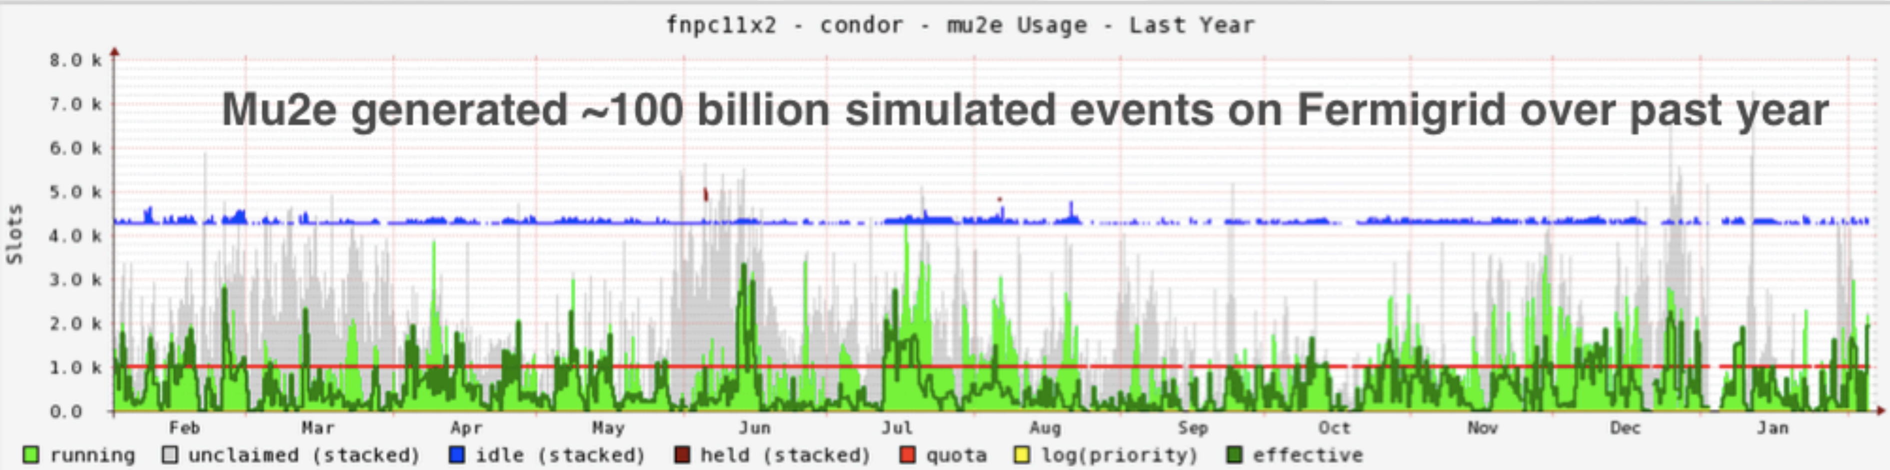

By utilizing existing tools, Mu2e successfully built their computing infrastructure with only a small fraction of the ~170-member collaboration

(less than 10 physicist FTE/year)

Challenge: Experiments in Transition

- Shared services model to ensure that all experiments, small and large, are able to make use of the facilities from **"cradle to grave" .**
- Always serving experiments in each stage.
- Different priorities: Early phase – rapid development; Late phase – stability).
- Different personnel: Students and post docs often not available in early stage.

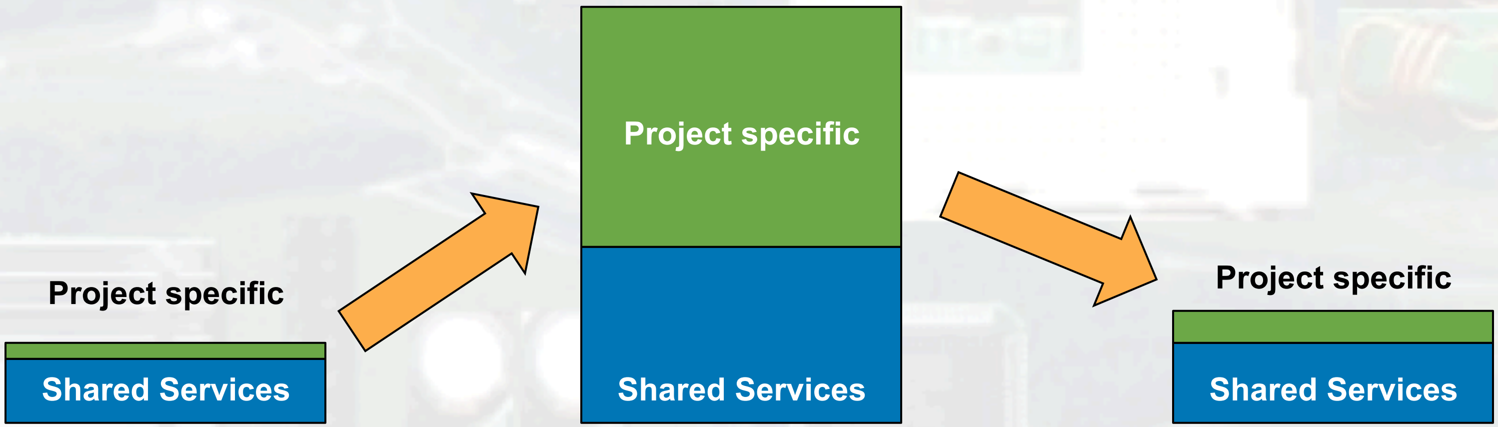

Early Period	Mature phase	Final data-taking and beyond
R&D, Ideas, Simulations LOI, TDR, Proposals	Construction, operations, analysis	Final analysis, data preservation and access

Challenge: Experiments in Transition

- Shared services model to ensure that all experiments, small and large, are able to make use of the facilities from **"cradle to grave" .**
- Always serving experiments in each stage.
- Different priorities: Early phase – rapid development; Late phase – stability).
- Different personnel: Students and post docs often not available in early stage.

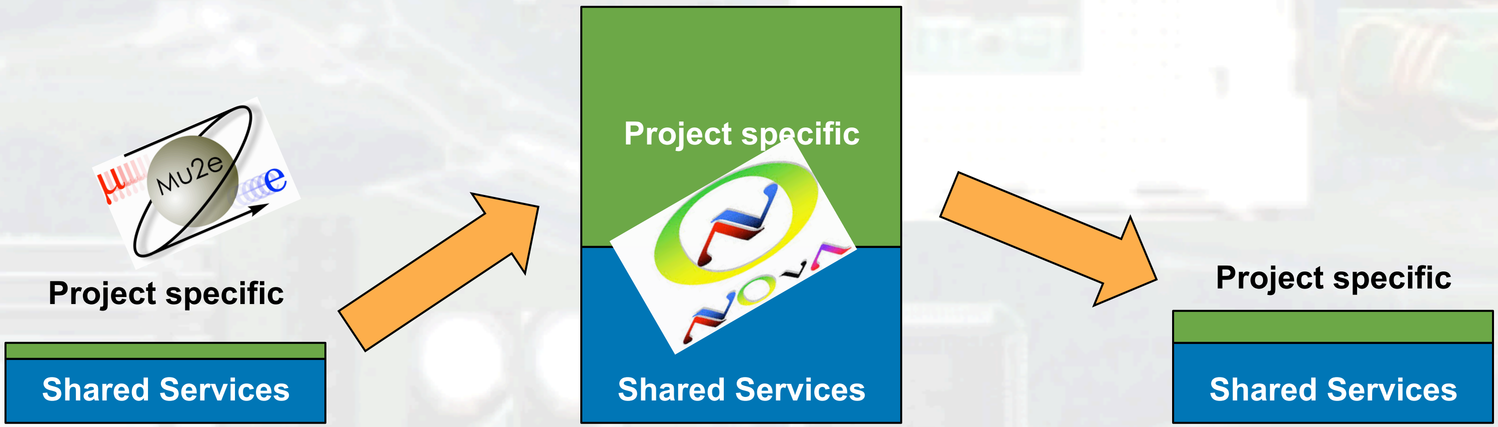

Early Period	Mature phase	Final data-taking and beyond
R&D, Ideas, Simulations LOI, TDR, Proposals	Construction, operations, analysis	Final analysis, data preservation and access

DUNE

- Deep Underground Neutrino Experiment will be of comparable size on its own to an LHC project.
- The current 35-ton prototype effort is using LarSoft and other FIFE tools and working efficiently in that system.
- DUNE may fit well into the Fermilab computing paradigm.
- However, DUNE may be of critical mass to drive computing tools and dictate rules for dedicated resources. (More like the LHC experiment model)
- Software and Computing challenges in the year 2025 and beyond are very large. Even global experiments (LHC, DUNE, …) may need to coordinate to meet them.

The Future of Jobs and Slots…

- Job requirements are changing
  - We now see IF jobs requiring > 2 GB memory or, utilizing multiple cores
  - Batch system evolving to allow jobs to request partitionable batch slots
  - Can ask for multiple processors, more memory
- Preparing for different HW architectures to appear
  - Job may need access to many-core systems: GPU, etc...
  - Small test cluster with GPUs, Intel Phi – but nothing for production yet
- Evolution of how we manage resource requests and accounting
  - Move from "slot" concept to "CPU-hours"
  - Some multi-unit resource requests will need multiplicative factor
- Goal: Expand resources transparently to commercial clouds.

Outlook

- As a whole, Intensity Frontier experiments have immense computing needs. Smaller, but significant relative to an LHC experiment.
- Providing computational resources efficiently to a diverse set of experiments is challenging.
- Centrally-managed services with support are available to (and heavily used by) Intensity Frontier experiments at Fermilab.
- Some tradeoffs:
  - Less control over the tools by a specific experiment.
  - But, less experimental effort to write and maintain tools translate to more effort for physics.
- Intensity Frontier physics is benefiting from this tool set.

Coming soon: neutrino oscillations, anomalous magnetic moment of the muon, charged lepton flavor violation, and hopefully some surprises!


Thank you very much!

Thanks to the CHEP organizers for the invitation!
Acknowledgements: A. Lyon, O. Gutsche, A. Norman, R. Kutschke, G. Perdue, M. Kirby, L. Sexton-Kennedy,S. Fuess, and others…

Interacting with art

<u>Types of MODULES:</u>
(All modules can read data from the event)

o Input source:
A source for data. E.g. a ROOT file, info from DAQ, or empty for start of simulated data

o Producers:
Create new event data from scratch or by running algorithms on existing data

o Filters:
Like producers, but can stop running of downstream modules

o Analyzers:
Cannot save to event. For, e.g. diagnostics plots

o Output module:
Writes data to output file (ROOT). Can specify conditions and have many files

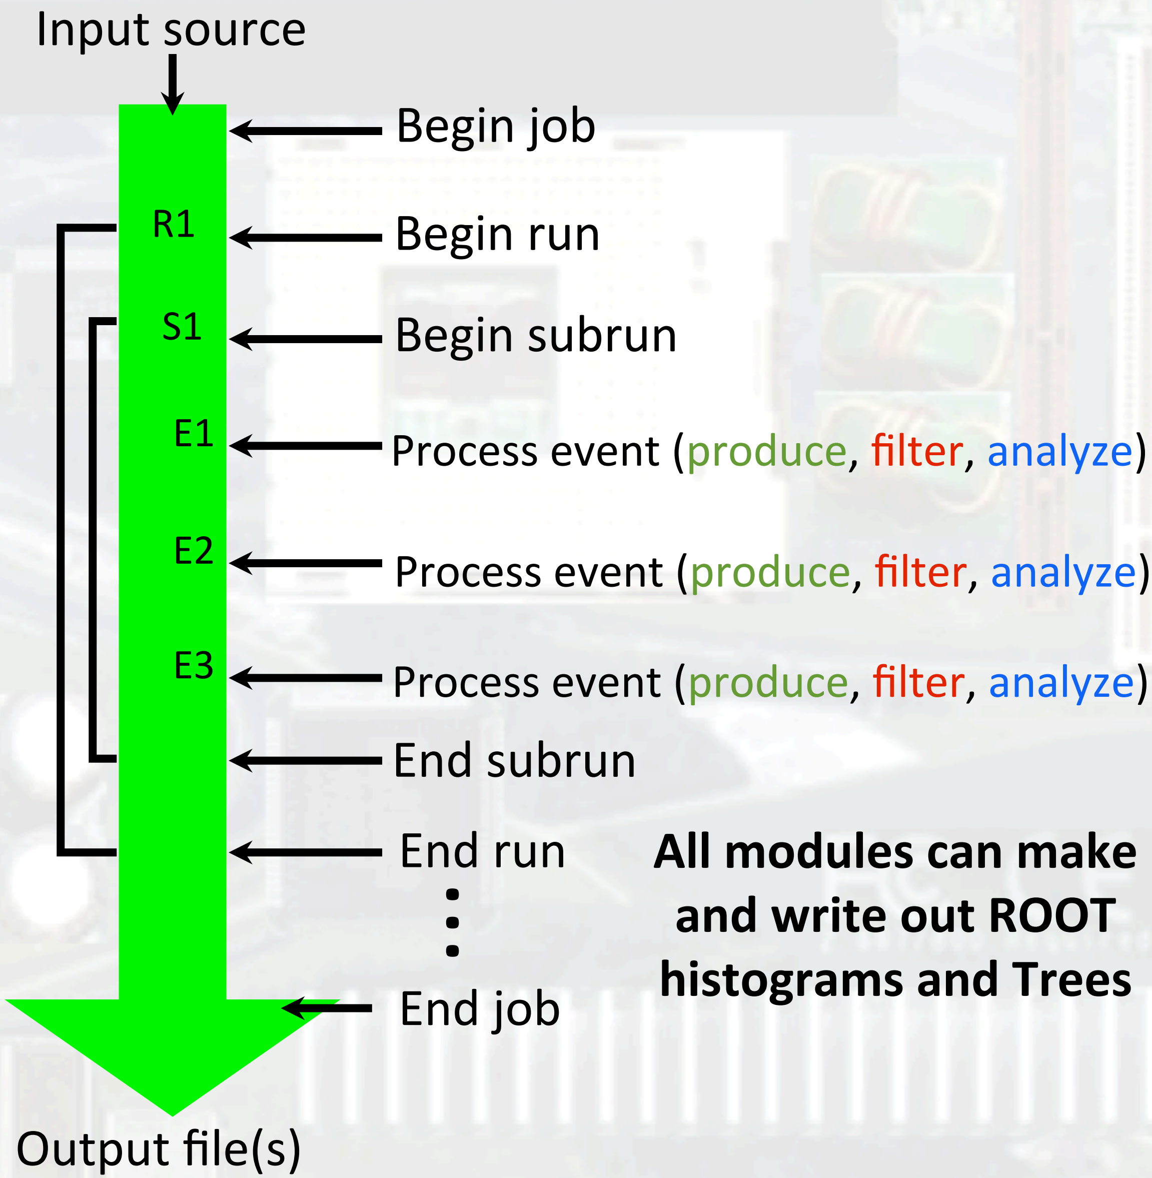

All modules can make and write out ROOT histograms and Trees

Slide content, A. Lyon – ArtG4 Seminar 2013

More art...

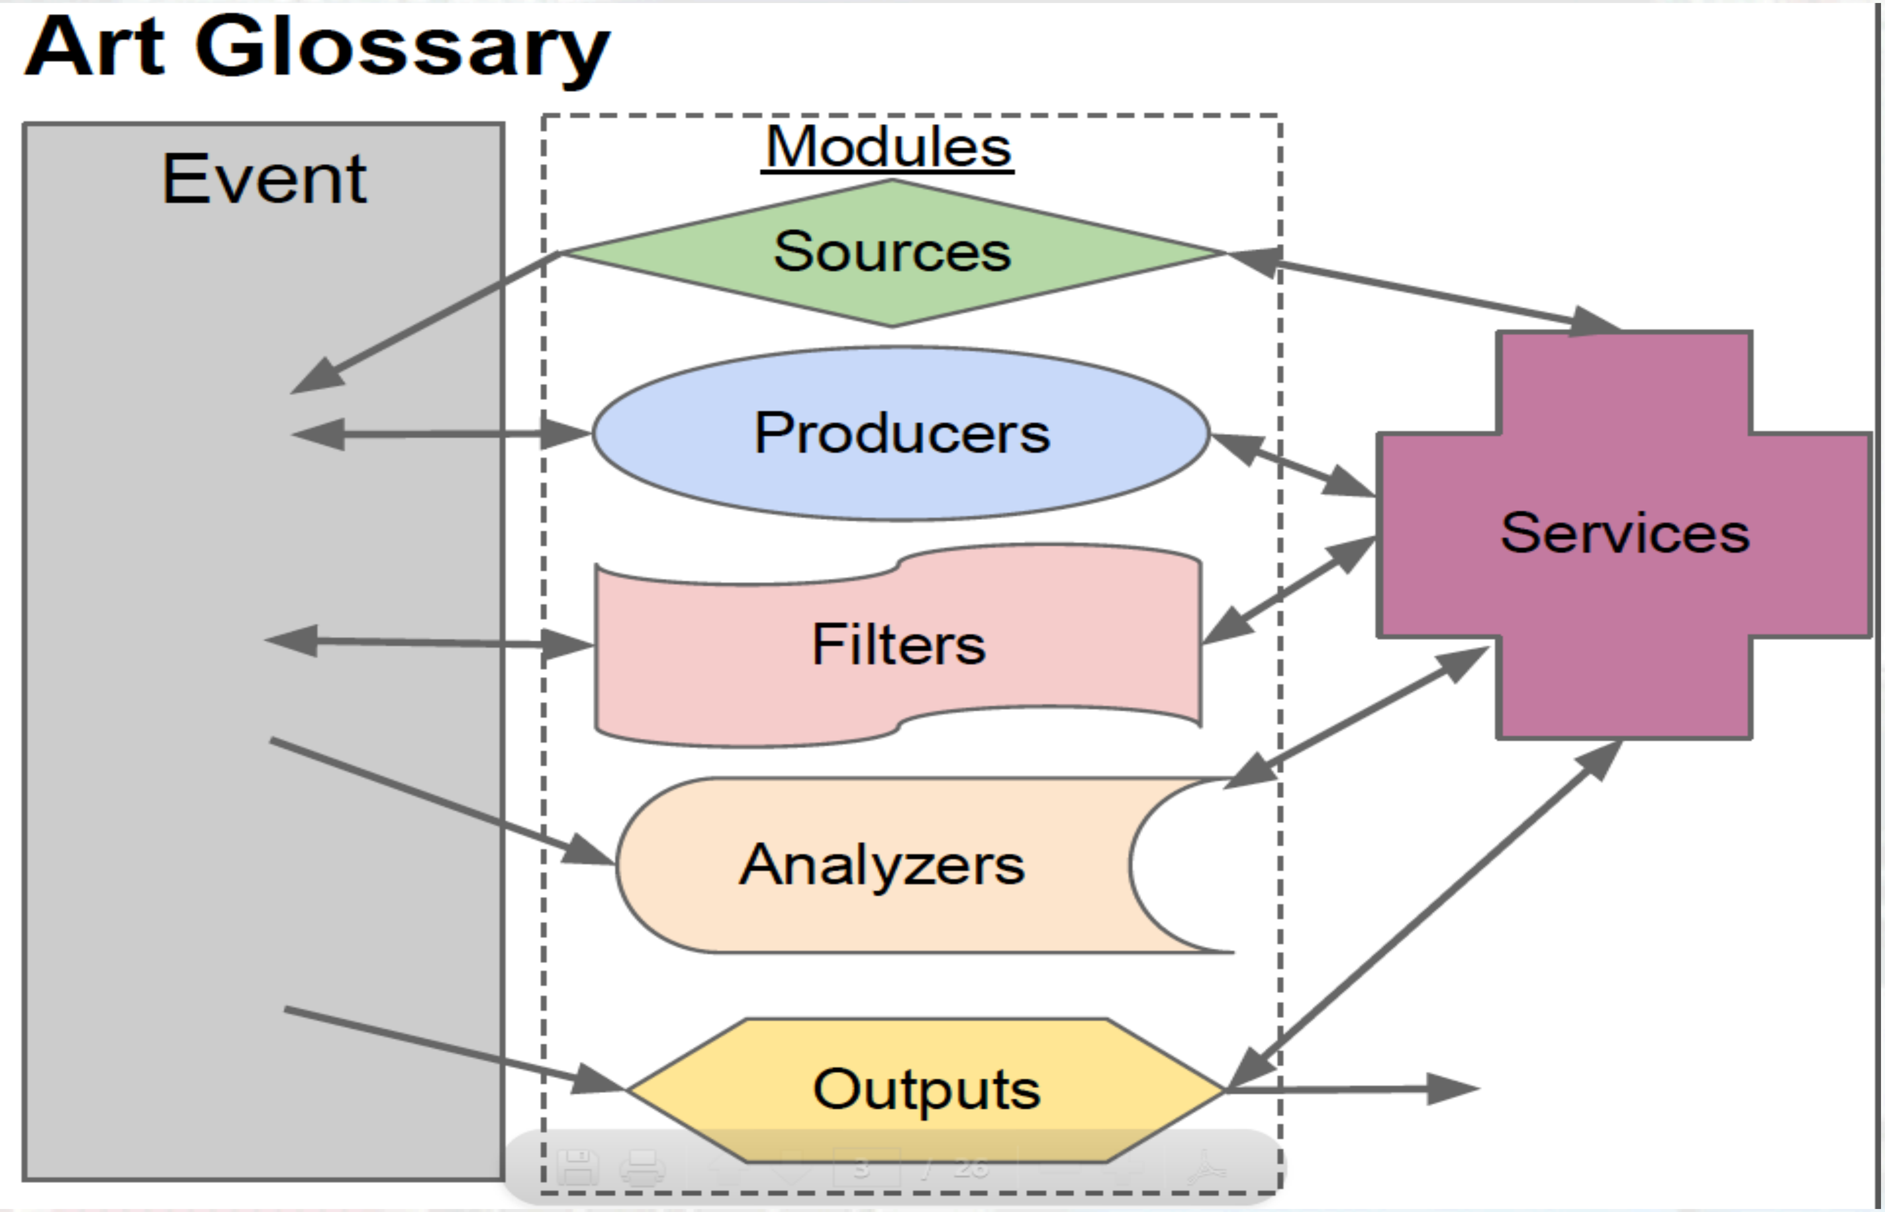

Art Glossary
Event
Modules
Sources
Producers
Filters
Analyzers
Outputs
Services

Example: FIFE Experiment CPU 2014

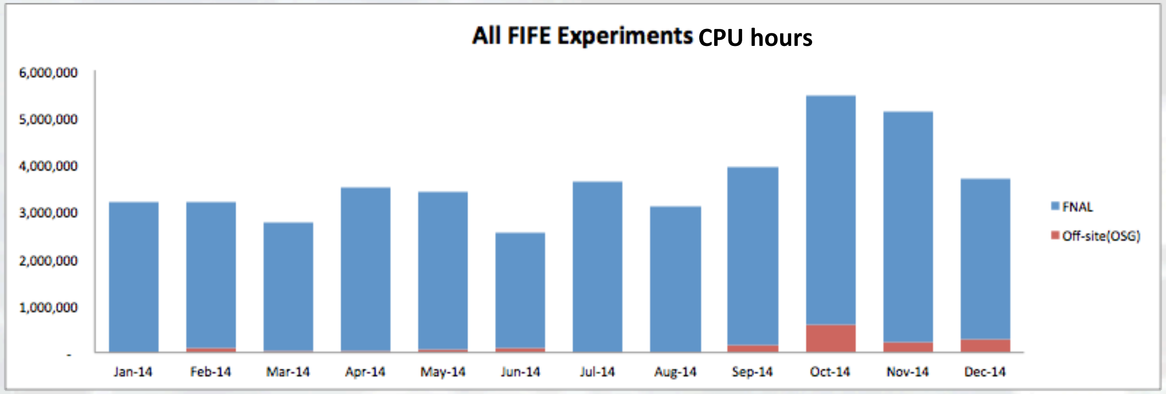

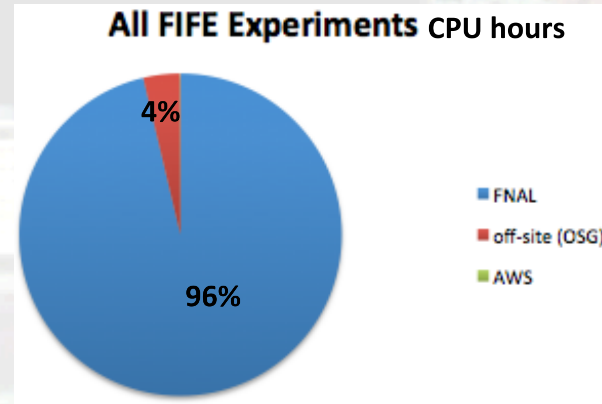

Off-site utilization is ramping up

In 2014 of 45M CPU hours, almost 2M hours were run off-site including first Amazon Web Services (AWS) usage of 20k CPU hours worth $3.3k as a pilot project

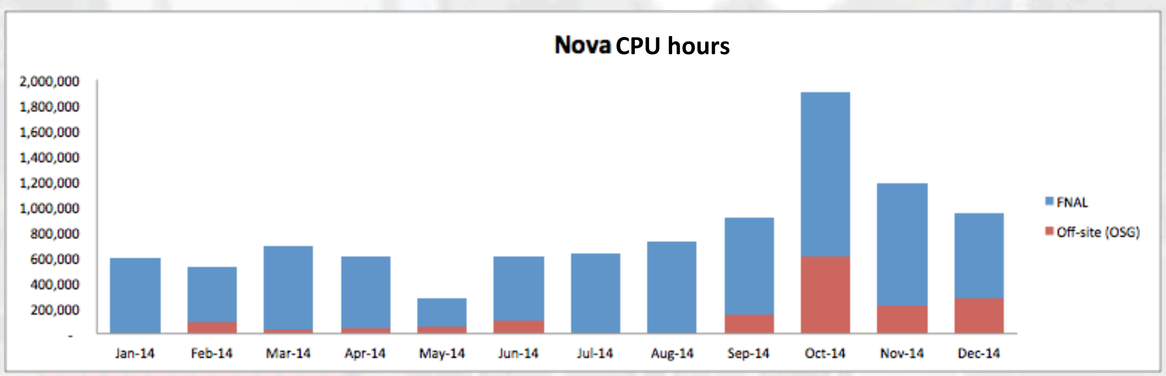

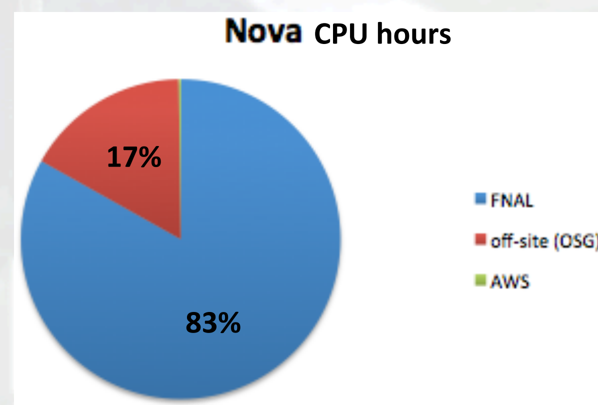

Slide content, O. Gutsche

Production Team

- A new "Production Team" was established in the computing division to help run production jobs for experiments.
- Production team members can quickly learn to run jobs for experiments that use art and FIFE tools.
- Physicists can focus more on the physics that goes into those jobs and analysis.
- Future: Production Management Database Project will streamline data set management using central tools.