

IF Usage of the OSG Overview

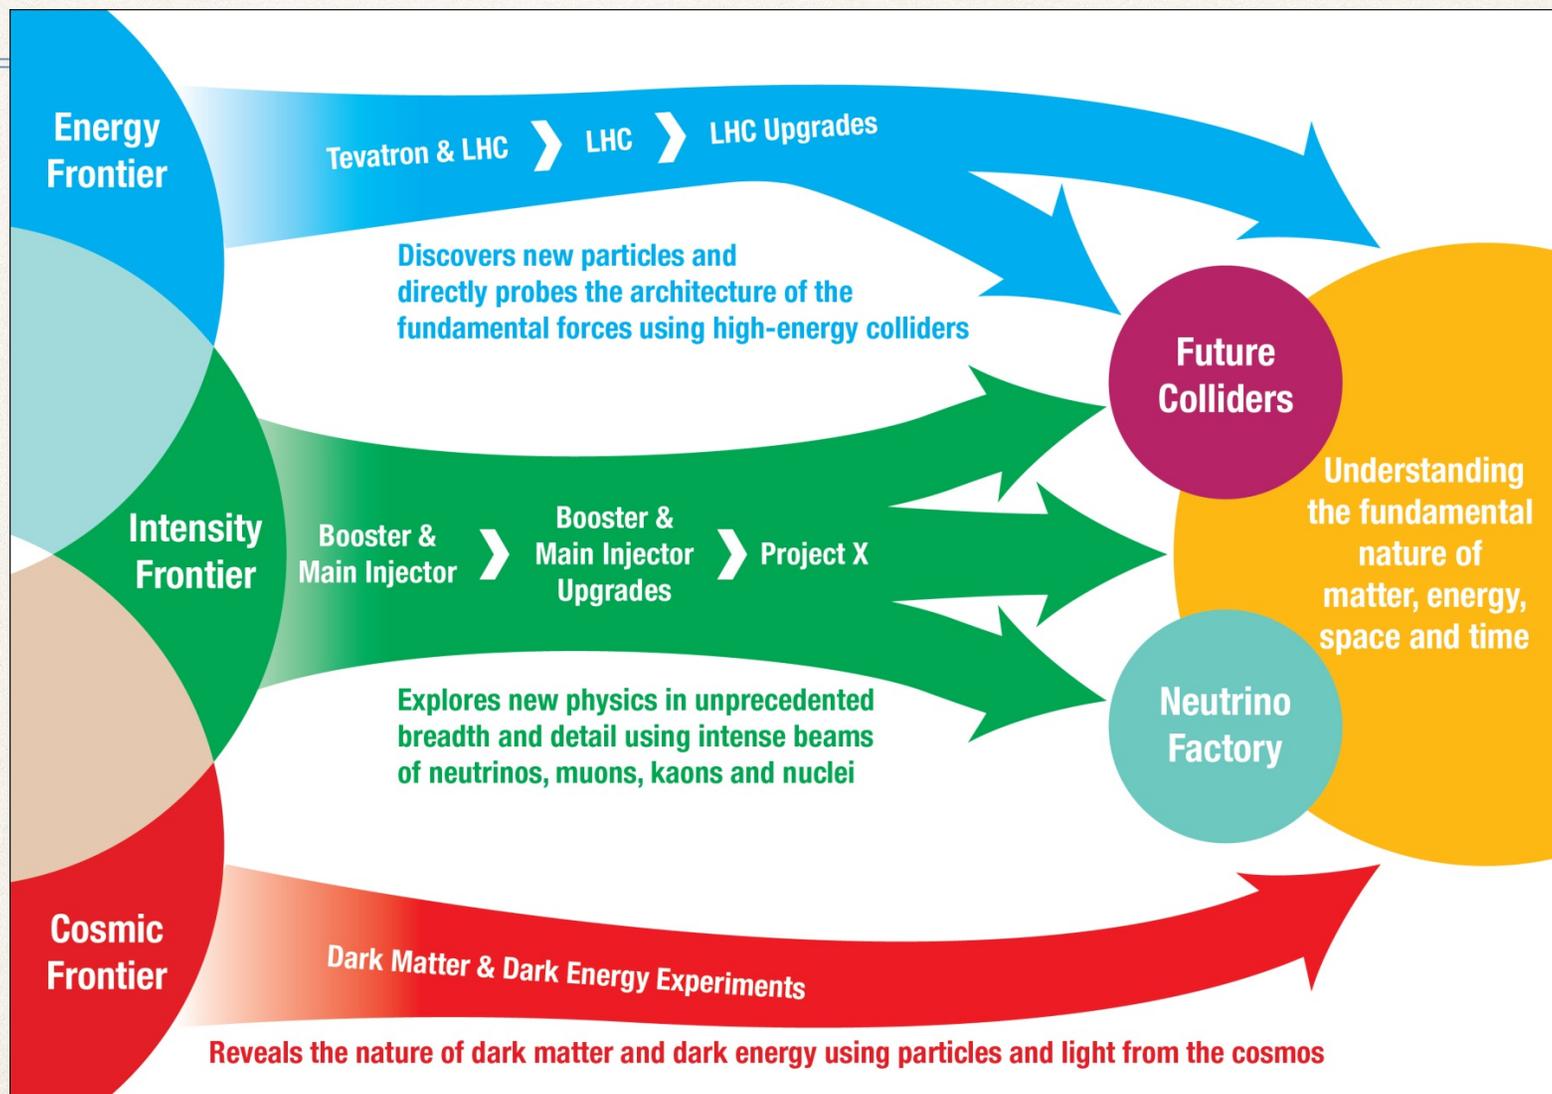
Mike Kirby, Fermilab SCD

Mar 23, 2015

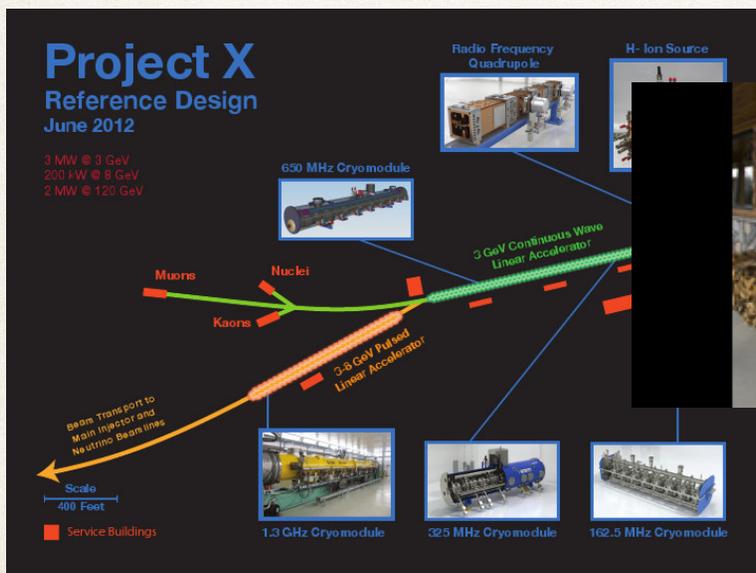
Outline

- ❖ scope of IF experiments in the coming decades
- ❖ computing needs for IF experiments next two years
- ❖ usage over the last year
- ❖ gaining access to the OSG for experiments
- ❖ tool requests and development from IF experiments
- ❖ establishment of as stakeholders within the OSG

Fermilab Frontier Roadmap



Diverse program of experiments



darkside
two-phase argon TPC for Dark Matter Direct Detection

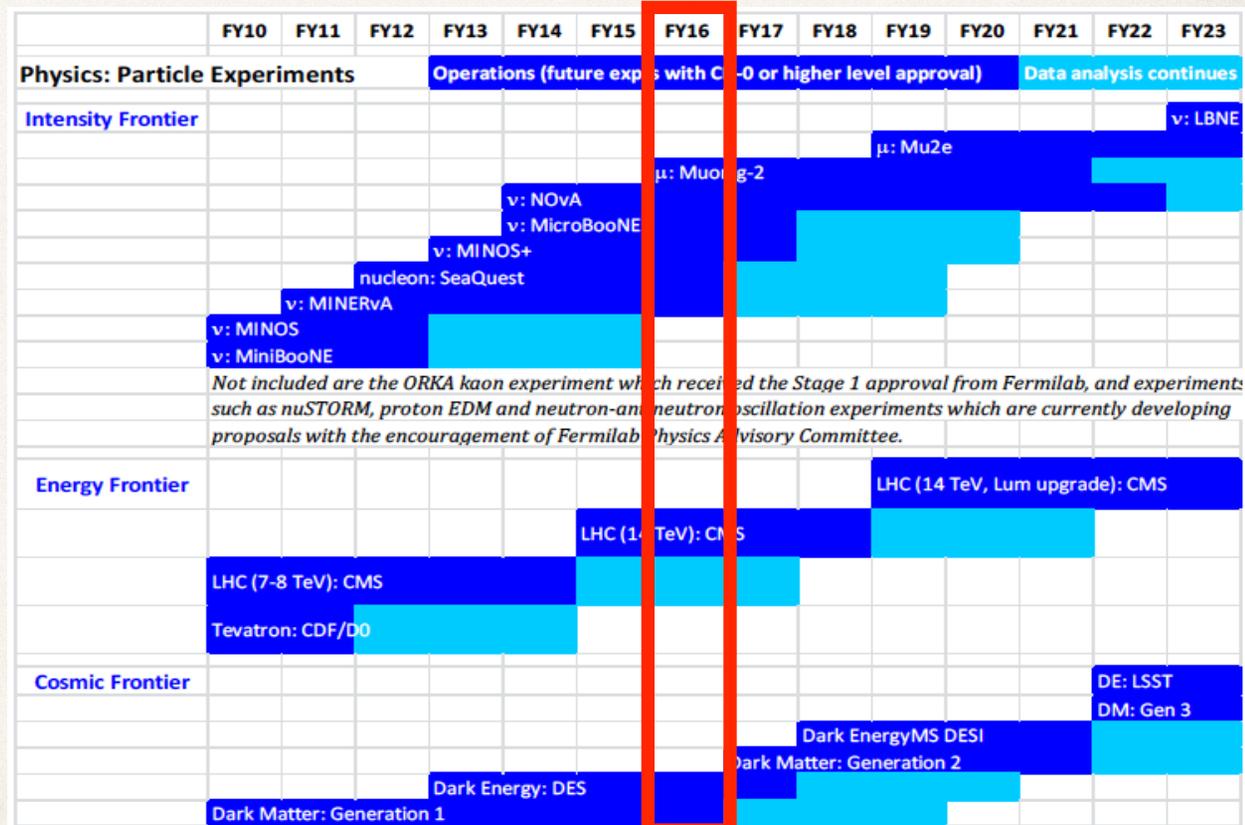
Diverse program of experiments

- ❖ numerous neutrino experiments
 - ❖ MINOS+, Minerva, NO ν A, LArIAT
 - ❖ MicroBooNE , ~~LBNE~~, ~~LBNE~~, DUNE, 35T (a.k.a. DUNEbuggy)
 - ❖ ArgoNeuT, SciBooNE
- ❖ Muon experiments - Muon g-2, mu2e
- ❖ Seaquest, Fermilab Test Beam Facility
- ❖ Belle II (more about this in Malachi Schram's talk)

Fermilab Experiment Schedule

- ❖ 8 major experiments in 3 frontiers running simultaneously in 2016
- ❖ sharing both beam and computing resources
- ❖ start of data taking:
NOvA (full det 2014)
MicroBooNE (now),
LArIAT (now), Muon g-2 (2016)

FY16



Computing requirements for IF experiments @ Fermilab

- * 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

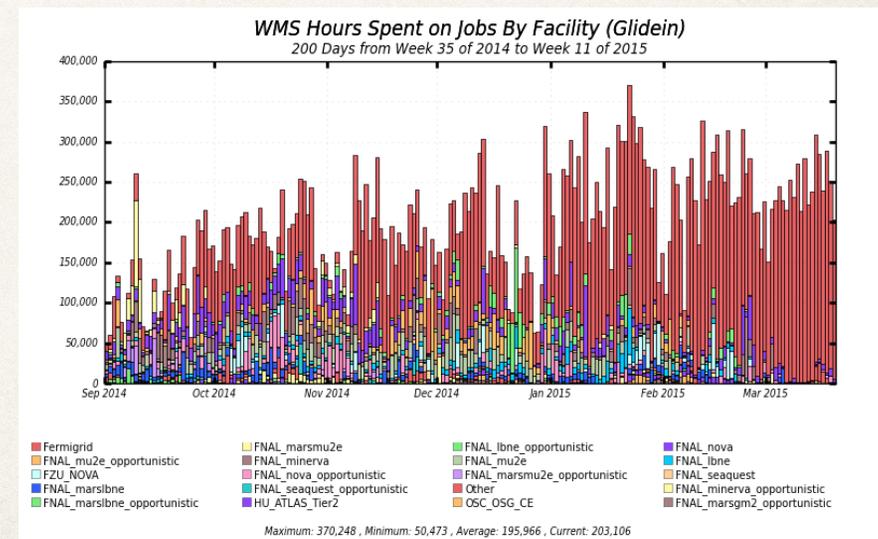
Fermilab Scientific Computing Review

Mil CPU Hours	FY14 Actual	FY15 Req	FY16 Req
MINERvA	9.29	22	28.6
MINOS	15.02	18	22
NOvA	10.18	16	16
Mu2e	9.64	14	20
DUNE	5.37	12.7	10.9
LArIAT	0.30	10	20
Muon g-2	0.38	5.2	7
MARS	6.1	4.7	6.1
Seaquest	0.65	1.8	1
MicroBooNE	1.33	0.5	1.2
Total	59.75	107	136

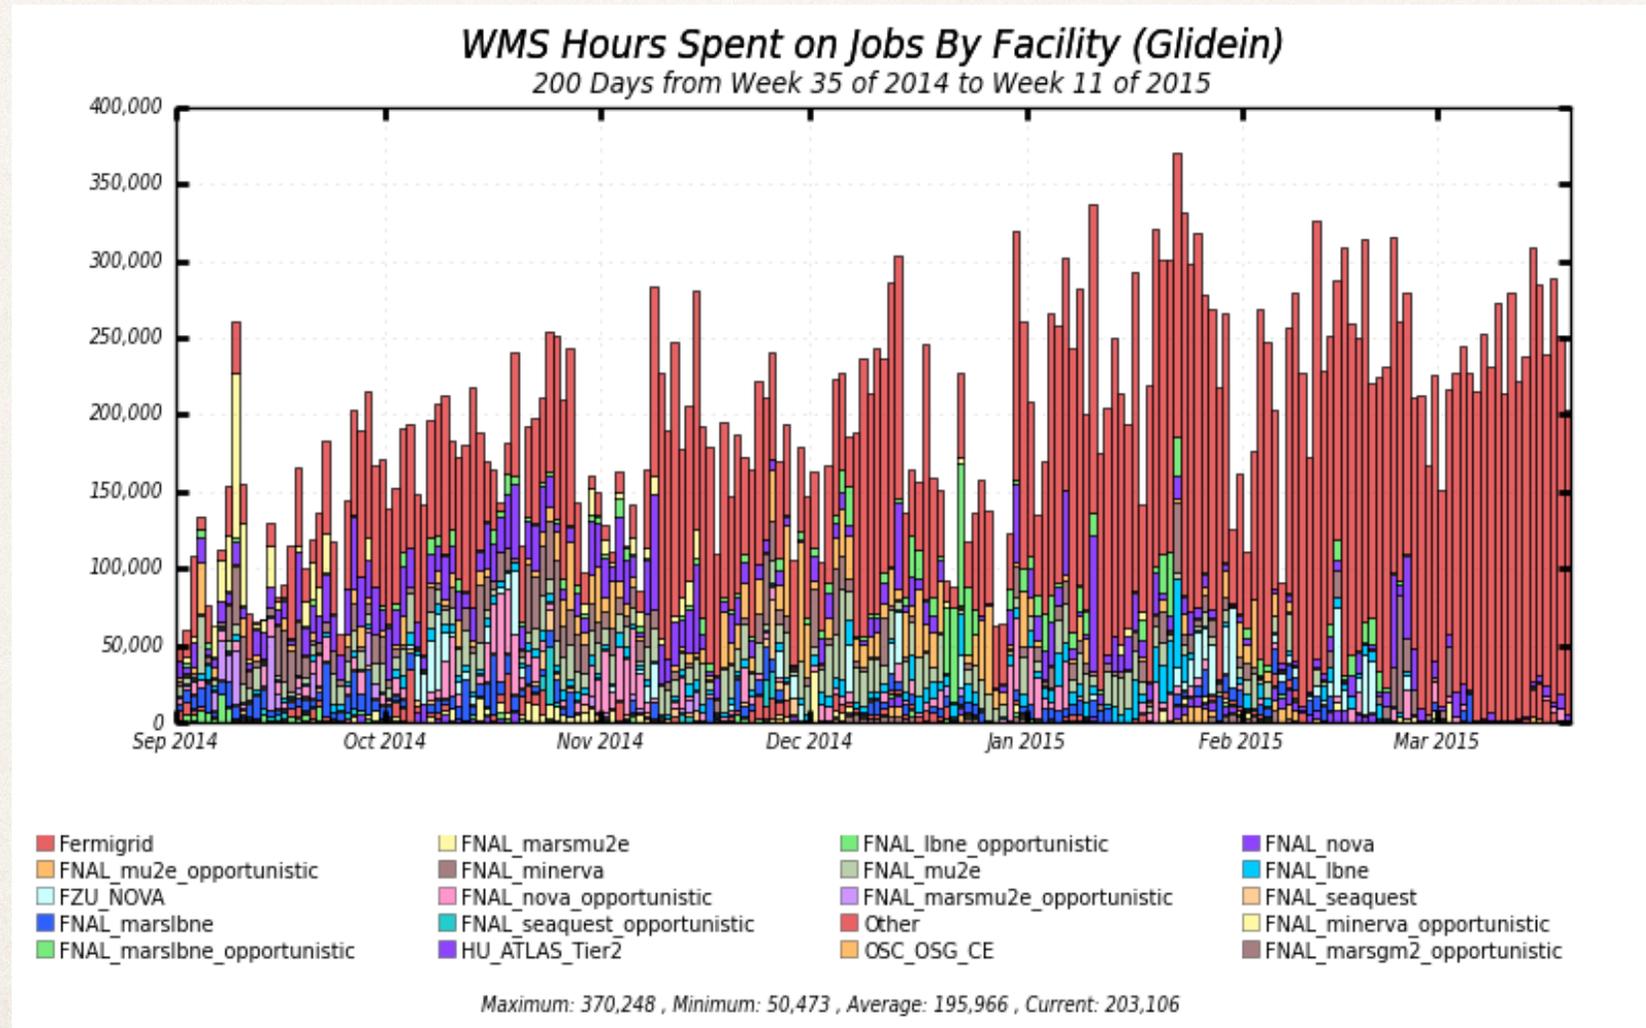
available onsite resources

- ❖ FermiGrid provides 82 M CPU h in FY14 and FY15
 - ❖ 47 M IF
 - ❖ 27 M CDF + Other
- ❖ integrating DØ nodes to increase capacity to 125 M CPU h in FY15/16
- ❖ refocusing experiments to OSG operations, if estimates are correct, they will exceed in FY16
- ❖ but doesn't account for the spiky nature of processing of data, simulation campaigns, and conference deadlines

Working extensively to provide access to OSG resources for the IF experiments

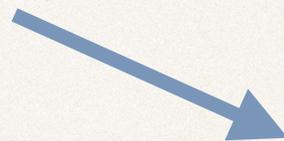
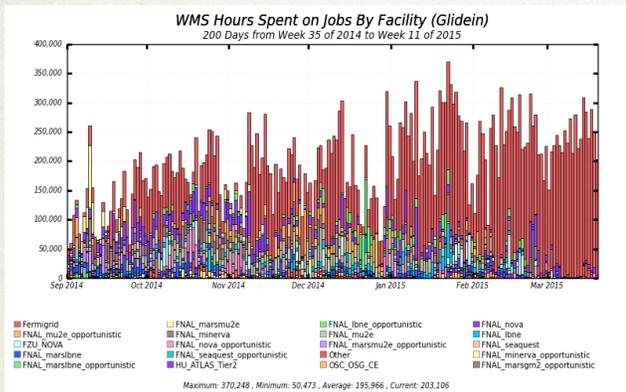


IF Experiment usage last 200 days

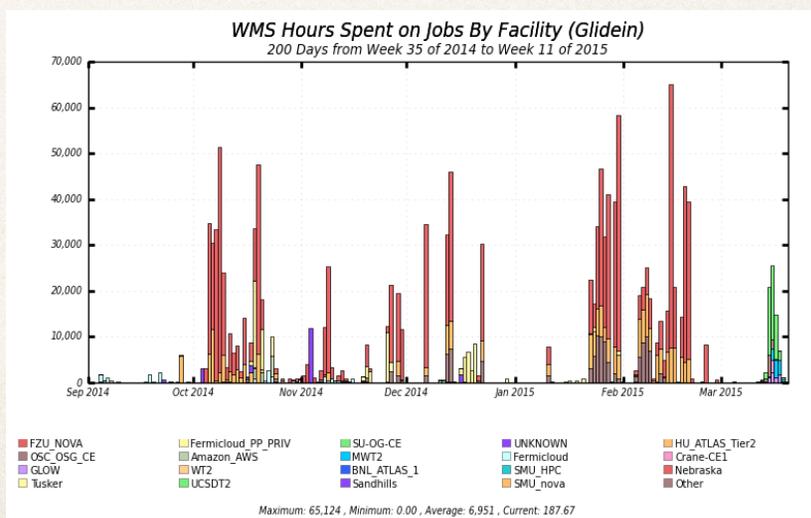
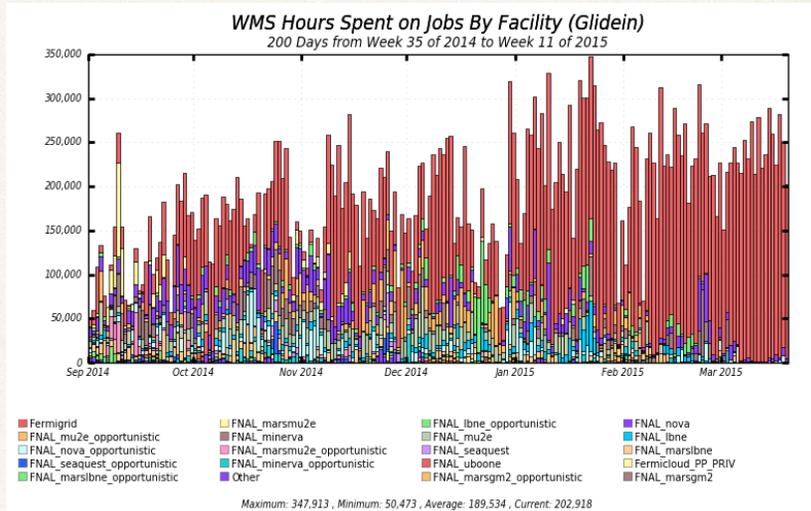


IF Experiment usage last 200 days

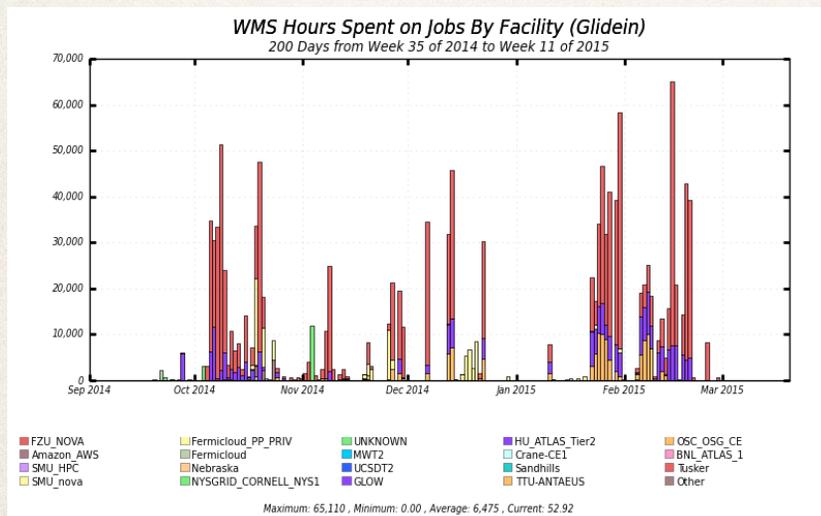
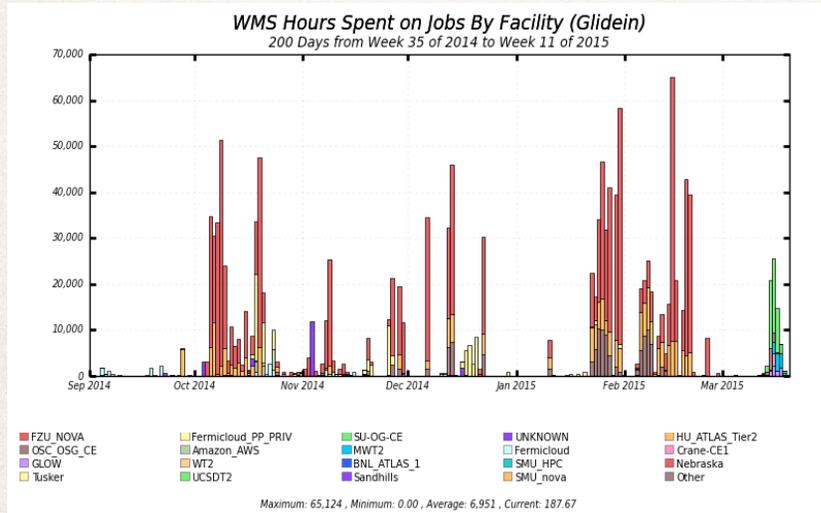
FermiGrid



OSG



IF Experiment usage last 200 days



- ❖ FIFE experiments used 71.5 M CPU h in FY14
 - ❖ 2.5 M CPU h offsite
- ❖ Almost all usage on the OSG from IF experiments has been from NOvA - 2.36 M CPU h)
 - ❖ this will change - Mu2e starting Simulation campaign 14 M CPU h April - Sept 2015
- ❖ Total usage 10 M CPU h in FY14

But why haven't the experiments already made the transition?

- ❖ Getting access to an OSG worker node isn't difficult
 - ❖ Running a hello_world.sh script doesn't take much
 - ❖ FIFE group and OSG exceed provide access through jobsub+GlideInWMS
 - ❖ but that also doesn't really get you a working environment for complex workflows
- ❖ As well, the luxurious resources at Fermilab don't match well to the environment of an OSG node
- ❖ what is difficult?
 - ❖ finding appropriate resources
 - ❖ debugging inconsistencies across sites
 - ❖ getting software distributed to worker node
 - ❖ certs/ access to storage elements
 - ❖ automated mapping of resources to specific resources

Challenges to IF experiments

- * FIFE has put in considerable effort to have the Fermilab VO accepted at more OSG sites
 - * numerous sites will only accept one group of Fermilab VO
 - * Now expanded to 12 sites accepting Fermilab VO completely
 - * some sites state that they accept Fermilab VO without restrictions, but are misconfigured
 - * worst/best case required > 1 month / 1 day to get Fermilab VO operational at site (on average about 1 week per issue)
 - * would hope that OSG sites would act with better sense of community
- * mismatch of requirements at sites
 - * mu2e recently starting probing sites in anticipation of their simulation campaign
 - * physicist ended up debugging problems
 - * blackhole caused by missing CVMFS
 - * mismatch of linux kernel
 - * mismatch mapping some proxies
- * We believe ability to match resources through classAds is critical
- * FIFE wants the OSG to appear “flat” from the physicist perspective

Resource Provisioning

- ❖ IF experiments are in interesting position with respect to computing funding
- ❖ very few PIs have the funding to construct OSG sites for their experiment
- ❖ increasing common for a PI on an IF experiment to receive funding for XX M CPU h at a local resource
 - ❖ also happening at for other frontiers as well
- ❖ ability to direct jobs to prioritize usage at specific sites to ensure the utilization of funding
- ❖ NOvA specifically has request to achieve this for several PIs
- ❖ want to work with the GlideInWMS and OSG teams to meet these requirements
- ❖ configurable priority direction to include OSG sites and paid Clouds - help to meet spiky needs

IF Experiments contribution to OSG tools - CVMFS

- ❖ IF experiments will also contribute back to the OSG toolset
- ❖ increase of experiments utilizing OASIS CVMFS instance led to several difficulties
- ❖ conflicts trying to update and sync software on OASIS -> remote CVMFS repositories
- ❖ distribution of large files for simulation tasks -> development of alien cache
- ❖ Huge thanks to sites for assistance configuring and testing!

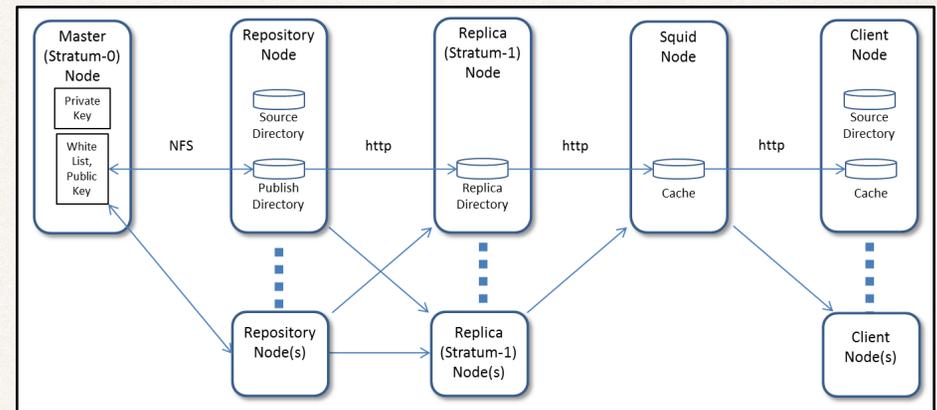


Diagram from Stu Fuess

- ❖ Have repositories located at sites remote from OASIS server
- ❖ independent CVMFS sync so no conflicts
- ❖ remote repo within the OASIS trust realm and so distributed for us

IF Experiments contribution to OSG tools - CVMFS

Alien Cache

- ❖ IF experiments will also contribute back to the OSG toolset
- ❖ increase of experiments utilizing OASIS CVMFS instance led to several difficulties
- ❖ conflicts trying to update and sync software on OASIS ->remote CVMFS repositories
- ❖ distribution of large files for simulation tasks -> development of alien cache
- ❖ Huge thanks to sites for assistance configuring and testing!



- ❖ utilize fast access storage elements located at site
- ❖ configure separate repository to deliver files to local SE (hadoop, dCache, Luster, etc)
- ❖ drastically improved performance certain workflows
- ❖ See Dave Dykstra's talk Tues 5:30 pm

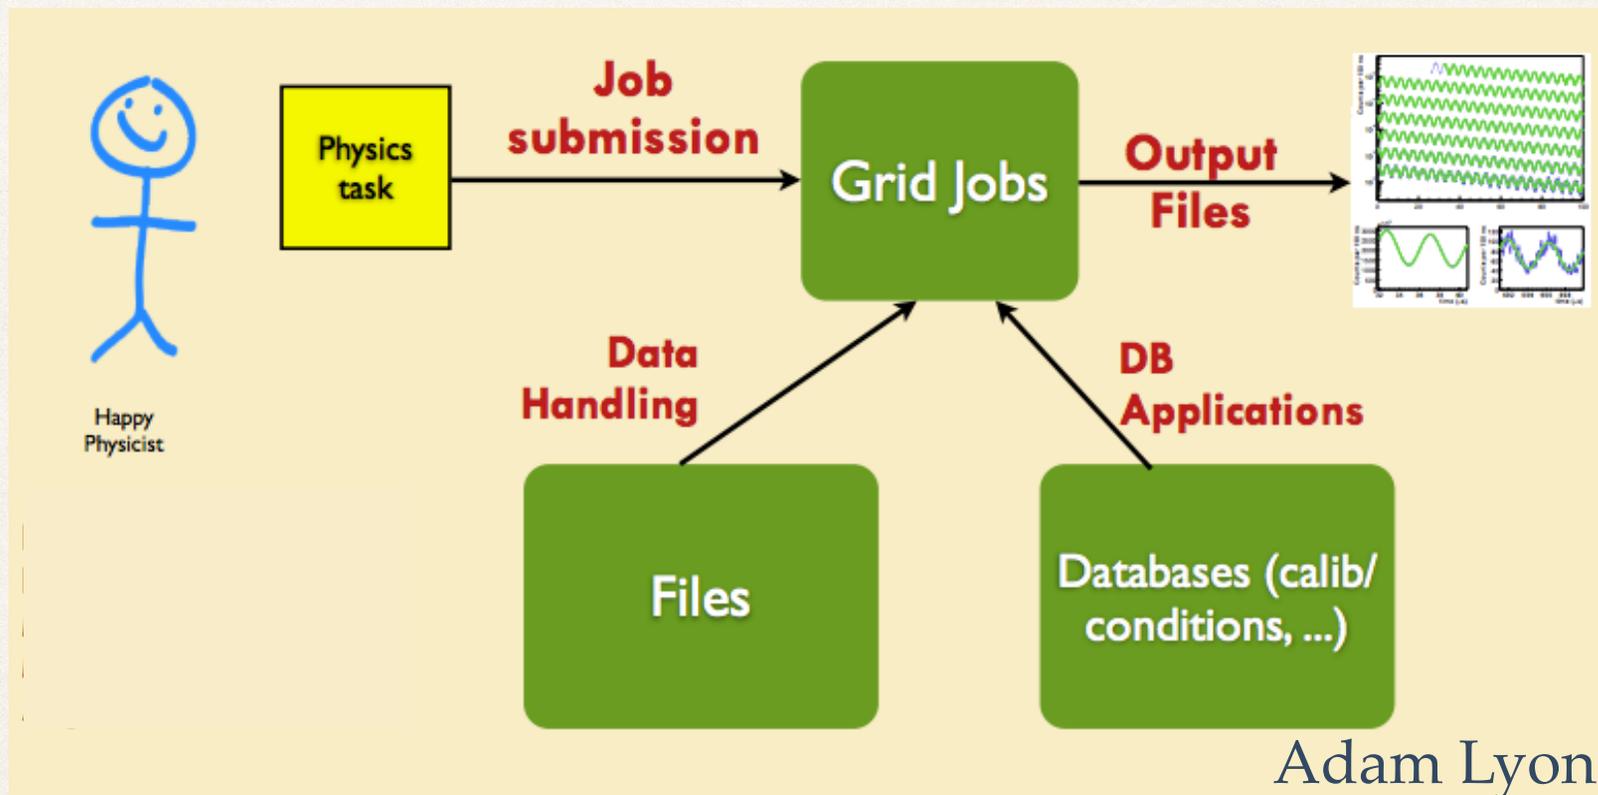
IF Experiments interface with OSG

- ❖ one issue going forward will be IF computing communication with the OSG
- ❖ the burden of addressing the concerns of 15+ disparate experiments would be considerable
- ❖ having the IF experiments consolidate into an organization which becomes a stakeholder in the OSG
- ❖ allows for the establishment of tools and give guidance to the OSG and back to the IF experiments
- ❖ in the future, large experiments like Belle2, DUNE, LSST will begin to provide OSG resources instead of just utilizing them
- ❖ most importantly, make it easier for experiments to incorporate the OSG into their computing model without the growing pains currently occurring

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

IF Experiments on the OSG



- ❖ ideally, this is what we want
- ❖ we know that the OSG can help this happen for IF exp

backup slides
