

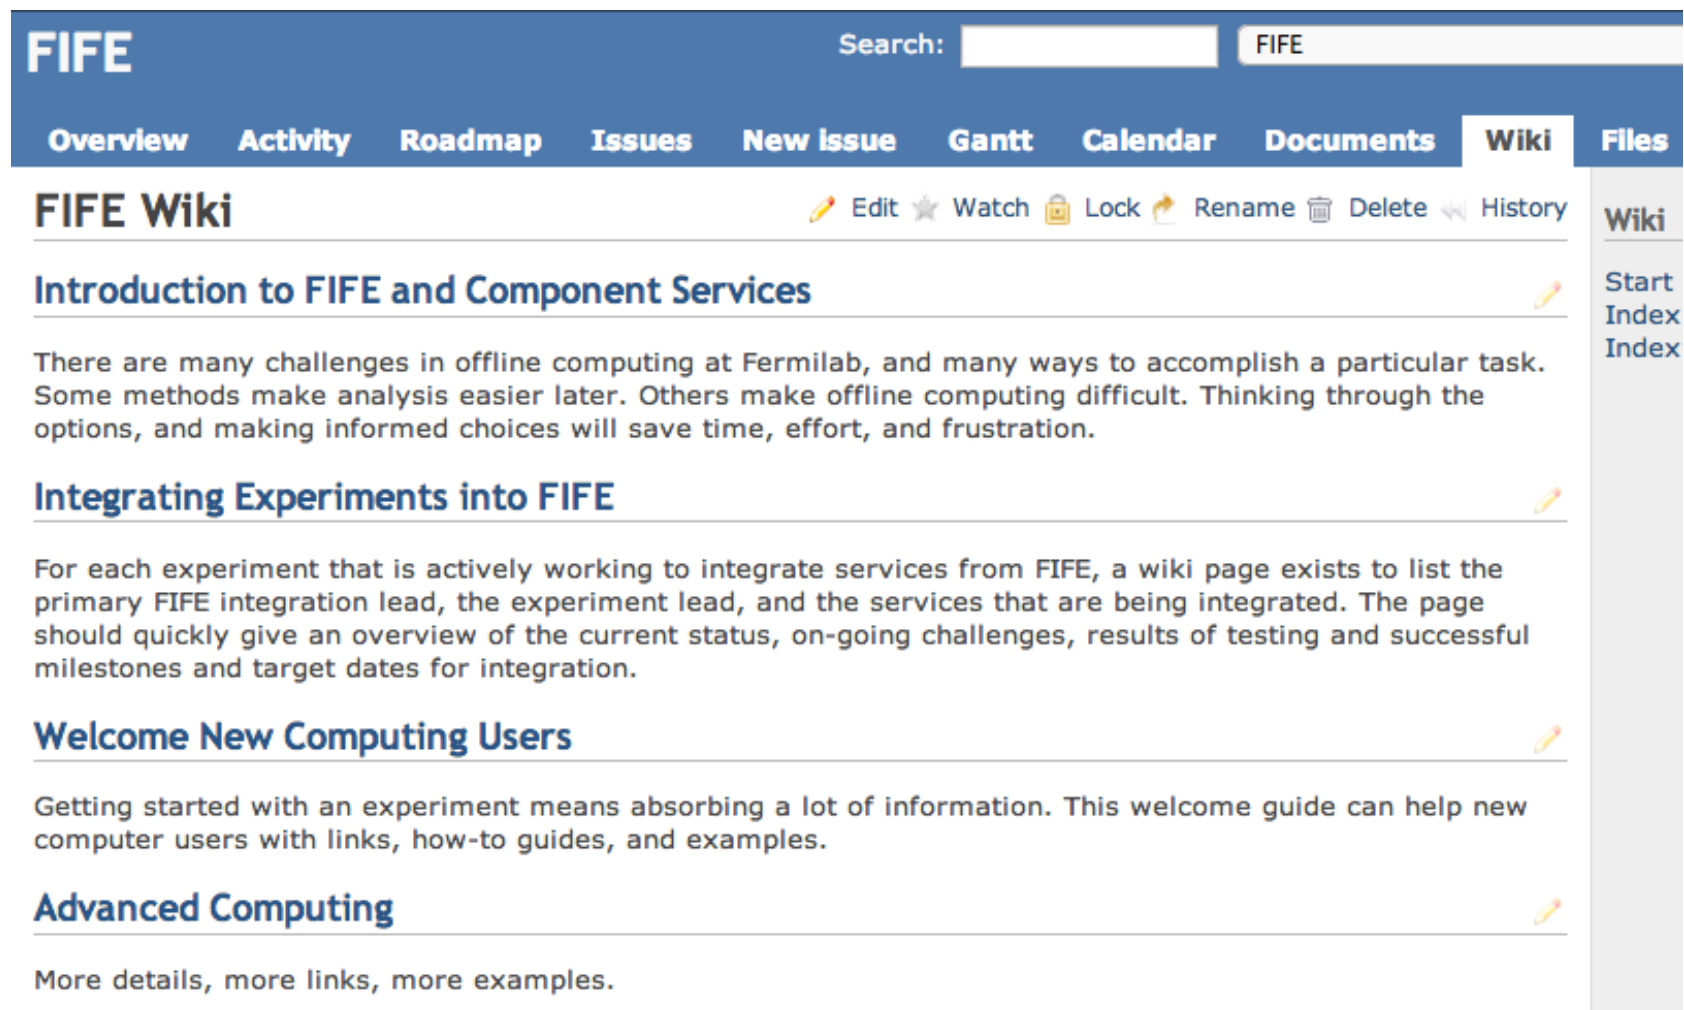


Introduction to FIFE

Mike Kirby/Ken Herner
FIFE Support Group

Documentation







- <https://cdcvcs.fnal.gov/redmine/projects/fife/wiki>




The screenshot shows the FIFE Wiki page within a Redmine application. The top navigation bar includes links for Overview, Activity, Roadmap, Issues, New Issue, Gantt, Calendar, Documents, Wiki (selected), and Files. A search bar is present with the text 'FIFE'. The main content area displays the 'FIFE Wiki' title with editing tools (Edit, Watch, Lock, Rename, Delete, History). Below the title are three article sections: 'Introduction to FIFE and Component Services', 'Integrating Experiments into FIFE', and 'Welcome New Computing Users'. Each section has a brief description and an edit icon. A sidebar on the right contains links for 'Wiki', 'Start', 'Index', and 'Index'.


FIFE Search:

Overview Activity Roadmap Issues New Issue Gantt Calendar Documents Wiki Files


FIFE Wiki  Edit  Watch  Lock  Rename  Delete  History

Introduction to FIFE and Component Services 


There are many challenges in offline computing at Fermilab, and many ways to accomplish a particular task. Some methods make analysis easier later. Others make offline computing difficult. Thinking through the options, and making informed choices will save time, effort, and frustration.

Integrating Experiments into FIFE 

For each experiment that is actively working to integrate services from FIFE, a wiki page exists to list the primary FIFE integration lead, the experiment lead, and the services that are being integrated. The page should quickly give an overview of the current status, on-going challenges, results of testing and successful milestones and target dates for integration.

Welcome New Computing Users 

Getting started with an experiment means absorbing a lot of information. This welcome guide can help new computer users with links, how-to guides, and examples.

Advanced Computing 

More details, more links, more examples.

Wiki
Start
Index
Index

Documentation

Introduction to FIFE and Component Services

Introduction to FIFE and Component Services

What is FIFE?

FIFE References:

FermiGrid

FermiGrid References

Open Science Grid Overview

Open Science Grid References

Jobsub

Jobsub Tools

JobSub Client-Server

JobSub References

Authentication

Authentication References

art framework

art references

OASIS/CVMFS

OASIS/CVMFS References

Data Management Overview

Data Management References

IF Data Handling Client Tools (ifdhc)

ifdhc References

SAM

SAM References

File Transfer System

FTS References

dCache

dCache References

FermiCloud

FermiCloud References:

Conditions database


Conditions Database References

Electronic Logbooks

Electronic Logbooks References

Useful links to other services from CD and SCD

- There is an overview of all of the topics, but also links for detailed documentation
- Please read and send us feedback - it's a work in progress
- we will focus on big picture for computing



The Muon $g-2$ Computing “Big Picture”

Adam Lyon
(Fermilab/Scientific Computing Division)

$g-2$ for Beginners
February 20, 2013



I'll discuss the following topics:

What is the Fermilab computing vision and how do we fit in?

**Grids, Virtualization, Clouds
Art, Multicore, GPUs**

**I'll tell you how this matters to you
(yes, we all just want to make plots!)**

Many opinions here are mine alone

“Hopefully I'll tell you a little more than what you already know; certainly a little more than what I already know.” - R. Mohapatra

Fermilab provides many computing services

Fermilab has hosted many *batch* farms, but we now wrap farms in the *GRID*.

Grid computing:

A common interface to many batch systems on many farms. Common infrastructure and assistance



Open Science Grid


Main US DOE Grid
HEP, Biology, Seismology

Not the same as NSF's
Supercomputer Xsede (Teragrid)



What is a Grid?

A step in the direction of “computing as a utility”



- You submit your job, and it advertises requirements.**
- Sites advertise capabilities.**
- A “broker” matches your job to the site and it runs.**


You don't care where your job runs, in principle

In practice it's been difficult to make this work smoothly


But the payoff? Opportunistic cycles! Even in the LHC era!

Our Fermilab Grid (Fermigrid)

**Fermigrid is a collection of farms at the lab all with an OSG interface:
CDF, DØ, CMS, General Purpose**



~25,000 slots! Separate farms cause some inefficiency



Some things to know about μBooNE Computing

μBooNE has 200 allocated slots on FermiGrid

potentially about to get more slots than that

you will compete against other folks from μBooNE first for those 200, and then with everyone for opportunistic slots

- fairshare

there are limits on jobs - 48 hours of wall time, memory limits, local disk storage

The Grid isn't easy!

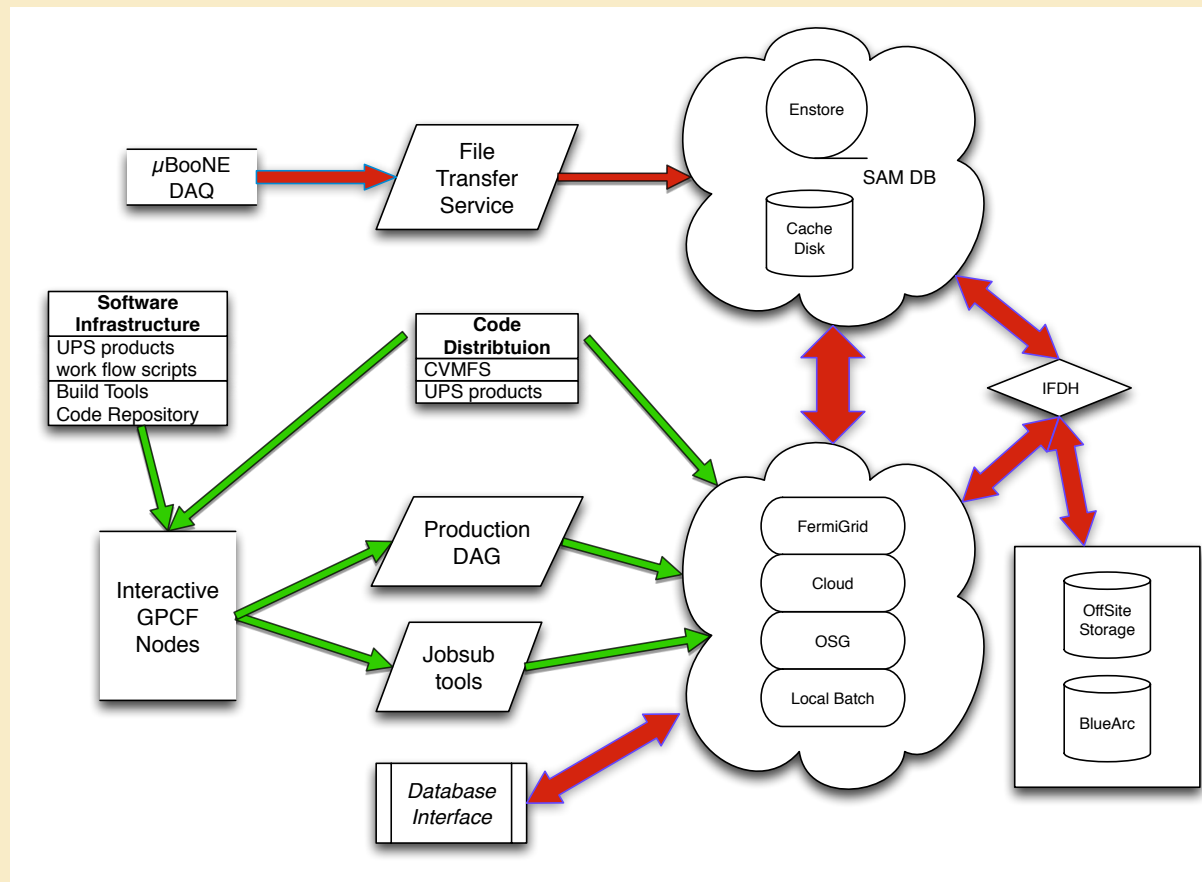
Making different sites accept Grid requirements (better if they are LHC sites)

Authentication and authorization – are you who you say you are? Do you have permission to run there?

Do your job requirements really match the site?

How do you get your application and data to the job? How are the results returned?


How do I get to data?




Some solutions to these problems


“Glidein-WMS” – a batch system “projected” onto the Grid that looks local

The outside world sees:



 = worker node running pilot

We see



We submit to unified batch system we manage made up of worker nodes distributed around the grid and activated on demand

We can enforce our policies (fair share, time limits), but are subject to underlying grid system limits too

Pilots prevent jobs from landing on unsuitable worker nodes

Looks like one big batch system to us - easier to manage.

The “JobSub” package does this for you



Local storage - Bluearc

Good news:

When used as designed, it works great

Bad news:



When used outside of its design, it kills computing for all of the IF experiments (hard to buy a robust, reasonably priced, multi-PB system)



Local storage - Bluearc

some general rules:

- never use “cp” command in a grid job. EVER!
- never open a file that is on BlueArc from grid
- copy the file onto the local grid disk, then process
- we'll get to the question of quotas shortly

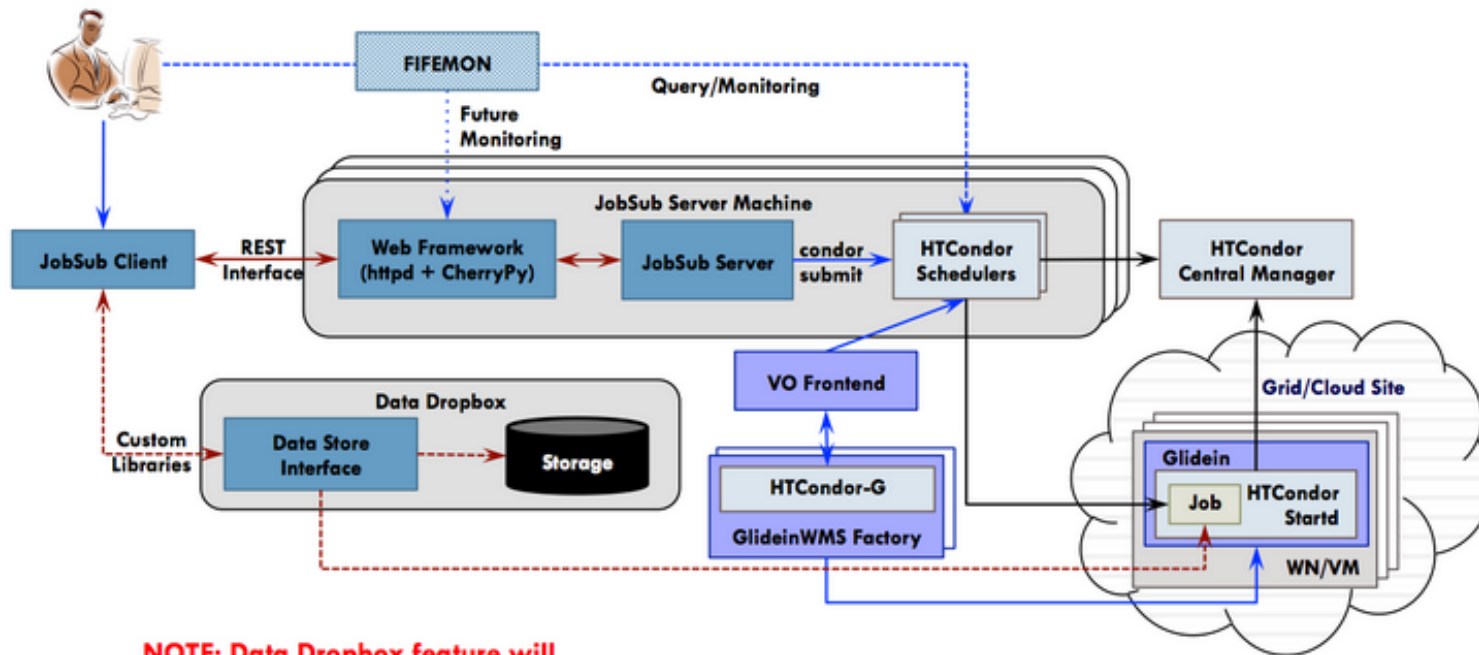


ifdh to move files around

<https://cdcvns.fnal.gov/redmine/projects/ifdhc/wiki/>

- **ifdh uses several different techniques to move files**
 - **cpn** - this is a throttled cp command - only 10 per exp
 - **gridftp** - this throttle the same way as cpn but it is a global access point - any grid site can copy it back - also comes back with the right ownership of files
 - **BestMan** - internal throttling using OSG software SRM copy to move files - the wave of the future
- **Always use ifdh to move files over then network on the grid**

jobsub and all it's options



NOTE: Data Dropbox feature will be implemented in future releases.

jobsub and all it's options

- https://cdcvcs.fnal.gov/redmine/projects/jobsub/wiki/Using_the_Client
- we're going to steal from Ken for a hello world example
- Thanks Ken!

jobsub hello world

log into gm2gpvm01.fnal.gov

```
> source /grid/fermiapp/products/common/etc/setups.sh  
> setup jobsub_client  
> get-cert
```

```
cp /afs/fnal.gov/files/home/room2/kherner/monitoring\_test.sh /gm2/app/users/<your username>
```

```
jobsub_submit.py -G lariat -M —OS=SL5,SL6 —resource-  
provides=usage_model=DEDICATED,OPPORTUNISTIC file:///gm2/app/users/<your username>/  
monirotng_test.sh
```

-G is your experiment GROUP

-M mails you when your jobs finish at <your username>@fnal.gov

—OS says which flavor of Scientific Linux you want, this requests both

—resource-provides tells what to look for in a computing resource, this time it's both dedicated and opportunistic

- last element is the file to copy in and run as the executable

Monitoring my jobs

- fifemon.fnal.gov/monitor/
- all kinds of stuff here, we'll just discuss it...

Moving your data to your job

Note that the LHC experiments tried “move your job to the data” and are migrating to this method instead

