



jobsub_lite Tutorial and Demo

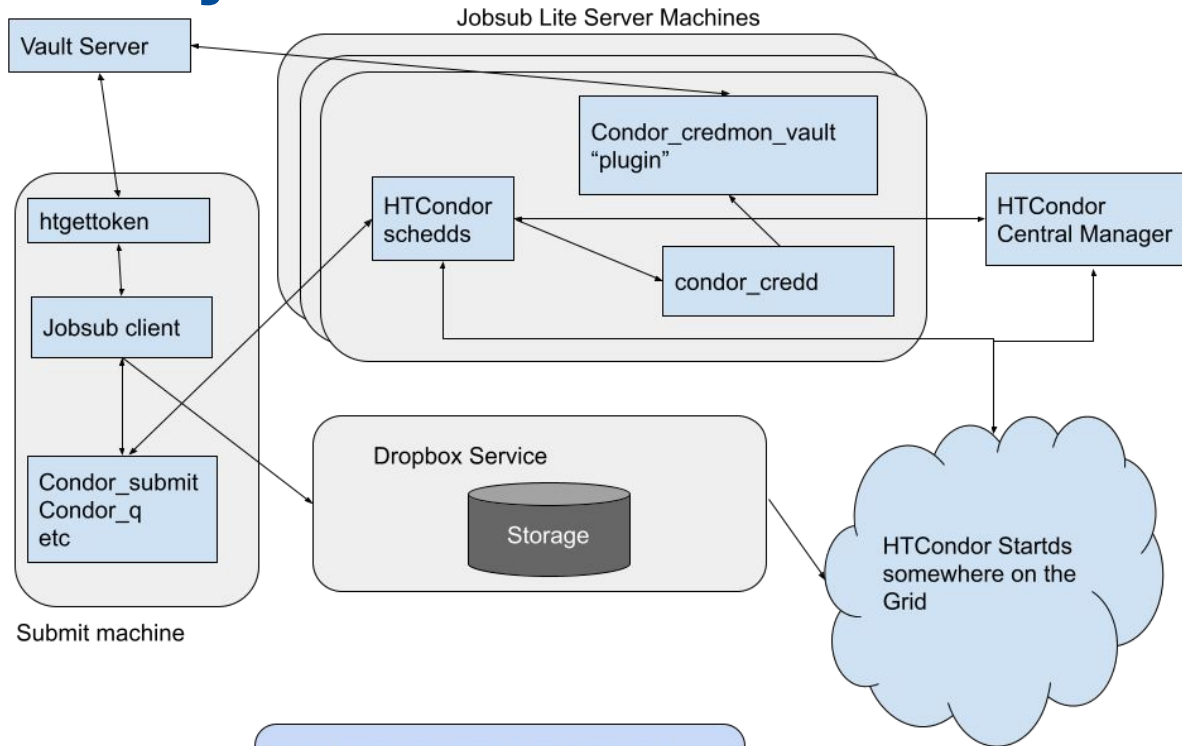
Shreyas Bhat and Marc Mengel
on behalf of the Jobsub Team

December 8, 2022, January 5, January 11, and
February 2, 2023

Intro to jobsub_lite

- New software for job submission and monitoring
- Tried to keep the most-used pieces of jobsub_client, strip out unnecessary parts
- Client-only, installed on experiment interactive nodes
- Much of the content of this talk has been written up here:
https://fifewiki.fnal.gov/wiki/Getting_started_with_jobsub_lite

jobsub_lite Layout with Tokens



Active/Large experiments get their own schedd and can't be brought down by other experiments errant users.

Original Image Credit: J. Boyd

What happens

- jobsub_lite executables present in \$PATH from login shell - NO SETUP NECESSARY!
- jobsub_lite takes user command, converts it to HTCondor submission file (Job Definition File), and uses HTCondor python bindings to submit the job to a remote schedd
- Currently deploying with shared schedds, but plan to transition to one schedd per large experiment

Tokens and Authentication

Token Authentication Flow

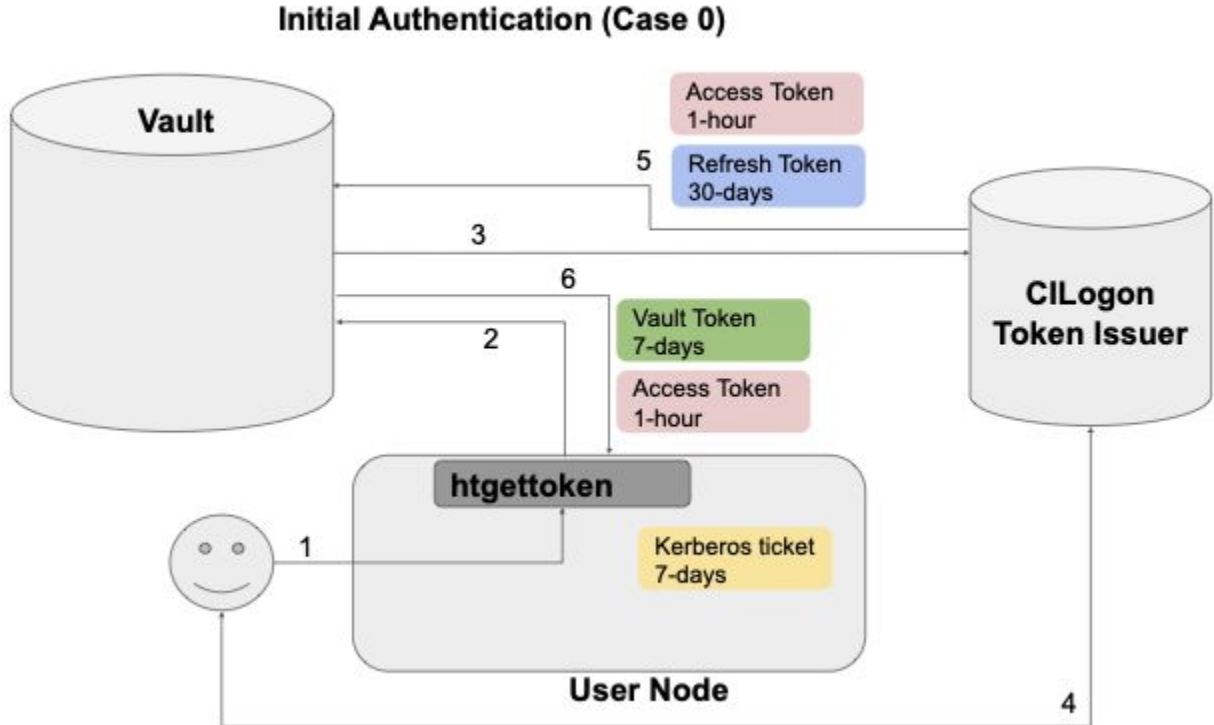


Image Credit: M. Altunay and D. Dykstra

First-time authentication

- Authentication happens for most grid operations - now X509 Proxy, soon tokens (jobsub_lite, ifdhc commands)
- Absence of vault or refresh token → Authenticate with CILogon

Attempting OIDC authentication with `https://htvaultprod.fnal.gov:8200`

Complete the authentication at:

`https://cilogon.org/device/?user_code=_redacted_user_code`

No web open command defined, please copy/paste the above to any web browser

Waiting for response in web browser

- Will need to copy/paste that link into browser

First-time authentication, continued



Consent to Attribute Release ▼

[htvaultst-fermilab-vault](#) requests access to the following information. If you do not approve this request, do not proceed.

- User Code: Z4H-CHD-WPJ
- Your CILogon user identifier
- Your name
- Your email address
- Your username and affiliation from your identity provider

Select an Identity Provider

Fermi National Accelerator Laboratory ⓘ ←

Remember this selection ⓘ

←

By selecting "Log On", you agree to the [privacy policy](#).

- Select "Fermi National Accelerator Laboratory"
- Then click "Log On"
- Log in with Services Credentials

For questions about this site, please see the FAQs or send email to help@cilogon.org.
Know your responsibilities for using the CILogon Service.
See acknowledgements of support for this site.

Further Notes about Authentication

- After initial authentication, as long as you use token-enabled grid tools for the same experiment/role at least every 30 days, you should *not* have to reauthenticate
- This is because refresh token (kept in vault) expires after 30 days of inactivity
- Tokens downloaded to user machine:
 - Vault Token: Used to authenticate to vault
 - Access (or Bearer) Token: SciToken (JWT) that is actually used for grid operations
- More information on SciTokens: <https://scitokens.org/>

Submit and Manage Simple Job

jobsub_submit

- Much easier than before. Just login, and jobsub_submit

```
$ jobsub_submit -G fermilab file:///usr/bin/printenv
Attempting to get token from https://fermicloud543.fnal.gov:8200 ... failed
Attempting kerberos auth with https://fermicloud543.fnal.gov:8200 ... succeeded
Attempting to get token from https://fermicloud543.fnal.gov:8200 ... failed
Attempting OIDC authentication with https://fermicloud543.fnal.gov:8200
```

Complete the authentication at:

```
https://cilogon.org/device/?user_code=<code>
```

No web open command defined, please copy/paste the above to any web browser

Waiting for response in web browser

```
Storing vault token in /tmp/vt_u10610
```

```
Storing bearer token in /tmp/bt_token_fermilab_Analysis_10610
```

```
Submitting job(s).
```

```
1 job(s) submitted to cluster 57106734.
```

```
Use job id 57106734.0@jobsub01.fnal.gov to retrieve output
```

jobsub_submit, continued

- Like before, -G/--group is required to submit job (and run all jobsub executables)
- Group dictates which token issuer is used to get a bearer token
- FOR NOW, jobsub will obtain a bearer token and VOMS-proxy (valid for ~one week) and send these to the job
 - Future - no VOMS proxy
- Demo

Manage jobs

- `jobsub_q`, `jobsub_hold`, `jobsub_release`, `jobsub_rm`, etc.
written as lightweight wrappers around `condor_*` commands
- Tried to keep backward-compatibility
- Examples on following slides

jobsub_q

```
$ jobsub_q -G fermilab
```

JOBSUBJOBID COMMAND	OWNER	SUBMITTED	RUNTIME	ST	PRIO	SIZE
57106962.0@jobsub01.fnal.gov simple.sh	sbhat	11/30 14:59	0+06:00:10	C	0	1953.1
57106973.0@jobsub01.fnal.gov simple.sh	sbhat	12/01 13:48	0+06:00:27	R	0	0.0

```
$ jobsub_q -G fermilab 57106973.0@jobsub01.fnal.gov
```

JOBSUBJOBID COMMAND	OWNER	SUBMITTED	RUNTIME	ST	PRIO	SIZE
57106973.0@jobsub01.fnal.gov simple.sh	sbhat	12/01 13:48	0+06:00:27	R	0	0.0

jobsub_hold

```
$ jobsub_hold -G fermilab 57106973.0@jobsub01.fnal.gov
```

```
Job 57106973.0 held
```

```
$ jobsub_q -G fermilab 57106973.0@jobsub01.fnal.gov
```

JOBSUBJOBID	OWNER	SUBMITTED	RUNTIME	ST	PRIO
57106973.0@jobsub01.fnal.gov	sbhat	12/01 13:48	0+06:00:27		H
0 0.0 simple.sh					

jobsub_release

```
$ jobsub_release -G fermilab 57106973.0@jobsub01.fnal.gov
```

```
Job 57106973.0 released
```

```
$ jobsub_q -G fermilab 57106973.0@jobsub01.fnal.gov
```

JOBSUBJOBID	OWNER	SUBMITTED	RUNTIME	ST	PRIO
57106973.0@jobsub01.fnal.gov	sbhat	12/01 13:48	0+06:00:27	I	
0 0.0 simple.sh					

jobsub_rm

```
$ jobsub_rm -G fermilab 57106973.0@jobsub01.fnal.gov
```

```
Job 57106973.0 marked for removal
```

```
$ jobsub_q -G fermilab 57106973.0@jobsub01.fnal.gov
```

JOBSUBJOBID	OWNER	SUBMITTED	RUNTIME	ST	PRIO
SIZE	COMMAND				

Singularity/Apptainer

- By default, jobs run in `fnal-wn-sl7:latest` singularity image
- Opt out by either:
 - Specifying singularity image:
“`--singularity-image=/path/to/singularity/image`”
 - Passing “`--no-singularity`”: Site-dependent. To truly get outside a singularity container, pass `--no-singularity` and request a site that you know does not run singularity containers
- Future: Will have `--apptainer-image` and `--no-apptainer` flags

Summary and Demo

- Now have `--onsite` and `--offsite` flags
 - To run everywhere, don't specify either – RECOMMENDED
- For more information, please pass the `--help` flag to each `jobsub` command
- Demo

DAGs

Submit DAGs

- jobsub_lite supports dagnabbit syntax to describe DAGs
- Example file mywork.dagnabbit:

```
<serial>
jobsub_submit  --mail_on_error $SUBMIT_FLAGS file://jobA.sh
jobsub_submit  --mail_on_error $SUBMIT_FLAGS file://jobB.sh
</serial>
<parallel>
jobsub_submit  --mail_on_error $SUBMIT_FLAGS file://jobC.sh
jobsub_submit  --mail_on_error $SUBMIT_FLAGS file://jobD.sh
</parallel>
<serial>
jobsub_submit  --mail_on_error $SUBMIT_FLAGS file://jobE.sh
</serial>
```

Submit DAGs, continued

- Submit DAG:

```
export SUBMIT_FLAGS="-G fermilab"  
jobsub_submit $SUBMIT_FLAGS --dag file://mywork.dagnabbit
```

- Demo

Tarfiles

-f and --tar-file-name

- All use Rapid Code Distribution Service (RCDS) via CVMFS by default
- *--tar-file-name*: specify TAR_FILE or DIRECTORY to be transferred to worker node
 - TAR_FILE will be accessible to the user job on the worker node via the environment variable \$INPUT_TAR_FILE
 - The unpacked contents will be in the same directory as \$INPUT_TAR_FILE
 - Successive --tar_file_name options will be in \$INPUT_TAR_FILE_1, \$INPUT_TAR_FILE_2, etc.
 - Use with `dropbox://` for pre-made tarfile, `tardir://` to specify directory to be tarred up

-f and --tar-file-name (2)

- *-f*: Copy INPUT_FILE file at runtime
- INPUT_FILE copied to directory \$CONDOR_DIR_INPUT on the execution node.
- Example :
-f /grid/data/minerva/my/input/file.xxx
copied to \$CONDOR_DIR_INPUT/file.xxx
- Specify as many -f INPUT_FILE_1 -f INPUT_FILE_2 args as you need.
- To copy file at submission time use -f dropbox://INPUT_FILE to copy the file

Demo

Fetching Logs

jobsub_fetchlog

- Works like before

```
$ jobsub_fetchlog -G fermilab --jobid 1557@jobsubdevgpvm01.fnal.gov
$ tar tzvf 1557@jobsubdevgpvm01.fnal.gov.tgz
-rw-rw-r-- mengel/mengel  2426 2022-10-18 16:18 simple.cmd
-rw-r--r-- mengel/mengel  1676 2022-10-18 16:18
lookaround.sh2022_10_18_15153647f934e7-f318-4fbe-9eef-aa0e00934401cluster.1557.0.log
-rw-r--r-- mengel/mengel   424 2022-10-18 16:18
lookaround.sh2022_10_18_15153647f934e7-f318-4fbe-9eef-aa0e00934401cluster.1557.0.err
-rw-r--r-- mengel/mengel  8117 2022-10-18 16:18
lookaround.sh2022_10_18_15153647f934e7-f318-4fbe-9eef-aa0e00934401cluster.1557.0.out
-rwxr-xr-x mengel/mengel  6418 2022-10-18 16:18 simple.sh
```

- Can also place in specified directory
- In progress: web-accessible logs for Fifemon/POMS
- Demo

Condor commands

Using Condor commands

- One major change with `jobsub_lite` is that users have access to condor commands
- We recommend users use the `jobsub_lite`-wrapped condor commands, as they handle authentication, but using HTCondor-provided condor commands is an option
- Demo

Production Jobs and Managed Tokens

Roles in token-world

- No VOMS-server signing proxies in the token-world
- Role = entry in “wlcg.groups” entry of token
 - This entry is mapped to “capability set” in LDAP/FERRY, which defines your “scopes” entry
 - “scopes” controls authorization
- Production tokens usually have access to read/write to ALL of an experiment’s dCache area, but this is configurable

Managed Tokens

- New service to push production vault tokens to interactive nodes, keep them refreshed
- Production users should *never* have to authenticate in CILogon
 - The pushed vault token is used to obtain bearer token
- Set in environment:

```
export
```

```
HTGETTOKENOPTS="--credkey=<account>/managedtokens/fifeutilgpvm01.fnal.gov"
```

```
export X509_USER_PROXY=/path/to/production/proxy
```

- Then, in `jobsub_*` command, pass `--role=production` (note lower-case "p")

Demo

File Transfer in Jobs (jobsub_lite and IFDHC)

File Transfer in jobs

- You will have `storage` permissions from your token:

```
$ httokencode /run/user/1733/bt_u1733
{
  "wlcg.ver": "1.0",
  "aud": "https://wlcg.cern.ch/jwt/v1/any",
  "scope": "storage.create:/dune/scratch/users/mengel storage.read:/dune ..."
  ...
}
```

- You can use `ifdh` commands as usual
- Most experiments/roles do not have `storage.modify` scope
 - cannot overwrite or remove files
- Different roles can have different scope
- For transition, `IFDH_FORCE=gsiftp` will use proxies

lfdh token demo

```
<bel-kwinith>$ setup ifdhc
```

```
<bel-kwinith>$ export IFDH_TOKEN_ENABLE=1
```

```
<bel-kwinith>$ export EXPERIMENT=dune
```

```
<bel-kwinith>$ ifdh ls /pnfs/dune/scratch/users/mengel
```

```
Attempting kerberos auth with https://htvaultprod.fnal.gov:8200 ...  
succeeded
```

```
Attempting to get token from https://htvaultprod.fnal.gov:8200 ...  
succeeded
```

```
Storing vault token in /tmp/vt_u1733
```

```
Storing bearer token in /tmp/bt_token_dune_Analysis_1733
```

```
/pnfs/dune/scratch/users/mengel/
```

```
/pnfs/dune/scratch/users/mengel/0504/
```

```
/pnfs/dune/scratch/users/mengel/bar.txt
```

```
/pnfs/dune/scratch/users/mengel/baz.txt
```

SAM

SAM

- The SAM webservice accepts tokens as well as proxies.
- Ifdhc (v2_6_6 and later) does token authentication to SAM
 - If BEARER_TOKEN_FILE is set to a token file (in jobs)
 - If IFDH_TOKEN_ENABLE=1
 - In future the latter will be the default
- Next sam_web_client release (samweb commands) will do token authentication
 - If BEARER_TOKEN_FILE is set to a token file

POMS

POMS and jobsub_lite

🏠 POMS

Campaign Search...

External Links

[POMS User Documentation](#)

[POMS ServiceNow Page](#)

[Samdev Logbook](#)

[Downtime Calendar](#)

Experiment

[Membership \(POMS\)](#)

User Data

[Uploaded Files](#)

[Held Launches](#)

Campaign Data

[Campaigns](#)


[Campaign Stages](#)

[Sample Campaigns](#)

Configure Work

[Compose Login/Setup](#)

Current POMS needs some added setup to use jobsub_lite (this will be rolled into an upcoming release). So to use it:

- Click “Compose Login/Setup”
- Clone  a current config as name_lite
- Add to the setup
 - for Production Users

```
export HTGETTOKENOPTS="--credkey=%(experiment)spro/managedtokens/fifeutilgpvm01.fnal.gov"
```

```
export _condor_SEC_CREDENTIAL_STORER=/bin/true;
```

```
htgettoken -i fermilab -a htvaultprod.fnal.gov;
```

```
setup jobsub_client v_lite
```

Poms and jobsub_lite (cont)

- For Analysis campaigns

```
export USER=%(experimenter)s;  
export XDG_CACHE_HOME=/tmp/%(experimenter)s;  
export HTGETTOKENOPTS="--vaulttokeninfile=$UPLOADS/vt_%(experiment)s_Analysis_$USER --credkey $USER";  
export _condor_SEC_CREDENTIAL_STORER=/bin/true;  
export BEARER_TOKEN_FILE=/tmp/bt$$;  
htgettoken -i %(experiment)s -a htvaultprod.fnal.gov;  
setup jobsub_client v_lite
```

- Also need to upload a vault token as well as proxy

```
. /grid/fermiapp/products/common/etc/setups  
setup poms_client v4_5_0  
upload_file --vaulttoken --experiment= EXP  
upload_file --proxy --experiment= EXP
```

- POMS otherwise works as usual
- Links to job logs don't work (yet) use `jobsub_fetchlog`.
- Demo

Deployment Overview

- November 2022: Iron out deployment details with mu2e
- December 2022:
 - Deploy to experiment “test” interactive nodes, get feedback
 - Announce to general users the go-live date
- January 2023: Run two more demos of jobsub_lite
- **February 1, 2023:**
 - **Go-live of jobsub_lite (see next slide)**
 - Plan changed from before
- April 19, 2023: Turn off jobsub servers → jobsub_client will no longer work

Phases of Rollout

- Now: All experiments that have access to `jobsub_lite` should be testing.
 - Production teams especially!
- February 1, 2023
 - Go-live! `jobsub_lite` available on all interactive nodes
 - POMS will use `jobsub_lite` by default
- February 15, 2023
 - `jobsub_client v_lite` will be made “current” - all users default to `jobsub_lite`
- March 15, 2023
 - `jobsub_servers` stop accepting new jobs
 - Fetching logs, managing jobs, still possible
- April 19, 2023
 - `jobsub_servers` shut down. FULLY `jobsub_lite`

Resources

- Github repo: https://github.com/marcmengel/jobsub_lite
- Roadmap: https://cdcvs.fnal.gov/redmine/projects/jobsub/wiki/Jobsub_lite_Tasks
- Getting Started: https://fifewiki.fnal.gov/wiki/Getting_started_with_jobsub_lite

Thank you!

The jobsub project team:

Shreyas Bhat, Joe Boyd, Lisa Goodenough, Marc Mengel, Nick Peregonow, Kevin Retzke, Dennis Box

Extra Slides

Moving and Deleting Files from Grid Jobs

- For a bearer token to authorize a user to move or remove a file, it must have the **storage.modify** scope on the path containing the file
- Due to security concerns, this is *not* granted to users by default; it must be requested on the `jobsub_submit` command line
- This is done with the “`--need-storage-modify <path>`” flag
- `jobsub_lite` will evaluate whether the `storage.modify` request is valid
- Similarly, “`--need-scope <scope>`” will request a scope be added to token

Moving and Deleting Files from Grid Jobs (2)

Examples:

- Request storage.modify on /pnfs/mu2e/scratch/users/username

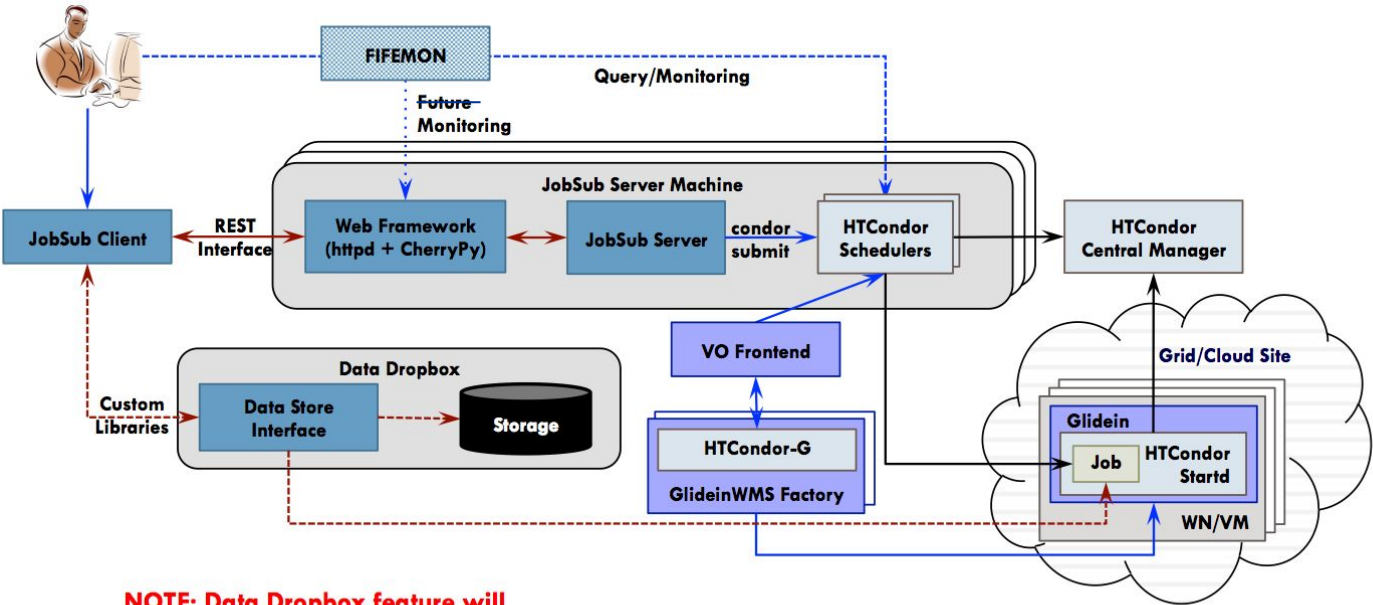
```
$ jobsub_submit -G mu2e --need-storage-modify /mu2e/scratch/users/username file:///bin/true
```

- Demo

Jobsub the Old

- FIFE batch job submission software for HTCondor batch system (<https://research.cs.wisc.edu/htcondor>)
- Two components
 - `jobsub_client`
 - `jobsub_server`
- `jobsub_client` generally installed on experiment interactive nodes
- `jobsub_server`, alongside HTCondor schedd run on separate machines (3 in production cluster)

Jobsub the Old (2)



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

Original Image Source: https://cdcvs.fnal.gov/redmine/projects/fife/wiki/Introduction_to_FIFE_and_Component_Services#Jobsub

Problems with Current Jobsub

- Interface too broad: Too many ways to do the same set of operations (e.g. tarball upload)
- Too much code customization:
 - “Nova” settings, “Minerva” settings, etc., in codebase
- >21k lines of code (not including packaging scripts, tests, etc.)
- i.e. Supporting current feature set too difficult for available effort
 - Also complicates building new features

jobsub_lite and HTCondor

- Idea is to keep jobsub_lite....light
- We will provide lightly-wrapped HTCondor executables (condor_submit, condor_q, etc.) on interactive nodes
- Most users: jobsub commands
- Advanced use-cases: condor commands
- To facilitate, jobsub_submit will have option to just create condor Job Definition File (no submission) to use with condor_submit

Benefit for Experiment Users

- Users will have full access to condor commands
 - No more passing through constraints through jobsub
 - Users will have access to condor JDFs
 - Don't have to wrap condor DAG commands
- With more focused interface, it should be easier to get new users started using jobsub_lite
- Horizontal scaling (adding more schedds) much easier