

## JobSub - Feature #24329

### jobsub\_lite

04/20/2020 08:05 PM - Dennis Box

<b>Status:</b>	New	<b>Start date:</b>	04/20/2020
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	Marc Mengel	<b>% Done:</b>	0%
<b>Category:</b>		<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>		<b>Spent time:</b>	0.00 hour
<b>Stakeholders:</b>			
<b>Description</b>			

### History

#### #1 - 04/20/2020 08:08 PM - Dennis Box

- Assignee set to Marc Mengel

#### #2 - 04/23/2020 10:32 AM - Marc Mengel

The prototype on branch 24329 now generates a submit file and wrapper script that looks like it would work for simple submissions, and prints the condor\_submit line it would use.

I need to go back through the templates for various options, especially the various dag-based options and especially the dagnabbit bits, and also the cvmfs tarball upload service calls.

Questions at the moment:

- I have a "get\_schedd()" call that currently just picks from jobsub01..3 at random; how should the thin-client jobsub pick a schedd? Should the available schedd's be listed in a config, or should there be a web page we should check, or???
- Do we want to support resilient DCache for tarball dropoff, or only the cvmfs publication service (which is how I have it at the moment).

#### #3 - 05/05/2020 02:57 PM - Dennis Box

Recipe for submitting using raw condor commands on jobsubdevgpvm01. According to condor\_q -better-analyze, these jobs will eventually run

1) make a proxy using kx509 or voms-proxy-init or whatever. In this example its called 'nova\_proxy'

2) scp nova\_proxy rexbatch@jobsubdevgpvm01:dbbox/submit\_dir/testjobs

3) ssh [rexbatch@jobsubdevgpvm01.fnal.gov](mailto:rexbatch@jobsubdevgpvm01.fnal.gov)

```
-bash-4.2$ cd dbbox/submit_dir/testjobs
```

```
-bash-4.2$ cat testjob.jdf
universe = vanilla
```

```
executable = system-info.sh
output = joboutput/out.$(cluster).$(process)
error = joboutput/err.$(cluster).$(process)
log = joboutput/log.$(cluster).$(process)
```

```
should_transfer_files = YES
when_to_transfer_output = ON_EXIT_OR_EVICT
```

```
#x509userproxy=/local/home/testuser/security/grid_proxy
x509userproxy=nova_proxy
```

```
Requirements =
```

```
queue 3
```

```
-bash-4.2$ condor_submit testjob.jdf
Submitting job(s)...
3 job(s) submitted to cluster 39107.
```

```
-bash-4.2$ condor_q rexbatch
```

-- Schedd: jobsubdevgpvm01.fnal.gov : <131.225.240.23:9615?... @ 05/05/20 14:54:50

OWNER BATCH\_NAME SUBMITTED DONE RUN IDLE TOTAL JOB\_IDS

rexbatch ID: 39107 5/5 14:42 \_ \_ 3 3 39107.0-2

Total for query: 3 jobs; 0 completed, 0 removed, 3 idle, 0 running, 0 held, 0 suspended

Total for all users: 93 jobs; 0 completed, 0 removed, 93 idle, 0 running, 0 held, 0 suspended

bash-4.2\$

bash-4.2\$ ls -lart joboutput/

total 20

drwxr-xr-x 3 rexbatch fife 4096 May 5 14:42 ..

-rw-r--r- 1 rexbatch fife 1113 May 5 14:42 log.39107.1

-rw-r--r- 1 rexbatch fife 1113 May 5 14:42 log.39107.0

drwxr-xr-x 2 rexbatch fife 4096 May 5 14:42 .

-rw-r--r- 1 rexbatch fife 1113 May 5 14:42 log.39107.2

bash-4.2\$ condor\_q -better-analyze 39107.0

-- Schedd: jobsubdevgpvm01.fnal.gov : <131.225.240.23:9615?...

The Requirements expression for job 39107.000 is

```
(TARGET.Arch "X86_64") && (TARGET.OpSys "LINUX") && (TARGET.Disk >= RequestDisk) && (TARGET.Memory >= RequestMemory) && (TARGET.HasFileTransfer)
```

Job 39107.000 defines the following attributes:

```
JobUniverse = 5
RequestDisk = ifThenElse(JobUniverse != 7,10000000,16000)
RequestMemory = ifThenElse(JobUniverse != 7,2000,10)
```

The Requirements expression for job 39107.000 reduces to these conditions:

Step	Matched	Condition
[0]	15923	TARGET.Arch "X86_64"
[1]	15923	TARGET.OpSys "LINUX"
[3]	14079	TARGET.Disk >= RequestDisk
[5]	9509	TARGET.Memory >= RequestMemory
[6]	8551	[3] && [5]

No successful match recorded.

Last failed match: Tue May 5 14:44:17 2020

Reason for last match failure: no match found

39107.000: Run analysis summary ignoring user priority. Of 2135 machines, 97 are rejected by your job's requirements 1721 reject your job because of their own requirements 0 match and are already running your jobs 0 match but are serving other users 317 are able to run your job

bash-4.2\$

#### #4 - 06/25/2020 04:27 PM - Marc Mengel

Notes from jobsub\_lite code review:

- Dictionary from arguments -- use vars(x) over x.\_\_dict\_\_
- Maybe move arg parser to separate file: many loc, not otherwise complicated.
- Possibly other splitouts? Tarfiles? dagnabbit parser?
- File versus params on command line -- if fixed lose if's about it and debug prints
- Add unit test/example comments to methods, esp. unit converter; would make it clearer.
- Document taking either upper/lower case.
- verify file: prefix in argparse(?) use type=func for argparse to check it.
- Look at using condor bindings to check what schedd's are there,
- and possibly also for submitting.
  - <https://htcondor.readthedocs.io/en/latest/apis/python-bindings/api/htcondor.html>
- Use rdcs round-robin DNS and drop picking server from list on tarball upload
- Review n ways to upload files: question: still support drobpox:/tardir: on -f parameters?
- Authentication: for current stuff need cigetcert/myproxy call;
- for future stuff need suitable tokens bits.
- Check more if dagnabbit should support nested <parallel> ... <serial> </serial> </parallel>

15. typo requirement in template files.
16.  Looks much easier to maintain, and easier to find things.
17. Do we want % formatting or .format() ? I'm in the habit of % formatting, bears discussion.
18.  Should we use os.path() versus formatted paths?
19. Put future suggestions in ticket; make subtickets on this ticket or obvious stuff just fix and commit.

Thanks to Vito, Bruno, and especially Shreyas for attending.

**#5 - 09/23/2020 10:49 AM - Marc Mengel**

So there is an example of using scitokens in a submit file:

[https://github.com/htcondor/scitokens-credmon/blob/master/examples/submit/scitoken\\_example/single\\_scitoken.submit](https://github.com/htcondor/scitokens-credmon/blob/master/examples/submit/scitoken_example/single_scitoken.submit)

which is supposed to let you say what token bits your job needs (?).  
Do we add these to our jobsub\_lite templates?

There are notes at [scitokens-credmon: Testing the credmon...](#)

Also <https://cdcvs.fnal.gov/redmine/projects/glideinwms/wiki/CredMonScitokensDocker>

**#6 - 09/23/2020 11:18 AM - Marc Mengel**

Also adding [Jobsub\\_lite\\_design\\_document](#) per Tanya's request.