

FIFE - Task #17434

Milestone # 16417 (Assigned): Prepare FIFE tools to handle HEP Cloud Submission

Task # 16423 (Closed): Prepare GPGGrid for HEP Cloud

Test monitoring with the refactored gpggrid

08/11/2017 10:37 AM - Joe Boyd

Status:	Closed	Start date:	08/11/2017
Priority:	Normal	Due date:	
Assignee:	Kevin Retzke	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Description			
Need to point probes at the test gpggrid cluster that has glideins and directly attached worker nodes and see what breaks in the monitoring.			

History

#1 - 08/11/2017 11:39 AM - Kevin Retzke

- Status changed from New to Assigned

Cluster info:

```
[rexbatch@htcjsdev01 ~]$ condor_status -schedd
Name Machine RunningJobs IdleJobs HeldJobs
htccdev01.fnal.gov htccdev01.fnal.gov 0 0 0
htcjsdev01.fnal.gov htcjsdev01.fnal.gov 3 5 9
htcjsdev02.fnal.gov htcjsdev02.fnal.gov 19 18 0
TotalRunningJobs TotalIdleJobs TotalHeldJobs
Total 22 23 9
```

```
[rexbatch@htcjsdev01 ~]$ condor_status -collector
Name Machine RunningJobs
IdleJobs HostsTotal
DFC@htccolldev01.fnal.gov htccolldev01.fnal.gov 0
0 0
DFC@htccolldev02.fnal.gov htccolldev02.fnal.gov 19
23 32
[rexbatch@htcjsdev01 ~]$ condor_status -negotiator
Name LastCycleEnd (Sec) Slots Submits Schedds Jobs Matches Rejections
NEGOTIATORFIFEQUOTAS@htccolldev01.fnal.gov 8/11 11:29 0 1 2 2 22 0
2
htccolldev01.fnal.gov 8/11 11:26 0 2 0 0 0 0
```

condor_status queries against the two collectors report different machines...

```
[ifmon@fermicloud147 ~]$ condor_status -pool htccolldev02
Name OpSys Arch State Activity LoadAv Mem ActvtyTime
slot1@gcotest01.fnal.gov LINUX X86_64 Unclaimed Idle 0.000 23949 0+21:02:41
slot1@gcotest02.fnal.gov LINUX X86_64 Owner Idle 0.060 23949 0+18:22:51
slot1@gcotest03.fnal.gov LINUX X86_64 Unclaimed Idle 0.140 23949 3+21:32:47
slot1@gcotest04.fnal.gov LINUX X86_64 Unclaimed Idle 0.010 23949 3+21:29:36
slot1@htcwndev01.fnal.gov LINUX X86_64 Unclaimed Idle 0.040 1995 0+13:09:00
slot1_1@htcwndev01.fnal.gov LINUX X86_64 Claimed Busy 0.000 1700 0+01:04:33
slot1@htcwndev02.fnal.gov LINUX X86_64 Unclaimed Idle 0.000 1995 3+02:56:33
slot1_1@htcwndev02.fnal.gov LINUX X86_64 Claimed Busy 0.010 1700 0+01:08:27
slot1@wnitb001.fnal.gov LINUX X86_64 Unclaimed Idle 0.010 3695 6+20:46:27
slot1@wnitb002.fnal.gov LINUX X86_64 Unclaimed Idle 0.460 3695 2+21:14:53
slot1@wnitb003.fnal.gov LINUX X86_64 Unclaimed Idle 0.020 7822 3+14:09:52
```

slot1@wnitb004.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.040	7822	3+13:24:16
slot1@wnitb005.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.020	7822	2+03:22:04
slot1@wnitb006.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.060	7822	2+00:57:39
slot1@wnitb007.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.050	7822	2+01:01:06
slot1@wnitb008.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.150	7822	2+00:55:00
slot1@wnitb009.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.200	7822	2+00:36:31
slot1@wnitb010.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.190	7822	2+00:32:17
slot1@wnitb011.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.000	7822	1+23:56:09
slot1@wnitb012.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.000	7822	1+23:52:53
slot1@wnitb013.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.320	7822	1+20:48:27
slot1@wnitb014.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.000	7822	1+20:52:05
slot1@wnitb015.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.060	7822	1+20:48:56
slot1@wnitb016.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.030	7822	0+16:43:26
slot1@wnitb017.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.020	7822	1+20:29:39
slot1@wnitb018.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.000	7822	1+20:30:33
slot1@wnitb019.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.010	7822	1+20:27:38
slot1@wnitb020.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.000	7822	1+20:29:26
slot1@wnitb021.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.020	7822	1+20:27:12
slot1@wnitb023.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.000	7822	1+20:25:28
slot1@wnitb024.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.080	7822	0+16:43:37
slot1@wnitb025.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.080	7822	1+20:23:27

Machines Owner Claimed Unclaimed Matched Preempting Drain

X86_64/LINUX 32 1 2 29 0 0 0

Total 32 1 2 29 0 0 0

[ifmon@fermicloud147 ~]\$ condor_status -pool htccolldev01

Name	OpSys	Arch	State	Activity	LoadAv	Mem
glidein_2_426874698@CRUSH-SUGWG-OSG-10-5-153-200	LINUX	X86_64	Claimed	Busy	19.660	25
00 0+00:50:48						
slot1@glidein_1_627302880@fermilab-518031.0-red-c1529.unl.edu	LINUX	X86_64	Claimed	Busy	0.310	14
95 0+01:04:46						
slot2@glidein_1_627302880@fermilab-518031.0-red-c1529.unl.edu	LINUX	X86_64	Claimed	Busy	0.000	14
95 0+01:04:47						
slot3@glidein_1_627302880@fermilab-518031.0-red-c1529.unl.edu	LINUX	X86_64	Claimed	Busy	0.000	14
95 0+01:04:48						
slot4@glidein_1_627302880@fermilab-518031.0-red-c1529.unl.edu	LINUX	X86_64	Claimed	Busy	0.000	14
95 0+01:04:48						
slot5@glidein_1_627302880@fermilab-518031.0-red-c1529.unl.edu	LINUX	X86_64	Claimed	Busy	0.000	14
95 0+01:04:49						
slot6@glidein_1_627302880@fermilab-518031.0-red-c1529.unl.edu	LINUX	X86_64	Claimed	Busy	0.000	14
95 0+01:04:49						
slot7@glidein_1_627302880@fermilab-518031.0-red-c1529.unl.edu	LINUX	X86_64	Claimed	Busy	0.000	14
95 0+01:04:48						
slot8@glidein_1_627302880@fermilab-518031.0-red-c1529.unl.edu	LINUX	X86_64	Claimed	Busy	0.000	14
95 0+01:04:40						
slot9@glidein_1_627302880@fermilab-518031.0-red-c1529.unl.edu	LINUX	X86_64	Claimed	Busy	0.000	14
95 0+01:04:43						
slot10@glidein_1_627302880@fermilab-518031.0-red-c1529.unl.edu	LINUX	X86_64	Claimed	Busy	0.000	14
95 0+01:04:45						
slot11@glidein_1_627302880@fermilab-518031.0-red-c1529.unl.edu	LINUX	X86_64	Claimed	Busy	0.000	14
95 0+01:04:43						
slot12@glidein_1_627302880@fermilab-518031.0-red-c1529.unl.edu	LINUX	X86_64	Claimed	Busy	0.000	14
95 0+01:04:47						
slot13@glidein_1_627302880@fermilab-518031.0-red-c1529.unl.edu	LINUX	X86_64	Claimed	Busy	0.000	14
95 0+01:04:47						
slot14@glidein_1_627302880@fermilab-518031.0-red-c1529.unl.edu	LINUX	X86_64	Claimed	Busy	0.000	14
95 0+01:04:45						
slot15@glidein_1_627302880@fermilab-518031.0-red-c1529.unl.edu	LINUX	X86_64	Claimed	Busy	0.000	14
95 0+01:04:48						
slot16@glidein_1_627302880@fermilab-518031.0-red-c1529.unl.edu	LINUX	X86_64	Claimed	Busy	0.000	14
95 0+01:04:39						
slot1@gcotest01.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.000	239
49 0+21:02:41						
slot1@gcotest02.fnal.gov	LINUX	X86_64	Owner	Idle	0.060	239
49 0+18:22:51						
slot1@gcotest03.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.140	239
49 3+21:32:47						
slot1@gcotest04.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.010	239
49 3+21:29:36						
slot1@htcwndev01.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.000	19
95 0+13:13:40						

slot1_1@htcwndev01.fnal.gov	LINUX	X86_64	Claimed	Busy	0.000	17
00 0+01:09:13						
slot1@htcwndev02.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.000	19
95 3+02:56:33						
slot1_1@htcwndev02.fnal.gov	LINUX	X86_64	Claimed	Busy	0.010	17
00 0+01:08:27						
slot1@wnitb001.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.010	36
95 6+20:46:27						
slot1@wnitb002.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.460	36
95 2+21:14:53						
slot1@wnitb003.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.020	78
22 3+14:09:52						
slot1@wnitb004.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.040	78
22 3+13:24:16						
slot1@wnitb005.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.020	78
22 2+03:22:04						
slot1@wnitb006.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.060	78
22 2+00:57:39						
slot1@wnitb007.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.050	78
22 2+01:01:06						
slot1@wnitb008.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.150	78
22 2+00:55:00						
slot1@wnitb009.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.200	78
22 2+00:36:31						
slot1@wnitb010.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.190	78
22 2+00:32:17						
slot1@wnitb011.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.000	78
22 1+23:56:09						
slot1@wnitb012.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.000	78
22 1+23:52:53						
slot1@wnitb013.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.320	78
22 1+20:48:27						
slot1@wnitb014.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.000	78
22 1+20:52:05						
slot1@wnitb015.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.060	78
22 1+20:48:56						
slot1@wnitb016.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.030	78
22 0+16:43:26						
slot1@wnitb017.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.020	78
22 1+20:29:39						
slot1@wnitb018.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.000	78
22 1+20:30:33						
slot1@wnitb019.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.000	78
22 1+20:31:32						
slot1@wnitb020.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.000	78
22 1+20:29:26						
slot1@wnitb021.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.020	78
22 1+20:27:12						
slot1@wnitb023.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.000	78
22 1+20:25:28						
slot1@wnitb024.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.080	78
22 0+16:43:37						
slot1@wnitb025.fnal.gov	LINUX	X86_64	Unclaimed	Idle	0.080	78
22 1+20:23:27						

	Machines	Owner	Claimed	Unclaimed	Matched	Preempting	Drain
X86_64/LINUX	49	1	19	29	0	0	0
Total	49	1	19	29	0	0	0

#2 - 08/11/2017 02:57 PM - Kevin Retzke

Collector statistics and Slot status being collected by dev probe on fermicloud147: <https://fifemon-pp.fnal.gov/dashboard/db/gpgrid2017>

Trouble talking to schedds and negotiator:

```

2017-08-11 14:55:19,056 [INFO] __main__ - querying pool htccolldev01.fnal.gov priorities
2017-08-11 14:55:19,064 [WARNING] root - Trouble communicating with pool htccolldev01.fnal.gov negotiator, re
trying in 10s.
2017-08-11 14:55:29,081 [WARNING] root - Trouble communicating with pool htccolldev01.fnal.gov negotiator, re
trying in 10s.
2017-08-11 14:55:39,092 [ERROR] root - Trouble communicating with pool htccolldev01.fnal.gov negotiator, givin
g up.
2017-08-11 14:55:39,094 [INFO] __main__ - querying pool htccolldev01.fnal.gov jobs

```

Traceback (most recent call last):

```
File "/home/ifmon/fifemon-probes/bin/condor/jobs.py", line 98, in get_jobs
    results = schedd.query(constraint, attrs)
IOError: Failed to fetch ads from schedd.
2017-08-11 14:55:39,124 [WARNING] condor.jobs - Trouble communicating with schedd htccdev01.fnal.gov, retrying in 10s.
```

condor_userprio works though, so maybe something with the Python libs?

```
condor_userprio -pool htccolldev01
Last Priority Update: 8/11 14:58
```

Group	Config	Use	Effective	Priority	Res	Total Usage	Time Since Requested
User Name	Quota	Surplus	Priority	Factor	In Use	(wghted-hrs)	Last Usage Resources
group_mu2e	8.00	no		100000.00	0	40.92	0+22:31
1							
timmm@fnal.gov			50000.00	100000.00	0	40.92	0+22:31
group_fermilab	0.00	Regroup		100.00	0	0.70	0+00:16
0							
kretzke@fnal.gov			51.55	100.00	0	0.70	0+00:16
<none>	0.00	yes		100000.00	92	4004.81	<now>
222							
group_minerva.dbox@fnal.gov			56374.90	100000.00	4	188.45	<now>
group_annie.dbox@fnal.gov			57739.60	100000.00	0	14.81	0+00:21
group_dune.timm@fnal.gov			1003981.50	100000.00	87	685.20	<now>
group_mu2e.timm@fnal.gov			1113860.00	100000.00	1	1171.69	<now>
osg@fnal.gov			8.49617e+17	1e+18	10	12.98	<now>

Number of users:	7	Regroup			102	2114.75	0+23:59

condor_q does not work, looks like authentication for firewall:

```
condor_q -pool htccolldev01 -global

-- Failed to fetch ads from: <131.225.154.92:9615?addrs=131.225.154.92-9615+[--1]-9615&noUDP&sock=592868_0725_3> : htcjsdev02.fnal.gov
AUTHENTICATE:1003:Failed to authenticate with any method
AUTHENTICATE:1004:Failed to authenticate using GSI
GSI:5003:Failed to authenticate. Globus is reporting error (851968:40). There is probably a problem with your credentials. (Did you run grid-proxy-init?)
AUTHENTICATE:1004:Failed to authenticate using FS

-- Failed to fetch ads from: <131.225.154.241:9618?addrs=131.225.154.241-9618+[--1]-9618&noUDP&sock=1157756_35e6_3> : htccdev01.fnal.gov
AUTHENTICATE:1003:Failed to authenticate with any method
AUTHENTICATE:1004:Failed to authenticate using GSI
GSI:5003:Failed to authenticate. Globus is reporting error (851968:81). There is probably a problem with your credentials. (Did you run grid-proxy-init?)
AUTHENTICATE:1004:Failed to authenticate using FS

-- Failed to fetch ads from: <131.225.154.88:9615?addrs=131.225.154.88-9615+[--1]-9615&noUDP&sock=462763_d1d0_3> : htcjsdev01.fnal.gov
AUTHENTICATE:1003:Failed to authenticate with any method
AUTHENTICATE:1004:Failed to authenticate using GSI
GSI:5003:Failed to authenticate. Globus is reporting error (851968:122). There is probably a problem with your credentials. (Did you run grid-proxy-init?)
AUTHENTICATE:1004:Failed to authenticate using FS
```

#3 - 08/11/2017 03:39 PM - Kevin Retzke

I was able to authenticate with existing gpgrid certificate (same used to query gpcollector01/02). Copied "Fifebatch" dashboard: <https://fifemon-pp.fnal.gov/dashboard/db/gpgrid2017-jobs>

#4 - 09/06/2017 12:07 PM - Kevin Retzke

Looks like the issue getting priorities is that the python libs try to talk directly to the negotiator, while condor_userprio will get them from the collector by default.

```
>>> ads=coll.locateAll(htcondor.DaemonTypes.Negotiator)
>>> n=htcondor.Negotiator(ads[0])
>>> n.getPriorities()
09/06/17 11:45:09 attempt to connect to <131.225.154.243:30075> failed: Connection refused (connect errno = 11
1).
```

```
[kretzke@fermicloud147 ~]$ condor_userprio -pool htccolldev01 -debug
```

```
Last Priority Update: 9/6 11:53
```

Group	Config	Use	Effective	Priority	Res	Total Usage	Time Since	Requested
User Name	Quota	Surplus	Priority	Factor	In Use	(wghted-hrs)	Last Usage	Resources
group_admin	0.00	Regroup		1.00	0	406.09	0+18:13	0
timmm@fnal.gov			8.25911e+17	1e+18	4	357.27	<now>	
<none>	0.00	yes		100000.00	0	4341.44	0+00:44	5

```
Number of users: 1 Regroup 4 357.27 0+23:59
```

```
[kretzke@fermicloud147 ~]$ condor_userprio -pool htccolldev01 -negotiator -debug
```

```
09/06/17 11:50:22 attempt to connect to <131.225.154.243:25110> failed: Connection refused (connect errno = 11
1).
```

```
failed to send GET_PRIORITY command to negotiator
```

Options: enable direct access to the negotiator, or change the probes to just query the collector (not clear how to do that through the Python libs, will investigate).

#5 - 09/06/2017 03:11 PM - Kevin Retzke

We can get the new Accounting classads from the collector, which appear to have all the priority information for "active" users.

```
$ condor_status -pool htccolldev01 -any -constraint 'MyType=="Accounting"' -af MyType Name PriorityFactor Priority
```

```
Accounting <none> 100000.0 50000.0
Accounting group_admin 1.0 1.492117047309875
Accounting group_admin.jeffderb@fnal.gov 1.0 1.492117047309875
Accounting group_admin.jeffderb@fnal.gov 1.0 1.92158305644989
Accounting group_annie 100000.0 50000.0
Accounting group_argoneut 100000.0 50000.0
Accounting group_cdf 100000.0 50000.0
Accounting group_cdf.dbox@fnal.gov 100.0 50.0
Accounting group_cdf.dbox@fnal.gov 100000.0 50000.0
Accounting group_cdms 100000.0 50000.0
Accounting group_chips 100000.0 50000.0
Accounting group_cms 9.999999843067494E+17 4.999999921533747E+17
Accounting group_coupp 100000.0 50000.0
Accounting group_darkside 100000.0 50000.0
Accounting group_des 100000.0 50000.0
Accounting group_dune 100000.0 110743.0
Accounting group_dune.prod 100000.0 50000.0
Accounting group_dune.timm@fnal.gov 9.999999843067494E+17 1.107429966921859E+18
Accounting group_dune.timm@fnal.gov 100000.0 50000.0
```

In Python:

```
>>> import htcondor
>>> coll=htcondor.Collector('htccolldev01')
>>> ads=coll.query(constraint='MyType=="Accounting"')
>>> [ad['Name'] for ad in ads]
['group_seaquest.production', 'group_seaquest', 'group_larl1', 'group_uboone.prod', 'group_minerva.production',
'group_minerva', 'group_mu2e', 'njp@fnal.gov', 'group_admin.jeffderb@fnal.gov', ...]
```

Need to figure out how to represent the two negotiators (currently on fifebatch it just flaps back and forth, depending on which negotiator reported last to the collector).

#6 - 10/12/2017 09:31 AM - Kevin Retzke

- Status changed from Assigned to Closed

There's some issues with the accounting classads being replaced based on the last negotiator to report to the collector (again...), but we have a working solution in place (ignoring 0 quotas from the "wrong" negotiator), and all the monitoring seems to be working as expected (minus the known issues with memory and CPU time that are being tracked in SNOW), so I'm closing this.