

GlideinWMS - Bug #21325

Frontend not recognizing entries in downtime

11/07/2018 10:47 AM - Parag Mhashilkar

Status:	Closed	Start date:	11/07/2018
Priority:	Normal	Due date:	
Assignee:	Lorena Lobato Pardavila	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	v3_4_3	Spent time:	0.00 hour
First Occurred:		Stakeholders:	HEPCloud
Occurs In:	v3_4_2		

Description

It appears that the frontend in glideinwms 3.4.2 does not correctly detect factory entries that are in downtime. Here is the evidence:

frontend cmssrv266 has but one group active, group_cms_all. 200 jobs in the queue, all idle.
It talks only to factory cmssrv280.

From the group log cms_all.info

```
[2018-11-07 09:43:57,669] INFO: Iteration at Wed Nov 7 09:43:57 2018
[2018-11-07 09:43:57,669] INFO: Querying schedd, entry, and glidein status using child processes.
[2018-11-07 09:43:58,130] INFO: All children terminated
[2018-11-07 09:43:58,134] INFO: Jobs found total 200 idle 200 (good 200, old(10min 200, 60min 200), grid 200, voms 0) running 0
[2018-11-07 09:43:58,134] INFO: Group glideins found total 1 limit 20000 curb 19000; of these idle 1 limit 20000 curb 19000 running
0
[2018-11-07 09:43:58,135] INFO: Frontend glideins found total 1 limit 170000 curb 167000; of these idle 1 limit 35000 curb 25000
[2018-11-07 09:43:58,135] INFO: Overall slots found total 747 limit 170000 curb 167000; of these idle 50 limit 35000 curb 25000
[2018-11-07 09:43:58,135] INFO: Updating usermap
[2018-11-07 09:43:58,135] INFO: Match
[2018-11-07 09:43:58,140] INFO: Active forks = 3, Forks to finish = 9
[2018-11-07 09:43:58,155] INFO: Active forks = 3, Forks to finish = 7
[2018-11-07 09:43:58,173] INFO: Active forks = 3, Forks to finish = 4
[2018-11-07 09:43:58,187] INFO: Active forks = 2, Forks to finish = 2
[2018-11-07 09:43:58,200] INFO: Active forks = 0, Forks to finish = 0
[2018-11-07 09:43:58,200] INFO: All children terminated - took 0.0647048950195 seconds
[2018-11-07 09:43:58,200] INFO: Total matching idle 200 (old 10min 200 60min 200) running 0 limit 20000
[2018-11-07 09:43:58,201] INFO:      Jobs in schedd queues      |      Slots      |      Cores      | Glidein Req |
Factory/Entry Information
[2018-11-07 09:43:58,201] INFO: Idle (match eff old uniq) Run ( here max ) | Total Idle Run Fail | Total Idle Run | Idle
MaxRun | State Factory
[2018-11-07 09:43:58,202] INFO: 13( 200 13 13 0) 0( 0 20000) | 0 0 0 0 | 0 0 0 | 1 4 | Up
CMSHTPC_T3_US_Bridges@gfactory_instance_cmssrv280@gfactory_service_cmssrv280@cmssrv280.fnal.gov
[2018-11-07 09:43:58,205] INFO: 13( 200 13 13 0) 0( 0 20000) | 0 0 0 0 | 0 0 0 | 1 4 | Up
CMSHTPC_T3_US_NERSC_Cori@gfactory_instance_cmssrv280@gfactory_service_cmssrv280@cmssrv280.fnal.gov
[2018-11-07 09:43:58,206] INFO: 13( 200 13 13 0) 0( 0 20000) | 0 0 0 0 | 0 0 0 | 1 4 | Up
CMSHTPC_T3_US_NERSC_Cori_KNL@gfactory_instance_cmssrv280@gfactory_service_cmssrv280@cmssrv280.fnal.gov
[2018-11-07 09:43:58,207] INFO: 13( 200 12 13 0) 0( 0 20000) | 1 1 0 0 | 8 8 0 | 1 4 | Up
CMSHTPC_T3_US_NERSC_Cori_shared@gfactory_instance_cmssrv280@gfactory_service_cmssrv280@cmssrv280.fnal.gov
[2018-11-07 09:43:58,208] INFO: 13( 200 13 13 0) 0( 0 20000) | 0 0 0 0 | 0 0 0 | 1 4 | Up
CMSHTPC_T3_US_NERSC_Edison@gfactory_instance_cmssrv280@gfactory_service_cmssrv280@cmssrv280.fnal.gov
[2018-11-07 09:43:58,210] INFO: 13( 200 13 13 0) 0( 0 20000) | 0 0 0 0 | 0 0 0 | 1 7 | Up
CMSHTPC_T3_US_NERSC_Edison_shared@gfactory_instance_cmssrv280@gfactory_service_cmssrv280@cmssrv280.fnal.gov
[2018-11-07 09:43:58,211] INFO: 13( 200 13 13 0) 0( 0 20000) | 0 0 0 0 | 0 0 0 | 1 4 | Up
CMSHTPC_T3_US_TACC@gfactory_instance_cmssrv280@gfactory_service_cmssrv280@cmssrv280.fnal.gov
[2018-11-07 09:43:58,212] INFO: 13( 200 13 13 0) 0( 0 20000) | 0 0 0 0 | 0 0 0 | 1 7 | Up
CMS_T1_US_FNAL_condce2@gfactory_instance_cmssrv280@gfactory_service_cmssrv280@cmssrv280.fnal.gov
[2018-11-07 09:43:58,213] INFO: 13( 200 13 13 0) 0( 0 20000) | 0 0 0 0 | 0 0 0 | 1 7 | Up
CMS_T1_US_FNAL_condce3@gfactory_instance_cmssrv280@gfactory_service_cmssrv280@cmssrv280.fnal.gov
[2018-11-07 09:43:58,215] INFO: 13( 200 13 13 0) 0( 0 20000) | 0 0 0 0 | 0 0 0 | 1 7 | Up
CMS_T1_US_FNAL_condce4@gfactory_instance_cmssrv280@gfactory_service_cmssrv280@cmssrv280.fnal.gov
[2018-11-07 09:43:58,216] INFO: 13( 200 13 13 0) 0( 0 20000) | 0 0 0 0 | 0 0 0 | 1 7 | Up
```

```

CMS_T1_US_FNAL_condce@gfactory_instance_cmssrv280@gfactory_service_cmssrv280@cmssrv280.fnal.gov
[2018-11-07 09:43:58,217] INFO: 13( 200 13 13 0) 0( 0 20000)| 0 0 0 0| 0 0 0| 1 7|Up
FNAL_HEPCLOUD_AWS_us-east-1a_m3.2xlarge@gfactory_instance_cmssrv280@gfactory_service_cmssrv280@cmssrv280.fnal.g
ov
[2018-11-07 09:43:58,218] INFO: 13( 200 13 13 0) 0( 0 20000)| 0 0 0 0| 0 0 0| 2 13|Up
FNAL_HEPCLOUD_AWS_us-east-1a_m3.xlarge@gfactory_instance_cmssrv280@gfactory_service_cmssrv280@cmssrv280.fnal.g
o
[2018-11-07 09:43:58,219] INFO: 13( 200 13 13 0) 0( 0 20000)| 0 0 0 0| 0 0 0| 1 4|Up
FNAL_HEPCLOUD_GOOGLE_us-central1-a_custom-16-32768@gfactory_instance_cmssrv280@gfactory_service_cmssrv280@cm
ssrv280.fnal.gov
[2018-11-07 09:43:58,221] INFO: 13( 200 13 13 0) 0( 0 20000)| 0 0 0 0| 0 0 0| 1 4|Up
FNAL_HEPCLOUD_GOOGLE_us-central1-a_custom-32-65536@gfactory_instance_cmssrv280@gfactory_service_cmssrv280@cm
ssrv280.fnal.gov
[2018-11-07 09:43:58,222] INFO: 13( 200 13 13 0) 0( 0 20000)| 0 0 0 0| 0 0 0| 8 40|Up
FNAL_HEPCLOUD_GOOGLE_us-central1-a_n1-standard-1@gfactory_instance_cmssrv280@gfactory_service_cmssrv280@cmssrv
280.fnal.gov
[2018-11-07 09:43:58,223] INFO:          Jobs in schedd queues          |          Slots          |          Cores          | Glidein Req |
Factory/Entry Information
[2018-11-07 09:43:58,223] INFO: Idle (match eff old uniq) Run ( here max ) | Total Idle Run Fail | Total Idle Run | Idle
MaxRun | State Factory
[2018-11-07 09:43:58,223] INFO: 208( 3200 207 208 0) 0( 0 320k) | 1 1 0 0| 8 8 0| 24 127 | Up
Sum of useful factories
[2018-11-07 09:43:58,224] INFO: 0( 0 0 0 0) 0( 0 0) | 0 0 0 0| 0 0 0| 0 0 | Down Sum of
down factories
[2018-11-07 09:43:58,224] INFO: 0( 0 0 0 0) 0( 0 0) | 0 0 0 0| 0 0 0| 0 0 | Down
Unmatched
[2018-11-07 09:43:58,257] INFO: Advertising global and singular requests for factory cmssrv280.fnal.gov
[2018-11-07 09:43:58,296] INFO: Advertising 16 glideresource classads to the user pool
[2018-11-07 09:43:58,300] INFO: There are 16 classads to advertise
[2018-11-07 09:43:58,449] INFO: Done advertising
[2018-11-07 09:43:58,451] INFO: iterate_one status: None
[2018-11-07 09:43:58,451] INFO: Writing stats

```

But from the factory cmssrv280:

```

[root@cmssrv280 local_hep_gwms_configs]# condor_status -any -constraint 'MyType=="glidefactory" -af Name
Glidein_in_downtime
CMSHTPC_T3_US_Bridges@gfactory_instance_cmssrv280@gfactory_service_cmssrv280 true
CMSHTPC_T3_US_NERSC_Cori@gfactory_instance_cmssrv280@gfactory_service_cmssrv280 true
CMSHTPC_T3_US_NERSC_Cori_KNL@gfactory_instance_cmssrv280@gfactory_service_cmssrv280 true
CMSHTPC_T3_US_NERSC_Cori_shared@gfactory_instance_cmssrv280@gfactory_service_cmssrv280 false
CMSHTPC_T3_US_NERSC_Edison@gfactory_instance_cmssrv280@gfactory_service_cmssrv280 false
CMSHTPC_T3_US_NERSC_Edison_shared@gfactory_instance_cmssrv280@gfactory_service_cmssrv280 false
CMSHTPC_T3_US_TACC@gfactory_instance_cmssrv280@gfactory_service_cmssrv280 true
CMS_T1_US_FNAL_condce2@gfactory_instance_cmssrv280@gfactory_service_cmssrv280 true
CMS_T1_US_FNAL_condce3@gfactory_instance_cmssrv280@gfactory_service_cmssrv280 true
CMS_T1_US_FNAL_condce4@gfactory_instance_cmssrv280@gfactory_service_cmssrv280 true
CMS_T1_US_FNAL_condce@gfactory_instance_cmssrv280@gfactory_service_cmssrv280 true
FIFE_T3_US_NERSC_Cori_KNL@gfactory_instance_cmssrv280@gfactory_service_cmssrv280 true
FIFE_T3_US_NERSC_Edison@gfactory_instance_cmssrv280@gfactory_service_cmssrv280 true
FNAL_HEPCLOUD_AWS_us-east-1a_m3.2xlarge@gfactory_instance_cmssrv280@gfactory_service_cmssrv280 true
FNAL_HEPCLOUD_AWS_us-east-1a_m3.xlarge@gfactory_instance_cmssrv280@gfactory_service_cmssrv280 true
FNAL_HEPCLOUD_GOOGLE_us-central1-a_custom-16-32768@gfactory_instance_cmssrv280@gfactory_service_cmssrv280 true
FNAL_HEPCLOUD_GOOGLE_us-central1-a_custom-32-65536@gfactory_instance_cmssrv280@gfactory_service_cmssrv280 true
FNAL_HEPCLOUD_GOOGLE_us-central1-a_n1-standard-1@gfactory_instance_cmssrv280@gfactory_service_cmssrv280 true

```

So all but three of the listed factory entries in the frontend log above are in downtime.

Yet the frontend is treating them as not in downtime and dividing the load equally among them all. End result is that I get only one idle glidein per entry submitted, for three glideins total, even though there are 200 idle jobs in the queue.

Have to figure out what is happening here, this appears to be a serious bug and leads to gross undersubmission of glideins.

Steve Timm

Related issues:

Related to GlideinWMS - Support #21537: Double-check functions that deal with... **Closed** **12/12/2018**

History

#1 - 11/29/2018 05:51 PM - Lorena Lobato Pardavila

- Status changed from New to Feedback
- Assignee changed from Lorena Lobato Pardavila to Marco Mambelli

The problem was a misleading comparison type of GLIDEIN_In_Downtime (string instead of boolean).
Double checked for other ClassAd attributes and corrected for GLIDEIN_REQUIRE_GLEXEC_USE and GLIDEIN_REQUIRE_VOMS.

Changes done in v35/21325

#2 - 12/07/2018 07:59 PM - Marco Mambelli

- Subject changed from Potential bug in 3.4.2 frontend--not recognizing entries in downtime to Frontend not recognizing entries in downtime
- Assignee changed from Marco Mambelli to Lorena Lobato Pardavila
- Occurs In v3_4_2 added

#3 - 12/07/2018 08:07 PM - Steven Timm

Will a back-patch be available for 3.4.2? We are not yet in a position to run without the switchboard.

Steve Timm

#4 - 12/10/2018 05:40 PM - Lorena Lobato Pardavila

- Assignee changed from Lorena Lobato Pardavila to Marco Mambelli

#5 - 12/12/2018 03:55 PM - Marco Mambelli

- Target version changed from v3_5 to v3_4_3

#6 - 12/12/2018 03:57 PM - Marco Mambelli

- Assignee changed from Marco Mambelli to Lorena Lobato Pardavila
- Stakeholders updated

#7 - 12/12/2018 04:20 PM - Lorena Lobato Pardavila

- Status changed from Feedback to Resolved

Opened [#21537](#) to investigate other function calls which deal with boolean attributes to avoid issues in the future.

#8 - 01/14/2019 03:59 PM - Marco Mascheroni

- Status changed from Resolved to Feedback

Please, review after fix for <https://cdcv.s.fnal.gov/redmine/issues/21527#note-10> (branch v35/21325_1)

#9 - 01/14/2019 08:18 PM - Lorena Lobato Pardavila

- Status changed from Feedback to Resolved

As pointed in <https://cdcv.s.fnal.gov/redmine/issues/21527#note-10>, there was a missing line where `require_voms_proxy` and `require_glidein_glexec_us` were being passed as string.
The trap was the following: By defining `bool([value])` to make sure that we were getting boolean, we didn't see that the value of those variables at that point were string and when:

```
bool(u'False')=True
```

Corrected and tested. We get the correct values now.

Merged into master.

For the record: Added to [#21512](#)

#10 - 01/26/2019 09:12 PM - Marco Mambelli

- Status changed from Resolved to Closed

#11 - 02/15/2019 10:14 AM - Lorena Lobato Pardavila

- Related to Support #21537: Double-check functions that deal with boolean ClassAd facing possible misleading behavior added