

GlideinWMS - Bug #23089

fermifactory02 monitoring is wrong somehow

08/08/2019 02:42 PM - Dennis Box

Status: New	Start date: 08/08/2019
Priority: High	Due date:
Assignee: Marco Mambelli	% Done: 0%
Category:	Estimated time: 0.00 hour
Target version: v3_6_5	Spent time: 0.00 hour
First Occurred:	Stakeholders:
Occurs In:	

Description

From: Steven C Timm <timmm@fnal.gov>
To: Dennis D Box <dbox@fnal.gov>, glideinwms-support <glideinwms-support@fnal.gov>
Subject: Re: fermifactory02 monitoring is wrong somehow
Thread-Topic: fermifactory02 monitoring is wrong somehow
Thread-Index: AQHVSK8x0357RuwcEEKPzcg1zH6dY6boDKUAgAPC1jeABGkQZYABGOURgAAH8S0=
Date: Thu, 8 Aug 2019 09:27:48 -0500

The full restart this time did clear out the high values that were left in the monitoring. But this is reproducible enough that a ticket should still be opened.

Steve

From: Steven C Timm <timmm@fnal.gov>
Sent: Thursday, August 8, 2019 9:10:29 AM
To: Dennis D Box <dbox@fnal.gov>; glideinwms-support <glideinwms-support@fnal.gov>
Subject: Re: fermifactory02 monitoring is wrong somehow

I would like to request that a ticket be opened to track this bug

OK now the queues are completely empty and the factory has only one glidein running.. yet the monitoring still shows thousands of cores running.

This is the last time we had any idle jobs in the queue at all, at 08:18 this morning

```
02@gfactory_service_fermifactory02@fermifactory02.fnal.gov
2019-08-08 08:18:41,164 - decision_engine - glide_frontend_element - GlideinRequestManifests - INFO - 0( 0 0 0 0)
0( 0 60000)| 12 0 10 0| 128 0 128| 0 0|Up INFINITY
CMSHTPC_T3_US_NERSC_Cori@gfactory_instance_fermifactory02@gfactory_service_fermifactory02@fermifactory02.fnal.gov
2019-08-08 08:18:41,172 - decision_engine - glide_frontend_element - GlideinRequestManifests - INFO - 0( 0 0 0 0)
0( 0 60000)| 1000 1 871 0| 8772 8 8722| 0 0|Up INFINITY
CMSHTPC_T3_US_NERSC_Cori_KNL@gfactory_instance_fermifactory02@gfactory_service_fermifactory02@fermifactory02.fnal.g
ov
2019-08-08 08:18:41,950 - decision_engine - glide_frontend_element - GlideinRequestManifests - INFO - 2( 2 2 0 0)
6( 0 15000)| 42 0 33 0| 252 0 252| 0 0|Up 0.0180
CMSHTPC_T3_US_Bridges@gfactory_instance_fermifactory02@gfactory_service_fermifactory02@fermifactory02.fnal.gov
2019-08-08 08:18:42,535 - decision_engine - glide_frontend_element - GlideinRequestManifests - INFO - 2( 2 1 0 0)
6( 0 15000)| 2 1 1 0| 24 16 8| 0 1|Up 0.0063
CMSHTPC_T3_US_SDSC_osg-comet-frontend@gfactory_instance_fermifactory02@gfactory_service_fermifactory02@fermifactory0
2.fnal.gov
```

Whenever there are idle jobs in the queue, the GlideinMonitorTotalClientclassads come back with the same wrong numbers they had before. When there are not idle jobs, within 30 minutes the glideclient and glidefactoryclient classads go away.

I have once again stopped the decision engines and the factory, waited for all classads to go away out of condor_status -any and then restarted.

We will see what happens when jobs next hit the queue.

Steve

From: Steven C Timm <timmm@fnal.gov>
Sent: Wednesday, August 7, 2019 4:19:57 PM
To: Dennis D Box <dbox@fnal.gov>; glideinwms-support <glideinwms-support@fnal.gov>
Subject: Re: fermifactory02 monitoring is wrong somehow

There appear to be two different problems going on:

1) the glidefactory classads are left with unrealistically high values of variables relating to GlideinMonitorTotalClient*

(see cklassad CMSHTPC_T3_US_NERSC_Cori_KNL@gfactory_instance_fermifactory02@gfactory_service_fermifactory02)

```
GlideinMonitorTotalClientMonitorCoresIdle = 8
GlideinMonitorTotalClientMonitorCoresRunning = 8722
GlideinMonitorTotalClientMonitorCoresTotal = 8772
GlideinMonitorTotalClientMonitorGlideIdle = 1
GlideinMonitorTotalClientMonitorGlideRunning = 871
GlideinMonitorTotalClientMonitorGlideTotal = 1000
GlideinMonitorTotalClientMonitorInfoAge = 13
```

The GlideinMonitorTotal variables in that same classad are correct

```
GlideinMonitorTotalStatusRunning = 28
GlideinMonitorTotalStatusRunningCores = 1904
```

There are only 28 glideins running, not 1000. There were at one point 8722 cores running but not anymore, the 1904 value is right.

I believe that something goes wrong with the TotalClientMonitor values above if

there are not glideclient classads being sent all the time. The decision engine only sends

those ads if there are idle jobs in the queue that match the group in question. About half the time at the moment there is nothing idle.

2) The frontend-like code in the decision engine refuses to request new glideins on the basis that a number of jobs in the group are already running. earlier in this email thread that number was accurately reported to be 8545, now it is reported to be 250 or so. Those numbers are accurate.

The problem is that the count of running jobs includes the number of jobs that are running on any entry at all.

of those 250 jobs, 50 of them or so are running at other sites in the OSG, 196 are running at NERSC, none at all are running at SDSC. Also as previously mentioned I have 3 groups defined in the DE for which the jobs almost completely overlap.

If I'm right, then once the amount of running jobs gets down to near zero (later tonight at 11 PM when the last NERSC glideins exit) then SDSC will finally get some glideins submitted again. We will see.

Steve

From: Steven C Timm <timmm@fnal.gov>
Sent: Sunday, August 4, 2019 8:55:44 PM
To: Dennis D Box <dbox@fnal.gov>; glideinwms-support <glideinwms-support@fnal.gov>
Subject: Re: fermifactory02 monitoring is wrong somehow

I restarted the factory and both decision engines that talk to it. Left it down long enough

so that all glideclient and glidefactoryclient classads drained out. On cold restart

we came to exactly the same position we were at before:

Found in channel cms_job_classification

=

	Frontend_Group	Job_Bucket_Criteria_Expr	Site_Bucket_Criteria_Expr	Totals
0	cms_nersc_passthrough	x509UserProxyVOName=='cms' and DESIRED_Sites.str.contains('T3_US_NERSC') and (REQUIRED_OS=='rhel6') and (WMCore_ResizeJob==True)	[u"GLIDEIN_CMSSite=='T3_US_NERSC' and GLIDEIN_Supported_VOs.str.contains('CMS') and GLIDEIN_REQUIRED_OS=='rhel6'"]	539
1	cms_nersc_passthrough_sl7	x509UserProxyVOName=='cms' and DESIRED_Sites.str.contains('T3_US_NERSC') and (REQUIRED_OS=='rhel7') and (WMCore_ResizeJob==True)	[u"GLIDEIN_CMSSite=='T3_US_NERSC' and GLIDEIN_Supported_VOs.str.contains('CMS') and GLIDEIN_REQUIRED_OS=='rhel7'"]	0
2	cms_xsede_passthrough	x509UserProxyVOName=='cms' and (DESIRED_Sites.str.contains('T3_US_TACC') or DESIRED_Sites.str.contains('T3_US_PSC'))	[u"(GLIDEIN_CMSSite=='T3_US_TACC' or GLIDEIN_CMSSite=='T3_US_PSC') and GLIDEIN_Supported_VOs.str.contains('CMS')"]	17010
3	cms_sdsc_passthrough	x509UserProxyVOName=='cms' and (DESIRED_Sites.str.contains('T3_US_SDSC'))	[u"(GLIDEIN_CMSSite=='T3_US_SDSC') and GLIDEIN_Supported_VOs.str.contains('CMS')"]	17010

=

2019-08-04 20:51:07,958 - decision_engine - glide_frontend_element - GlideinRequestManifests - INFO - 512(523 512 495 0) 1077(0 60000) | 13 0 11 0 | 128 0 128 | 41 193 | Up 0.0180
CMSHTPC_T3_US_NERSC_Cori@gfactory_instance_fermifactory02@gfactory_service_fermifactory02@fermifactory02.fnal.gov

2019-08-04 20:51:07,965 - decision_engine - glide_frontend_element - GlideinRequestManifests - INFO - 12(523 0 11 0) 1077(0 60000) | 928 142 726 0 | 13736 4806 8864 | 0 1 | Up 0.8120
CMSHTPC_T3_US_NERSC_Cori_KNL@gfactory_instance_fermifactory02@gfactory_service_fermifactory02@fermifactory02.fnal.gov

2019-08-04 20:51:12,011 - decision_engine - glide_frontend_element - GlideinRequestManifests - INFO - 0(0 0 0 0) 0(0 15000) | 1458 6 1358 0 | 2800 7 2793 | 0 0 | Up INFINITY
CMSHTPC_T3_US_Bridges@gfactory_instance_fermifactory02@gfactory_service_fermifactory02@fermifactory02.fnal.gov

2019-08-04 20:51:15,031 - decision_engine - glide_frontend_element - GlideinRequestManifests - INFO - 17072(17072 17072 16990 0) 8545(0 15000) | 26 0 24 0 | 48 0 48 | 1 2 | Up 0.0188

CMSHTPC_T3_US_SDSC_osg-comet-frontend@gfactory_instance_fermifactory02@gfactory_service_fermifactory02@fermifactory02.fnal.gov

So here we are again--17072 jobs pending and only 2 glideins sent out.

Factory limit is 144.

This is a very weird corner case.. it may be because the groups overlap to a large extent.

i.e. the jobs that match NERSC, SDSC, and Bridges are all in 3 different groups but the groups overlap significantly

Steve

From: Dennis D Box <dbox@fnal.gov>

Sent: Friday, August 2, 2019 11:25:08 AM

To: Steven C Timm <timmm@fnal.gov>; glideinwms-support <glideinwms-support@fnal.gov>

Subject: Re: fermifactory02 monitoring is wrong somehow

I notice that gfactory processes are running on fermifactory02:

```
[root@fermifactory02 google_json]# ps auxww | grep "gfactory"
```

```
gfactory 1771738 1.0 0.0 375988 38792 ? S< Jul25 122:41 python /usr/sbin/glideFactory.py /var/lib/gwms-factory/work-dir
```

```
gfactory 1771743 0.2 0.0 376596 38116 ? S< Jul25 25:40 /bin/python /usr/sbin/glideFactoryEntryGroup.py 1771738 60 5
```

```
/var/lib/gwms-factory/work-dir
```

```
CMSHTPC_T3_US_Bridges:CMSHTPC_T3_US_NERSC_Cori:CMSHTPC_T3_US_NERSC_Cori_KNL:CMSHTPC_T3_US_NERSC_Cori_KNL_SL7:CMSHTPC_T3_US_NERSC_Cori_shared:CMSHTPC_T3_US_SDSC_osg-comet-frontend:CMSHTPC_T3_US_TACC:CMS_T1_US_FNAL_condce:CMS_T1_US_FNAL_condce2:CMS_T1_US_FNAL_condce3:CMS_T1_US_FNAL_condce4:DUNE_T3_US_NERSC_Cori:DUNE_T3_US_NERSC_Cori_KNL:DUNE_T3_US_NERSC_Cori_KNL_SL7:DUNE_T3_US_NERSC_Cori_shared:FIFE_T3_US_NERSC_Cori:FIFE_T3_US_NERSC_Cori_KNL:FIFE_T3_US_NERSC_Cori_shared:FNAL_HEPCLOUD_AWS_us-east-1a_m3_2xlarge:FNAL_HEPCLOUD_AWS_us-east-1a_m3_xlarge:FNAL_HEPCLOUD_AWS_us-west-2a_m3_xlarge:FNAL_HEPCLOUD_GOOGLE_us-central1-a_custom-16-32768:FNAL_HEPCLOUD_GOOGLE_us-central1-a_custom-32-65536:FNAL_HEPCLOUD_GOOGLE_us-central1-a_n1-standard-1 0
```

```
gfactory 2704836 0.0 0.0 0 0 ? Z< 11:17 0:00 [python] <defunct>
```

```
gfactory 2704838 0.0 0.0 0 0 ? Z< 11:17 0:00 [python] <defunct>
```

```
gfactory 2705064 112 0.0 375568 27584 ? R< 11:18 0:01 /bin/python /usr/sbin/glideFactoryEntryGroup.py 1771738 60 5
```

```
/var/lib/gwms-factory/work-dir
```

```
CMSHTPC_T3_US_Bridges:CMSHTPC_T3_US_NERSC_Cori:CMSHTPC_T3_US_NERSC_Cori_KNL:CMSHTPC_T3_US_NERSC_Cori_KNL_SL7:CMSHTPC_T3_US_NERSC_Cori_shared:CMSHTPC_T3_US_SDSC_osg-comet-frontend:CMSHTPC_T3_US_TACC:CMS_T1_US_FNAL_condce:CMS_T1_US_FNAL_condce2:CMS_T1_US_FNAL_condce3:CMS_T1_US_FNAL_condce4:DUNE_T3_US_NERSC_Cori:DUNE_T3_US_NERSC_Cori_KNL:DUNE_T3_US_NERSC_Cori_KNL_SL7:DUNE_T3_US_NERSC_Cori_shared:FIFE_T3_US_NERSC_Cori:FIFE_T3_US_NERSC_Cori_KNL:FIFE_T3_US_NERSC_Cori_shared:FNAL_HEPCLOUD_AWS_us-east-1a_m3_2xlarge:FNAL_HEPCLOUD_AWS_us-east-1a_m3_xlarge:FNAL_HEPCLOUD_AWS_us-west-2a_m3_xlarge:FNAL_HEPCLOUD_GOOGLE_us-central1-a_custom-16-32768:FNAL_HEPCLOUD_GOOGLE_us-central1-a_custom-32-65536:FNAL_HEPCLOUD_GOOGLE_us-central1-a_n1-standard-1 0
```

```
[root@fermifactory02 google_json]#
```

But systemd thinks the factory is not running:

```
❏ gwms-factory.service - GWMS Factory Service
```

```
Loaded: loaded (/usr/lib/systemd/system/gwms-factory.service; disabled; vendor preset: disabled)
```

```
Active: inactive (dead)
```

```
Docs: http://glideinwms.fnal.gov/doc.prd/factory/index.html
```

```
GlideinWMS - Factory
```

```
glideinwms.fnal.gov
```

Overview. The main task of the Glidein Factory (or WMS Factory) is to advertise itself, listen for requests from Frontend clients and submit glideins. Look at the picture below for a schematic view.

```
[root@fermifactory02 google_json]#
```

Perhaps killing and restarting the factory would help?

Dennis

On 8/1/19 4:33 PM, Steven C Timm wrote:

all glideinwms developers currently have login to fermifactory02.
Some of the monitoring appears to be confused.. in particular it seems that the

ClientMonitor fields are quite confused from the glideclient/glidefactoryclient/glidefactory classads

```
GlideinMonitorTotalClientMonitorCoresIdle = 408
GlideinMonitorTotalClientMonitorCoresRunning = 325
GlideinMonitorTotalClientMonitorCoresTotal = 748
GlideinMonitorTotalClientMonitorGlideIdle = 6
GlideinMonitorTotalClientMonitorGlideRunning = 25
GlideinMonitorTotalClientMonitorGlideTotal = 36
GlideinMonitorTotalClientMonitorInfoAge = 14
GlideinMonitorTotalClientMonitorJobsIdle = 77
GlideinMonitorTotalClientMonitorJobsRunHere = 25
GlideinMonitorTotalClientMonitorJobsRunning = 515
```

The above numbers are for entry CMSHTPC_T3_US_NERSC_Cori_KNL and are an underestimate of what is actually running.. there

are 117 glideins active on that entry and they are all from the same glideclient group cms_nersc_passthrough on decision engine cmsde01

Look if you will at fermifactory02

condor_status -any -constraint 'MyType=="glidefactory"&&EntryName=="CMSHTPC_T3_US_NERSC_Cori_KNL"' and glidefactoryclient classad:

```
CMSHTPC_T3_US_NERSC_Cori_KNL@gfactory_instance_fermifactory02@gfactory_service_fermifactory02@HEPCloud-cmsde01-fnal.gov.cms_nersc_passthrough
```

and glideclient classad:

```
693452_CMSHTPC_T3_US_NERSC_Cori_KNL@gfactory_instance_fermifactory02@gfactory_service_fermifactory02@HEPCloud-cmsde01-fnal.gov.cms_nersc_passthrough
```

All are underestimating the total number of glideins running.

The standing requests are the following from cmsde01 (in 3 different groups, cms_nersc_passthrough for the 1st 2, cms_xsede_passthrough, and cms_sdsc passthrough)

```
2019-08-01 16:28:01,877 - decision_engine - glide_frontend_element - GlideinRequestManifests - INFO - 0( 0 0 0 0) 0( 0 60000)| 0 0 0 0| 0 0 0| 0 0|Up INFINITY
CMSHTPC_T3_US_NERSC_Cori@gfactory_instance_fermifactory02@gfactory_service_fermifactory02@fermifactory02.fnal.gov
2019-08-01 16:28:01,883 - decision_engine - glide_frontend_element - GlideinRequestManifests - INFO - 296( 296 290 296 0) 563( 0 60000)| 36 6 25 0| 748 408 325| 18 85|Up 0.0480
CMSHTPC_T3_US_NERSC_Cori_KNL@gfactory_instance_fermifactory02@gfactory_service_fermifactory02@fermifactory02.fnal.gov
2019-08-01 16:28:04,030 - decision_engine - glide_frontend_element - GlideinRequestManifests - INFO - 5898( 5898 5870 5898 0) 3462( 0 15000)| 1060 28 995 0| 1820 333 1486| 23 68|Up 0.8730
CMSHTPC_T3_US_Bridges@gfactory_instance_fermifactory02@gfactory_service_fermifactory02@fermifactory02.fnal.gov
2019-08-01 16:28:05,863 - decision_engine - glide_frontend_element - GlideinRequestManifests - INFO - 5898( 5898 5898 0) 3462( 0 15000)| 25 0 24 0| 24 0 24| 1 2|Up 0.0125
CMSHTPC_T3_US_SDSC_osg-comet-frontend@gfactory_instance_fermifactory02@gfactory_service_fermifactory02@fermifactory02.fnal.gov
```

In the last hour since I've been describing this problem the request for Cori_KNL went up significantly (from 3 idle 16 running to 19 idle 26 running in last hour) but I would expect that the request for SDSC should be much higher too and it is only requesting one idle glidein and 2 running.

Steve Timm

History

#1 - 08/08/2019 05:39 PM - Marco Mambelli

- Target version changed from v3_4_6 to v3_4_7

- Assignee set to Marco Mambelli

#2 - 09/16/2019 05:26 PM - Marco Mambelli

- Assignee changed from Marco Mambelli to Lorena Lobato Pardavila

#3 - 09/16/2019 05:27 PM - Marco Mambelli

- Target version changed from v3_4_7 to v3_6_1

#4 - 09/16/2019 05:28 PM - Marco Mambelli

- Priority changed from Normal to High

#5 - 10/07/2019 01:51 PM - Lorena Lobato Pardavila

- Assignee changed from Lorena Lobato Pardavila to Marco Mambelli

#6 - 10/29/2019 11:13 AM - Marco Mambelli

- Target version changed from v3_6_1 to v3_6_2

#7 - 03/03/2020 11:00 AM - Marco Mambelli

- Target version changed from v3_6_2 to v3_6_3

#8 - 05/07/2020 12:01 PM - Marco Mambelli

- Target version changed from v3_6_3 to v3_6_4

#9 - 09/17/2020 02:14 PM - Marco Mambelli

- Target version changed from v3_6_4 to v3_6_5