

GlideinWMS - Support #22999

Make Factory compatible w/ older 3.4 Frontends - Revert back to send REQUIRE_VOMS/REQUIRE_GLEXEC_USE attributes as strings

07/27/2019 12:51 AM - Marco Mambelli

Status:	Closed	Start date:	07/27/2019
Priority:	High	Due date:	
Assignee:	Lorena Lobato Pardavila	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	v3_4_6	Spent time:	0.00 hour
Stakeholders:	OSG		

Description

As per request from OSG we should maintain Factory compatibility w/ older 3.4 Frontends, including the buggy 3.4.2 and older. As from a preliminary analysis in [\[#22928\]](#) this probably implies going back to send the REQUIRE_VOMS/REQUIRE_GLEXEC_USE attributes as a string instead of boolean. The example that was breaking was when you set these attributes to "NEVER", because in this case, it is added checking to the match expression:

```
<attr name="GLIDEIN_Glexec_Use" glidein_publish="True" job_publish="True" parameter="False" type="string" value="NEVER"/>
```

We didn't catch this in our tests as exposed in the ticket [#22928](#), those attributes were not present in the compatibility testing setup. (For the record, they should have, considering that they were in important tickets for the release)

This should give back 3.4.2 compatibility and not break anything else (3.4.5 and 3.5).

This change will be only for the 3.4 series (at least any workaround to be compatible w/ buggy behavior).

Ticket tasks:

- test 3.4.5 Factory against different Frontend versions (at least 3.4.5 and 3.4.2) to verify the incompatibility
- verify the actual cause
- change the Factory code to be compatible w/ any 3.4 Frontend (repeat the tests above, this time all should work)
- identify the things to change back for [\[#23046\]](#)

Below are the requests from OSG

Hi Marco,

I still prefer what I proposed (option B) so that frontend users have the option to revert to an older 3.4 version no matter what 3.4 version the factory is on. Imagine that the release of 3.4.6 has a critical bug for some VO and for whatever reason the only version not affected by this bug is 3.4.2. If we went ahead with plan A, then that VO will be dead in the water due to the bug and factory/FE version incompatibility until you can look at the issue and provide a workaround.

I think compatibility between minor versions is a must wherever possible and supporting compatibility of old/new methods between major versions is important for simplifying transitions for users.

Thanks,

Brian

On 7/26/19 12:53 PM, Marco Mambelli wrote:

Hi Brian,

I don't know if the email was clear.

The incompatibility was already in 3.4.5 that is currently in OSG and CMS production.

People using these attributes had already to update their Frontends when after the upgrade of the Factory their Frontends were not matching jobs:

- FIFE frontend at Fermilab
- HCC

Was involuntary but the incompatibility has been already introduced and is there in production.
3.5 is behaving like 3.4.5 respect to that, nothing new introduced

For 3.4.6 and following (3.5.1, ...) should I keep the current behavior A (send booleans as booleans) or change back to B (send booleans as strings) ?

Factory 3.4.4, 3.4.5 and 3.5 do A

All older Factories do B

3.4.5 and 3.5 Frontends (and future ones) will work with both A and B

Frontends < 3.4.3 work only with B

Going back to B is still your suggestion?

Marco

On Jul 26, 2019, at 11:10 AM, Brian Lin <blin@cs.wisc.edu> wrote:

Hi Marco,

Instead of introducing version incompatibilities, I think it would make sense to go back to sending strings with boolean values inside. Then in 3.5.1, you change the frontends to prefer booleans but handle the string versions so that they are compatible with older factories. Then after we are confident that all FEs are on 3.5 (+ maybe a few months of slush time), you can complete the string -> boolean transition on the factory side. Then when all factories are updated, you can remove the string support.

Does that sounds reasonable?

- Brian

On 7/26/19 10:18 AM, Marco Mambelli wrote:

Hi Brian,

here the summary about the Frontend problems after the Factory upgrade.

More details are in the ticket:

<https://cdcvcs.fnal.gov/redmine/issues/22928>

The problem is with 3.4.2 Frontends and 3.4.3 or bigger Factories (3.4.5, because 3.4.3 and 3.4.4 did not go into production)

It happens when the REQUIRE_VOMS/REQUIRE_GLEXEC_USE attributes are used.

This is related to a series of tickets about booleans vs strings in Frontend and Factory.

The final result is correct but creates an incompatibility with FEs < 3.4.4.

For 3.4.6 we could go back sending strings with boolean values inside from The Factory (for the REQUIRE_VOMS/REQUIRE_GLEXEC_USE attributes) and that would be backward compatible and not cause problems.

But Factory 3.4.5 is out in production and is the one causing the incompatibility (triggering the Frontend bug).

Do you have any advise?

Thank you,

Marco

PS GlideinWMS 3.5 (no to be confused w/ 3.4.5), currently in OSG upcoming testing, includes all code in 3.4.5, so it will trigger the bug as well, same as 3.4.5, no more, no less

Related issues:

Related to GlideinWMS - Support #22928: Fermilab Frontend not communicating w...

Closed

07/12/2019

Precedes GlideinWMS - Support #23046: Revert the changes in #22999, at least ...

Accepted

07/28/2019

History

#1 - 07/31/2019 10:26 AM - Lorena Lobato Pardavila

- Priority changed from Normal to High

- Assignee set to Lorena Lobato Pardavila

#2 - 08/02/2019 10:11 AM - Marco Mambelli

- Related to Support #22928: Fermilab Frontend not communicating w/ upgraded Factory added

#3 - 08/02/2019 10:32 AM - Marco Mambelli

- Stakeholders updated

- Description updated

- Subject changed from Revert back to send REQUIRE_VOMS/REQUIRE_GLEXEC_USE attributes as strings to Make Factory compatible w/ older 3.4 Frontends - Revert back to send REQUIRE_VOMS/REQUIRE_GLEXEC_USE attributes as strings

#4 - 08/02/2019 10:33 AM - Marco Mambelli

- Precedes Support #23046: Revert the changes in #22999, at least any workaround to be compatible w/ buggy behavior in 3.4 added

#5 - 08/02/2019 05:31 PM - Lorena Lobato Pardavila

- Status changed from New to Work in progress

#6 - 08/05/2019 07:19 PM - Lorena Lobato Pardavila

- Description updated

#7 - 08/05/2019 07:35 PM - Lorena Lobato Pardavila

- Description updated

For testing, I have created several machines with different versions of glideinwms¹:

- Frontend: 3.2.22, 3.4.2 and 3.4.5.
- Factory: 3.4.2 and 3.4.5

For the record: I had to give up for the moment Frontend 3.2.22 connected with Factory 3.4.5 as I was getting credential authentication issues that need to be investigated and for this ticket is not urgent

- I verified the actual case:

1. I have configured all the environments first having set GLIDEIN_GLEXEC_USE to "OPTIONAL" and had jobs running fine.
2. Changed GLIDEIN_GLEXEC_USE to "NEVER", killed old glideins, waited and checked that condor_status was empty
3. Submitted new jobs. No job is running, because

```
match_expr = '(%s) and (glidein["attrs"].get("GLIDEIN_REQUIRE_GLEXEC_USE", "False") == "False")'
```

% match_expr is failing: boolean false != string "False". You can check doing:

```
condor_status -any -l | grep GLIDEIN_REQUIRE_GLEXEC_USE
GlideClientMatchingGlideinCondorExpr = "((True) and (getGlideinCpusNum(glidein) >= int(job.get(\"RequestCpus\", 1)))) and ((True) and (glidein[\"attrs\"].get(\"GLIDEIN_REQUIRE_GLEXEC_USE\", \"False\") == \"False\"))"
GLIDEIN_REQUIRE_GLEXEC_USE = false
GlideClientMatchingGlideinCondorExpr = "((True) and (getGlideinCpusNum(glidein) >= int(job.get(\"RequestCpus\", 1)))) and ((True) and (glidein[\"attrs\"].get(\"GLIDEIN_REQUIRE_GLEXEC_USE\", \"False\") == \"False\"))"
GLIDEIN_REQUIRE_GLEXEC_USE = false
```

4. Entry is down as expected

```
[2019-08-05 17:23:25,643] INFO: glideinFrontendElement:1763: All children terminated - took 0.0414960384369 seconds
[2019-08-05 17:23:25,644] INFO: glideinFrontendElement:522: Total matching idle 2 (old 10min 2 60min 0) running 0 limit 10000
[2019-08-05 17:23:25,645] INFO: glideinFrontendElement:1927: Jobs in schedd queues
| Slots | Cores | Glidein Req | Factory/Entry Information
[2019-08-05 17:23:25,645] INFO: glideinFrontendElement:1928: Idle (match eff old uniq) Run ( here max ) | Total Idle Run Fail | Total Idle Run | Idle MaxRun | State Factory
[2019-08-05 17:23:25,645] INFO: glideinFrontendElement:1916: 0( 0 0 0 0) 0( 0 1000) | 0 0 0 0 | 0 0 0 0 | 0 0 | Up
```

```

ITB_FC_CE2@gfactory_instance@gfactory_service@fermicloud115.fnal.gov
[2019-08-05 17:23:25,647] INFO: glideinFrontendElement:1916:      0(      0      0      0      0)      0(      0 10
000) |      0      0      0      0 |      0      0      0 |      0      0 | Up  ITB_FC_HTC_SIN_CE2@gfactory_instan
ce@gfactory_service@fermicloud115.fnal.gov
[2019-08-05 17:23:25,648] INFO: glideinFrontendElement:1927:      Jobs in schedd queues
|      Slots      |      Cores      |      Glidein Req |      Factory/Entry Information
[2019-08-05 17:23:25,648] INFO: glideinFrontendElement:1928: Idle (match eff old uniq) Run ( here max )
| Total Idle Run Fail | Total Idle Run | Idle MaxRun | State Factory
[2019-08-05 17:23:25,649] INFO: glideinFrontendElement:1916:      0(      0      0      0      0)      0(      0 20
000) |      0      0      0      0 |      0      0      0 |      0      0 | Up  Sum of useful factories
[2019-08-05 17:23:25,649] INFO: glideinFrontendElement:1916:      0(      0      0      0      0)      0(      0
0) |      0      0      0      0 |      0      0      0 |      0      0 | Down Sum of down factories
[2019-08-05 17:23:25,649] INFO: glideinFrontendElement:1916:      2(      2      2      2      2)      0(      0
0) |      0      0      0      0 |      0      0      0 |      0      0 | Down Unmatched
[2019-08-05 17:23:25,690] INFO: glideinFrontendElement:791: Advertising global and singular requests for f
actory fermicloud115.fnal.gov
[2019-08-05 17:23:25,694] DEBUG: glideinFrontendInterface:1068: In create Advertize work
[2019-08-05 17:23:25,695] DEBUG: glideinFrontendInterface:1110: Checking Credential file /etc/gwms-fronten
d/mm_proxy ...
[2019-08-05 17:23:25,696] DEBUG: glideinFrontendInterface:1164: Advertizing credential /etc/gwms-frontend/
mm_proxy with (0 idle, 0 max run) for request ITB_FC_CE2@gfactory_instance@gfactory_service

```

1. Next step was to switch to the Factory 3.4.5 sending the parameter as string. It took around 2 mins to update the classad on the Frontend but finally was sent as string.

```

$ condor_status -any -l | grep GLIDEIN_REQUIRE_GLEXEC_USE
GlideClientMatchingGlideinCondorExpr = "((True) and (getGlideinCpusNum(glidein) >= int(job.get(\"RequestCpus\"
, 1))) and ((True) and (glidein[\"attrs\"]).get(\"GLIDEIN_REQUIRE_GLEXEC_USE\", \"False\") == \"False\"))"
GLIDEIN_REQUIRE_GLEXEC_USE = "False"
GlideClientMatchingGlideinCondorExpr = "((True) and (getGlideinCpusNum(glidein) >= int(job.get(\"RequestCpus\"
, 1))) and ((True) and (glidein[\"attrs\"]).get(\"GLIDEIN_REQUIRE_GLEXEC_USE\", \"False\") == \"False\"))"
GLIDEIN_REQUIRE_GLEXEC_USE = "False"

```

And you can double-check in the Factory:

```

[root@fermicloud115 ~]# condor_status -any -l | grep GLIDEIN_REQUIRE_GLEXEC_USE | tail -n 1
GLIDEIN_REQUIRE_GLEXEC_USE = "False"

```

1. And I can confirm that with the change that I did in the Factory, now the entry is up and working:

```

[2019-08-05 17:27:25,814] INFO: glideinFrontendElement:1763: All children terminated - took 0.0416989326477 se
conds
[2019-08-05 17:27:25,815] INFO: glideinFrontendElement:522: Total matching idle 2 (old 10min 2 60min 0) runnin
g 0 limit 10000
[2019-08-05 17:27:25,815] INFO: glideinFrontendElement:1927:      Jobs in schedd queues
|      Slots      |      Cores      |      Glidein Req |      Fac
tory/Entry Information
[2019-08-05 17:27:25,816] INFO: glideinFrontendElement:1928: Idle (match eff old uniq) Run ( here max )
| Total Idle Run Fail | Total Idle Run | Idle MaxRun | Sta
te Factory
[2019-08-05 17:27:25,816] INFO: glideinFrontendElement:1916:      1(      2      1      1      0)      0(      0 10000)
|      0      0      0      0 |      0      0      0 |      1      2 | Up
ITB_FC_CE2@gfactory_instance@gfactory_service@fermicloud115.fnal.gov
[2019-08-05 17:27:25,818] INFO: glideinFrontendElement:1916:      1(      2      1      1      0)      0(      0 10000)
|      0      0      0      0 |      0      0      0 |      1      2 | Up  ITB_FC_HTC_SIN_CE2@gfactory_instan
ce@gfactory_service@fermicloud115.fnal.gov
[2019-08-05 17:27:25,819] INFO: glideinFrontendElement:1927:      Jobs in schedd queues
|      Slots      |      Cores      |      Glidein Req |      Factory/Entry Information
[2019-08-05 17:27:25,819] INFO: glideinFrontendElement:1928: Idle (match eff old uniq) Run ( here max )
| Total Idle Run Fail | Total Idle Run | Idle MaxRun | State Factory
[2019-08-05 17:27:25,819] INFO: glideinFrontendElement:1916:      2(      4      2      2      0)      0(      0 20000)
|      0      0      0      0 |      0      0      0 |      2      4 | Up  Sum of useful factories
[2019-08-05 17:27:25,819] INFO: glideinFrontendElement:1916:      0(      0      0      0      0)      0(      0
0) |      0      0      0      0 |      0      0      0 |      0      0 | Down Sum of down factories
[2019-08-05 17:27:25,820] INFO: glideinFrontendElement:1916:      0(      0      0      0      0)      0(      0
0) |      0      0      0      0 |      0      0      0 |      0      0 | Down Unmatched
[2019-08-05 17:27:25,841] INFO: glideinFrontendElement:791: Advertising global and singular requests for facto
ry fermicloud115.fnal.gov
[2019-08-05 17:27:25,845] DEBUG: glideinFrontendInterface:1068: In create Advertize work
[2019-08-05 17:27:25,845] DEBUG: glideinFrontendInterface:1110: Checking Credential file /etc/gwms-frontend/mm
_proxy ...

```

And glideins are working:

```

$ condor_status
Name                               OpSys      Arch      State      Activity      LoadAv Mem      ActvtyTi

```

```
me
glidein_757510_582306921@fermicloud378.fnal.gov LINUX      X86_64 Claimed   Busy           1.190 1853 0+00:00:
03
glidein_763380_390701100@fermicloud378.fnal.gov LINUX      X86_64 Unclaimed Benchmarking  1.040 1853 0+00:00:
03
```

	Machines	Owner	Claimed	Unclaimed	Matched	Preempting	Drain
X86_64/LINUX	2	0	1	1	0	0	0
Total	2	0	1	1	0	0	0

Note: This test was done with the versions that were giving the issue. Have tested the changes with the other versions and we can confirm that works with any 3.4 Frontends.

Gonna do the last test to double-check having the attribute set in global and in group. To make sure that nothing could break. I might try also with GLIDEIN_In_Downtime (due to [#21898](#))

#8 - 08/05/2019 07:38 PM - Lorena Lobato Pardavila

- Description updated

#9 - 08/06/2019 06:06 PM - Lorena Lobato Pardavila

- Assignee changed from Lorena Lobato Pardavila to Marco Mambelli

- Status changed from Work in progress to Feedback

Double checked to have the attribute set in global and in a group with different values.

1. GLIDEIN_GLEXEC_USE to "NEVER" globally and GLIDEIN_GLEXEC_USE to "OPTIONAL" in the group. Works as it should with any change because group attributes have preference over global, so no bug is triggered.
1. GLIDEIN_GLEXEC_USE to "OPTIONAL" globally and GLIDEIN_GLEXEC_USE to "NEVER" in the group. In this case, we triggered the bug and did the same tests exposed above. It was corrected with the same changes (pass the value as string) and I can confirm that it worked properly afterward.

Changes are done in v34/22999. Ready for review.

#10 - 08/06/2019 10:49 PM - Marco Mambelli

Changes are OK but you need to branch off branch_v3_4, not master, otherwise, it cannot go in 3.4.6 because it has all the 3.5 changes. You can start from branch_v3_4 and cherry pick, otherwise, it may be easier to redo the changes manually since are limited. Once you've done that you can merge. (to branch_v3_4)

#11 - 08/07/2019 11:47 AM - Lorena Lobato Pardavila

- Assignee changed from Marco Mambelli to Lorena Lobato Pardavila

- Status changed from Feedback to Resolved

Done. Changes are done now in v34/22999_1.

Resolving ticket.

#12 - 08/16/2019 10:02 AM - Marco Mambelli

- Status changed from Resolved to Closed