

## GlideinWMS - Feature #6556

### Frontend massively overprovisions sites due to no grouping of entries

06/26/2014 11:30 AM - Igor Sfiligoi

<b>Status:</b>	Closed	<b>Start date:</b>	06/26/2014
<b>Priority:</b>	Low	<b>Due date:</b>	
<b>Assignee:</b>	Parag Mhashilkar	<b>% Done:</b>	0%
<b>Category:</b>		<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	v3_6_x	<b>Spent time:</b>	0.00 hour
<b>Stakeholders:</b>	CMS OSG		

#### Description

The OSG/CMS GFactory ops has received several complaints from sites who are annoyed by the number of idle glideins waiting in their queues. This is especially noticeable for those who have many CEs in front of their batch system.

I think the core issue is due to the fact that the FE is not aware that all those CEs point to the same physical resource. It thus requests N x as many idle glideins compared to a site with just on CE. And this gets multiplied by the number of factories used as well, making it even worse.

We need a mechanism to tell the FE that certain entries (even accross different factories) are equivalent, and that it should scale down the requests accordingly.

#### History

##### #1 - 06/26/2014 11:30 AM - Igor Sfiligoi

Marked as high priority since we are getting complaints from sites right now.

##### #2 - 06/26/2014 11:34 AM - Igor Sfiligoi

My proposal would be to implement this fully in the FE.

The FE config would have a list of entry attributes.

For matchmatching/counting purposes, all the entries with the same attribute values would be treated as equivalent (independently of which GF they come from).

The final numbers are then divided by N before requests are sent out for the single entires.

PS: By default, there are no attributes defined, so no grouping.

Once we agree on the list of attributes in OSG-land, we add the list in the default config file (but do not hardcode).

##### #3 - 06/26/2014 11:52 AM - Anthony Tiradani

Hi Igor,

This sounds very much like a policy issue rather than a code issue. It seems extremely hacky to code around some implicit assumptions that are made by some "agreed upon" attributes that factory operators must know about and put in the factory config. Shouldn't you be able to create a match expression that handles this?

Tony

##### #4 - 06/26/2014 11:56 AM - Igor Sfiligoi

Nope.

There is no way to implement this, currently.

Each entry is treated as a completely independent entity right now.

There is absolutely no way to make correlations from the config point of view.

##### #5 - 06/27/2014 09:49 AM - Burt Holzman

- Target version changed from v3\_2\_6 to v3\_2\_x

I don't agree with this approach. It adds a lot of complexity to configs and makes it even harder to understand what the system is doing.

The factory configuration can be tuned to handle these situations (in general) by turning down the thresholds for entries that share resources.

Pushing to 3.2.x for now.

**#6 - 06/27/2014 10:40 AM - Igor Sfiligoi**

- Target version changed from v3\_2\_x to v3\_2\_6

From the operational point of view, micromanaging this kind of things in the factory is not sustainable. And I am talking with my OSG+CMS GFactory ops lead hat on!

Not to mention the fact that the number of entries per site can change from day to day, both due to some of the CEs going in downtime (e.g. for maintenance) and due to one or more factories going in downtime.

If you want to propose a different solution, I am all ears. But we do need a solution.

**#7 - 06/27/2014 11:33 AM - Burt Holzman**

- Target version changed from v3\_2\_6 to v3\_2\_x

This requires micromanagement on the factory side no matter what you do -- it's the factory configuration that would decide how to set the logical groupings. As a solution I'd be more supportive of the factory configuration being capable of setting limits per logical group, but exposing this to the frontend seems wrong.

**#8 - 06/27/2014 11:35 AM - Igor Sfiligoi**

The frontend needs to know about it.

If nothing else, there is no coordination between the factories. So the frontend needs to do the right thing on its side.

**#9 - 02/23/2015 01:18 PM - Brian Bockelman**

- Priority changed from High to Low

Complaints here have leveled off. Decreasing the priority; we can revisit in the next group meeting if necessary.

**#10 - 04/26/2018 04:44 PM - Marco Mambelli**

- Target version changed from v3\_2\_x to v3\_4\_x

**#11 - 09/06/2018 03:25 PM - Marco Mambelli**

- Target version changed from v3\_4\_x to v3\_5\_x

**#12 - 10/04/2019 05:12 PM - Marco Mambelli**

- Target version changed from v3\_5\_x to v3\_6\_x

**#13 - 08/31/2020 05:10 AM - Marco Mascheroni**

- Status changed from New to Closed