

GlideinWMS - Feature #4070

Bug # 4489 (New): Reduce the memory consumption of the factory

Reduce glideinEntryGroup memory footprint

06/12/2013 11:05 AM - Burt Holzman

Status:	Assigned	Start date:	06/12/2013
Priority:	Normal	Due date:	
Assignee:	Parag Mhashilkar	% Done:	0%
Category:	Factory	Estimated time:	0.00 hour
Target version:	v3_x	Spent time:	0.00 hour
Stakeholders:			

Description

For a test 2.7.1 factory with a lot of entries (but a single frontend), the factoryEntryGroup takes 240 MB of memory. On a production factory, it's more like 1.4 GB.

This is fine, except we fork and collect a lot of work, and the proportional set size per child might be growing over time as memory gets more and more fragmented.

Reducing the memory footprint significantly would help a lot.

History

#1 - 06/12/2013 01:20 PM - Igor Sfiligoi

- Status changed from Assigned to Feedback

- Assignee changed from Igor Sfiligoi to Parag Mhashilkar

I think the easiest thing to do is to push most of the "find work" activity into a different process, and then pickle back the results.

I have written the code in
branch_v3plus_igor_4070

Parag: Can you have a look? And if you have a test factory handy, try it out?
(I don't have one at this time)

#2 - 06/13/2013 01:38 PM - Parag Mhashilkar

Sorry for the delay, will try it out after v3.1. Don't want to disturb my existing v3.1 setup.

#3 - 06/13/2013 05:41 PM - Igor Sfiligoi

- Status changed from Feedback to Assigned

- Assignee changed from Parag Mhashilkar to Burt Holzman

I think my approach is not good enough;
testing at SDSC production factory show a saving in the order of 20%, not much more.
(for either/both VZISE and RSS as reported by top, from ~1.9G/1.6G -> 1.6G/1.4G)

Bottom line, not worth the additional complexity.

#4 - 06/13/2013 07:02 PM - Igor Sfiligoi

- Assignee changed from Burt Holzman to Parag Mhashilkar

I am getting convinced that the memory footprint is due mostly to the post-entry monitoring.

The
my_entries
structure persists between iterations, and my bet is that it is to be blamed for much of the memory use.

Assigning the ticket to Parag, since he is likely the most familiar with that part.

#5 - 08/05/2013 10:23 AM - Burt Holzman

- Parent task set to #4489