

GlideinWMS - Feature #2191

Need to speed up frontend

11/23/2011 09:25 AM - John Weigand

Status:	Closed	Start date:	01/31/2012
Priority:	Normal	Due date:	
Assignee:	Igor Sfiligoi	% Done:	0%
Category:	Frontend	Estimated time:	0.00 hour
Target version:	v2_5_5	Spent time:	6.00 hours
Stakeholders:			
Description			
From Igor			
Hi all.			
Even with the forking, and thus using 5 CPUs, the frontend is taking 10+ mins to make an iteration (see below). That's a long time! And we are not only at a third of what CMS plans to scale.			
We should spend some time in optimizing the code for speed.			
Cheers, Igor			
<pre>[1957] frontend@glidein-collector ~/frontstage/frontend_UCSD-v5_2\$ grep 'Sum of useful' log/group_main/frontend.20111122.info.log tail [2011-11-22T18:35:58-07:00 26804] 9489(77600 9456 9398 0) 1621k(17420 7320k) 17790 343 17441 969 2445k Up Sum of useful factories [2011-11-22T18:44:10-07:00 26804] 9574(91661 9504 9453 0) 1604k(16948 7320k) 17739 772 16961 981 2421k Up Sum of useful factories [2011-11-22T18:53:03-07:00 26804] 9420(77796 9376 9409 0) 1602k(16867 7320k) 17654 768 16881 956 2418k Up Sum of useful factories [2011-11-22T19:01:09-07:00 26804] 9371(76814 9331 9369 0) 1602k(16816 7320k) 17304 478 16822 957 2418k Up Sum of useful factories [2011-11-22T19:09:34-07:00 26804] 9315(76112 9279 9222 0) 1604k(16786 7320k) 17113 303 16802 955 2421k Up Sum of useful factories [2011-11-22T19:18:04-07:00 26804] 9277(76375 9236 9247 0) 1604k(16766 7320k) 17015 239 16770 960 2420k Up Sum of useful factories [2011-11-22T19:27:04-07:00 26804] 9253(76080 9200 9229 0) 1605k(16777 7320k) 16977 175 16792 956 2422k Up Sum of useful factories [2011-11-22T19:35:34-07:00 26804] 9193(75608 9144 9183 0) 1595k(16683 7320k) 16929 240 16685 940 2407k Up Sum of useful factories [2011-11-22T19:43:53-07:00 26804] 9165(76221 9123 9153 0) 1594k(16655 7320k) 16924 250 16668 941 2406k Up Sum of useful factories [2011-11-22T19:56:54-07:00 26804] 23461(1768k 23302 9042 0) 1593k(16634 7320k) 16900 224 16671 3133 2425k Up Sum of useful factories</pre>			
One more data point:			
<pre>[2031] frontend@glidein-collector ~/frontstage/frontend_UCSD-v5_2\$ grep 'Sum of useful' log/group_main/frontend.20111122.info.log tail [2011-11-22T19:09:34-07:00 26804] 9315(76112 9279 9222 0) 1604k(16786 7320k) 17113 303 16802 955 2421k Up Sum of useful factories [2011-11-22T19:18:04-07:00 26804] 9277(76375 9236 9247 0) 1604k(16766 7320k) 17015 239 16770 960 2420k Up Sum of useful factories [2011-11-22T19:27:04-07:00 26804] 9253(76080 9200 9229 0) 1605k(16777 7320k) 16977 175 16792 956 2422k Up Sum of useful factories [2011-11-22T19:35:34-07:00 26804] 9193(75608 9144 9183 0) 1595k(16683 7320k) 16929 240 16685 940 2407k Up Sum of useful factories [2011-11-22T19:43:53-07:00 26804] 9165(76221 9123 9153 0) 1594k(16655 7320k) 16924 250 16668 941 2406k Up Sum of useful factories [2011-11-22T19:56:54-07:00 26804] 23461(1768k 23302 9042 0) 1593k(16634 7320k) 16900 224 16671 3133 2425k Up Sum of useful factories</pre>			

[2011-11-22T20:16:19-07:00 26804] 23247(1760k 23154 22995 0) 1593k(16718 7320k) | 16867 128 16734 | 3642 2425k Up
Sum of useful factories
[2011-11-22T20:35:20-07:00 26804] 22840(1741k 22294 22598 0) 1594k(16657 7320k) | 17806 824 16792 | 3421 2426k Up
Sum of useful factories
[2011-11-22T20:53:20-07:00 26804] 22640(1711k 21844 22396 0) 1591k(16646 7320k) | 18117 1192 16856 | 0 2420k Up
Sum of useful factories
[2011-11-22T21:14:00-07:00 26804] 22859(1682k 22319 22035 0) 1578k(16484 7320k) | 17365 801 16547 | 3293 2402k Up
Sum of useful factories

Subtasks:

Bug # 2440: Improve frontend performance

Assigned

History

#1 - 11/23/2011 09:38 AM - Parag Mhashilkar

Igor, can you also include the full (info+debug+err) logs for couple of frontend's group cycle? It would be interested to see which consumes most of the cycles.

#2 - 12/05/2011 09:57 AM - Parag Mhashilkar

- Assignee set to Parag Mhashilkar

#3 - 12/06/2011 10:14 AM - Igor Sfiligoi

Here we go:

```
[0810] frontend@glidein-frontend ~/frontstage/frontend_UCSD-v5_3$ grep -v '|' log/group_main/frontend.20111206.info.log |grep -v "Voms proxy" |tail -100
[2011-12-06T07:00:39-07:00 13599] Counting subprocess created
[2011-12-06T07:08:19-07:00 13599] All children terminated
[2011-12-06T07:08:19-07:00 13599] Total matching idle 35040 (old 34889) running 18566
[2011-12-06T07:08:19-07:00 13599] Advertizing 264 requests
[2011-12-06T07:09:56-07:00 13599] Done advertizing
[2011-12-06T07:09:56-07:00 13599] Advertising 264 glideresource classads to the user pool
[2011-12-06T07:09:57-07:00 13599] Done advertising glideresource classads
[2011-12-06T07:09:58-07:00 13599] Writing stats
[2011-12-06T07:09:58-07:00 13599] Sleep
[2011-12-06T07:10:58-07:00 13599] Iteration at Tue Dec 6 07:10:58 2011
[2011-12-06T07:10:58-07:00 13599] Querying schedd, entry, and glidein status using child processes.
[2011-12-06T07:12:14-07:00 13599] All children terminated
[2011-12-06T07:12:14-07:00 13599] Jobs found total 53341 idle 34777 (old 34759, voms 34719) running 18564
[2011-12-06T07:12:15-07:00 13599] Glideins found total 19124 idle 693 running 18333 limit 25000 curb 20000
[2011-12-06T07:12:15-07:00 13599] Using 10 proxies
[2011-12-06T07:12:15-07:00 13599] Match
[2011-12-06T07:12:15-07:00 13599] Counting subprocess created
[2011-12-06T07:20:08-07:00 13599] All children terminated
[2011-12-06T07:20:08-07:00 13599] Total matching idle 34777 (old 34719) running 18564
[2011-12-06T07:20:09-07:00 13599] Advertizing 264 requests
[2011-12-06T07:20:23-07:00 13599] WARNING: Advertizing failed for 1 requests. See debug log for more details.
[2011-12-06T07:20:23-07:00 13599] Advertising 264 glideresource classads to the user pool
[2011-12-06T07:20:23-07:00 13599] Done advertising glideresource classads
[2011-12-06T07:20:24-07:00 13599] Writing stats
[2011-12-06T07:20:24-07:00 13599] Sleep
[2011-12-06T07:21:24-07:00 13599] Iteration at Tue Dec 6 07:21:24 2011
[2011-12-06T07:21:24-07:00 13599] Querying schedd, entry, and glidein status using child processes.
[2011-12-06T07:22:33-07:00 13599] All children terminated
[2011-12-06T07:22:33-07:00 13599] Jobs found total 53075 idle 34780 (old 34762, voms 34552) running 18295
[2011-12-06T07:22:34-07:00 13599] Glideins found total 18940 idle 827 running 18002 limit 25000 curb 20000
[2011-12-06T07:22:34-07:00 13599] Using 10 proxies
[2011-12-06T07:22:34-07:00 13599] Match
[2011-12-06T07:22:34-07:00 13599] Counting subprocess created
[2011-12-06T07:30:38-07:00 13599] All children terminated
[2011-12-06T07:30:38-07:00 13599] Total matching idle 34780 (old 34552) running 18295
[2011-12-06T07:30:38-07:00 13599] Advertizing 264 requests
[2011-12-06T07:31:02-07:00 13599] WARNING: Advertizing failed for 1 requests. See debug log for more details.
[2011-12-06T07:31:02-07:00 13599] Advertising 264 glideresource classads to the user pool
[2011-12-06T07:31:02-07:00 13599] Done advertising glideresource classads
[2011-12-06T07:31:03-07:00 13599] Writing stats
[2011-12-06T07:31:03-07:00 13599] Sleep
[2011-12-06T07:32:03-07:00 13599] Iteration at Tue Dec 6 07:32:03 2011
[2011-12-06T07:32:03-07:00 13599] Querying schedd, entry, and glidein status using child processes.
[2011-12-06T07:33:07-07:00 13599] All children terminated
[2011-12-06T07:33:08-07:00 13599] Jobs found total 52743 idle 34654 (old 34636, voms 34518) running 18089
```

```

[2011-12-06T07:33:08-07:00 13599] Glideins found total 19138 idle 1090 running 17972 limit 25000 curb 20000
[2011-12-06T07:33:09-07:00 13599] Using 10 proxies
[2011-12-06T07:33:09-07:00 13599] Match
[2011-12-06T07:33:09-07:00 13599] Counting subprocess created
[2011-12-06T07:40:58-07:00 13599] All children terminated
[2011-12-06T07:40:58-07:00 13599] Total matching idle 34654 (old 34518) running 18089
[2011-12-06T07:40:59-07:00 13599] Advertizing 264 requests
[2011-12-06T07:41:20-07:00 13599] WARNING: Advertizing failed for 1 requests. See debug log for more details.
[2011-12-06T07:41:20-07:00 13599] Advertising 264 glideresource classads to the user pool
[2011-12-06T07:41:21-07:00 13599] Done advertizing glideresource classads
[2011-12-06T07:41:21-07:00 13599] Writing stats
[2011-12-06T07:41:22-07:00 13599] Sleep
[2011-12-06T07:42:22-07:00 13599] Iteration at Tue Dec 6 07:42:22 2011
[2011-12-06T07:42:22-07:00 13599] Querying schedd, entry, and glidein status using child processes.
[2011-12-06T07:43:11-07:00 27842] WARNING: Failed to talk to schedd. See debug log for more details.
[2011-12-06T07:43:25-07:00 13599] All children terminated
[2011-12-06T07:43:26-07:00 13599] Jobs found total 52429 idle 34396 (old 34396, voms 34348) running 18033
[2011-12-06T07:43:26-07:00 13599] Glideins found total 18876 idle 941 running 17848 limit 25000 curb 20000
[2011-12-06T07:43:26-07:00 13599] Using 10 proxies
[2011-12-06T07:43:26-07:00 13599] Match
[2011-12-06T07:43:26-07:00 13599] Counting subprocess created
[2011-12-06T07:51:10-07:00 13599] All children terminated
[2011-12-06T07:51:10-07:00 13599] Total matching idle 34396 (old 34348) running 18033
[2011-12-06T07:51:10-07:00 13599] Advertizing 264 requests
[2011-12-06T07:51:30-07:00 13599] WARNING: Advertizing failed for 1 requests. See debug log for more details.
[2011-12-06T07:51:30-07:00 13599] Advertising 264 glideresource classads to the user pool
[2011-12-06T07:51:31-07:00 13599] Done advertizing glideresource classads
[2011-12-06T07:51:32-07:00 13599] Writing stats
[2011-12-06T07:51:32-07:00 13599] Sleep
[2011-12-06T07:52:32-07:00 13599] Iteration at Tue Dec 6 07:52:32 2011
[2011-12-06T07:52:32-07:00 13599] Querying schedd, entry, and glidein status using child processes.
[2011-12-06T07:53:19-07:00 28313] WARNING: Failed to talk to schedd. See debug log for more details.
[2011-12-06T07:53:35-07:00 13599] All children terminated
[2011-12-06T07:53:36-07:00 13599] Jobs found total 52209 idle 34061 (old 34061, voms 34029) running 18148
[2011-12-06T07:53:36-07:00 13599] Glideins found total 18600 idle 636 running 17898 limit 25000 curb 20000
[2011-12-06T07:53:36-07:00 13599] Using 10 proxies
[2011-12-06T07:53:36-07:00 13599] Match
[2011-12-06T07:53:36-07:00 13599] Counting subprocess created
[2011-12-06T08:01:08-07:00 13599] All children terminated
[2011-12-06T08:01:08-07:00 13599] Total matching idle 34061 (old 34029) running 18148
[2011-12-06T08:01:08-07:00 13599] Advertizing 264 requests
[2011-12-06T08:01:58-07:00 13599] Done advertizing
[2011-12-06T08:01:58-07:00 13599] Advertising 264 glideresource classads to the user pool
[2011-12-06T08:01:59-07:00 13599] Done advertizing glideresource classads
[2011-12-06T08:02:00-07:00 13599] Writing stats
[2011-12-06T08:02:00-07:00 13599] Sleep
[2011-12-06T08:03:00-07:00 13599] Iteration at Tue Dec 6 08:03:00 2011
[2011-12-06T08:03:00-07:00 13599] Querying schedd, entry, and glidein status using child processes.
[2011-12-06T08:03:45-07:00 28770] WARNING: Failed to talk to schedd. See debug log for more details.
[2011-12-06T08:04:01-07:00 13599] All children terminated
[2011-12-06T08:04:01-07:00 13599] Jobs found total 52015 idle 34191 (old 34191, voms 33923) running 17824
[2011-12-06T08:04:02-07:00 13599] Glideins found total 18053 idle 550 running 17438 limit 25000 curb 20000
[2011-12-06T08:04:02-07:00 13599] Using 10 proxies
[2011-12-06T08:04:02-07:00 13599] Match
[2011-12-06T08:04:02-07:00 13599] Counting subprocess created

```

nothing much in err and debug logs (just timeouts talking to daemons)

Here is what was the status:

```
[0811] frontend@glidein-frontend ~/frontstage/frontend_UCSD-v5_3$ condor_status -schedd
```

Name	Machine	TotalRunningJobs	TotalIdleJobs	TotalHeldJobs
cmsfnal01.fnal.gov	cmsfnal01.	0	0	0
glidein-2.t2.ucsd.edu	glidein-2.	463	586	115
submit-2.t2.ucsd.edu	submit-2.t	17154	33889	1089
	TotalRunningJobs		TotalIdleJobs	TotalHeldJobs

Total	17617	34475	1204
-------	-------	-------	------

```
[0812] frontend@glidein-frontend ~/frontstage/frontend_UCSD-v5_3$ condor_status -total -state
```

Total	Owner	Unclaimed	Claimed	Preempting	Matched	Backfill
-------	-------	-----------	---------	------------	---------	----------

Benchmarking	13	0	13	0	0	0
--------------	----	---	----	---	---	---

Busy	9806	0	0	9806	0	0	0
Idle	537	0	489	48	0	0	0
Killing	7	0	0	0	7	0	0
Retiring	7854	0	0	7854	0	0	0
Vacating	7	0	0	0	7	0	0
Total	18224	0	502	17708	14	0	0

#4 - 01/12/2012 10:48 PM - Igor Sfiligoi

- Status changed from New to Assigned

- Target version set to v2_5_5

- % Done changed from 0 to 70

- Estimated time set to 6.00 h

Fixed the matchmaking time.

Addressed it in three steps:

- 1) Add clustering, just like the Condor AutoClusters
- 2) Use the job_format_list to filter out attributes that are not used for matchmaking
- 3) Use arrays to build the set of job, and change indexes from tuples to longint

The net result is that the CMS matchmaking time went from 12mins to under 2mins!
(the 12mins forced me to do it... it was getting ridiculously high!)

Before:

```
[2012-01-12T09:57:33-07:00 25241] Counting subprocess created
[2012-01-12T10:10:23-07:00 25241] All children terminated
[2012-01-12T10:14:27-07:00 25241] Counting subprocess created
[2012-01-12T10:27:03-07:00 25241] All children terminated
```

After:

```
[2012-01-12T20:28:42-07:00 27882] Counting subprocess created
[2012-01-12T20:30:31-07:00 27882] All children terminated
[2012-01-12T20:34:05-07:00 27882] Counting subprocess created
[2012-01-12T20:35:46-07:00 27882] All children terminated
[2012-01-12T20:39:45-07:00 27882] Counting subprocess created
[2012-01-12T20:41:31-07:00 27882] All children terminated
```

The full cycle is still quite long, but it is now due mostly to the time it takes to get the classads from Condor, and then advertizing back.

Commits:

branch_v2plus:

```
commit:971470670f847b50b3f2e40e1adf0cbd02cdf092
commit:6e9cf76a75b4f87ff9372b64b9e6cfeec492e8c7
commit:ad1719a0364e8797013167b039f481e7de78fee7
commit:67278818f25d90d0383e472e9df0fcdf73a149eb
```

master:

```
commit:fd5e452a1004c530b28547b7dc139adcbf51219b
commit:cf7c26e94fc57b4efe319ccb9c9c93301e02dc95
commit:881a77ec53a281ee4cebb3d455f2856dd7779b9c
commit:9e6b26bed5e89b86a471bd073d58491e6ecab23c
```

#5 - 01/31/2012 02:49 PM - Parag Mhashilkar

- Status changed from Assigned to Resolved

- Assignee changed from Parag Mhashilkar to Igor Sfiligoi

Closing this one for v2.5.5 and opening a sub ticket for future release.

#6 - 02/28/2012 01:37 PM - Parag Mhashilkar

- Status changed from Resolved to Closed