

glideinWMS - Support #20302

CMS Scale tests: frontend scalability issues

07/10/2018 08:17 AM - Marco Mascheroni

Status:	Resolved	Start date:	07/10/2018
Priority:	Normal	Due date:	
Assignee:	Marco Mascheroni	% Done:	0%
Category:	Frontend	Estimated time:	0.00 hour
Target version:	v3_5	Spent time:	0.00 hour
Stakeholders:			

Description

Recently CMS did a series of scale tests that highlighted some potential scalability issues in the frontend. I put some printouts in the functions that fork the workers, and here are the results. I had not time to investigate more, but I think this is a good starting point:

1. there are 2 different sets of `condr_q` [1.1] and [1.2] that probably could be merged.
 1. 1.1 and they taking very different time scales (seconds vs minutes) maybe projected vs unprojected (Antonio's hypothesis)
2. It is obvious that "subprocess_count_dt" is taking most of the time [1.3].
3. subprocess_count_dt could be better parallelized:
 1. the number of instances is less than the number of matchmakers (7 vs 10)
 2. The time consumed per instance is very different for couple of them, look at [1.3.1] (less than 1 second vs [1.3.2] (less than 1 hour) vs the rest (~ 1.5 hours)

[2018-06-30 10:10:12,512] INFO: Iteration at Sat Jun 30 10:10:12 2018

[1.1]

[2018-06-30 10:10:12,512] INFO: Querying schedd, entry, and glidein status using child processes.

[2018-06-30 10:10:13,861] INFO: DDEBUG: 1310202;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;1.3404071331

[2018-06-30 10:10:13,875] INFO: DDEBUG: 1310206;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;1.35142397881

[2018-06-30 10:10:13,885] INFO: DDEBUG: 1310207;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;1.36031794548

[2018-06-30 10:10:15,338] INFO: DDEBUG: 1310197;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;2.82039284706

[2018-06-30 10:10:15,339] INFO: DDEBUG: 1310201;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;2.81942415237

[2018-06-30 10:10:15,344] INFO: DDEBUG: 1310199;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;2.82512307167

[2018-06-30 10:10:15,344] INFO: DDEBUG: 1310208;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;2.81894803047

[2018-06-30 10:10:15,346] INFO: DDEBUG: 1310200;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;2.82667779922

[2018-06-30 10:10:15,349] INFO: DDEBUG: 1310196;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;2.83266615868

[2018-06-30 10:10:15,352] INFO: DDEBUG: 1310198;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;2.83415913582

[2018-06-30 10:10:15,363] INFO: DDEBUG: 1310205;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;2.84001111984

[2018-06-30 10:10:15,402] INFO: DDEBUG: 1310204;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;2.88005900383

[2018-06-30 10:10:15,414] INFO: DDEBUG: 1310203;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;2.89220905304

[2018-06-30 10:10:16,455] INFO: DDEBUG: 1310194;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;3.9400370121

[2018-06-30 10:10:16,475] INFO: DDEBUG: 1310195;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;3.9575510025

[2018-06-30 10:10:16,491] INFO: DDEBUG: 1310193;<bound method glideinFrontendElement.query_factory of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;3.97680807114

[1.2]

[2018-06-30 10:11:12,665] INFO: DDEBUG: 1310218;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;60.1305930614
[2018-06-30 10:11:15,698] INFO: DDEBUG: 1310209;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;63.1720750332
[2018-06-30 10:11:19,188] INFO: DDEBUG: 1310217;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;66.6543180943
[2018-06-30 10:11:23,373] INFO: DDEBUG: 1310216;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;70.8397378922
[2018-06-30 10:11:28,105] INFO: DDEBUG: 1310212;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;75.5762519836
[2018-06-30 10:11:35,226] INFO: DDEBUG: 1310215;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;82.6942749023
[2018-06-30 10:11:41,466] INFO: DDEBUG: 1310211;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;88.9382030964
[2018-06-30 10:11:48,499] INFO: DDEBUG: 1310214;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;95.967936039
[2018-06-30 10:11:55,852] INFO: DDEBUG: 1310210;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;103.324620008
[2018-06-30 10:12:03,638] INFO: DDEBUG: 1310219;<bound method glideinFrontendElement.get_condor_status of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;111.101990223
[2018-06-30 10:12:04,033] INFO: DDEBUG: 1310213;<bound method glideinFrontendElement.get_condor_q of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;111.502775908
[2018-06-30 10:12:12,276] INFO: All children terminated
[2018-06-30 10:12:23,590] INFO: Jobs found total 1000000 idle 1000000 (good 1000000, old(10min 1000000, 60min 763250), grid 1000000, voms 1000000) running 0
[2018-06-30 10:12:23,670] INFO: Group glideins found total 9376 limit 550000 curb 500000; of these idle 9090 limit 20000 curb 20000 running 0
[2018-06-30 10:12:23,670] INFO: Frontend glideins found total 9456 limit 550000 curb 500000; of these idle 9263 limit 550000 curb 500000
[2018-06-30 10:12:23,670] INFO: Overall slots found total 0 limit 550000 curb 500000; of these idle 0 limit 550000 curb 500000
[2018-06-30 10:12:23,670] INFO: Updating usermap
[2018-06-30 10:12:23,670] INFO: Match
[2018-06-30 10:12:24,413] INFO: DDEBUG: 1310318;<bound method glideinFrontendElement.subprocess_count_glidein of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;0.674268960953
[2018-06-30 10:12:24,571] INFO: DDEBUG: 1310319;<bound method glideinFrontendElement.subprocess_count_glidein of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;0.755457878113
[2018-06-30 10:12:24,628] INFO: Active forks = 10, Forks to finish = 18
[2018-06-30 10:12:24,649] INFO: DDEBUG: 1310320;<bound method glideinFrontendElement.subprocess_count_glidein of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;0.732969045639
[2018-06-30 10:12:24,782] INFO: DDEBUG: 1310321;<bound method glideinFrontendElement.subprocess_count_glidein of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;0.75870013237
[2018-06-30 10:12:24,784] INFO: Active forks = 10, Forks to finish = 16
[2018-06-30 10:12:24,821] INFO: DDEBUG: 1310322;<bound method glideinFrontendElement.subprocess_count_glidein of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;0.692282915115
[2018-06-30 10:12:24,897] INFO: DDEBUG: 1310324;<bound method glideinFrontendElement.subprocess_count_glidein of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;0.661998987198
[2018-06-30 10:12:24,998] INFO: Active forks = 10, Forks to finish = 14
[2018-06-30 10:12:25,045] INFO: DDEBUG: 1310325;<bound method glideinFrontendElement.subprocess_count_glidein of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;0.697236776352
[2018-06-30 10:12:25,072] INFO: DDEBUG: 1310326;<bound method glideinFrontendElement.subprocess_count_glidein of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;0.618963956833
[2018-06-30 10:12:25,151] INFO: DDEBUG: 1310327;<bound method glideinFrontendElement.subprocess_count_glidein of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;0.613980054855
[2018-06-30 10:12:25,197] INFO: DDEBUG: 1310328;<bound method glideinFrontendElement.subprocess_count_glidein of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;0.569120883942
[2018-06-30 10:12:25,530] INFO: DDEBUG: 1310329;<bound method glideinFrontendElement.subprocess_count_real of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;0.819628953934

[1.3]

[1.3.1] [2018-06-30 10:12:25,728] INFO: DDEBUG: 1310332;<bound method glideinFrontendElement.subprocess_count_dt of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;0.732350111008

[2018-06-30 10:12:25,870] INFO: Active forks = 10, Forks to finish = 12

[2018-06-30 10:12:26,039] INFO: Active forks = 6, Forks to finish = 6

[1.3.2] [2018-06-30 11:03:22,405] INFO: DDEBUG: 1310330;<bound method glideinFrontendElement.subprocess_count_dt of

<__main__.glideinFrontendElement instance at 0x1d49ea8>>;3057.61987305

[2018-06-30 11:03:23,020] INFO: Active forks = 5, Forks to finish = 5

[2018-06-30 11:42:46,833] INFO: DDEBUG: 1310336;<bound method glideinFrontendElement.subprocess_count_dt of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;5420.79418206

[2018-06-30 11:42:47,385] INFO: Active forks = 4, Forks to finish = 4

[2018-06-30 11:43:05,031] INFO: DDEBUG: 1310331;<bound method glideinFrontendElement.subprocess_count_dt of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;5440.1055119

[2018-06-30 11:43:05,405] INFO: Active forks = 3, Forks to finish = 3

[2018-06-30 11:44:10,885] INFO: DDEBUG: 1310333;<bound method glideinFrontendElement.subprocess_count_dt of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;5505.09298015

[2018-06-30 11:44:11,533] INFO: Active forks = 2, Forks to finish = 2

[2018-06-30 11:44:21,240] INFO: DDEBUG: 1310334;<bound method glideinFrontendElement.subprocess_count_dt of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;5515.36989713

[2018-06-30 11:44:21,674] INFO: Active forks = 1, Forks to finish = 1

[2018-06-30 11:44:30,103] INFO: DDEBUG: 1310335;<bound method glideinFrontendElement.subprocess_count_dt of <__main__.glideinFrontendElement instance at 0x1d49ea8>>;5524.13912106

[2018-06-30 11:44:30,434] INFO: Active forks = 0, Forks to finish = 0

[2018-06-30 11:44:30,435] INFO: All children terminated - took 5526.76403117 seconds

History

#1 - 09/21/2018 09:56 AM - Marco Mascheroni

- Assignee set to Marco Mascheroni

- Target version set to v3_5

#2 - 11/21/2018 10:05 AM - Marco Mascheroni

- Target version changed from v3_5 to v3_5_1

#3 - 03/22/2019 05:19 PM - Marco Mascheroni

- Target version changed from v3_5_1 to v3_5

- Status changed from New to Work in progress

#4 - 04/17/2019 01:08 PM - Marco Mascheroni

- Assignee changed from Marco Mascheroni to Marco Mambelli

- Status changed from Work in progress to Feedback

- Category set to Frontend

Marco, feel free to reassign this to Lorena/Dennis for review, I'll let you decide.

#5 - 05/03/2019 07:19 PM - Marco Mambelli

- Assignee changed from Marco Mambelli to Marco Mascheroni

#6 - 05/15/2019 04:52 AM - Marco Mascheroni

- Status changed from Feedback to Resolved