

openQA Project - action #39980

Cache locks assets when worker dies in critical section

2018-08-20 06:44 - EDiGiacinto

Status:	Resolved	Start date:	2018-10-02
Priority:	High	Due date:	
Assignee:	EDiGiacinto	% Done:	100%
Category:	Concrete Bugs	Estimated time:	0.00 hour
Target version:	Done		
Difficulty:			
Description			
<p>When an asset is being downloaded by the cache (since now is bit more race-free) it acquires a lock that is meant to avoid that other instances (in the same machine) start to download the same asset. If a worker dies or crashes in the critical section, the same asset is locked and the other instances will wait for the lock to be released.</p> <p>See also: https://progress.opensuse.org/issues/39833</p> <p>This happened once in osd already, openqaworker6 now is stuck on:</p> <pre>Aug 20 08:38:12 openqaworker6 worker[13371]: [info] CACHE: Being downloaded by another worker, sleeping. Aug 20 08:38:17 openqaworker6 worker[13371]: [info] CACHE: Being downloaded by another worker, sleeping. Aug 20 08:38:22 openqaworker6 worker[13371]: [info] CACHE: Being downloaded by another worker, sleeping. Aug 20 08:38:27 openqaworker6 worker[13371]: [info] CACHE: Being downloaded by another worker, sleeping. Aug 20 08:38:32 openqaworker6 worker[13371]: [info] CACHE: Being downloaded by another worker, sleeping. Aug 20 08:38:37 openqaworker6 worker[13371]: [info] CACHE: Being downloaded by another worker, sleeping. Aug 20 08:38:42 openqaworker6 worker[13371]: [info] CACHE: Being downloaded by another worker, sleeping. Aug 20 08:38:47 openqaworker6 worker[13371]: [info] CACHE: Being downloaded by another worker, sleeping. Aug 20 08:38:52 openqaworker6 worker[13371]: [info] CACHE: Being downloaded by another worker, sleeping.</pre> <p>And jobs will result stuck in running.</p> <p>My first experiments where using IPC memory to acquire semaphores, and use them with SEM_UNDO exactly to cover this case, so if a process exits abnormally the lock is released automatically https://progress.opensuse.org/issues/34597#note-4. But that approach was not used in favor of keeping SQLite - now we need to implement some monitoring checks to unlock the cache in such situations (e.g. checking pid that were downloading are still alive, but imho it's racy as could lead to instances unlocking in wrong timings and potentially able to corrupt downloads)</p>			
Subtasks:			
action # 41891: Worker stalled during flock			Resolved
Related issues:			
Related to openQA Project - action #39833: [tools] When a worker is abruptly ...		Resolved	2018-08-16
Related to openQA Project - action #34597: Race condition causing problems wi...		Resolved	2018-05-11
Related to openQA Project - action #40871: Asset removed during job runtime		Resolved	2018-09-11

History

#1 - 2018-08-20 06:44 - EDiGiacinto

- Related to action #39833: [tools] When a worker is abruptly killed, jobs get blocked - CACHE: Being downloaded by another worker, sleeping added

#2 - 2018-08-20 10:11 - EDiGiacinto

- Description updated

#3 - 2018-08-20 13:17 - EDiGiacinto

- Description updated

#4 - 2018-09-12 10:34 - EDiGiacinto

Is still being worked in <https://github.com/os-autoinst/openQA/pull/1783> but still not ready - the general idea we agreed upon: move the cache to a separate service, and the workers just filling downloads requests to it.

As in our codebase we use already Minion for other things (GRU), come up naturally to base off this work on it.

Currently the service is able to replace the cache component inside the workers, but further work is required to integrate it with the exiting codebase, make the worker code use it and also to adapt all cache tests that are already part of our testsuite.

#5 - 2018-09-17 07:16 - coolo

- Target version set to Current Sprint

this part of the current story

#6 - 2018-09-17 08:11 - EDiGiacinto

Update:

I'm seeing races using minion locks with the sqlite backend while setting up stress tests:

```
ok 1 - Asset downloaded
[INFO] Downloading sle-12-SP3-x86_64-0368-200_3@64bit.qcow2 from http://localhost:38075/tests/922756/asset/hdd/sle-12-SP3-x86_64-0368-200_3@64bit.qcow2
[Fri Sep 14 16:57:41 2018] [debug] GET "/tests/922756/asset/hdd/sle-12-SP3-x86_64-0368-200_3@64bit.qcow2" (3be6ac10)
[Fri Sep 14 16:57:41 2018] [debug] Routing to a callback
[Fri Sep 14 16:57:41 2018] [debug] 200 OK (0.000446s, 2242.152/s)
[DEBUG] CACHE: Health: Real size: 0, Configured limit: 53687091200
[INFO] CACHE: updating the t/cache.d/cache/sle-12-SP3-x86_64-0368-200_3@64bit.qcow2 with andi $a3, $t1, 41399 and 3
[DEBUG] CACHE: Asset download successful to t/cache.d/cache/sle-12-SP3-x86_64-0368-200_3@64bit.qcow2, Cache size is: 3
ok 2 - Asset downloaded
not ok 3 - Asset downloaded
# Failed test 'Asset downloaded'
# at t/25-cache-service.t line 120.
[INFO] Downloading sle-12-SP3-x86_64-0368-200_4@64bit.qcow2 from http://localhost:38075/tests/922756/asset/hdd/sle-12-SP3-x86_64-0368-200_4@64bit.qcow2
[Fri Sep 14 16:57:47 2018] [debug] GET "/tests/922756/asset/hdd/sle-12-SP3-x86_64-0368-200_4@64bit.qcow2" (fca d800b)
[Fri Sep 14 16:57:47 2018] [debug] Routing to a callback
[Fri Sep 14 16:57:47 2018] [debug] 200 OK (0.000439s, 2277.904/s)
[DEBUG] CACHE: Health: Real size: 0, Configured limit: 53687091200
[INFO] CACHE: updating the t/cache.d/cache/sle-12-SP3-x86_64-0368-200_4@64bit.qcow2 with andi $a3, $t1, 41399 and 4
[DEBUG] CACHE: Asset download successful to t/cache.d/cache/sle-12-SP3-x86_64-0368-200_4@64bit.qcow2, Cache size is: 4
not ok 4 - Asset downloaded
# Failed test 'Asset downloaded'
# at t/25-cache-service.t line 120.
not ok 5 - Asset downloaded
# Failed test 'Asset downloaded'
# at t/25-cache-service.t line 120.
[INFO] Downloading sle-12-SP3-x86_64-0368-200_8@64bit.qcow2 from http://localhost:38075/tests/922756/asset/hdd/sle-12-SP3-x86_64-0368-200_8@64bit.qcow2
[Fri Sep 14 16:57:52 2018] [debug] GET "/tests/922756/asset/hdd/sle-12-SP3-x86_64-0368-200_8@64bit.qcow2" (0762f4d6)
[Fri Sep 14 16:57:52 2018] [debug] Routing to a callback
[Fri Sep 14 16:57:52 2018] [debug] 200 OK (0.000549s, 1821.494/s)
[DEBUG] CACHE: Health: Real size: 0, Configured limit: 53687091200
[INFO] CACHE: updating the t/cache.d/cache/sle-12-SP3-x86_64-0368-200_8@64bit.qcow2 with andi $a3, $t1, 41399 and 8
[DEBUG] CACHE: Asset download successful to t/cache.d/cache/sle-12-SP3-x86_64-0368-200_8@64bit.qcow2, Cache size is: 8
not ok 6 - Asset downloaded
```

To trigger it, i filed requests to the cache service with very small assets - it seems there is a race when we detect the task as completed or not, but still to dissect if comes from my changes or the Minion backend.

#7 - 2018-10-08 07:19 - EDiGiacinto

Copied from the GH PR:

What is missing still:

- Digging into failures of other tests (compile check, and warning that makes the cache-service fail, which passes locally) and finish adapting logging messages changes that aren't matching anymore in fullstack (the cache messages now are in the minion worker, and fails for that reason)
- Documentation of the modules
- systemd services and package spec adjustments
- Acceptance testing
- Improve error messages reporting on worker side, but otoh the Cache service is exposing also Minion::Admin plugin
- Bump cache limit size (?) and polishing

#8 - 2018-10-08 07:34 - EDiGiacinto

- Related to action #34597: Race condition causing problems with the worker cache added

#9 - 2018-10-18 08:01 - coolo

- Category set to Concrete Bugs

- Status changed from New to In Progress

- Assignee set to EDiGiacinto

part of the worker story being worked on

#10 - 2018-10-18 08:02 - EDiGiacinto

- Related to action #40871: Asset removed during job runtime added

#11 - 2018-10-18 08:06 - EDiGiacinto

Update (what is missing still):

- Investigate/Fix GRU tests that failed after rebase
- Acceptance testing
- Fix bugs that were met during testing (e.g. Job should be flagged as incomplete when cache service is down, and on Leap42.3 the cache is failing in updating asset status, while on my TW this doesn't happen)

#12 - 2018-11-14 10:08 - EDiGiacinto

- Status changed from In Progress to Feedback

<https://github.com/os-autoinst/openQA/pull/1783> has been merged, setting to Feedback

#13 - 2018-11-27 09:32 - coolo

- Status changed from Feedback to Resolved

- Target version changed from Current Sprint to Done