

Klever - Bug #8941

Native Scheduler has a race when cancelling job decisions

06/09/2018 05:35 PM - Evgeny Novikov

Status:	Closed	Start date:	06/09/2018
Priority:	Urgent	Due date:	
Assignee:	Ilja Zakharov	% Done:	0%
Category:	Scheduling	Estimated time:	0.00 hour
Target version:	1.0	Published in build:	
Detected in build:	svn		
Platform:			
Description			
See details within the attached log file. Fortunately the issue seems to be extremely seldom since we didn't observe it so far.			

History

#1 - 06/10/2018 10:46 AM - Evgeny Novikov

Perhaps related issue when starting Native Scheduler after starting verification job decision:

```
root INFO> Create session for user "service" at Klever Bridge "localhost:8998"
root INFO> Scheduler base initialization has been successful
root INFO> Resource manager is live now
root INFO> Submit information about the workload to Bridge
root INFO> Initialize pool with 4 processes to run tasks and jobs
root INFO> Start scheduler loop
root INFO> Start scheduling iteration with statuses exchange with the server
root INFO> Scheduler has 0 pending, 0 processing, 0 finished and 0 error tasks
root INFO> Scheduler has 0 pending, 0 processing, 0 finished and 0 error jobs and 0 cancelled
root INFO> Going to prepare execution of the job 28cd81eca4a7996cd3c595a24a40abc3
root INFO> Going to check that deadlock are impossible
root INFO> Start scheduling new tasks
root INFO> Going to check that job '28cd81eca4a7996cd3c595a24a40abc3' does not introduce deadlocks
root INFO> Going to start 0 new tasks and 1 jobs
root INFO> Try to start job 28cd81eca4a7996cd3c595a24a40abc3
root INFO> Do not wait because of statuses changing
2018-06-10 10:42:54,882 SchedulerClient INFO> Going to solve a verification job
2018-06-10 10:42:54,883 SchedulerClient INFO> Start task execution with the following options: ['/home/novikov/work/klever/addons/benchexec/bin/runexec', '--memlimit', '1000MB', '--no-container', '/home/novikov/work/klever/scheduler/client/../../core/bin/klever-core']
root INFO> Start scheduling iteration with statuses exchange with the server
root INFO> Scheduler has 0 pending, 0 processing, 0 finished and 0 error tasks
root INFO> Scheduler has 1 pending, 0 processing, 0 finished and 0 error jobs and 1 cancelled
root INFO> Going to check that deadlock are impossible
root INFO> Submit information about the workload to Bridge
root INFO> Start scheduling new tasks
root INFO> Going to start 0 new tasks and 0 jobs
root INFO> Increase data exchange period from 1s to 2s
26443: Cancelling process 26442
26443: Cancellation of 26442 is successful, exiting
2018-06-10 10:42:56,619 SchedulerClient INFO> Runexec exited successfully but it is not possible to read job
exit code, aborting
2018-06-10 10:42:56,620 SchedulerClient INFO> Job solution has finished with exit code 1
2018-06-10 10:42:56,620 SchedulerClient INFO> Exiting with exit code 1
root INFO> Start scheduling iteration with statuses exchange with the server
root INFO> Scheduler has 0 pending, 0 processing, 0 finished and 0 error tasks
root INFO> Scheduler has 1 pending, 0 processing, 0 finished and 0 error jobs and 0 cancelled
root INFO> Skip the next data exchange iteration with Bridge
root INFO> Going to prepare execution of the job 28cd81eca4a7996cd3c595a24a40abc3
root ERROR> An error occurred:
Traceback (most recent call last):
  File "/home/novikov/work/klever/scheduler/bin/./schedulers/__init__.py", line 397, in launch
    self.__process_future(self.process_job_result, self.__jobs[job_id], job_id)
  File "/home/novikov/work/klever/scheduler/bin/./schedulers/__init__.py", line 733, in __process_future
    item["status"] = handler(identifier, item["future"]) if "future" in item else None
  File "/home/novikov/work/klever/scheduler/bin/./schedulers/native.py", line 221, in process_job_result
```

```
    return self._check_solution(identifiier, future, mode='job')
File "/home/novikov/work/klever/scheduler/bin/../schedulers/native.py", line 401, in _check_solution
    return self._postprocess_solution(identifiier, future, mode)
File "/home/novikov/work/klever/scheduler/bin/../schedulers/native.py", line 449, in _postprocess_solution
    reserved_space = utils.dir_size(work_dir)
File "/home/novikov/work/klever/scheduler/bin/../utils/__init__.py", line 231, in dir_size
    raise ValueError('Expect existing directory but it is not: {}'.format(dir))
ValueError: Expect existing directory but it is not: scheduler/jobs/28cd81eca4a7996cd3c595a24a40abc3
root INFO> Going to cancel execution of the job 28cd81eca4a7996cd3c595a24a40abc3
root ERROR> Uncaught exception
Traceback (most recent call last):
  File "/home/novikov/work/klever/scheduler/bin/../schedulers/__init__.py", line 397, in launch
    self.__process_future(self.process_job_result, self.__jobs[job_id], job_id)
  File "/home/novikov/work/klever/scheduler/bin/../schedulers/__init__.py", line 733, in __process_future
    item["status"] = handler(identifiier, item["future"] if "future" in item else None)
  File "/home/novikov/work/klever/scheduler/bin/../schedulers/native.py", line 221, in process_job_result
    return self._check_solution(identifiier, future, mode='job')
  File "/home/novikov/work/klever/scheduler/bin/../schedulers/native.py", line 401, in _check_solution
    return self._postprocess_solution(identifiier, future, mode)
  File "/home/novikov/work/klever/scheduler/bin/../schedulers/native.py", line 449, in _postprocess_solution
    reserved_space = utils.dir_size(work_dir)
  File "/home/novikov/work/klever/scheduler/bin/../utils/__init__.py", line 231, in dir_size
    raise ValueError('Expect existing directory but it is not: {}'.format(dir))
ValueError: Expect existing directory but it is not: scheduler/jobs/28cd81eca4a7996cd3c595a24a40abc3
```

During handling of the above exception, another exception occurred:

```
Traceback (most recent call last):
  File "/home/novikov/work/klever/scheduler/bin/native-scheduler", line 38, in <module>
    scheduler_impl.launch()
  File "/home/novikov/work/klever/scheduler/bin/../schedulers/__init__.py", line 498, in launch
    self.terminate()
  File "/home/novikov/work/klever/scheduler/bin/../schedulers/native.py", line 255, in terminate
    super(Scheduler, self).terminate()
  File "/home/novikov/work/klever/scheduler/bin/../schedulers/__init__.py", line 631, in terminate
    self.__process_future(self.cancel_job, self.__jobs[job_id], job_id)
  File "/home/novikov/work/klever/scheduler/bin/../schedulers/__init__.py", line 733, in __process_future
    item["status"] = handler(identifiier, item["future"] if "future" in item else None)
  File "/home/novikov/work/klever/scheduler/bin/../schedulers/native.py", line 232, in cancel_job
    return self._cancel_solution(identifiier, future, mode='job')
  File "/home/novikov/work/klever/scheduler/bin/../schedulers/native.py", line 425, in _cancel_solution
    return self._postprocess_solution(identifiier, future, mode)
  File "/home/novikov/work/klever/scheduler/bin/../schedulers/native.py", line 440, in _postprocess_solution
    del self._job_processes[identifiier]
KeyError: '28cd81eca4a7996cd3c595a24a40abc3'
```

#2 - 06/13/2018 04:47 PM - Ilija Zakharov

- Status changed from New to Resolved

This is a strange bug. First, I cannot reproduce it at all. Second, the same bug I fixed previously, so maybe you tried an outdated version. Third, you ran scheduler changing configuration turning on keeping dirs flag (which is a debug one but not the production option). Anyway I found a bug related to stopping the whole scheduler with ctrl+C and fixed it and added additional checks. If you will be able to reproduce something similar then additional efforts will be required. Fixes are in 8941-scheduler-race.

#3 - 06/13/2018 11:02 PM - Evgeny Novikov

- Status changed from Resolved to Closed

Let's hope the issue has gone. I merged the branch to master in [5732f00f](#).

But schedulers may need a more accurate synchronization discipline since there are quite much data shared between multiple processes. For instance, there may be a classical data race in the following code fragment:

```
if exists(a)
    do_something_with(a)
```

Files

info.zip	24.9 KB	06/09/2018	Evgeny Novikov
----------	---------	------------	----------------