

Torque

qstat -f accounting logs sometimes report "resources_used.vmem=0"

Issue: qstat -f accounting logs sometimes report "resources_used.vmem=0". This is due to how Linux reports memory usage

Demonstration:

In this test, I simply submit sleep jobs (of 5, 6, 7, and 8 seconds long, respectively), each requesting 1GB of virtual memory, and then print out the accounting end record once it completes. In a nutshell: you see the 2nd job reports "{resources_used.mem=2428kb resources_used.vmem=235592kb}", while the other 3 all report zeros for both. For cgroups, they all list mem and swap correctly.

```
[rmckay@ibm02 ~]$ pbs_server --version
Version: 6.0.2
Commit: 158053bc87018dc9525124f7788530cbc214dcb0
[rmckay@ibm02 ~]$
[rmckay@ibm02 ~]$
[rmckay@ibm02 ~]$
[rmckay@ibm02 ~]$ echo sleep 5 | qsub -l vmem=1gb && sleep 5
7693.ibm02.ac
[rmckay@ibm02 ~]$ qstat
Job ID Name User Time Use S Queue
-----
7693.ibm02.ac STDIN rmckay 00:00:00 E batch
[rmckay@ibm02 ~]$ sudo grep ";E;7693"
/var/spool/torque/server_priv/accounting/20160826
08/26/2016 14:45:07;E;7693.ibm02.ac;user=rmckay group=company
jobname=STDIN queue=batch ctime=1472244300 qtime=1472244300
etime=1472244300 start=1472244301 owner=rmckay@ibm02.blah
exec_host=ibm07.ac/0 Resource_List.pmem=768mb Resource_List.vmem=1gb
Resource_List.walltime=01:00:00 session=30003 total_execution_slots=1
unique_node_count=1 end=1472244307 Exit_status=0 resources_used.cput=0
resources_used.energy_used=0 resources_used.mem=0kb
resources_used.vmem=0kb resources_used.walltime=00:00:05
[rmckay@ibm02 ~]$
[rmckay@ibm02 ~]$
[rmckay@ibm02 ~]$ echo sleep 6 | qsub -l vmem=1gb && sleep 6
7694.ibm02.ac
[rmckay@ibm02 ~]$ qstat
Job ID Name User Time Use S Queue
-----
7693.ibm02.ac STDIN rmckay 00:00:00 C batch
7694.ibm02.ac STDIN rmckay 00:00:00 C batch
[rmckay@ibm02 ~]$ sudo grep ";E;7694"
/var/spool/torque/server_priv/accounting/20160826
08/26/2016 14:45:32;E;7694.ibm02.ac;user=rmckay group=company
jobname=STDIN queue=batch ctime=1472244325 qtime=1472244325
```

Torque

```
etime=1472244325 start=1472244326 owner=rmckay@ibm02.blah
exec_host=ibm07.ac/0 Resource_List.pmem=768mb Resource_List.vmem=1gb
Resource_List.walltime=01:00:00 session=30026 total_execution_slots=1
unique_node_count=1 end=1472244332 Exit_status=0 resources_used.cput=0
resources_used.energy_used=0 resources_used.mem=2428kb
resources_used.vmem=235592kb resources_used.walltime=00:00:06
```

```
[rmckay@ibm02 ~]$
```

```
[rmckay@ibm02 ~]$
```

```
[rmckay@ibm02 ~]$ echo sleep 7 | qsub -l vmem=1gb && sleep 7
```

```
7695.ibm02.ac
```

```
[rmckay@ibm02 ~]$ qstat
```

```
Job ID Name User Time Use S Queue
```

```
-----
```

```
7693.ibm02.ac STDIN rmckay 00:00:00 C batch
```

```
7694.ibm02.ac STDIN rmckay 00:00:00 C batch
```

```
7695.ibm02.ac STDIN rmckay 00:00:00 C batch
```

```
[rmckay@ibm02 ~]$ sudo grep ";E;7695"
```

```
/var/spool/torque/server_priv/accounting/20160826
```

```
08/26/2016 14:46:00;E;7695.ibm02.ac;user=rmckay group=company
```

```
jobname=STDIN queue=batch ctime=1472244352 qtime=1472244352
```

```
etime=1472244352 start=1472244353 owner=rmckay@ibm02.blah
```

```
exec_host=ibm07.ac/0 Resource_List.pmem=768mb Resource_List.vmem=1gb
```

```
Resource_List.walltime=01:00:00 session=30052 total_execution_slots=1
```

```
unique_node_count=1 end=1472244360 Exit_status=0 resources_used.cput=0
```

```
resources_used.energy_used=0 resources_used.mem=0kb
```

```
resources_used.vmem=0kb resources_used.walltime=00:00:07
```

```
[rmckay@ibm02 ~]$
```

```
[rmckay@ibm02 ~]$
```

```
[rmckay@ibm02 ~]$ echo sleep 8 | qsub -l vmem=1gb && sleep 8
```

```
7696.ibm02.ac
```

```
[rmckay@ibm02 ~]$ qstat
```

```
Job ID Name User Time Use S Queue
```

```
-----
```

```
7693.ibm02.ac STDIN rmckay 00:00:00 C batch
```

```
7694.ibm02.ac STDIN rmckay 00:00:00 C batch
```

```
7695.ibm02.ac STDIN rmckay 00:00:00 C batch
```

```
7696.ibm02.ac STDIN rmckay 00:00:00 C batch
```

```
[rmckay@ibm02 ~]$ sudo grep ";E;7696"
```

```
/var/spool/torque/server_priv/accounting/20160826
```

```
08/26/2016 14:46:28;E;7696.ibm02.ac;user=rmckay group=company
```

```
jobname=STDIN queue=batch ctime=1472244379 qtime=1472244379
```

```
etime=1472244379 start=1472244379 owner=rmckay@ibm02.blah
```

```
exec_host=ibm07.ac/0 Resource_List.pmem=768mb Resource_List.vmem=1gb
```

```
Resource_List.walltime=01:00:00 session=30078 total_execution_slots=1
```

```
unique_node_count=1 end=1472244388 Exit_status=0 resources_used.cput=0
```

```
resources_used.energy_used=0 resources_used.mem=0kb
```

```
resources_used.vmem=0kb resources_used.walltime=00:00:08
```

```
[rmckay@ibm02 ~]$ qstat -f
```

```
Job Id: 7693.ibm02.ac
```

Torque

Job_Name = STDIN
Job_Owner = rmckay@ibm02.blah
resources_used.cput = 00:00:00
resources_used.energy_used = 0
resources_used.mem = 0kb
resources_used.vmem = 0kb
resources_used.walltime = 00:00:05
job_state = C
queue = batch
server = ibm02.ac
Checkpoint = u
ctime = Fri Aug 26 14:45:00 2016
Error_Path = ibm02.ac:/home/rmckay/STDIN.e7693
exec_host = ibm07.ac/0
Hold_Types = n
Join_Path = n
Keep_Files = n
Mail_Points = a
mtime = Fri Aug 26 14:45:07 2016
Output_Path = ibm02.ac:/home/rmckay/STDIN.o7693
Priority = 0
qtime = Fri Aug 26 14:45:00 2016
Rerunable = True
Resource_List.pmem = 768mb
Resource_List.vmem = 1gb
Resource_List.walltime = 01:00:00
session_id = 30003
substate = 59
Variable_List = PBS_O_QUEUE=batch,PBS_O_HOME=/home/rmckay,
PBS_O_LOGNAME=rmckay,
PBS_O_PATH=/home/rmckay/.local/bin:/home/rmckay/bin:/usr/local/bin:/o
pt/src/build/moab/9.0.2.h1/bin:/usr/local/bin:/usr/bin:/usr/local/sbin
:/usr/sbin,PBS_O_MAIL=/var/spool/mail/rmckay,PBS_O_SHELL=/bin/bash,
PBS_O_LANG=en_US.UTF-8,PBS_O_WORKDIR=/home/rmckay,PBS_O_HOST=ibm02.
ac,
PBS_O_SERVER=ibm02.ac
euser = rmckay
egroup = company
hashname = 7693.ibm02.ac
queue_rank = 353
queue_type = E
etime = Fri Aug 26 14:45:00 2016
exit_status = 0
submit_args = -l vmem=1gb
start_time = Fri Aug 26 14:45:01 2016
start_count = 1
fault_tolerant = False
comp_time = Fri Aug 26 14:45:07 2016
job_radix = 0

Torque

```
total_runtime = 5.642798
submit_host = ibm02.ac
request_version = 1
req_information.task_count.0 = 1
req_information.lprocs.0 = 1
req_information.memory.0 = 786432kb
req_information.swap.0 = 1048576kb
req_information.thread_usage_policy.0 = allowthreads
req_information.hostlist.0 = ibm07.ac:ppn=1
req_information.task_usage.0.task.0.cpu_list = 0
req_information.task_usage.0.task.0.mem_list = 0
req_information.task_usage.0.task.0.cores = 0
req_information.task_usage.0.task.0.threads = 1
req_information.task_usage.0.task.0.host = ibm07.ac
```

```
Job Id: 7694.ibm02.ac
Job_Name = STDIN
Job_Owner = rmckay@ibm02.blah
resources_used.cput = 00:00:00
resources_used.energy_used = 0
resources_used.mem = 2428kb
resources_used.vmem = 235592kb
resources_used.walltime = 00:00:06
job_state = C
queue = batch
server = ibm02.ac
Checkpoint = u
ctime = Fri Aug 26 14:45:25 2016
Error_Path = ibm02.ac:/home/rmckay/STDIN.e7694
exec_host = ibm07.ac/0
Hold_Types = n
Join_Path = n
Keep_Files = n
Mail_Points = a
mtime = Fri Aug 26 14:45:32 2016
Output_Path = ibm02.ac:/home/rmckay/STDIN.o7694
Priority = 0
qtime = Fri Aug 26 14:45:25 2016
Rerunable = True
Resource_List.pmem = 768mb
Resource_List.vmem = 1gb
Resource_List.walltime = 01:00:00
session_id = 30026
substate = 59
Variable_List = PBS_O_QUEUE=batch,PBS_O_HOME=/home/rmckay,
PBS_O_LOGNAME=rmckay,
PBS_O_PATH=/home/rmckay/.local/bin:/home/rmckay/bin:/usr/local/bin:/o
pt/src/build/moab/9.0.2.h1/bin:/usr/local/bin:/usr/bin:/usr/local/sbin
:/usr/sbin,PBS_O_MAIL=/var/spool/mail/rmckay,PBS_O_SHELL=/bin/bash,
```

Torque

PBS_O_LANG=en_US.UTF-8,PBS_O_WORKDIR=/home/rmckay,PBS_O_HOST=ibm02.
ac,
PBS_O_SERVER=ibm02.ac
euser = rmckay
egroup = company
hashname = 7694.ibm02.ac
queue_rank = 354
queue_type = E
etime = Fri Aug 26 14:45:25 2016
exit_status = 0
submit_args = -l vmem=1gb
start_time = Fri Aug 26 14:45:26 2016
start_count = 1
fault_tolerant = False
comp_time = Fri Aug 26 14:45:32 2016
job_radix = 0
total_runtime = 6.643986
submit_host = ibm02.ac
request_version = 1
req_information.task_count.0 = 1
req_information.lprocs.0 = 1
req_information.memory.0 = 786432kb
req_information.swap.0 = 1048576kb
req_information.thread_usage_policy.0 = allowthreads
req_information.hostlist.0 = ibm07.ac:ppn=1
req_information.task_usage.0.task.0.cpu_list = 0
req_information.task_usage.0.task.0.mem_list = 0
req_information.task_usage.0.task.0.cores = 0
req_information.task_usage.0.task.0.threads = 1
req_information.task_usage.0.task.0.host = ibm07.ac

Job Id: 7695.ibm02.ac

Job_Name = STDIN
Job_Owner = rmckay@ibm02.blah
resources_used.cput = 00:00:00
resources_used.energy_used = 0
resources_used.mem = 0kb
resources_used.vmem = 0kb
resources_used.walltime = 00:00:07
job_state = C
queue = batch
server = ibm02.ac
Checkpoint = u
ctime = Fri Aug 26 14:45:52 2016
Error_Path = ibm02.ac:/home/rmckay/STDIN.e7695
exec_host = ibm07.ac/0
Hold_Types = n
Join_Path = n
Keep_Files = n

Torque

Mail_Points = a
mtime = Fri Aug 26 14:46:00 2016
Output_Path = ibm02.ac:/home/rmckay/STDIN.o7695
Priority = 0
qtime = Fri Aug 26 14:45:52 2016
Rerunable = True
Resource_List.pmem = 768mb
Resource_List.vmem = 1gb
Resource_List.walltime = 01:00:00
session_id = 30052
substate = 59
Variable_List = PBS_O_QUEUE=batch,PBS_O_HOME=/home/rmckay,
PBS_O_LOGNAME=rmckay,
PBS_O_PATH=/home/rmckay/.local/bin:/home/rmckay/bin:/usr/local/bin:/o
pt/src/build/moab/9.0.2.h1/bin:/usr/local/bin:/usr/bin:/usr/local/sbin
:/usr/sbin,PBS_O_MAIL=/var/spool/mail/rmckay,PBS_O_SHELL=/bin/bash,
PBS_O_LANG=en_US.UTF-8,PBS_O_WORKDIR=/home/rmckay,PBS_O_HOST=ibm02.
ac,
PBS_O_SERVER=ibm02.ac
euser = rmckay
egroup = company
hashname = 7695.ibm02.ac
queue_rank = 355
queue_type = E
etime = Fri Aug 26 14:45:52 2016
exit_status = 0
submit_args = -l vmem=1gb
start_time = Fri Aug 26 14:45:53 2016
start_count = 1
fault_tolerant = False
comp_time = Fri Aug 26 14:46:00 2016
job_radix = 0
total_runtime = 7.632436
submit_host = ibm02.ac
request_version = 1
req_information.task_count.0 = 1
req_information.lprocs.0 = 1
req_information.memory.0 = 786432kb
req_information.swap.0 = 1048576kb
req_information.thread_usage_policy.0 = allowthreads
req_information.hostlist.0 = ibm07.ac:ppn=1
req_information.task_usage.0.task.0.cpu_list = 0
req_information.task_usage.0.task.0.mem_list = 0
req_information.task_usage.0.task.0.cores = 0
req_information.task_usage.0.task.0.threads = 1
req_information.task_usage.0.task.0.host = ibm07.ac

Job Id: 7696.ibm02.ac
Job_Name = STDIN

Torque

Job_Owner = rmckay@ibm02.blah
resources_used.cput = 00:00:00
resources_used.energy_used = 0
resources_used.mem = 0kb
resources_used.vmem = 0kb
resources_used.walltime = 00:00:08
job_state = C
queue = batch
server = ibm02.ac
Checkpoint = u
ctime = Fri Aug 26 14:46:19 2016
Error_Path = ibm02.ac:/home/rmckay/STDIN.e7696
exec_host = ibm07.ac/0
Hold_Types = n
Join_Path = n
Keep_Files = n
Mail_Points = a
mtime = Fri Aug 26 14:46:28 2016
Output_Path = ibm02.ac:/home/rmckay/STDIN.o7696
Priority = 0
qtime = Fri Aug 26 14:46:19 2016
Rerunable = True
Resource_List.pmem = 768mb
Resource_List.vmem = 1gb
Resource_List.walltime = 01:00:00
session_id = 30078
substate = 59
Variable_List = PBS_O_QUEUE=batch,PBS_O_HOME=/home/rmckay,
PBS_O_LOGNAME=rmckay,
PBS_O_PATH=/home/rmckay/.local/bin:/home/rmckay/bin:/usr/local/bin:/o
pt/src/build/moab/9.0.2.h1/bin:/usr/local/bin:/usr/bin:/usr/local/sbin
:/usr/sbin,PBS_O_MAIL=/var/spool/mail/rmckay,PBS_O_SHELL=/bin/bash,
PBS_O_LANG=en_US.UTF-8,PBS_O_WORKDIR=/home/rmckay,PBS_O_HOST=ibm02.
ac,
PBS_O_SERVER=ibm02.ac
euser = rmckay
egroup = company
hashname = 7696.ibm02.ac
queue_rank = 356
queue_type = E
etime = Fri Aug 26 14:46:19 2016
exit_status = 0
submit_args = -l vmem=1gb
start_time = Fri Aug 26 14:46:19 2016
start_count = 1
fault_tolerant = False
comp_time = Fri Aug 26 14:46:28 2016
job_radix = 0
total_runtime = 8.634951

Torque

```
submit_host = ibm02.ac
request_version = 1
req_information.task_count.0 = 1
req_information.lprocs.0 = 1
req_information.memory.0 = 786432kb
req_information.swap.0 = 1048576kb
req_information.thread_usage_policy.0 = allowthreads
req_information.hostlist.0 = ibm07.ac:ppn=1
req_information.task_usage.0.task.0.cpu_list = 0
req_information.task_usage.0.task.0.mem_list = 0
req_information.task_usage.0.task.0.cores = 0
req_information.task_usage.0.task.0.threads = 1
req_information.task_usage.0.task.0.host = ibm07.ac
```

```
[rmckay@ibm02 ~]$
```

Explanation: When running sleep jobs we notice that short sleep jobs will use 0 vmem. Whereas longer jobs (30 seconds) begin to use more vmem. We have verified the usage by looking at the `proc/<job pid>/stat` file while the job is running and the totals reported in `qstat -f` for mem and vmem agree with what is reported by the os in stat files.

So the discrepancy in different jobs can be attributed to the Linux OS and how it is handling memory or the updating of the stat files.

Unique solution ID: #1149

Author: Shawn Hoopes

Last update: 2017-01-12 22:26