Launching parallel EUTERPE job with srun vs mpiexec

N. Shukla, A. Marani & D. Molinari

High-Performance Computing Department
CINECA, Casalecchio di Reno
Bologna Italy

Email: n.shukla@cineca.it

www.cineca.it || Bologna December 15th 2021

Launching parallel EUTERPE job with srun vs mpiexec

Methodology

- Using srun to launch euterpe on sparse nodes
- Using mpiexec on sparse nodes
- Using mpiexec on sparse nodes but with the flag

  --hostfile ./nodelist

For each case we run 99 jobs each one running Euterpe on 32 nodes up to fill the entire cluster and check if the job starts correctly or is found in hang after few minutes of run.

Results: Using srun to launch euterpe on sparse nodes

Total jobs analized: 99

- ✓ Total jobs successful : 99
- ✳ Total nodes con job in hang: 0

Results: Using mpiexec on sparse nodes

Run1: Total jobs analized: 99

- ✓ Total jobs successful : 16
- ✳ Total nodes con job in hang: 83

Run 2: Total jobs analized: 99

- ✓ Total jobs successful : 51
- ✳ Total nodes con job in hang: 48

Run 3: Total jobs analized: 99

- ✓ Total jobs successful : 92
- ✳ Total nodes con job in hang: 7

Summary:

Total jobs analized: 297
Total jobs successful : 159
Total nodes con job in hang: 138

Results: Using mpiexec on sparse nodes

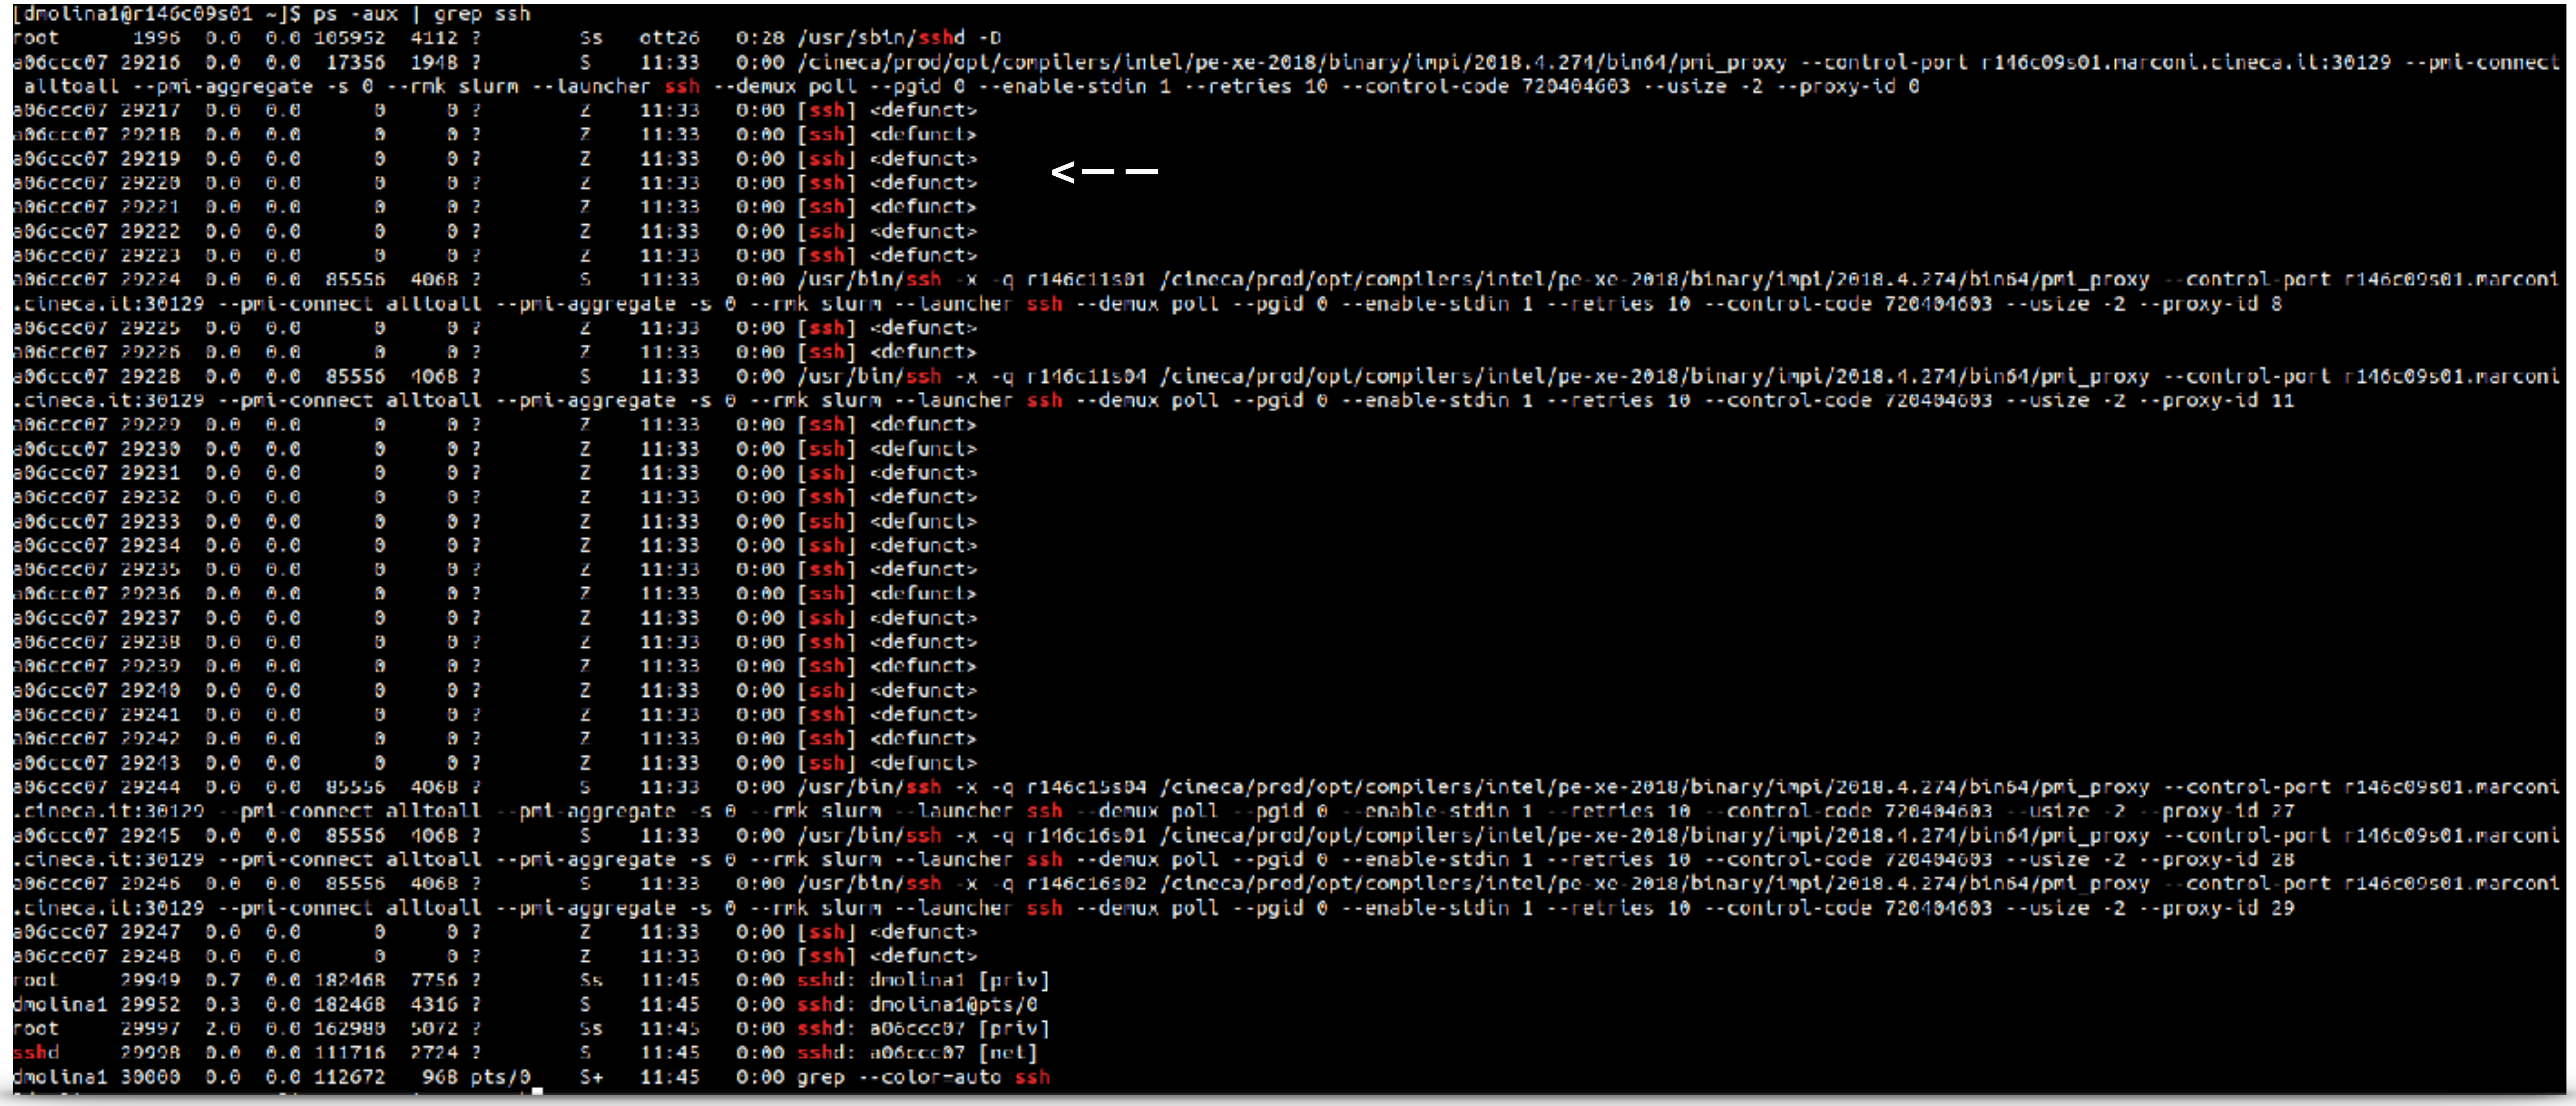

```
[dmolina1@r146c09s01 ~]$ ps -aux | grep ssh
root       1996  0.0  0.0 105952  4112 ?        Ss   ott26   0:28 /usr/sbin/sshd -D
a06ccc07 29216  0.0  0.0  17356  1948 ?        S    11:33   0:00 /cineca/prod/opt/compilers/intel/pe-xe-2018/binary/impi/2018.4.274/bin64/pmi_proxy --control-port r146c09s01.marconi.cineca.it:30129 --pmi-connect
 alltoall --pmi-aggregate -s 0 --rmk slurm --launcher ssh --demux poll --pgid 0 --enable-stdin 1 --retries 10 --control-code 720404603 --usize -2 --proxy-id 0
a06ccc07 29217  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29218  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29219  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>            <--
a06ccc07 29220  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29221  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29222  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29223  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29224  0.0  0.0  85556  4068 ?        S    11:33   0:00 /usr/bin/ssh -x -q r146c11s01 /cineca/prod/opt/compilers/intel/pe-xe-2018/binary/impi/2018.4.274/bin64/pmi_proxy --control-port r146c09s01.marconi
.cineca.it:30129 --pmi-connect alltoall --pmi-aggregate -s 0 --rmk slurm --launcher ssh --demux poll --pgid 0 --enable-stdin 1 --retries 10 --control-code 720404603 --usize -2 --proxy-id 8
a06ccc07 29225  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29226  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29228  0.0  0.0  85556  4068 ?        S    11:33   0:00 /usr/bin/ssh -x -q r146c11s04 /cineca/prod/opt/compilers/intel/pe-xe-2018/binary/impi/2018.4.274/bin64/pmi_proxy --control-port r146c09s01.marconi
.cineca.it:30129 --pmi-connect alltoall --pmi-aggregate -s 0 --rmk slurm --launcher ssh --demux poll --pgid 0 --enable-stdin 1 --retries 10 --control-code 720404603 --usize -2 --proxy-id 11
a06ccc07 29229  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29230  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29231  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29232  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29233  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29234  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29235  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29236  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29237  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29238  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29239  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29240  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29241  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29242  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29243  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29244  0.0  0.0  85556  4068 ?        S    11:33   0:00 /usr/bin/ssh -x -q r146c15s04 /cineca/prod/opt/compilers/intel/pe-xe-2018/binary/impi/2018.4.274/bin64/pmi_proxy --control-port r146c09s01.marconi
.cineca.it:30129 --pmi-connect alltoall --pmi-aggregate -s 0 --rmk slurm --launcher ssh --demux poll --pgid 0 --enable-stdin 1 --retries 10 --control-code 720404603 --usize -2 --proxy-id 27
a06ccc07 29245  0.0  0.0  85556  4068 ?        S    11:33   0:00 /usr/bin/ssh -x -q r146c16s01 /cineca/prod/opt/compilers/intel/pe-xe-2018/binary/impi/2018.4.274/bin64/pmi_proxy --control-port r146c09s01.marconi
.cineca.it:30129 --pmi-connect alltoall --pmi-aggregate -s 0 --rmk slurm --launcher ssh --demux poll --pgid 0 --enable-stdin 1 --retries 10 --control-code 720404603 --usize -2 --proxy-id 28
a06ccc07 29246  0.0  0.0  85556  4068 ?        S    11:33   0:00 /usr/bin/ssh -x -q r146c16s02 /cineca/prod/opt/compilers/intel/pe-xe-2018/binary/impi/2018.4.274/bin64/pmi_proxy --control-port r146c09s01.marconi
.cineca.it:30129 --pmi-connect alltoall --pmi-aggregate -s 0 --rmk slurm --launcher ssh --demux poll --pgid 0 --enable-stdin 1 --retries 10 --control-code 720404603 --usize -2 --proxy-id 29
a06ccc07 29247  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
a06ccc07 29248  0.0  0.0      0     0 ?        Z    11:33   0:00 [ssh] <defunct>
root      29949  0.7  0.0 182468  7756 ?        Ss   11:45   0:00 sshd: dmolina1 [priv]
dmolina1 29952  0.3  0.0 182468  4316 ?        S    11:45   0:00 sshd: dmolina1@pts/0
root      29997  2.0  0.0 162980  5072 ?        Ss   11:45   0:00 sshd: a06ccc07 [priv]
sshd      29998  0.0  0.0 111716  2724 ?        S    11:45   0:00 sshd: a06ccc07 [net]
dmolina1 30000  0.0  0.0 112672   968 pts/0    S+   11:45   0:00 grep --color=auto ssh
```

Resolving The Problem: Wait for the parent process to complete

Results: Using mpiexec but with a sleep of 30 seconds before launching mpiexec

Run 4: Total jobs analized: 99

- ✓ Total jobs successful : 99
- ✴ Total nodes con job in hang: 0

Results: Using mpiexec on sparse nodes but with the flag --hostfile ./hostfile

Run 4: Total jobs analized: 99

- ✓ Total jobs successful : 99
- ✳ Total nodes con job in hang: 0

Summary run

We repeated the mpiexec run to gain statistic, because it was the only one showing jobs in hang state

Case	Total jobs	Successful	Hang
mpiexec	297	159	138
srun	99	99	0
mpiexec --hostfile ./nodelist	99	99	0
sleep(30); mpiexec	99	99	0

Summary

- Launching with srun things goes fine.
- launching with mpiexec, the job may go in hang due to ssh processes denied because of the concurrency between mpiexec and slurm notifying the nodes to be in the job.
- The effect are ssh processes in state "defunct" the master node as in the snapshot attached.

Please use srun

Regarding the mpiexec runs we also found some workarounds

- Launching mpiexec with the flag --hostfile ./nodelist
- adding a brief sleep before the mpiexec command