# GROMACS - Bug #2981

## segfault in opencl build

06/18/2019 03:30 AM - Mark Abraham

| | | | |
|---|---|---|---|
| **Status:** | New | | |
| **Priority:** | Normal | | |
| **Assignee:** | | | |
| **Category:** | mdrun | | |
| **Target version:** | | | |
| **Affected version - extra info:** | | **Difficulty:** | uncategorized |
| **Affected version:** | git master | | |

**Description**

On an agent config gcc-8 openmp simd=avx2_256 gpuhw=amd opencl-1.2 clFFT-2.14 buildfftw host=bs_gpu01,bs_gpu01, an unrelated change produced

```
GROMACS:      gmx mdrun, version 2020-dev-20190617-632bc08-local
Executable:   /home/jenkins/workspace/Matrix_PreSubmit_master/756cb28b/gromacs/bin/gmx
Data prefix:  /home/jenkins/workspace/Matrix_PreSubmit_master/756cb28b/gromacs (source tree)
Working dir:  /mnt/workspace/Matrix_PreSubmit_master/756cb28b/regressiontests/complex/nbnxn_pme_or
der5
Command line:
  gmx mdrun -ntmpi 2 -ntomp 2 -notunepme

Reading file topol.tpr, VERSION 2020-dev-20190617-632bc08-local (single precision)
Changing nstlist from 10 to 100, rlist from 0.9 to 0.999

Using 2 MPI threads
Using 2 OpenMP threads per tMPI thread

On host bs-gpu01 2 GPUs selected for this run.
Mapping of GPU IDs to the 2 GPU tasks in the 2 ranks on this node:
  PP:0,PP:1
PP tasks will do (non-perturbed) short-ranged interactions on the GPU

NOTE: The number of threads is not equal to the number of (logical) cores
      and the -pin option is set to auto: will not pin threads to cores.
      This can lead to significant performance degradation.
      Consider using -pin on (and -pinoffset in case you run multiple jobs).
Segmentation fault (core dumped)
```

Note that only PME order 4 is supported on GPUs, so this issue is not in the PME-on-GPU code.