

General information

We have 8 KNLs in the SDV right now.

All KNLs nodes have their own local NVMe device installed.

Node allocation

Nodes can be allocated through the PBS based batch system that is also used for the DEEP Cluster and Booster. You can start an interactive session on our KNLs like this:

```
qsub -l nodes=2:ppn=256:kn1,walltime=00:10:00 -I    (this will randomly choose 2 of our 8 KNLs)
qsub -l nodes=kn10X:ppn=256:kn1,walltime=00:10:00 -I (you can also choose a specific KNL like this)
```

To use the knls (4 nodes) reserved for the DEEP-ER people (every day from 8:30am-5:00pm) please use:

```
qsub -q deeper-knl -l nodes=X:ppn=Y:kn1,walltime=...
```

When using a batch script, you have to adapt the -l (-q) option within your script.

Compiling

Use the -xMIC-AVX512 flag instead of -mmic.

Check actual vectorisation with -qopt-report=5 -qopt-report-phase=vec → info given in *.optrpt files

5 things to consider when using KNL

- i. Make sure to use the fast MCDRAM:
 - When MCDRAM is in cache mode:
 - No changes are needed.
 - When MCDRAM is in flat mode:
 - If the total memory footprint of the application is smaller than the size of MCDRAM: numactl ?m 1 ./my_application.out (Allocations that don't fit into MCDRAM make the application fail.)
 - If the total memory footprint of the application is larger than the size of MCDRAM: numactl ?p 1 ./my_application.out (Allocations that don't fit into MCDRAM spill over to DDR)
 - To make a manual choice of what should be allocated in the MCDRAM: Use the memkind library.
- i. Verify that the pinning is as you wish:
 - Start job on KNL node(s).
 - Log in on KNL.
 - Invoke htop.
 - Check the load distribution.
 - Remark: Each core can execute 1, 2 or 4 threads. On KNL ? unlike on KNC ? already one thread per core can lead to optimal performance.
- i. Use VTune/Advisor to analyse the performance:
 - Start job on KNL node(s).
 - Log in on KNL.
 - 'module load VTune / Advisor'.
 - Run ampxe-gui / advixe-gui.
 - Follow instructions.
 - Remark: If you run into erros of the sort ?sepd not available? please contact the administrator. Both tools rely on a kernel module to access hardware counter.
- i. Provide hints to the compiler:
 - Check *optrpt for info on vectorisation.
 - If you find ?unaligned...? → add alignment in your code by adding "#pragma vector aligned" before the loop.
 - If a loop does not vectorise although it clearly should, you can add "#pragma simd" before the loop.
 - Re-check *.optrpt.

- Re-check in VTune / Advisor
- i. Verify the performance via benchmarks:
 - Set up JUBE for your code.
 - Benchmark the various versions with proper timing.
 - Be aware: VTune / Advisor sometimes give estimates that are a little off. It's imperative to check the actual performance.