
# KNL Performance Report
# NAMD

13 February 2017

# 1. Compilation, Setup and Input

## Compilation

NAMD was compiled on both the ARCHER KNL system and ARCHER Xeon system using the instructions found online at:

https://github.com/ARCHER-CSE/build-instructions/blob/master/NAMD/build_namd_2.11_knl.md

(this KNL compile procedure was also used for the Xeon system).

Compiler and library versions used were:

| Compiler/Library | ARCHER Xeon | ARCHER KNL |
|---|---|---|
| **Intel Compiler** | 15.0.2.164 | 17.0.0.098 |
| **FFTW** | 3.3.4.5 | 3.3.4.10 |

## Setup

- The ARCHER KNL nodes were used in "quad_100" configuration with all the MCDRAM configured as an additional cache level.
- In all cases, jobs were run on fully-populated nodes with the number of processes and threads varied to ensure all Hyperthreads were occupied.

## Benchmark Case

We used the STMV benchmark from the NAMD website at:

http://www.ks.uiuc.edu/Research/namd/performance.html

This system consists of 1 million atoms and the input was modified to run for 1000 simulation steps with 2 fs timestep, 12 angstrom cutoff + PME every 3 steps.

## 2. Performance Results

The plot below shows a performance comparison between KNL and Xeon system for best performing configuration in each case. This corresponds to two MPI processes per node with two or four hyperthreads for KNL (64 or 128 threads per MPI process) and single MPI process per node with two hyperthreads for Xeon (48 threads per MPI process).

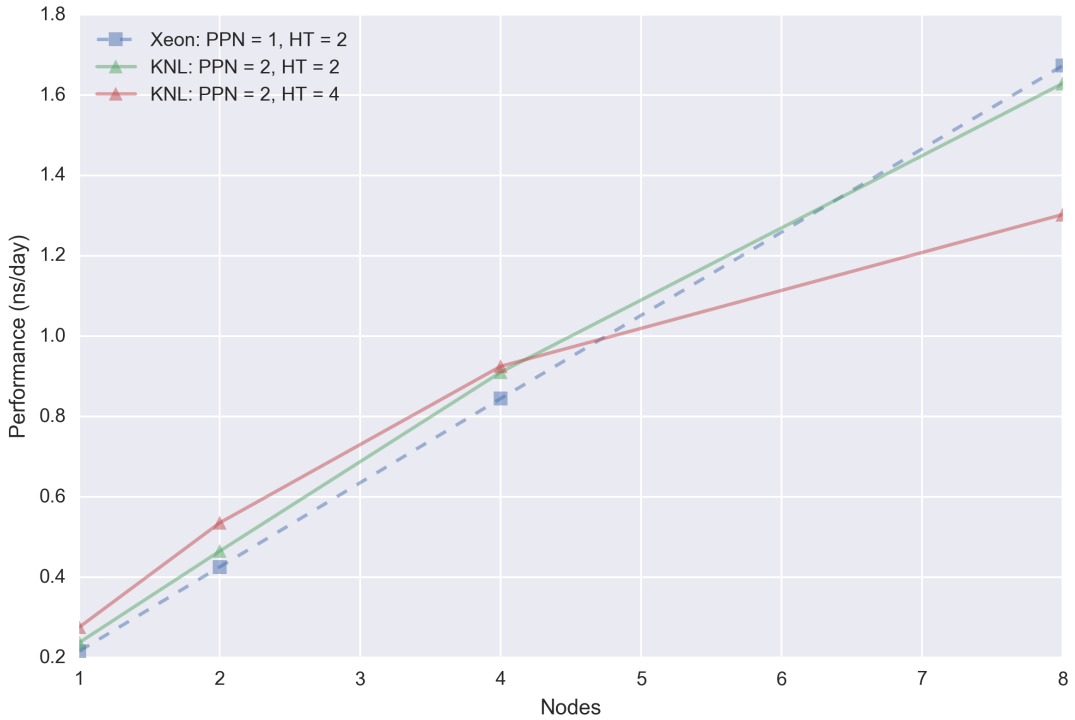


In general, each KNL node (with a single KNL processor) provides similar performance to a Xeon node (with two Xeon processors). For 4 nodes or less on the KNL system it is best to use 4 hyperthreads and beyond 4 nodes the best performance is gained by using 2 hyperthreads. For the Xeon system it is always best to use the maximum number of 2 hyperthreads. The table below compares KNL to Xeon performance (for the best performing configuration in each case) as a function of number of nodes:

| Nodes | KNL (ns/day) | Xeon (ns/day) | KNL:Xeon |
|---|---|---|---|
| **1** | 0.27 | 0.22 | 1.28 |
| **2** | 0.54 | 0.43 | 1.26 |
| **4** | 1.00 | 0.84 | 1.19 |
| **8** | 1.63 | 1.67 | 0.97 |

It is possible to improve the KNL performance on 8 nodes fractionally by changing the configuration to use 4 processes per node along with 2 hyperthreads and this puts the performance almost identical to 8 Xeon nodes.

Performance differences with number of hyperthreads on a single node:

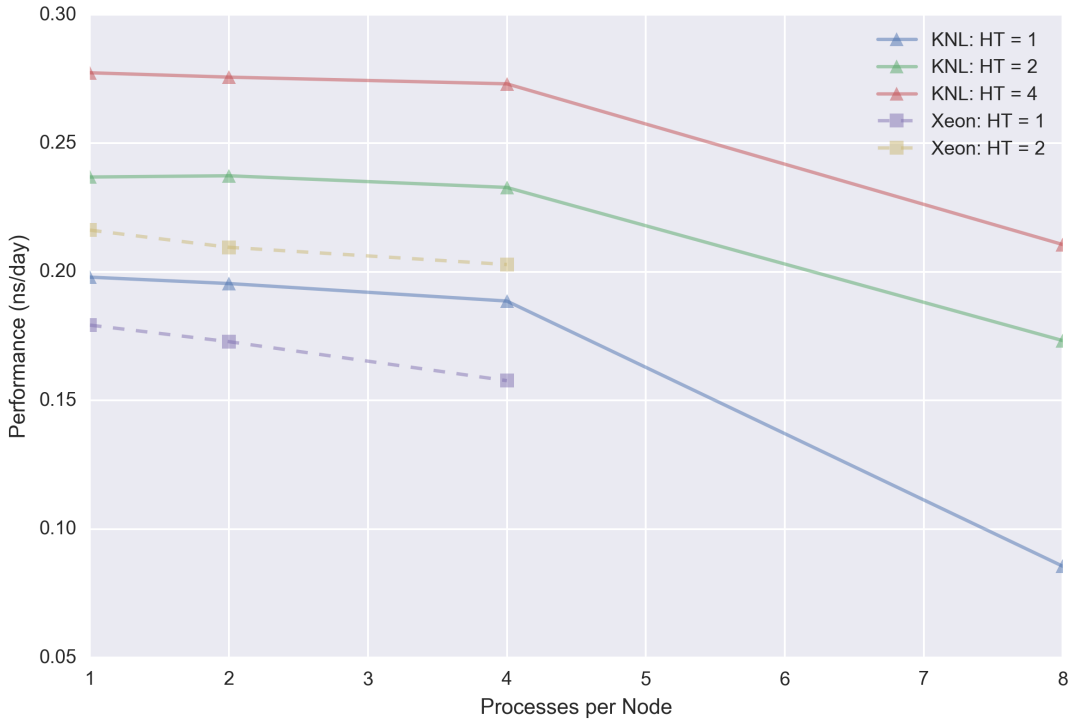


In general, NAMD performance improves as the number of hyperthreads is increased and as the number of MPI processes per node is decreased (and hence the number of threads per process is increased). It is particularly important on the KNL system to use four processes per node or less to get good performance.

Performance difference for single hyperthread with different numbers of threads per process (i.e. different numbers of MPI processes per node):

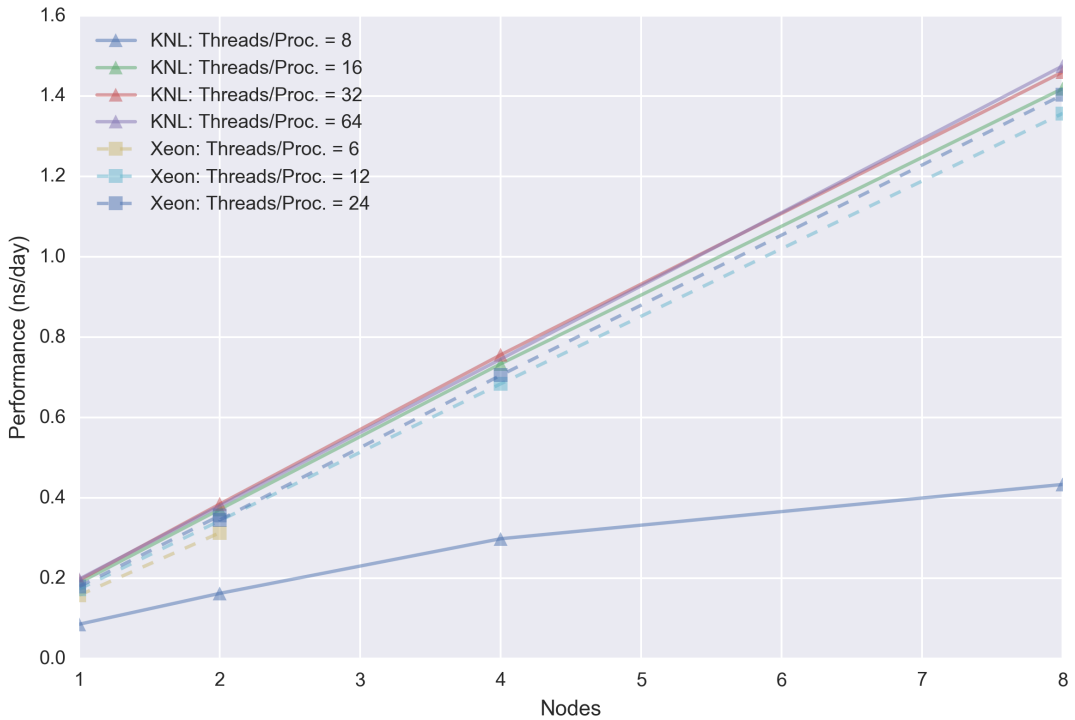


This plot shows that the maximum number of threads per process (i.e. minimum number of processes per node) should be used to achieve maximum performance.

# 3. Summary and Conclusions

- Below 8 nodes the KNL nodes provide 20-25% better performance than Xeon nodes for this NAMD benchmark. At 8 nodes, the performance of KNL and Xeon is very similar.
- For both KNL and Xeon, the maximum number of hyperthreads should be used to get best performance. On KNL the difference between 2 and 4 hyperthreads is minimal.
- Generally, for both KNL and Xeon, the maximum number of threads per process should be used. This corresponds to a single MPI process per node. For KNL on 8 nodes, this rule is broken with the best performance seen with 4 MPI processes per node and 2 hyperthreads.
- Benchmarking is required on larger KNL systems to understand how the performance varies as the number of nodes is increased.

## 4. Full Performance Results

| System | Nodes | Hyperthreads | Processes | Processes per Node | Threads per Process | Performance days/ns |
|---|---|---|---|---|---|---|
| KNL | 1 | 1 | 1 | 1 | 64 | 5.05 |
| KNL | 1 | 2 | 1 | 1 | 128 | 4.22 |
| KNL | 1 | 4 | 1 | 1 | 256 | 3.60 |
| KNL | 1 | 1 | 2 | 2 | 32 | 5.12 |
| KNL | 1 | 2 | 2 | 2 | 64 | 4.21 |
| KNL | 1 | 4 | 2 | 2 | 128 | 3.62 |
| KNL | 1 | 1 | 4 | 4 | 16 | 5.30 |
| KNL | 1 | 2 | 4 | 4 | 32 | 4.30 |
| KNL | 1 | 4 | 4 | 4 | 64 | 3.66 |
| KNL | 1 | 1 | 8 | 8 | 8 | 11.69 |
| KNL | 1 | 2 | 8 | 8 | 16 | 5.77 |
| KNL | 1 | 4 | 8 | 8 | 32 | 4.75 |
| KNL | 2 | 1 | 2 | 1 | 64 | 2.59 |
| KNL | 2 | 2 | 2 | 1 | 128 | 2.18 |
| KNL | 2 | 4 | 2 | 1 | 256 | 1.95 |
| KNL | 2 | 1 | 4 | 2 | 32 | 2.60 |
| KNL | 2 | 2 | 4 | 2 | 64 | 2.15 |
| KNL | 2 | 4 | 4 | 2 | 128 | 1.87 |
| KNL | 2 | 1 | 8 | 4 | 16 | 2.69 |
| KNL | 2 | 2 | 8 | 4 | 32 | 2.17 |
| KNL | 2 | 4 | 8 | 4 | 64 | 1.90 |
| KNL | 2 | 1 | 16 | 8 | 8 | 6.16 |
| KNL | 2 | 2 | 16 | 8 | 16 | 3.21 |
| KNL | 2 | 4 | 16 | 8 | 32 | 2.88 |
| KNL | 4 | 1 | 4 | 1 | 64 | 1.32 |
| KNL | 4 | 2 | 4 | 1 | 128 | 1.12 |
| KNL | 4 | 4 | 4 | 1 | 256 | 1.41 |
| KNL | 4 | 1 | 8 | 2 | 32 | 1.32 |
| KNL | 4 | 2 | 8 | 2 | 64 | 1.10 |
| KNL | 4 | 4 | 8 | 2 | 128 | 1.08 |
| KNL | 4 | 1 | 16 | 4 | 16 | 1.36 |
| KNL | 4 | 2 | 16 | 4 | 32 | 1.11 |
| KNL | 4 | 4 | 16 | 4 | 64 | 1.00 |
| KNL | 4 | 1 | 32 | 8 | 8 | 3.35 |
| KNL | 4 | 2 | 32 | 8 | 16 | 2.31 |
| KNL | 4 | 4 | 32 | 8 | 32 | 2.95 |
| KNL | 8 | 1 | 8 | 1 | 64 | 0.68 |
| KNL | 8 | 2 | 8 | 1 | 128 | 0.81 |
| KNL | 8 | 1 | 16 | 2 | 32 | 0.68 |
| KNL | 8 | 2 | 16 | 2 | 64 | 0.61 |
| KNL | 8 | 4 | 16 | 2 | 128 | 0.77 |
| KNL | 8 | 1 | 32 | 4 | 16 | 0.70 |
| KNL | 8 | 2 | 32 | 4 | 32 | 0.58 |
| KNL | 8 | 4 | 32 | 4 | 64 | 0.60 |
| KNL | 8 | 1 | 64 | 8 | 8 | 2.31 |
| KNL | 8 | 2 | 64 | 8 | 16 | 2.06 |
| KNL | 8 | 4 | 64 | 8 | 32 | 2.22 |
| Xeon | 1 | 1 | 1 | 1 | 24 | 5.58 |
| Xeon | 1 | 2 | 1 | 1 | 48 | 4.62 |
| Xeon | 1 | 1 | 2 | 2 | 12 | 5.79 |
| Xeon | 1 | 2 | 2 | 2 | 24 | 4.77 |
| Xeon | 1 | 1 | 4 | 4 | 6 | 6.34 |

| | | | | | | |
|---|---|---|---|---|---|---|
| Xeon | 1 | 2 | 4 | 4 | 12 | 4.93 |
| Xeon | 2 | 1 | 2 | 1 | 24 | 2.81 |
| Xeon | 2 | 2 | 2 | 1 | 48 | 2.35 |
| Xeon | 2 | 1 | 4 | 2 | 24 | 2.90 |
| Xeon | 2 | 2 | 4 | 2 | 48 | 2.41 |
| Xeon | 2 | 1 | 8 | 4 | 6 | 3.19 |
| Xeon | 2 | 2 | 8 | 4 | 12 | 2.63 |
| Xeon | 4 | 1 | 4 | 1 | 24 | 1.42 |
| Xeon | 4 | 2 | 4 | 1 | 48 | 1.18 |
| Xeon | 4 | 1 | 8 | 2 | 12 | 1.46 |
| Xeon | 4 | 2 | 8 | 2 | 24 | 1.22 |
| Xeon | 4 | 1 | 16 | 4 | 6 | 1.61 |
| Xeon | 4 | 2 | 16 | 4 | 12 | 1.49 |
| Xeon | 8 | 1 | 8 | 1 | 24 | 0.71 |
| Xeon | 8 | 2 | 8 | 1 | 48 | 0.60 |
| Xeon | 8 | 1 | 16 | 2 | 12 | 0.74 |
| Xeon | 8 | 2 | 16 | 2 | 24 | 0.62 |
| Xeon | 8 | 1 | 32 | 4 | 6 | 1.61 |
| Xeon | 8 | 2 | 32 | 4 | 12 | 1.49 |