

NAMD benchmarking on 4 nodes cluster with Intel Xeon Phi 7120P Coprocessors

by Ashish Kumar Singh

This blog explores the application performance analysis of NAMD (NANoscale Molecular Dynamics) for large data sets on cluster of PowerEdge R730 servers with Intel Xeon Phi 7120Ps. All the runs were carried out with Hyper Threading (logical processors) disabled. IB verbs version of NAMD was used for all the runs.

Test Cluster Configuration:

The test cluster consisted of four PowerEdge R730 servers with two Intel Xeon Phi 7120P co-processors each. Each PowerEdge R730 had two Intel Xeon E5-2695v3 @ 2.3GHz CPU and eight 16GB DIMMS of 2133MHz making it a total of 128GB of memory per server. Each PowerEdge R730 consisted of one Mellanox FDR Infiniband HCA card in the low-profile x8 PCIe Gen3 slot (Linked with CPU2).

Compute node configuration

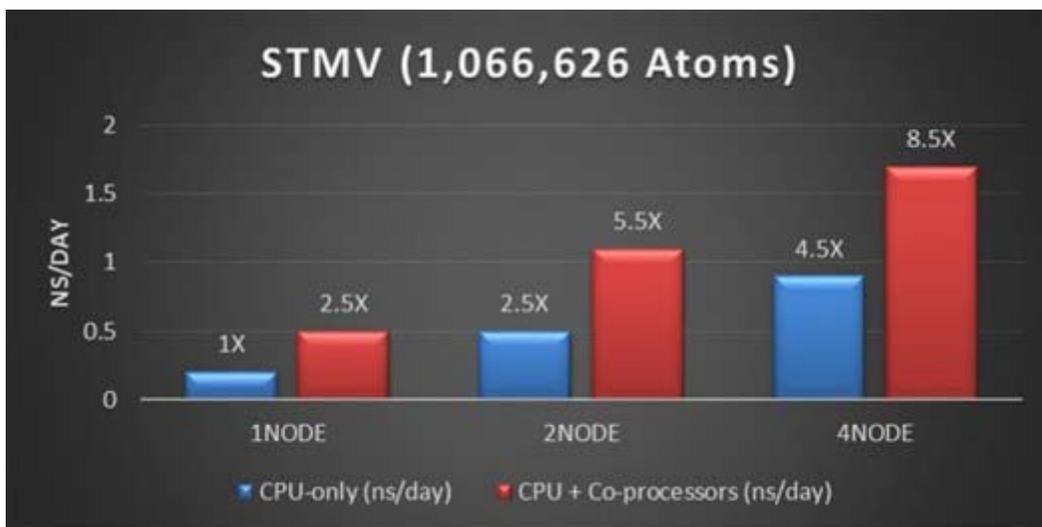
Component	Value
<i>Server</i>	<i>R730</i>
<i>Processor</i>	<i>Intel Xeon E5-2695v3 @2.3GHz</i>
<i>Memory</i>	<i>8x 16GB @2133MHz</i>
<i>Infiniband</i>	<i>Mellanox Connect-X3 FDR Adaptor (CX354A)</i>
<i>Cluster Size</i>	<i>4 Servers</i>
<i>OS</i>	<i>RHEL6.5</i>
<i>MPSS version</i>	<i>3.4.1</i>
<i>Coprocessor</i>	<i>2 x 7120P</i>
<i>Intel Compiler</i>	<i>Intel Parallel Studio XE 2015</i>
<i>NAMD</i>	<i>2.10</i>
<i>FTW</i>	<i>2.1.5</i>
<i>TCL</i>	<i>8.6.2</i>

The BIOS options selected for this blog are as below:

System BIOS Options	Settings
<i>Memory Settings > Snoop Mode</i>	<i>Early Snoop</i>
<i>Processor Settings > Logical Processor</i>	<i>Disabled</i>
<i>Processor Settings > QPI Speed</i>	<i>Maximum Data Rate</i>
<i>Processor Settings > Configurable TDP</i>	<i>Nominal</i>
<i>System Profile Settings > System Profile</i>	<i>Performance</i>

NAMD (NANoscale Molecular Dynamics) is a parallel, object-oriented simulation package written using the Charm++ parallel programming model, designed for high performance simulation of large biomolecular systems. Charm++ is developed with simplified parallel programming and also provides automatic load balancing, which is crucial to the performance of NAMD.

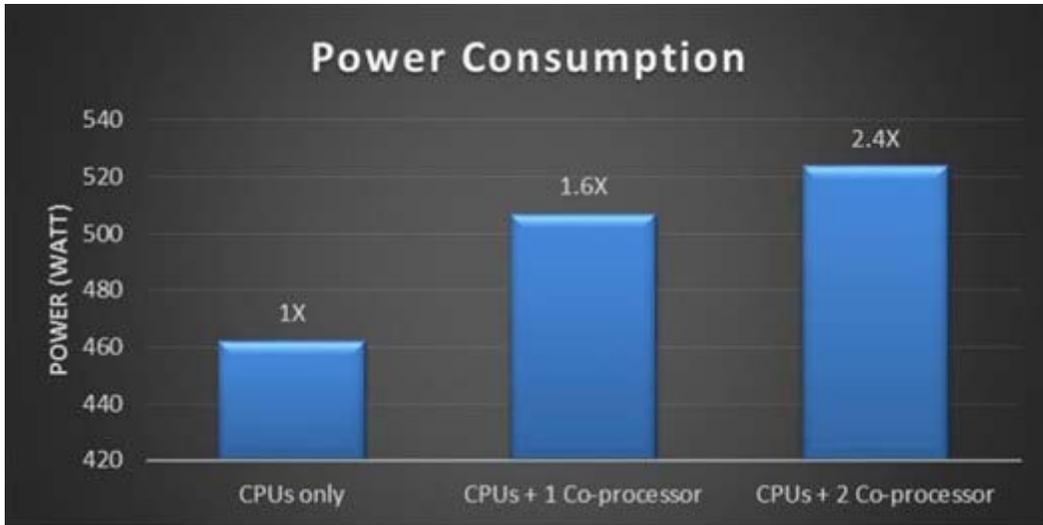
All the runs with STMV (virus) benchmark were run with ibverbs version of NAMD. The performance analysis with STMV benchmark shown below. STMV (Satellite Tobacco Mosaic Virus) is a small, icosahedral plant virus. On single node, we observed performance improvement of 2.5 times on CPUs with Intel Xeon Phi configuration in comparison to CPUs-only configuration.



STMV (1,066,626 Atoms)			
	1NODE	2NODE	4NODE
CPU-only (ns/day)	0.2	0.5	0.9
CPU + Co-processors (ns/day)	0.5	1.1	1.7

STMV showed performance of 0.2ns/day with CPUs-only configuration. With CPUs and two Intel Xeon Phi performance was 0.5ns/day, which showed performance increase of 2.5 times. While on a four-node cluster with the CPUs and Intel Xeon Phi 7120P performance increase was 8.5 times. Scaling from one node to four-node resulted in almost 3.5 times scale-up.

The Power analysis was done for single node among CPUs-only configuration, CPUs with one Intel Xeon Phi 7120P configuration and CPUs with two Intel Xeon Phi 7120P configuration. With CPUs and two Intel Xeon Phi configuration, the power consumption increased along with the performance per watt, which was 2.4 times in comparison to CPU-only configuration. The power efficiency increase showed in below picture.



Power consumption (Watt)		
CPUs-only	CPUs with one Intel Xeon Phi	CPUs with two Intel Xeon Phi
462	506.8	523.6

Performance per Watt		
CPUs-only	CPUs with one Intel Xeon Phi	CPUs with two Intel Xeon Phi
0.00043	0.00069	0.00103

Conclusion:

With CPUs and two Intel Xeon Phi 7120Ps, the STMV benchmark demonstrated increase of 2.5 times in performance and 2.4 times in power efficiency when compared to CPUs-only configuration, resulting in a powerful and energy efficient HPC platform.