

LAMMPS Four Node Comparative Performance Analysis on Skylake Processors

Author: Joseph Stanfield

The purpose of this blog is to provide a comparative performance analysis of the Intel® Xeon® Gold 6150 processor (architecture code named “Skylake”) and the previous generation Xeon® E5-2697 v4 processor using the LAMMPS benchmark. The Xeon® Gold 6150 CPU features 18 cores or 36 when utilizing hyper threading. Intel significantly increased the L2 cache per core from 256 KB on previous generations of Xeon to 1 MB. The new processor also touts 24.75 MB of L3 cache and a six channel DDR4 memory interface.

[LAMMPS](#), or Large Scale Atom/Molecular Massively Parallel Simulator, is an open-source molecular dynamics program originally developed by Sandia National Laboratories, Temple University, and the United States Department of Energy. The main function of LAMMPS is to model particles in a gaseous, liquid, or solid state.

Test cluster configuration:

	Dell EMC PowerEdge C6420	Dell EMC PowerEdge C6320
CPU	2x Xeon® Gold 6150 18c 2.7 GHz (Skylake)	2x Xeon® E5-2697 v4 16c 2.3 GHz (Broadwell)
RAM	12x 16GB @2666 MT/s	8x 16GB @2400 MT/s
HDD	1TB SATA	1 TB SATA
OS	RHEL 7.3	RHEL 7.3
InfiniBand	EDR ConnectX-4	EDR ConnectX-4

BIOS Options	Settings
System Profile	<i>Performance Optimized</i>
Logical Processor	<i>Disabled</i>
Virtualization Technology	<i>Disabled</i>

The LAMMPS version used for testing release was lammps-6June-17. The in.eam dataset was used for the analysis on both configurations. In.eam is a dataset that simulates a metallic solid, Cu EAM potential with 4.95 Angstrom cutoff (45 neighbors per atom), NVE integration. The simulation was executed using 100 steps with 32,000 atoms.

The first series of benchmarks conducted were to measure performance in units of timesteps/s. The test environment consisted of four servers interconnected with InfiniBand EDR, and tests were run on a single node, two nodes, and four nodes with LAMMPS, three times for each configuration. Average results from a single node showed 106 time steps per second while a two node result nearly doubled performance with 216 time steps per second. This trend remained consistent as the environment was scaled to four nodes as seen in Figure 1.

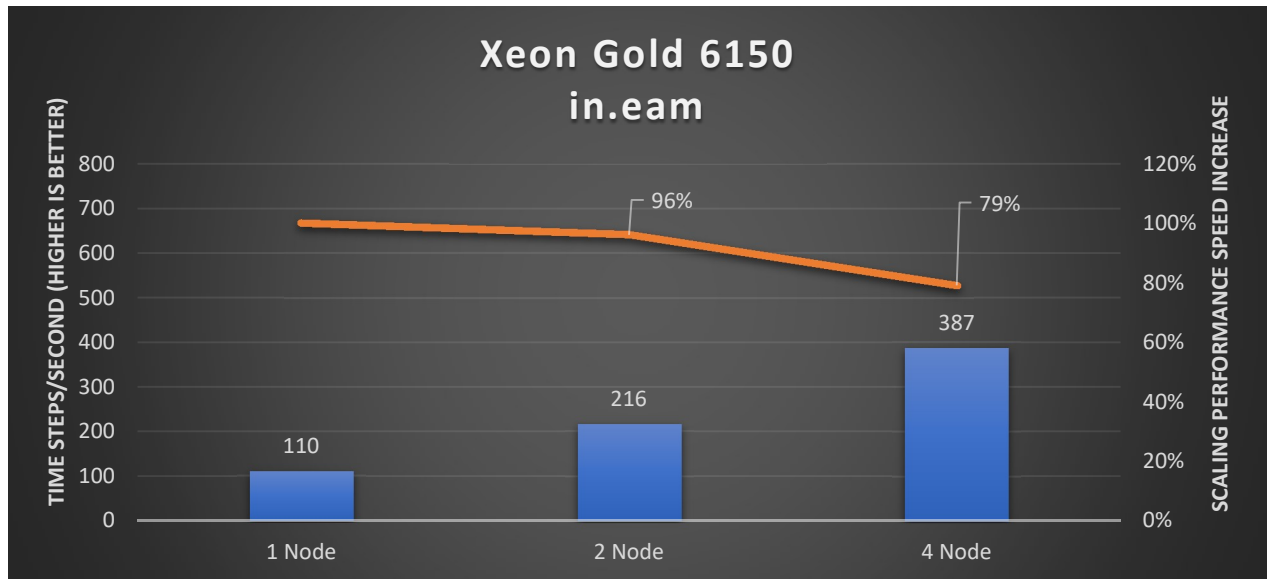


Figure 1.

The second series of benchmarks were run to compare the Xeon® Gold against the previous generation, Xeon® E5 v4. The same dataset, in.eam, was used with 32,000 atoms and 100 steps per run. As you can see below in Figure 2, the Xeon® Gold CPU outperforms the Xeon® E5 v4 by about 120% with each test, but the performance increase drops slightly as the cluster is scaled.

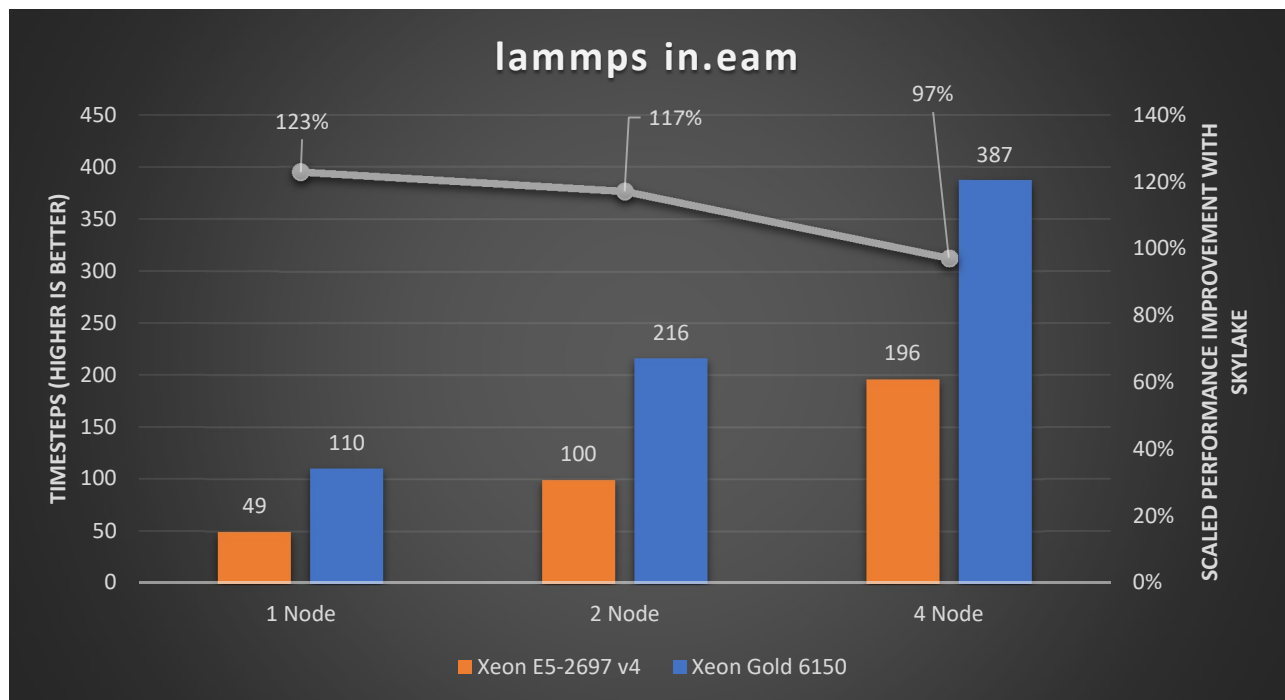


Figure 2.

Conclusion

In this blog, we analyzed and presented the performance of a Dell EMC PowerEdge C6420 cluster scaling from a single node to four nodes running the LAMMPS benchmark. Results show that performance of LAMMPS scales linearly with the increased number of nodes.

A comparative analysis was also conducted with the previous generation Dell EMC PowerEdge C6320 server with a Xeon® E5 v4 (Broadwell) processor. As with the first test, as node count was increased the linear scaling of the application was observed on Xeon® E5 v4, results similar to the Xeon® Gold.. But the Xeon® Gold processor outperformed the previous generation CPU by about 120% each run.

Resources

Intel LAMMPS Recipe: <https://software.intel.com/en-us/articles/recipe-lammps-for-intel-xeon-phi-processors>

LAMMPS USER-INTEL Package: http://lammps.sandia.gov/doc/accelerate_intel.html