

Comparing Haswell processor models for HPC applications

Garima Kochhar, September 2014

This blog evaluates four Haswell processor models (Intel® Xeon® E5-2600 v3 Product Family) comparing them for performance and energy efficiency on HPC applications. This is part three in a three part series. [Blog one](#) provided HPC results and performance comparisons across server generations, comparing Ivy Bridge (E5-2600 v2), Sandy Bridge (E5-2600) and Westmere (X5600) to Haswell. The [second blog](#) discussed the performance and energy efficiency implications of BIOS tuning options available on the new Dell Haswell servers.

In this study we evaluate processor models with different core counts, CPU frequencies and Thermal Design Power (TDP) ratings and analyze the differences in performance and power. Focusing on HPC applications, we ran two benchmarks and four applications on our server. The server in question is part of Dell's PowerEdge 13th generation (13G) server line-up. These servers support DDR4 memory at up to 2133 MT/s and Intel's latest E5-2600 v3 series processors (architecture code-named Haswell). Haswell is a net new micro-architecture when compared to the previous generation Sandy Bridge/Ivy Bridge. Haswell based processors use a 22nm process technology, so there's no process-shrink this time around. Note the "v3" in the Intel product name – that is what distinguishes a processor as one based on Haswell micro-architecture. You'll recall that "E5-2600 v2" processors are based on the Ivy Bridge micro-architecture and plain E5-2600 series with no explicit version are Sandy Bridge based processors. Haswell processors require a new server/new motherboard and DDR4 memory. The platform we used is a standard dual-socket rack server with two Haswell-EP based processors. Each socket has four memory channels and can support up to 3 DIMMs per channel (DPC).

Configuration

Table 1 below details the applications we used and Table 2 describes the test configuration on the new 13G server.

Table 1 - Applications and benchmarks

Application	Domain	Version	Benchmark
Stream	Memory bandwidth	v5.9	Triad
HPL	Computation - solve a dense system of linear equations	From Intel MKL	Problem size 90% of total memory
Ansys Fluent	Computational fluid dynamics	v15.0	truck_poly_14m
LS-DYNA	Finite element analysis	v7_0_0_79069	car2car with endtime=0.02
WRF	Weather Research and Forecasting	v3.5.1	Conus 2.5km
MILC	Quantum chromo dynamics	v7.7.3, v7.7.11	Input data file from Intel

Table 2 - Server configuration

Components	Details
Server	PowerEdge R730xd prototype
Processor	2 x Intel® Xeon® E5-2693 v3 – 2.6/2.2 GHz, 14c, 145W 2 x Intel® Xeon® E5-2680 v3 – 2.5/2.1 GHz, 12c, 120W 2 x Intel® Xeon® E5-2660 v3 – 2.6/2.2 GHz, 10c, 105W 2 x Intel® Xeon® E5-2640 v3 – 2.6/2.2 GHz, 8c, 90W * Frequency noted as “Rated base/ AVX base GHz”
Memory	128GB - 8 x 16GB 2133 MHz DDR4 RDIMMs
Hard drive	1 x 300GB SAS 6Gbps 10K rpm
RAID controller	PERC H330 mini
Operating System	Red Hat Enterprise Linux 6.5 x86_64
Kernel	2.6.32-431.el6.x86_64
BIOS settings	As noted per test
MPI	Intel® MPI 4.1.3.049
Math Library	Intel® MKL 11.1.3.174
Compilers	Intel® 2013_sp1.3.174 - v14.0.3.174 Build 20140422

All the results shown here are based on single-server performance. The following metrics were used to compare performance.

- Stream – Triad score as reported by the stream benchmark.
- HPL – GFLOP/second as reported by the benchmark.
- Fluent – Solver rating as reported by the application.
- LS DYNA – Elapsed Time as reported by the application.
- WRF – Average time step computed over the last 719 intervals for Conus 2.5km.
- MILC – Time as reported by the application.

Power was recorded during the tests on a power meter attached to the server. The average steady state power is used as the power metric for each benchmark.

Energy efficiency (EE) computed as Performance per Watt (performance/power).

Results

Figure 1 plots the performance of the four processor models (SKUs) for the benchmarks and applications used in this study. The BIOS was set to Early Snoop memory mode (ES), DAPC system profile (Turbo enabled, C-states enabled), and Logical Processor (Hyper-Threading) was turned off. All other BIOS options were at Dell defaults. The baseline used for comparison is the E5-2660 v3 10c 2.6 GHz processor.

From the graph, it can be seen that the memory bandwidth for the first three processor models (E5-2697 v3, E5-2690 v3, E5-2660 v3) was about the same. These SKUs can support memory at 2133 MT/s. The E5-2640 v3 has lower memory bandwidth since the maximum memory speed supported is 1866 MT/s.

All the other applications show a steady performance improvement with higher bin processor models. For codes that have per-core licenses, the improvement with higher bin processors that have more cores is not commensurate with the increase in number of cores. For example, when comparing the 10c SKU to the 12c SKU, adding 20% more cores (20 cores vs. 24 cores) allows Fluent running truck_poly_14m a 17% performance improvement.




Figure 1 - Performance - ES.DAPC

Figure 2 plots the relative energy efficiency of the test cases in Figure 1 using the 10c E5-2660 v3 SKU as the baseline.

Since the 14c, 12c and 10c SKUs have very similar memory bandwidth @ 2133 MT/s and the higher end processors have higher TDP and consume more power, the Stream EE follows the inverse of the TDP. (EE is performance/power).

HPL shows an improvement in energy efficiency with higher bin processors, the improvements in performance out-weigh the additional power consumed by the higher wattage processors.

For all the other applications, the energy efficiency is within 5% for each SKU and varies per application. Fluent and LS-DYNA share similar characteristics with the 12c and 8c SKUs measuring slightly EE than the 14c and 10c SKUs. WRF and MILC have similar trends with the lower end SKUs showing better EE than the higher end SKUs.




Figure 2 - Energy Efficiency - ES.DAPC

Figure 3 plots similar results for performance and energy efficiency on a different BIOS configuration. In these tests the BIOS was set to Cluster On Die Snoop mode, Performance profile (Turbo enabled, Cstates disabled) and Logical Processor disabled. [Recall](#) that the Cluster On Die (COD) mode is only supported on SKUs that have two memory controller per processor, i.e. 10 or more cores. The 8c E5-2640 v3 does not support COD mode.

The relative performance and energy efficiency patterns shown in Figure 3 for COD.Perf match those of the ES.DAPC mode (Figures 1 and 2). We know from [Blog 2](#) that COD.Perf performs 1-3% better than ES.DAPC for the applications and data sets used in this study. This improvement is seen across the different processor models given that the relative performance between SKUs stays similar for ES.DAPC and COD.Perf BIOS settings.




Figure 3 - Performance, Energy Efficiency - COD.DAPC

Figure 4 plots the idle and peak power consumption across the different BIOS settings. (Note this data was gathered on an early prototype unit running beta firmware.) The power measurements shown here are for comparative purposes across SKUs and *not* an absolute indicator of the server's power requirements. The text values within the bar graph show relative values using the 10c E5-2660 v3 as a baseline.

The idle power of the system is similar irrespective of the processor model used. This is good and demonstrates the energy efficiency of the base system. As desired, a higher wattage processor does not consume additional power when the system is idle.

As expected, the peak power draw measured during HPL initialization is greater for the higher bin CPUs that have higher TDP.




Figure 4 - Power idle and peak - ES.DAPC

Conclusion

There is a clear performance up-side for all the applications and datasets studied here when using higher bin/higher core count processors. The goal of this study was to quantify these performance improvements as a guide to choosing the best processor model for a workload or cluster. This blog concludes our three part series on the impact of the new Dell 13G servers with Intel Haswell processors on HPC applications.