

Measuring Performance of Intel Broadwell Processors with High Performance Computing Benchmarks

Authors: Ashish Kumar Singh, Mayura Deshmukh and Neha Kashyap

The increasing demand for more compute power pushes servers to be upgraded with higher and more powerful hardware. With the release of the new Intel® Xeon® processor E5-2600 v4 family of processors (architecture codenamed “Broadwell”), Dell has refreshed the 13th generation servers to benefit from the increased number of cores and higher memory speeds thus benefiting a wide variety of HPC applications.

This blog is part one of “Broadwell performance for HPC” blog series and discusses the performance characterization of Intel Broadwell processors with High Performance LINPACK (HPL) and STREAM benchmarks. The next three blogs in the series will discuss the BIOS tuning options and the impact of Broadwell processors on Weather Research Forecast (WRF), NAMD, ANSYS® Fluent®, CD-adapco® STAR-CCM+®, OpenFOAM, LSTC LS-DYNA® HPC applications as compared to the previous generation processor models.

In this study, performance was measured across five different Broadwell processor models listed in Table2 along with 2400 MT/s DDR4 memory. This study focuses on HPL and STREAM performance for different BIOS profiles across all five Broadwell processor models and compares the results to previous generations of Intel Xeon processors. The platform we used is a PowerEdge R730, which is a 2U dual socket rack server with two processors. Each socket has four memory channels and can support up to 3 DIMMs per channel (DPC). For our study, we used 2 DPC for a total of 16 DDR4 DIMMs in the server.

Broadwell (BDW) is a tick in Intel’s tick-tock principle as the next step in semiconductor fabrication. It is a 14nm processor with the same microarchitecture as the Haswell-based (HSW, Xeon E5-2600 v3 series) processors with the same TDP range. Broadwell E5-2600 v4 series processors support up to 22 cores per socket with up to 55MB of LLC, which is 22% more cores and LLC than Haswell. Broadwell supports DDR4 memory with max memory speed of up to 2400 MT/s, 12.5% higher than the 2133 MT/s that is supported with Haswell.

Broadwell introduces a [new snoop mode option in the BIOS memory setting](#), Directory with Opportunistic Snoop Broadcast (DIR+OSB), which is the default snoop mode for Broadwell. In this mode, the memory snoop is spawned by the Home Agent and a directory is maintained in the DRAM ECC bits. DIR+OSB mode allows for low local memory latency, high local memory bandwidth and I/O directory cache to reduce directory update overheads for I/O accesses. The other three snoop modes: Home Snoop (HS), Early Snoop (ES), and Cluster-on-Die (COD) are similar to what was available with Haswell. The Cluster-on-die (COD) is only supported on processors that have two memory controllers per processor. The Dell BIOS on systems that support both Haswell and Broadwell will display the supported snoop modes based on the processor model populated in the system.

Table 1 describes the other new features available in the Dell BIOS on systems that support Broadwell processors.

Table1: New BIOS features with Intel Xeon E5 v4 processor family (Broadwell)

BIOS feature	Description
Snoop Mode > Directory with Opportunistic Snoop Broadcast (DIR+OSB)	Directory with Opportunistic Snoop Broadcast, available on select processor models, works well for workloads of mixed NUMA optimization. It offers a good balance of latency and bandwidth.
System Profile Settings > Write Data CRC	When set to enabled, the DDR4 data bus issues are detected and corrected during 'write' operations. Two extra cycles are required for CRC bit generation which impacts the performance. Read-only unless System Profile is set to Custom.
System Profile Settings > CPU Power Management > Hardware P States	If supported by the CPU, Hardware P States is another performance-per-watt option that relies on the CPU to dynamically control individual core frequency. Read-only unless System Profile is set to Custom.
System Profile Settings > C States > Autonomous	Autonomous is a new BIOS option for C States in addition to the previous options, Enable and Disable. Autonomous (if Hardware controlled is supported), processor can operate in all available Power States to save power, but may increase memory latency and frequency jitter.

Intel Broadwell supports Intel® Advanced Vector Extensions 2 (Intel AVX2) vector technology, which allows a processor core to execute 16 FLOPs per cycle. HPL is a benchmark that solves a dense linear system. The HPL problem size (N) was chosen to be 177408 along with a block size (NB) of 192. The theoretical peak value of HPL was calculated using the AVX base frequency, which is lower than rated base frequency of the processor model. [Broadwell processors consume more power when running Intel® AVX2 workloads than non-AVX workloads](#). Starting with the Haswell product family Intel provides two frequencies for each SKU. Table 2 lists the rated base and AVX base frequencies of each Broadwell processor used for this study. Since HPL is an AVX-enabled workload, we would calculate HPL theoretical maximum performance with AVX base frequency as (AVX base frequency of processor * number of cores * 16 FLOP/cycle)

Table 2: Base frequencies of Intel Broadwell Processors

Base Frequencies of Intel Broadwell processors			
Processors	Rated base frequency (GHz)	AVX base frequency (GHz)	Theoretical Maximum Performance (GFLOPS)
E5-2699 v4, 22 core, 145W	2.2	1.8	634
E5-2698 v4, 20 core, 135W	2.2	1.8	576
E5-2697A v4, 16 core, 145W	2.6	2.2	563
E5-2690 v4, 14 core, 135W	2.6	2.1	470
E5-2650 v4, 12 core, 105W	2.2	1.8	346

Table 3 gives more information about the hardware configuration and the benchmarks used for this study.

Table 3: Server and Benchmark details for Intel Xeon E5 v4 processors

Server	PowerEdge R730
Processor	As described in table 2
Memory	16 x 16GB DDR4 @ 2400 MT/s (Total=256GB)
Power Supply	2 x 1100W
Operating System	RHEL 7.2 (3.10.0-327.el7.x86_64)

BIOS options	System profile – Performance and Performance Per Watt (DAPC) Logical Processor – Disabled Power Supply Redundant Policy – Not Redundant Power Supply Hot Spare Policy – Disabled I/O Non-Posted Prefetch - Disabled Snoop modes – OSB, ES, HS and COD Node Interleaving - Disabled
BIOS Firmware	2.0.0
iDRAC Firmware	2.30.30.02
HPL	From Intel Parallel Studio 2016 update 1
MPI	Intel MPI - 5.1.2
STREAM	5.4

Intel Broadwell Processors




Figure1: HPL performance characterization

Figure 1 shows HPL characterization of all five Intel Broadwell processors used for this study, with the PowerEdge R730 platform. Table 2 shows the TDP values for each of the Broadwell processors. The text value in each bar shows

the efficiency of that processor. The “X” value on top of each bar shows the performance gain over 12 core Broadwell processor. The HPL performance improvement with top bin Broadwell processor is not correspondingly increasing as number of cores. For example, adding 83% more cores in top bin 22 core than 12 core Broadwell processor, allows HPL a 57% performance improvement. The line pattern on the graph shows the HPL performance per core. Since the HPL performance is not accelerating as per number of cores, the performance per core has decreased by 8 to 15 % for 20 and 22 core processors respectively.




Figure2: STREAM (Triad) Performance characterization

The STREAM benchmark calculates the memory bandwidth by counting only the bytes that the user program requested to be loaded or stored. This study uses the results reported by the TRIAD function of the stream bandwidth test.

Figure 2 plots the STREAM (TRIAD) performance for all Broadwell processors used for this study. The bars show the memory bandwidth in GB/s for each of the processors. As per the graph, memory bandwidth across all Broadwell processors is approximately same. Since, the memory bandwidth across all Broadwell processors are same, the memory bandwidth per core is decreasing due to more number of cores.

BIOS Profiles comparison

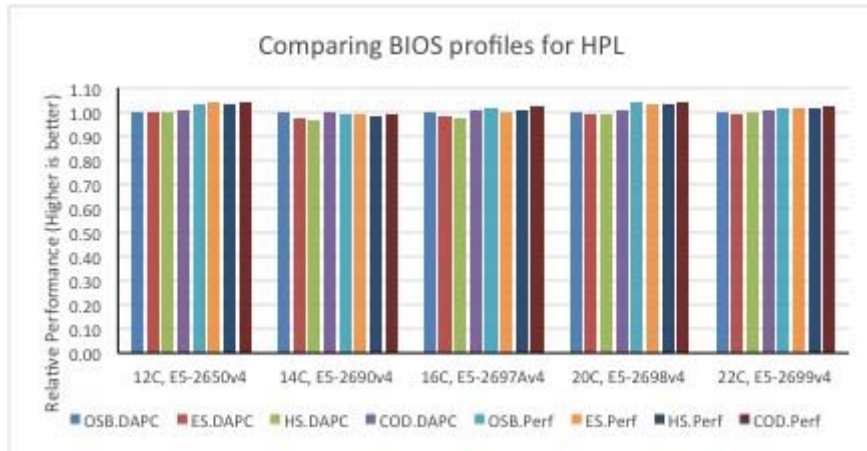


Figure 3: Comparing BIOS profiles with HPL

Figure 3 plots HPL performance with two [BIOS system profile options](#) for all four snoop modes across all five Broadwell processors. As Directory + Opportunistic Snoop Broadcast (DIR+OSB) snoop mode performs well for all workloads and DAPC system profile balances performance and energy efficiency, these options are set as default in the BIOS and so has been chosen as the baseline.

From this graph, Cluster-on-Die (COD) memory mode with the “Performance” System Profile setting performs 2 to 4 % better than other BIOS profile combinations across all Broadwell processors. The Cluster-on-die (COD) is only supported on processors that have two memory controllers per processor, i.e. 12 or more cores.

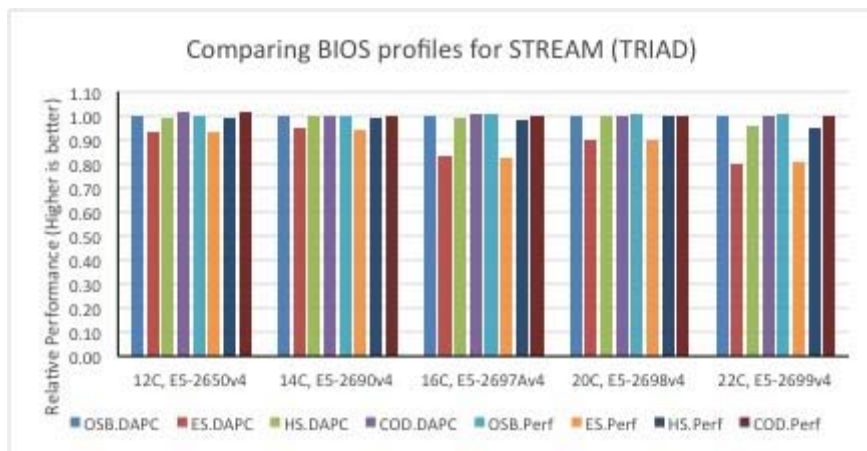


Figure 4: Comparing BIOS profiles with STREAM (TRIAD)

Figure 4 shows the STREAM performance characteristics with two [BIOS system profile options](#) for all the snoop modes. Opportunistic snoop Broadcast (OSB) snoop mode along with DAPC system profile is chosen as the baseline for this study. Memory Bandwidth with each BIOS profile combination except Early snoop (ES) mode with both system

profiles are almost same. The memory bandwidth with Early snoop (ES) mode for both system profiles is lower by 8 to 20 % and the difference is more apparent for 22 core processor up to 25%. The Early Snoop (ES) mode have less Requester Transaction IDs (RTIDs) distributed across all the cores, while other snoop modes gets higher RTIDs, that is higher number of credits for local and remote traffic at the home agent.

Generation over Generation comparison

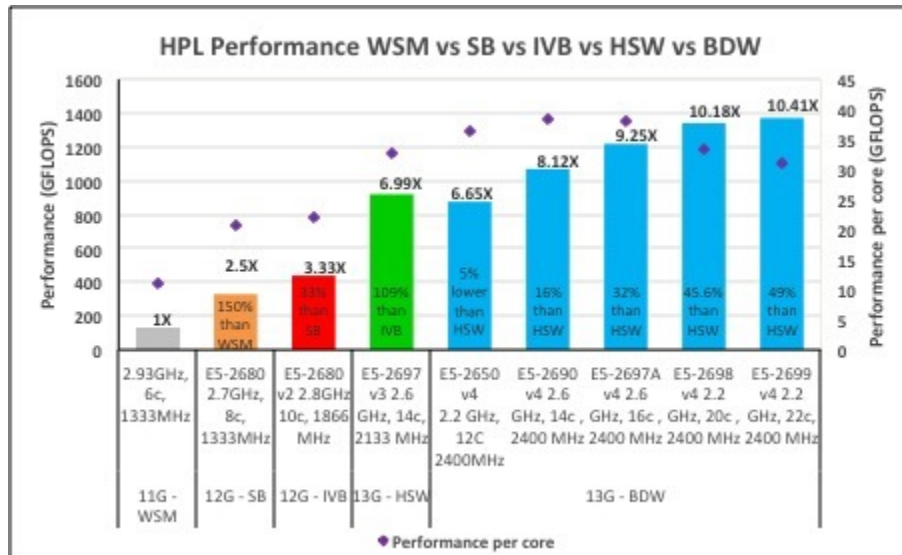


Figure 5: Comparing HPL Performance across multiple generations of Intel processors

Figure 5 plots generation over generation performance comparison for HPL with Intel Westmere (WSM), Sandy Bridge (SB), Ivy-Bridge (IVB), Haswell (HSW) and Broadwell (BDW) Processors. The percentages on the bars shows the HPL performance improvement than their previous generation processor. The graph shows that the 14 core Broadwell processor with similar frequencies performs 16% better than 14 core Haswell processor for the HPL benchmark. Broadwell processors measure better power efficiencies than the Haswell processors. The top bin 22 core Broadwell processor performance is 49% better than 14 core Haswell processor. The purple diamonds in the graph show the performance per core. The “X” value on top of every bar shows acceleration over 6 core WSM processor.

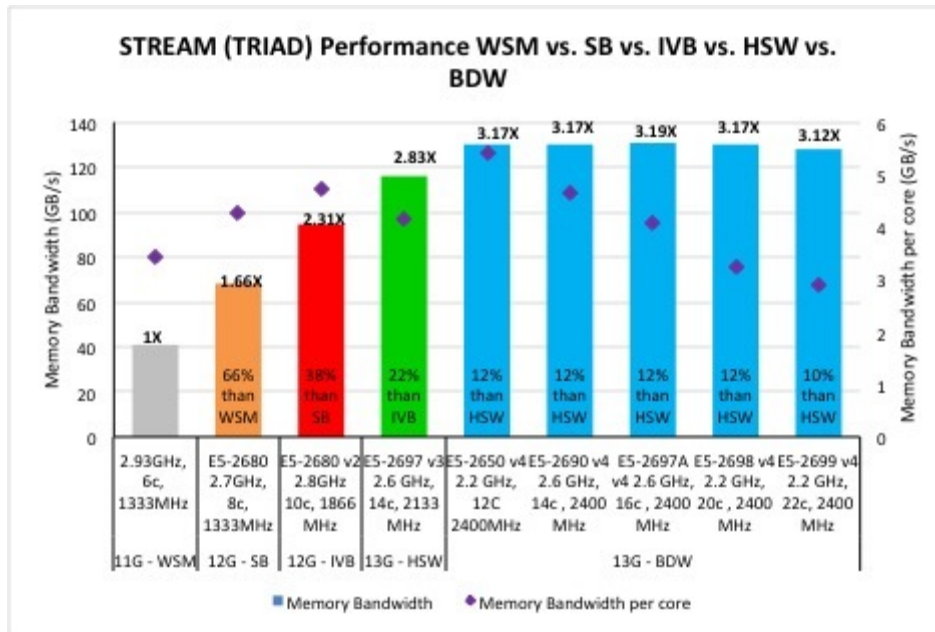


Figure 6: Generation over generation comparison with STREAM

Figure 6 plots performance comparison of STREAM (TRIAD) for multiple generations of Intel processors. From the graph, it can be seen that the memory bandwidth on the system has increased over generations. The theoretical maximum memory frequency increased by 12.5% in Broadwell over Haswell (2133 MT/s to 2400 MT/s) and this translates into 10 to 12% better measured memory bandwidth as well. However, the maximum core-count per socket has increased by up to 22% in Broadwell over Haswell, and so the memory bandwidth per core depends on the specific Broadwell SKU. The 20 core and 22 core BDW processors support only ~3 GB/s per core and that is likely to be very low for most HPC applications, the 16core BDW is on par with the 14core HSW at ~4 GB/s per core.

Conclusion

The performance of all Broadwell processor used for this study is higher for both HPL and STREAM benchmarks. There is ~12% increase in measured memory bandwidth for Broadwell processors compared to Haswell processors. Broadwell processors measure better power efficiencies than the Haswell processors. In conclusion, Broadwell processors may fulfill the demands of more compute power for HPC applications.