

# Eeny Meeny Miney Mo, should I go with 20 cores?

Nishanth Dandapanthula, June 2014

---

Intel's Xeon E5-2600 v2 product family processors (architecture code named Ivy Bridge) have been available in the server market for a few months now. Ivy bridge processor based systems provide better performance when compared to previous generation processor families such as Sandy Bridge (Xeon E5-2600) and Westmere (Xeon X5600). This can be attributed to several factors, such as increased core counts because of the 22nm process technology, higher clock rates, increased system memory speeds and larger last level cache. This performance improvement is shown in several studies mentioned in [1](#), [2](#) and [3](#).

So, once the decision to move to a new platform or new processor technology has been made, what next? How should these new systems be configured? There are so many choices for the processor itself – different options with different core counts, processor frequency, TDP and, of course, price. **Which processor model is optimal for a specific workload?** This blog provides quantitative data and analysis to help answer this question by comparing the performance and power profile of different processor models across a variety of HPC applications.

Once the decision has been made to choose a particular processor SKU, there is another important decision which needs to be made. This is in regard to the choice between single rank and dual rank memory modules. As mentioned in studies [4](#) and [5](#), the choice and configuration of memory modules impacts the performance and power consumption statistics of an environment. This blog also describes a performance comparison done between single and dual rank memory modules for several HPC applications. These studies were done in the Dell engineering lab in December 2013 and results are actual measured results.

Table 1 describes the configuration of the HPC test cluster and the benchmarks used for this analysis. Six different types of Ivy Bridge processors were studied along with two types of memory modules, single ranked and dual ranked.

BIOS options were set to reflect Maximum Performance settings, details are in Table 1. A four node cluster was interconnected with InfiniBand. All of the results shown in this blog were obtained by fully subscribing the four servers, i.e., all cores on all servers were in use.

**Table 1 Test bed configuration and Benchmarks**

<b>Servers</b>	4 x Dell PowerEdge C6220 II sleds in one chassis
<b>Processors per server</b>	2 x E5-2643 v2 @ 3.5 GHz 6c 130 W
	2 x E5-2667 v2 @ 3.3 GHz 8c 130 W
	2 x E5-2680 v2 @ 2.80 GHz 10c 115 W
	2 x E5-2670 v2 @ 2.5 GHz 10c 115W
	2 x E5-2690 v2 @ 3.00 GHz 10c 130 W
	2 x E5-2697-v2 @ 2.7 GHz 12c 130W
<b>Memory per server*</b>	1 DPC 8 x 16 GB 1866 MT/s Dual Rank

	1 DPC 8 x 8GB 1866 MT/s Single Rank	
<b>Interconnect</b>	Mellanox ConnectX3 FDR adapter	Mellanox SX 6036 FDR Switch
<b>BIOS Options</b>	Turbo	Enabled
	C States	Disabled
	Hyper Threading	Disabled
	Node Interleaving	Disabled
	Power Management	Max Performance
<b>OFED version</b>	Mellanox OFED 1.5.3-3	
<b>OS</b>	RHEL 6.4 (kernel 2.6.32-358.el6.x86_64)	
<b>BIOS version</b>	2.1.2	
<b>BMC version</b>	2.25	
<b>Benchmarks</b>	HPL 2.0	N set to 85 % of total memory available NB = 168
	NAS PB 3.3.1	LU Class D
	Stream for Memory Bandwidth	N = 160000000
	Open Foam 2.2.1	Dam break 3D
	ANSYS Fluent v15.0	Truck_111m Truck_poly_14m
	LS DYNA 6.1.0	Top crunch 3 Vehicle Collision

\*1 DPC is 1 DIMM per channel.

## Results and Analysis

Figure 1 Performance across four nodes using multiple IVB processors and dual rank DIMMs




Figure 1 compares the different processor models for each benchmark from a pure performance standpoint. The graph plots cluster-level relative performance for each benchmark using the Xeon E5-2643 v2 processor (3.5 GHz 6 core) as the baseline. Dual Rank DIMMs were used for this study.


From the graph it can be seen that best performance is obtained with the 12 core E5-2697 v2 (2.7G Hz 12 core) processor in all the cases except for LS-DYNA. HPL is a very compute intensive applications and does the best with E5-2697 v2 processors because of the sheer number of cores. The theoretical peak performance [6](#) for HPL is also the highest for E5-2697 v2 processor compared to all the others. NAS parallel benchmark's LU benchmark is also very processor intensive and thus shows the same behavior as HPL. Similarly, ANSYS Fluent is more sensitive to core speed and number of cores than memory bandwidth, thus replicating the behavior where the E5-2697 v2 processor performs the best.

LS-DYNA is a structural and fluid analysis simulation software used in manufacturing, crash testing, aerospace industry, automobile industry etc. It has high CPU and high memory bandwidth requirements. On a given server, a processor with fewer cores will result in higher memory bandwidth per core since the total memory bandwidth is fixed. The E5-2697 v2 with 12 cores has lower memory bandwidth per core when compared to the 8 core E5-2667 v2. From the results, this 8c model appears to be the sweet spot for LS DYNA. Why then would the E5-2643 v2 with only 6c and therefore maximum memory bandwidth per core not be the best choice? In the case of the 6c E5-2643 v2, the additional memory bandwidth per core is more than the application's requirements and thus wasted providing no additional gain.

The other aspect is processor frequency. Once the memory bandwidth requirements are met, faster processors will perform better. This is clearly illustrated in the 10c case with the 3.0 GHz processors performing better than the 2.8 GHz parts, which in turn perform better than the 2.5 GHz model.

Thus in the LS DYNA scenario, the E5-2667 v2 performs the best, *providing the right balance of memory bandwidth and processor speed.*

Figure 2 Performance per Watt across four nodes using multiple IVB processors and dual rank DIMMs

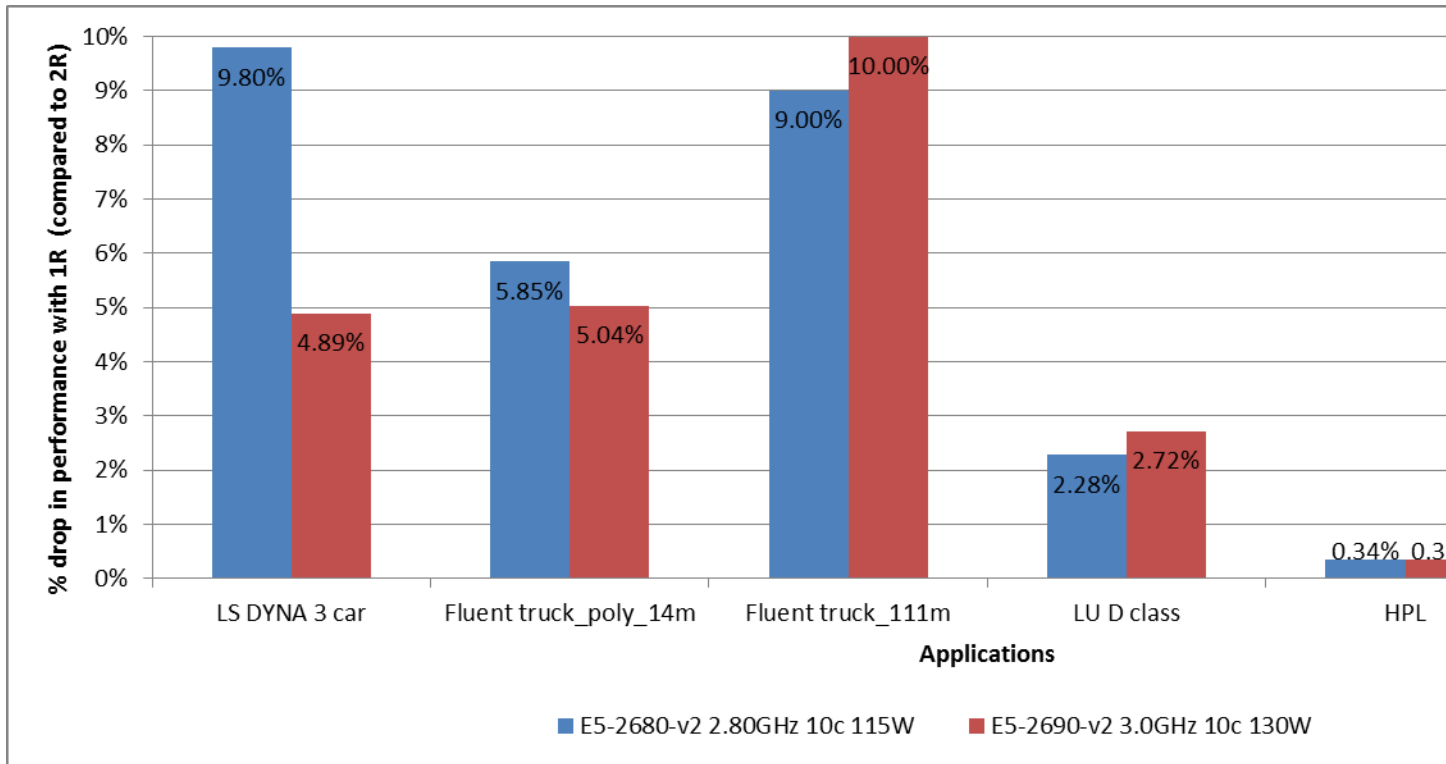


Performance is one metric, but most administrators and data centers are also concerned with power consumption. We use a “performance per watt” metric to quantify energy efficiency. For HPL, this metric translates to “GFlops per Watt” and for LS-DYNA, Fluent, LU and Open Foam this is “(Number of jobs per day) per Watt”. The power consumption was measured using metered PDUs and represents the actual power consumption of the four-node cluster for each test.

Figure 2 plots the energy efficiency of each processor model for the different benchmarks. As before, the 6c 3.5 GHz E5-2643 v2 processor is used as the baseline. Dual rank DIMMs were used for this testing. From the figure, the 2.7 GHz 12c E5-2697 v2 processor has the best energy efficiency in all the cases except for LS-DYNA. Even though the E5-2697 v2 processor has a TDP of 130 Watts and the power consumption is high, the performance obtained as a result of higher core counts compensates in the performance to power ratio thus providing the best energy efficiency. In case of LS-DYNA, the E5-2670

v2 10 core part is doing the best in terms of performance per watt but the E5-2667 v2 processor does the best in terms of pure performance. The difference in this case is less than 5 %. This could be because of lower wattage of the E5-2670 v2 processor (115 Watts).

Figure 3 Performance drop when using single rank memory modules



As mentioned previously, choice of memory modules can affect application performance. The study done in this blog shows that Dual rank memory modules give the best performance. Figure 3 shows the percentage drop in performance when single rank DIMMs are used when compared to a scenario where dual rank DIMMs are used. Recall that both single rank and dual rank configurations used single DIMM per channel. HPL is very compute intensive and the dependency on memory is very limited. So, we see the least drop in performance when single rank memory modules are used in case of HPL. In all other scenarios single rank DIMMs perform ~3 – 10 % lower than dual rank DIMMs.

Figure 4 Performance per core across four nodes for multiple processors using dual rank memory modules




Figure 4 shows the application performance in terms of performance per core. The 3.5 GHz E5-2643 v2 processor has the highest clock speed among all the processors used and gives the best performance per core. In case of a scenario where the application being used has a per core licensing fee, a better choice would be to pick the E5 2643 v2 processor which has fewer cores with high clock speed, rather than the E5-2697 v2 processor, which has double the number of cores but almost 23% lower clock speed than the 3.5 GHz E5 2643 v2 processor.

### Conclusion

In conclusion, in terms of performance and performance per Watt, The E5-2697 v2 processor along with dual rank DIMMs provides the best performance for the HPC applications and benchmark data sets evaluated here. However, the selection should be different if there is a licensing cost per core involved with the application (Example: LS-Dyna, Fluent). Finally, this decision depends on the application itself and its characteristics.

### References

1. [http://en.community.dell.com/techcenter/high-performance-computing/b/general\\_hpc/archive/2013/11/07/performance-comparison-of-intel-xeon-e5-2600-and-e5-2600-v2-series-processors.aspx](http://en.community.dell.com/techcenter/high-performance-computing/b/general_hpc/archive/2013/11/07/performance-comparison-of-intel-xeon-e5-2600-and-e5-2600-v2-series-processors.aspx)
2. [http://en.community.dell.com/techcenter/high-performance-computing/b/general\\_hpc/archive/2012/03/06/hpc-performance-on-the-12th-generation-12g-poweredge-servers.aspx](http://en.community.dell.com/techcenter/high-performance-computing/b/general_hpc/archive/2012/03/06/hpc-performance-on-the-12th-generation-12g-poweredge-servers.aspx)
3. [ftp://ftp.dell.com/Manuals/Common/poweredge-r420\\_White%20Papers96\\_en-us.pdf](ftp://ftp.dell.com/Manuals/Common/poweredge-r420_White%20Papers96_en-us.pdf)
4. [http://www.dell.com/downloads/global/products/pedge/poweredge\\_12th\\_generation\\_server\\_memory.pdf](http://www.dell.com/downloads/global/products/pedge/poweredge_12th_generation_server_memory.pdf)
5. [http://en.community.dell.com/techcenter/high-performance-computing/b/hpc\\_storage\\_and\\_file\\_systems/archive/2012/10/09/unbalanced-memory-performance.aspx](http://en.community.dell.com/techcenter/high-performance-computing/b/hpc_storage_and_file_systems/archive/2012/10/09/unbalanced-memory-performance.aspx)
6. <http://www.netlib.org/utk/people/JackDongarra/faq-linpack.html>